

Computing Resource Reporting format specification (version 1.2)

Computing Resource Reporting (CRR) is an alternative to the BDII for description of the computing resources of the WLCG sites. The CRR file (json format) should be generated at the site and made accessible through the http protocol.

“Computing Resource” is the main concept for modeling of the description of the computing resources of the WLCG sites. It combines "identical" computing resources (i.e. same hardware, same software) and is uniquely identified in the scope of a single site. Same hardware assumes that the hardware of a given ‘computingresource’ can be considered identical regarding computing power and common conversion factor for translating wallclock time into wallclock work can be applied for the resource. Same software assumes common low-level resource management framework.

A given computing resource is normally organized into shares, each of those having common values for attributes like maximum wallclock time , maximum physical memory per processor, list of VOs which can use a particular share. The access to the computing resource is normally (though not necessary) enabled via computing element services. In the specification they are referenced as ‘accesspoints’.

Normally all access points provide access to all shares of the resource. Though the specification foresees the possibility to describe situations when a particular set of access points provides access to some subset of the shares of a given resource.

This specification does not currently contain all necessary attributes to describe GPU resources. Will be extended later when we have better understanding what it might require.

Format

```
“computingresources” : {
  <computing_resource_identifier> :{
    ‘site’ : <site_gocdb/oim_identifier>,
    ‘type’ : <low-level_resource_managemnt_framework_type>,
    ‘jobmanager’ : < lowlevel_resource_managemnt
                    _framework_implementation>,
    ‘jobmanager_version’: < lowlevel_resource_managemnt
                    _framework_implementation_version>,
    ‘coprocessor’: <GPU_enabled>, //Indicates whether GPU is enabled, if not
    can be ommited
    ‘status’: <status_of_the_resource>,
    ‘number_logical_cpus’: #,
    ‘capacity_HS06’: #,
    ‘os’ : <WN_operation_system>,
    ‘default_walltime_minutes’:#,
    ‘default_cputime_minutes’:#,
    ‘memory_GB’:#,
    ‘assigned_vos’: [...],
    ‘publication_time’: <UTC_publication_time_stamp>,
    ‘comment’: <free_form_string>,
    ‘shares’: {
```

```

    <share_identifier> :{
        'type' : <share_type>,
        'max_wall_minutes' : #,
        'max_main_memory_GB' : #,
        'assigned_vos' : [], // needed only if it is different of the value
                             of the attribute defined at the resource level
    }
    ...
}
'accesspoints' : {
    <accesspoint_identifier> : {
        'endpoint_urls' : [],
        'flavour' : <accesspoint_implementation>,
        'version' : <accesspoint_implementation_version>,
        'shares' : [] //needed only if this accesspoint does not provide
                     access to all shares of the resource. In this case the list should
                     contain the identifiers of the subset of shares accessible
                     through this accesspoint
    },
    ....
}
...
}
}

```

Definitions

- **computingresource**: computeresource concept combines "identical" computing resources (i.e. same hardware, same software) and is uniquely identified in the scope of a single site. Same hardware assumes that the hardware of a given computingresource can be considered identical regarding computing power and common conversion factor for translating wallclock time into wallclock work can be applied for the resource. Same software assumes common low-level resource management framework.
 - **sitename**: the GOCDDB/OIM name of the site hosting the computing service . Value type is string.
 - **type**: defines the type of the resource management framework. Examples of currently considered values are : 'batch', 'vacuum' . Value type is string.
 - **jobmanager, jobmanager_version**: the type and the version of the resource management framework or LRMS in case the type of the resource management framework is 'batch'. In the latter case is identical to the GLUE2ManagerProductName and GLUE2ManagerProductVersion
 - **status**: the maturity of the service in terms of quality of the software/hardware components. fixed enumeration values development,pre-production,testing,production. See GLUE2QualityLevel definitions. Value type is string.
 - **number_logical_cpus**: number of logical CPUs of the resource. Value type is number.

