Information System usage by the EGI Operational Tools

D.Cesini (INFN/IGI)

On behalf of the EGI-JRA1 activity

Towards an Integrated Information System Workshop

Amsterdam 01/12/2011

Outline

- The EGI-JRA1 Activity
  - Developed tools and other tasks
- The GOCDB
  - Tool presentation
  - Tool usage statistics
- The ops current usage of the Information System and GOCDB
- Conclusion

Tools developed by EGI-JRA1

- Operation Portal (CNRS)
- EGI Helpdesk (KIT)
- GOCDB (RAL)
- Accounting Repository (RAL)
- Accounting Portal (CESGA)
- SAM/MyEGI (CERN/SRCE)
- Metrics Portal (CESGA)

Other activity tasks

- Message Broker Configuration to support tools and operation (AUTH)

- Accounting for different resource types (LUH/INFN/RAL)
  - Billing
  - Accounting of application usage
  - Accounting of data usage
  - Accounting of capacity and cloud computing usage

JRA1 Tools & Dependencies

EGI-JRA1 Tool

- **Depend ON InfoSYS**
- **Depend ON GOCDB**

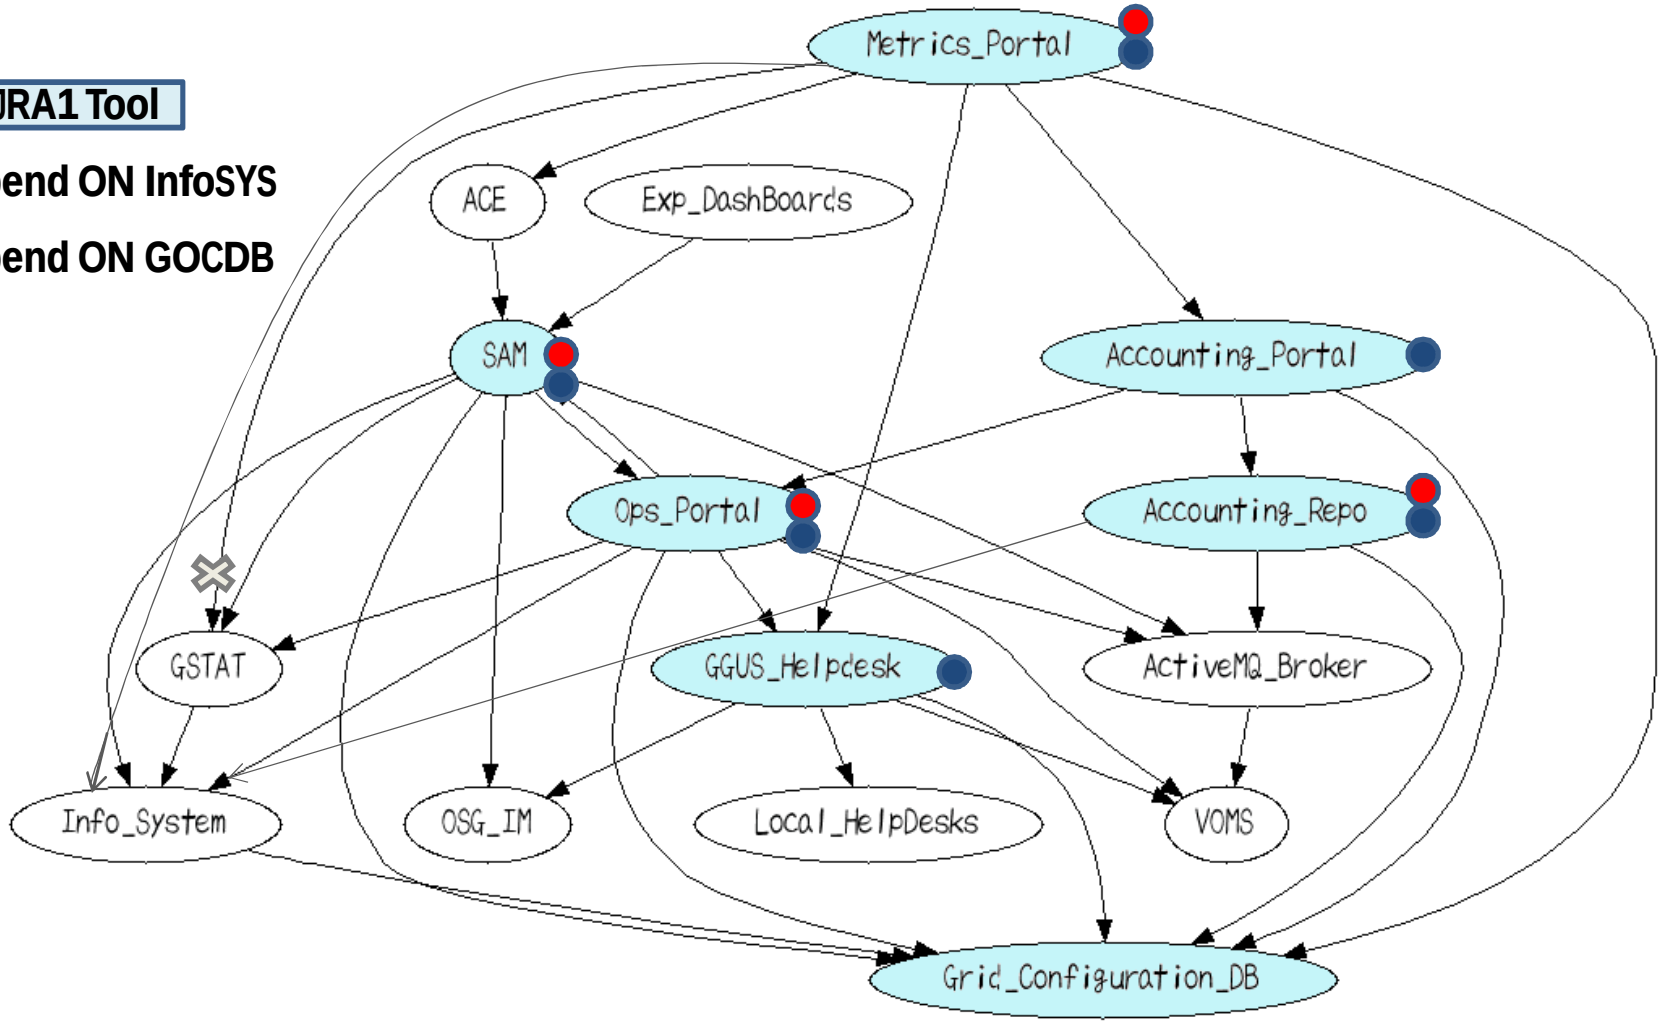

Source MS704 + fixes

GOCDB

GOCDB

- EGI relies on a central configuration database to record “static” information
- Contributed by the resource providers
  - service instances that they are running
  - the individual contact, role and status information for those responsible for particular services

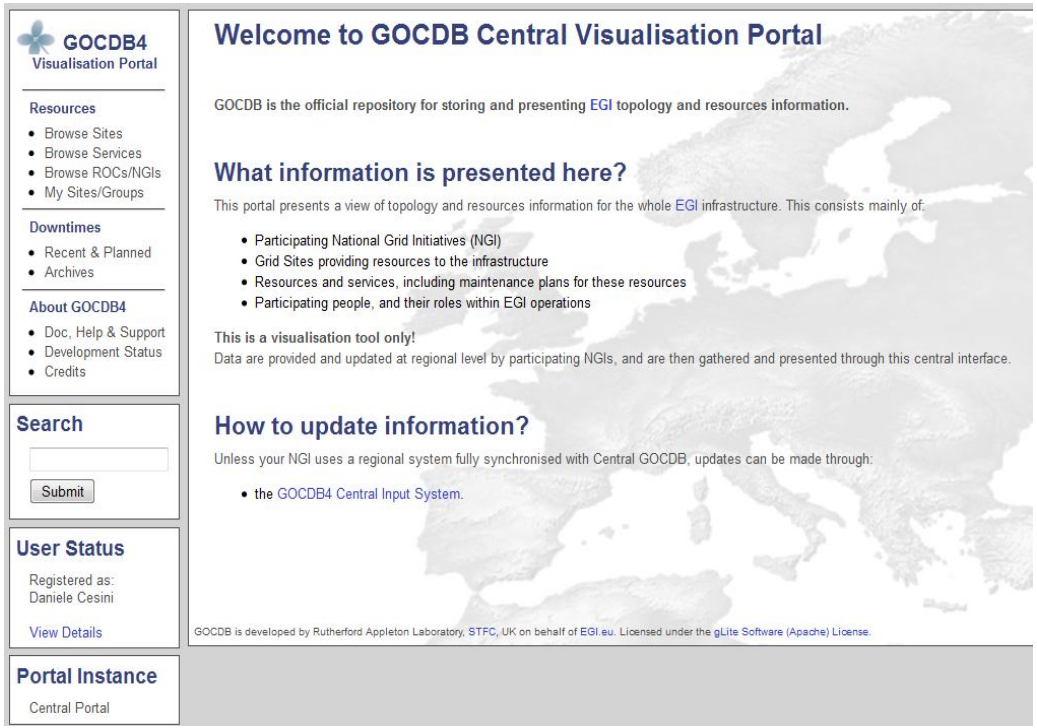

GOCDB

- GOCDB is currently hosted centrally and can store both EGI data and Locally scoped data not intended to be visible to EGI
  - Regional Standalone GOCDB available/maintained for NGI willing to install it
  - Regional Publishing GOCDB not available - is a long term development
- GOCDB comes along in two flavors, a *read-only* "Central Visualisation Portal" and a *writeable* "Input System"

