



Network Functions Virtualisation (NFV) Release 2; Testing; NFVI Compute and Network Metrics Specification

Disclaimer

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

Reference

DGS/NFV-TST008

Keywords

network, NFV, NFVI, SDN, testing

ETSI

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

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

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

Important notice

The present document can be downloaded from:

<http://www.etsi.org/standards-search>

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

Users of the present document should be aware that the document may be subject to revision or change of status.

Information on the current status of this and other ETSI documents is available at

<https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx>

If you find errors in the present document, please send your comment to one of the following services:

<https://portal.etsi.org/People/CommiteeSupportStaff.aspx>

Copyright Notification

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.

© European Telecommunications Standards Institute 2017.

All rights reserved.

DECT™, **PLUGTESTS™**, **UMTS™** and the ETSI logo are Trade Marks of ETSI registered for the benefit of its Members.

3GPP™ and **LTE™** are Trade Marks of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.

oneM2M logo is protected for the benefit of its Members

GSM® and the GSM logo are Trade Marks registered and owned by the GSM Association.

Contents

Intellectual Property Rights	4
Foreword.....	4
Modal verbs terminology.....	4
Introduction	4
1 Scope	5
2 References	5
2.1 Normative references	5
2.2 Informative references.....	5
3 Definitions and abbreviations.....	6
3.1 Definitions.....	6
3.2 Abbreviations	6
4 Time and Time Intervals for Metrics.....	6
5 Framework for Metric Definitions	6
6 Compute Metrics	7
6.1 Background	7
6.2 Name	7
6.3 Parameters	8
6.4 Scope	8
6.5 Units of Measure	8
6.6 Definition	8
6.7 Method of Measurement	8
6.8 Sources of Error	9
6.9 Discussion	9
7 Network Metrics.....	9
7.1 Background	9
7.2 Name	9
7.3 Parameters	10
7.4 Scope	10
7.5 Units of Measure	10
7.6 Definition	10
7.7 Method of Measurement	11
7.8 Sources of Error	11
7.9 Discussion	11
8 Memory Metrics.....	11
8.1 Background	11
8.2 Name	11
8.3 Parameters	12
8.4 Scope	12
8.5 Units of Measure	12
8.6 Definition	12
8.7 Method of Measurement	12
8.8 Sources of Error	12
8.9 Discussion	13
9 Follow-on Activities.....	13
Annex A (informative): Examples of Scope Specification for Metrics	14
A.1 Description	14
Annex B (informative): Authors & contributors.....	15
History	16

Intellectual Property Rights

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

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

Foreword

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

Modal verbs terminology

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

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

Introduction

Although many metrics for the performance and utilization of the Network Function Virtualisation Infrastructure (NFVI) components have been in wide use for many years, there were no independent specifications to support consistent metric development and interpretation. The present document provides the needed specifications for key NFVI metrics.

1 Scope

The present document specifies detailed and vendor-agnostic key operational performance metrics at different layers of the NFVI, especially processor usage and network interface usage metrics. These metrics are expected to serve as references for processed and time-aggregated measurement values for performance management information that traverses the Or-Vi and/or Vi-Vnfm reference points of the NFV architectural framework. The present document contains normative provisions.

2 References

2.1 Normative references

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

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

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

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

Not applicable.

2.2 Informative references

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

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

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

[i.1] ETSI GS NFV-INF 003 (V.1.1.1) (12-2014): "Network Functions Virtualisation (NFV); Infrastructure; Compute Domain".

[i.2] Linux/UNIX system programming training, Linux man-pages: "TOP(1)".

NOTE: Available at http://man7.org/linux/man-pages/man1/top.1.html#2._SUMMARY%C2%A0Display.

[i.3] O'Reilly Linux Dev Center: "Exploring the /proc/net/ Directory".

NOTE: Available at <http://www.onlamp.com/pub/a/linux/2000/11/16/LinuxAdmin.html>.

[i.4] RHEL™ 6.8 Deployment Guide: "E.2.18. /proc/meminfo".

NOTE: Available at https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/s2-proc-meminfo.html.

[i.5] ETSI GS NFV 003: "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV".

[i.6] IETF RFC 7348: "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks".

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the terms and definitions given in ETSI GS NFV 003 [i.5] apply.

3.2 Abbreviations

For the purposes of the present document, the abbreviations given in ETSI GS NFV 003 [i.5] and the following apply:

CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
HZ	Hertz of the system clock, an operating system parameter
KiB	Kibibytes
NFVI	Network Function Virtualisation Infrastructure
OS	Operating System
RAM	Random Access Memory
VIM	Virtual Infrastructure Manager
VXLAN	Virtual eXtensible Local Area Network (IETF RFC 7348 [i.6]).

