





WLCG Space Accounting in the SRM-less world

CHEP 2018, Sofia, 09.07.2018

Julia Andreeva on behalf of the team contributing to the WLCG
Storage Space Accounting project
CERN

Motivation

- Lack of global storage space accounting for the WLCG infrastructure
- Existing solutions are working in the scope of a single experiment or a given infrastructure (EGI APEL-based)
- Looking for a common solution which can work across experiments, Grid infrastructures and can be used for operations and accounting purposes
- Gradually get rid of the SRM dependencies



Directions of work (1)


- Enable storage space topology description
- Enable possibility to query storage space accounting information for all kinds of storage implementation
- Implementation of the WLCG Storage Space Accounting service (WSSA) which implies data collection, storage, processing and visualization

Directions of work (2)

- Enable storage space topology description
- Enable possibility to query storage space accounting information for all kinds of storage implementation

First two goals are being addressed by the joined effort of the WLCG Data Steering group in collaboration with storage providers and the WLCG Accounting Task Force. The technical proposal is described in detail in Storage Resource Reporting (SRR) document. Presented at approved at the Grid Deployment Board. The initial proposal considered disk storage space accounting. Implementation is in progress.

Storage topology description

- The goal is to provide description of the storage services with space quotas (storage shares) which have to be accounted separately. Storage shares should be independent with no physical overlap in terms of space. Aggregation of occupied/free space of storage shares of a given storage service should provide a complete view of the storage service capacity.
- According to SRR every storage should provide an URL with the file in JSON format.
- URL will be recorded as an attribute of the storage service in GocDB/OIM
- It is foreseen that Computing Resource Information Catalogue (CRIC) which is currently under development will collect and record storage topology description and will serve as an information provider for global WLCG view. For more details see  [presentation](#)

“Space quota” level resource reporting

- Requirement : Storage systems should provide total used and total free space for all distinct *space quotas* available to the experiment through a non-SRM protocol.
 - Query frequency - order of minutes
 - Accuracy – order of tens of GB depending on storage implementation
 - Freshness – tens of minutes
- At least one non-SRM protocol (gridFTP, HTTP and xrootd) should be enabled for storage occupancy queries (used/free if possible)
- An alternative solution is a JSON file as for topology description complemented with accounting data


SRR implementation

- Implementation by all storage providers is ongoing.
 - EOS – done. Already enabled at CERN
 - DPM – enabled in DPM 1.10. Needs special configuration performed by the sites. Will be followed up by SRR deployment campaign
 - dCache – work in progress
 - Other storage providers are also planning to work on SRR implementation

