

## Storage Resource Reporting Specification V6

```
    {"storageservice": {
      "name": "FZK-ATLAS", // [1]
      "id": "global unique", // [2]
      "servicetype": "org.wlwg.se", // [3] WLCG agreement on
service types!
      "implementation": "dcache",
      "implementationversion": "2.13.51",
      "capabilities": ["data.access.file",
"data.access.streaming", "data.management.storage"], // [4]
      "qualitylevel": "production", // [5]
      "storagecapacity": {"online": {"totalsize": 11, "usedsize":
2, "reservedsize": 3},
      "offline": {"totalsize": 21, "usedsize": 4, "reservedsize":
5},
      "nearline": {"totalsize": -1, "usedsize": -1,
"reservedsize": -1 }
      }, [6]
      "message": "optional comment",
      "latestupdate": 1502719262 // timestamp when the storage
description has been published

      "storageendpoints": [
        {"name": "atlassrm", // [1]
          "id": "global unique", // [2]
          "endpointurl": "http://atlassrm-
fzk.gridka.de:8443/srm/managerv2",
          "interfacetype": "srm", // [7]
          "interfaceversion": "2.2",
          "capabilities": ["data.management.transfer",
"data.management.storage"], // [4]
          "qualitylevel": "production",
          "assignedshares": ["all"], // [8]
          "message": ""
        },
        {"name": "atlasgsiftp",
          "id": "global unique",
          "endpointurl": "gsiftp://f01-060-114-
e.gridka.de:2811",
          "interfacetype": "gsiftp",
          "interfaceversion": "1.0.0",
          "capabilities": ["data.management.transfer"],
          "qualitylevel": "production",
          "assignedshares": ["ATLASDATADISK"],
          "message": ""
        },
        {...}
      ]
      "storageshares": [
        {"name": "ATLASDATADISK" // [1]
          "id": "global unique", // [2]
          "policyrules": ["defaultpermission", "acl1",
"acl2"], // [9]
          "servingstate": "open", // [10]
          "accessmode": ["read/0", "delete/2"], // [11]
          "accesslatency": "online", // [12]
          "retentionpolicy": "replicated", // [13]
          "lifetime": {"default": "1 week", "maximum": ""},
          "expiration": "never"}, // [14]
```

```

        "timestamp" : 1502719262 // timestamp when the
measurement has been taken
        "totalsize": 26789588699,[15]
        "usedsize": 123564878, [15]
        "numberoffiles": -1,
        "path": ["/storage/data/atlasdatadisk/"],
"/disks/atlasdata/"],
        "assignedendpoints": ["all"], // [8]
        "vos" : ["atlas"],// list of the names of the VOs
which can use the storage share
        "message": ""
    },
    {"name": "ATLASSCRATCHDISK"
      "id": FZK-ATLAS.S2,
      "policyrules": ["scratchpermission"],

      "servingstate": "open",
      "accessmode": ["pilot read/0", "pilot write/3"],
      "lifetime": {"default": "1 week", "maximum": "1
year", "expiration": "release"}, // [14]
      "timestamp" : 1502719262 // timestamp when the
measurement has been taken
      "totalsize": 6789588699,
      "usedSize": 23564878,
      "numberoffiles": -1,
      "path": ["/storage/data/scratch/"],
      "assignedendpoints": ["atlassrm"],
      "vos": ["atlas"],// list of the names of the VOs
which can use the storage share
      "message": ""
    },
    {"name": "ATLASGROUPDISK"
      "id": FZK-ATLAS.S3,
      "policyrules": ["defaultpermission"],
      "servingstate": "closed",
      "accessmode": ["read/0"],
      "accesslatency": "online",
      "timestamp" : 1502719262 // timestamp when the
measurement has been taken
      "totalsize": 46789588699,
      "usedsize": 23564878,
      "numberoffiles": 1568747,
      "path": ["/storage/data/atlasgroup/"],
      "assignedendpoints": ["atlasgsiftp"],
      "vos" : ["atlas"],// list of the names of the VOs
which can use the storage share
      "message": ""
    },
    {...}
  ]
  "datastores": [
    {"name": "ATLAS Tape", // [1]
      "id": "global unique", // [2]
      "datastoretype": "tape", //[16]
      "accesslatency": "offline",
      "totalsize": 45784567485956578,
      "vendor": "some vendor name",
      "bandwith": {"network": "", "disk": "", "etc": ""},
// [17]

      "message": ""
    },
  ],

