



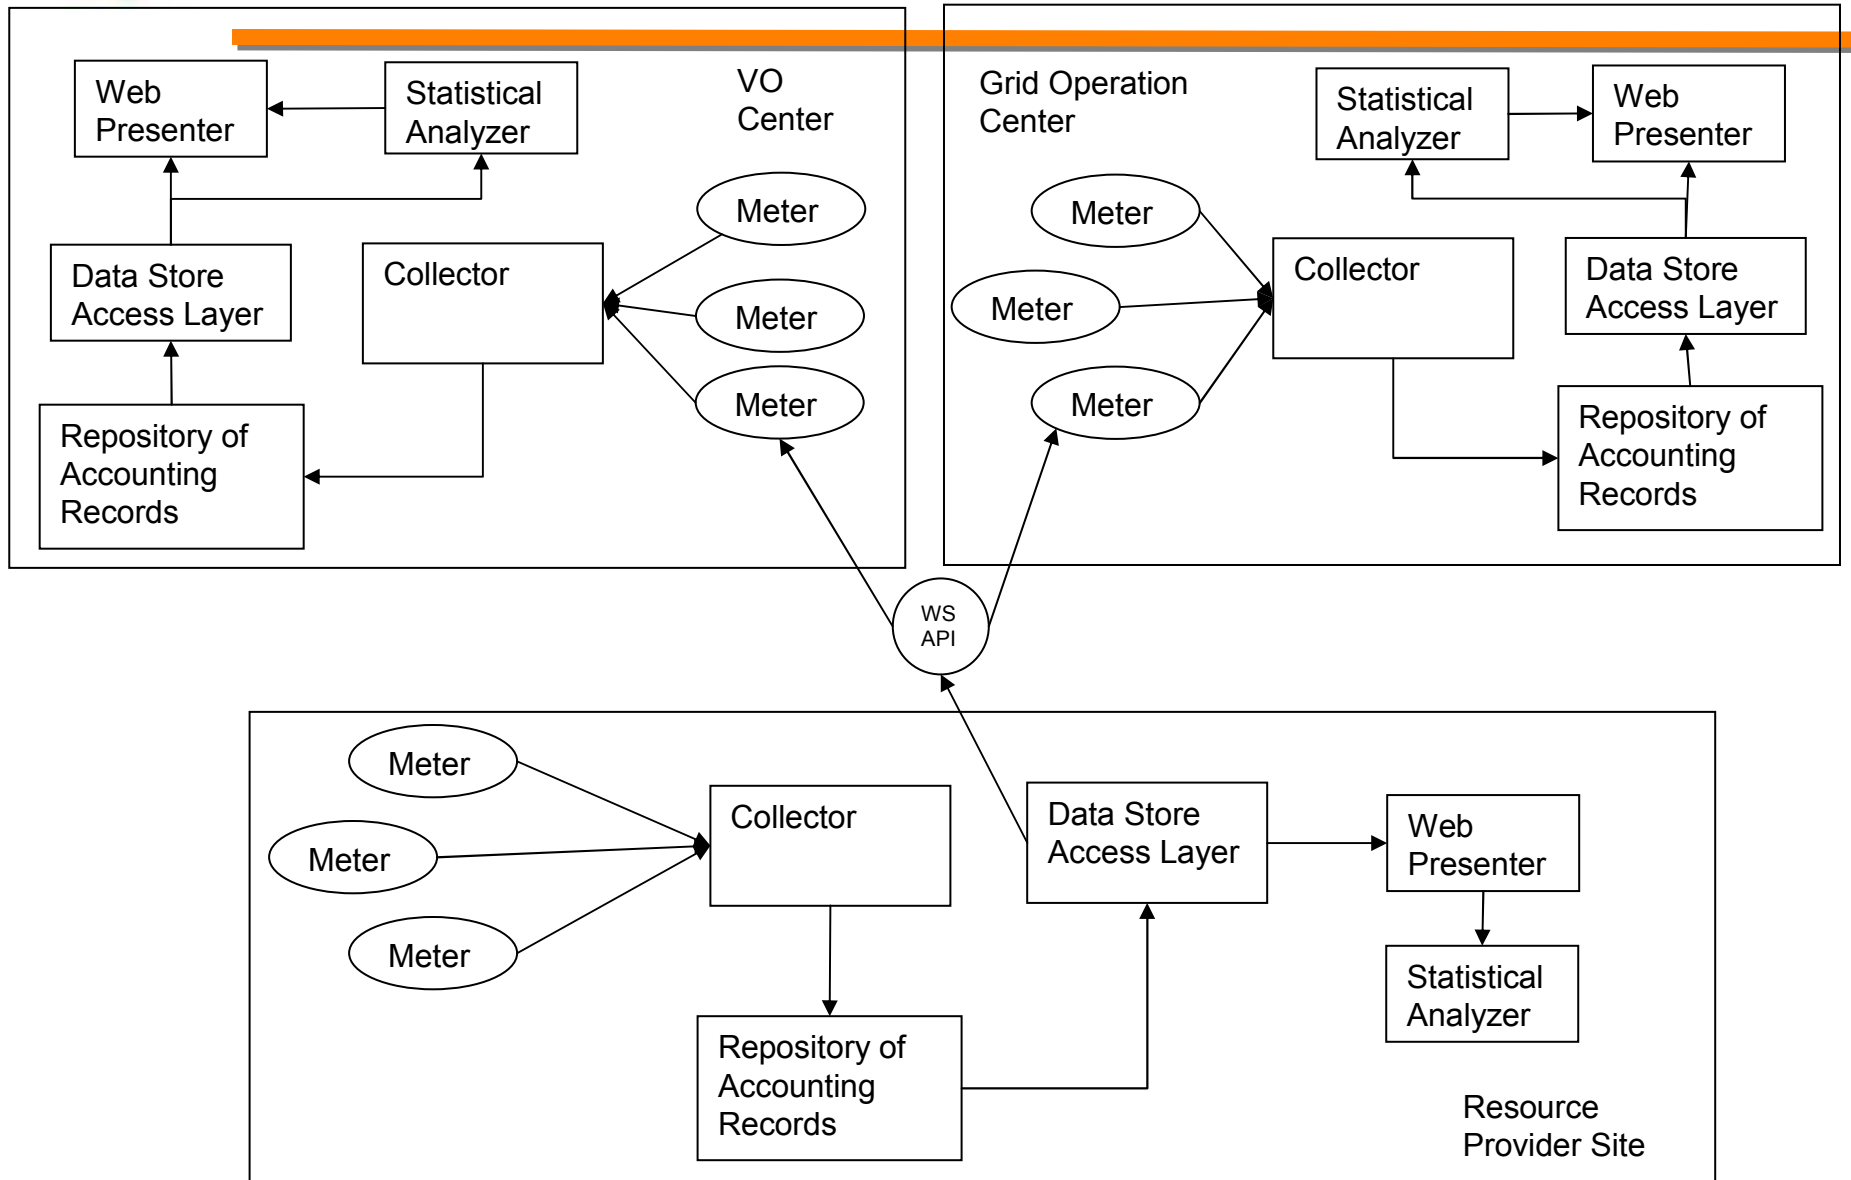
Open Science Grid


OSG Accounting System Design Overview

Sudhir Borra

Philippe Canal


Matteo Melani





Collectors

- Main Functionalities
 - Keep track of which meter are reporting
 - Receive information from the Meters
 - Correlate or complete the information
 - For example looking up which VO a UID is mapped to.
 - Apply filters and policies
 - Store the information in the ‘local accounting data store’



Meters

- Ideally they would be the data producers themselves
- Fall back solution is the typical log parsing with all its drawbacks
 - When is it run? Where is the log file? How long is it kept? Does it really contain all the information needed? etc.
- Another alternative is active query
 - condor_q for example
 - but there are performance issues in some cases
- In either cases, the ‘Meter’ should only depend
 - on the ‘tool’ (condor, pbs, etc.)
 - and the simple, lightweight Accounting Interface Library
 - (and should be eventually distributed with the tool)



Meter – Collector Connection

- A Meter uses a secure channel to send usage information to the Collector, the usage data header contains:
 - creation time of the record
 - time interval covered by the accounting data
 - Resource/Service ID
 - UID or DN
 - Optional Global Request ID
 - Structure of the accounting data that follow and size

- Main information entered (mostly) via a name/value pair
 - The accounting connection library will format this into an XML files



Meter – Collector Connection

- Each Meter instance needs the Collector hostname or IP address and a Token.
- A token is a string that the Meter instance and the Collector perform mutual authentication based on a secret-key protocol
- The tokens are contained in a file managed by the Collector
- A token uniquely identify a Meter instance such that the Collector can distinguish when a Meter instance has **never started** or when there is **no usage information** to report.
- System administrator can add Meter Instances to the Collector configuration file. Each token must be a unique string for the site.



Avoiding Data-loss

- One of our goals is to reduce as much as possible the amount of data loss due to transient problems.
- Another goal is to make the Accounting Interface as simple to use as possible
- In consequence we will provide a library which will establish and manage the network connections
- In case where the connection between two elements is broken the information that would have been sent is record in a local data file.
- As soon as the connection is restored the local data file is send to the next element.
 - Example of elements pairs:
 - Meter-Collector
 - Collector – Local Accounting store



Conclusion

- More Information
 - [Project Definition](#), Requirements and Design Documents
 - [OSG Accounting TWiki](#) and [mailing list](#) (osg-accounting@openscience.org)
- Any Questions, Comments, etc?