



Science and
Technology
Facilities Council

Improving ICAT Performance

Kevin Phipps

May 2023 (ICAT F2F meeting in Berlin)

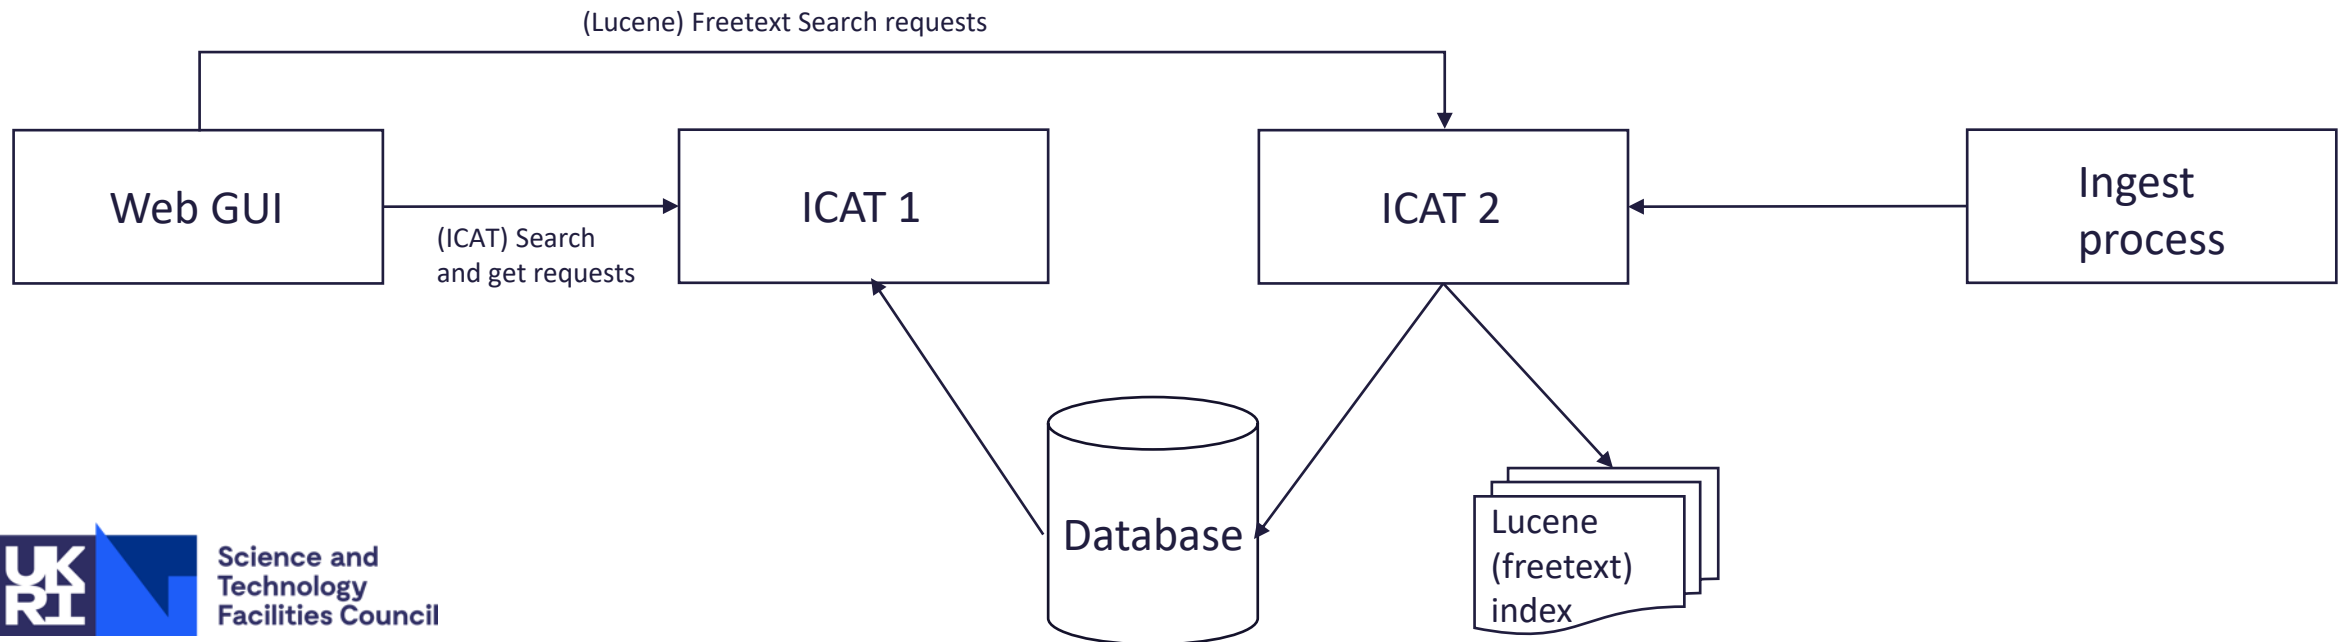
Introduction

- A few ways that I'm aware of to:
 - scale ICAT to handle more requests
 - improve the performance of ICAT
 - investigate performance problems

NOTE: These slides were not presented at the meeting due to lack of time but I include them here as I created the material and I hope it is useful to somebody 😊

Scaling method 1: “unofficial” clustering

- This is how the Diamond ICAT has been run for many years
- Install 2 ICATs pointing to the same database
- Use one ICAT for ingest and the other for recall
- Slight oddity: freetext search requests have to go to the ingest ICAT unless you set up the Lucene component on a separate machine



Scaling method 2: “official” clustering

- This was created many years ago as the official way to create a cluster
- The ICATs are able to pass messages to each other where necessary
- WARNING: I don't know of anyone who has tried it apart from the original developer!
- But worth looking into and testing if you need this functionality

```
github.com/icatproject/icat.server/blob/master/src/main/config/run.properties.example
52 !lucene.entitiesIoIndex = DataFile Dataset Investigation InvestigationUser DataFileParameter
53
54 # List members of cluster
55 !cluster = http://vm200.nubes.stfc.ac.uk:8080 https://smfisher:8181
56
57 # Optional key which must match that of the IDS server if the IDS is in use and has a key fo
58 !key = ???
```

repo.icatproject.org/site/icat/server/5.0.0/installation.html

Installing a group of ICATs

If your facility depends upon a single ICAT instance then ingestion of data can be held up by a user making a request to the single instance.

To avoid this it is suggested that you install multiple servers each running a Glassfish with an ICAT but all on the same network. Finally ingestion can be directed to one node and the other nodes can be loaded as needed.

All machines must use the same database and icat.lucene server

Authentication can either be carried out on one machine - which has the advantage that you only have a single set of restful authenticators then more than one equivalent authenticator server may be specified to make the system more resilient.

