

Glue2.1 Cloud extensions New Accelerators attributes

Alessandro Paolini

EGI.eu



www.egi.eu

This work by EGI.eu is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

Current situation in GLUE2

- The CLOUD services do not publish any information related to the VO enabled on them
- Used classes and attributes defined for the GRID services
 - much information are incomplete and have been adjusted to publish something coherent and consistent with the schema.

- New classes and attributes for cloud computing services
- New attributes for GPGPUs hardware
 - Now called Accelerators

Main entities: relation to cloud computing

- The **location** is **related** to zero or more **cloud computing services**
- The **contact** is **related** to zero or more **cloud computing services**
- An **AdminDomain** **manages** zero or more **cloud computing services**

- main logical unit, and aggregation point for several entities together modeling a computing infrastructure capability in a cloud system
- capable of executing **CloudComputingInstance** on its associated resources
 - Correspond to ComputingActivities in the grid scenario
- The resources behind the CloudComputingService are described via the **CloudComputingManager**, **CloudComputingInstanceType**, **CloudComputingImage** and **Benchmark** entities
 - Correspond to ComputingManager, ExecutionEnvironment, ApplicationEnvironment, ApplicationHandle and Benchmark in the grid scenario

- The governing policies and status of the resources are given by the **CloudComputingShare** elements
 - Correspond to ComputingShare in the grid scenario
- The **CloudComputingInstance** of a **CloudComputingService** are submitted and controlled via a **CloudComputingEndpoint**
 - Correspond to ComputingEndpoint in the grid scenario

- Attributes:
 - Virtual Machines (VM): total, running, pending, suspended, halted
 - Link to the Acceptable User Policy (AUP)
- Associations:
 - **CloudComputingEndpoint.ID**: a **CloudComputingService** is associated with zero or more endpoints (interfaces)
 - **CloudComputingShare.ID**: a **CloudComputingService** offers zero or more computing shares
 - **CloudComputingManager.ID**: A **CloudComputingService** offers zero or more computing manager.
 - **Extension.Key, Contact.ID, Location.ID, Service.ID**

- represents an endpoint which is used to create, control and monitor cloud computing activities called Virtual Machines instances
- provides attributes that MAY be used to publish summary information about VM instantiated via a particular Endpoint

- Attributes:
 - Name and version of the supported contextualization mechanism
- Associations:
 - **CloudComputingService.ID**: A Cloud endpoint is part of a Cloud Computing Service
 - **CloudComputingShare.ID**: A Cloud endpoint MAY pass activities to zero or more Cloud computing shares
 - **CloudComputingInstance.ID**: A Cloud endpoint has accepted and is managing zero or more Cloud Activities
 - **AccessPolicy.ID**: A computing endpoint has associated zero or more AccessPolicies

- high-level concept introduced to model a utilisation target for a pool of resources, defined by a homogeneous set of configuration parameters and characterized by single status information
- carries information about "policies" (limits) defined over all or a subset of resources and describes their dynamic status (load)
- stores also a set of **CloudComputingImage** and **CloudComputingInstanceType**
 - used to define respectively the virtual OS and the virtual hardware resources of the **CloudComputingInstance** running on the share
 - Such virtual OS and hardware resources are provided by the Share with a given Service Level Agreement (SLA)
- In case of the same **CloudComputingInstanceType** and **CloudComputingImage** are offered under different SLAs, multiple **CloudComputingShares** shall be created, one for each SLA

- Attributes:
 - Virtual Machines (VM): total, running, pending, suspended, halted, maximum number (upper limit in the share)
 - For a single VM: maximum number of CPU, **Accelerators** and RAM
 - Network info

CloudComputingShare: associations

- **CloudComputingEndpoint.ID**: a cloud computing share MAY be consumed via one or more computing endpoints
- **CloudComputingInstanceType.ID**: a cloud computing share is defined on one or more computing resources
- **CloudComputingService.ID**: a cloud computing share participates in a computing service
- **CloudComputingInstance.ID**: a cloud computing share is being consumed by zero or more computing activities
- **CloudComputingImage.ID**: a CloudComputingShare provides zero or more VM Image templates
- **ToCloudStorageService.ID**: Link to the storage share used to store instances templates, VM images and/or attached disks
- **MappingPolicy.ID**: A share has zero or more mapping policies.