https://wiki.egi.eu/wiki/GOCDB/Release4/Architecture

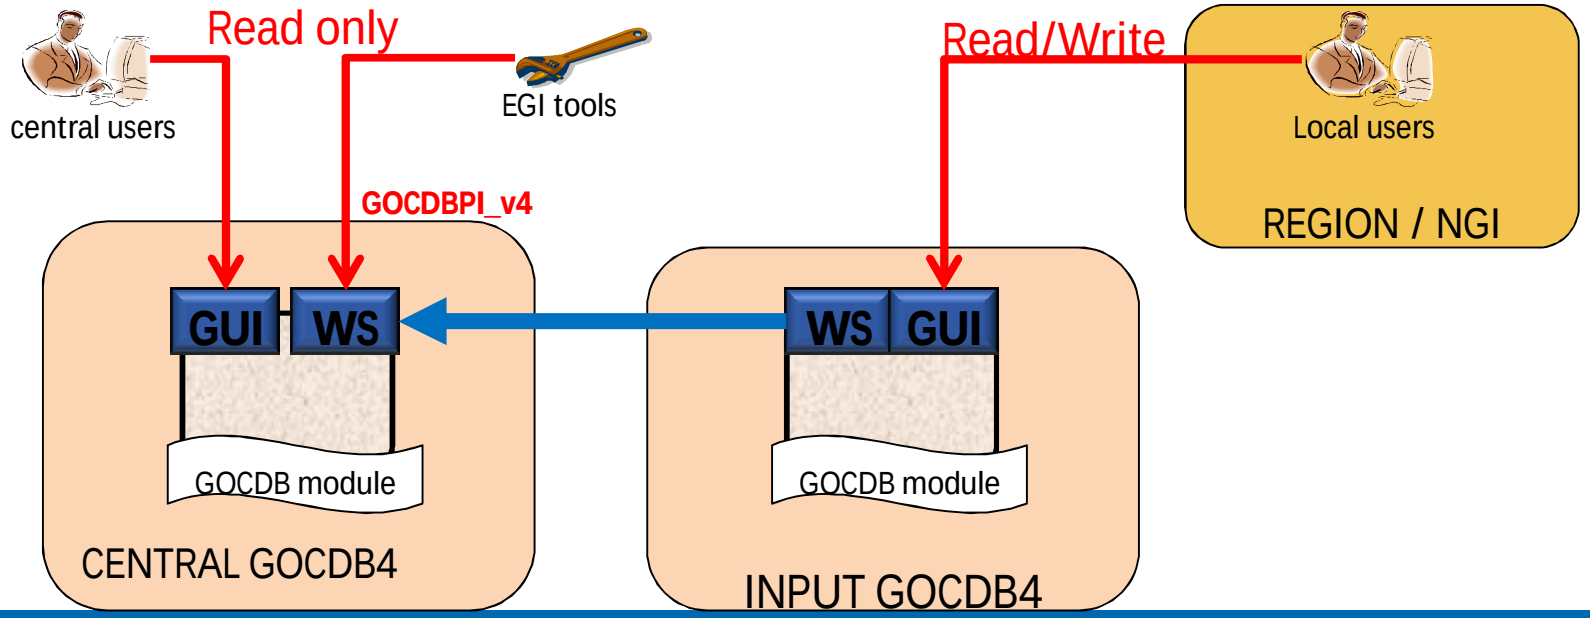

GOCDB Architecture

- A three tiered web application
  - a web GUI for manually recording Grid topology information
  - a REST style API for querying that data in XML
  - Data stored in an Oracle database

- GOCDB does not provide a writable programmatic interface for create/update/delete
  - Rather, the main CRUD interface is the web portal which provides control to users in defining their topology

GOCDB Data Model

- GOCDB records topology information and also its associated state
  - i.e. recording what sites/services/downtimes are meant to be functioning at a particular snapshot in time
- The information is largely hierarchical;
  - a parent NGI (group object) aggregates many child Site objects which in turn aggregate child Service Endpoints.
- Downtime objects are linked to individual service endpoints (not sites)
  - If all of the site's services are in downtime, then the whole site is effectively in downtime.
  - Each of the different entity types have relevant attributes, including people memberships, contact information, service types, downtime start/end dates etc.

GOCDB Permission & Role Model

- The permissions model is hierarchical
  - users with higher level permissions on parent objects (e.g. NGIs) can grant permissions to other users over sibling and child objects (e.g. Sites)
  - following the creation of a user with a top level ‘NGI manager role,’ NGIs can manage and propagate their own users, Sites, and Service Endpoints without involving the GOCDB administrators
  - the roles/permissions model is currently being re-developed to facilitate finer grained permissions
    - In the newly evolving AAA model, the hierarchical permission model is less obvious
      - https://wiki.egi.eu/w/images/b/b7/FinerGrainedGOCDB_rolesVeraProposal2.xls

GOCDB Stateful Data and State Transition Rules

- GOCDB has predefined state transitions rules
  - sites have an associated ‘Production’ and ‘Certification’ status.
  - A Site would normally moves from ‘Candidate -> Uncertified -> Certified’ during the certification process
  - other transitions are strictly forbidden.
    - Similar state transition rules also exist when declaring and editing downtimes and when granting/revoking roles

- https://wiki.egi.eu/wiki/GOCDB/Input_System_User_Documentation

GOCDB Historical Data and Auditing

- Data are never actually deleted from the DB
  - Instead, objects have an associated timestamp which indicates whether that object is currently ‘live’ or is deleted
  - By providing timestamp parameters to the PI methods, the PI can be used to query for historical/audit data
    - e.g. list a site’s certification status history
    - who, why and when those changes were applied

GOCDB Usage Statistics (1/3)

PI Statistics on Visualization Portal

- In October 2011 there were nearly 6 million separate queries for the PI:
  - which averages ~2.3 queries per second
  - The query rate has a background steady state averaging 0.4 MB/sec
  - cyclic daily peaks
    - reflecting requests from automated scripts
- Given this high usage, caching the results of popular and expensive queries is essential
  - Selected queries are executed every 30mins or so and the XML results are cached and served by Apache.

GOCDB Usage Statistics (2/3)

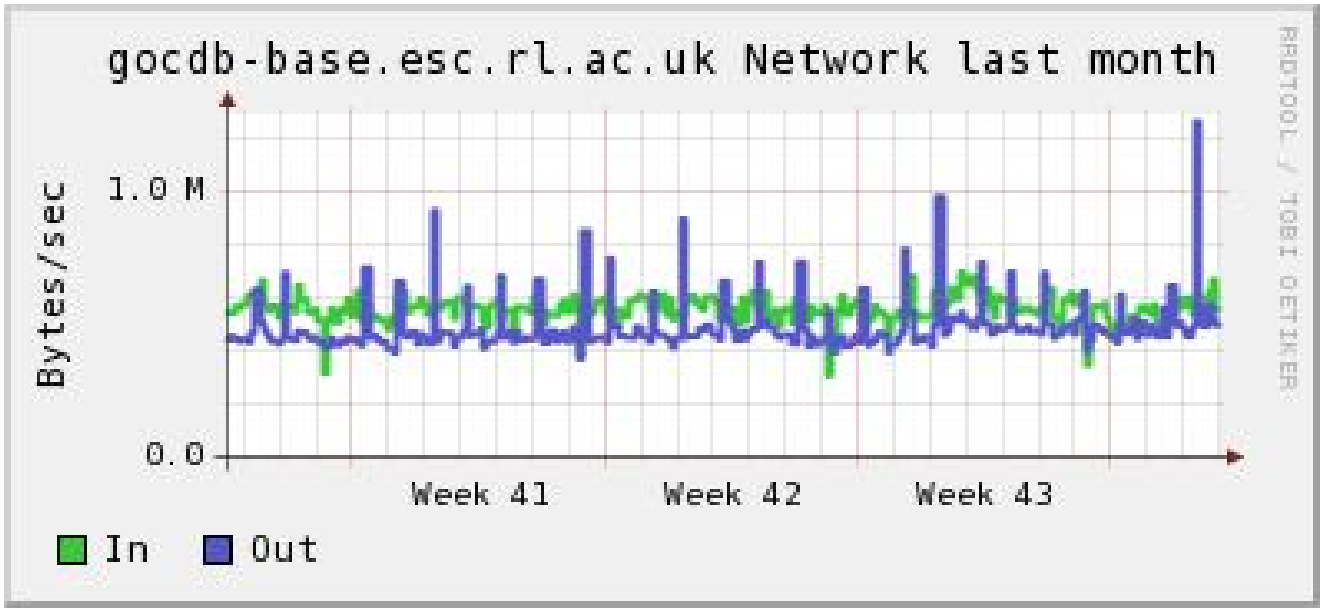

- The number of valid PI queries starting with URL fragment '/gocdbpi':
  - Sep=4.3M
  - Oct=5.9M

