



Mplify Standard

Mplify 132

**Edge Computing IaaS Subscriber Building Blocks
and Attributes**

October 2025

Disclaimer

© Mplify Alliance 2025. All Rights Reserved.

The information in this publication is freely available for reproduction and use by any recipient and is believed to be accurate as of its publication date. Such information is subject to change without notice and Mplify Alliance (Mplify) is not responsible for any errors. Mplify does not assume responsibility to update or correct any information in this publication. No representation or warranty, expressed or implied, is made by Mplify concerning the completeness, accuracy, or applicability of any information contained herein and no liability of any kind shall be assumed by Mplify as a result of reliance upon such information.

The information contained herein is intended to be used without modification by the recipient or user of this standard. Mplify is not responsible or liable for any modifications to this standard made by any other party.

The receipt or any use of this standard or its contents does not in any way create, by implication or otherwise:

- a) any express or implied license or right to or under any patent, copyright, trademark or trade secret rights held or claimed by any Mplify member which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor
- b) any warranty or representation that any Mplify members will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s)
- c) embody any or all of the ideas, technologies, or concepts contained herein; nor
- d) any form of relationship between any Mplify member and the recipient or user of this standard.

Implementation or use of specific Mplify standards, specifications, or recommendations will be voluntary, and no Member shall be obliged to implement them by virtue of participation in Mplify Alliance. Mplify is a global alliance of network, cloud, cybersecurity, and enterprise organizations working together to accelerate the AI-powered digital economy through standardization, automation, certification, and collaboration. Mplify does not, expressly or otherwise, endorse or promote any specific products or services.

Table of Contents

1	List of Contributing Members	8
2	Abstract.....	8
3	Terminology and Abbreviations	9
4	Compliance Levels	13
5	Numerical Prefix Conventions.....	13
6	Introduction.....	14
7	Edge Computing IaaS Building Blocks.....	17
7.1	Processor Building Block	17
7.2	Memory Building Block	17
7.3	Storage Building Block	18
7.4	Access Network Building Block	18
7.5	Host OS Building Block	18
7.6	Hypervisor Building Block.....	18
7.7	Guest OS Building Block	18
7.8	Container Engine Building Block.....	19
7.8.1	Container Orchestrator Building Block	19
7.9	Virtual Machine Building Block	19
7.9.1	Virtual Machine Orchestrator Building Block.....	19
7.10	Edge IaaS Offering Examples	19
8	Definition of Edge Computing IaaS Subscriber Service Building Block Service Attributes	22
8.1	Edge IaaS Identifier	22
8.2	Edge IaaS Type.....	23
8.3	Edge IaaS Installation Place	23
8.4	Building Block State.....	24
8.5	Building Block Identifier String.....	24
8.6	Processor Building Block Service Attributes	24
8.6.1	Manufacturer Identifier Service Attribute	26
8.6.2	Manufacturer Model Identifier Service Attribute	26
8.6.3	Processor Architecture Service Attribute.....	27
8.6.4	Clock Speed Max Service Attribute	27
8.6.5	Clock Speed Min Service Attribute	27
8.6.6	Number of Cores Service Attribute	28
8.6.7	Number of Threads Service Attribute.....	28
8.6.8	Level 1 Cache Size Service Attribute	28
8.6.9	Level 2 Cache Size Service Attribute	29
8.6.10	Instructions per Clock Cycle Service Attribute	29
8.6.11	Quantity Service Attribute	29
8.7	Memory Building Block	29
8.7.1	Memory Amount Service Attribute	30
8.8	Storage Building Block	30
8.8.1	Storage Type Service Attribute.....	31
8.8.2	Storage Retention Duration Service Attribute	31

8.8.3	Storage Time to Access Service Attribute	31
8.8.4	Storage Amount Service Attribute	31
8.8.5	Replication Service Attribute.....	31
8.9	Access Network Building Block	31
8.9.1	Subscriber Application Awareness Service Attribute.....	33
8.9.2	Cloud UNI.....	34
8.9.2.1	Cloud UNI Service Attributes.....	34
8.9.3	Cloud UNI Access Links	36
8.9.3.1	Cloud UNI Access Links Service Attributes	36
8.9.4	Edge Computing IaaS Cloud VC Service Attributes.....	36
8.9.5	Edge Computing IaaS Cloud VC End Points	38
8.9.6	Cloud Application Interface.....	39
8.9.6.1	Cloud Application Interface Service Attributes	39
8.10	Edge IaaS Access Service Attributes.....	41
8.10.1	Subscriber Management Access	41
8.10.2	Edge IaaS Access URL.....	41
8.10.3	Edge IaaS Access Identifier.....	42
8.10.4	Edge IaaS Access Password.....	42
8.11	Host OS.....	42
8.11.1	Host OS Vendor Identifier.....	42
8.11.2	Host OS Name	43
8.11.3	Host OS Version	43
8.12	Hypervisor	43
8.12.1	Hypervisor Vendor.....	43
8.12.2	Hypervisor Name	44
8.12.3	Hypervisor Version.....	44
8.13	Guest OS.....	44
8.13.1	Guest OS Publisher	45
8.13.2	Guest OS Name.....	45
8.13.3	Guest OS Version	45
8.14	Container Engine	45
8.14.1	Container Engine Vendor	46
8.14.2	Container Engine Name.....	46
8.14.1	Container Engine Version.....	46
8.15	Virtual Machine Requirements.....	47
8.15.1	Number of Virtual Processors.....	47
8.15.2	Number of Threads	47
8.15.3	Number of Cores.....	47
8.15.4	Number of Physical CPUs	48
8.15.5	Virtual Memory Amount	48
8.15.6	Host OS.....	48
8.16	Edge IaaS Power Requirements	48
8.17	Edge IaaS Physical Characteristics.....	48
8.18	Other Service Attributes for Edge IaaS	48
8.18.1.1	Security Policies.....	48
8.18.1.1.1	Security Policy Identifier.....	49
8.18.1.1.2	Security Policy Name	49
8.18.1.1.3	Security Policy Constraints.....	50
8.18.1.1.4	Security Policy Subject.....	50
8.18.1.2	Region Name	50
8.18.1.3	Availability Zone Name.....	51

8.18.2	Edge IaaS Service Level Specification	51
8.18.3	Time Interval Sequence	52
8.18.4	S Parameter	53
8.18.5	IaaS Availability Performance Metric	53
8.18.6	One-way Mean Packet Delay Performance Metric	53
8.18.7	Building Block Common Reporting Parameters	54
8.18.8	Processor Performance Reported Parameters	54
8.18.8.1	<i>Processor Utilization</i>	56
8.18.8.2	<i>Processor Utilization Period</i>	56
8.18.9	Memory Performance Service Attributes	56
8.18.9.1	<i>Memory Utilization</i>	58
8.18.9.2	<i>Memory Utilization Period</i>	58
8.18.10	Storage	58
9	References	59
Appendix A	Service Attributes by IaaS Type	60
Appendix B	Acknowledgements	65

List of Figures

Figure 1 – Edge Computing IaaS Components.....	15
Figure 2 – Building Blocks for Dedicated Offering Examples.....	20
Figure 3 – Building Blocks for Shared Offering Examples.....	21
Figure 4 – Access Network Building Block Implementation Example.....	32
Figure 5 – Subscriber Application Awareness Implementation Example	34

List of Tables

Table 1 – Terminology.....	11
Table 2 – Abbreviations.....	12
Table 3 – Numerical Prefix Conventions.....	13
Table 4 – Edge IaaS Type Building Block Building Block Service Attribute	22
Table 5 – Edge IaaS Type Building Block Service Attribute	23
Table 6 – Edge IaaS Installation Place Building Block Service Attribute	24
Table 7 – Processor Building Block Service Attributes	25
Table 8 – Memory Resource Service Attributes	30
Table 9 – Storage Resource Service Attributes	30
Table 10 – Access Network Building Block Components	33
Table 11 – Cloud UNI Subscriber Application Awareness Service Attributes	35
Table 12 – Edge Computing IaaS Cloud VC Application Awareness Service Attributes	38
Table 13 – Cloud Application Interface Service Attributes.....	40
Table 14 – Edge IaaS Access Service Attributes.....	41
Table 15 – Host OS Service Attributes.....	42
Table 16 – Hypervisor Service Attributes	43
Table 17 – Guest OS Service Attributes	44
Table 18 – Container Engine Service Attributes	46
Table 19 – Virtual Machine Service Attributes	47
Table 20 – Security Policy Service Attributes.....	49
Table 21 – PM Entry for the Edge Computing IaaS Availability Performance Metric.....	53
Table 22 – Building Block Common Performance Parameters and Definitions	54
Table 23 – Processor Performance Reported Parameters	55
Table 24 – Memory Performance Service Attribute Parameters	57
Table 25 – Storage Performance Service Attribute Parameters.....	58
Table 26 – Service Attribute to Edge IaaS Type Mapping	64

1 List of Contributing Members

The following members of Mplify participated in the development of this standard and have requested to be included in this list.

- Bell Canada
- Ciena
- Verizon

2 Abstract

This standard defines Edge Computing Infrastructure as a Service (IaaS) and its Building Blocks and Service Attributes whose values are agreed to by a Cloud Service Provider and a Cloud Subscriber.

3 Terminology and Abbreviations

This section defines the terms used in this standard. In many cases, the normative definitions to terms are found in other standards. In these cases, the third column is used to provide the reference that is controlling, in other Mplify or external standards.

In addition, terms defined in MEF 61.1 [5], MEF 68 [7], MEF 69.1 [8], and MEF 118 [11] are included in this standard by reference and are not repeated in the table below.

Term	Definition	Reference
Affinity	Virtual Machines located on the same host	This document
Anti-affinity	Virtual Machines located on different hosts	This document
Availability Zone	A cluster of multiple Edge Locations equipped with independent power, cooling, and network infrastructure housed in separate physical facilities.	This document
Average Percent Utilization	The average percentage resource usage, calculated over the Calculation Interval, of a particular Building Block.	This document
Building Block	One of the functions that make up an Edge IaaS offering	This document
Building Block Identifier String	An identifier for each building block that is unique within the Edge IaaS instance	This document
Cache	An auxiliary memory from which high-speed retrieval is possible.	This document
Calculation Interval	The interval over which the Performance Service Attributes are determined. The Calculation Interval is always $\geq$ to the Measurement Period and an integer multiplier of the Measurement Period.	This document
Clock Speed	The rate at which a computer's central processing unit (CPU) executes instructions.	This document
Computing	Computing is any goal-oriented activity requiring, benefiting from, or creating computing machinery.	This document
Container Engine	Performs virtualization enabling the ability to run multiple isolated container instances on the OS.	This document
Container Orchestrator	An orchestrator that automatically provisions, deploys, scales, and manages containerized applications without worrying about the underlying infrastructure.	This document
Cores	Separate processing units on a microprocessor.	This document
Data Sovereignty	The ability of the Subscriber to ensure that their data remains within a specified jurisdiction e.g., data must stay within a country or in a building.	This document
Dedicated Edge Computing IaaS Offering	IaaS offerings that belong to one and only one Subscriber.	This document
Edge Computing Subscriber IaaS	A type of Edge Computing Service that delivers processing, storage, networks, and other computing resources. Used interchangeably with Edge IaaS.	This document