- The **CloudComputingManager** is responsible for the local control of resources
- piece of software, firmware or hardware which creates, runs and manages Virtual Machines
- The layer is not exposed directly to external clients or to the Virtual Machines themselves
- also known as Hypervisor
- The class provides aggregated information on controlled resources, hypervisors limits and also describes local storage extents accessible to the Virtual Machines

CloudComputingManager: attributes

- Total number of Physical CPUs and **Accelerators**
 - SHOULD represent the total installed capacity
- Total amount of RAM
 - SHOULD represent the total installed capacity
- VMs: max/min of virtual CPUs/Accelerators/RAM per instance
- Type of network, CPU and Accelerators virtualization
- format of virtual disk images supported
- Failover, LiveMigration, VMBackupRestore

CloudComputingManager: associations

- **CloudComputingService.ID**: A computing manager participates in a computing service
- **CloudComputingInstanceType.ID**: A computing manager manages one or more execution environments
- **Benchmark.ID**: A computing manager has zero or more associated benchmarks. This benchmarks are referred to the virtual resources (RAM, CPU, disk, network) provided to the VMs

- The class characterizes the relative performance either about an **Execution Environment** providing computing capacity or about a **CloudComputingInstanceType** providing cloud computing capacity
- Associations:
 - ExecutionEnvironment.ID, ComputingManager. ID
 - **CloudComputingInstanceType.ID**: a benchmark MAY be related to a cloud computing instance type
 - **CloudComputingManager.ID**: a benchmark MAY be related to a cloud computing resource

CloudComputingInstanceType

- The class describes the hardware environment of the VM, i.e. the amount of RAM, CPU, disk and network resources the VM OS will see and manage
- The resources provided to the VM are virtual resources
- The performances of the provided resources are specified via the Benchmarks associated to the Instance Type

CloudComputingInstanceType: attributes

- Template ID and reference to the Market Place
- Platform, vCPU, RAM, disk, Ephemeral storage
- Network

CloudComputingInstanceType: associations

- **CloudComputingManager.ID**: Cloud Computing Instance Type is managed by a Cloud computing manager
- **CloudComputingShare.ID**: Cloud Computing Instance Type is served by a set of computing shares
- **CloudComputingInstance.ID**: Zero or more cloud computing instances runs this Cloud Computing Instance Type
- **CloudComputingEndpoint.ID**: Cloud Computing Instance Type is available on a set of Cloud Computing Endpoints
- **CloudServicePrice.ID**: The price metric associated to the resources provided by this template. It contains a different metric for each resource (Computing, Memory, Network IN/OUT)

- The entity is used to describe a set of homogeneous virtual accelerator devices.
 - Generally a virtual accelerator device corresponds to physical one installed on the host.
- A cloud computing instance may be associated with one or more virtual accelerators
- Attributes: type, number, vendor, model, version, speed and memory

- The class describes the software environment of the VM:
 - which OS is booting at VM startup
 - which pre-installed software is available on it
- Each application installed on the OS is identified by a name (the InstalledSoftware attribute)
 - these names are not defined within the schema
 - SHOULD be assigned in a way which allows applications to be uniquely identified
- The OS template may require a certain amount of resources (CPU, RAM and GPU) to run

CloudComputingImage: attributes

- Template and Market Place Ids
- OS: platform, family, name, version
- Disk size
- Recommended CPUs, Accelerators and RAM
- Minimum numbers of CPUs, Accelerators and RAM
- Information about user access to the VM
 - credentials injected during contextualization, pre-defined username/password, pre-defined RSA key
- Default username and password
- Custom software installed on the instance

CloudComputingImage: associations

- **CloudComputingShare.ID**: an OS template is available to one or more computing shares
- **CloudComputingInstance.ID**: an OS template is used by one or more cloud computing activities
- **CloudComputingEndpoint.ID**: an OS template is available on a set of Cloud Computing Endpoints
- **CloudServicePrice.ID**: the price metric associated to the resources provided by this template. It contains a different metric for each resource (OS License, Application license, etc...)
- **ToCloudStorageService**: link to the OS disk location in the storage service

- The class represents a single VM (but possibly multi-VM) instance
- The attributes give the instance properties and state as seen by the hypervisor:
 - Type, State, ID and local ID
 - Error, Owner, Local Owner, Execution node
 - Total CPU and accelerators time consumed so far by the VM
 - amount of RAM currently used by the VM
 - Submission, start and end time
 - Submission host and client name