GOCDB Usage Statistics (3/3)

User interactions on the read/write input portal

- Usage of the read/write input portal can also be considered as high
  - The total number of separate page requests, excluding noise such as css and image file requests is:
    - Sep=8566
    - Oct=8974 (for Oct, this averages 289 user interactions per day).
  - A partial selection of different page requests for October and September are:

TOTAL GET/POST requests (excluding noise) (Sep=8566, Oct=8974)

POST new downtimes (Sep=391, Oct=457)

POST edit downtimes (Oct=175, Sep=148)

View object (Sep=3854, Oct=3975)

GOCDB Documentation

- **Main Entry Point: https://goc.egi.eu**

- https://wiki.egi.eu/wiki/GOCDB
- https://wiki.egi.eu/wiki/GOCDB/Documentation_Index
- https://wiki.egi.eu/wiki/GOCDB/Release4

GOCDB BACKUP

GOCDB PROM (1/3)

- PROM is a proprietary Object Relational Mapping (ORM) style solution for persisting data as objects in a relational database, and importantly, for recording parent/child relationships (links) between those objects *without* defining DDL schema constraints at the database-schema level (relationships are traditionally enforced in a relational database using PK/FK constraints).

- By excluding DDL schema constraints, a PROM database can accommodate schema changes *without* affecting existing software or data.

- For GOCDB, the intention is to allow NGIs to add their own custom data objects to a local installation whilst leaving the core GOCDB objects intact.

- While the PROM methodology certainly does provide flexibility, it also increases complexity. This is especially true when querying for object relationships. This is because queries must repeatedly join across the core ‘admin’ PROM tables in order to determine those relationships before those objects can be returned as linked object-graphs.

GOCDB PROM (2/3)

- A PROM database requires a standard set of core ‘admin’ tables that are specifically used to record information about the different types of data objects, and which of those objects are related (linked). For example, for every new data object, a new entry is recorded in the PROM ‘object list’ table (recording the object id and other meta-data such as the object’s table name).

- Similarly, parent/child relationships between objects are stored as individual entries in the PROM ‘object link’ table (each link entry associates a parent object of a particular type to a child object of a particular type).

- The rules for linking objects are defined in a separate ‘link type’ table. The core PROM admin tables include TOBJECTS, TOBJECTTYPES, TOBJECTLINKS, and TLINK_TYPES. For GOCDB, the custom ‘data’ tables include USER, SITE, SERVICE_ENDPOINT, DOWNTIME GROUP etc (see the logical and physical entity model diagrams).

GOCDB PROM (3/3)

- To ease working with a PROM database, a database agnostic ‘IPromAPI’ interface has been defined (introduced in GOCDB v4.2) which can be implemented for different databases as required (currently only implemented in PHP for Oracle).

- The IPromAPI interface is a CRUD style API and simplifies insertion/deletion/linking/updating of objects. Importantly, it also enforces the relationship rules as defined in the ‘link type’ table.

- Since object relationships are not defined in DDL, the ‘Physical PROM DDL’ is distinctly different from the ‘Logical Object Model.’ For more details on PROM see: https://wiki.egi.eu/wiki/GOCDB/PROM

Logical GOCDB Object Model (simplified)

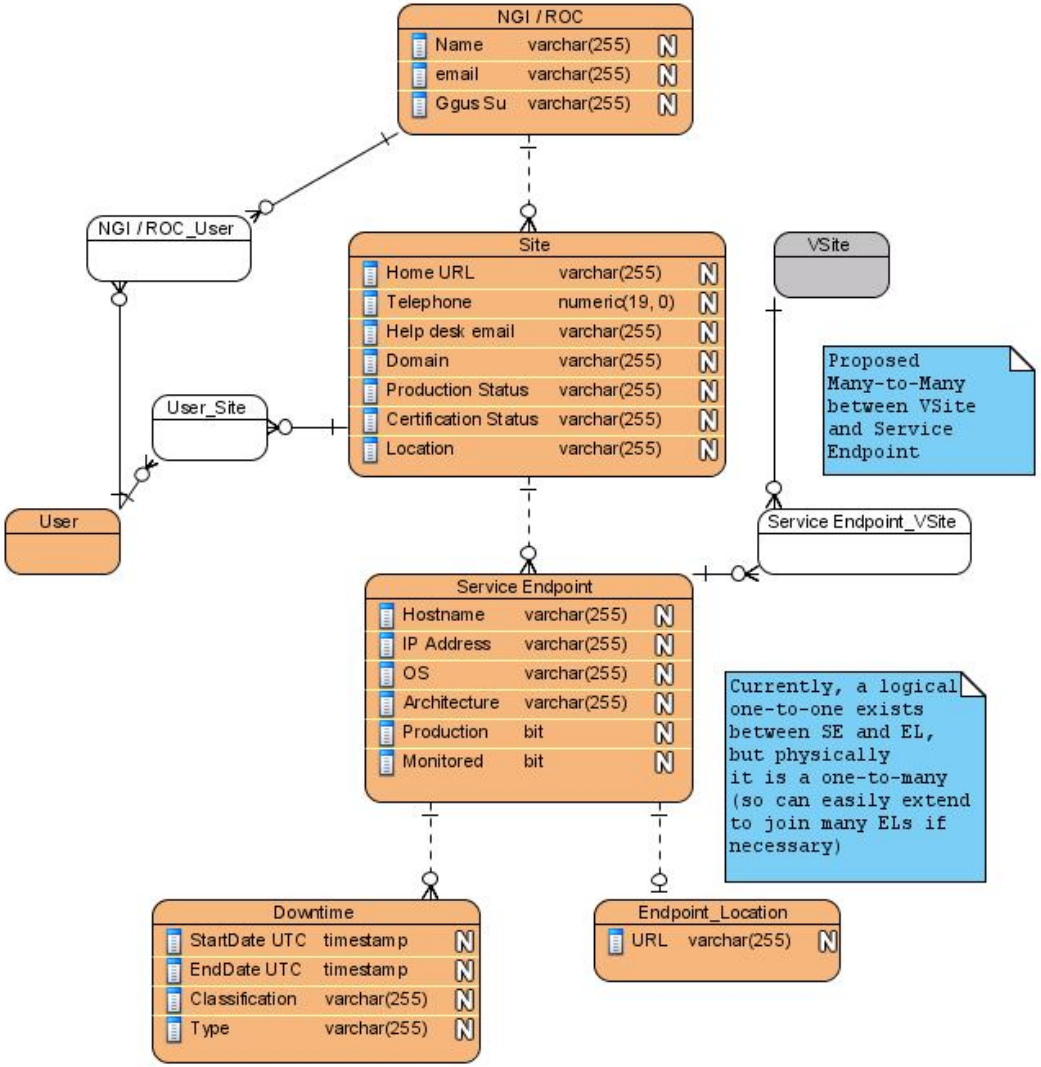

The relationships shown here are *logical* – they are not actually defined in a DDL schema with PK/FK constraints.

In a PROM database, relationships are enforced by the PROM API and the available link types recorded in the PROM admin tables (admin tables not shown, see the GOCDB physical entity model).