Edge IaaS	A type of Edge Computing Service that delivers processing, storage, networks, and other computing resources.	This document
Ephemeral	The temporary storage of data that will be lost if the system restarts or a process ends.	This document
Historical Report Time Interval	The maximum interval for which historical data is reported. Example values are in milliseconds, seconds, days, weeks and months.	This document
Host OS	Identifies any OS functioning as the host within an Edge IaaS offering.	This document
Hypervisor	A type of computer software, firmware or hardware that creates and runs virtual machines.	This document
Infrastructure as a Service	A cloud service provided by the Cloud Service Provider to the Cloud Subscriber that delivers processing, storage, networks, and other fundamental computing resources where the Subscriber is able to deploy and run arbitrary software.	Adapted from NIST SP 800-145 [14]
Installation Place	A place on Earth, which may or may not be fixed, described using one or more Installation Place Representations.	Mplify 150 [13]
Instructions per Clock Cycle	The number of machine instructions executed per clock cycle supported by the processor	This document
Maximum Percent Utilization	The maximum percentage of the Processor that is being used by the Subscriber over the Calculation Interval.	This document
Measured Percent Utilization	The measured processor utilization percentage.	This document
Persistent	The storage of data that will not be lost if the system restarts or a process ends.	This document
Report Interval	The interval over which the Performance Service Attributes are reported. The Report Interval is $\geq$ to the Calculation Interval and an integer multiplier of the Calculation Interval.	This document
Report Start Time	The time that the report started in UTC.	This document
Reported Average Percent Utilization	The average for each Calculation Interval included in the Report Interval.	This document
Reported Maximum Percentage Utilization	The maximum for each Calculation Interval included in the Report Interval.	This document
Service Level Specification	The technical details of the service level, in terms of performance objectives, agreed between the Service Provider and the Subscriber.	This document derived from MEF 61.1 [5]

Service Provider	An organization that provides services to Subscribers. In this document, “Service Provider” means “Cloud Service Provider”.	MEF 10.4 [4]
Shared Edge Computing IaaS Offering	An IaaS solution that is shared among multiple Subscribers, that includes a Virtualization layer to allow the sharing of Processor, Memory, Storage, and Access Network	This document
Storage	Storage defines the storage resource Building Block Service Attributes including Amount, Replication, Storage Duration, and Performance	This document
Storage Duration	The length of time, either Ephemeral or Persistent, that data is stored for.	This document
Subscriber	The end-user of a service. In this document, “Subscriber” should be read as meaning “Cloud Subscriber”.	MEF 10.4 [4]
Subscriber Management Access	Provides limited access to manage Edge IaaS Building Blocks like Container Engines, Hypervisors, Operating Systems, and Hardware Platforms.	This document
Threads	Sequences of instructions given to the CPU.	This document
Uniform Resource Identifier	A unique sequence of characters that identifies a logical or physical resource	This document
Virtualization	Technology that you can use to create virtual representations of servers, storage, networks, and other physical machines	This document
Virtual Machine	In computing, a virtual machine is the virtualization or emulation of a computer system.	This document

Table 1 – Terminology

Term	Definition	Reference
CoS	Class of Service	This document
CSP	Cloud Service Provider	This document
DPU	Data Processing Unit	This document
GPU	Graphics Processing Unit	This document
IaaS	Infrastructure as a Service	NIST SP 800-145 [14]
IdP	Identify Provider	MEF 118.1 [11]
IPC	Instructions Per Clock Cycle	This document
OS	Operating System	This document
PM	Performance Monitoring	This document
QoS	Quality of Service	This document
SA	Service Attribute	This document
SLS	Service Level Specification	This document
SP	Service Provider	This document
SP VC	Service Provider Virtual Connection	This document
URI	Uniform Resource Identifier	This document
UTC	Coordinated Universal Time	This document
UTF-8	Unicode Transformation Format – 8-bit	This document

Table 2 – Abbreviations

4 Compliance Levels

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**NOT RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this standard are to be interpreted as described in BCP 14 (RFC 2119 [1], RFC 8174 [3]) when, and only when, they appear in all capitals, as shown here. All key words must be in bold text.

Items that are **REQUIRED** (contain the words **MUST** or **MUST NOT**) are labeled as [Rx] for required. Items that are **RECOMMENDED** (contain the words **SHOULD** or **SHOULD NOT**) are labeled as [Dx] for desirable. Items that are **OPTIONAL** (contain the words **MAY** or **OPTIONAL**) are labeled as [Ox] for optional.

A paragraph preceded by [CRa]< specifies a conditional mandatory requirement that **MUST** be followed if the condition(s) following the "<" have been met. For example, "[CR1]<[D38]" indicates that Conditional Mandatory Requirement 1 must be followed if Desirable Requirement 38 has been met. A paragraph preceded by [CDb]< specifies a Conditional Desirable Requirement that **SHOULD** be followed if the condition(s) following the "<" have been met. A paragraph preceded by [COc]< specifies a Conditional Optional Requirement that **MAY** be followed if the condition(s) following the "<" have been met.

5 Numerical Prefix Conventions

This standard uses the prefix notation to indicate multiplier values as shown in Table 3.

Decimal		Binary	
Symbol	Value	Symbol	Value
k	10 ³	Ki	2 ¹⁰
M	10 ⁶	Mi	2 ²⁰
G	10 ⁹	Gi	2 ³⁰
T	10 ¹²	Ti	2 ⁴⁰
P	10 ¹⁵	Pi	2 ⁵⁰
E	10 ¹⁸	Ei	2 ⁶⁰
Z	10 ²¹	Zi	2 ⁷⁰
Y	10 ²⁴	Yi	2 ⁸⁰

Table 3 – Numerical Prefix Conventions

6 Introduction

This Mplify standard defines Edge Computing Infrastructure as a Service (IaaS), its Building Blocks, and its attributes such that the buyer (i.e. the Subscriber) of the Edge Computing IaaS Service can compare the different Service offerings using a common set of terminology and definitions of functionalities. This document defines IaaS as a cloud service provided by the Cloud Service Provider to the Cloud Subscriber that delivers processing, storage, networks, and other fundamental computing resources where the Subscriber is able to deploy and run arbitrary software.

The Cloud Services Architecture defined in MEF 68 [7] is used as the base for Edge Computing Subscriber IaaS described within this standard. The scope of this standard is limited to the Service Attributes whose values are agreed on by the Cloud Subscriber and the Cloud Service Provider. Cloud Service Provider is defined in MEF 68 [7] as an organization that provides Cloud Services to Cloud Subscribers. This includes hyperscalers, telcos, or other organizations providing Cloud Services. A Cloud Subscriber is defined in MEF 68 [7] as the end-user (i.e., a person or organization) that maintains a business relationship with and uses Cloud Services from a Cloud Service Provider. Within this standard the term Service Provider (SP) indicates a Cloud Service Provider, and the term Subscriber indicates a Cloud Subscriber.

Edge Computing IaaS is made up of several components, Figure 1 shows the architectural components covered by the Building Block Service Attributes in this standard.




Figure 1 – Edge Computing IaaS Components

These components are derived from MEF 68 [6] as follows:

- Cloud UNI - the demarcation point marking the boundary between the responsibility of the Cloud Subscriber and the Cloud Service Provider.
- Edge Computing IaaS Cloud VC End Point - a logical construct at a Cloud UNI, or Cloud Application Interface to which a subset of the Cloud Service Protocol Data Units (PDUs) passing over that interface is mapped.
- Edge Computing IaaS Cloud VC - an association of two or more Edge Computing IaaS Cloud VC End Points.
- Cloud Application Interface - the demarcation point between a Cloud Application and one or more Edge Computing IaaS Cloud VC End Points. This represents the other termination point for the Edge Computing IaaS Cloud VC.
- Edge Cloud Application - a component of a Cloud Service such as compute, memory, storage, or software applications, that does not provide connectivity to the Subscriber.

The specific Subscriber's requirements may not be met unless the Processor, Memory, Storage, and Access Network are located close to the Subscriber. This is known as Edge Computing Subscriber IaaS. Edge Computing Subscriber IaaS is defined as a Digital Service that provides Edge IaaS capabilities with attribute values that cannot be met with a traditional cloud service and in which Cloud and computing resources may be distributed to meet Subscriber requirements.

Subscriber requirements that make the use of Edge Computing IaaS desirable include:

- Data Sovereignty – the ability of the Subscriber to ensure that their data remains within a specified jurisdiction e.g., data must stay within a country or in a building.
- Access Network Performance – the performance of the Access Network defined by metrics such as Packet Delay, Inter-Packet Delay Variation, and Packet Loss Ratio

The SP may only be able to meet the Subscriber's requirements by implementing an Edge Computing IaaS.

In addition to the Cloud Application per MEF 68 [7] that provides virtualization, this standard also defines attributes that describe the behavior that the Subscriber expects from the service. This includes defining Service Level Specifications (SLs) for the Cloud Application (i.e., Compute, Memory, and Storage areas) along with Cloud Application Interface. Within this document virtualization is defined as technology that you can use to create virtual representations of servers, storage, networks, and other physical machines.

In addition to the Service Attributes described in section 8, the Building Blocks that make up the Edge Computing Subscriber IaaS are detailed in section 7.

Priority of applications is supported on the Cloud UNI and Edge Computing IaaS Cloud VC. The following bullets describe how priority is given:

- Cloud UNI needs to be application-aware
 - How each application is identified must be agreed to by the Subscriber and Cloud SP
- Cloud Application Interface needs to be application-aware
 - How each application is identified must be agreed to by the Subscriber and Cloud SP

7 Edge Computing IaaS Building Blocks

Edge Computing IaaS offerings are made up of a number of Building Blocks. Building Blocks are defined as one of the functions that make up an Edge IaaS offering. These Building Blocks are combined in various ways to create different Edge IaaS offerings. The Building Blocks included in this standard are:

- Processor – see section 7.1
- Memory – see section 7.2
- Storage – see section 7.3
- Access Network – see section 7.4
- Host OS – see section 7.5
- Hypervisor – see section 7.6
- Guest OS – see section 7.7
- Container Engine – see section 7.8
- Virtual Machine – see section 7.9

Each of the Building Blocks is defined in the following sections.

7.1 Processor Building Block

The Processor Building Block identifies the processing unit capability of the Edge IaaS offering. Processor Service Attribute Values define the Processing Unit Building Block Service Attributes such as Clock Speed, number of Cores, number of Threads, amount of Cache, Power requirements, and Instructions Per Clock Cycle (IPC) or may specify a manufacturer and model number of a Processing Unit that brings with it certain values for those Building Block Service Attributes as shown in section 8.6. Cache is defined as an auxiliary memory from which high-speed retrieval is possible. Clock Speed is defined as the rate at which a computer's central processing unit (CPU) executes instructions. Cores are defined as separate processing units on a microprocessor. Threads are defined as sequences of instructions given to the CPU.

7.2 Memory Building Block

The Memory Building Block identifies the memory included in an Edge IaaS offering. Memory defines Service Attributes including the amount as specified in section 8.7. For the purposes of this standard, Memory in this Building Block is defined separately from the Processor Cache Memory in the Processor Building Block.

7.3 Storage Building Block

The Storage Building Block identifies the amount and type of Storage included in an Edge IaaS offering. Storage defines the storage resource Service Attributes including Amount, Replication, Storage Duration, and Performance as specified in section 8.8. Storage maintains a quantity of information and can be created of a certain size for persistent storage of information which can be created and dynamically attached to/detached from Processing. Storage Duration has two options, Ephemeral and Persistent. Ephemeral is defined as the temporary storage of data that will be lost if the system restarts, or a process ends. Persistent is defined as the storage of data that will not be lost if the system restarts, or a process ends. HDD, SSD and Tape storage provide persistent storage

7.4 Access Network Building Block

The Access Network Building Block identifies the connection between the Cloud UNI and the Service Provider Edge IaaS location. The components of the Access Network Building Block are the Cloud UNI, Cloud UNI Access Links, the Subscriber Cloud Virtual Connection (VC), the Cloud Application Interface, and the Edge Computing IaaS Cloud VC End Points. The Service Attributes and the requirements of the Edge IaaS Access Network Building Block are discussed in section 8.9.

7.5 Host OS Building Block

The Host OS Building Block is used to identify if a Host OS is required, and which Host OS is required. The Host OS identifies any OS functioning as the host within an Edge IaaS offering. The Host OS is the operating system on which a Type 2 Hypervisor or Container Engine runs as specified in section 8.11.

