

Storage Resource Reporting: Discussion on Topology Description

Dimitrios Christidis

Resources

Proposal document

<https://docs.google.com/document/d/1yzCvKpxsbcQC5K9MyvXc-vBF1HGpBk4vhjw3MEXoXf8/>

Extended example

https://twiki.cern.ch/twiki/pub/LCG/AccountingTaskForce/storage_service_v4.txt

Basic Structure

```
{  
  "storageservice": {  
    ... ③  
    "storageendpoints": [  
      ... ②  
    ],  
    "storageshares": [  
      ... ①  
    ]  
  }  
}
```

Storage Shares (1)

```
{  
  "name": "ATLASDATADISK",  
  "id": "global unique",  
  "policyrules": [ "defaultpermission", "acl1", "acl2" ],  
  "servingstate": "open",  
  "accessmode": [ "read/0", "delete/2" ],  
  "accesslatency": "online",  
  "retentionpolicy": "replicated",  
  "lifetime": { "default": "1 week", "maximum": "", "expiration": "never" },  
  "timestamp": 1502719262,  
  "totalsize": 26789588699,  
  "usedsize": 123564878,  
  "numberoffiles": -1,  
  "path": [ "/storage/data/atlasdatadisk/", "/disks/atlasdata/" ],  
  "assignedendpoints": [ "all" ],  
  "vos": [ "atlas" ],  
  "message": ""  
}
```

Storage Shares (2)

- name: For ATLAS and LHCb it *must* match the name of the corresponding SRM space token
 - Convention for others?
- usedsize, totalsize: Expressed in bytes
- numberoffiles: Optional but recommended
- timestamp: Of measurement, Unix time

Storage Shares (3)

- assignedendpoints: Match with entries from storageendpoints, can use "all"
- path: Potentially problematic
 - Why is it a list?
 - What if it doesn't match the location expected by the VO?
- vos: Not standardised

Storage Shares (4)

DPM:

```
"vos": [  
    "atlas/Role=lcgadmin",  
    "atlas/cz/Role=production",  
    "atlas/Role=production",  
    "atlas/cz",  
    "atlas",  
    ...  
]
```

EOS:

```
"vos": [  
    "atlas"  
]
```

→ What does it mean if there are multiple VOs?

Storage Shares (5)

- id: *Must not exist*
- policyrules, accessmode: Not defined
- servicingstate: open, closed or draining
- accesslatency: online, nearline or offline
- retentionpolicy: none, intermediate, replicated

Storage Shares (6)

- message: Purpose not defined

→ How are they supposed to be set by sites?

Storage Shares (7)

Mandatory:

name, usedsize, totalsize, path,
timestamp, vos, assignedendpoints

Optional:

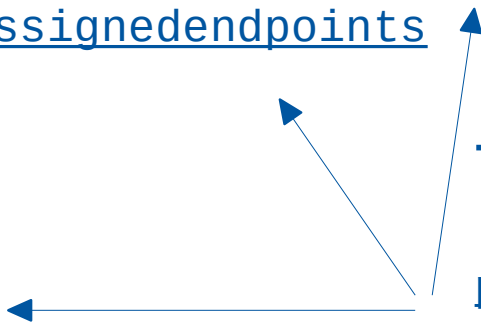
numberoffiles

Discarded:

id

TBD:

policyrules, servingstate,
accessmode, accesslatency,
retentionpolicy, lifetime, message



Storage Shares (8)

- How will CRIC differentiate between internal and leaf shares? Between Grid and non-Grid? Between pledged and non-pledged?

Storage Endpoints (1)

```
{  
  "name": "atlassrm",  
  "id": "global unique",  
  "endpointurl": "http://atlassrm-fzk.gridka.de:8443/srm/managerv2",  
  "interfacetype": "srm",  
  "interfaceversion": "2.2",  
  "capabilities": [ "data.management.transfer", "data.management.storage" ],  
  "qualitylevel": "production",  
  "assignedshares": [ "all" ],  
  "message": ""  
}
```

Storage Endpoints (2)

- id: *Must not exist*
- message: Purpose not defined
- storageendpoints: Match with entries from assignedendpoints, can use "all"
- name: No restrictions
 - Convention?

Storage Endpoints (3)

- qualitylevel: development, testing, pre-production or production
- interfacetype: srm, gridftp, xroot, https, file
 - http? dav?
- interfaceversion: Purposed not defined
 - Does it apply to all protocols?
 - What is the use case?

Storage Endpoints (4)

- capabilities: Some examples, not defined
 - How to set priorities?
- endpointurl:
 - Preference over the scheme?
e.g. srm:// vs httpg://, davs:// vs https://
 - Trailing slashes?

Storage Endpoints (5)

Mandatory:

name, endpointurl,
interfacetype, assignedshares

Discarded:

id

Optional:

TBD:

interfaceversion, capabilities,
qualitylevel, message