Physical GOCDB Entity Model (1/2)

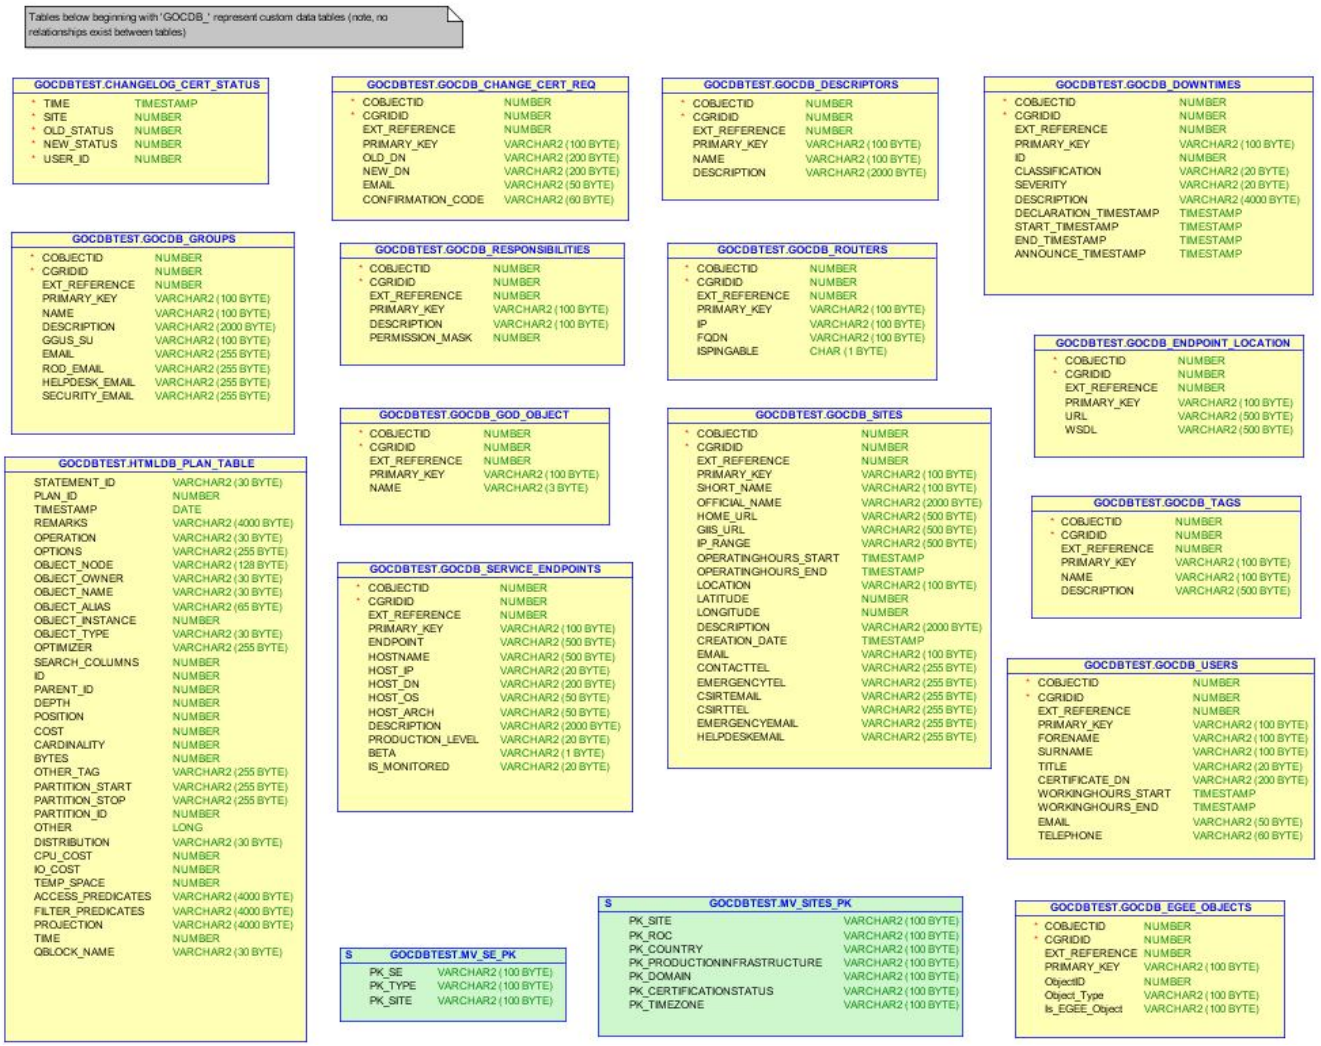

Note, PK/FK relationships do not exist between tables in a PROM database. Instead, they are enforced by the application using the PROM API.

The tables in the upper half of the diagram are the GOCDB specific data tables, while the tables in the lower half are the PROM admin tables.

Physical GOCDB Entity Model (2/2)

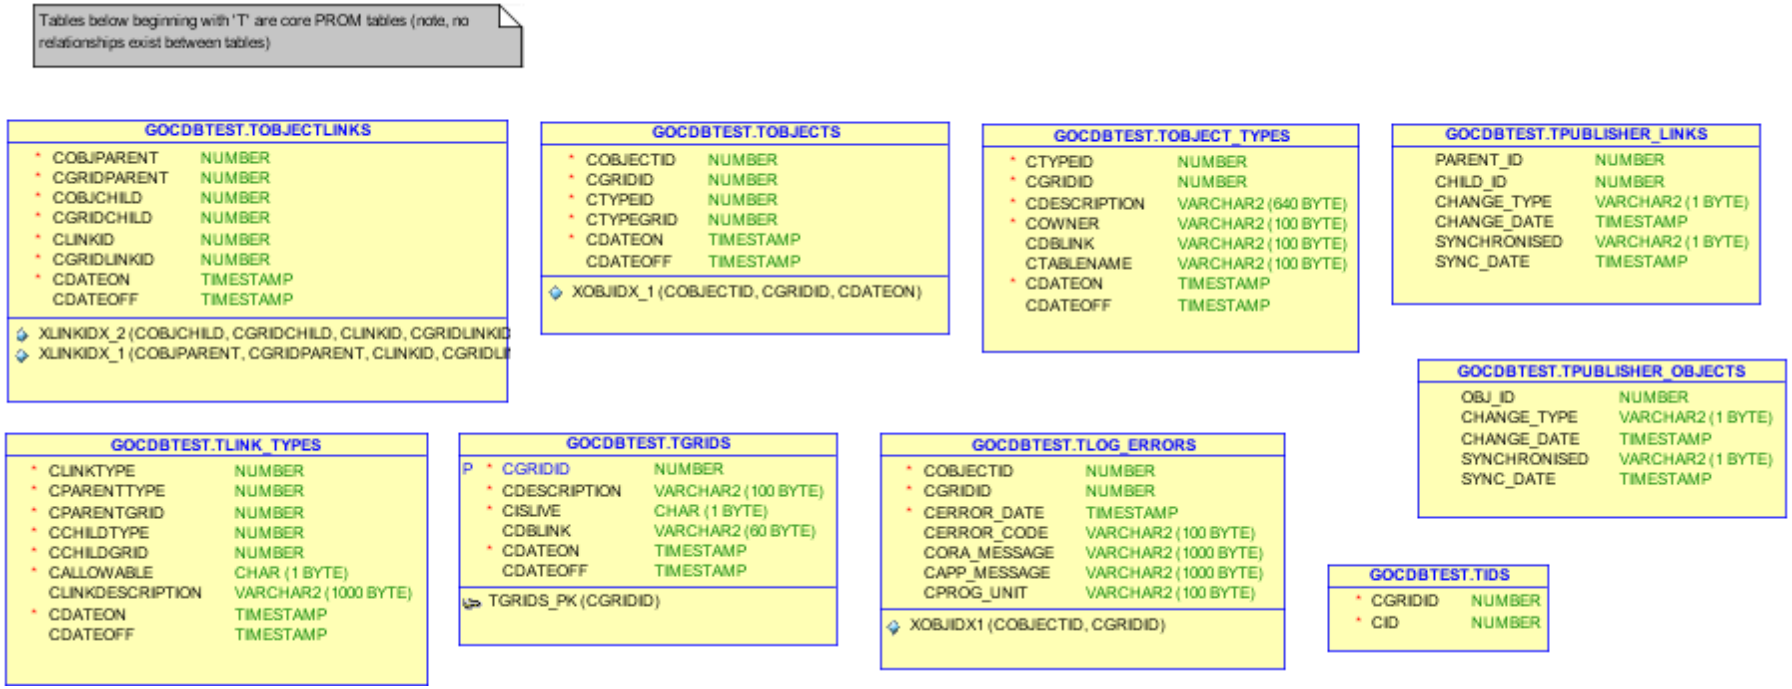

GOCDB PI

- The PI is a set of RESTful GET request methods for querying GOCDB data formatted in XML over HTTPS
- Each method can be parameterized in order to narrow the search results, e.g.:
  - select all SEs belonging to Site X
  - show all sites belonging to NGI X
- Some popular methods can be scoped, i.e.
  - https://goc.egi.eu/gocdbpi/private/?method=get_site&scope=Local (returns all Local scoped sites)

