

Monitoring and Information Services

Brian Bockelman

Information in the OSG

- It turns out that getting information in and out of grids - then turning this into usable data - is very hard work.
- We'll discuss available OSG accounting, monitoring, and information services.
- Lots of pictures, links, and demonstrations.

Outline

- Registration: OIM
- Accounting: Gratia
- Monitoring:
 - Monitoring Services: RSV
 - Monitoring jobs: ???
- Information Services: GIP

OIM

- OIM = OSG Information Management
- This is where sites, resources, and services are registered.
- You can also look up a list of all the sites on the OSG.
- Not designed for matching jobs like ReSS

OIM

- Make sure your site is registered here; otherwise, some OSG components may not work correctly.
- <http://oim.grid.iu.edu>
- This is how GOC will be contacting your site.

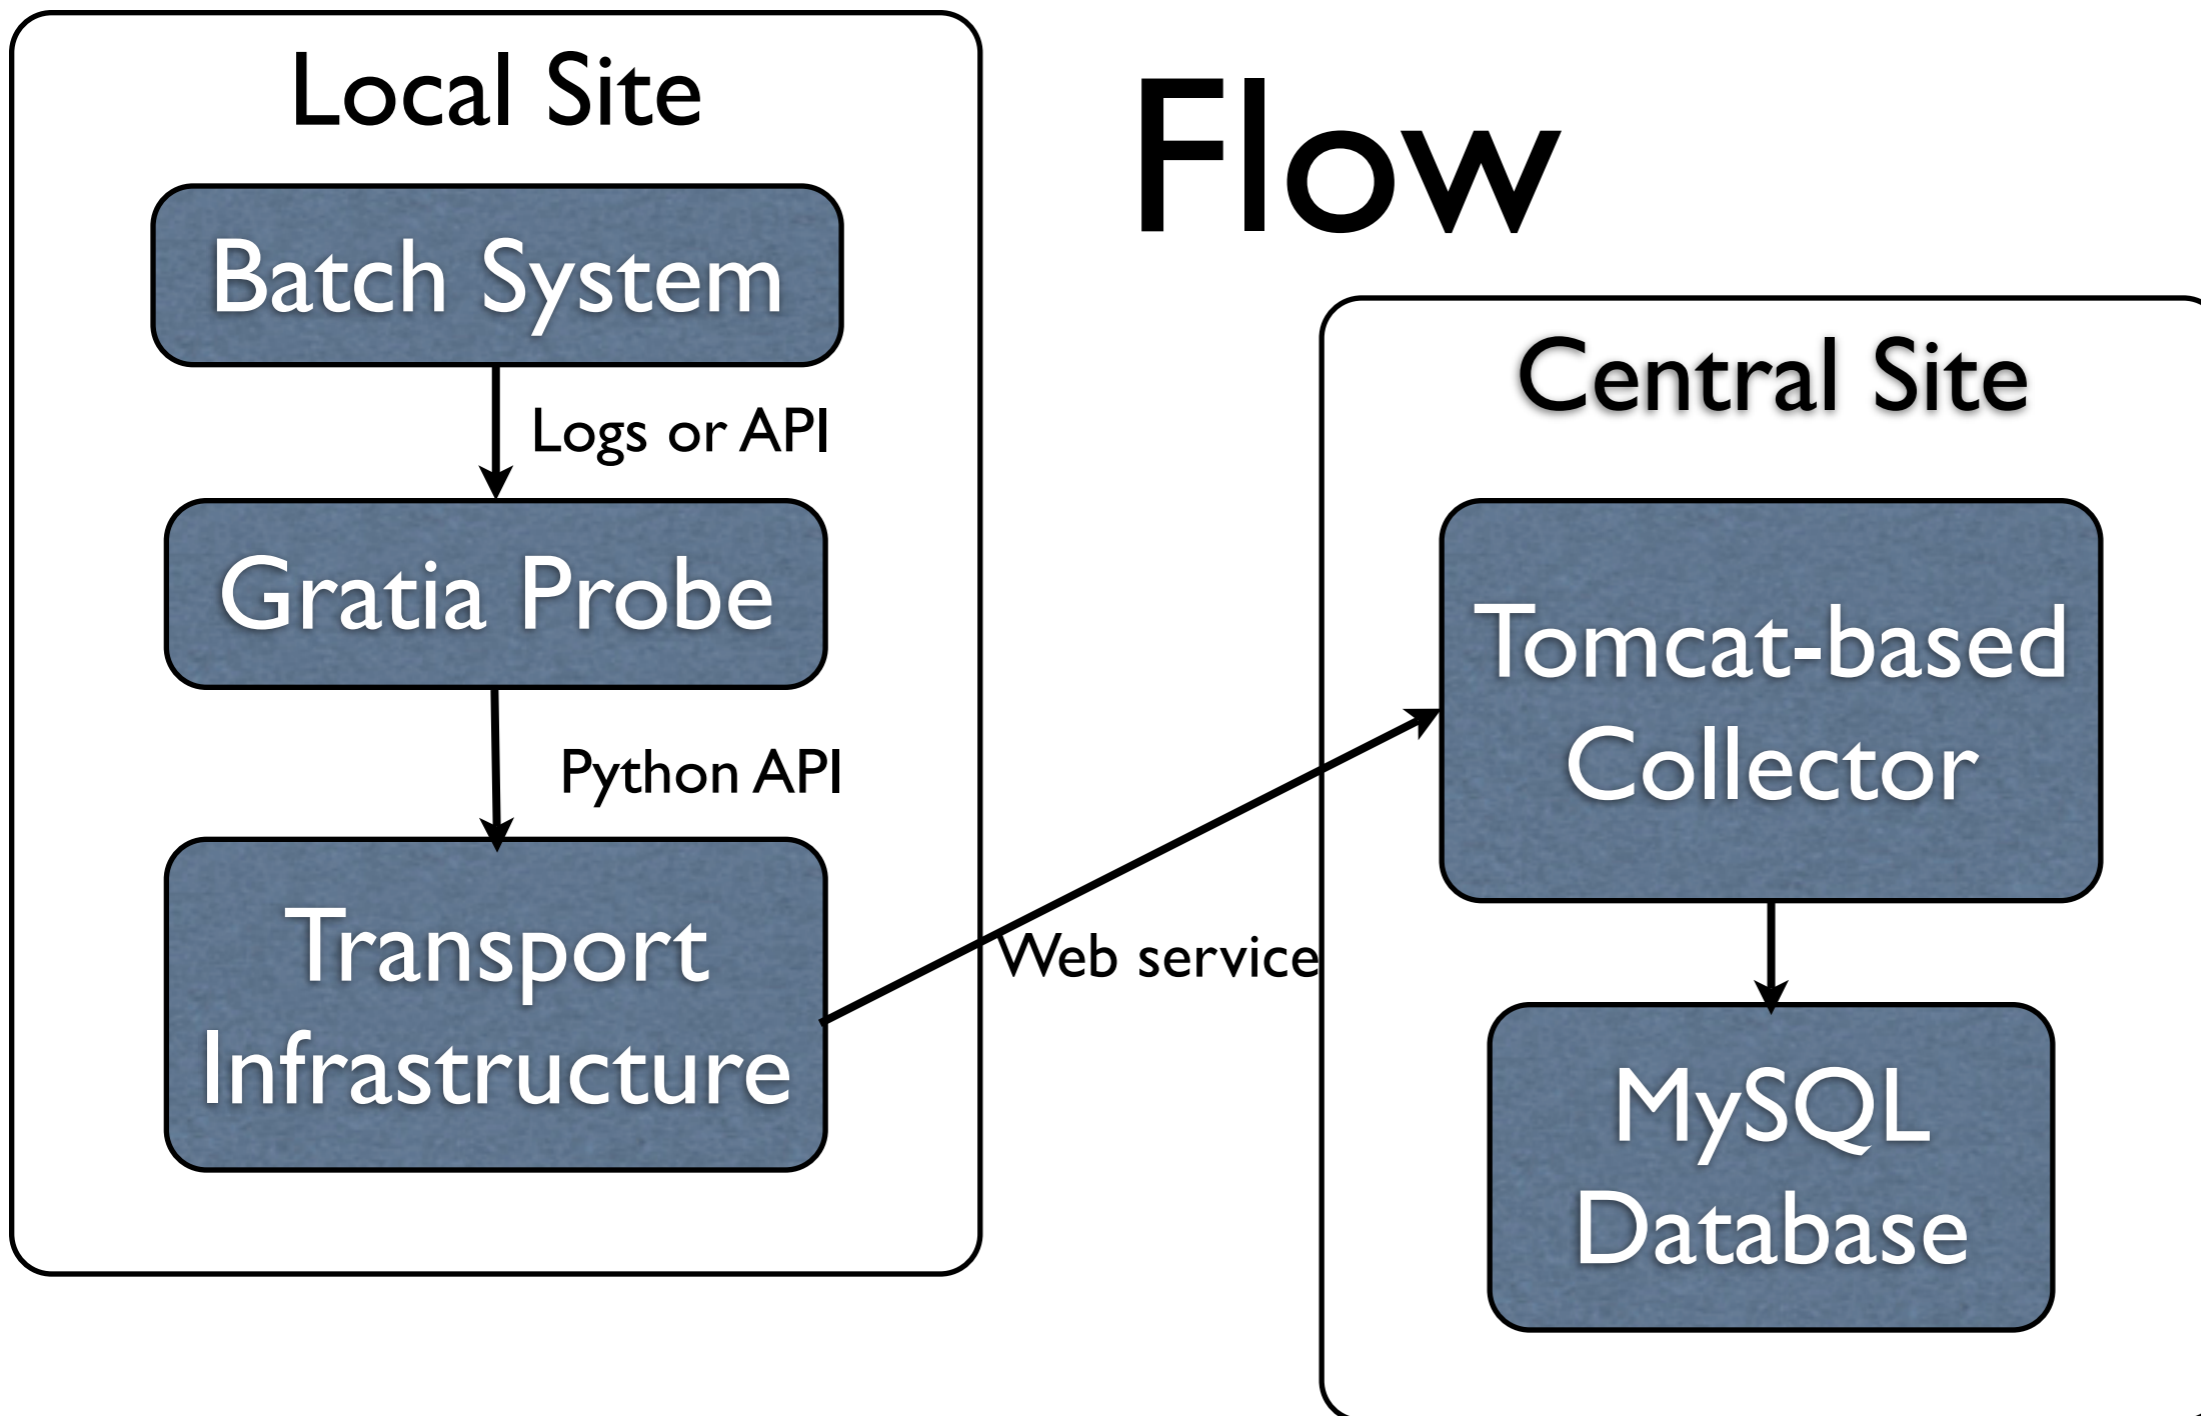
Accounting

- In the OSG, we aim to do accounting for
 - Batch systems
 - Storage or transfers.
- We don't do end-to-end (client) accounting.
- We do rely on batch systems for information.

Accounting System

- The accounting system used is provided by Gratia, a FNAL project.
- Gratia is composed of a central database - the “collector”, a transportation infrastructure, and a set of scripts which probe various systems.

Gratia Information Flow



Gratia Infrastructure

- One strong requirement is that no information is lost.
- Once the job information is sent through the Python API, you are guaranteed to have it make it to the collector
- But we can't guarantee the latency!

Gratia vs Email

- Like email, Gratia has a store-and-forward design.
- Messages are stored to local disk.
- The local infrastructure attempts to re-send the messages until the next server in the line confirms receipt.
- Like email, even if one server is down, messages still get transported.
- But there could be considerable latency!

Gratia Infrastructure

- The Gratia setup is also similar to some of the Logging/Bookkeeping of EGEE or the JMS standard.
- It is a bit more specialized: there's a pre-defined XML format and associated database schema for the collector components.
- Has been reused by other OSG projects!