7.6 Hypervisor Building Block

The Hypervisor Building Block identifies the type of Hypervisor used for the Edge IaaS offering. The Hypervisor Building Block defines the Hypervisor requirements and Service Attributes as specified in section 8.12. Hypervisors are defined as a type of computer software, firmware or hardware that creates and runs virtual machines. Hypervisors are available as Type 1 and Type 2. A Type 1 Hypervisor does not include a Host OS. A Type 2 Hypervisor always contains a Host OS that is not selectable and may not be published or made known to the Subscriber. In most cases, Hypervisors are described by the vendor's name and version for the Hypervisor. This information infers if there is a Host OS of a specific type used for Type 2 Hypervisors.

7.7 Guest OS Building Block

The Guest OS Building Block identifies the OS functioning as a guest within the Edge IaaS offering. One or more Guest OSs as specified in section 8.13 run on a Type 1 Hypervisor. If a Type 2 Hypervisor is used, the Guest OS runs on the Virtual Machine. Guest OSs provide the OS for the Subscriber applications that are instantiated on the Edge IaaS service.

7.8 Container Engine Building Block

The Container Engine Building Block identifies if there is a Container Engine included in the Edge IaaS. A Container Engine runs on an OS and performs virtualization enabling the ability to run multiple isolated container instances on the OS. A Container Engine supports Containers such as Docker Container Engine or Podman as specified in section 8.14. A Container Engine is defined as a function that performs virtualization enabling the ability to run multiple isolated container instances on the OS. Container Engines always require an OS and can run on either a Host OS or Guest OS.

7.8.1 Container Orchestrator Building Block

The Container Orchestrator Building Block supports the operation of the Container Engine. A Container Orchestrator, such as Kubernetes, is used to automatically provision, deploy, scale, and manage containerized applications without worrying about the underlying infrastructure. If the Container Orchestrator is provided by the SP, the attributes of the Orchestrator and Orchestration need to be agreed to between the Subscriber and SP. These attributes are outside the scope of this version of the document.

7.9 Virtual Machine Building Block

The Virtual Machine Building Block describes the attributes of the Virtual Machine. A Virtual Machine is defined as in computing, a virtual machine is the virtualization or emulation of a computer system as specified in section 8.15.

7.9.1 Virtual Machine Orchestrator Building Block

The Virtual Machine Building Block supports a Virtual Machine orchestrator. This orchestrator controls the configuration and operation of the Virtual Machine. The Virtual Machine Orchestrator that may be the Container Orchestrator.

7.10 Edge IaaS Offering Examples

Examples of the use of Building Blocks to construct Edge IaaS offerings are shown in this section.




Figure 2 – Building Blocks for Dedicated Offering Examples

Figure 2 shows examples of how the Building Blocks can be combined to provide different Dedicated Edge IaaS offerings. Dedicated Edge IaaS offerings belong to one and only one Subscriber. The Subscriber can add any software to the platform. Dedicated Example 2 shows a Processor, Memory, Storage, and Access Network, a Type 1 Hypervisor, and a Guest OS. The Subscriber can add any software that runs on the Guest OS.

The examples in Figure 2 do not limit what other combinations the Subscriber and SP may agree to.




Figure 3 – Building Blocks for Shared Offering Examples

Figure 3 shows how the Building Blocks can be combined to provide Shared Edge IaaS offerings. Shared Example 1 shows where the Subscriber and SP have agreed on the Processor, Memory, Storage, and Access Network attributes along with a Type 2 Hypervisor, Virtual Machine and Guest OS. The Subscriber can add anything that they desire to that. Shared Example 2 shows where the Subscriber and SP have agreed on Processor, Memory, Storage, and Access Network attributes and a Container Engine. Other Building Blocks can be added to these examples as agreed to between the Subscriber and the SP. The provider virtualization layer shown in the figure does not include any Service Attributes. The provider virtualization encompasses virtualization of the Processor, Memory, Storage, and Access Network and is controlled by the SP.

8 Definition of Edge Computing IaaS Subscriber Service Building Block Service Attributes

The Service Attributes for Edge IaaS are defined in section 8. These Building Block Service Attributes are used to describe the capabilities of an Edge IaaS offering. All Building Blocks have Service Attributes and there will be additional Building Block Service Attributes that do not tie directly to a Building block. As an example, Processor Service Attributes include options like Performance Acceleration while QoS Attributes may not impact just one Building Block.

- [R1]** Dedicated and Shared Edge Computing IaaS offerings **MUST** specify the Processor, Memory, Storage and Access Network Building Blocks.

The Processor, Memory, Storage, and Access Network Building Blocks are always present in an Edge Computing IaaS offering. These Building Blocks are key to the behavior of the Edge Computing IaaS.

- [R2]** When providing a Shared offering, the Service Provider **MUST** provide the Subscriber with information on the Virtual Machine, Host OS, Guest OS, and/or Container that are provided
- [O1]** The Building Blocks other than those specified in [R1] and [R2] and their associated Service Attributes **MAY** be included by the Service Provider in their Edge Computing IaaS offering.
- [R3]** If an optional Building Block is included in the Edge Computing IaaS, the value of the Building Block Enabled Service Attribute **MUST** be *TRUE* for that Building Block.
- [R4]** If the value of the Building Block Enabled Service Attribute is *FALSE*, the values of the remaining Service Attributes for that Building Block **MUST** be *Not Applicable*.
- [R5]** If a Building Block is included in the Edge Computing IaaS offering, it **MUST** support the values for Service Attributes as stated in this document.

8.1 Edge IaaS Identifier

The Edge IaaS Identifier is unique within the Service Provider for the Edge IaaS instance. Building Block Identifiers (described in each Service Attribute list) are associated with an Edge IaaS Identifier.

Service Attribute Name	Description	Possible Values
Edge IaaS Identifier	The unique identifier for the Edge IaaS instance.	String

Table 4 – Edge IaaS Type Building Block Building Block Service Attribute

- [R6] The value of the Edge IaaS Identifier Building Block Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R7] The value of the Edge IaaS Identifier Building Block Service Attribute **MUST** contain no more than 53 UTF-8 characters.
- [R8] The Edge IaaS Identifier **MUST** be unique from all other Edge IaaS Identifiers for a given Cloud Service Provider.

8.2 Edge IaaS Type

The Edge IaaS Type Building Block Service Attribute is used to specify if the Edge IaaS offering is a Dedicated or a Shared as defined in Table 5 . Based on the value of the Edge IaaS Type Service Attribute, other Building Blocks' Service Attributes values will be constrained as indicated in the relevant sections.

Service Attribute Name	Description	Possible Values
Edge IaaS Type	The type of Edge IaaS offering, either Shared or Dedicated. Determines if the Edge IaaS is used by a single Subscriber (Dedicated) or shared among multiple Subscribers (Shared).	One of: <i>Dedicated</i> <i>Shared</i>

Table 5 – Edge IaaS Type Building Block Service Attribute

The Edge IaaS is dedicated to a single Subscriber when the value of the Edge IaaS Type Building Block Service Attribute is Dedicated (described as Dedicated throughout this standard). When the value of the Edge IaaS Type Service Attribute is Shared (described as Shared throughout this standard), Virtualization hides the attributes of the Hardware Platform from the Subscriber.

8.3 Edge IaaS Installation Place

The Edge IaaS Installation Place defines where the Edge IaaS instance is located at the Subscriber premises using Installation Place terminology from Mplify 150 [13]. This allows the place to be identified using multiple formats at one time.

Service Attribute Name	Description	Possible Values
Edge IaaS Installation Place	The Installation Place Identifier or Representation of the Edge IaaS instance when the Edge IaaS instance is installed at the Subscriber premises	<i>String</i> that is an Installation Place Identifier or Representation or <i>None</i>

Table 6 – Edge IaaS Installation Place Building Block Service Attribute

[R9] The value of the Edge IaaS Installation Place Service Attribute **MUST** be *None* when the Edge IaaS instance is not installed at the Subscriber premises.

8.4 Building Block State

The Building Block State Service Attribute is used to indicate if a specific Building Block is included in the Edge IaaS offering or not. The values of this Service Attribute is either *Enabled* (when the Service Attribute is included in the Edge IaaS offering) or *Disabled* (when the Building Block is not included in the Edge IaaS offering). This value only applies to Building Blocks other than the four mandatory Building Blocks. This implies that the four mandatory Building Blocks are always enabled. If the Building Block State is *Disabled*, then the remaining Building Block Service Attributes are not specified.

8.5 Building Block Identifier String

A Building Block Identifier String is used to uniquely identify the Building Block instance within the Edge IaaS instance.

[R10] A Building Block Identifier String **MUST** be one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.

[R11] A Building Block Identifier String **MUST** contain no more than 53 UTF-8 characters.

For example a Processor Building Block Identifier value could be *A123*, and the Memory Building Block Identifier value could be *B456*.

8.6 Processor Building Block Service Attributes

The Processor Service Attributes in this section are only used for a physical processor as summarized in Table 7 .

Service Attribute Name	Summary Description	Possible Values
Processor Building Block Identifier	The identifier of the Processor Building Block	Unique <i>Building Block Identifier String</i>

Service Attribute Name	Summary Description	Possible Values
Processor Manufacturer Identifier	The name of the processor manufacturer	<i>String</i> identifying the manufacturer or <i>Unspecified</i>
Manufacturer Model Identifier	Manufacturer model identifier	<i>String</i> identifying a model number of the manufacturer or <i>Unspecified</i>
Processor Architecture	Specifies the architecture and machine language of the processor.	One of: <i>X86</i> <i>ARM</i> <i>GPU</i> <i>Other</i> <i>Unspecified</i>
Clock Speed Range	Specifies the minimum and maximum clock speed of the processor	2-tuple $\langle csmin, csmax \rangle$ of real numbers > 0 each in GHz or <i>Unspecified</i>
Number of Cores	Specifies the number of cores supported by the processor	An integer > 0 or <i>Unspecified</i>
Number of Threads	Specifies the number of threads supported by the processor	An integer > 0 or <i>Unspecified</i>
Level 1 Cache Size	Specifies the size of the level 1 cache supported by the processor	An integer > 0 in Mbytes or <i>Unspecified</i>
Level 2 Cache Size	Specifies the size of the level 2 cache supported by the processor	An integer > 0 in Mbytes or <i>Unspecified</i>
Instructions Per Clock Cycle	Specifies the number of machine instructions executed per clock cycle supported by the processor	An integer > 0 in thousands or <i>Unspecified</i>
Quantity	The number of processors in an Edge Computing IaaS instance	An integer > 0

Table 7 – Processor Building Block Service Attributes

[R12] When the value of the Processor Manufacturer Identifier Service Attribute is *String* and the value of the Manufacturer Model Identifier Service Attribute is

String, then the values for the remaining Processor Service Attributes except for the value of the Quantity Service Attribute **MUST** equal *Unspecified*

- [R13] If the value of the Processor Manufacturer Identifier Service Attribute is *String*, and the value of the Manufacturer Model Identifier is *Unspecified*, the value of the remaining Service Attributes except for the value of Quantity **MUST** be *Unspecified*.
- [R14] If the value of the Processor Manufacturer Identifier Service Attribute is *Unspecified*, the value of the Manufacturer Model Identifier Service Attribute **MUST** be *Unspecified*.
- [R15] When the value of the Processor Manufacturer Identifier Service Attribute is *Unspecified* and the value of the Manufacturer Model Identifier Service Attribute is *Unspecified*, then the values for the remaining Processor Service Attributes **MUST NOT** be *Unspecified*

The Quantity Service Attribute is used to indicate the number of Processors that exist in an Edge Computing IaaS instance.

8.6.1 Manufacturer Identifier Service Attribute

The Processor Manufacturer Identifier Service Attribute is used to specify the name of the Processor Manufacturer. Its value is either a string identifying the Processor Manufacturer or *Unspecified*.