- **capacity_HS06**: capacity of the resource in HS06. Value type is number.
- **os**: version of the OS pre-installed on the resource worker nodes. Value type is string.
- **coprocessor**: indicates whether GPU is enabled, in this case the value is 'GPU'. Otherwise can be omitted. Value type is string.
- **default_walltime_minutes**: default max wallclock time of the resource in minutes. Value type is number.
- **default_cputime_minutes**: default max CPU time of the resource in minutes. Value type is numbers
- **memory_gb**: default memory per processor in GB. Value type is number.
- **assigned_vos**: the list of authorized user groups.Can be defined on the resource level, however can be overloaded on the level of a particular share (see below). Value type is list of strings.
- **publication_time**: UTC time stamp in seconds since epoch of the time when the data block has been published. Value is number.
- **shares**: shares section describes possible split of the resource in subsets, each of which can be different in terms of authorized user groups, max memory and wallclock or other attributes which are not currently considered. Each share should be uniquely identified in the scope of a single resource.
 - **type**: the type of the share depends on the jobmanager implementation, the most common type is 'queue'. Type of the share can be omitted. Value type is string.
 - **max_walltime_minutes**: as defined in the GLUE2ComputingShareMaxWalltime, the max. obtainable wallclock limit for a single-core job. The unit is minutes! Value type is number.
 - **max_cputime_minutes**: the max obtainable CPU limit for a single-core job. Value type is number.
 - **max_main_memory_GB**: as defined by the GLUE2ComputingShareMaxMainMemory, the max physical memory per processor. The unit is GBs! Value type is number.
 - **assigned_vos**: the list of authorized user groups for a particular queue or subset of resources. Can be omitted if it does not differ from the one defined on the resource level. Value is a list of strings.
- **accesspoints**: accesspoints section describes the way the resource or a particular share can be accessed. Generally it contains the description of the computing elements. Accesspoints are considered identical if they share accesspoints attributes described below and are different only by the endpoint url. The identical accesspoints can be grouped in an accesspoints set. Every accesspoints set should be uniquely identified in the scope of a single resource.
 - **endpoint_urls**: the list of endpoint urls contained in the accesspoints set. Value is a list of strings.
 - **flavour**: implementation of accesspoint. For example 'ARC-CE', 'CREAM-CE', 'HTCondor-CE'. Value is a string.
 - **version**: version of the accesspoint implementation. Value is a tstring.
 - **shares**: list of shares which can be accessed via accesspoint set. Should be omitted if all shares of the resource can be accessed by the accesspoint set.Value is a list of strings.