Gratia Infrastructure

- The collectors can also form a network topology.
- For example, you could have a central collector for a large facility/region which forwards records to the OSG collector.
- Would provide web interfaces for just your facility

Gratia Probes

- The Gratia probes are customized for each batch system.
- Some (Condor) have APIs which allow the probes to query the history directly.
- Others (PBS, LSF, SGE) require the probes to read logfiles: a bit more error prone.

Gratia Probes

- Gratia Probes are fairly easy to configure:
 - Make sure you install them with pacman post-CE install
 - They get configured by the configure-osg.py script.
 - You may need to tell them where to find batch system log files for PBS.

Gratia Probes

- We also have transfer and storage probes
 - For dCache, this reports every data-level transfer made to the central DB.
 - For dCache, also reports pool and space reservation usage.
- We are also working on a Globus GridFTP probe

Gratia Probes on OSG CE

- Short demonstration: where Gratia Probes are on the OSG CE, how they run, where to find the final probe configuration.

Gratia Web Interfaces

- Gratia's database provides a great opportunity to build web interfaces
- Indeed, there is the "official" web interface and other web interfaces which build off that

Gratia Web Interfaces

- <http://t2.unl.edu/gratia/>
- <http://t2.unl.edu/gratia/xml/>
- <http://gratia-osg.fnal.gov:8880/gratia-reporting/>

Gratia Summary

- This is an important component!
- You can get daily reports about this, if you'd like
- I encourage you to look at this constantly
- If not that, I encourage you to check at least once that you set it up right!

Monitoring

- I want to cover two monitoring concepts:
 - Overall batch system monitoring
 - Job-level monitoring

First, the bad news

- In the OSG, we have no pre-built job-level monitoring :(
- However, tune in this afternoon for examples of what CMS and ATLAS can do!

Then, the Good News

- RSV (Resource Validation Service) is a great site monitoring tool.
- We talked about this at length yesterday, but I want to show some pieces of it on the file system.

Information Services

- If you have your site registered, the GIP (Generic Information Provider) will upload batch system and (possibly) storage information to the central LDAP server
- Sample output (lots!): http://is.grid.iu.edu/data/cemon_processed_osg/red.unl.edu.processed?which=Nebraska

GLUE Schema

- A common language to describe grid objects and their relationships
- Provides descriptions of the Site, CE, SE, batch queues, SRM endpoints, Gridftp servers, etc.

Glue Schema Diagram

- Don't try to memorize this

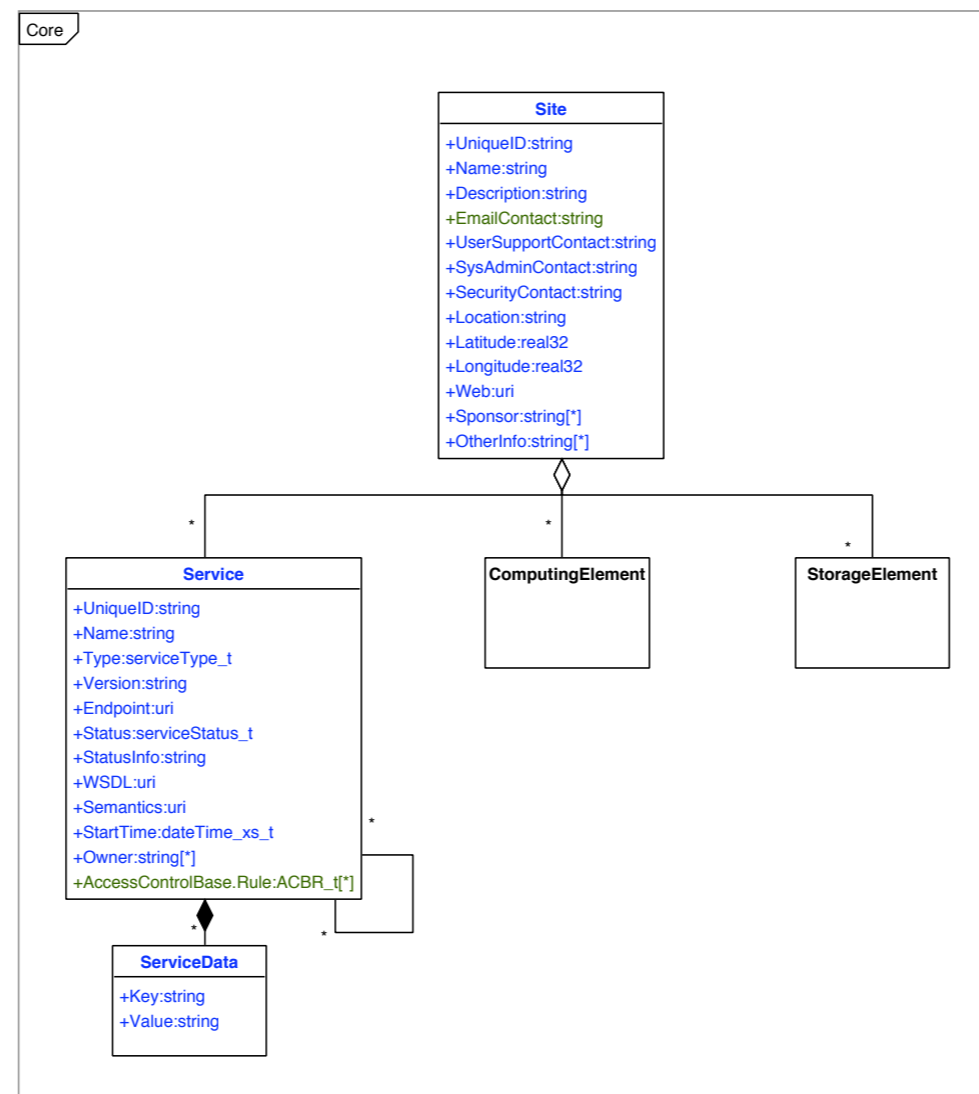


Figure 1: Core (black=unmodified, red=deprecated, blue=new in 1.2, green=new in 1.3)