- [R16] If the value of the Processor Manufacturer Identifier Service Attribute is not *Unspecified*, then its value **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R17] If the value of the Processor Manufacturer Identifier Building Block Service Attribute is not *Unspecified*, then its value **MUST** contain no more than 53 UTF-8 characters.
- [R18] If the Edge IaaS Type Building Block Service Attribute is *Shared*, the Manufacturer Identifier Service Attribute **MUST** be set to *Unspecified*.

Requirement [R18] means that the value of the Manufacturer Identifier Service Attribute is determined by the Service Provider when the value of the Edge IaaS Type Service Attribute is *Shared*.

8.6.2 Manufacturer Model Identifier Service Attribute

When the value of the Processor Manufacturer Identifier Service Attribute is specified, the Manufacturer Model Identifier Service Attribute is either a string to identify a model number of the manufacturer or *Unspecified*.

- [R19] If the value of the Manufacturer Model Identifier Service Attribute is not *Unspecified*, then its value **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R20] If the value of the Manufacturer Model Identifier Service Attribute is not *Unspecified*, then its value **MUST** contain no more than 53 UTF-8 characters.

Note: it is possible that the value of the Manufacturer Model Identifier Service Attribute can be *Unspecified* even when the value of the Processor Manufacturer Identifier Service Attribute is a string identifying a manufacturer.

8.6.3 Processor Architecture Service Attribute

The Processor Architecture Service Attribute is used to specify the architecture and machine language of the processor. Its value is one of:

- *X86*
- *ARM*
- *Other*
- *Unspecified*

If the value is *Other*, the Subscriber and SP need to agree on the architecture and machine language of the processor.

- [R21] If the value of the Manufacturer Model Identifier Service Attribute is not *Unspecified*, then the value of the Processor Architecture Service Attribute **MUST** be *Unspecified*.
- [R22] If Processor Architecture value is *Unspecified*, then the Processor Manufacturer Identifier **MUST NOT** have values of *Unspecified*.

Note that the implication of [R19] and [R22] is that a given processor product from a given manufacturer will have a known processor architecture.

8.6.4 Clock Speed Max Service Attribute

The Clock Speed Max Service Attribute is used to specify the maximum clock speed of the processor. The value of the Clock Speed Max Service Attribute is either a real number > 0 in GHz or *Unspecified*.

Note that a given processor product from a given manufacturer will have a known clock speed.

8.6.5 Clock Speed Min Service Attribute

The Clock Speed Min Service Attribute is used to specify the minimum clock speed of the processor. The value of the Clock Speed Min Service Attribute is either a real number > 0 in GHz or *Unspecified*.

- [D1] If the value of Clock Speed Max Service Attribute is not *Unspecified*, then the value of Clock Speed Min Service Attributes **SHOULD** be a real number >0 in GHz.
- [R23] If the value of Clock Speed Min Service Attribute is not *Unspecified*, the value **MUST** be $\leq$ to the value of Clock Speed Max Service Attribute.

Note that a given processor product from a given manufacturer will have a known clock speed.

8.6.6 Number of Cores Service Attribute

The Number of Cores Service Attribute is used to specify the number of cores supported by the processor. Its value is either an integer > 0 or *Unspecified*.

- [R24] If the value of the Manufacturer Model Identifier Service Attribute is not *Unspecified*, then the value of the Number of Cores Service Attribute **MUST** be *Unspecified*.

When the value of the Edge IaaS Type Service Attribute = *Dedicated*, a given processor product from a given manufacturer has a known number of cores.

8.6.7 Number of Threads Service Attribute

The Number of Threads Service Attribute is used to specify the number of threads that the processor supports. Its value is an integer > 0 or *Unspecified*.

- [R25] If the value of the Manufacturer Model Identifier Service Attribute is not *Unspecified*, then the value of the Number of Threads Service Attribute **MUST** be *Unspecified*.

When the value of the Edge IaaS Type Service Attribute = *Dedicated*, a given processor product from a given manufacturer has a known number of threads.

8.6.8 Level 1 Cache Size Service Attribute

A Level 1 Cache is a memory cache that is directly built into the Processor that is used for storing the Processor's recently accessed information. Level 1 Cache is not related to Memory as defined in section 8.7

The Level 1 Cache Size Service Attribute is used to specify the size of the level 1 cache that the processor supports. Its value is an integer > 0 in bytes or *Unspecified*.

- [R26] If the value of the Manufacturer Model Identifier Service Attribute is not *Unspecified*, then the value of the Level 1 Cache Size Service Attribute **MUST** be *Unspecified*.

When the value of the Edge IaaS Type Service Attribute = *Dedicated*, a given processor product from a given manufacturer has a known L1 cache size.

8.6.9 Level 2 Cache Size Service Attribute

A Level 2 Cache is a cache memory that is located outside of the Processor chip core, although it is typically found on the same chip package. Level 2 Cache is not related to the Memory defined in section 8.7.

The Level 2 Cache Size Service Attribute is used to specify the size of the Level 2 Cache that the processor supports. Its value is an integer > 0 in bytes or *Unspecified*.

[R27] If the value of the Manufacturer Model Identifier Service Attribute is not *Unspecified*, then the value of the Level 2 Cache Size Service Attribute **MUST** be *Unspecified*.

When the value of the Edge IaaS Type Service Attribute = *Dedicated*, a given processor product from a given manufacturer has a known Level 2 Cache size.

8.6.10 Instructions per Clock Cycle Service Attribute

The Instructions per Clock Cycle Service Attribute is used to specify the number of machine instructions executed per clock cycle by the processor. Its value is an integer > 0 or *Unspecified*.

[R28] If the value of the Manufacturer Model Identifier Service Attribute is not *Unspecified*, then the value of the Instructions per Clock Cycle Service Attribute **MUST** be *Unspecified*.

When the value of the Edge IaaS Type Service Attribute = *Dedicated*, a given processor product from a given manufacturer has a known number of Instructions per Clock Cycle.

8.6.11 Quantity Service Attribute

The Quantity Service Attribute is used to specify the number of processors that are included in the Edge Computing IaaS.

[R29] The value of the Quantity Service Attribute **MUST** be an integer > 0 .

The implication of [R29] is that the Processor Building Block is always present.

8.7 Memory Building Block

The Memory Building Block comprises two Service Attributes that are only used for physical Memory. They are summarized in Table 8 and defined in this section.

Service Attribute Name	Summary Description	Possible Values
Memory Building Block Identifier	The identifier of the Memory Building Block	Unique within the Edge IaaS Offering <i>String</i>

Service Attribute Name	Summary Description	Possible Values
Memory Amount	The amount of memory	An integer > 0 in GBytes

Table 8 – Memory Resource Service Attributes

8.7.1 Memory Amount Service Attribute

The value of the Memory Amount Service Attribute in the Edge IaaS Instance is > 0 GBytes.

[R30] Regardless of the values of the Processor attributes, the Memory Amount Service Attribute **MUST** be specified.

The implication of 8.7.1 is that the Memory Building Block is always present.

8.8 Storage Building Block

The Storage Building Block comprises six Service Attributes that are only used for physical Storage. They are summarized in Table 9 and defined within this section.

Note: In Storage, the terms Archival and Replication are used. These represent two different methods of storing the Subscriber data. When data is archived, it resides on a storage medium, sometimes located at a different location than other stored data, and it may take an extended time, such as days, to retrieve the data from the archive. Replicated data is a copy of the Subscriber data that is stored and can be accessed in real-time. Often replicated data is stored in a different location from the Subscriber data. The Storage Building Block is always Persistent storage.

Service Attribute Name	Summary Description	Possible Values
Storage Building Block Identifier	The identifier of the Storage Building Block	Unique identifier within the Edge IaaS Offering <i>String</i>
Storage Type	The type of storage	One of: <i>File</i> <i>Block</i> <i>Object</i>
Storage Retention Duration	The duration that data is retained	An integer > 0 days
Storage Time to Access	The time to access the Subscriber data	An integer > 0 in milliseconds
Storage Amount	The amount of storage available.	A real number > 0 in GBytes
Storage Replication	Indicates if replication of stored data is supported within an Edge Availability Zone.	One of: <i>No</i> <i><Yes, Installation Place></i>

Table 9 – Storage Resource Service Attributes

The day described in Storage Retention Duration is defined as midnight to 11:59 PM in the same time zone as the Edge IaaS location.

8.8.1 Storage Type Service Attribute

The Storage Type Service Attribute is used to specify the type of storage. Its value is either *File*, *Block*, or *Object*.

8.8.2 Storage Retention Duration Service Attribute

The Storage Retention Duration Service Attribute is used to specify the number of days that the data is retained. Its value is an integer > 0 in days.

8.8.3 Storage Time to Access Service Attribute

The Storage Time to Access Service Attribute is used to specify the time required to access the Subscriber data that has been stored. Its value is an integer > 0 in milliseconds.

8.8.4 Storage Amount Service Attribute

The Storage Amount Building Block Service Attribute is used to specify the amount of storage available. Its value is an integer > 0 in GBytes.

[R31] The Storage Amount **MUST** be specified.

The implication of 0 is that the Storage Building Block is always present.

8.8.5 Replication Service Attribute

The Replication Service Attribute is used to specify if the stored data is replicated by the Service Provider. Its value is either *No* or *<Yes, Installation Place>*.

8.9 Access Network Building Block

Access Network Building Block is made up of the Cloud UNI and Cloud UNI Access Links, the Edge Computing IaaS Cloud Virtual Connection (VC), the Edge Computing IaaS VC End Points, and the Cloud Application Interface. Each of these plays a role in the connection between the Subscriber and the IaaS instance. The Access Network Building Block is provided by the Cloud SP. Connectivity from the Subscriber location to the Cloud UNI may be provided by the Subscriber, or a 3rd Party.

The values of the Cloud VC Service Attributes are set to match the type of connection between the Cloud UNI and the Cloud Application Interface. This connection can be as simple as local LAN connection or could be via the Internet or other backhaul technology (WAN connection). When it is a WAN connection, the values of the Service Attributes are significant. When it is a local LAN connection, the values of the Service Attributes may not be significant.




Figure 4 – Access Network Building Block Implementation Example

Figure 4 shows an example of the Access Network Building Block and its use to connect the Cloud Subscriber to the Cloud Service Provider infrastructure. In this example, there are two Subscriber’s applications (e.g., Point of Sale and video surveillance) that are resident on a single IaaS instance. Each Subscriber application is associated with a single Cloud Application Interface and the multipoint Cloud VC connects the Cloud UNI to each of the Cloud Application Interfaces. There are other possible implementations that are not included in the examples in this document.

The Access Network Building Block provides connectivity using the Internet Protocol (IP) and is defined using the IP Service Attributes in MEF 61.1 [5], MEF 61.1.1 [6], and this document. Table 10 summarizes the 5 Access Network Building Block components and indicates the sub-section where the detailed Service Attributes for each are described.

Building Block Components	Component Description	Reference
Cloud UNI	The demarcation point marking the boundary between the responsibility of the Cloud Subscriber and the Cloud Service Provider.	See Table 11
Cloud UNI Access Link	An individual connection between the Cloud Subscriber and the Cloud SP that forms part of a Cloud UNI.	See section 8.9.3
Edge Computing IaaS Cloud VC	An association of two or more Edge Computing IaaS Cloud VC End Points.	See Table 12
Edge Computing IaaS Cloud VC End Points	A logical construct at a Cloud UNI or Cloud Application Interface to which a subset of the Edge IaaS PDUs passing over that interface is mapped.	See section 8.9.5
Cloud Application Interface	The demarcation point between the Edge IaaS instance and one or more Edge Computing IaaS Cloud VC End Points.	See Table 13

Table 10 – Access Network Building Block Components

8.9.1 Subscriber Application Awareness Service Attribute

Nested within the Edge Computing IaaS application are one or more Subscriber applications. All of these are considered Cloud Applications as defined in MEF 68. Subscriber Application Awareness is the ability to identify a specific Subscriber Application (e.g., Point of Sale or video surveillance). The methods used to identify Subscriber Applications are beyond the scope of this document but must be agreed upon by the Subscriber and Cloud SP. Subscriber Application Awareness is a capability that is represented by the Subscriber Application Awareness Service Attribute that is either *True* or *False*. If the value is *True* it applies to at least the Cloud UNI and Cloud Application Interface. It may apply to the Edge Computing IaaS Cloud VC End Points. If the value is *False*, then Subscriber Application Awareness does not apply.