- https://wiki.egi.eu/wiki/GOCDB/PI/Technical_Documentation

GOCDB PI Methods

get_site	Returns site information including contacts, grouped by site
get_site_list	Returns a list of sites with minimal associated information
get_site_contacts	Returns a list of persons (and associated info) having a role at site level, grouped per site
get_site_security_info	Returns security contact information for sites
get_roc_list	Returns a list of NGIs with minimal associated information
get_subgrid_list	Returns a list of Subgrids (i.e. registered sub-parts of an NGI) with minimal associated information
get_roc_contacts	Returns NGI contact details, including NGI contact mail address and list of NGI staff
get_egee_contacts	Returns a list of contacts for staff that have a role a EGI level
get_downtime	Returns a list of EGI downtimes for sites and nodes
get_service_endpoint	Returns a list of service endpoints (single node x single service) and associated information
get_service_types	Returns a list of valid service types and associated description
get_user	Returns a user or a list of users with associated details and roles
get_downtime_to_broadcast	Returns the list of downtimes recently declared with notification settings for CIC portal downtime notification service
get_cert_status_changes	Returns a list of changes to certification statuses
get_cert_status_date	Returns a list of current certification statuses for Production sites and the date they entered that status

Status of Regionalisation Developments

- GOCDB developments in 2011 have focused on
  - a significant amount of re-development/refactoring
  - addressing new requirements for the central instance that have emerged throughout the year.
  - (more info on https://wiki.egi.eu/wiki/GOCDB/Release4/Development)
- Fixing the Atomic PROM API and addition of data-scoping/tagging are essential for the regional-publishing model described at: https://wiki.egi.eu/wiki/GOCDB/Release4/Regionalisation
- There are still three outstanding developments to be tackled before we can really start in earnest on regionalization:
  - Virtual Sites
  - a finer grained permission/authorization model
  - replacing the current XML output module as it cannot create nested XML collections (is necessary)
  - These are detailed at: https://wiki.egi.eu/wiki/GOCDB/Release4/Development#Development_Roadmap

Operations Portal

Operations Portal

Single access point to almost all operational information. Widely used in the day-by-day run of the GRID.

- Broadcast tool
- Operational Dashboard
- VO Information
  - IDCard
  - User Traking
  - Resources
- USCT Management Dashboard
- Security Dashboard
- COD Dashboard
- VO Oriented Dashboard (in progress...)

