ETSI GS NFV-TST 008 V3.5.1 (2021-12)

Network Functions Virtualisation (NFV) Release 3; 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

RGS/NFV-TST008ed351

Keywords

metrics, network, NFV, NFVI, 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 - 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:
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 prevailing version of an ETSI deliverable is the one made publicly available in PDF format at www.etsi.org/deliver.

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

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.

Contents

Intellectual Property Rights 5
Foreword....5
Modal verbs terminology....5
Introduction 5
1 Scope 6
2 References 6
2.1 Normative references 6
2.2 Informative references....6
3 Definition of terms, symbols and abbreviations....7
3.1 Terms....7
3.2 Symbols....8
3.3 Abbreviations 8
4 Time and Time Intervals for Metrics....8
5 Framework for Metric Definitions 9
6 Compute Metrics 9
6.1 Background 9
6.2 Name 10
6.3 Parameters 10
6.4 Scope....10
6.5 Units of Measure 10
6.6 Definition 10
6.7 Method of Measurement 11
6.8 Sources of Error 11
6.9 Discussion 11
7 Network Metrics....12
7.1 Background 12
7.2 Name 12
7.3 Parameters 12
7.4 Scope....12
7.5 Units of Measure 13
7.6 Definition 13
7.7 Method of Measurement 13
7.8 Sources of Error 13
7.9 Discussion 14
8 Memory Metrics....14
8.1 Background 14
8.2 Name 14
8.3 Parameters 15
8.4 Scope....15
8.5 Units of Measure 15
8.6 Definition 15
8.7 Method of Measurement 16
8.8 Sources of Error 16
8.9 Discussion 16
9 Storage Metrics 16
9.1 Background 16
9.2 Name 16
9.3 Parameters 17
9.4 Scope....17
9.5 Units of Measure 17
9.6 Definition 17

9.7 Method of Measurement 18
9.8 Sources of Error 18
9.9 Discussion 18

10 Follow-on Activities..19

Annex A (informative): Examples of Scope Specification for Metrics 20

A.1 Description 20

A.2 Memory Scope 21

History 23

Intellectual Property Rights

Essential patents

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

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

Trademarks

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

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

Foreword

This Group 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 Functions 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 (V1.1.1) (12-2014): "Network Functions Virtualisation (NFV); Infrastructure; Compute Domain".

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

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

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

[i.3] TecMint: "Exploring /proc File System in Linux", January 3, 2015.

NOTE: Available at https://www.tecmint.com/exploring-proc-file-system-in-linux/.

[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.

[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".

[i.7] "free(1) Linux User Commands", published by man7.org.

NOTE: Available at http://man7.org/linux/man-pages/man1/free.1.html.

[i.8] collectd codebase, file: "/src/memory.c", published by GitHub®, Inc.

NOTE: Available at https://github.com/collectd/collectd/blob/collectd-5.7/src/memory.c#L325.

[i.9] ETSI GS NFV-IFA 003 (V2.4.1): "Network Functions Virtualisation (NFV) Release 2; Acceleration Technologies; vSwitch Benchmarking and Acceleration Specification".

[i.10] Debian Wiki: "Hugepages".

NOTE: Available at https://wiki.debian.org/Hugepages.

[i.11] ETSI GS NFV-INF 010 (V1.1.1) (12-2014): "Network Functions Virtualisation (NFV); Service Quality Metrics".

[i.12] "Understanding CPU Steal Time - when should you be worried?", Scout application performance monitoring blog, Derek Haynes, July 25, 2013.

NOTE: Available at https://scoutapm.com/blog/understanding-cpu-steal-time-when-should-you-be-worried.

[i.13] "Understanding Linux Container Scheduling", Squarespace Engineering blog, June 27, 2017.

NOTE: Available at https://engineering.squarespace.com/blog/2017/understanding-linux-container-scheduling.

[i.14] "procfs-diskstats", October 30, 2020.

NOTE: Available in Linux Kernel at https://github.com/torvalds/linux/blob/master/Documentation/ABI/testing/procfs-diskstats or at https://www.kernel.org/doc/Documentation/ABI/testing/procfs-diskstats.

[i.15] "proc(5) - Linux manual page", man7.org, March 22, 2021.

NOTE: Available at https://man7.org/linux/man-pages/man5/proc.5.html.

[i.16] "collectd Wiki - Plugin: Disk", January 6, 2017.

NOTE: Available at https://collectd.org/wiki/index.php/Plugin:Disk.

[i.17] "collectd Wiki - Plugin: DF", June 11, 2015.

NOTE: Available at https://collectd.org/wiki/index.php/Plugin:DF.

[i.18] "SNIA Solid State Storage (SSS) Performance Test Specification (PTS) Version 2.0.1", February 5, 2018.

NOTE: Available at https://www.snia.org/sites/default/files/technical_work/PTS/SSS_PTS_2.0.1.pdf.

[i.19] "OPNFV StorPerf Project, Blog: Storage Performance Guidelines", Mark Beierl, October 12, 2018.

NOTE: Available at https://wiki.opnfv.org/x/UQO8AQ.

3 Definition of terms, symbols and abbreviations

3.1 Terms

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

3.2 Symbols

Void.

3.3 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
IOPS	Input or Output operations Per Second
kB	kiloBytes
KiB	Kibibytes
NFVI	Network Functions Virtualisation Infrastructure
OS	Operating System
OSC	Operating System Container
RAM	Random Access Memory
VIM	Virtual Infrastructure Manager
VM	Virtual Machine
VNF	Virtual Network Function
VNFC	Virtual Network Function Component
VXLAN	Virtual eXtensible Local Area Network

NOTE: See 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. The compute domain may contain a very large number of NFV compute nodes [i.1], and each node needs to execute a process to synchronize its hardware and system clocks to a source of accurate time-of-day, preferably traceable to an international time standard.

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	Each node in 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	Each node in 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-IFA 003 [i.9].

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 (VIMs) 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 250 Hz 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 eight 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 nanoseconds, 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/or the 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.

The main purpose of NFVI employing Hypervisors and Virtual Machines (VMs) (or other forms of virtualisation layers, such as Operating System Containers, OSCs) is to provide a predictable computing environment for VNFs. ETSI GS NFV-INF 010 [i.11] describes a set of Service Quality Metrics on the NFVI from the VNF point-of-view. One particular metric that was intended to constrain the predictability of computing environments is the VM Stall metric, defined in clause 5.1 of ETSI GS NFV-INF 010 [i.11]. Essentially, this metric is intended to measure both the frequency and duration of events where the VM was unable to execute its processes as planned, and execution is deferred to a later time.

Execution context is one of the Parameters for the Processor Usage or Utilization metric defined in clause 6.6 above, and [i.2] lists a set of states available in Linux OS. The key execution context of interest here is CPU steal time, or "st: time stolen from this vm by the hypervisor" [i.2]. Thus, for the Scope of a specific VM and the parameter CPU steal time, the Usage and Utilization metrics report the total time and percentage of time during the Measurement Interval that the VM Stall condition was present. While these results are not in the form of event frequency and duration originally envisioned in ETSI GS NFV-INF 010 [i.11], they should still be useful for trouble-shooting performance-related issues.

Ideally, CPU steal time Utilization should be zero percent. If a measurement indicates CPU steal time Utilization much greater than zero, there are several possible causes, including that the VM's resource allocation is too low or that the host is overbooked and the VM cannot access the configured resources [i.12].