Update on OSG's Network Datastore Service

Shawn McKee

WLCG Network and Transfer Metrics Working Group Report
8th July 2015

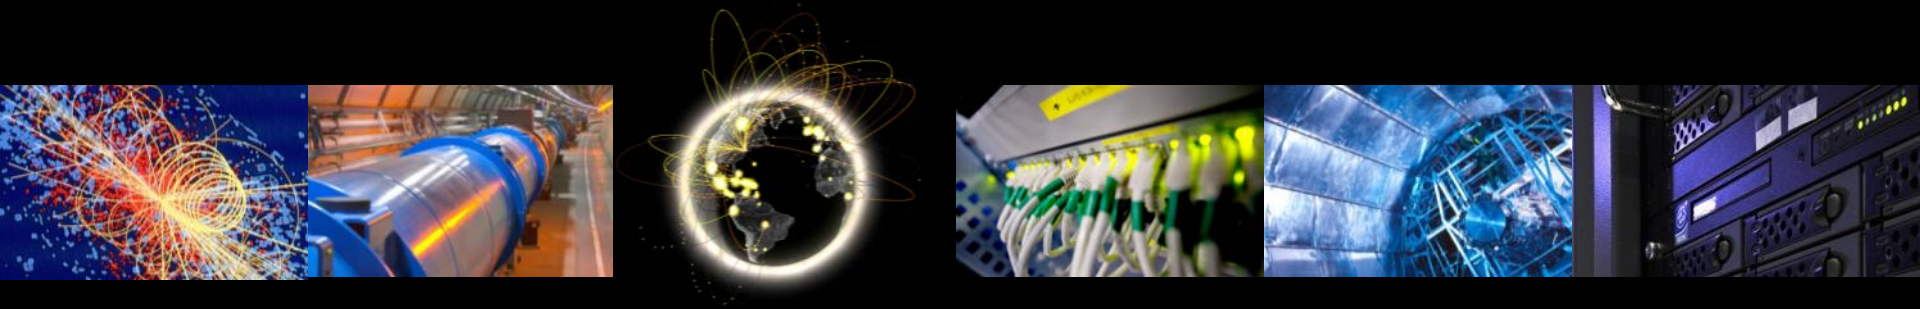

WLCG Worldwide LHC Computing Grid

Open Science Grid

Overview of OSG Network Service

- Open Science Grid (OSG) has deployed a network service for WLCG (and LHCONE). It consists of:
  - A datastore based upon Esmond (new MA in perfSONAR v3.4)
  - A GUI using MaDDash
  - A service monitoring component built on OMD
  - A "mesh-creation-configuration" utility built on registered information in OIM and GOCDB
- Status:
  - Datastore is "almost" production (targeting end July)
  - MaDDash and OMD working well.
  - Mesh-config: plans are to package to allow others to use.

OSG Datastore Motivation

- OSG intends to provide network metric data for its members and WLCG
  - The data is gathered from WLCG/OSG perfSONAR instances
  - Stored indefinitely on OSG hardware
  - Made available via API
- The primary use-cases
  - Network problem identification and localization
  - Network-related decision support
  - Network baseline: set expectations and identify weak points for upgrading

OSG Network Datastore Diagram

- OSG is gathering relevant metrics from the complete set of OSG and WLCG perfSONAR instances
- This data will be available via an API, must be visualized and must be organized to provide the **"OSG Networking Service"**
- **Operating now**
- **Targeting a production service by end of July**