Figure 5 – Subscriber Application Awareness Implementation Example

In the example, functionality exists on the Cloud SP side of the Cloud UNI and at the Cloud Application Interface that determines the Subscriber Application. Ingress IP Packets are received from the Cloud Subscriber at the Cloud UNI and egress IP Packets are received at the Cloud Application Interface. The Subscriber Application is identified by the Cloud SP.

The value of the Cloud Application Awareness Service Attribute is either *TRUE* or *FALSE* and is always provided.

- [R32]** If the Cloud Application Awareness Service Attribute value is *TRUE*, then both the Cloud UNI and the Cloud Application Interface **MUST** be Subscriber application-aware.
- [O2]** If the Cloud Application Awareness Service Attribute value is *TRUE*, the Cloud VC **MAY** be Subscriber application-aware.

8.9.2 Cloud UNI

The Cloud UNI is the demarcation point marking the boundary between the responsibility of the Cloud Subscriber and the Cloud Service Provider. The Cloud UNI can be Subscriber application-aware as defined in section 8.9.1 and can prioritize traffic based on the Subscriber Application. It can also prioritize IP Packets based on DSCP markings received from the Subscriber.

8.9.2.1 Cloud UNI Service Attributes

The following MEF 61.1/61.1.1 UNI Service Attributes, their values, and associated requirements are applicable to a Cloud UNI:

- UNI Identifier
- List of UNI Access Links

- UNI Ingress Bandwidth Profile Envelope
- UNI Egress Bandwidth Profile Envelope
- UNI List of Control Protocols
- UNI Routing Protocols
- UNI Reverse Path Forwarding

This document does not impose any additional constraints on the MEF 61.1 and 61.1.1 Service Attribute values nor specify any additional requirements.

Additional attributes that are specific to the Cloud UNI that are not included in MEF 61.1 or 61.1.1 are shown in the following table.

Service Attribute Name	Summary Definition	Possible Values
Cloud UNI Subscriber Application Awareness	Indicates if the Cloud UNI is Subscriber application-aware.	One of: • <i>TRUE</i> • <i>FALSE</i>
Cloud UNI Subscriber Application Identifier Method	The method used to identify the Subscriber Application.	A list of one or more of the following: • <i>Source IP Address</i> • <i>Destination IP Address</i> • <i>Source Port Number</i> • <i>Destination Port Number</i> • <i>DSCP Value</i> • <i>AppID</i>
Cloud UNI Subscriber Application Policy	The definition of a Cloud Subscriber Application policy as agreed to by the Subscriber and SP.	One of the following per application: • <i>No Policy</i> • <i>Policy definition</i>

Table 11 – Cloud UNI Subscriber Application Awareness Service Attributes

- [R33]** If the value of Cloud UNI Subscriber Application Awareness is *TRUE*, Cloud UNI Subscriber Application Identifier Method and Cloud UNI Subscriber Application Policy values **MUST** be provided.
- [R34]** If the value of Cloud UNI Subscriber Application Awareness is *FALSE*, Cloud UNI Subscriber Application Identifier Method and Cloud UNI Subscriber Application Policy values **MUST NOT** be provided.

Requirement [R34] means that when the value of the Cloud UNI Subscriber Application Awareness is *FALSE*, the other two Service Attributes and their values in Table 11 are not specified.

8.9.3 Cloud UNI Access Links

A Cloud UNI Access Link is defined as an individual connection between the Cloud Subscriber and the Cloud SP that forms part of a Cloud UNI. The concept of Cloud UNI Access Links is derived from MEF 61.1 [5].

8.9.3.1 Cloud UNI Access Links Service Attributes

The following MEF 61.1 and 61.1.1 UNI Access Link Service Attributes, their values, and associated requirements are applicable to Cloud UNI Access Links:

- UNI Access Link Identifier
- UNI Access Link Connection Type
- UNI Access Link L2 Technology
- UNI Access Link IPv4 Connection Addressing
- UNI Access Link IPv6 Connection Addressing
- UNI Access Link DHCP Relay
- UNI Access Link Prefix Delegation
- UNI Access Link BFD
- UNI Access Link IP MTU
- UNI Access Link Ingress Bandwidth Profile
- UNI Access Link Egress Bandwidth Profile
- UNI Access Link Reserved VRIDs

This document does not impose any additional constraints on the Service Attribute values nor specify any additional requirements beyond those in MEF 61.1 and 61.1.1

There are no additional attributes for a Cloud UNI Access Link.

8.9.4 Edge Computing IaaS Cloud VC Service Attributes

The Service Attributes for the Edge Computing IaaS Cloud VC are defined in this section of the document. This includes Service Attributes that apply if the Edge Computing IaaS Cloud VC is Subscriber application-aware.

Edge Computing IaaS Cloud VC Service Attributes are derived from Subscriber IPVC Service Attributes in MEF 61.1 [5]. Additional Service Attributes have been added to support Subscriber Application Awareness. Cloud Application Priority can be identified using any method agreed to by the Cloud Subscriber and the Cloud Service Provider (e.g. H, M, L, or Bronze, Silver, Gold)

The following MEF 61.1 and 61.1.1 IPVC Service Attributes, their values, and associated requirements are applicable to Edge Computing IaaS Cloud VCs:

- IPVC Identifier
- IPVC Topology (always Cloud Access)
- IPVC End Point List
- IPVC Packet Delivery
- IPVC Maximum Number of IPv4 Routes
- IPVC Maximum Number of IPv6 Routes
- IPVC DSCP Preservation
- IPVC List of Class of Service Names
- IPVC Service Level Specification
- IPVC MTU
- IPVC Path MTU Discovery
- IPVC Fragmentation
- IPVC Cloud
- IPVC Reserved Prefixes

[R35] The value of IPVC Topology **MUST** be Cloud Access.

Additional attributes that are specific to the Edge Computing IaaS Cloud VC that are not included in MEF 61.1 or 61.1.1 are shown in the following table.

Service Attribute Name	Summary Description	Possible Values
Edge Computing IaaS Cloud VC Subscriber Application Awareness	Indicates if the Cloud VC is Subscriber application-aware.	One of: • <i>True</i> • <i>False</i>
Edge Computing IaaS Cloud VC Application Ingress Policy	The ingress policy used to identify an application.	String
Edge Computing IaaS Cloud VC Application Egress Policy	The egress policy used to identify an application.	String
Edge Computing IaaS Cloud VC Application Priority	The policy that identifies the priority of a given application.	String
Edge Computing IaaS Cloud VC Application Priority to DSCP Mapping	The mapping between an application priority and one or more DSCP values.	String

Table 12 – Edge Computing IaaS Cloud VC Application Awareness Service Attributes

[R36] If the value of Edge Computing IaaS Cloud VC Application Awareness is *TRUE*, the Cloud Application Ingress Policy, Cloud Application Egress Policy, Cloud Application Priority values **MUST** be provided.

[R37] If the Value of Edge Computing IaaS Cloud VC Application Awareness is *FALSE*, the Cloud Application Ingress Policy, the Cloud Application Egress Policy, the Cloud Application Priority values **MUST NOT** be provided.

Requirement [R37] means that when the value of the Edge Computing IaaS Cloud VC Application Awareness is *FALSE*, the other four Service Attributes and their values in Table 12 are not specified.

8.9.5 Edge Computing IaaS Cloud VC End Points

IPVC End Points are defined in MEF 61.1 [5]. Their values, and associated requirements are applicable to Edge Computing IaaS Cloud VC End Points:

- IPVC End Point Identifier
- IPVC End Point EI Type
- IPVC End Point EI
- IPVC End Point Role
- IPVC End Point Prefix Mapping
- IPVC End Point Maximum Number of IPv4 Routes

- IPVC End Point Maximum Number of IPv6 Routes
- IPVC End Point Ingress Class of Service Map
- IPVC End Point Egress Class of Service Map
- IPVC End Point Ingress Bandwidth Profile Envelope
- IPVC End Point Egress Bandwidth Profile Envelope

8.9.6 Cloud Application Interface

A Cloud Application Interface, as described in MEF 68 [7], is the demarcation point between the Edge IaaS instance and one or more Edge Computing IaaS Cloud VC End Points. It performs Cloud UNI-like functionality but provides a demarcation point between the Edge Computing IaaS Cloud VC and the Building Blocks that make up the IaaS.

8.9.6.1 Cloud Application Interface Service Attributes

The Cloud Application Interface Service Attributes are defined in this section.

Cloud Application Interface Service Attributes are derived from UNI Service Attributes in MEF 61.1 [5] with additional attributes defined within this document for Application Awareness implementations.

Service Attribute Name	Summary Definition	Possible Values
Cloud Application Interface Identifier	A string that is used to allow the Cloud Subscriber and Cloud Service Provider to uniquely identify the Cloud Application Interface for operations purposes.	String
Cloud Application Interface Ingress Bandwidth Profile Envelope	Bandwidth Profile Envelope used for an ingress Bandwidth Profile.	<i>None</i> or a set of parameters as specified in MEF 61.1 [5] section 17.3.
Cloud Application Interface Egress Bandwidth Profile Envelope	Bandwidth Profile Envelope used for an egress Bandwidth Profile.	<i>None</i> or a set of parameters as specified in MEF 61.1 [5] section 17.3.
Cloud Application Interface List of Control Protocols	Indication of IP Control Protocols that are not forwarded transparently by the SP.	See MEF 61.1 [5] section 12.6.
Cloud Application Interface Routing Protocols	List of Routing Protocols used across the Cloud Application Interface.	See MEF 61.1 [5] section 12.7.
Cloud Application Interface Reverse Path Forwarding	Indicates whether RPF checks are used by the SP at the Cloud Application Interface.	One of: • <i>Enabled</i> • <i>Disabled</i>
Cloud Application Interface Subscriber Application Awareness	Indicates if the Cloud Application Interface is Subscriber application-aware.	One of: • <i>True</i> • <i>False</i>
Cloud Application Interface Ingress Policy	The ingress policy used to identify an application.	String
Cloud Application Interface Egress Policy	The egress policy used to identify an application.	String
Cloud Application Interface Priority	The policy that identifies the priority of a given application.	String
Cloud Application Interface Priority to DSCP Mapping	The mapping between an application priority and one or more DSCP values.	String

Table 13 – Cloud Application Interface Service Attributes

[R38] If the value of Cloud Application Interface Subscriber Application Awareness Service Attribute is *True*, then values for Cloud Application Interface Ingress Policy, Cloud Application Interface Egress Policy, and Cloud Application Interface Priority Service Attributes **MUST** be provided.

- [R39]** If the value of Cloud Application Interface Subscriber Application Awareness Service Attribute is *False*, then values for Cloud Application Interface Ingress Policy, Cloud Application Interface Egress Policy, and Cloud Application Interface Priority Service Attributes **MUST NOT** be provided.

8.10 Edge IaaS Access Service Attributes

Access by the Subscriber to the Edge IaaS is described by the Edge IaaS Access Service Attributes. These attributes are provided by the SP to the Subscriber and inform the Subscriber how to gain access to their Edge IaaS instance. This access method is not used to access applications added by the Subscriber or for automated API access.

The Cloud SP provides the Subscriber the information listed Table 14 to enable the Subscriber accessing its Edge IaaS instance.