CloudComputingInstance: associations

- **CloudComputingEndpoint.ID**: A Cloud computing activity is submitted to a computing endpoint
- **CloudComputingShare.ID**: A Cloud computing activity is mapped into a computing share
- **CloudComputingInstanceType.ID**: A Cloud computing activity is executed in an execution environment
- **ToCloudStorageService**: Link to the attached disks location in the storage service. The OS disk is included in this list

- The entity models the resources price in the cloud environment.
 - This entity represents a single price voice of the final price for the activity.
 - This voice is related to a given consumed resource which is specified in the attributes of the entity.
- The price for the given resource is represented by a fixed price or consumption model
- Different costs models may be applied to different users according to the scope of the Computing Activity
- With the **Scope** attribute, this entity MAY be used in alternative to **User Domain** and **Access Policy** entities, to discriminate which category of users have access to the resource

CloudServicePrice: attributes

- Referring resource (CPU, Memory, Disk, Software Licenses, ...)
- Scope
- Fixed fee to be paid for the service usage and related billing period
- Consumption fee to be paid for the service and related billing period

CloudServicePrice: associations

- **CloudComputingInstanceTypeID**: it MAY be associated to a Computing Instance Type
- **CloudComputingImageID**: it MAY be associated to a OS template
- **StorageServiceID**: it MAY be associated to a Storage service

- The class represents the case where a virtual disk is created into the Storage Service for VM usage.
- The disk may be attached to the VM (visible as a disk device by the VM OS) or available via other export protocols (NFS share, iSCSI, etc...)
- The attributes of this entity refer to the link of the storage resource into the Storage System and the VM environment:
 - Reference to this disk to be used for attaching/detaching it to a VM
 - The link to the local storage into the VM environment
 - The link to the storage into the remote Storage Service

ToCloudStorageService: associations

- **CloudComputingInstance.ID**: it MAY be associated to a cloud computing instance
- **CloudComputingImage.ID**: it MAY be associated to a OS template
- **CloudComputingService.ID**: is associated to a cloud computing service
- **StorageService.ID**: is associated to a storage service

- **StorageAccessProtocol:**
 - Added the relation to **ToCloudComputingService**
- **ToCloudComputingService**
 - The class describes a network connection between a Storage Service and a Cloud Computing Service
 - Attributes: **NetworkInfo**, **Bandwidth**
 - **StorageAccessProtocol.ID**: the storage service MAY be accessed via an access protocol by a certain computing service
 - **CloudComputingService.ID**: it is associated to a cloud computing service
 - **StorageService.ID**: it is associated to a storage service

- ComputingShare
 - MaxAcceleratorSlotsPerJob, FreeAcceleratorSlots, UsedAcceleratorSlots
- ComputingManager
 - Total physical accelerators, used and total slots
- ExecutionEnvironment: association with the new class AcceleratorEnvironment.ID:
 - An execution environment MAY contains one or more accelerator environments

- The entity is used to describe an homogeneous set of accelerator processors
- It is associated with one or more **execution environments**
- Attributes:
 - Type, number of physical and logical, vendor, model, version, speed and memory

- CloudComputingInstanceState_t (open enumeration)

Value	Description
bes:failed	(a terminal state): the activity has failed due to some system error/failure event, such as failure of a computational resource that the activity was running on
bes:finished	(a terminal state): the activity has terminated successfully. Successful termination implies that the activity exited of its own accord rather than due to some failure in the BES or of the computational resources on which the activity was running. Note that a successfully terminating activity MAY nevertheless return an error code as its return value
bes:pending	the service has created a record for an activity but not yet instantiated it on a suitable computational resource or enabled it to start execution on such a resource
bes:running	the activity is executing on some computational resource
bes:terminated	(a terminal state): the client – which might be some system administrator (and hence not necessarily the client who originated the request to create the activity) – has issued a TerminateActivity request

CloudComputingInstanceType_t (closed enumeration)

Value	Description
collectionelement	A VM submitted as part of a collection of individual VMs which do not communicate among them
parallelement	A VM submitted as part of a collection of individual VM which communicate among them (ex. part of a computing cluster)
single	An individual stand-alone VM

CloudComputingManagerType_t (open enumeration)

Value	Description
openstack	OpenStack
opennebula	OpenNebula
wnodes	WNoDeS
synnefo	Synnefo
cloudstack	Apache CloudStack
vcloud	VMWare vCloud

CloudResourceName_t (closed enumeration)