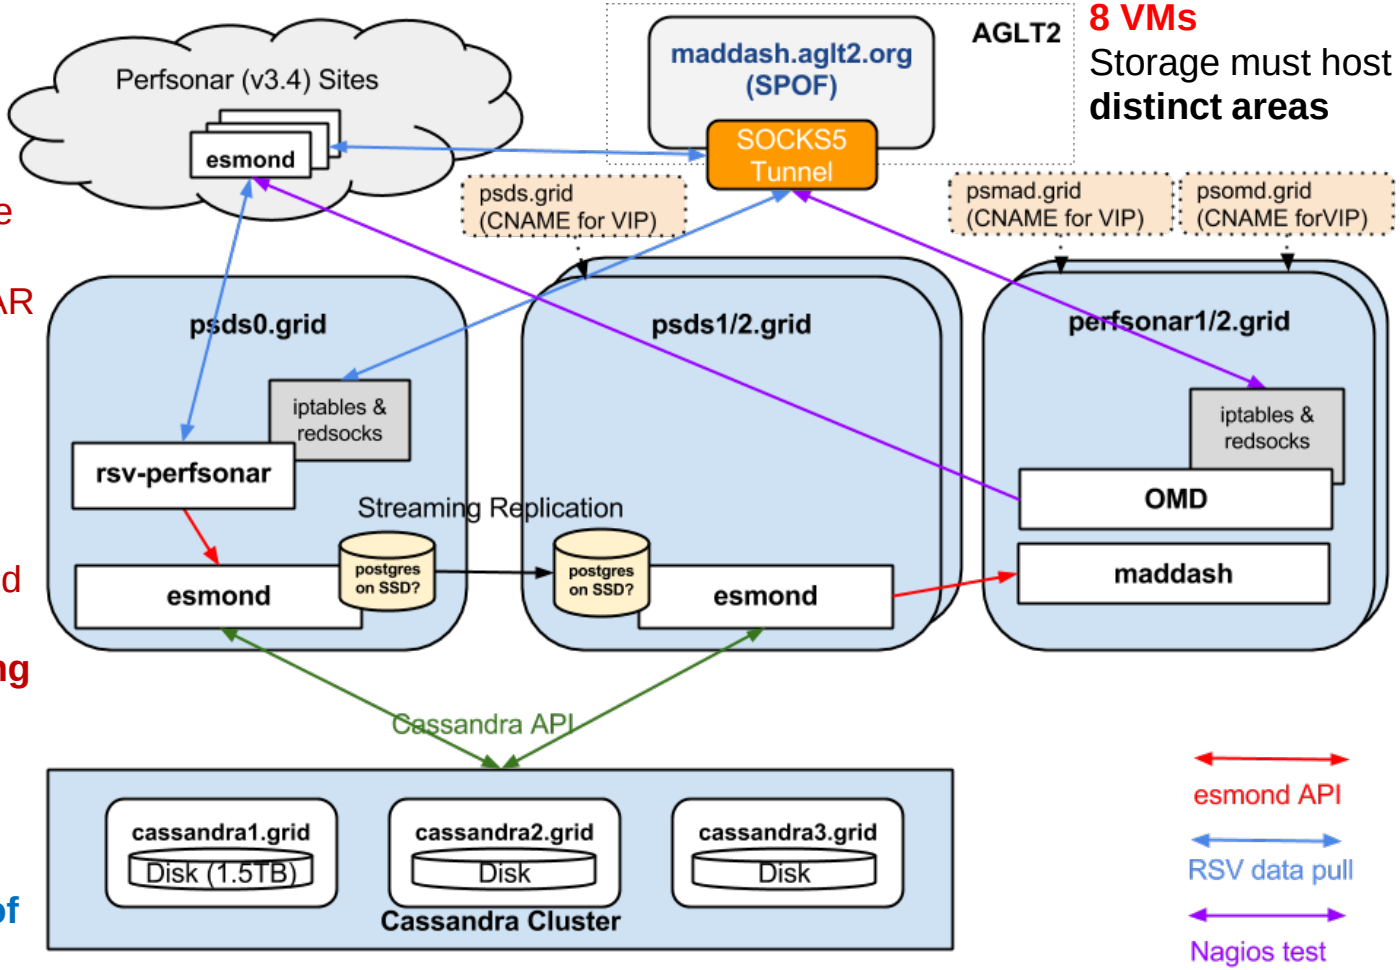

8 VMs
Storage must host **7 distinct areas**

Proposed Datastore Extension

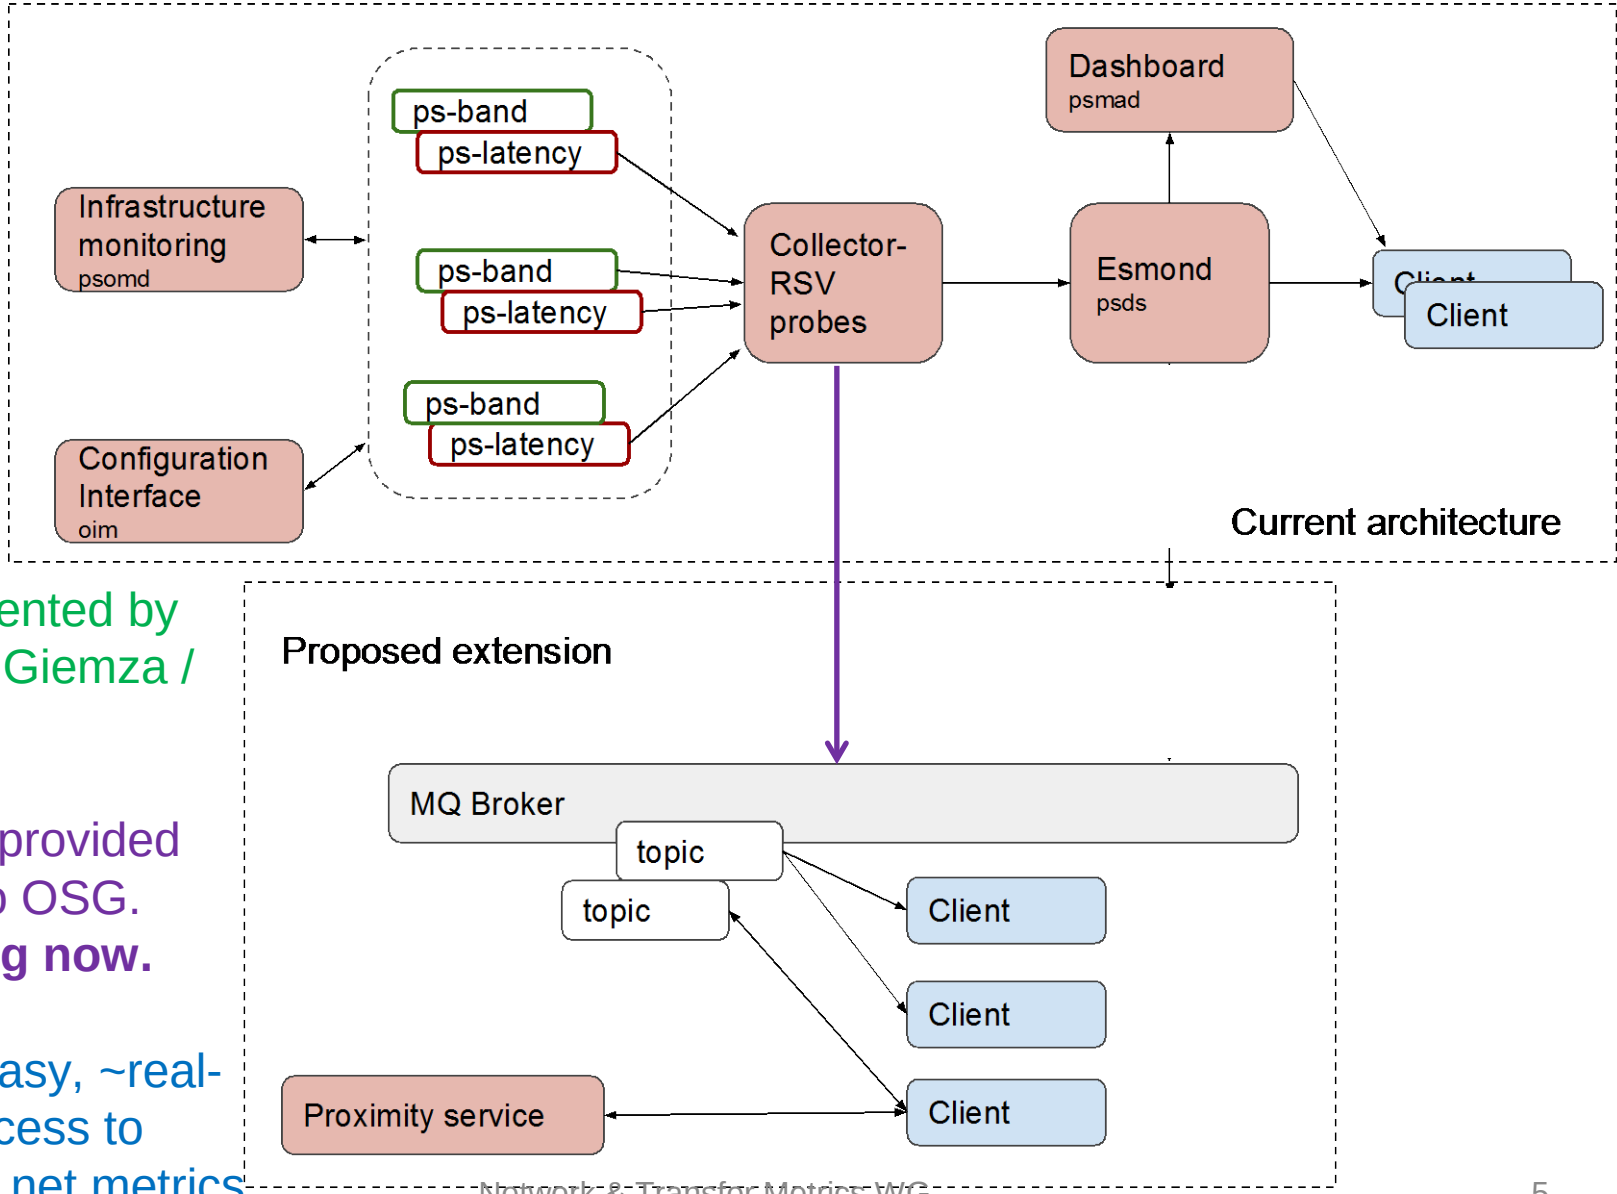

Implemented by Henryk Giemza / LHCb

Marian provided patch to OSG. **Working now.**

Goal: easy, ~real-time access to specific net metrics

Status

- OSG Network Datastore is on-target to be a production service in OSG by the end of the month
  - An OSG sub-group with members from OSG networking, operations and technology as well as Andy Lake/ESnet have met weekly to identify issues and keep the timeline
  - Data is being published to a MQ at CERN (50Hz, 1GB/hour). Need to agree to an SLA between OSG and CERN and get needed MQ cluster running for "production" at CERN