Service Attribute Name	Summary Description	Possible Values
Edge IaaS Access URL	A URL for the Edge IaaS instance with possible values following the URI format specified in Section 2.7.2 of RFC 7230 [2].	String
Edge IaaS Access Identifier	An identifier used for authentication to the Edge IaaS instance.	1-53 Characters
Edge IaaS Access Password	Password used for authentication to the Edge IaaS instance.	1-53 Characters

Table 14 – Edge IaaS Access Service Attributes

A possible flow is to connect to the URL, enter the Edge IaaS Access Identifier and the Edge IaaS Access Password. It is also recommended that 2 factor authentication, like email and password be used prior to access being provided. Another possible flow is to use an Identity Provider (IdP) as the method to gain access. This could be as it is defined in MEF 118.1 [11] or other mechanisms.

8.10.1 Subscriber Management Access

Subscriber Management Access provides limited access to manage Edge IaaS Building Blocks like Container Engines, Hypervisors, Operating Systems, and Hardware Platforms. The access method and requirements to manage the Building Blocks are beyond the scope of this document. It is recommended that these building blocks be secured but the methodology is not defined in this document.

8.10.2 Edge IaaS Access URL

The Edge IaaS Access URL used to access the Edge IaaS.

- [R40]** The value of the Edge IaaS Access URL Service Attribute **MUST** be provided by the Service Provider to the Subscriber.

8.10.3 Edge IaaS Access Identifier

The identifier used by the Subscriber to access the Edge IaaS such as a user name.

- [R41] The value of the Edge IaaS Access Identifier Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R42] The value of the Edge IaaS Access Identifier Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.10.4 Edge IaaS Access Password

The password, used in combination with the Edge IaaS Access Identifier, that allows the Subscriber to access the Edge IaaS.

- [R43] The value of the Edge IaaS Access Password Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R44] The value of the Edge IaaS Access Password Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.11 Host OS

The Host OS Service Attributes are defined in Table 15 within this section.

Service Attribute Name	Summary Description	Possible Values
Host OS Building Block State	Indicates if this Building Block is included in the Edge IaaS	<i>Enabled</i> <i>Disabled</i>
Host OS Building Block Identifier	The identifier of the Host OS Building Block	Unique <i>Building Block Identifier String</i>
Host OS Vendor Identifier	The manufacturer of the Host OS	1–53-character string identifying the manufacturer
Host OS Name	The model of the Host OS	1–53-character string identifying the model
Host OS Version	The version of the Host OS	1–53-character string identifying the version

Table 15 – Host OS Service Attributes

8.11.1 Host OS Vendor Identifier

The name of the Host OS manufacturer.

- [R45] The value of the Host OS Vendor Identifier Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.

- [R46] The value of the Host OS Vendor Identifier Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.11.2 Host OS Name

The Host OS Vendor's name for the OS.

- [R47] The value of the Host OS Name Identifier Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.

- [R48] The value of the Host OS Name Identifier Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.11.3 Host OS Version

The version of the Host OS.

- [R49] The value of the Host OS Version Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.

- [R50] The value of the Host OS Version Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.12 Hypervisor

The Hypervisor Service Attributes are defined in Table 16 within this section.

Service Attribute Name	Summary Description	Possible Values
Hypervisor Building Block State	Indicates if this Building Block is included in the Edge IaaS.	<i>Enabled</i> <i>Disabled</i>
Hypervisor Building Block Identifier	The identifier of the Hypervisor Building Block	<i>Unique Building Block Identifier String</i>
Hypervisor Vendor	The manufacturer of the Hypervisor	1–53-character string identifying the manufacturer
Hypervisor Name	The name of the Hypervisor	1–53-character string identifying the name
Hypervisor Version	The version of the Hypervisor	1–53-character string identifying the version

Table 16 – Hypervisor Service Attributes

8.12.1 Hypervisor Vendor

The name of the Hypervisor manufacturer.

- [R51] The value of the Hypervisor Vendor Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R52] The value of the Hypervisor Vendor Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.12.2 Hypervisor Name

The Hypervisor Vendor’s name for the Hypervisor.

- [R53] The value of the Hypervisor Name Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R54] The value of the Hypervisor Name Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.12.3 Hypervisor Version

The Hypervisor Vendor’s version for the Hypervisor.

- [R55] The value of the Hypervisor Version Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R56] The value of the Hypervisor Version Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.13 Guest OS

The Guest OS Service Attributes are defined in Table 17 within this section.

Service Attribute Name	Summary Description	Possible Values
Guest OS Building Block State	Indicates if this Building Block is included in the Edge IaaS.	<i>Enabled</i> <i>Disabled</i>
Guest OS Building Block Identifier	The identifier of the Guest OS Building Block	Unique <i>Building Block Identifier String</i>
Guest OS Publisher	The publisher of the Guest OS	1–53-character string identifying the publisher
Guest OS Name	The name of the Guest OS	1–53-character string identifying the name
Guest OS Version	The version of the Guest OS	1–53-character string identifying the version

Table 17 – Guest OS Service Attributes

8.13.1 Guest OS Publisher

The name of the Guest OS publisher.

- [R57] The value of the Guest OS Publisher Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R58] The value of the Guest OS Publisher Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.13.2 Guest OS Name

The Guest OS Vendor's name for the OS.

- [R59] The value of the Guest OS Name Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R60] The value of the Guest OS Name Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.13.3 Guest OS Version

The Guest OS Vendor's version for the OS.

- [R61] The value of the Guest OS Version Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R62] The value of the Guest OS Version Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.14 Container Engine

The Container Engine Service Attributes are defined in Table 18 within this section.

Service Attribute Name	Summary Description	Possible Values
Container Engine Building Block State	Indicates if this Building Block is included in the Edge IaaS.	<i>Enabled</i> <i>Disabled</i>
Container Engine Building Block Identifier	The identifier of the Container Engine Building Block	Unique <i>Building Block Identifier String</i>
Container Engine Vendor	The manufacturer of the Container Engine	1–53-character string identifying the manufacturer
Container Engine Name	The name of the Container Engine	1–53-character string identifying the name
Container Engine Version	The version of the Container Engine	1–53-character string identifying the version

Table 18 – Container Engine Service Attributes

8.14.1 Container Engine Vendor

The name of the Container Engine maker.

[R63] The value of the Container Engine Vendor Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.

[R64] The value of the Container Engine Vendor Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.14.2 Container Engine Name

The Container Engine Vendor’s name for the Container Engine.

[R65] The value of the Container Engine Name Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.

[R66] The value of the Container Engine Name Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.14.1 Container Engine Version

The Container Engine Vendor’s version for the Container Engine.

[R67] The value of the Container Engine Version Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.

[R68] The value of the Container Engine Version Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.15 Virtual Machine Requirements

The Virtual Machine Service Attributes are defined in this section.

Service Attribute Name	Summary Description	Possible Values
Virtual Machine Building Block State	Indicates if this Building Block is included in the Edge IaaS.	<i>Enabled</i> <i>Disabled</i>
Virtual Machine Building Block Identifier	The identifier of the Virtual Machine Building Block	Unique <i>Building Block Identifier String</i>
Number of Virtual Processors	The number of virtual Processors included in the Virtual Machine	An integer > 0
Number of Threads	Number of Threads in the virtual Processor	An integer > 0
Number of Cores	Number of Cores in the virtual Processor	An integer > 0
Number of Physical CPUs	Number of physical CPUs that make up the virtual Processor.	An integer > 0
Virtual Memory Amount	Amount of virtual memory	An integer > 0 in GBytes
Host OS	The Host OS used in the Virtual Machine	See section 8.11 for details

Table 19 – Virtual Machine Service Attributes

Note: A Virtual Machine Building Block contains the Virtual CPU.

8.15.1 Number of Virtual Processors

Indicates the number of virtual Processors included in the Virtual Machine.

[R69] If the value of the Virtual Machine Building Block State Service Attribute is *Enabled*, the Number of Virtual Processors **MUST** be specified.

8.15.2 Number of Threads

Indicates the number of threads included in the virtual Processors.

[R70] If the value of the Number of Virtual Processors Service Attribute is >0, the Number of Threads **MUST** be specified.

8.15.3 Number of Cores

Indicates the number of cores included in the virtual Processors.

[R71] If the value of the Number of Virtual Processors Service Attribute is >0, the Number of Cores **MUST** be specified.

8.15.4 Number of Physical CPUs

Indicates the number of physical CPUs included in the Virtual Machine.

- [R72] If the value of the Number of Virtual Processors Service Attribute is >0 , the Number of physical CPUs **MUST** be specified.

8.15.5 Virtual Memory Amount

Indicates the amount of virtual memory included in the Virtual Machine.

- [R73] If the value of the Virtual Machine Building Block State Service Attribute is *Enabled*, the Virtual Memory Amount **MUST** be specified.

8.15.6 Host OS

Indicates the Host OS that is included in the Virtual Machine. See section 8.11 for details.

- [R74] If the value of the Virtual Machine Building Block State Service Attribute is *Enabled*, the Host OS **MUST** be specified as defined in section 8.11.

8.16 Edge IaaS Power Requirements

Edge IaaS Power Requirements Service Attributes are defined in section 7.14 in MEF 130 [12].

- [R75] The value of the Edge IaaS Power Service Attributes **MUST** be specified by the SP when the computing resources are being placed in space controlled by the Subscriber.

8.17 Edge IaaS Physical Characteristics

The Edge IaaS physical characteristics are defined in MEF 130 [12]. When an Edge IaaS is located outside of the SP network, the characteristics specified in MEF 130 [12] apply. The installation of any physical space required, racks, climate control, or other attributes to install the Edge IaaS is beyond the scope of this document.

8.18 Other Service Attributes for Edge IaaS

The Edge IaaS Service Attributes that are not addressed by a given building block are defined in this section.

8.18.1.1 Security Policies

Policies are created that define the operation of entities like Virtual Machines, Container Engines, and Hardware Platforms. Security Framework is defined in MEF 118.1 [11].

- [R76] Security Policies **MUST** be enabled.

Service Attribute Name	Summary Description	Possible Values
Security Policy Identifier	A unique identifier for the Policy	1-53 characters
Security Policy Name	The name of the Policy	1-53 characters
Security Policy Constraints	Security Policy Constraints are typically expressed as a predicate based on the context that the object is operating in, and/or the state of the object.	1-53 characters
Security Policy Subject	The identifier of the entity on which security is being configured.	One of: <i>UNI</i> <i>Application Interface</i>

Table 20 – Security Policy Service Attributes

8.18.1.1.1 Security Policy Identifier

A unique identifier for the Security Policy.

- [R77] The value of the Security Policy Identifier Service Attribute **MUST** be specified.
- [R78] The Security Policy Identifier **MUST** be unique among the values of the Policy Service Attribute for a given Subscriber of Edge IaaS provided by the SP.
- [R79] The value of the Security Policy Identifier Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R80] The value of the Security Policy Identifier Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.18.1.1.2 Security Policy Name

The name of the Security Policy.

- [R81] The value of the Security Policy Name Service Attribute **MUST** be specified.
- [R82] The Security Policy Name **MUST** be unique among the values of the Security Policy Service Attribute for a given Subscriber of Edge IaaS provided by the SP.
- [R83] The value of the Security Policy Name Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R84] The value of the Security Policy Name Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.18.1.1.3 Security Policy Constraints

Security Policy Constraints are typically expressed as a predicate based on the context that the object is operating in, and/or the state of the object.

- [R85] The value of the Security Policy Constraints Service Attribute **MUST** be specified.
- [R86] The value of the Security Policy Constraints Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R87] The value of the Security Policy Constraints Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.18.1.1.4 Security Policy Subject

The identifier of the entity (UNI or Cloud Application Interface) on which security policy is being configured.

- [R88] The value of the Security Policy Subject Service Attribute **MUST** be specified.
- [R89] The value of the Security Policy Subject Service Attribute **MUST** be either *UNI* or *Cloud Application Interface*.

8.18.1.2 Region Name

A Region could be a neighborhood, city, country, part of a country, (e.g., U.S.-East), a continent, (e.g., Europe), or a geographical country grouping, (e.g., APAC). The value of the Region Name Service Attribute is the name of the geographical location that hosts the Edge Computing IaaS provides coverage or is physically located.