4 Time and Time Intervals for Metrics

Coherent compute domains [i.1] usually need access to a clock with accurate time-of-day (or simply time) and sources of periodic interrupts. Time sources are accessed to provide timestamps for events and log entries that document the recent history of the compute environment. Periodic interrupts provide a trigger to increment counters and read current conditions in the compute and networking environments.

With the foundation of time, date, and periodic interrupts, a measurement system can determine the beginning and end of time intervals, which is a fundamental aspect of metrics that involve counting and collecting events.

Table 4-1 specifies requirements applicable to time, date, and periodic interrupts.

Table 4-1: Requirements applicable to time, date and periodic interrupts

General-Time-01	The compute domain shall be able to take readings from (or access) a clock with accurate time-of-day and calendar date.
General-Time-02	The compute domain shall have a source of periodic interrupts available which are derived from the time-of-day clock, with configurable period (a parameter of metrics that use this feature).

When the results from measurement systems are collected and reported by management systems, the management systems may provide an additional time and date reading associated with the operation to collect the results, using their own time source.

5 Framework for Metric Definitions

The metric definitions in the present document are primarily based on the fact that the resources of the NFVI have utilization and performance characteristics that can be assessed by measurement processes. The resources may be implemented in hardware, software (such as virtual resources), or a combination of both. The measurement processes are primarily implemented in software (such as in the kernel or user space), but may be assisted by features of the hardware.

The measured NFVI resources and the measurement processes shall be completely specified in the dimensions of model numbers, firmware versions, software versions, and any other aspects that influence the results (such as physical location of the components within a datacentre's racks and shelves). For example, the fixed frequency of the physical CPU clock in Hz, which governs the rate that the CPU executes instructions, is one important descriptor of the NFVI.

Clock Speed may depend on other CPU settings, such as energy-saving power control. For one list of NFVI platform descriptors, see clause 5.1 of ETSI GS NFV-INF 003 [i.1].

For each metric it specifies, the present document provides the following elements:

- Background
- Name
- Parameters (input factors)
- Scope of coverage
- Unit(s) of measure
- Definition
- Method of Measurement
- Sources of Error
- Discussion

NOTE: The present document specifies well-known metrics, and assumes that Virtual Infrastructure Managers (VIM) will control and expose the metrics as specified here, or will be enhanced to collect and convey the metrics with the required framework elements, which are Name, Parameters, Scope, Units of measure, and the source of the measurement (where the metric was measured, which may be synonymous with the Scope).

6 Compute Metrics

6.1 Background

The Central Processing Unit (CPU) is an essential component of every coherent compute domain. Each CPU is a limited resource in terms of the instructions per second it can execute. It is valuable to monitor the utilization of the CPU resource to fulfil the goals of maintaining continued and efficient operations, and for troubleshooting abnormal behaviour to find root causes. For many uses, it is helpful to categorize the CPU's execution time into multiple execution contexts, such as system and user contexts. A compute node may include additional processors beyond the main CPU; the metrics specified in this clause can also be used to measure and report the usage of such processors.

VNFs also have a view of CPU resources in terms of execution time they have used during a measurement interval. However, the configured instantiation of the VNF determines how to map its view of virtual processor resource usage to actual hardware CPU resources available and used. For example, a VNF's processes may be pinned to one or more CPU cores, or the VNF may be sharing the resources of many CPU cores with other VNFs.

6.2 Name

There are two variants of this metric:

- Processor Usage
- Processor Utilization

The two variants allow reporting this metric as a percentage. The metric is a function of the scope, set of reported contexts, measurement interval and other factors.

6.3 Parameters

The following parameters shall be supported for this metric:

- **Tick interval:** the period of timed interrupts when the processor's execution context can be recorded. Note that this parameter is an integral part of the method of measurement. The tick interval is sometimes called a "jiffy". The tick interval is controlled by a system parameter, "HZ", whose default value shall be 100 for measurements complying with the present document.
- **Set of execution contexts:** the desired set of processor states with reported utilization. For example, the simplest set includes two states: active and idle. Time in the active context can be calculated as a sum of states with more specific definitions, such as the active states user and system. A commonly chosen set of four states is user, system, wait, and idle. See [i.2] for a list of seven states available in Linux OS.
- **End time:** the termination of the measurement interval (time and date).
- **Measurement interval:** the duration of the observation by the measurement system to assess the metric.

