



# APEL & MySQL

Alison Packer

Richard Sinclair

# APEL

- **A**ccounting **P**rocessor for **E**vent **L**ogs
- extracts job information by parsing batch system (PBS, LSF, SGE or Condor) event log files and blah accounting log files
- LDAP queries of the CE obtain the CPU performance figures for the worker node clusters.
- inserts results into a local MySQL database.



# The APEL accounting database

- MySQL database installed as well as APEL client software.
- Following tables:

BlahdRecords

EventRecords

GkRecords

LcgProcessedFiles

LcgRecords

MessageRecords

RepublishInfo

SpecRecords



# Archiving old records

Archive old records to reclaim space:

1. Get the latest successful publishing date -

```
mysql> select * from RepublishInfo
```

Shows last MeasurementDate records published

2. Choose dates to archive, then dump the records processed before that date to file...



# Saving the data

- `mysqldump -u <DBUSERNAME> -p  
<DATABASE> <TABLE> --  
where="<DATEFIELD> < ' <YYYY-MM-  
DD> ' " > <FILE>.sql`

- **Example:**

```
# mysqldump -u accounting -p  
accounting LcgRecords --  
where="MeasurementDate < '2011-05-  
01' " > Lcgrecords.sql
```



# Saving the data contd.

- For all other tables there is also a processed flag to add to the mysqldump command:

```
# mysqldump -u accounting -p accounting  
EventRecords --where="EventDate <  
'YYYY-MM-DD' and Processed = 1" >  
EventRecords.sql
```

```
# mysqldump -u accounting -p accounting  
BlahdRecords --where="ValidFrom <  
'YYYY-MM-DD' and Processed = 1" >  
BlahdRecords.sql
```



```
# mysqldump -u accounting -p accounting  
MessageRecords --where="ValidFrom <  
'YYYY-MM-DD' and Processed = 1" >  
MessageRecords.sql
```

```
# mysqldump -u accounting -p accounting  
GkRecords --where="ValidFrom < 'YYYY-  
MM-DD' and Processed = 1" >  
GkRecords.sql
```



# Delete old records from the tables

- Backup the filesystem where your archived mysqldump files were saved
- Check the backup was successful





# Restore from sql files

- Create a new database (e.g. called accounting\_restore)
- Restore records from file to new database, e.g.

```
# mysql -u accounting -p  
accounting_restore < LcgRecords.sql
```

etc.



# Delete the records

```
mysql> delete from EventRecords where  
EventDate < 'YYYY-MM-DD' and Processed  
= 1;
```

```
mysql> delete from BlahdRecords where  
ValidFrom < 'YYYY-MM-DD' and Processed  
= 1;
```

```
mysql> delete from LcgRecords where  
MeasurementDate < 'YYYY-MM-DD';
```



# Delete contd.

```
mysql> delete from MessageRecords where  
ValidFrom < 'YYYY-MM-DD' and Processed  
= 1;
```

```
mysql> delete from GkRecords where  
ValidFrom < 'YYYY-MM-DD' and Processed  
= 1;
```



# Optimize tables to reclaim space

- Optimizing tables reclaims space which is not automatically reclaimed by deleting the old records:

```
mysql> optimize table EventRecords;
```

```
mysql> optimize table LcgRecords;
```

```
mysql> optimize table BlahdRecords;
```

```
mysql> optimize table MessageRecords;
```

```
mysql> optimize table GkRecords;
```



# The SpecRecords table

- When APEL parser is run, it connects to the local GIS and retrieves the latest SpecInt values for every cluster/subcluster.
- This information is stored in the SpecRecords table



# SpecRecords table - wrong value

- Ensure parser configuration is correct for retrieving values from GIS, check this in

```
/etc/glite-apel-  
batchsystem/parser-config.xml
```

in `<CPUProcessor>` **section**



# CPUProcessor Info

CPU Scaling value is either retrieved from LDAP query or may be defined manually.

Either (SpecInt:SpecFloat ):

```
<DefaultCPUSpec>300:200</DefaultCPUSpec>
```

Or:

```
<GIIS>ldap://site-bdii.gridpp.rl.ac.uk</GIIS>
```

N.B. if both, the default value will be used in preference to querying the LDAP server.



# Wrong value & records published

- Ensure publishing is now with correct values
- Know the dates where the records had incorrect SpecInt2000 values – see EGI Accounting Portal
- Update the fields of incorrect records in LcgRecords table and republish ...





# Update Incorrect Records

Check the records which need updating, example for September, 2011:

```
mysql> select * from LcgRecords  
where MeasurementDate >= '2011-09-  
01' and MeasurementDate < '2011-  
10-01';
```

Check the SpecInt2000 field in the output for those with the incorrect value



# Update SpecInt2000 errors

```
mysql> update LcgRecords set  
SpecInt2000 = <NEWVALUE>  
where SpecInt2000 =  
<OLDVALUE> and  
MeasurementDate >= '2011-09-  
01' and MeasurementDate <  
'2011-10-01';
```



# Republish Records

- Run the apel publisher in "gap" publishing mode, set in the `publisher-config.xml` file

- **example:**

```
<Republish recordStart="2011-09-01" recordEnd="2011-09-30">gap</Republish>
```



# Corrupt Database

- You can check a table is OK by running:

```
mysql> check table <TABLENAME>;
```

- If there are errors you can run:

```
mysql> repair table <TABLENAME>;
```

**This can take a LONG time!**



# Links

- [APEL-SUPPORT@JISCMAIL.AC.UK](mailto:APEL-SUPPORT@JISCMAIL.AC.UK)