Glue Schema Diagram

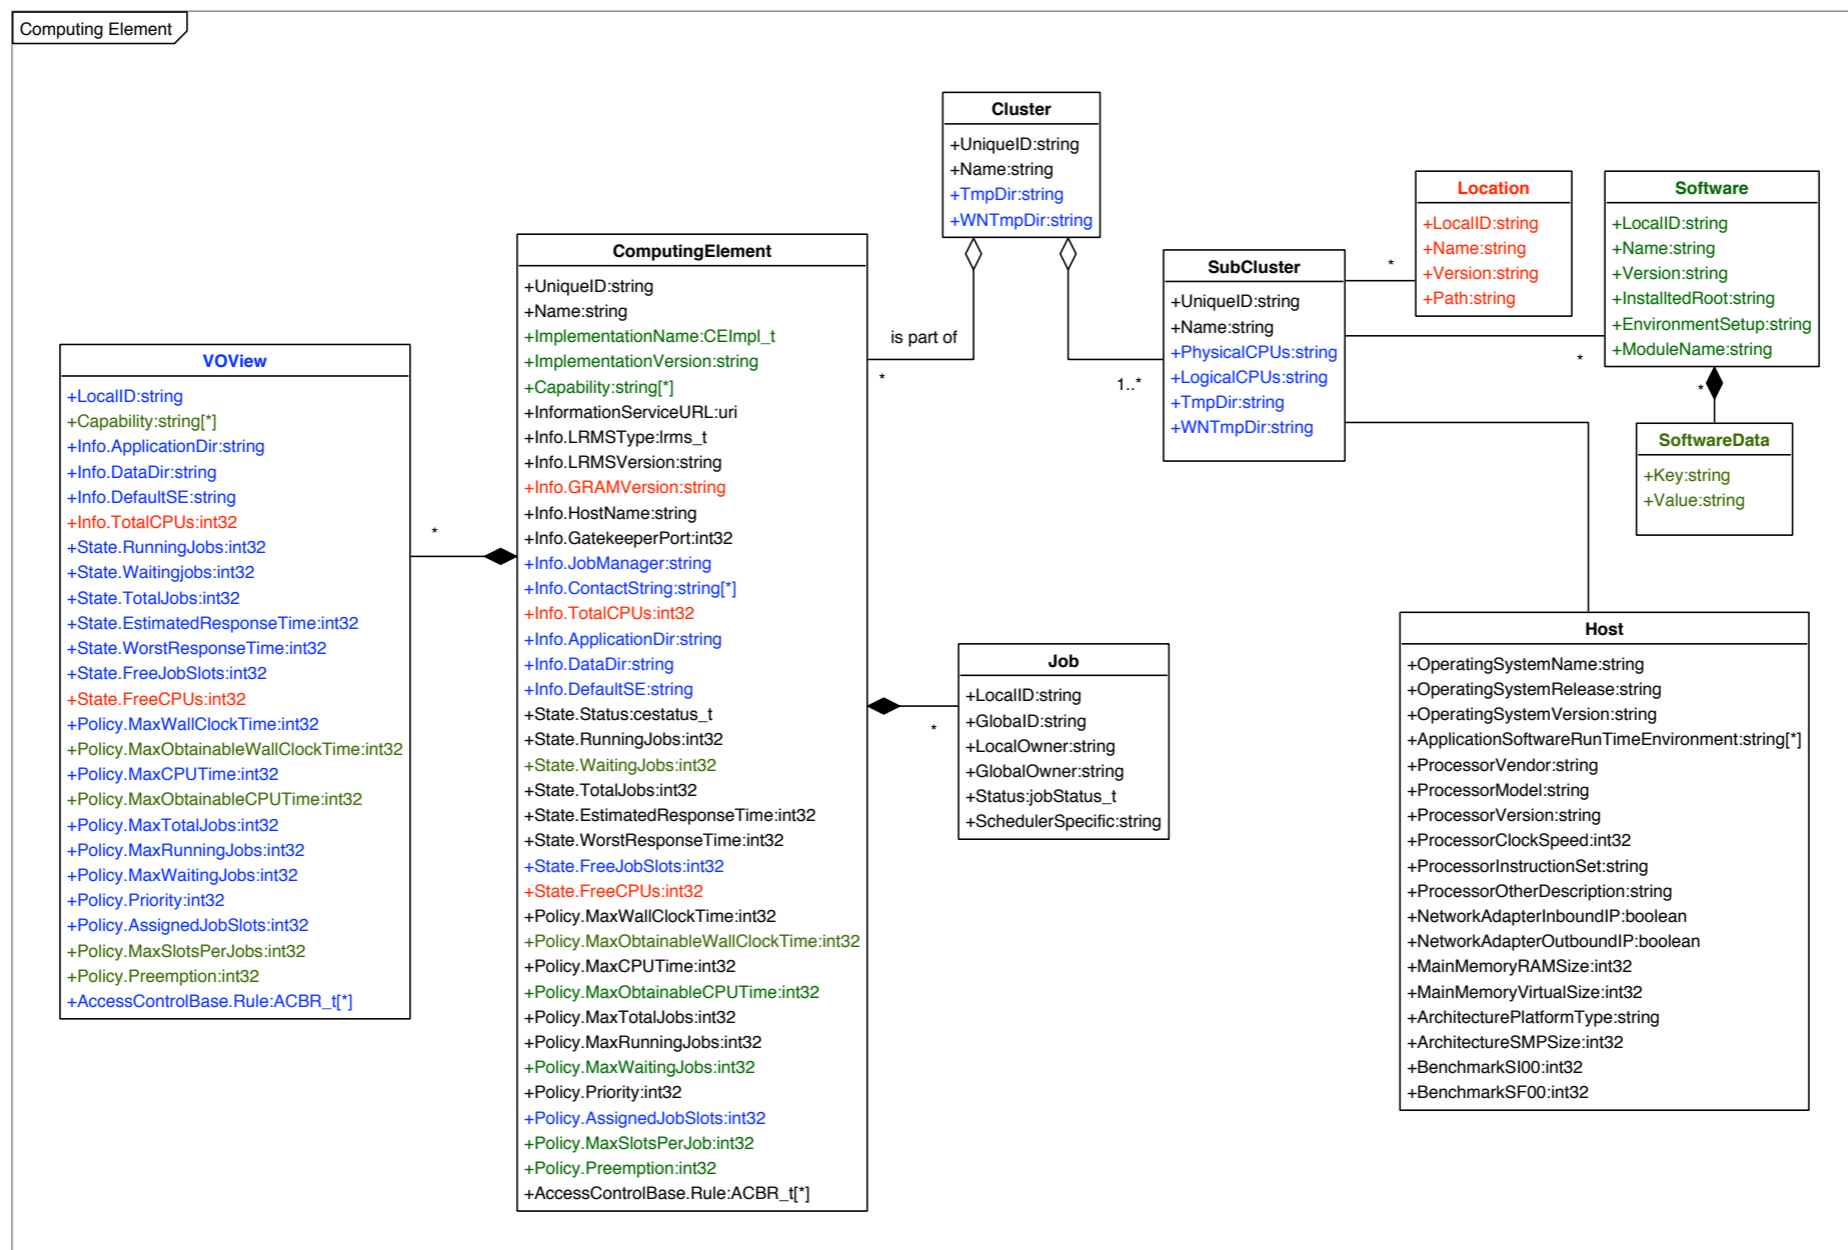