6.4 Scope

The list of one or more compute resources which shall be included in the values reported, and whether the resource is physical or virtual. Annex A gives examples of the scope usage for physical and virtual processor metrics.

6.5 Units of Measure

Processor usage results shall be reported as time in seconds, and utilization shall be reported as the ratio of total time in one or more execution contexts to the total time in the measurement interval, expressed as a percentage.

6.6 Definition

The compute resource usage at time T (for a given scope, context, and measurement interval) shall be assessed as indicated in table 6.6-1.

Table 6.6-1: Requirements for Processor Metric variants

Requirement number	Metric Variant Name	Requirement
Compute-01	Processor Usage	Reported values of this metric variant shall quantify the total time that one or more compute resources (according to the defined scope) execute instructions in the specified execution context during the measurement interval
Compute-02	Processor Utilization	Reported values of this metric variant shall quantify the ratio of time of Processor Usage to the time in the measurement interval, expressed as a percentage

6.7 Method of Measurement

The measurement process shall interrupt the processor periodically, determine the execution context which was present when the timed interrupt occurred, and allocate the entire time interval since the previous interrupt to the usage total of that execution context. The total times for each desired context shall be accumulated throughout the specified measurement interval. On the completion of a measurement interval, the measured times shall be summed and the usage and utilization shall be reported.

6.8 Sources of Error

The sources of error for this metric are listed below:

- 1) The method only reads the execution context when the timed interrupt occurs. Therefore, a context which is present only briefly between successive timed interrupts will not be observed and accumulate no measured usage.
- 2) The interval between successive time interrupts (between ticks, or the length of a jiffy) is a configurable system control setting in most operation systems. Use of a long interval with respect to processor context swapping will result in usage assessment with large quantization and to frequently miss contexts that have short duration.
- 3) The measurement process itself uses processor resources, and the resources consumed increase with the number of timed interrupts, or ticks per measurement interval. In Linux systems, the system parameter that controls the tick interval or duration of a jiffy is called HZ.
- 4) The accurate generation of timed interrupts on the intended interval boundaries is a function of many factors, including the activity of other interrupts and virtualisation of the process that generates the tick interval. There can be slightly less of slightly more than the configured number of ticks in each second, and this causes error in both the processor usage and utilization measurements.
- 5) Synchronization of the time-of-day clock with an external reference usually ensures sufficiently accurate timestamps and measurement intervals, but loss of synchronization for an extended period will cause time accuracy to suffer.

6.9 Discussion

The processor usage measurement for virtual processors is a topic of ongoing study and advanced development. Methods which produce more accurate usage measurements are expected in the future.

7 Network Metrics

7.1 Background

Network Interfaces exist at many points on the communication paths through the NFVI. It is valuable to monitor the traffic load of arbitrary interfaces, both the physical interfaces with hardware resources and their virtual counterparts. Additional information can assist troubleshooting and resolution, so interfaces also indicate problems with attempted transmission or reception of packets.

7.2 Name

There are four fundamental metrics of traffic load and communication issues:

- 1) Packet count.
- 2) Octet count.
- 3) Dropped Packet count.
- 4) Errored Packet count.

These four metrics apply to a single direction of transmission. Therefore, transmit and receive versions of the four metrics make a total of eight for each interface.

7.3 Parameters

The following parameters shall be supported for the four metrics:

- **Interface:** the name of a single interface where communication metrics are monitored.
- **Measurement time:** the point in time when the counter(s) were read (time and date).
- **Interface Speed:** the nominal frequency of the physical interface bit clock in Hz, which governs the rate that the interface operates. Virtual interfaces may not have a meaningful value for this parameter.
- **Interface Status:** the operational state of the interface indicating readiness for use, usually expressed as "up" or "down".

7.4 Scope

The list of one or more interface resources which are included in the values reported, and whether the resource is physical or virtual.

7.5 Units of Measure

The Packet interface metrics shall be reported in units of packets as defined by the layer at which the interface is operating, such as the link layer. Examples include the Ethernet layer and VXLAN.

The Octet interface metric shall be reported in units of 8-bit bytes contained in successfully communicated packets as defined by the layer at which the interface is operating, such as link layer. Examples include the Ethernet layer and VXLAN.

7.6 Definition

The values of each metric at time T (for a given scope and measurement time) shall be assessed as the total counts meeting the criteria as specified in table 7.6-1, since the interface was registered with the overall management process.