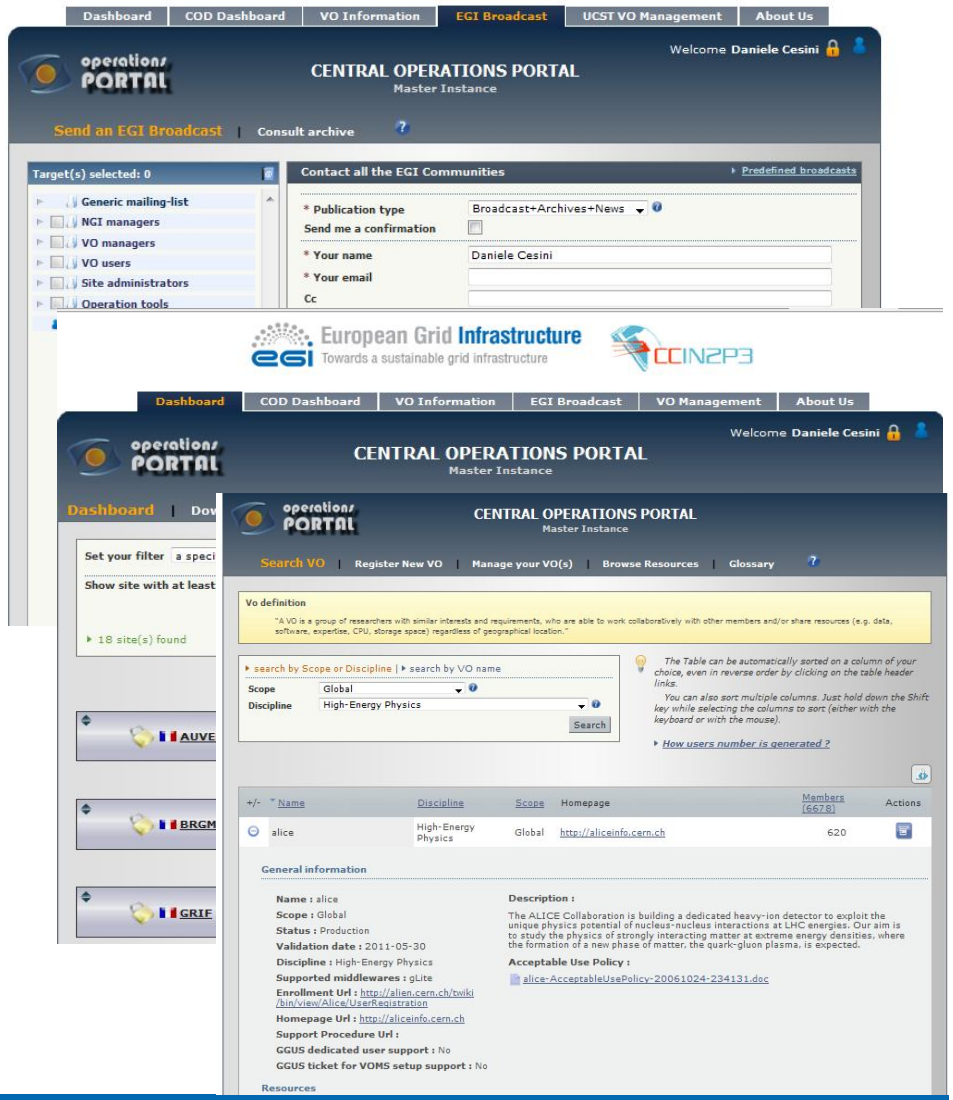

Operations Portal Architecture

Regional and Central Instances synchronized though Lavoisier

More info:
MS701 :
https://documents.egi.eu/document/39

MS705:
https://documents.egi.eu/document/27

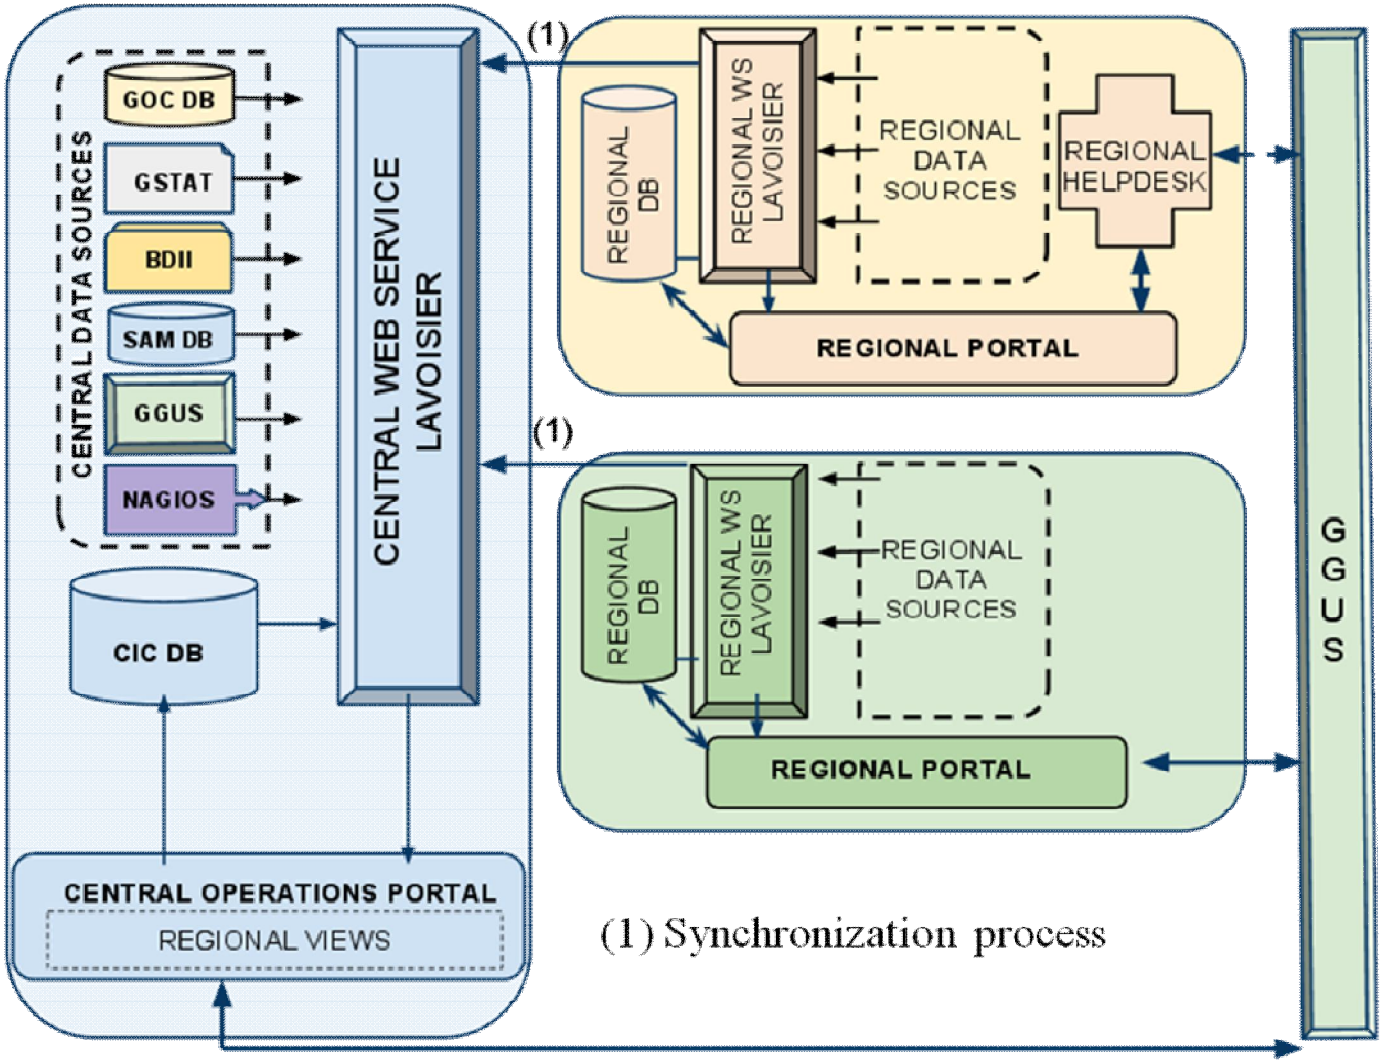

Ops Portal and Infosys

- Direct ldap queries to one Top-BDII
  - Service topology: where ? What ? And which VO ?
  - Every 2 hours
  - Cache mechanism: not replaced in case of failure
  - GOC DB information insufficient (VO Information)

- Query to GSTAT (per site and per VO)
  - status of the sBDII - CPU (site, vo) - Jobs (site, vo) - Storage (site, vo)
  - Every 6 min
  - Cache mechanism: not replaced in case of failure

Ops Portal and Infosys

- Information coming from BDII is used
  - By Downtime Notification System
  - Into the Resource Distribution browser
    - site and VO view
- Information coming from GSTAT is used
  - By the VO Module
  - By the Dashboard

Ops Portal and GOCDB

- GOCDB as official source of the following information:
  - services/sites/ngi topology
  - contact info
  - production status

- This information is enhanced by GSTAT or BDII queries

Ops Portal Requirement to InfoSys

- Continue to offer the same level of information available today by the three tools
- Enhance this information if possible
  - more information at the service level
    - storage distribution on specific SE

- More APIs
  - To query the BDII ldap commands are not enough flexible and easy usable
  - in GSTAT not possible to have a global view (egi_ngi , egee_roc, wlcg_tier)

SAM Framework

Service Availability Monitoring (SAM)

SAM monitoring framework for RCs and services

- one of the main data sources for the Operations Portal
- data source to create Availability/Reliability statistics
- components:
  1. test submission framework: based on the NAGIOS system set up and customized by the NAGIOS Configurator (NCG)
  2. databases for storage of information about topology, metrics and results
  3. visualization tool GUI: MyEGI

SAM architecture

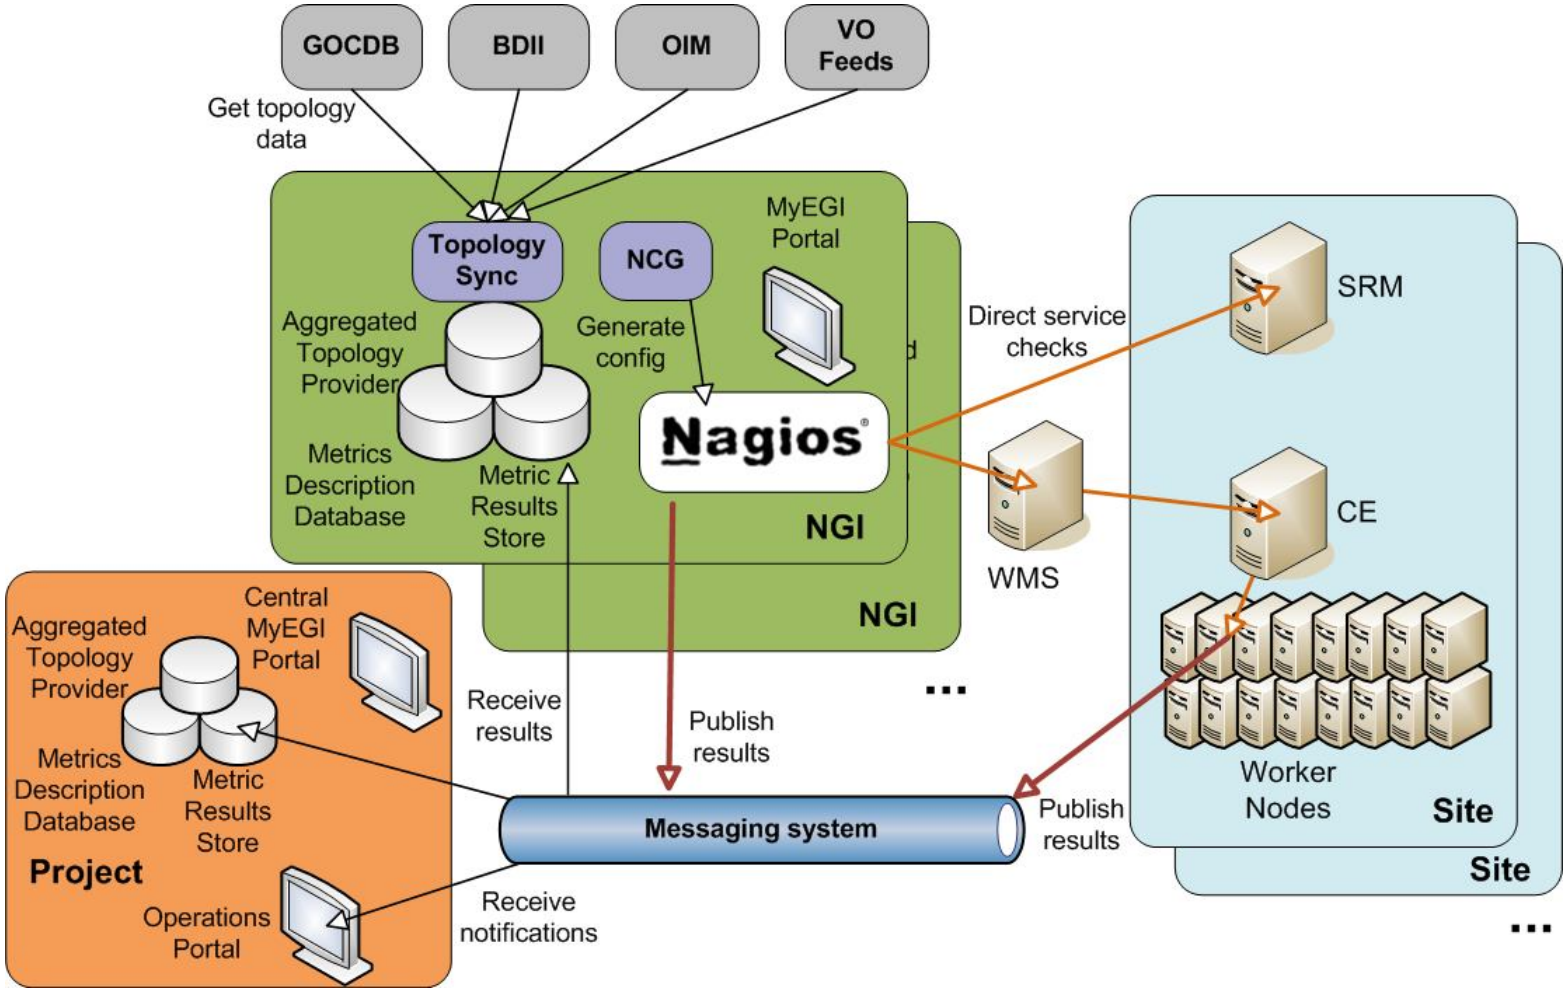

SAM and The Information System

- BDII, GSTAT and REBUS are used by SAM
- From BDII (direct ldap queries):
  - service URIs
  - supported VOs for each service
  - services supporting different MPI implementations;
- From GSTAT:
  - CPU information (PhyCPU, LogCPU, ksi2k, HEPSPEC06);
- From REBUS:
  - Tier and federations (Tier, Federation, FederationAccountingName, Site, Infrastructure, Country).