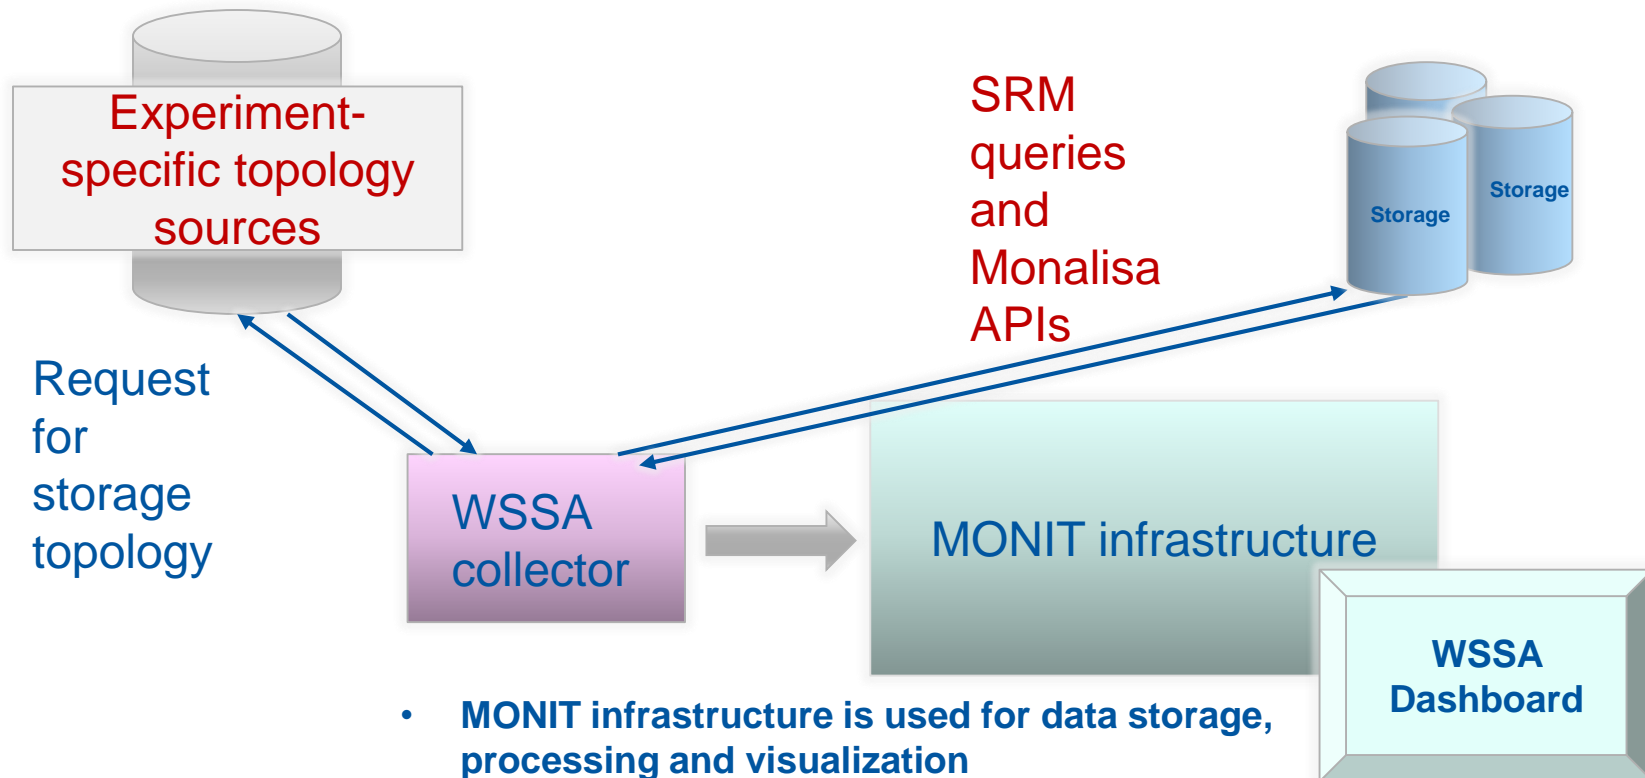
Directions of work (3)

- CRIC and SRR are work in progress
 - Deployment campaign for SRR might take quite some time
 - The system is constructed in a way to make progress using existing sources and to switch transparently to CRIC and SRR as they are in place
-
- Implementation of the WLCG Storage Space Accounting service (WSSA) which implies data collection, storage, processing and visualization