Table 7.6-1: Requirements for Network Metrics

Requirement number	Metric Name	Requirement
Network-01	Packet Count	Reported values of this metric shall quantify the number of successfully transmitted or received packets from the individual interface's perspective.
Network-02	Octet Count	Reported values of this metric shall quantify the number of 8-bit bytes that constitute all successfully transmitted or received packets from the individual interface's perspective.
Network-03	Dropped Packet Count	Reported values of this metric shall quantify the number of discarded packets due to lack of available processing resources, such as memory buffer space.
Network-04	Errored Packet Count	For received packets, reported values of this metric shall quantify the number of packets detected with various forms of corruption during transmission, including errors detected with checks on length (too long, too short), and errors detected with integrity checks (e.g. based on cyclic redundancy checks, CRC). For transmitted packets, reported values of this metric shall quantify the number of packet transmission attempts with errors, including loss of link carrier during transmission, aborted transmission (due to multi-access collisions), and other detectable forms of transmission failure.

7.7 Method of Measurement

The interface driver shall interrupt its management process when one or more transmission or reception events take place, and to increment the relevant counters. When a measurement is performed at time T, the present communications counter values of each of the desired interfaces shall be recorded for both the transmit and receive directions.

7.8 Sources of Error

The sources of error for these metrics are listed below:

- 1) Counter roll-over: All counters have a finite range, and end-of-range roll-over will likely be encountered during the life of a busy interface. This is especially true for the Octets count and Gbit/s interfaces. Mitigations include reading the counters often enough to detect roll-over algorithmically and provide a wider range for the counts.
- 2) Hardware discard: Some received packets will be discarded at the physical layer, and may not be included in the counts of Dropped packets. In other words, not all lost packets are reflected in error counts.
- 3) Resource limitations: Since interface counts are managed according to the priority of their interrupts, circumstance may occur when all interrupts cannot be served in a sufficiently timely fashion, and information (including packets) may be lost.
- 4) Synchronization of the time-of-day clock with an external reference usually ensures sufficiently accurate timestamps, but loss of synchronization for an extended period will cause time accuracy to suffer.

7.9 Discussion

When interfaces are activated or created, they are registered with the management processes (in the Linux kernel, for example). Each registered interface effectively defines the packet type and header format and contents. For example, an Ethernet definition would reference the Ethernet header and trailer.

For the Linux OS, [i.3] describes the information available directly from the kernel and how to access it from the command line.

8 Memory Metrics

8.1 Background

Memory is a key resource the NFVI. It is valuable to monitor the utilization and management of system and user memory and their virtual counterparts. This information can assist troubleshooting and resolution.

8.2 Name

There are four fundamental metrics of memory utilization:

- 1) Memory buffered
- 2) Memory Cached
- 3) Memory free
- 4) Memory Slab

These four metrics comprise the total of used memory. Therefore, the sum of these metrics can be subtracted from the Total Memory to obtain the current value of free memory.

8.3 Parameters

The following parameters shall be supported for these four metrics:

- **Measurement time:** the point in time when the values were read (time and date).
- **Total Memory:** the Random Access Memory, RAM, available to the compute node measured.
- **Swap space:** the configured memory available for processes to share through swapping.

8.4 Scope

The list of one or more compute node resources which are included in the values reported, and whether the resource is physical or virtual.

8.5 Units of Measure

The memory values shall be reported in units of Kibibytes, or KiB, where 1 KiB equals 1 024 Bytes.

NOTE: That some systems will report values and display the units of kB (equal to 1 000 Bytes), such as the Linux `/proc/meminfo`. However, the values are actually in units of KiB.

8.6 Definition

The values of each metric at time T (for a given scope and measurement time) shall be assessed as the total KiB meeting the criteria as specified in table 8.6-1.

Table 8.6-1: Requirements for Memory Metrics

Requirement number	Metric Name	Requirement
Memory-01	Memory Buffered	Reported values of this metric variant shall quantify the amount of temporary storage for raw disk blocks
Memory-02	Memory Cached	Reported values of this metric variant shall quantify the amount of RAM used as cache memory
Memory-03	Memory Free	Reported values of this metric variant shall quantify the amount of RAM unused
Memory-04	Memory Slab	Reported values of this metric variant shall quantify the amount of memory used as a data structure cache by the kernel

8.7 Method of Measurement

When a measurement is performed at time T, each of the desired metrics and their present values shall be recorded for reporting.

8.8 Sources of Error

The sources of error for these metrics are listed below:

- 1) Synchronization of the time-of-day clock with an external reference usually ensures sufficiently accurate timestamps, but loss of synchronization for an extended period will cause time accuracy to suffer.

8.9 Discussion

For the Linux OS, [i.3] describes the information available directly from the kernel and how to access it from the command line. [i.4] describes the information available in /proc/meminfo. Note that the methods of measurement for storage systems vary widely and depend on the implementation, therefore storage metrics are not considered for standardization at this time.