SAM and GOCDB

- SAM uses the GOCDB by querying the following feeds:
  - https://goc.egi.eu/gocdbpi/private/?method=get_site
  - https://goc.egi.eu/gocdbpi/private/?method=get_service_endpoint
  - https://goc.egi.eu/gocdbpi/private/?method=get_downtime
  - https://goc.egi.eu/gocdbpi/private/?method=get_service_types
- The Information System is needed to get extra information not provided by GOCDB
  - which VOs support each service and services supporting MPI implementations

SAM and Configuration

- Frequency of information retrieval = 30mins (configurable)
- A caching mechanism is used internally
  - Information is stored and updated every time it is modified.
- Infosys endpoints are in the configuration files (there is no discovery)

SAM Requirements to the InfoSys

- Current InfoSys provides sufficient information and we have no new requirements
- All the information provided is needed by SAM
  - Replacing one of the InfoSys would require getting the respective information from elsewhere
- It would be convenient to have a single place for all the information
- Also we would benefit if each InfoSys tool has a well defined scope (avoiding overlaps)

EGI Helpdesk (GGUS)

- EGI Helpdesk
  - distributed helpdesk with central coordination: Global Grid User Support (GGUS)

GGUS Architecture

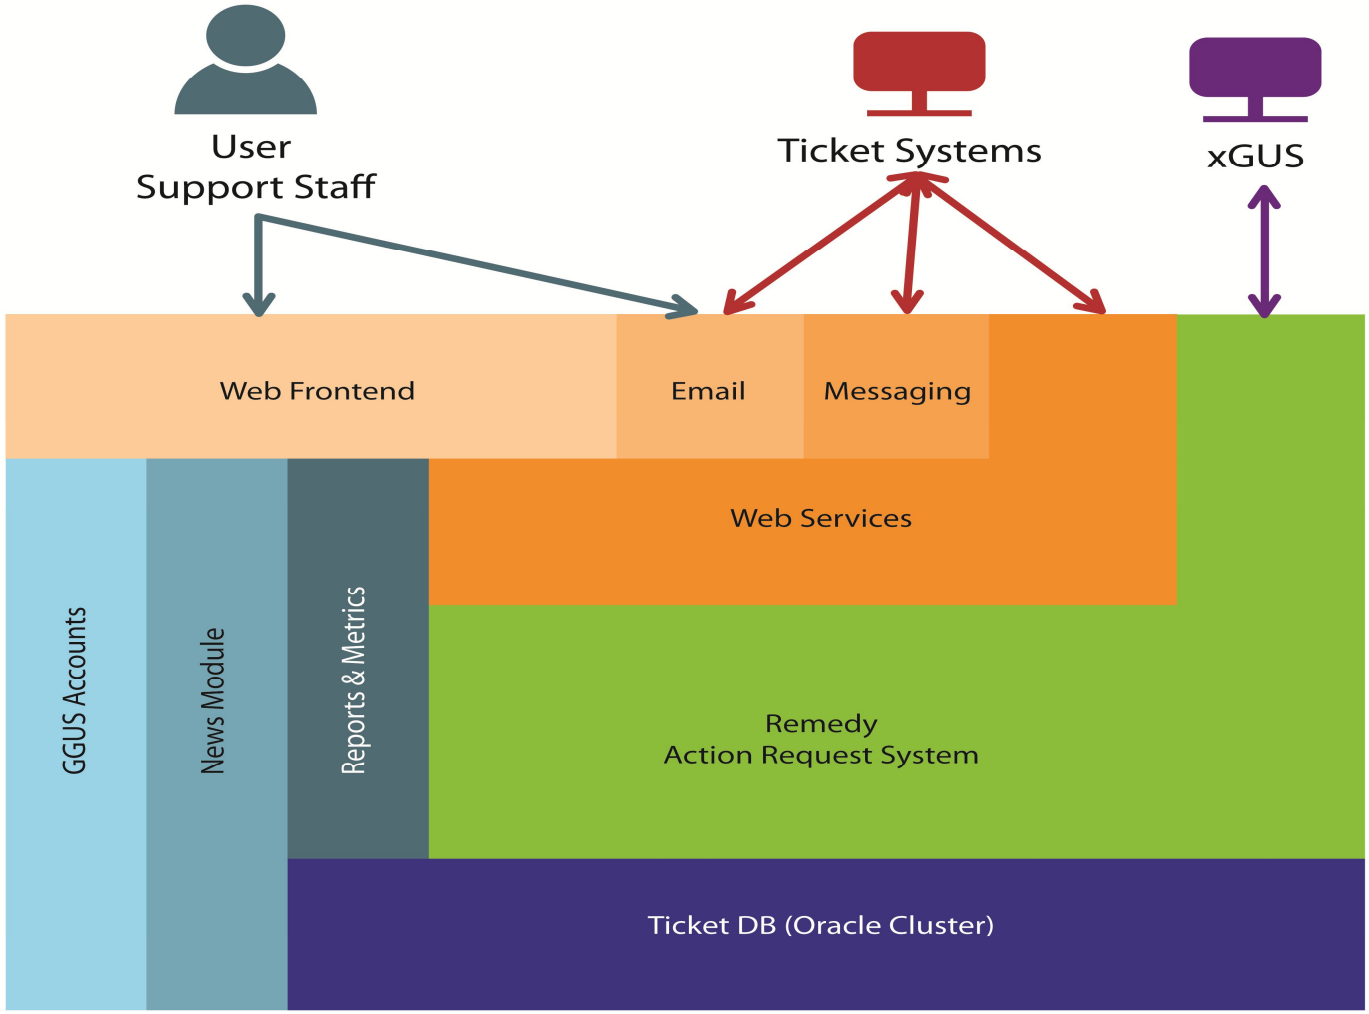

GGUS and GOCDB

- GGUS has NO dependencies on the Information System

- Has dependencies on GOCDB
  - Information retrieved every night through a Cronjob:
    - site data
    - contact mail address
    - site names