Examples

1. Manchester example

```
"computingresources" : {
  "UKI-NORTHGRID-MAN-HEP-condor-arc-1": {
    "site": "UKI-NORTHGIRD-MAN-HEP",
    "type": "batch",
    "jobmanager": "condor",
    "jobmanager_version": "8.6.18",
    "status": "production",
    "number_logical_cpus": 1648,
    "capacity_HS06": 17963,
    "os" : "CentOS7.6",
    "default_walltime_minutes": 4300,
    "default_cputime_minutes": 4300,
    "default_memory_GB": 2,
    "assigned_vos": ["ATLAS", "ska", "icecube", "pheno"],
    "comment" : "free form string...",
    "publication_time":1540306937,
    "shares": {
      "long" : {
        "type": "queue",
        "max_walltime_minutes": 4300,
        "max_cputime_minutes":4300,
        "max_main_memory_GB": 3,
        "assigned_vos": ["atlas"]
      },
      "medium": {
        "type": "queue",
        "max_walltime_minutes": 1444,
        "max_cputime_minutes":1444,
        "max_main_memory_GB": 3,
        "assigned_vos": ["ska", "icecube", "pheno"]
      },
      "gpu": {
        "type": "queue",
        "max_walltime_minutes": 1444,
        "max_cputime_minutes":1444,
        "max_main_memory_GB": 12,
        "assigned_vos": ["atlas", "icecube"],
      },
    },
    "accesspoints": {
      "ARC5.4":{
        "endpoint_urls" :["gsiftp://ce01.tier2.hep.manchester.ac.uk:2811"],
        "flavour": "ARC-CE",
        "version": "5.4",
```

```

    },
  },
},
"UKI-NORTHGRID-MAN-HEP-condor-arc-2": {
  "site": "UKI-NORTHGIRD-MAN-HEP",
  "type": "batch",
  "jobmanager": "condor",
  "jobmanager_version": "8.6.18",
  "status": "production",
  "number_logical_cpus": 1024,
  "capacity_HS06": 11161,
  "os" : "CentOS7.6",
  "default_walltime_minutes": 1444,
  "default_cputime_minutes": 1444,
  "deafault_memory_GB": 2,
  "assigned_vos": ["atlas", "ska", "icecube", "pheno"],
  "comment" : "free form string...",
  "publication_time":1540306937,
  "shares": {
    "long" : {
      "type": "queue",
      "max_walltime_minutes": 4300,
      "max_cputime_minutes":4300,
      "max_main_memory_GB": 4,
      "assigned_vos": ["atlas"],
    },
    "medium": {
      "type": "queue",
      "max_walltime_minutes": 1444,
      "max_cputime_minutes":1444,
      "max_main_memory_GB": 4,
      "assigned_vos": ["ska", "icecube", "pheno"],
    },
  },
  "gpu": {
    "type": "queue",
    "max_walltime_minutes": 1444,
    "max_cputime_minutes":1444,
    "max_main_memory_GB": 12,
    "assigned_vos": ["atlas", "icecube"],
  },
},
"accesspoints": {
  "ARC5.4":{
    "endpoint_urls" :["gsiftp://ce02.tier2.hep.manchester.ac.uk:2811"],
    "flavour": "ARC-CE",
    "version": "5.4",
  },
},

```

```

    },
  },
  "UKI-NORTHGRID-MAN-HEP-torque-cream": {
    "site": "UKI-NORTHGIRD-MAN-HEP",
    "type": "batch",
    "jobmanager": "torque",
    "jobmanager_version": "2.5",
    "status": "production",
    "number_logical_cpus": 796,
    "capacity_HS06": 8676,
    "os": "SL6.8",
    "default_walltime_minutes": 4300,
    "default_cputime_minutes": 4300,
    "default_memory_GB": 2,
    "assigned_vos": ["ATLAS", "ska", "icecube", "pheno"],
    "comment": "free form string...",
    "publication_time": 1540306937,
    "shares": {
      "long": {
        "type": "queue",
        "max_walltime_minutes": 4300,
        "max_cputime_minutes": 4300,
        "max_main_memory_GB": 2,
        "assigned_vos": ["atlas"],
      },
    },
    "accesspoints": {
      "CREAM1.16.4": {
        "endpoint_urls": ["https://ce03.tier2.hep.manchester.ac.uk:8443"],
        "flavour": "CREAM-CE",
        "version": "1.16.4",
      },
    },
  },
}

```

2. CERN vcycle example

```

"computingresources" : {
  "vcycle-cern.lhcb.cern": {
    "site": "CERN-PROD",
    "type": "vacuum",
    "jobmanager": "Vcycle",
    "jobmanager_version": "01.00+pre12",
    "number_logical_cpus": 500,
    "default_walltime_minutes": xxx,
    "default_cputime_minutes": xxx,
    "memory_GB": xxx,
  },
}

```

```
"capacity_HS06":XXX,  
"assigned_vos": ["LHCB"],  
"publication_time":1545148291,  
"shares": {  
  "lhcb-vm-prod" : {  
    "max_walltime_minutes": 5000,  
  },  
  "lhcb-squid" : {  
    "max_walltime_minutes": 525600,  
  },  
  "lhcb-vm-dev" : {  
    "max_walltime_minutes": 5000,  
  },  
}
```

Accesspoints section is not needed in case of vcycle

```
}  
}
```