The icat.servers are linked by specifying a cluster parameter in the run.properties file for each machine.

You could then set up an Apache front end to do load balancing. This will probably just connect to the same set of servers.

Performance improvements: ICAT rules

- ICAT rules can be expensive (time-wise) to evaluate
- Use PublicTables and PublicSteps to allow the authorisation system to take “shortcuts”
- PublicTables:
Any tables which contain no data which is sensitive can be readable by everyone
eg. InvestigationTypes, DatasetTypes, DatafileFormats
- PublicSteps:
Allow related entities to be INCLUDED without re-evaluating the rules when the parent entity has already been authorised

Performance improvements

- If you have slow queries, consider whether the standard ICAT indices cover the query that you are doing
- If necessary, add additional indices
- We do this in an ICAT we run for the Central Laser Facility:

```
12
13 -- eCAT2 does most searches using the dataset name and type so index these
14 create index DATASET_NAME_AND_TYPE on DATASET (NAME, TYPE_ID);
15 -- indexes to speed up the queries made every minute to check for new data
16 create index DATAFILE_MODTIME on DATAFILE (MOD_TIME);
17 create index DSPARAMETER_MODTIME on DATASETPARAMETER (MOD_TIME);
18
```

Performance problems

- On Oracle, an “Optimizer” analyses each query that comes in and selects an “Execution Plan” to retrieve that data with the lowest resource usage
- Once it recognises regular patterns of queries, it will tend to find a plan that returns the data quickly and continue to use that plan
- Occasionally, it decides to switch to an alternative plan which results in slower queries but it doesn’t seem to recognise this
- With help from our DBAs, they locate the previous plan used for that query and “pin” it so that Oracle switches back to the previous plan.
- From memory, we seem to have this problem every couple of years
- TIP: if you need the actual SQL that ICAT is sending to the database, add the lines below to persistence.xml and the SQL will be in the Payara server log:

```
<property name="eclipselink.logging.level" value="FINE"/>  
<property name="eclipselink.logging.level.sql" value="FINE"/>  
<property name="eclipselink.logging.parameters" value="true"/>
```



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