9 Follow-on Activities

The present document focuses on specifications of metrics reported by the NFVI and virtualised versions of platform resources. Additional metrics could be provided in the future, such as Storage Metrics.

Normative metric specifications are also needed for externally observable characteristics of the NFVI, for example, in terms of the communications speed and capacity provided.

Annex A (informative): Examples of Scope Specification for Metrics

A.1 Description

As described in clause 5 of the present document, each metric specification is composed of a framework with nine elements. One framework element is the Scope applicable to each measurement, which provides the list of physical or virtual resources intended to be measured. An example of each type of metric (physical and virtual) is illustrated in figure A.1-1.

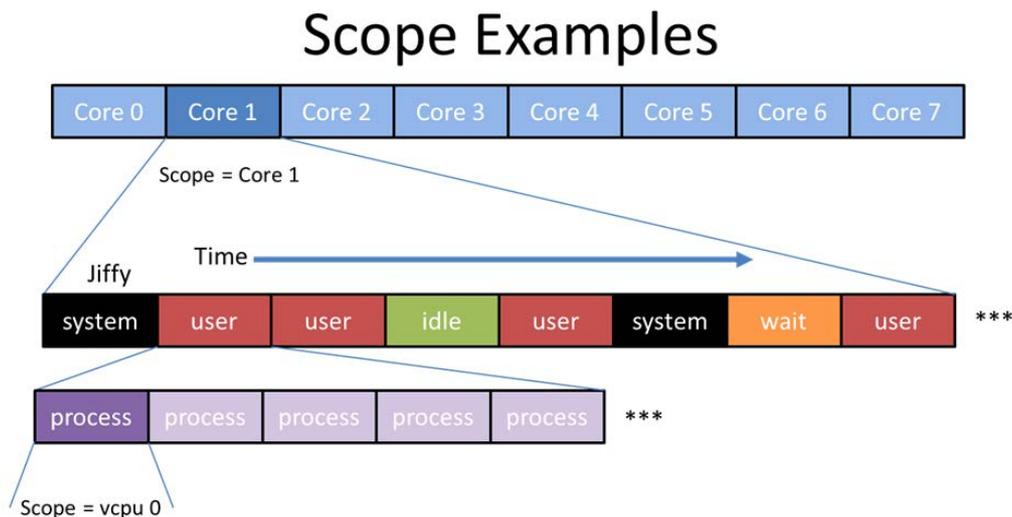


Figure A.1-1: Examples of Scope coverage for Physical and Virtual Processor Usage Metrics and Measurements

The top of figure A.1-1 depicts a compute resource with eight physical processor cores. As an example of a physical processor measurement Scope, figure A.1-1 indicates coverage of a single processor core (Core 1).

Within the single processor core, processing time is divided among different processor contexts [i.2]. The context is measured at each tick interval (called a Jiffy). The time spent in each context is accumulated and reported at the end of each measurement interval.

The user context is further divided among the processes spawned by different applications. VNFC (synonymous with Virtual Machines (VM) or OS containers (OSC), or a Libvirt domain) will occupy one or more processes.

Last, figure A.1-2 illustrates a Libvirt domain specification with a single virtual CPU (vcpu) that has been statically placed on Core 1 of the physical processor. Therefore, the scope of process identification(s) associated with the VNFC of this domain describes its compute resource usage within the user context (sometimes called the guest context in some virtualisation methods), and comprises the coverage of the processor usage measurement for a VM or OSC.

```
Libvirt domain spec for CPU allocation
<domain>
...
  <vcpu placement='static' cpuset="1" current="1"></vcpu>
  <vcpus>
    <vcpu id='0' enabled='yes' hotpluggable='no' order='1' />
  </vcpus>
  ...
</domain>
```

Figure A.1-2: Example of Scope coverage partly based on Libvirt domain specification

In summary, the Scope element provides a way to narrow-down the resources to which a metric will be applied, and the coverage of measurements using the metric definition.

Annex B (informative): Authors & contributors

The following people have contributed to the present document:

Rapporteur:

Mr. Al Morton, AT&T Labs

Other contributors:

Ms. Maryam Tahhan, Intel Corporation

Mr. Pierre Lynch, Ixia Technologies

Mr. Gergely Csatari, Nokia

Mr. Bruno Chatras, Orange

Mr. Joerg Aelken, Ericsson Telefonaktiebolaget LM

Mr. Klaus Seggelke, Spirent Communications

History

Document history		
V2.1.1	May 2017	Publication