

Workload Management: How can UGE use PostgreSQL for the spooling Database?

How can UGE use PostgreSQL for the spooling Database?

For clusters running UGE, its plain text spool is rarely a performance bottleneck.

For busy enough clusters the BerkeleyDB is meant to be a faster way to spool, but it has stability issues at the time of writing (May 2015). A spool running on Postgres can be used instead, as follows:

1- Install the needed postgres packages, then initialize and start postgresQL:

```
# yum install postgresql-server postgresql
```

```
# service postgresql initdb
```

```
# service postgresql start
```

```
# su - postgres
```

```
$ createuser -S -D -R -l -P -E ugeadmin
```

```
$ createdb -O ugeadmin ugespool
```

2- Edit "/cm/shared/apps/uge/var/cm/inst_template.conf" and change the following parameters:

```
SPOOLING_METHOD="postgres"
```

```
PG_SPOOLING_ARGS="host=master dbname=ugespool user=ugeadmin"
```

Notes:

- The spooling arguments should be changed according to the Postgresql database configuration.
- The cm-uge package provides the inst_template.conf.template file. This file should be copied to inst_template.conf so that it gets recognized by the UGE installer which gets called implicitly by the wlm-setup utility.

Workload Management: How can UGE use PostgreSQL for the spooling Database?

3- Set up UGE using wlm-setup

```
# wlm-setup -s -w uge -a <path to UGE tarballs>
```

Unique solution ID: #1266

Author: Frank Furter

Last update: 2015-05-22 19:51