





# WLCG Space Accounting

MB 20.09.2018

Julia Andreeva

# Motivation

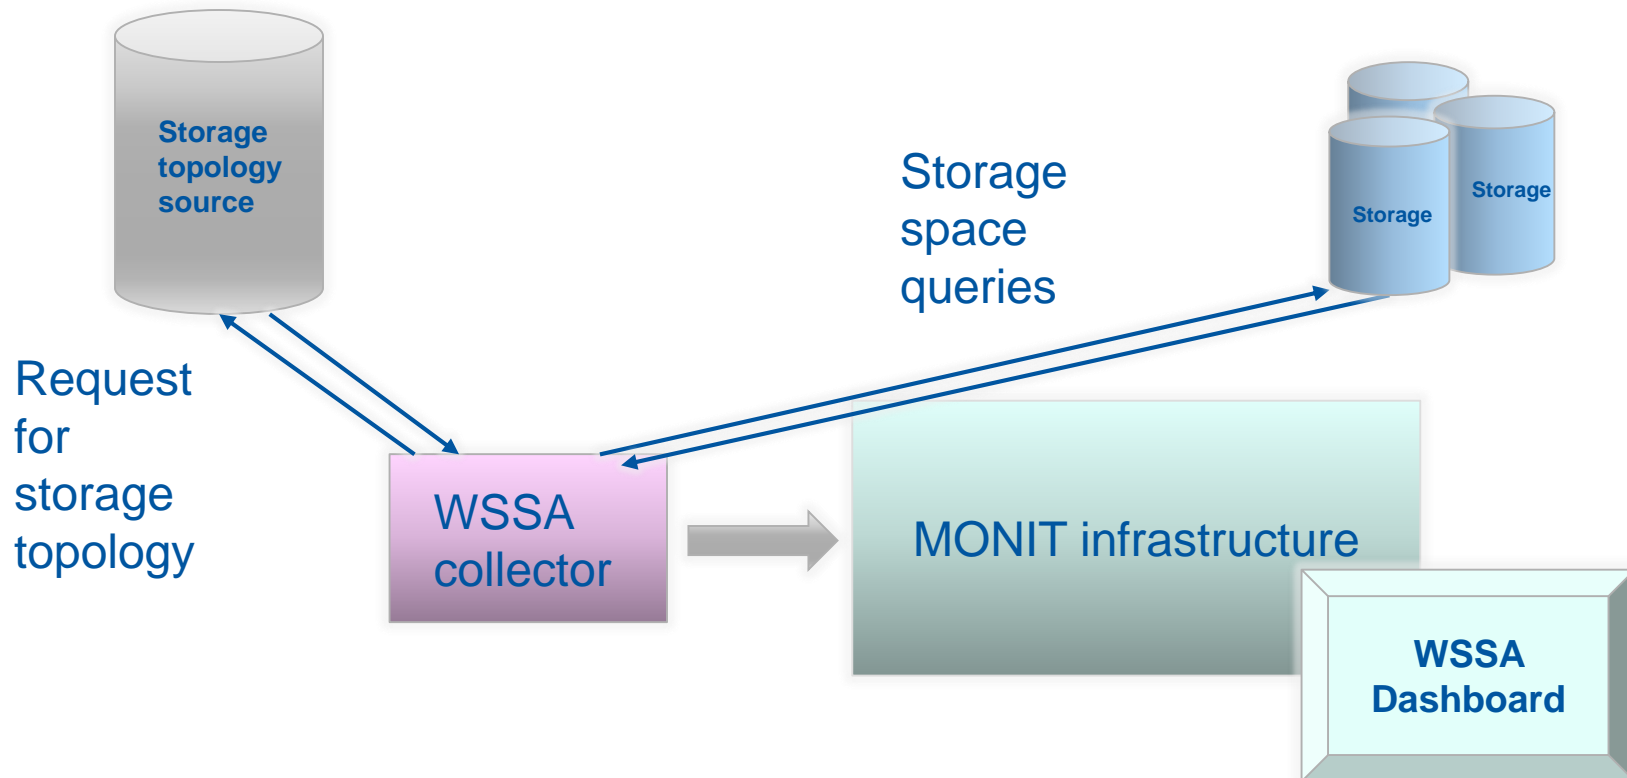
- Lack of global storage space accounting for the WLCG infrastructure
- Existing solutions are working in the scope of a single experiment or a given infrastructure (EGI APEL-based)
- Looking for a common solution which can work across experiments, Grid infrastructures and can be used for operations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