Accounting in EGI-JRA1

Accounting

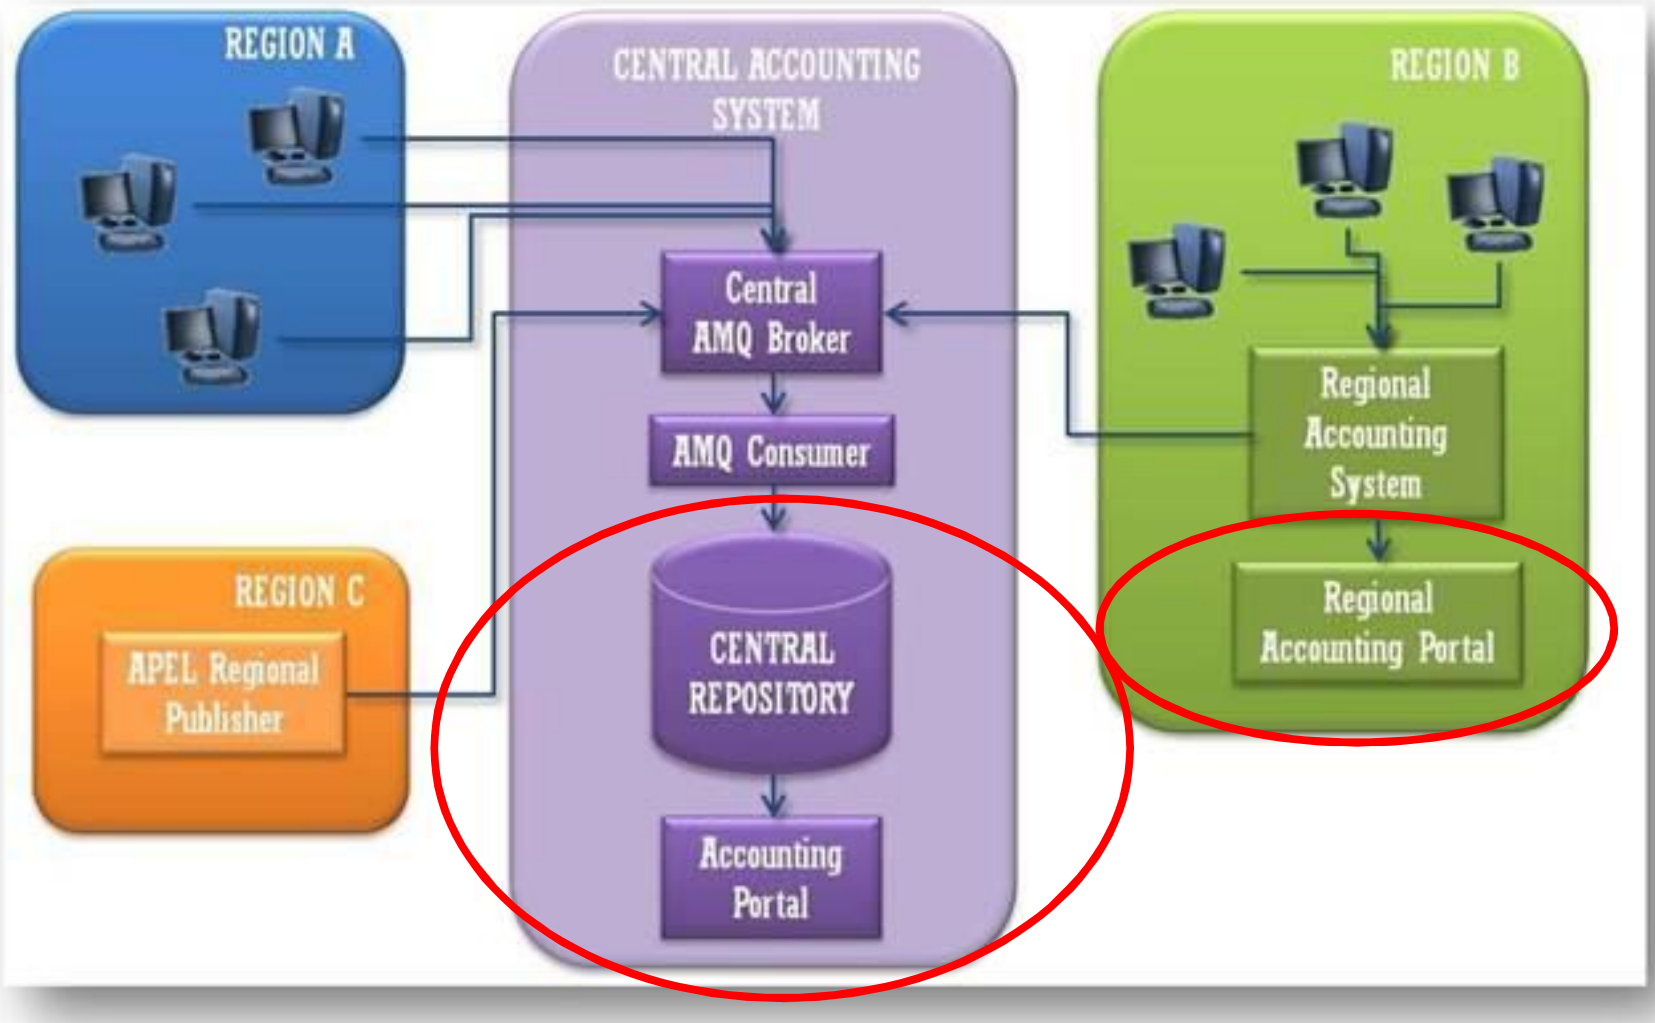
REGION A
CENTRAL ACCOUNTING SYSTEM
REGION B
Central AMQ Broker
AMQ Consumer
Regional Accounting System
REGION C
APEL Regional Publisher
CENTRAL REPOSITORY
Regional Accounting Portal
Accounting Portal

Accounting Repository and Portal

Accounting Repository (STFC)

- usage of compute resources within the production infrastructure
- based on gLite-APEL

Accounting Portal (FCTSG)

- GUI for access to data from the Accounting Repository

Accounting Repository and Infosys

- APEL retrieves batch scaling or benchmark values for a cluster from the BDII
- Uses them to normalize CPU accounting data

- Two ldap queries:
  - First search is the highest priority search.
  - If both searches return a SpecInt value for a cluster, the value from the first search is used; the second is disgarded.

Accounting Repository and Infosys

First search

Find a GIIS entry where
“objectclass=GlueCE”
(GlueCEUniqueID=lcgce06.gridpp.rl.ac.uk:2119/jobmanager-lcgpbs-grid1000M)

Get the “GlueForeignKey” attribute
(GlueClusterUniqueID=sl5-1000.gridpp.rl.ac.uk)

Remove “GlueClusterUniqueID=“ from the start
(sl5-1000.gridpp.rl.ac.uk)

Extract “SpecInt” from
“GlueCECapability”
(CPUScalingReferenceSI00=1000)

Second search

Find a GIIS entry where
objectclass=GlueSubCluster
(GlueSubClusterUniqueID=lcgce06.gridpp.rl.ac.uk)

Get the GlueChunkKey
(GlueClusterUniqueID=lcgce06.gridpp.rl.ac.uk)

Get the GlueSubClusterUniqueID
(lcgce06.gridpp.rl.ac.uk) Get SpecInt from

GlueHostBenchmarkSI00
(1000)

Accounting Repository and Infosys

- At most sites the accounting records are built locally
  - Normalization information could probably be stored elsewhere by APEL configuration
- At a few sites (eg France) accounting is published by one site for many
  - In this case the normalization values are retrieved remotely using the BDII.
  - Any alternative to this would either have to gather normalization data during the parsing phase on the CE, or establish some other remote publishing technique.

APEL and GOCDB

- APEL's uses GOCDB to retrieve information necessary to build
  - the EGI topology information for summaries
  - the ACLs for the APEL AMQ broker.

Accounting Portal and Infosys/GOCDB

- BDII is not used by the Accounting Portal
- GOCDB is used to retrieve topological information:
  - site information
  - tier status
  - country/ngi mappings, etc..
  - This information is gathered every 3 hours.

Metrics Portal

Metrics Portal

- Collects a set of metrics from different resources to help in measuring project performance
- Keeps track of the project evolution by displaying historical values of the metrics in a single place.
- It also provides web interfaces to inject the metrics into the database

Metrics Portal and Infosys

- The metrics portal was supposed to get information from GSTAT, but resulted to be not flexible
  - the first problem was separating EGI+EGEE nodes
  - there are metrics not covered by GSTAT
  - unnecessarily strict with the publication
  - and many sites disappear
- Now use BDII directly
- Information is collected for all EGI sites and then aggregated by NGI based on GOCDB
  - some NGIs are further partitioned into countries

Metrics Portal and Infosys

- For each site it gets weekly:
  - Sites using MPI
  - Disk and Tape Storage
  - Logical CPUs
  - HEPSPEC06 capacity
    - get SI2K, since HEPSPEC06 is reported with issues and inconsistencies

- No caching is used
  - If the infosys is not reached, the script sends an email to the maintainer
  - The update can be then done manually later

Metrics Portal and GOCDB

- Many of the above metrics are not usually covered by GOCDB, but if GOCDB carried all the needed information sufficiently updated the Metrics Portal would probably use it instead of BDII

- From GOCDB retrieves data for mapping sites to NGIs
  - for some NGIs the mapping is refined to countries
    - e.g. Ibergrid -> Spain+Portugal
    - NGI_NGDF -> Finland, Denmark, Sweden and Norway

Metrics Portal future Usage of the InfoSys and Requirements

- The future uses depend on the metrics evolution
- If NGI wide BDIIs are available it would be possible to use them

Requirements:

- Better coverage for HEPSPEC06 information
- Info on Cloud resources should be made as seamless as possible to have a unified view of the total composition the infrastructure (Grid+Cloud)

Summary Table

	BDII	GSTAT	GOCDB	BDII query Frequency	Cached	IS Clients
SAM	Yes	Yes	Yes	30 mins	Yes	BDII: ldapsearch (will use ARC Lib and Unicore CLI) GOCDB PI
Ops Portal	Yes	Yes	Yes	BDII: 2hrs GSTAT: 6mins	Yes	BDII: ldapsearch GOCDB PI
Acc Repo	No	No	Yes	-	-	GOCDB PI
Acc Portal	No	No	Yes	GOCDB: 3hrs	-	GOCDB PI
Metrics Portal	Yes	Yes	Yes	weekly	-	BDII: ldapsearch GOCDB PI
GGUS	No	No	Yes	Every night	-	GOCDB PI
GOCDB	No dependecies					

Conclusion

- From an EGI-JRA1 perspective BDII, GSTAT and GOCDB are all needed as source of information
  - BDII and GSTAT are mainly used to complement GOCDB info for VO mapping
- No big requirements on the Information System
  - Improvement of some features, i.e.:
    - better coverage for HEPSPEC06 information
    - More information at service level (SE)
  - APIs to query the BDII would be great
  - Stability of information could be improved but is obtained with internal caching mechanisms
  - Accounting data punctual normalization can provide new requirements depending on how it is implemented
- It would be convenient to have a single place for all the information