- [R90] An Edge Computing IaaS instance in each Region **MUST** be geographically isolated and operate independently from other Regions.
- [D2] An Edge Computing IaaS instance in a Region **SHOULD** support multiple Availability Zones.
- [R91] When ordering Edge Computing IaaS, the Region where the IaaS is located **MUST** be specified.
- [R92] The value of the Region Name Service Attribute **MUST** be unique within the SP domain.
- [R93] The value of the Region Name Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R94] The value of the Region Name Service Attribute **MUST** contain no more than 53 UTF-8 characters.

Each Installation Place and Edge Availability Zone is located in a Region; each Region can have multiple Edge Computing IaaS Installation Places and Edge Availability Zones.

8.18.1.3 Availability Zone Name

An Availability Zone is defined as a cluster of two or more interconnected Edge IaaS Instances within a Region. The Availability Zone Name Service Attribute value is a name that is unique within the Region.

- [R95] An Availability Zone **MUST** be equipped with independent power and cooling systems.
- [R96] Edge IaaS Instances in an Availability Zone **MUST** connect to each other over multiple network paths to achieve high availability.
- [R97] The value of the Availability Zone Name Service Attribute **MUST** be unique within the Region.
- [R98] The value of the Availability Zone Name Service Attribute **MUST** be a string consisting of one or more UTF-8 characters in the range of 32–126 (0x20 to 0x7e), inclusive.
- [R99] The value of the Availability Zone Name Service Attribute **MUST** contain no more than 53 UTF-8 characters.

8.18.2 Edge IaaS Service Level Specification

The Edge IaaS Service Level Specification Service Attribute (SLS) contains the technical details of the service level, in terms of Performance Objectives, agreed between the Subscriber and SP as part of the Service Level Agreement.

The value of the Edge IaaS Service Level Specification Service Attribute (SLS) is either *None* or a list of 3-tuple of the form $\langle ts, D, PM \rangle$ where:

- *ts* has a value that is the date and time for the start of the SLS expressed as a UTC time to the second.¹
- *D* has a value that is a time duration specified to the second, e.g., 1 month, 2 weeks, that is used in conjunction with *ts* to specify time intervals for determining when Performance Objectives are met. Note that the units for *D* are not constrained; in particular, 1 month is an allowable value for *D*, corresponding to a calendar month, e.g., from midnight on the 10th of one month up to but not including midnight the 10th of the following month.

¹ Note that *ts* is the start of the SLS and might not be the time that a service is first turned up, for example when the value of SLS is modified after the service is first turned up.

- *PM* is a non-empty list where each element in the list consists of a Performance Metric Name, a list of parameter values specific to the definition of the Performance Metric, and a Performance Objective.

[R100] If the value of the Edge IaaS Service Level Specification Service Attribute is not *None* then it **MUST** be a non-empty list of 3-tuple of the form $\langle ts, D, PM \rangle$.

A Performance Metric is a quantitative characterization of the IaaS service level experienced by the Subscriber. If any Performance Metrics are not included in the PM List within the SLS then the following sections (8.18.3, 8.18.4, 8.18.5, and 8.18.6) do not apply for that Performance Metric. This document specifies the following Performance Metrics, all of which can apply to an Edge IaaS instance:

- Availability Performance Metric (section 8.18.5)
- One-way Mean Packet Delay Performance Metric (section 8.18.6)

[CR1]<[R100] If *PM* contains multiple entries with a given Performance Metric Name, then for any two such entries, at least one of the parameter values **MUST** be different.

As an example of **[CR1]<[R100]**, One-way Mean Packet Delay Performance Metric could appear twice, once for each direction between the Cloud Service UNI and the Cloud Application Interface.

These Performance Metrics describe the performance that can be expected by the Subscriber. Methods for the SP and the Subscriber to monitor the performance to estimate this Subscriber experience are beyond the scope of this document.

The value of a Performance Metric that is compared against the corresponding Performance Objective is based on a given time interval from the Time Interval Sequence and the instances of Maintenance Intervals that apply during the given time interval. The Time Interval Sequence is defined in section 8.18.3.

The One-way Mean Packet Delay Performance Metric uses a parameter denoted by *S* as defined in section 8.18.4.

The One-way Mean Packet Delay Performance Metric might be important to the Subscriber when they are determining if the IaaS Edge needs to be located at their location.

8.18.3 Time Interval Sequence

The Time Interval Sequence is defined as $\{T_k, k = 0, 1, \dots\}$ where²

$$T_k = [ts + kD, ts + (k + 1)D)$$

² Numeric ranges are denoted as follows: z in $[x, y]$ means $x \leq z \leq y$, z in $[x, y)$ means $x \leq z < y$, z in $(x, y]$ means $x < z \leq y$, and z in (x, y) means $x < z < y$.

Each element of $\{T_k\}$ is used for assessing the success of the IaaS in meeting the Performance Objectives of the SLS.

Note that the use of the right parentheses in the above equation means that the elements of $\{T_k\}$ do not overlap. Even though the last second in a given time interval, T_k , is not fully contained in the given time interval, it is deemed to be in the given time interval for the purposes of computing the value of a Performance Metric.

8.18.4 S Parameter

The S parameter value is either *Up* or *Down*. When the value is *Up*, the Performance Metric applies to Packets that ingress at the Cloud Service UNI that are accepted. When the value is *Down*, the Performance Metric applies to Packets that ingress at the Cloud Application Interface that are accepted.

8.18.5 IaaS Availability Performance Metric

Each second in an interval, T_k is classified as either Available or Unavailable.

Let:

- $TS(T_k)$ = the number of seconds in T_k ,
- $AS(T_k)$ = the number of Available Seconds in T_k ,
- $\overline{AV}(T_k)$ = the value of the IaaS Availability Performance Metric for T_k .

Then the value of the IaaS Availability Performance Metric for T_k , expressed as percentage, is given by

$$\overline{AV}(T_k) = 100 \times \left(\frac{AS(T_k)}{TS(T_k)} \right)$$

Table 21 shows what is contained in a PM entry for the IaaS Availability Performance Metric.

Item	Value
Performance Metric Name	IaaS Availability
Performance Objective	$\hat{A}$ expressed as a percentage

Table 21 – PM Entry for the Edge Computing IaaS Availability Performance Metric

[CR2]<[R100] The SLS **MUST** define the Edge Computing IaaS Availability Performance Objective as met over T_k for a PM entry of the form in Table 21 if and only if $\overline{AV}(T_k) \geq \hat{A}$.

8.18.6 One-way Mean Packet Delay Performance Metric

The delay for a IP packet traversing the Edge Computing IaaS Cloud VC is defined in MEF 61.1 [5]. The one-way packet delay for an IP Data Packet that flows between Cloud UNI and Cloud

Application Interface is defined as the time elapsed from the reception of the first bit of the packet at Cloud UNI until the transmission of the last bit of the first corresponding egress packet at Cloud Application Interface. If the packet is erroneously duplicated as it traverses the network, the delay is based on the first copy that is delivered. The One-way Mean Packet Delay $\bar{d}(T_k, S)$ is defined as the arithmetic mean of one-way packet delay for IP packets for a given ordered pair of Cloud UNI and Cloud Application Interface, a given CoS Name, and a given time period T_k .

[CR3]<[R100] The SLS **MUST** define the One-way Mean Packet Delay Performance Objective, $\hat{d}$, as met over T_k for a *PM* entry of the form in MEF 61.1 [5] if and only if $\bar{d}(T_k, S) \leq \hat{d}$.

8.18.7 Building Block Common Reporting Parameters

The following parameters are used in sections 8.18.8, 8.18.9, and 8.18.10. These common parameters are defined in Table 22 .

Parameter Name	Summary Description	Possible Values
Calculation Interval	The interval over which the Performance Service Attributes are determined. The Calculation Interval is always $\geq$ to the Measurement Period and an integer multiplier of the Measurement Period.	An integer > 0 in milliseconds
Historical Report Time Interval	The maximum interval for which historical data is reported. Example values are milliseconds, seconds, days, weeks and months.	An integer > 0 in milliseconds
Report Interval	The interval over which the Performance Service Attributes are reported. The Report Interval is always $\geq$ to the Calculation Interval and an integer multiplier of the Calculation Interval.	An integer > 0 in milliseconds

Table 22 – Building Block Common Performance Parameters and Definitions

8.18.8 Processor Performance Reported Parameters

The Processor Performance Reported Parameters that are measured, calculated, and reported for Processors are defined in this section.

[R101] The Processor Performance Reported Parameters value **MUST** be either *None* or the list of parameters in Table 23 .

If the value of the Processor Performance Reported Parameters is *None*, then no reported parameters is agreed to by the Cloud Subscriber and the Cloud SP.

Parameter Name	Summary Description	Possible Values
Processor Utilization Period	The time period in milliseconds used to calculate processor utilization	≥ 1 and ≤ 1000
Measured Utilization	The measured processor utilization percentage.	≥ 1 and ≤ 100
Calculation Interval	See section 8.18.7	$\geq$ Measurement Period and an integer multiple of the Measurement Period
Average Utilization	The average percentage calculated by the SP of the Processor that is being used by the Subscriber over the Calculation Interval.	≥ 1 and ≤ 100
Maximum Utilization	The maximum percentage of the Processor that is being used by the Subscriber over the Calculation Interval.	≥ 1 and ≤ 100
Report Interval	See section 8.18.7	$\geq$ Calculation Interval and an integer multiple of the Calculation Interval
Reported Average Utilization	The average for each Calculation Interval included in the Report Interval	≥ 1 and ≤ 100
Reported Maximum Utilization	The maximum for each Calculation Interval included in the Report Interval	≥ 1 and ≤ 100
Report Start Time	The time that the Report started in UTC	HH:MM:SS:ms
Historical Report Time Interval	See section 8.18.7	• milliseconds • Seconds • Minutes • Hours • Days • Weeks • Months • Years

Table 23 – Processor Performance Reported Parameters

The Processor Utilization performance metric is measured by sampling the processor utilization over a period of time. For the purpose of this standard, that period of time is no greater than one second. The sample taken each time period identifies the processor utilization percentage. This time period sample is reported as the Processor Utilization at the end of each sample period. The method used to determine processor utilization is beyond the scope of this standard.

8.18.8.1 Processor Utilization

The Processor Utilization is the percentage of the Processor Utilization Period that the processor is determined to be in use. The method used to determine the Processor Utilization is beyond the scope of this document.

8.18.8.2 Processor Utilization Period

The Processor Utilization Period is the measurement period for Processor Utilization in milliseconds.

8.18.9 Memory Performance Service Attributes

The Memory Performance Service Attribute parameters that are measured, calculated, and reported for Memory are defined in this section.

[CR4]<[R100] The Memory Service Attribute value **MUST** be either *None* or the list of parameters in Table 24 .

Parameter Name	Summary Description	Possible Values
Memory Utilization Period	The time period in milliseconds used to calculate memory utilization	≥ 1 and ≤ 1000
Memory Utilization	The measured memory utilization percentage.	≥ 1 and ≤ 100
Calculation Interval	See section 8.18.7	$\geq$ measurement period and an integer multiplier of the measurement period
Average Memory Utilization	The average percentage calculated by the SP of the Memory that is being used by the Subscriber over the Calculation Interval.	≥ 1 and ≤ 100
Maximum Memory Utilization	The maximum percentage of the Memory that is being used by the Subscriber over the Calculation Interval.	≥ 1 and ≤ 100
Report Interval	See section 8.18.7	$\geq$ Calculation Interval and an integer multiplier of the Calculation Interval
Reported Average Memory Utilization	The average for each Calculation Interval included in the Report Interval	≥ 1 and ≤ 100
Reported Maximum Memory Utilization	The maximum for each Calculation Interval included in the Report Interval	≥ 1 and ≤ 100
Report Start Time	The time that the Report started in UTC	HH:MM:SS:ms
Historical Report Time Interval	See section 8.18.7	• milliseconds • Seconds • Minutes • Hours • Days • Weeks • Months • Years

Table 24 – Memory Performance Service Attribute Parameters

The Memory Utilization performance metric is measured by sampling the percentage of memory in use over a period of time. For the purpose of this standard, that period of time is no greater than one second. The sample taken each time period identifies the memory utilization percentage. This time period sample is reported as the Memory Utilization at the end of each sample period. The method used to determine the memory utilization percentage is beyond the scope of this standard.