WSSA Data Flow




WSSA Current Implementation



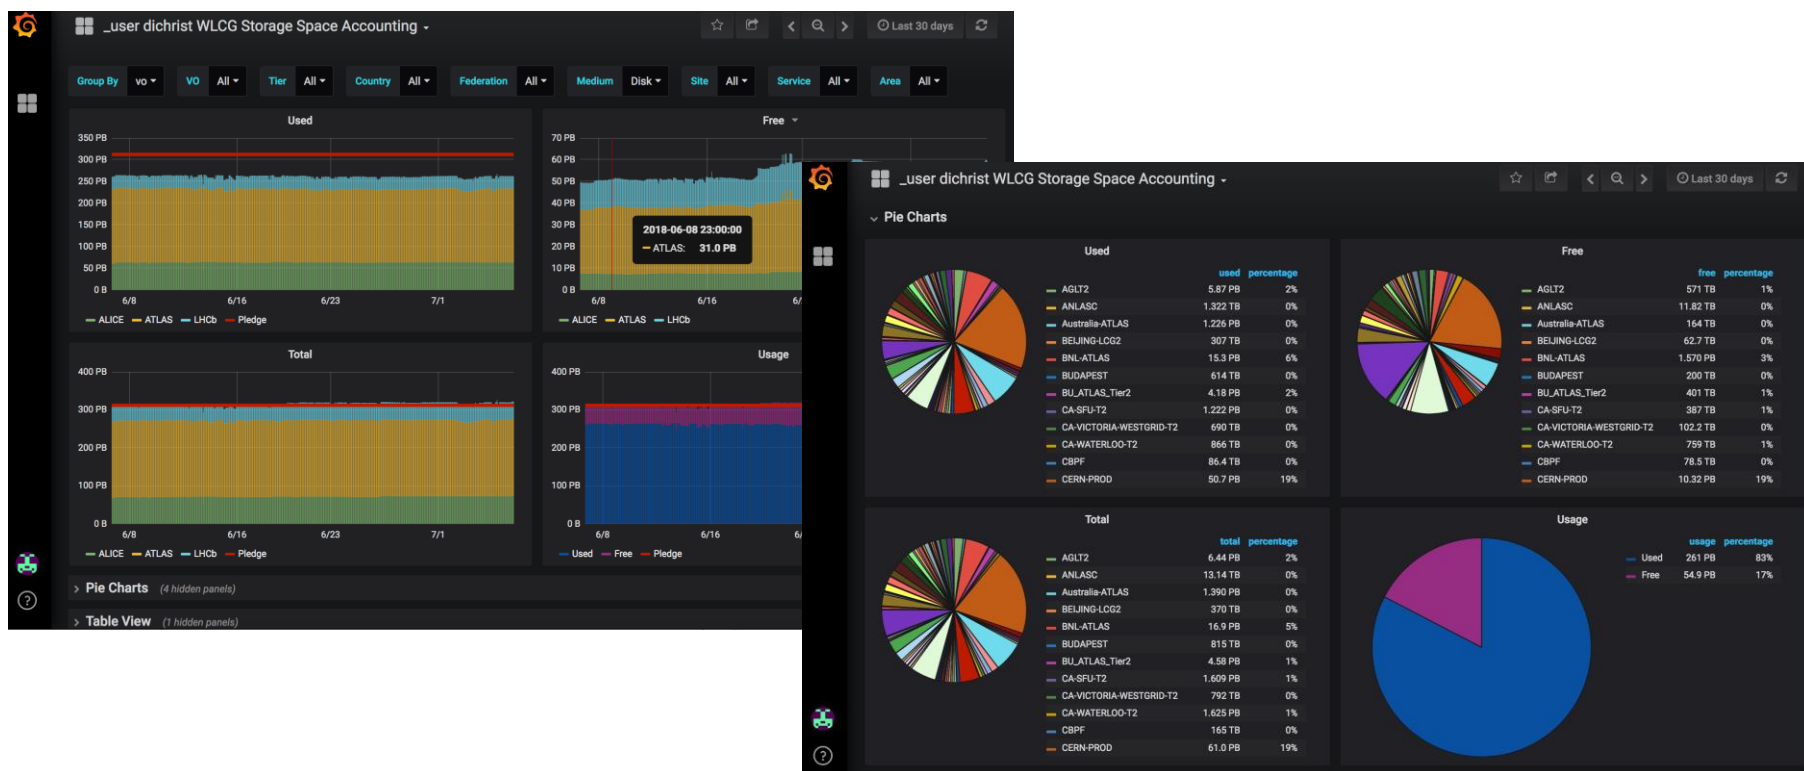
- **MONIT infrastructure** is used for data storage, processing and visualization
- Data is stored on HDFS, Elasticsearch and InfluxDB
- **WSSA Dashboard** is implemented in Grafana and uses InfluxDB storage backend

WSSA Implementation



WSSA status (disk storage)

- Space accounting for disk storage is provided in WSSA prototype for ALICE, ATLAS and LHCb. Work for CMS is ongoing.

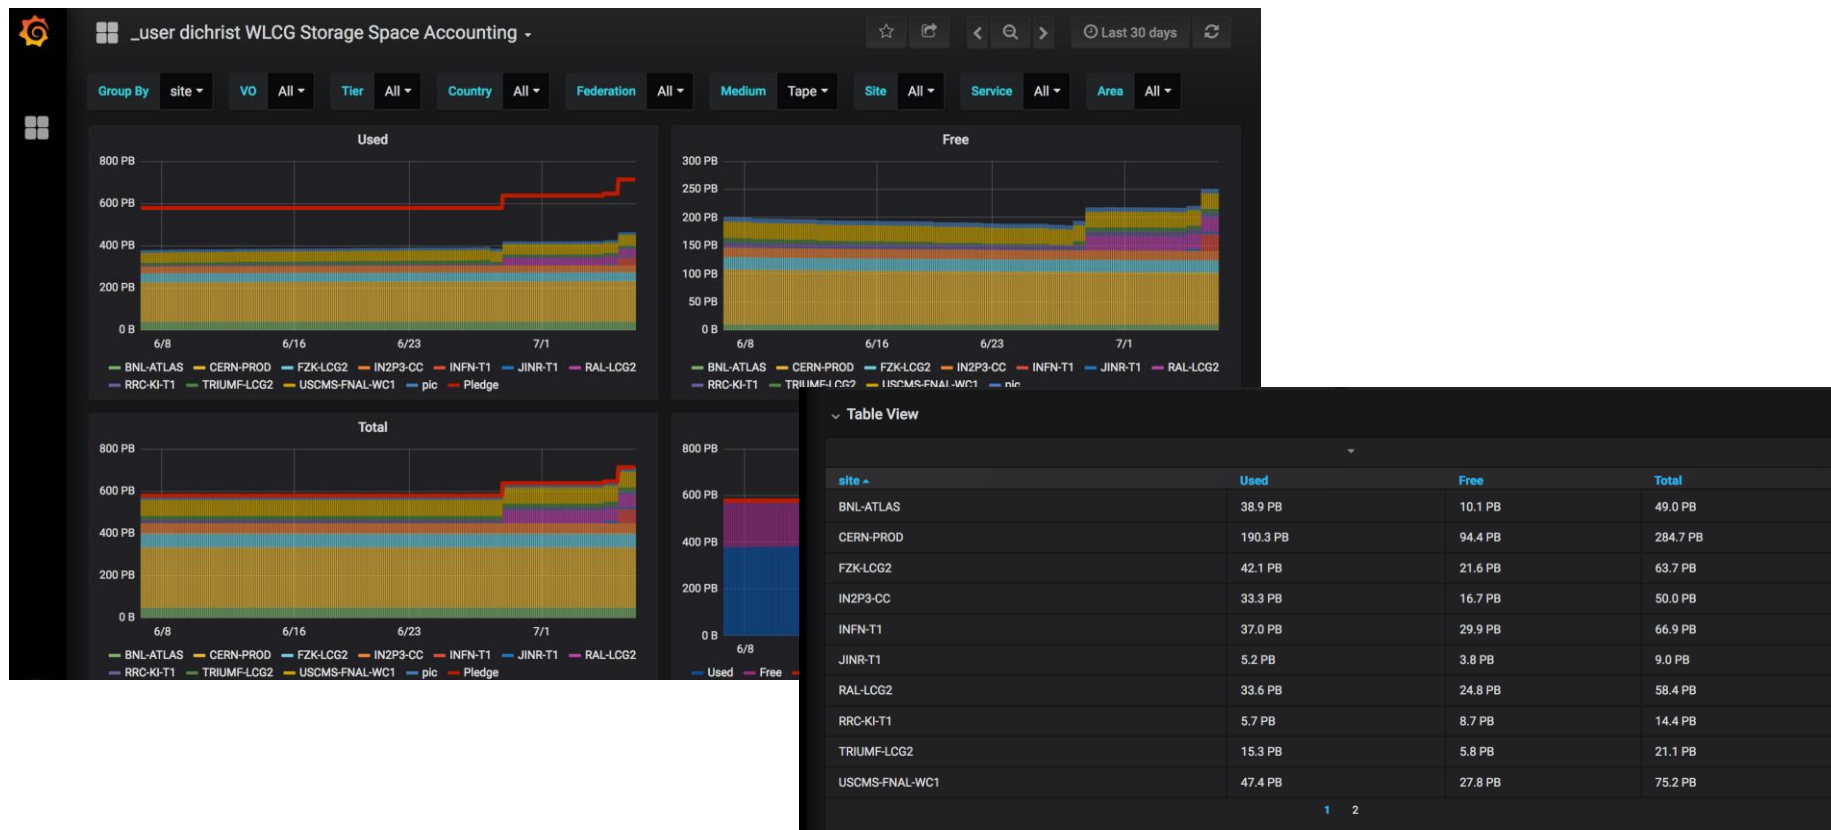


Storage Space Accounting for tapes


- In collaboration with the WLCG Archival Storage Working Group the progress has been done for tape storage accounting
- File format very similar to SRR
- Pull mechanism over HTTP
- Can support any number of metrics, not just limited to the accounting data, but for other tape metrics which can help to understand and optimize tape usage for the LHC workflows

WSSA Status for tape storage

- Implemented for 11 sites out of 14 running tape storage. 3 missing ones are on the way



WSSA. Disk vs tape



Next steps

- Enable WSSA for CMS
- SRR implementation and deployment which requires effort of many people of the storage development teams , WLCG operations and sites.
- Deployment of WSSA to production (planned for autumn this year)
- Integration of the WLCG monthly accounting report generation with WSSA

Collaborative effort

- This work requires contribution of many people for development, deployment and validation.
- Would like to highlight important contribution of Oliver Keeble and Alessandro Di Girolamo (authors of SRR proposal) and Dimitrios Christidis (WSSA service developer)

Conclusions

- WSSA provides a global view of storage occupancy for the WLCG infrastructure both for disk and tape storage
- Current implementation is using available methods for getting information from the primary sources but we are moving forward to deploy SRM-free storage accounting
- WSSA prototype is stable and ready for validation. Will move to production early autumn
- Please, give a try and provide your feedback:
<http://cern.ch/go/B8vw>