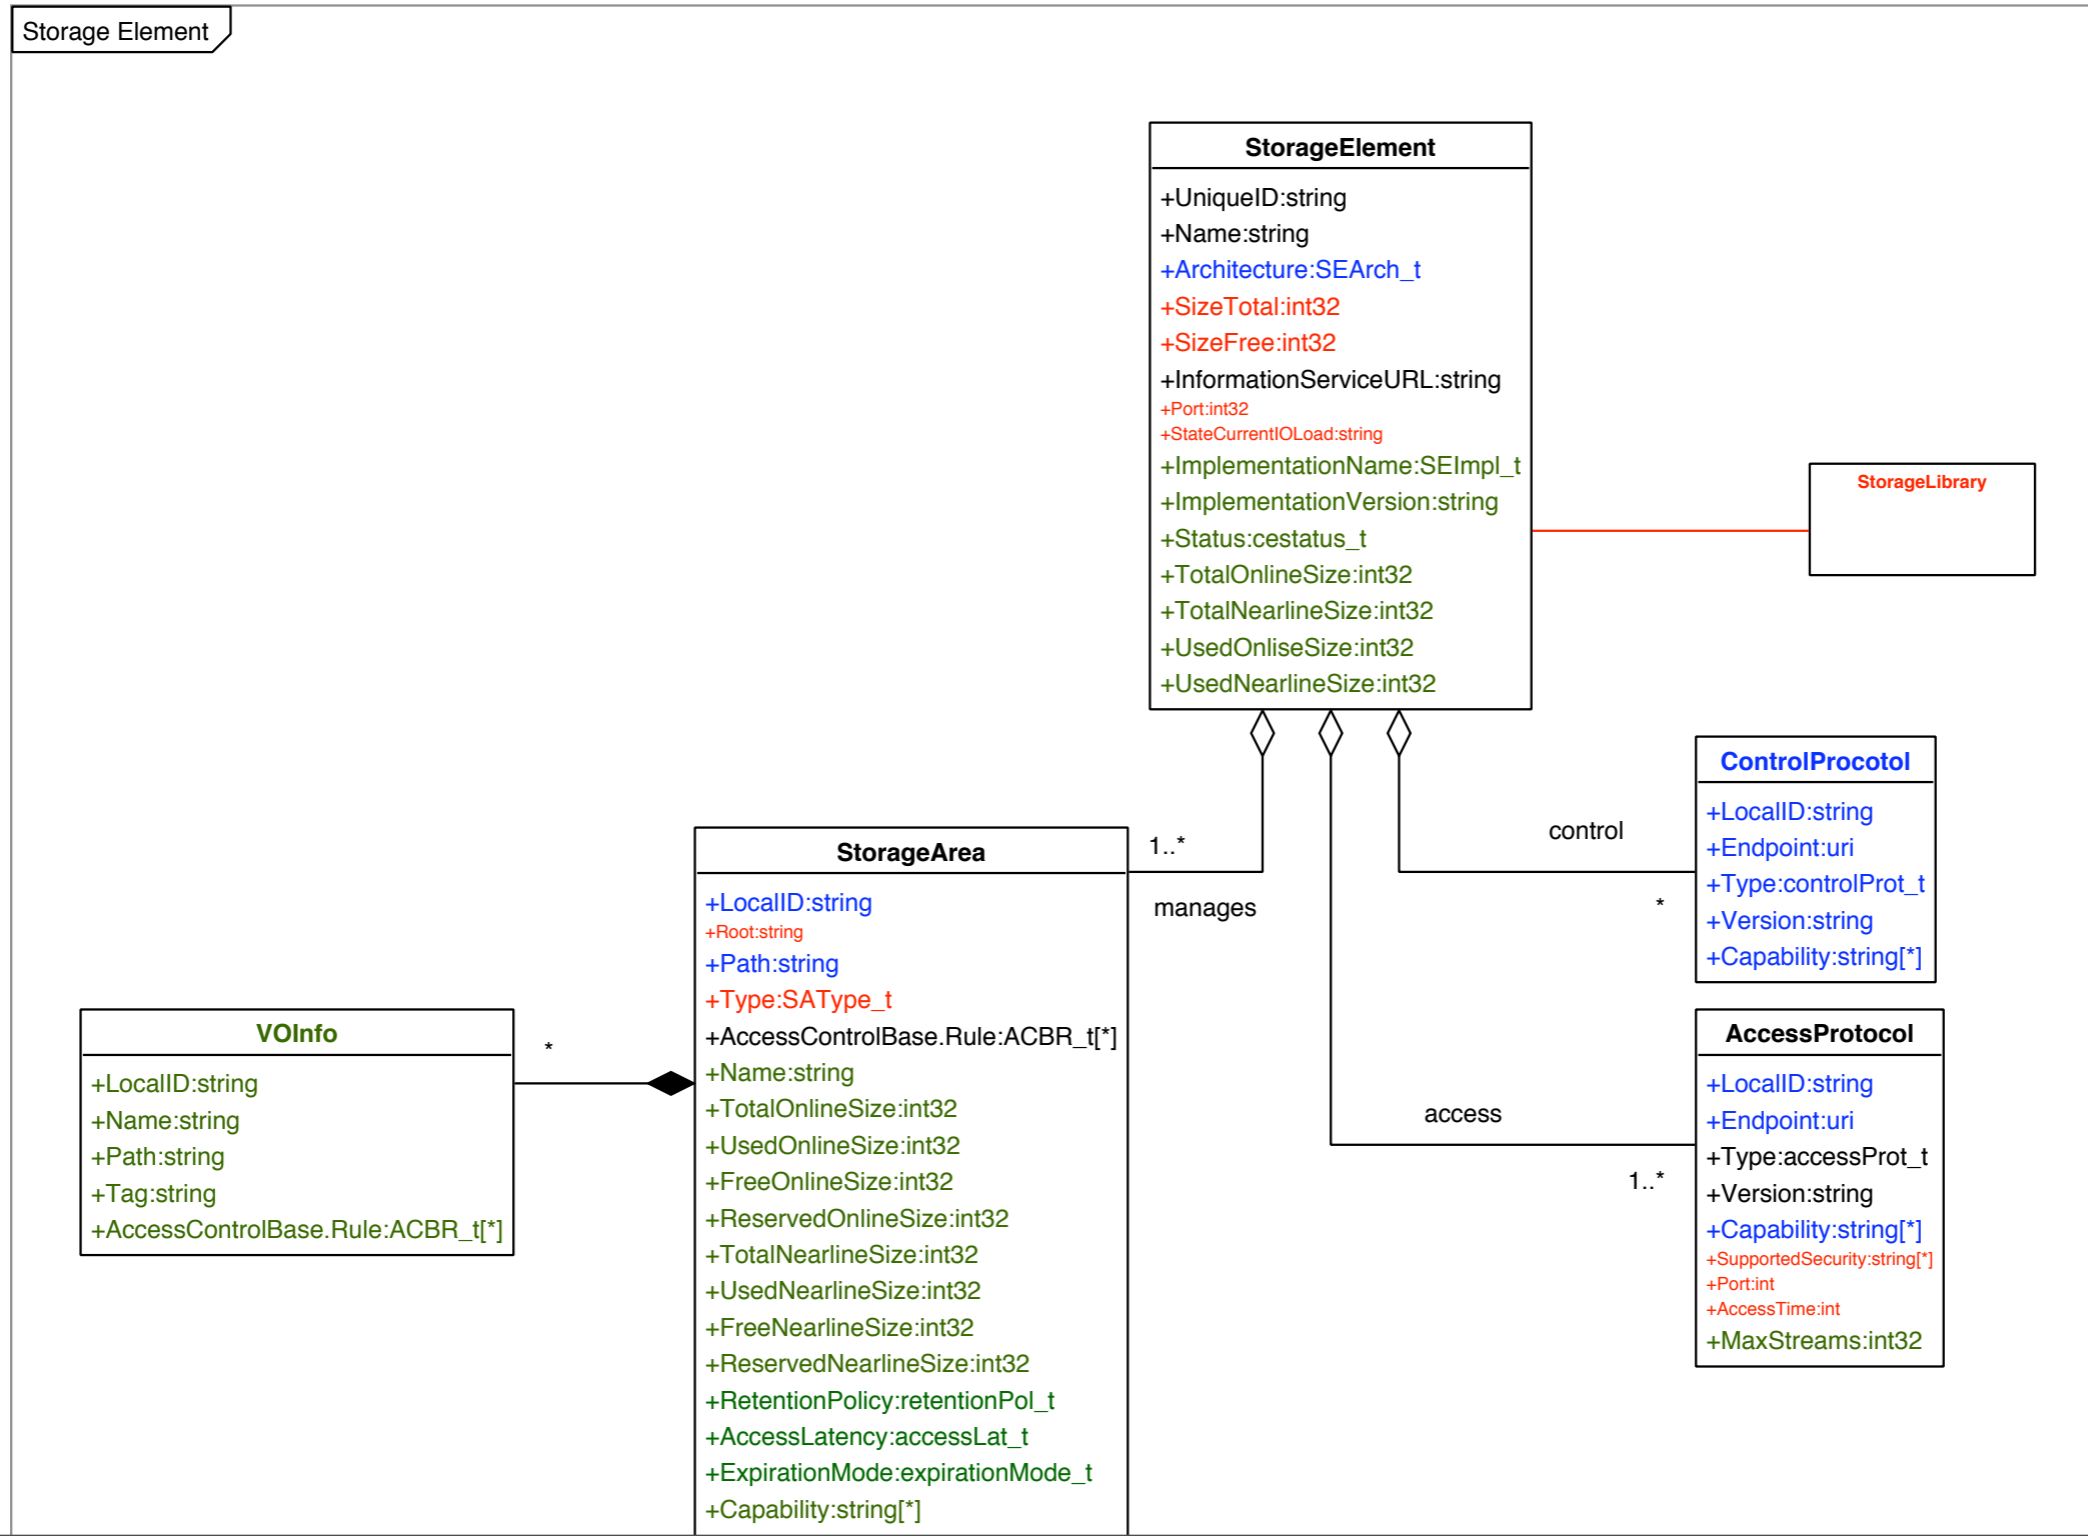