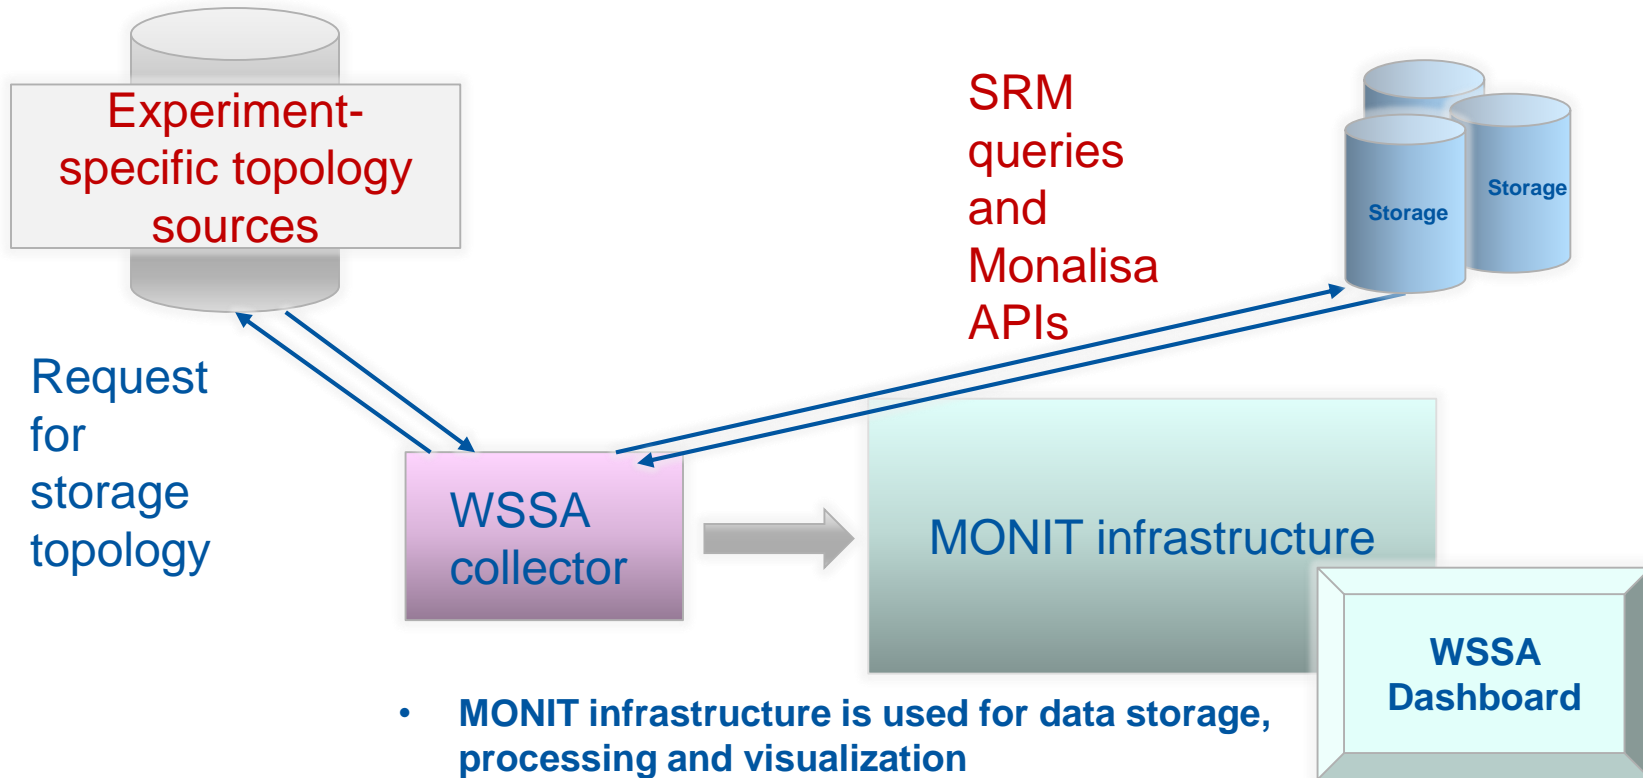
# Directions of work

- 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**

# WSSA Data Flow

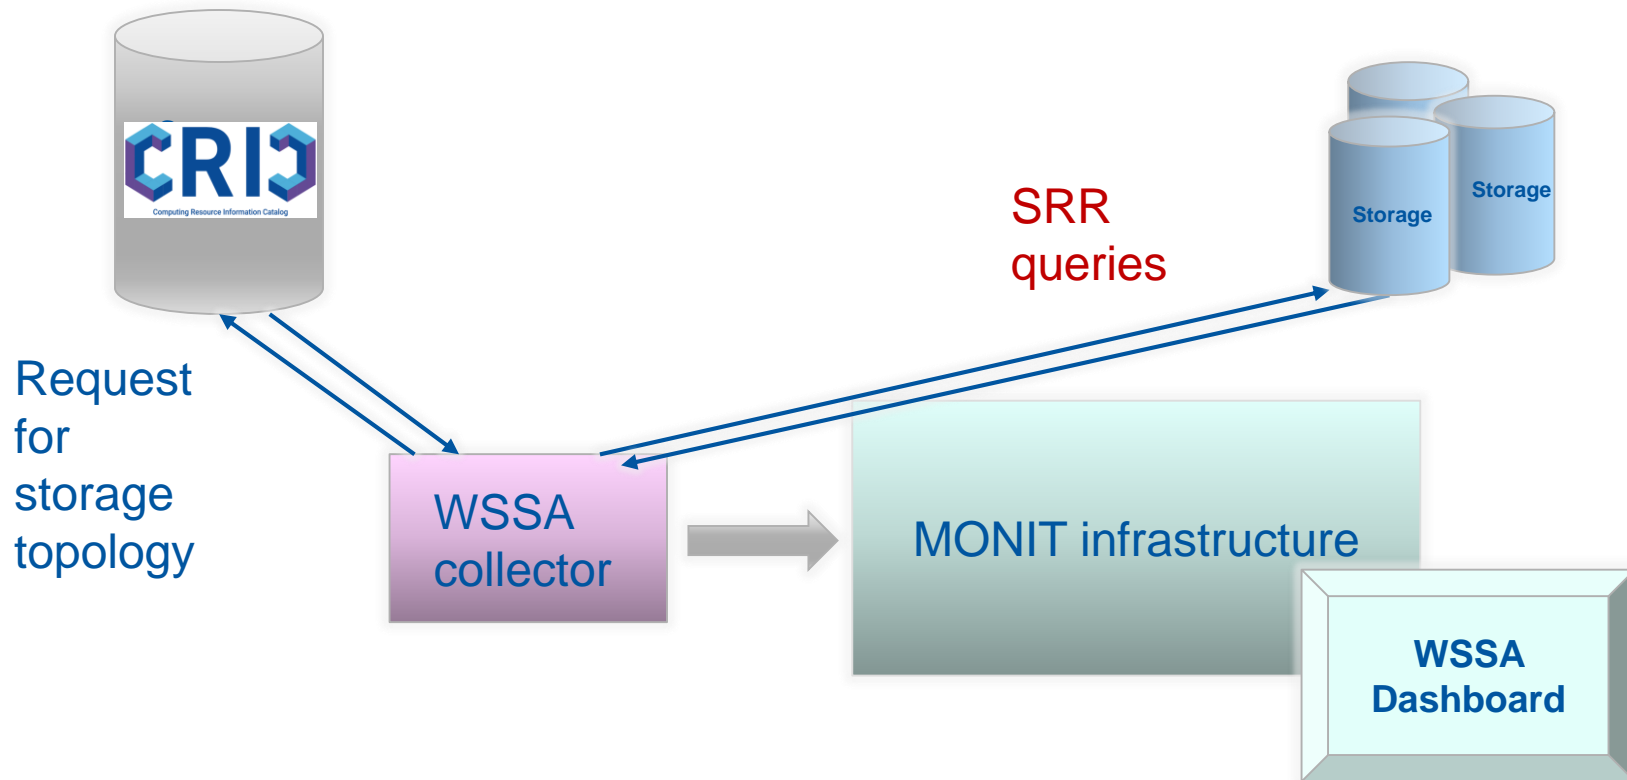


# WSSA Current Implementation



- **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



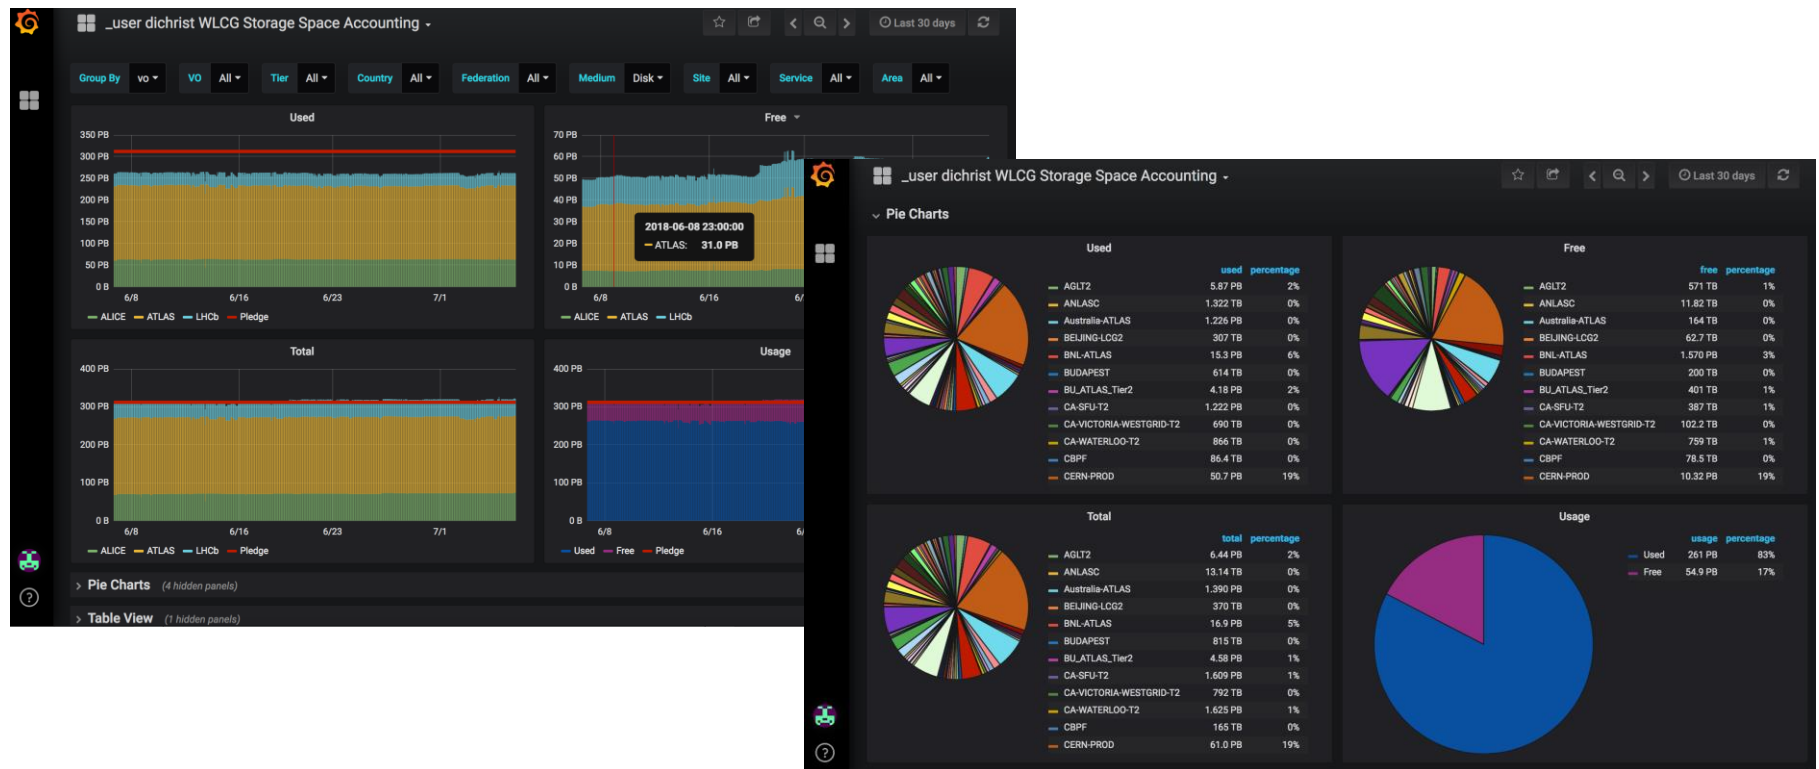
# Status

- WSSA prototype is running since spring 2018. First round of validation is done. Will move to production by the end of October.
- Provides a global view of storage occupancy for the WLCG infrastructure both for disk (ALICE, ATLAS and LHCb) and tape (all 4 experiments) storage
- Current implementation is using available methods for getting information from the primary sources (SRM for ATLAS and LHCb, queries to ALICE storage space accounting system based on xrootd queries) but we are moving forward to deploy SRM-free storage accounting
- For CMS we intended to use Dynamo, however comparison of Dynamo data with sites where EGI storage accounting is enabled, showed that Dynamo numbers are very low and might not be reliable. So either we use EGI data for sites where it is available, or wait for SRR implementation at the sites
- There are some known limitations of the UI based on Grafana, which can not be solved without deploying custom plugins on the MONIT infrastructure. Monitoring team does not support this idea. These limitations are not a showstopper