```

```

        ]
    }
}

```

Notes and definitions:

-----

**Depending on storage implementation, not all attributes described in the specification can be provided. The minimum set of attributes is marked by yellow colour**

[1] names are human-readable, to be chosen by the site managers, should be unique in the scope of storagendpoint or storageshares section

[2] id's (global unique) are generated, assigned by CRIC. no need to provide it on the SE

[3] servicetype enumeration values should be based on WLCG agreement. a possible convention could be something like this:

- org.wlcg.se: generic wlcg storage element
- org.wlcg.ce: generic wlcg CE
- org.wlcg.atlas.se: an atlas specific se (if such exists)
- org.wlcg.fts
- org.wlcg.pilotfactory

[4] capabilities enumeration values should be based on WLCG agreement. some possible values:

- data.access.files: capacity of providing access to flat files (direct io)
- data.access.streaming: capacity of providing streaming to data
- data.access.relational: capacity of providing access to a relational data source
- data.access.xml: capacity of providing access to an XML data source
- data.management.replica: capacity of managing the creation of file replicas upon request
- data.management.storage: capacity of managing a storage resource
- data.management.transfer: capacity of managing a transfer of files from the start to the completion
- data.management.3rdpartytransfer: capacity of managing a third party transfer of files from the start to the completion
- data.naming.resolver: capacity of resolving one name to another
- data.transfer: capacity of moving a file from one network location to another.

[5] qualitylevel values:

- development: The component/service/interface is under development, it is an early prototype system
- testing: The component/service/interface has completed the development phase and is under testing
- pre-production: The component/service/interface it is being used in real world scenarios as part of pre-production testing
- production: The component/service/interface is considered stable for real world scenarios

[6] storagecapacity values are to be calculated by CRIC

[7] interfacetype enumeration values should be based on WLCG agreement.

- possible values: srm, gridftp, xroot, https, file

[8] assignedshares, assignedendpoints: use either the keyword "all" or the name of the

assigned share/endpoint.

[9] policyrules: TODO. more details needed, maybe a simple policy language based on pre-defined permissions and roles to avoid hypercomplication.

[10] servingstate enumeration values:

- open: the share is in operation
- closed: the share is not available
- draining: the share is scheduled for maintenance/downtime and runs in a restricted mode

[11] accessmode values to be discussed. need to agree on e.g. howto specify the pilot read/webUI read, etc..

[12] accesslatency enumeration values:

- online: Files are always stored on a medium with a "minimal access time". Disk storage.
- nearline: Files are not always "immediately available". Typically a fully automated tape robot, but also a remote storage system could fit this qualification. Such a facility will need an unspecified amount of time to make a copy of the file available
- offline: Files may be stored on an offline component of the storage system, for example a tape library that is not connected to an automated tape robot.

[13] retentionpolicy key describes the reliability of storage, indicates the probability of the storage system loosing a file. possible values:

- none: no protection
- intermediate: some sort of replication is available
- replicated: the stored data is secured against data lost

[14] lifetime nested object describes the optional lifetime of files and the action the storage system may take upon the expiration of lifetime values. Empty values mean infinite lifetime.

Enumeration values for the expiration key:

- release: files on expiration will be removed by the storage system
- warn: storage system will not serve expired files instead it will return a warning.

Expired files may be deleted but should always be recoverable e.g. from backups.

- never: indicates that the files are stored with infinite expiration time.

[15] - both 'totalsize' and 'usedsize' refer to the logical file sizes, not taking replication factor into account.

[16] Datastore type values:

- disk: the storage is provided by magnetic disks
- tape: the storage is provided by magnetic tapes
- cloud: the storage is hosted in the cloud

[17] bandwith: TODO, describe the network bandwidth, IOPS, disk bandwidth in some object

**General comment regarding total and occupied size of the storage described via SRR**

The part describing total and occupied size of the storage is required for WLCG accounting and is mainly focused on the pledged resources. WSSA (WLCG Storage Space Accounting) collector should be able to generate from SRR at least used and total space per site per LHC experiment. This implies that published storage shares should not overlap in terms of space.

If site administrators do not want to expose information not related to the WLCG accounting via SRR, they should be able to configure SRR generation in a way that information non-related to the WLCG accounting is excluded.