8.18.9.1 Memory Utilization

The Memory Utilization is the percentage of Memory Amount that is determined to be in use during the Memory Utilization Period. The method used to determine the Memory Utilization is beyond the scope of this document.

8.18.9.2 Memory Utilization Period

The time in milliseconds that is specified as the sample period for the Memory Utilization.

8.18.10 Storage

The Storage Performance Service Attribute parameters that are measured, calculated, and reported for Storage are defined in this section.

Parameter Name	Summary Description	Possible Values
Storage Utilization Period	The time period in minutes used to calculate storage utilization	≥ 1 and ≤ 1440
Calculation Interval	See section 8.18.7	$\geq$ to measurement period and an integer multiplier of the measurement period
Report Interval	See section 8.18.7	$\geq$ to Calculation Interval and an integer multiplier of the Calculation Interval
Reported Storage Utilization	The value for each Calculation Interval included in the Report Interval	≥ 1 and ≤ 100
Reported Storage Utilization Period	The time period of the measured Reported Percent Storage Utilization	• Minutes • Hours • Days • Weeks • Months • Years
Storage Utilization	The amount of storage in use in Mbytes	Numeric value > 0

Table 25 – Storage Performance Service Attribute Parameters

The Storage Utilization performance metric is measured by sampling the percentage of storage in use over a period of time. For the purpose of this standard, that period of time is no greater than one day. The sample taken each time period identifies the storage utilization percentage. This time period sample is reported as the Storage Utilization at the end of each sample period. The method used to determine the storage utilization percentage is beyond the scope of this standard.

9 References

- [1] IETF RFC 2119, *Key words for use in RFCs to Indicate Requirement Levels*, by S. Bradner, March 1997
- [2] IETF RFC 7230, *Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing*, by R. fielding, Ed., J. Reschke, June 2014, Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.
- [3] IETF RFC 8174, *Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words*, by B. Leiba, May 2017, Copyright (c) 2017 IETF Trust and the persons identified as the standard authors. All rights reserved.
- [4] MEF 10.4, *Subscriber Ethernet Service Attributes*, December 2018
- [5] MEF 61.1, *IP Building Block Service Attributes*, May 2019
- [6] MEF 61.1.1, *Amendment to MEF 61.1: UNI Access Link Trunks, IP Addresses, and Mean Time to Repair Performance Metric*, July 2022
- [7] MEF 68, *Cloud Services Architecture*, October 2021
- [8] MEF 69.1, *Subscriber IP Service Definitions*, February 2022
- [9] MEF 70.2, *SD-WAN Service Attributes and Service Framework*, October 2023
- [10] MEF 95, *MEF Policy Driven Orchestration (PDO)*, July 2021
- [11] MEF 118.1, *Zero Trust Framework for MEF Services*, July 2024
- [12] MEF 130, *Device Physical and Environmental Building Block Service Attributes*, February 2023
- [13] Mplify 150, *Installation Place and Service Site Management Business Requirements and Use Cases*, June 2024
- [14] NIST Special Publication (SP) 800-145, *The NIS Definition of Cloud Computing*, Peter Mell, Timothy Grance, September 2011

Appendix A Service Attributes by IaaS Type

This appendix contains a table that describes the relationship between all Service Attributes in this document and the two types of IaaS, Shared, Dedicated with two types of Dedicated, with Processor Manufacturer Specified, and without Processor Manufacturer Specified.

Service Attribute	Dedicated		Shared
	Manufacturer Specified	Manufacturer Unspecified	
Processor Service Attributes			
Processor Building Block Identifier	Specified	Specified	Specified
Processor Manufacturer Identifier	Specified	Unspecified	Unspecified
Manufacturer Model Identifier	Specified	Unspecified	Unspecified
Processor Architecture	Unspecified	Specified	Unspecified
Clock Speed Range	Unspecified	Specified	Unspecified
Number of Cores	Unspecified	Specified	Unspecified
Number of Threads	Unspecified	Specified	Unspecified
Level 1 Cache Size	Unspecified	Specified	Unspecified
Level 2 Cache Size	Unspecified	Specified	Unspecified
Instructions Per Clock Cycle	Unspecified	Specified	Unspecified
Quantity	Specified	Specified	Specified
Memory Service Attributes			
Memory Building Block Identifier	Specified	Specified	Specified
Memory Amount	Specified	Specified	Specified
Storage Attributes			
Storage Building Block Identifier	Specified	Specified	Specified
Storage Type	Specified	Specified	Specified
Storage Retention Duration	Specified	Specified	Specified
Storage Time to Access	Specified	Specified	Specified
Storage Amount	Specified	Specified	Specified
Storage Replication	Specified	Specified	Specified
Access Network Building Block Components			
Cloud UNI	Specified	Specified	Specified
Cloud UNI Access Link	Specified	Specified	Specified
Edge Computing IaaS Cloud VC	Specified	Specified	Specified
Edge Computing IaaS Cloud VC End Points	Specified	Specified	Specified
Cloud Application Interface	Specified	Specified	Specified
Cloud UNI Service Attributes			

Service Attribute	Dedicated		Shared
	Manufacturer Specified	Manufacturer Unspecified	
Cloud UNI Subscriber Application Awareness	Specified	Specified	Specified
Cloud UNI Subscriber Application Identifier Method	Specified	Specified	Specified
Cloud UNI Subscriber Application Policy	Specified	Specified	Specified
Edge Computing IaaS Cloud VC Service Attributes			
Edge Computing IaaS Cloud VC Subscriber Application Awareness	Specified	Specified	Specified
Edge Computing IaaS Cloud VC Subscriber Application Ingress Policy	Specified	Specified	Specified
Edge Computing IaaS Cloud VC Subscriber Application Egress Policy	Specified	Specified	Specified
Edge Computing IaaS Cloud VC Subscriber Application Priority	Specified	Specified	Specified
Edge Computing IaaS Cloud VC Subscriber Application Priority to DSCP Mapping	Specified	Specified	Specified
Cloud Application Interface Service Attributes			
Cloud Application Interface Identifier	Specified	Specified	Specified

Service Attribute	Dedicated		Shared
	Manufacturer Specified	Manufacturer Unspecified	
Cloud Application Interface Ingress Band-width Profile Envelope	Specified	Specified	Specified
Cloud Application Interface Egress Band-width Profile Envelope	Specified	Specified	Specified
Cloud Application Interface List of Control Protocols	Specified	Specified	Specified
Cloud Application Interface Routing Protocols	Specified	Specified	Specified
Cloud Application Interface Reverse Path Forwarding	Specified	Specified	Specified
Cloud Application Interface Subscriber Application Awareness	Specified	Specified	Specified
Cloud Application Interface Ingress Policy	Specified	Specified	Specified
Cloud Application Interface Egress Policy	Specified	Specified	Specified
Cloud Application Interface Priority	Specified	Specified	Specified
Cloud Application Interface Priority to DSCP Mapping	Specified	Specified	Specified
Edge IaaS Access Service Attributes			
Edge IaaS Access URL	Specified	Specified	Specified
Edge IaaS Access Identifier	Specified	Specified	Specified
Edge IaaS Access Password	Specified	Specified	Specified
Host OS Service Attributes			
Host OS Building Block State	Specified	Specified	Specified
Host OS Building Block Identifier	Specified	Specified	Specified
Host OS Vendor Identifier	Specified	Specified	Specified
Host OS Name	Specified	Specified	Specified
Host OS Version	Specified	Specified	Specified
Hypervisor Service Attributes			
Hypervisor Building Block State	Specified	Specified	Specified

Service Attribute	Dedicated		Shared
	Manufacturer Specified	Manufacturer Unspecified	
Hypervisor Building Block Identifier	Specified	Specified	Specified
Hypervisor Vendor	Specified	Specified	Specified
Hypervisor Name	Specified	Specified	Specified
Hypervisor Version	Specified	Specified	Specified
Guest OS Service Attributes			
Guest OS Building Block State	Specified	Specified	Specified
Guest OS Building Block Identifier	Specified	Specified	Specified
Guest OS Publisher	Specified	Specified	Specified
Guest OS Name	Specified	Specified	Specified
Guest OS Version	Specified	Specified	Specified
Container Engine Service Attributes			
Container Engine Building Block State	Specified	Specified	Specified
Container Engine Building Block Identifier	Specified	Specified	Specified
Container Engine Vendor	Specified	Specified	Specified
Container Engine Name	Specified	Specified	Specified
Container Engine Version	Specified	Specified	Specified
Virtual Machine Service Attributes			
Virtual Machine Building Block State	Specified	Specified	Specified
Virtual Machine Building Block Identifier	Specified	Specified	Specified
Number of Virtual Processors	Specified	Specified	Specified
Number of Threads	Specified	Specified	Specified
Number of Cores	Specified	Specified	Specified
Number of Physical CPUs	Specified	Specified	Specified
Virtual Memory Amount	Specified	Specified	Specified
Host OS	Specified	Specified	Specified
Security Policy Service Attributes			
Security Policy Identifier	Specified	Specified	Specified
Security Policy Name	Specified	Specified	Specified
Security Policy Constraints	Specified	Specified	Specified
Security Policy Subject	Specified	Specified	Specified
Edge Availability Zone Service Attributes			
Availability Zone Name	Specified	Specified	Specified
List of Location Identifiers	Specified	Specified	Specified
Affinity Type	Specified	Specified	Specified
Entity Identifier	Specified	Specified	Specified
Edge IaaS SLS Service Attributes			

Service Attribute	Dedicated		Shared
	Manufacturer Specified	Manufacturer Unspecified	
Edge IaaS SLS	Specified	Specified	Specified
Availability	Specified	Specified	Specified
One-way Mean Packet Delay	Specified	Specified	Specified
Processor Reporting Service Attributes			
Percent Utilization Period	Specified	Specified	Specified
Measured Utilization	Specified	Specified	Specified
Calculation Interval	Specified	Specified	Specified
Average Utilization	Specified	Specified	Specified
Maximum Utilization	Specified	Specified	Specified
Report Interval	Specified	Specified	Specified
Reported Average Utilization	Specified	Specified	Specified
Reported Maximum Utilization	Specified	Specified	Specified
Report Start Time	Specified	Specified	Specified
Historical Report Time Interval	Specified	Specified	Specified
Memory Reporting Service Attributes			
Memory Utilization Period	Specified	Specified	Specified
Memory Utilization	Specified	Specified	Specified
Calculation Interval	Specified	Specified	Specified
Average Memory Utilization	Specified	Specified	Specified
Maximum Memory Utilization	Specified	Specified	Specified
Report Interval	Specified	Specified	Specified
Reported Average Memory Utilization	Specified	Specified	Specified
Reported Maximum Memory Utilization	Specified	Specified	Specified
Report Start Time	Specified	Specified	Specified
Historical Report Time Interval	Specified	Specified	Specified
Storage Reporting Service Attributes			
Storage Utilization Period	Specified	Specified	Specified
Calculation Interval	Specified	Specified	Specified
Report Interval	Specified	Specified	Specified
Reported Storage Utilization	Specified	Specified	Specified
Reported Storage Utilization Period	Specified	Specified	Specified
Storage Volume Utilization	Specified	Specified	Specified

Table 26 – Service Attribute to Edge IaaS Type Mapping

Appendix B Acknowledgements

- Mike **BENCHECK**
- Basil **NAJEM**
- Dan **PITT**
- Ralph **SANTITORO**
- Larry **SAMBERG**
- Isabelle **MORENCY**