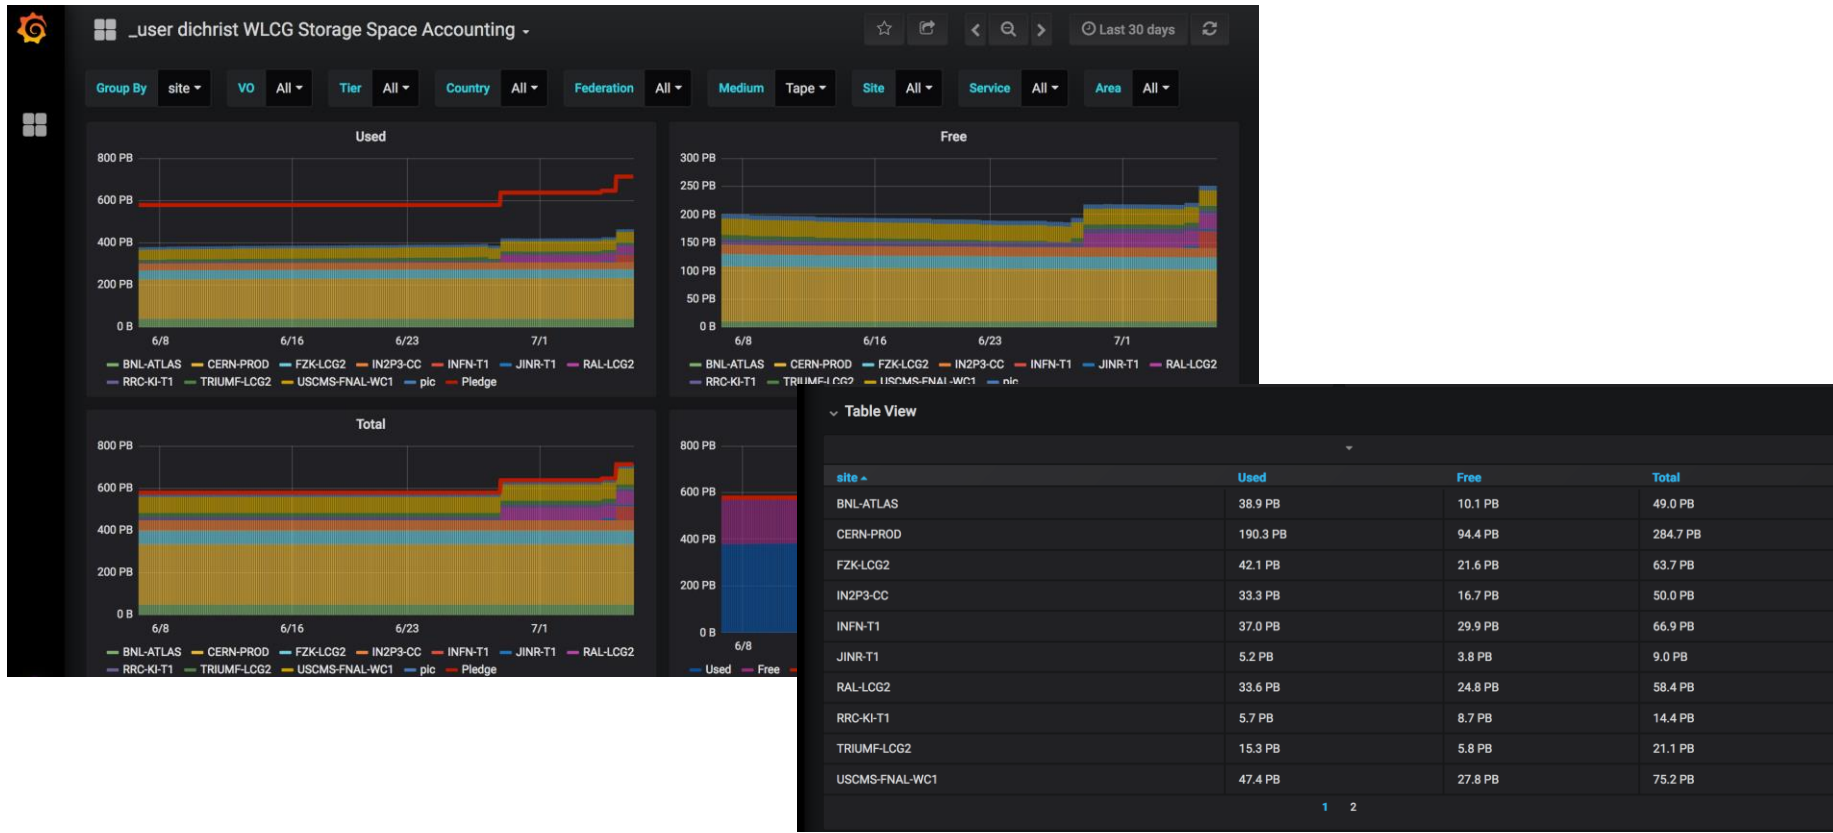
# WSSA status (disk storage)