  - Still some issues to resolve in data-coverage. Tools in place to validate data and corresponding coverage.
  - OSG hardware configuration discussion after this meeting to optimize production setup for datastore.

Summary, Questions or Comments?

- We would welcome testers for the datastore and feedback on the API and MQ data delivery.
  - See references for some details on what exists now.

Questions?

References

- Network Documentation https://www.opensciencegrid.org/bin/view/Documentation/NetworkingInOSG
- Deployment documentation for OSG and WLCG hosted in OSG
  https://twiki.opensciencegrid.org/bin/view/Documentation/DeployperfSONAR
- New 3.4 MA guide http://software.es.net/esmond/perfsonar_client_rest.html
- Modular Dashboard and OMD Prototypes
  - http://maddash.aglt2.org/maddash-webui

    https://maddash.aglt2.org/WLCGperfSONAR/check_mk
- OSG Production instances for OMD, MaDDash and Datastore
  - http://psmad.grid.iu.edu/maddash-webui/
  - https://psomd.grid.iu.edu/WLCGperfSONAR/check_mk/
  - http://psds.grid.iu.edu/esmond/perfsonar/archive/?format=json
- Mesh-config in OSG https://oim.grid.iu.edu/oim/meshconfig
- Use-cases document for experiments and middleware https://docs.google.com/document/d/1ceiNlTUJCwSuOuvbEHZnZp0XkWkwdkPQTQic0VbH1mc/edit