Figure 2: Computing Resources (black=unmodified, red=deprecated, blue=new in 1.2, green=new in 1.3)

Glue Schema Diagram



The GIP and GLUE

- The GIP's purpose is to describe the site information using the GLUE schema.
- To accomplish this, it uses a combination of the static information provided in config.ini
- And dynamic information which it queries from your cluster

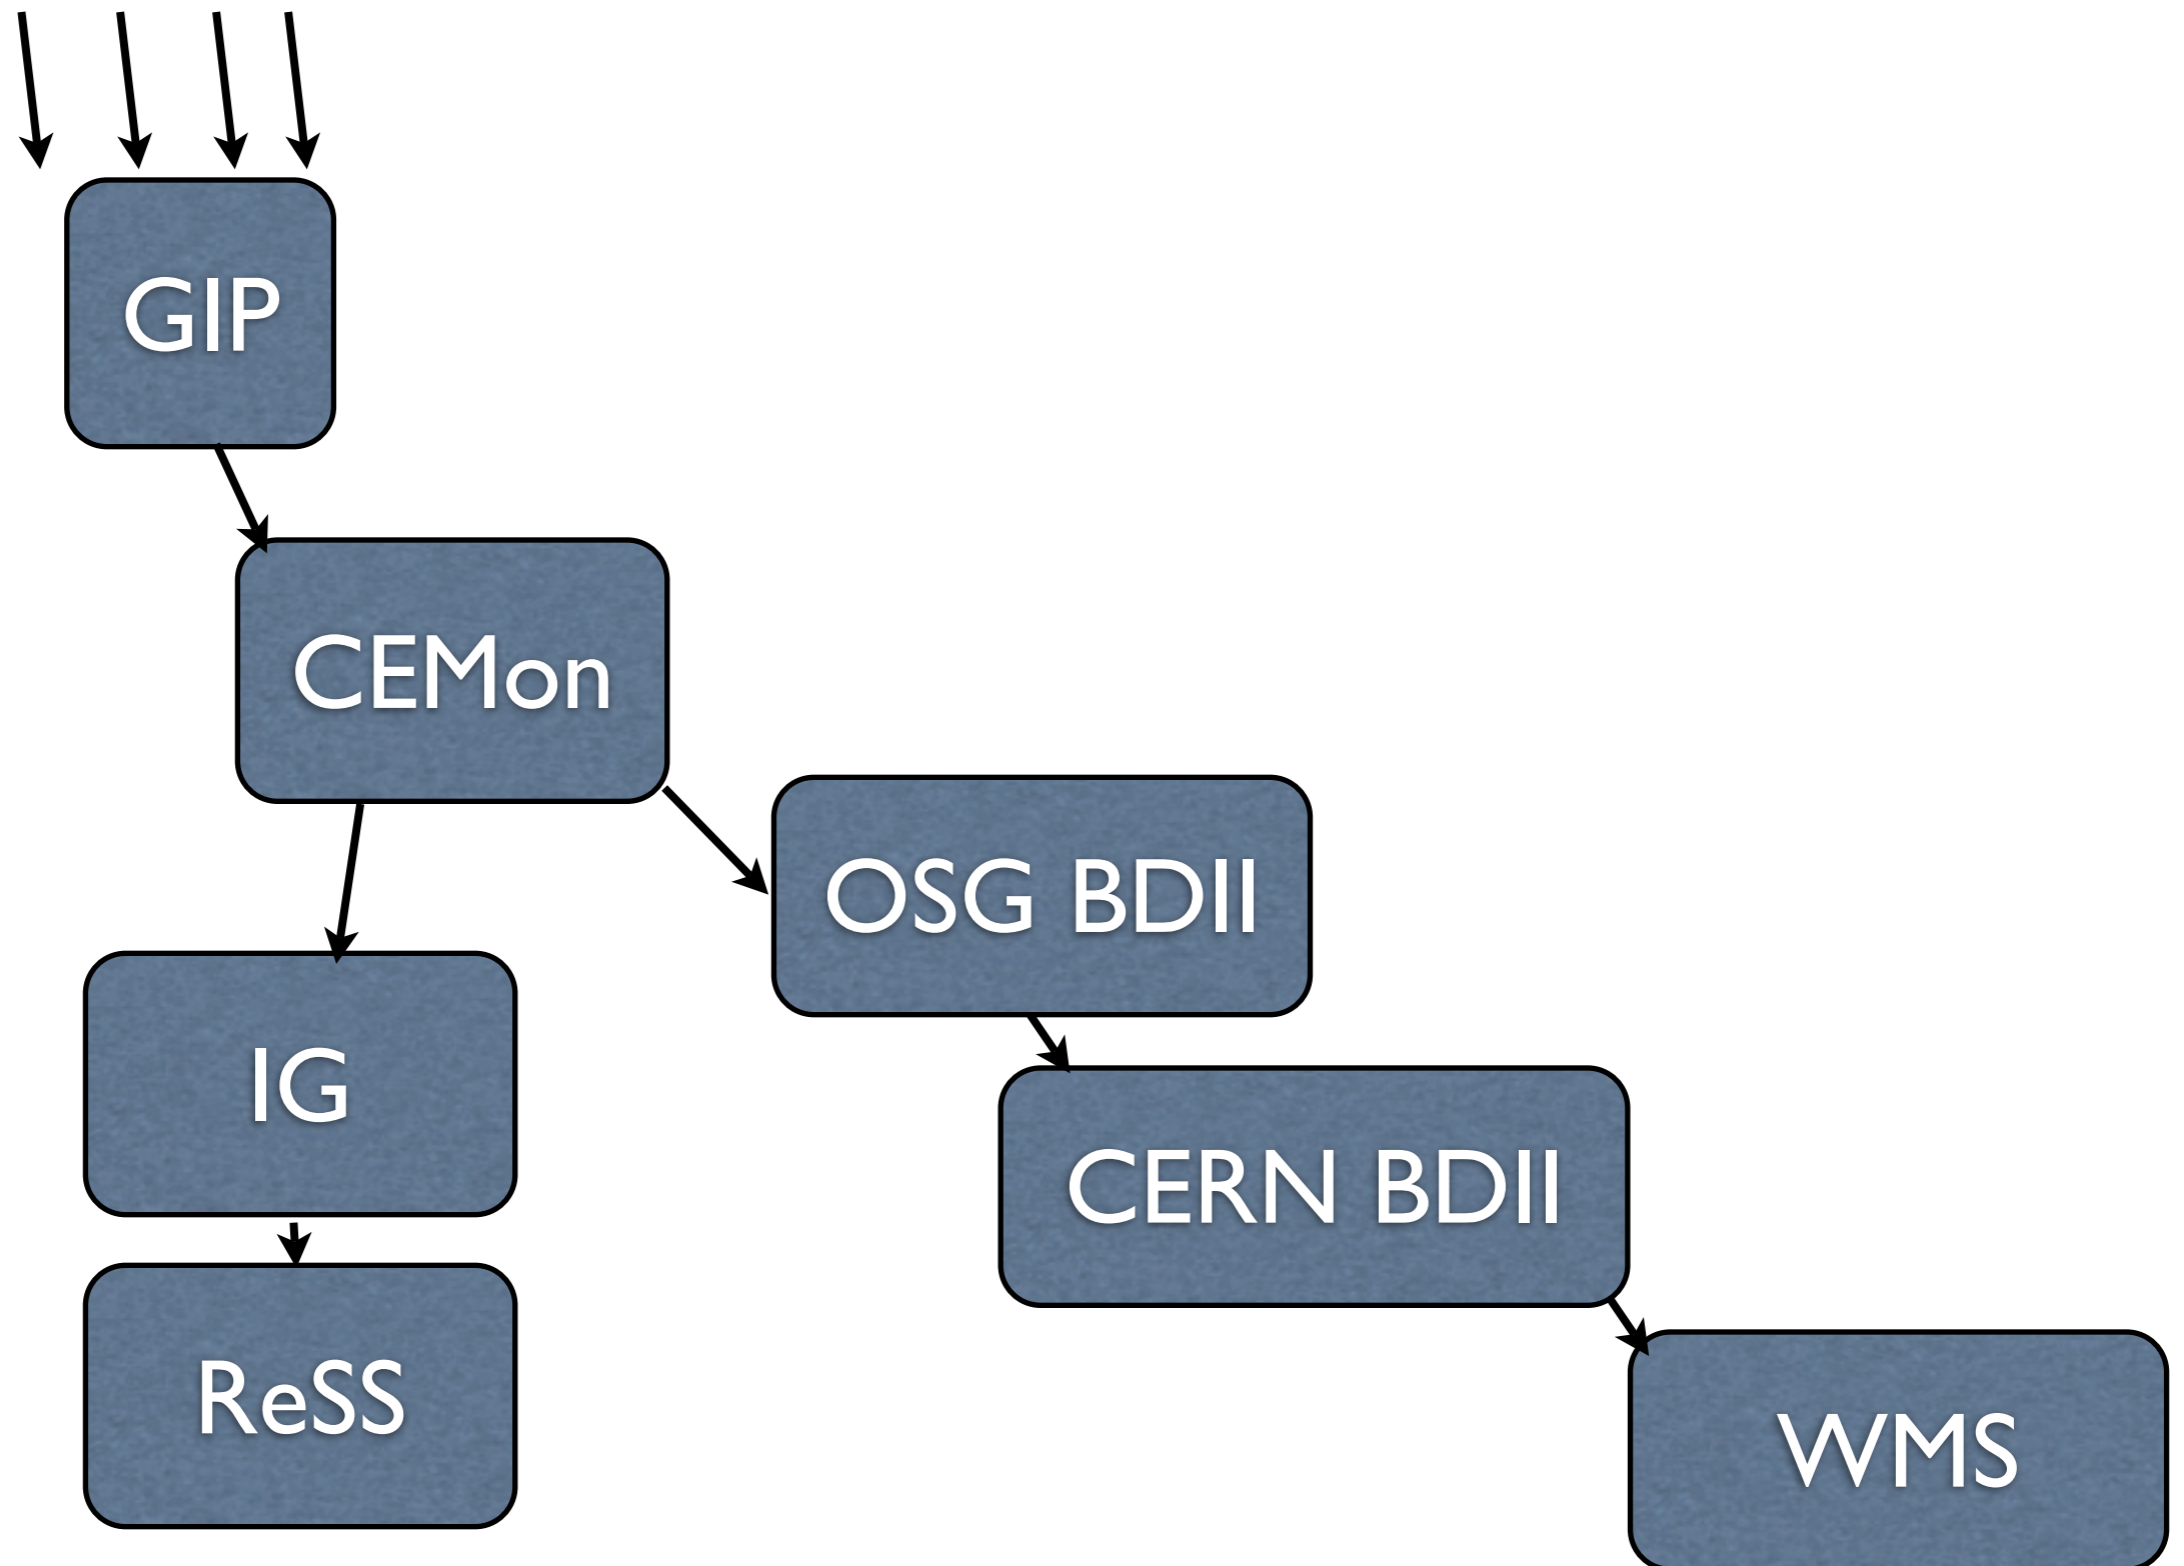
The GIP

- So, for example, the GIP will read in config.ini that your site admin's name is bob@example.com
- But it will read from Condor that you have 10 free batch slots

GLUE information

- The GLP's output is used in the OSG and EGEE:
- In EGEE, this informs the decisions of the WMS
- In OSG, this informs the decisions of ReSS.

Information Path



GIP Parts Demonstration

- Time Permitting, a demonstration of the filesystem layout of the current GIP

Putting Everything Together

- <http://t2.unl.edu/status>