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

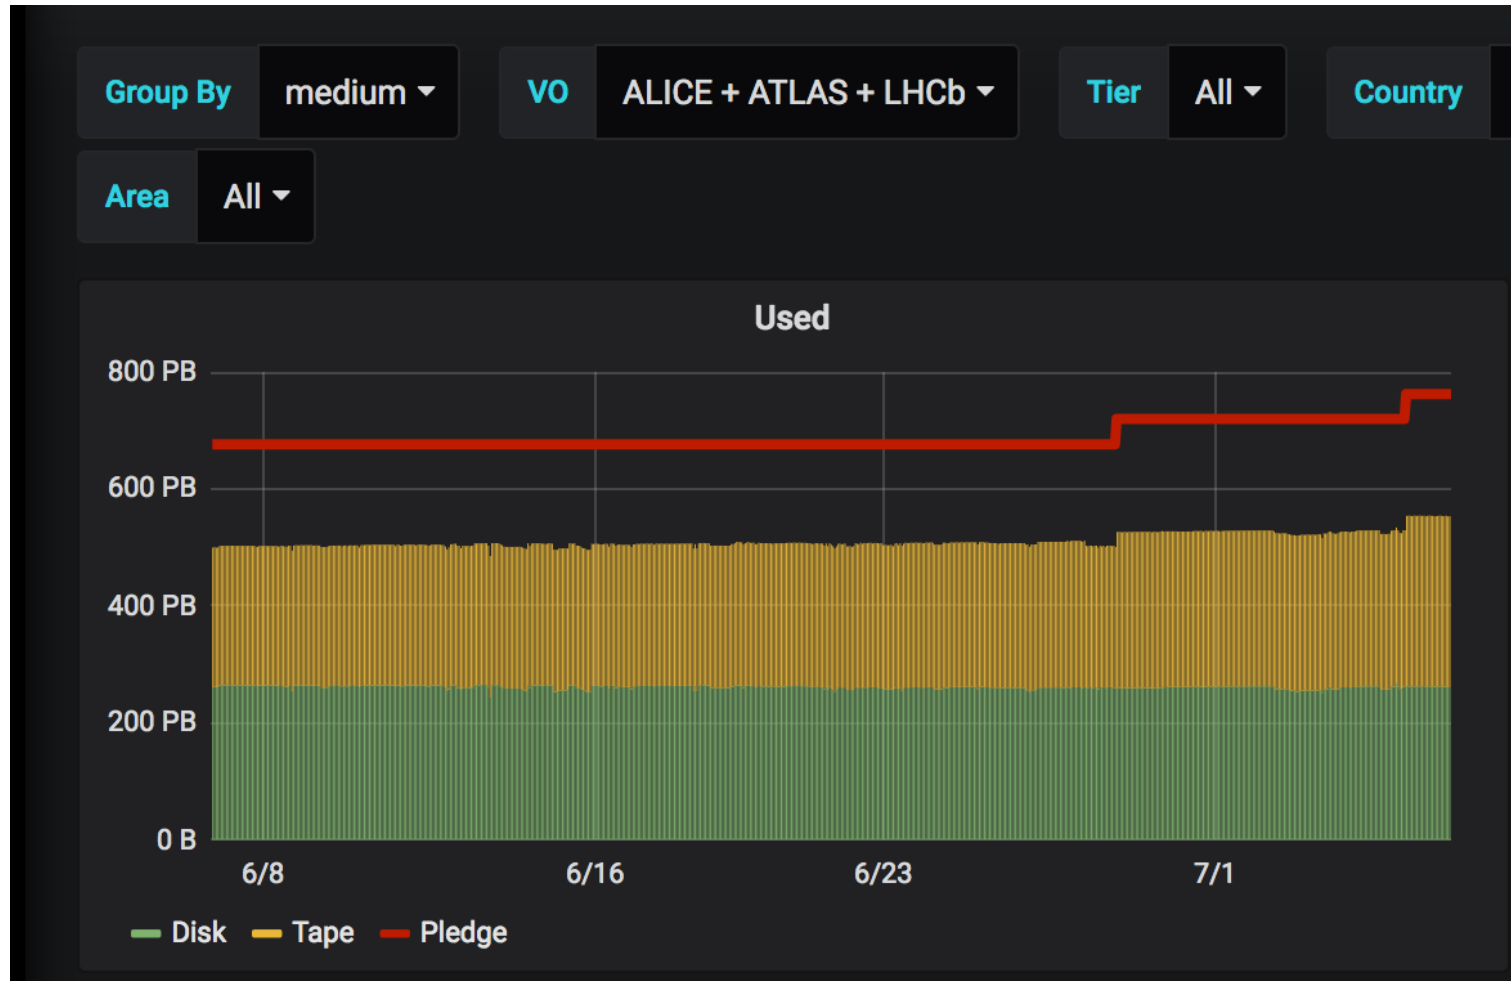


# WSSA Status for tape storage

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



# WSSA. Disk vs tape



# Still to be done

- Deployment based on puppet
- Documentation
- Move to Grafana public which is not behind CERN SSO, not sure whether this is needed
- Switch to CRIC for storage topology when CRIC is ready. Switch to CRIC will allow to distinguish between pledged and non-pledged storage shares and therefore show both pledged and non-pledged storage resources in SSA. Currently only pledged ones are accounted by SSA.
- Integrate SSA information in generation of monthly accounting reports performed by EGI accounting portal
- Follow up SRR implementation/deployment and gradually integrate as available

