

Table of Contents

WLCG sites Tape Metrics.....	1
Introduction.....	1
List of tape related metrics - for each VO.....	1
Reporting mechanism.....	2
CERN example.....	2

WLCG sites Tape Metrics

Introduction

This page contains documentation related to the WLCG sites Tape Metrics that have been agreed and should be periodically reported by all sites.

The current tape metrics dash board can be found here [↗](#) (CERN credentials required - click on CERN SSO).

List of tape related metrics - for each VO

- **Amount of data stored** (in Bytes, rounded to decimal GB)
 - ◆ The amount of actual data of the given VO stored on tape (including multiple copies). Only referenced data is included, holes due to deletion and remaining free space on tapes are excluded.
 - ◆ Published as `usedsize`.
 - ◆ Corresponds to "Used" in CRIC.
 - ◆ 1 GB = 1000000000 bytes [↗](#)
 - ◆ Example of rounding to GB using `round(number[, ndigits])` function: `round(23125065994338304, -9) = 23125066000000000`
- **Amount of space occupied** (in Bytes, rounded to decimal GB)
 - ◆ The amount of tape space occupied by the VO data (including multiple copies). It is the sum of the capacities of all the tapes allocated to the experiment - data, holes after deleted data and free remaining space are treated equally.
 - ◆ Published as `occupiedsize`
 - ◆ Corresponds to "Total" in CRIC.
 - ◆ 1 GB = 1000000000 bytes [↗](#)
 - ◆ Example of rounding to GB using `round(number[, ndigits])` function: `round(62715387334098944, -9) = 62715387000000000`
- **READ bytes 24h** (in Bytes) = amount of data READ by the given VO in the past 24 hours
 - ◆ Example: 530599479
- **WRITE bytes 24h** (in Bytes) = amount of data WRITTEN by the given VO in the past 24 hours
 - ◆ Example: 3927215943923
- **Average Tape Remounts 24h** (count, rounded to 2 decimal points): period = past 24h
 - ◆ Formula to use: for a given VO, sum how many times each tape was mounted in the past 24 hours (group by tape), calculate the average
 - ◆ Formula in shell script: `awk '{print $COLUMN_WITH_TAPE}' PAST_24_HOURS_TAPE.VO.LOG | sort | uniq -c | awk '{ sum += $1; n++ } END { if (n > 0) print sum / n; }'`
 - ◆ Example: For VOx, in the past 24 hours, tape T1 was mounted once, T2 5 times and T3 twice. Average remount (rounded to 2 decimal points) is: $(1 + 5 + 2) / 3 = 8 / 3 = 2.67$
 - ◆ Ideal value = 1. It means that for a given VO, for the past 24 hours, all tapes were mounted only once as all the (file) requests were properly grouped together.
- **Total Mounts 24h** (count) : period = past 24h
 - ◆ This is a raw value of the all mounts in the past 24 hours to calculate the above mentioned average. It is needed in order to let the dashboard GUI to properly calculate average for multiple sites.
- **Unique Mounts 24h** (count) : period = past 24h
 - ◆ This is a raw value of unique mounts in the past 24 hours to calculate the above mentioned average. It is needed in order to let the dashboard GUI to properly calculate average for multiple sites.

Reporting mechanism

Data from all sites will be collected by a *pull* mechanism as part of the WLCG Space Accounting system[?].

Each site should provide a URL with data stored in the following valid JSON file format:

```
{
  "storageservice": {
    "name": "SITE-Tape",
    "implementation": "APPLICATION_NAME",
    "implementationversion": "APPLICATION_VERSION",
    "latestupdate": TIMESTAMP_OF_THE_REPORT,
    "storageshares": [
      {
        "name": "VO_NAME",
        "usedsize": BYTES_ROUNDED_TO_GB,
        "occupiedsize": BYTES_ROUNDED_TO_GB,
        "readbytes24h": BYTES,
        "writebytes24h": BYTES,
        "avgtaperemounts": COUNT_2_DECIMAL,
        "totalmounts24h": COUNT,
        "uniquemounts24h": COUNT,
        "timestamp": TIMESTAMP_OF_THE_SUMMARY,
        "vos": ["vo_name"]
      },
      :
      :
      :
      {
        "name": "VO_NAME",
        "usedsize": BYTES_ROUNDED_TO_GB,
        "occupiedsize": BYTES_ROUNDED_TO_GB,
        "readbytes24h": BYTES,
        "writebytes24h": BYTES,
        "avgtaperemounts": COUNT_2_DECIMAL,
        "totalmounts24h": COUNT,
        "uniquemounts24h": COUNT,
        "timestamp": TIMESTAMP_OF_THE_SUMMARY,
        "vos": ["vo_name"]
      }
    ]
  }
}
```

CERN example

As of 3 October 2018, the JSON file reported by CERN looked like this:

```
{
  "storageservice": {
    "name": "CERN-PROD-Tape",
    "implementation": "CASTOR",
    "implementationversion": "2.1.17-36",
    "latestupdate": 1538573402,
    "storageshares": [
      {
        "name": "ALICE",
        "usedsize": 34626717000000000,
        "occupiedsize": 34525146000000000,
        "readbytes24h": 732092431,
        "writebytes24h": 30812221455,
        "totalmounts24h": 13,
        "uniquemounts24h": 10,
      }
    ]
  }
}
```

TapeMetricsJSON < HEPTape < TWiki

```
"avgtaperemounts": 1.30,  
"timestamp": 1538559894,  
"vos": ["alice"]  
},  
{  
  "name": "ATLAS",  
  "usedsize": 77645926000000000,  
  "occupiedsize": 79737751000000000,  
  "readbytes24h": 201113384166,  
  "writebytes24h": 66111329391336,  
  "totalmounts24h": 111,  
  "uniquemounts24h": 96,  
  "avgtaperemounts": 1.16,  
  "timestamp": 1538559894,  
  "vos": ["atlas"]  
},  
{  
  "name": "CMS",  
  "usedsize": 77480294000000000,  
  "occupiedsize": 87940521000000000,  
  "readbytes24h": 71073577934446,  
  "writebytes24h": 14439964639374,  
  "totalmounts24h": 2002,  
  "uniquemounts24h": 1254,  
  "avgtaperemounts": 1.60,  
  "timestamp": 1538559894,  
  "vos": ["cms"]  
},  
{  
  "name": "LHCb",  
  "usedsize": 27784286000000000,  
  "occupiedsize": 27752429000000000,  
  "readbytes24h": 0,  
  "writebytes24h": 11720658493973,  
  "totalmounts24h": 16,  
  "uniquemounts24h": 7,  
  "avgtaperemounts": 2.29,  
  "timestamp": 1538559894,  
  "vos": ["lhcb"]  
}  
]  
}  
}
```

-- VladimirBahyl - 2018-01-19

This topic: HEPTape > TapeMetricsJSON

Topic revision: r10 - 2020-11-13 - OliverKeeble



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback