



Science and
Technology
Facilities Council

GOCDB Technical Road map

Greg Corbett

gocdb-admins@mailman.egi.eu



Science and
Technology
Facilities Council

Grid Operations Centre Database (GOCDB)

GOCDB is a repository for storing and presenting topology and resource information, i.e. service endpoints, site and service contacts and downtime schedules, etc

Used within multiple nationally and internationally distributed digital research infrastructures, such as the IRIS community, the EGI Federation and the Worldwide LHC Computing Grid (WLCG).

It is a definitive information source, with the emphasis on infrastructure and service operators maintaining their own information.

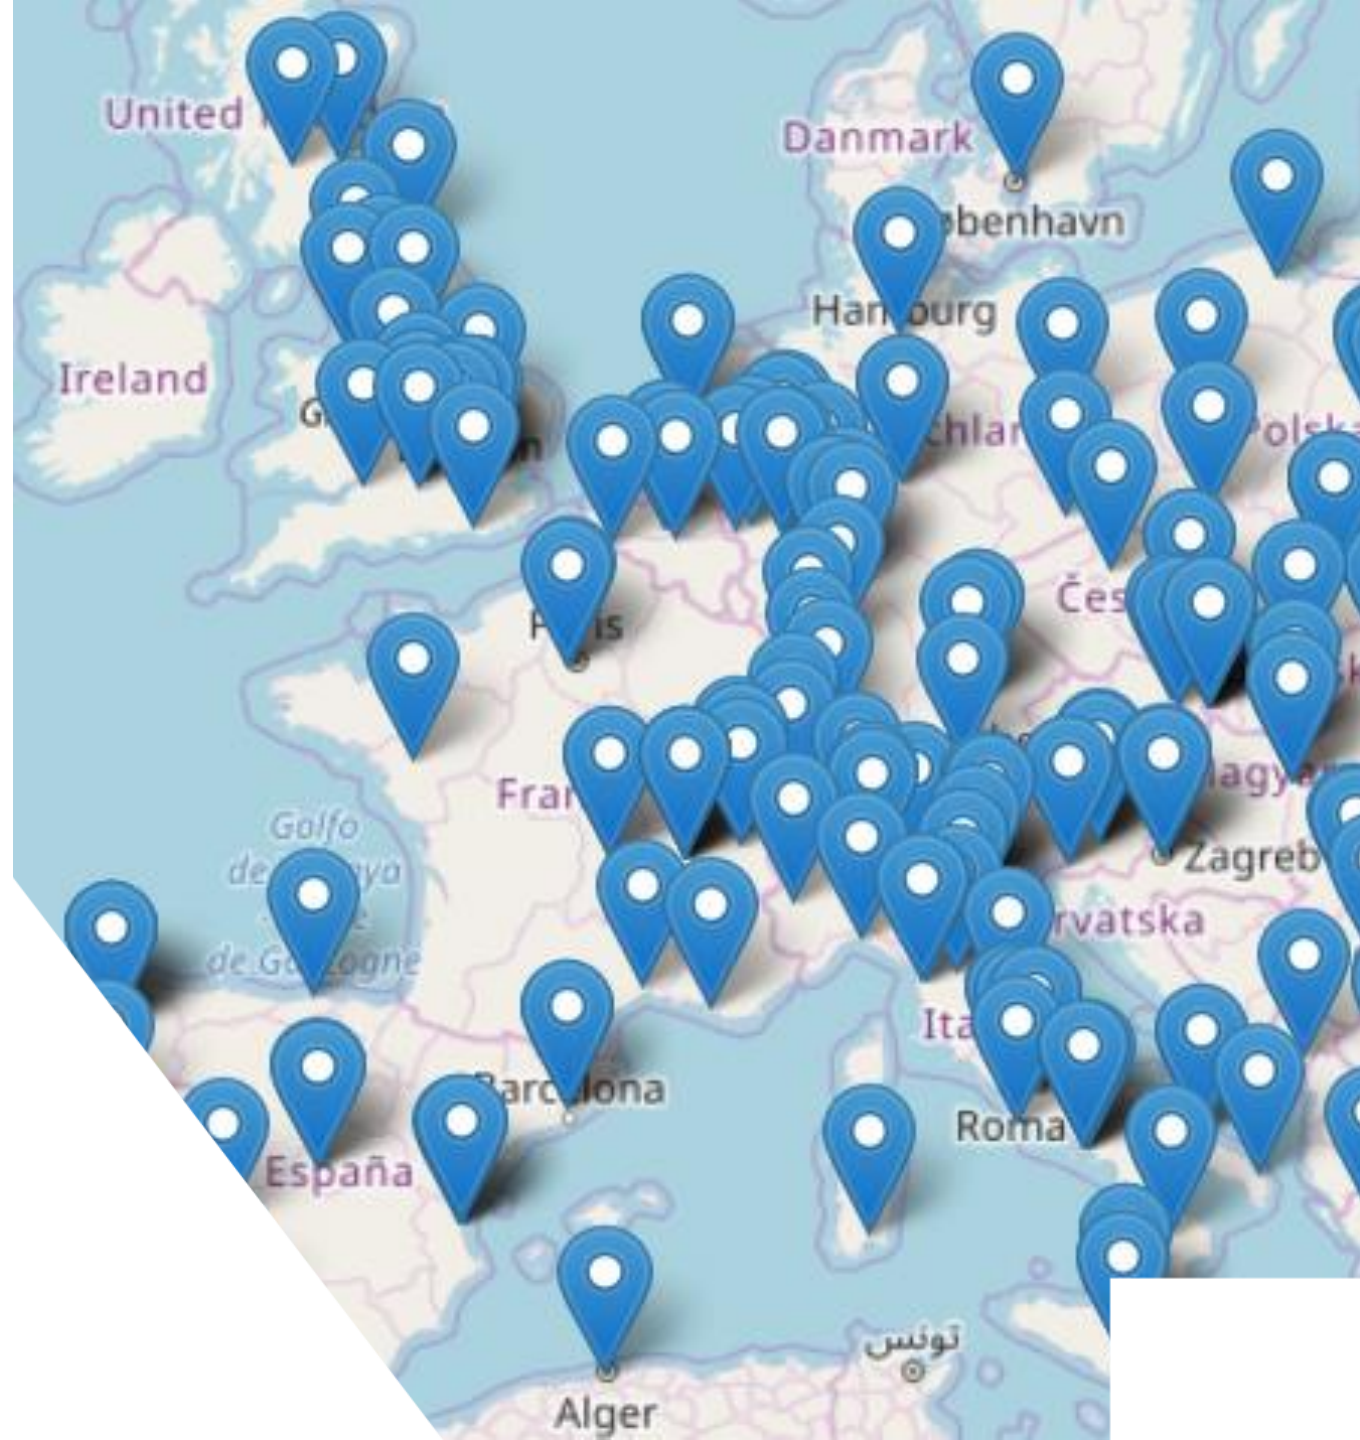
Agenda

1 Recent Releases

2 Upcoming Releases / Work

3 2022/23

4 2023/24



Recent Releases

5.7.6 – October [1]

Bug Fixes around MariaDB support only

Will enable GOCDB to move off of its current Oracle database, motivated by licensing fees.

Recent Releases

5.8.0 – November. Large feature release, code only [2].

Including, but not limited to:

- Write API expansion [3]
- Separation of DNs in extension properties with ‘<’ and ‘>’ [4]
- Enhancements that will allow us to easily differentiate different instances (i.e. production and pre-production) [5]
- Gender neutral or empty title options for users [6]
- Allowing users to remove their old “Home Site” [7]

Upcoming Releases

5.8.1 – December [8]

Minor bug fixes as well as code quality / look and feel improvements

Upcoming Releases

5.9.0 – March 2022 [9]

Support in codebase for authentication with Indigo Identity Access Management technologies via OIDC [10]. Precursor to integrating with WLCG IAMs and other IdPs (EOSC AAI)

Identity Linking within GOCDB [11]

Prevent downtimes starting in 24 hours from being deleted [12]

Identity Linking within GOCDB

Currently, EGI Check-In identifiers (6fa...3fb@egi.eu) appear in the (user) CERTDN field of the API.

The work to support Identity Linking fixes this bug by showing the ID string of the user in a correctly labelled XML tag (see below).

```
<CERTDN/>  
<EGICHECKIN>6fa...3fb@egi.eu</EGICHECKIN>
```

```
<CERTDN>/C=UK/.../CN=greg corbett</CERTDN>  
<EGICHECKIN/>
```

```
<CERTDN>/C=UK/.../CN=greg corbett</CERTDN>  
<EGICHECKIN>6fa...3fb@egi.eu</EGICHECKIN>
```

Identity Linking within GOCDB

If your service relies on the CERTDN field in the API containing non X.509 DN values, you will need to extend your service to also process the EGICHECKIN tag.

We plan on deploying GOCDB 5.9.0 into preproduction by the end of Jan 2022 to allow services time to make these extensions.

We have identified sites hosting services that access user identifiers via the API and will get in touch with them after the deployment to preproduction.

We can provide sample API output now, if helpful to service teams.

Database Migration

GOCDDB currently uses an ORACLE database.

This is a risk to the long-term sustainability of the service due to the license fees required.

The GOCDDB team have avoided paying this in the past due to other larger Oracle use-cases in our organisation bearing the cost.

This will not continue however as this use cases are reaching end of life.

We will be migrating the database to a GaleraDB service with STFC.

Database Migration

We have already verified the data migration process does not change the data itself.

GOCDDB 5.7.6 ensured current behaviour is preserved

Other instances of GOCDDB have been running GOCDDB on MySQL / Maria DB style databases for some time, and development work is initially tested against on these platforms as well.

Currently implementing a new failover pipeline.

2022/23

OIDC Based Integration with multiple IdPs

- Motivated by both EGI's and WLCG's aim of using OIDC-based end user authentication mechanisms across the core services.
- Reworking existing portal integration with the EGI Check-In service
- Allowing access to the API via non-X.509 credentials for the first time.

Improvements to User Privacy

- Restricting access to personal information to:
 - registered users with a role over any entity
 - robots using an "API identifier" that has been registered with GOCDB by a registered user (with a role) as needing access to personal data.

2023/24

Preparing GOCDB for SL7 End of Life

- Updating internal dependencies
- Testing deployments
- Failover



Science and
Technology
Facilities Council

Questions?



Science and
Technology
Facilities Council

Thank you



Science and Technology Facilities Council



@STFC_Matters



Science and Technology Facilities Council

Links

- [1] <https://github.com/GOCDB/gocdb/milestone/5>
- [2] <https://github.com/GOCDB/gocdb/milestone/1>
- [3] https://wiki.ege.eu/wiki/GOCDB/Write_API/Technical_Documentation#Supported_methods_in_GOCDB_5.8
- [4] <https://github.com/GOCDB/gocdb/pull/221>
- [5] <https://github.com/GOCDB/gocdb/pull/215>
- [6] <https://github.com/GOCDB/gocdb/pull/250>
- [7] <https://github.com/GOCDB/gocdb/pull/268>
- [8] <https://github.com/GOCDB/gocdb/milestone/7>
- [9] <https://github.com/GOCDB/gocdb/milestone/6>
- [10] <https://github.com/GOCDB/gocdb/pull/260/files>
- [11] <https://github.com/GOCDB/gocdb/pull/294>
- [12] <https://github.com/GOCDB/gocdb/issues/269>