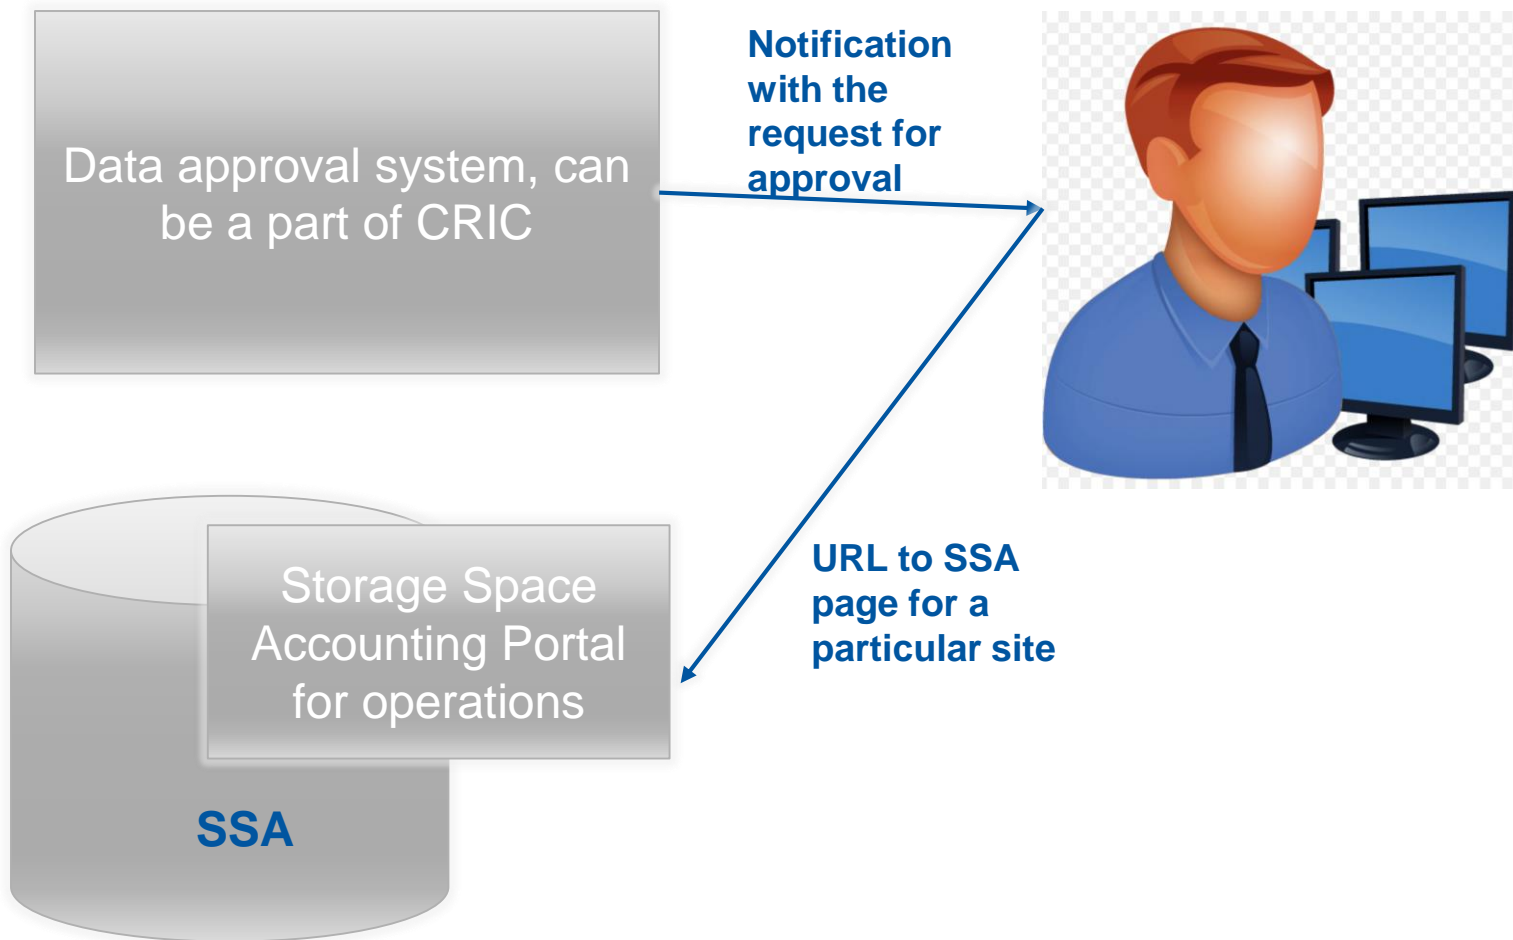
# Validation

- Performed validation comparing SSA results with ATLAS storage accounting Dashboard and EGI storage space accounting for the sites where EGI storage space accounting is deployed. Good agreement.

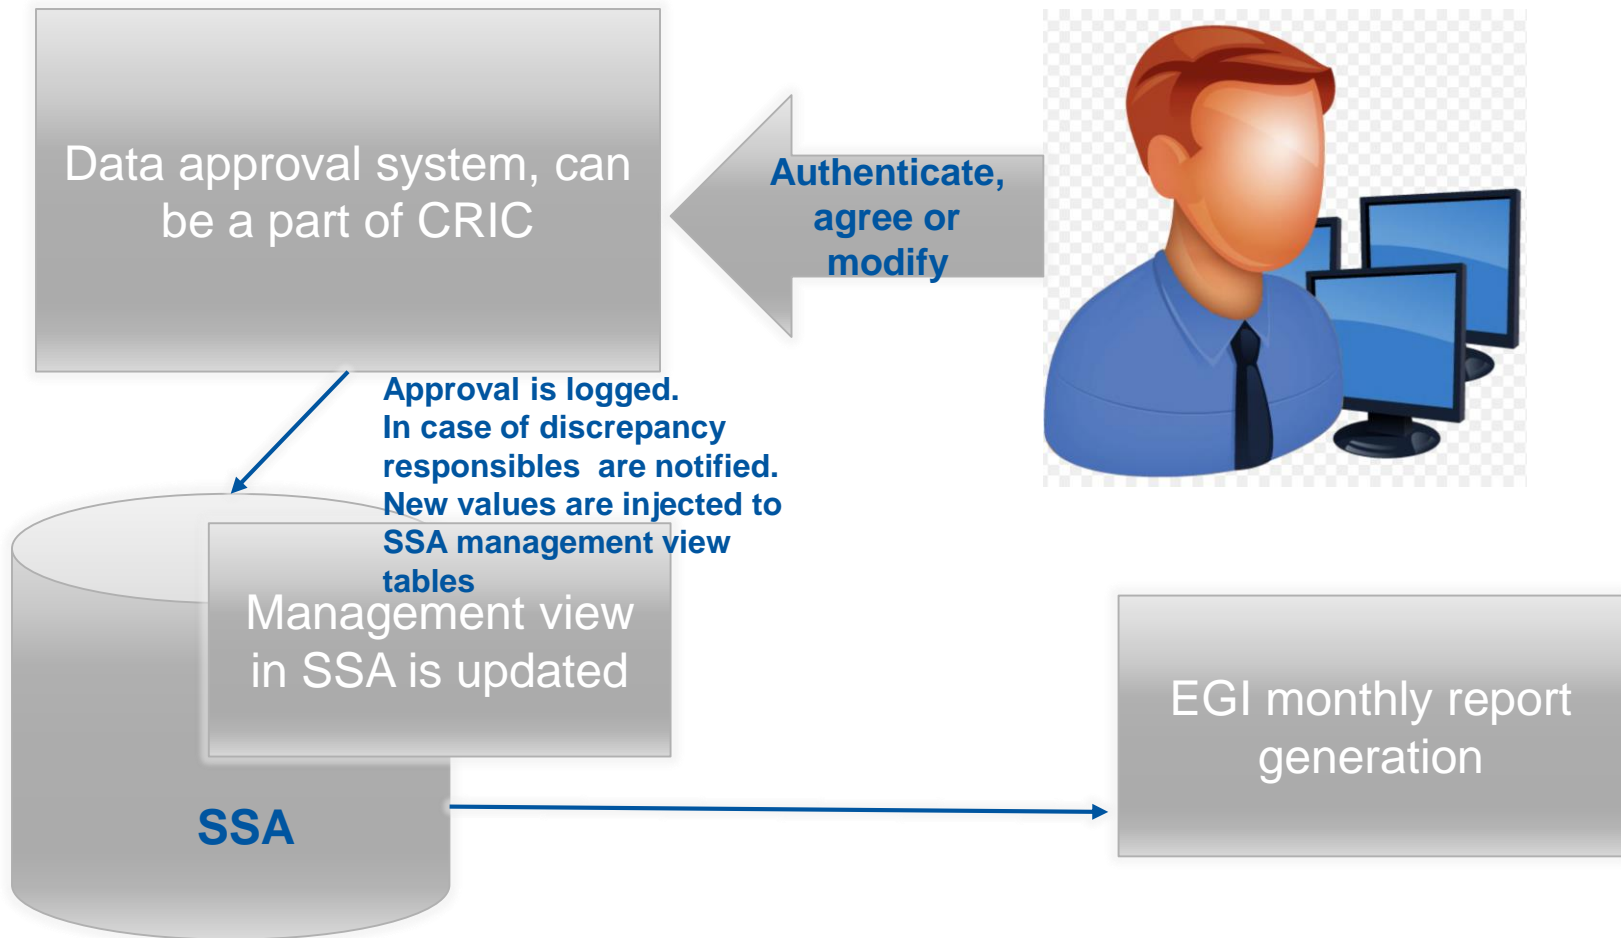
# Proposal for data approval workflow (1)

- While discussing WSSA at the WLCG Accounting Task Force meeting, Brian Bockelman suggested that we provide a way for sites to approve accounting values and create a sort of “Management view” for storage space accounting

# Proposal for data approval workflow. Possible scenario.(1)



# Proposal for data approval workflow. Possible scenario. (2)






# Advantages

- No mail exchange between site admins and WLCG project office
- Corrections are applied immediately
- The whole procedure is properly logged
- Easy to follow up on any inconsistencies between automatically generated numbers and those which are provided by site admins
- Similar workflow can be implemented for other monthly reports metrics (availability and accounting)

# Backup slides

# 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