Value	Description
flavour	Resource flavour, comprehensive of disk, CPU, ram, memory and network
network	Network usage (amount of kilobytes inbound/outbound from the VM network interfaces)
networkin	Network inbound usage (amount of kilobytes inbound from the VM network interfaces)
networkout	Network outbound usage (amount of kilobytes outbound from the VM network interfaces)
cpu	Number of vCPU associated to the VM. A vCPU seen by the OS is considered in use, even if it is in idle mode
gpu	Number of vGPU associated to the VM. A vGPU seen by the OS is considered in use, even if it is in idle mode
memory	RAM associated to the VM in MBs. Memory associated to the VM, and seen by the OS, is accounted even if free in the OS space.
storage	Storage associated to the VM in GBs. It includes all the type of storage associated to the VM, both block storage or ephemeral storage or OS disk image storage
ephemeralstorage	Ephemeral storage associated to the VM in GBs
blockstorage	Block storage associated to the VM in GBs. The entire amount of the block storage device disk is accounted, even if storage is not used by the OS
osdiskstorage	Storage associated to the VM Image OS disk in GB. The entire amount of the block storage device disk is accounted, even if storage is not used by the OS.
softwarelicense	License for the usage of software on the VM image disk. This shall include all the licensing prices for all the software installed in the VM besides the OS
oslicense	License for the usage of the OS on the VM image disk.

ContextualizationName_t

Value	Description
cloudinit	Ubuntu CloudInit
bash	Generic bash script

CPUVirtualizationT_t

Value	Description
emulation	Full virtualization with emulation
paravirtualization	Para-virtualization
hardware	Hardware assisted virtualization

DiskVirtualizationT_t

Value	Description
qcow2	QCOW2
raw	RAW Image
vmdk	VMDK

Open enumeration

AcceleratorType_t (open enum.)

Value	Description
GPU	GPU architecture
MIC	Intel MIC architecture
FPGA	Generic FPGA accelerator

AcceleratorSlot_t

- The accelerator slot type is a tuple with the following fields:
 - AcceleratorType_t: the accelerator architecture type
 - UInt32: the number of slots

AcceleratorNumber_t:

- a tuple with the following fields:
 - AcceleratorType_t: the accelerator architecture type
 - UInt32: the number of physical or virtual accelerator card

AcceleratorVirtualizationT_t (open enum.)

Value	Description
emulation	Full virtualization with emulation
paravirtualization	Para-virtualization
hardware	Hardware assisted virtualization
vgpu	NVidia vGPU

AcceleratorVirtualization_t

- The accelerator virtualization is a tuple with the following fields:
 - AcceleratorType_t: the accelerator architecture type
 - AcceleratorVirtualizationT_t: the type of virtualization

NetVirtualizationT_t (open enum.)

Value	Description
none	No network isolation
vswitch	vSwitch isolation
eatables	EBtables isolation
vlan	VLAN isolation
firewall	Firewall isolation

HostAccessInfo_t (closed enum.)

Value	Description
none	No default credentials are present in the machine, access credentials shall be injected by the middleware
passwd	Pre-set username/password
rsa	Pre-set RSA key set

Period_t (open enum.)

Value	Description
once	The metric is calculated once in the lifetime of the service
asaccounted	The metric is calculated with the same granularity as it is accounted
daily	The metric is calculated daily
monthly	The metric is calculated monthly
yearly	The metric is calculated yearly

ResourceScope_t (open enum.)

Value	Description
any	Any scope (commercial or non-commercial)
commercial	Access by commercial organizations for commercial usage
noncommercial	Access by no-profit or public organizations for non-commercial usage
research	Access for research by any organization, commercial or non-commercial
noncommercialresearch	Access for research by no-profit or public organizations for non-commercial usage
Institutions	Access for public organization, institutions or governments for any scope
euonly	Access restricted to users of EU based countries
owncountryonly	Access restricted to users of the same country of the provider

- The naming of **ToCloudStorageService** is maybe not ideal since it suggests that it relates to a **CloudStorageService** object, which isn't true
 - **CloudComputingServiceToStorageService** (seems too long) ?
 - **CloudToStorageService** (to keep the name short) ?
- in the **StorageService** definition is not included the relation to **ToComputingService** class:
 - do we need to include the association to **ToCloudComputingService**?
- Same question for **StorageService** and **ToStorageService** / **ToCloudStorageService** classes
- should we rename the class **CloudVirtualAccelerator** in **CloudComputingVirtualAccelerator** ?

Thank you for your attention.

Questions?



www.egi.eu

This work by EGI.eu is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).