

Table of Contents

Issues and workarounds.....	1
Pending issues - no workaround yet.....	1
Known EMI-1 issues and already implemented workarounds.....	1
LB fills /tmp with sockets -- https://ggus.eu/ws/ticket_info.php?ticket=80343	1
LB fill /var/tmp with RTM notification files -- https://ggus.eu/ws/ticket_info.php?ticket=80341	1
Nordugrid CEs crashing Condor's Gridmanager continuously -- https://ggus.eu/ws/ticket_info.php?ticket=80215	2
Cron job deletes /var/proxycache when no activity -- https://ggus.eu/ws/ticket_info.php?ticket=78295	2
Condor 7.4.2 not truncating Matchlog log file.....	2
Missing job to clean up LBProxy database entries -- https://ggus.eu/ws/ticket_info.php?ticket=75881	3
LCMAPS logging problems -- https://ggus.eu/ws/ticket_info.php?ticket=80065	3
Log file generation/rotation is not compliant with grid policies -- https://ggus.eu/tech/ticket_show.php?ticket=77161	3
High memory usage of in glite-wms-workload daemon -- https://ggus.eu/ws/ticket_info.php?ticket=76149	3

Issues and workarounds

Pending issues - no workaround yet

WMS log file location is still split between `/var/log/wms` and `/var/log/glite`. Everything should go to `/var/log/glite`. Supposed to be fixed in Update 11 but it's not (<https://savannah.cern.ch/bugs/?86682>; problem re-open in https://ggus.eu/ws/ticket_info.php?ticket=80065). The `gservices-wms-admintools` RPM will have to be updated with the new path when the log location changes, log rotation to be checked as well (`altlogrotate` component + `yaim`-generated entries in `logrotate.d`).

Critical security information is missing from `/var/log/messages` -- https://ggus.eu/tech/ticket_show.php?ticket=79938. Ticket open by security team. If this is fixed, then we probably don't need TSM archiving of log files as critical information will be collected by the security team via `rsyslog`.

WMS does not support Argus yet as the central authorization service -- expected for EMI2 (https://savannah.cern.ch/task/?func=detailitem&item_id=23030 and <https://savannah.cern.ch/bugs/?90760>)

WMS does not report correctly job exit code -- potentially affects Atlas software installation -- https://ggus.eu/ws/ticket_info.php?ticket=79932

Alarm on LB mysql database size not functional -- <https://cern.service-now.com/service-portal/view-incident.do?n=INC111804> -- re-enable in `wms_lemon.tpl` when fixed or workaround found.

ICE deadlocks cause ICE job submission to stop, ICE must be restarted -- https://ggus.eu/tech/ticket_show.php?ticket=81161 Recovery procedure in `WmsOperationsAndTroubleshooting`.

WMS LogMonitor inefficiency - old jobs block progress on new jobs -- https://ggus.eu/tech/ticket_show.php?ticket=82586

Known EMI-1 issues and already implemented workarounds

LB fills /tmp with sockets --

https://ggus.eu/ws/ticket_info.php?ticket=80343

Fixed in EMI2: <https://savannah.cern.ch/bugs/?92708>

As a workaround we have a Lemon actuator restarting `glite-lb-locallogger` when there are more than 1500 sockets in `/tmp`, cf. exception 30120 (`dglogd_high`) in `prod/cluster/gridwms/wms_lemon`. This deletes all `dglogd` sockets in `/tmp`.

LB fill /var/tmp with RTM notification files --

https://ggus.eu/ws/ticket_info.php?ticket=80341

RTM monitoring causes bloat in `/var/tmp` when slow, causes LB notification agent to use a lot of memory and eventually crash (at around 400 notif files, 10 MB each = ~ 4GB)

This is now monitored by Lemon (metric 5430 and exception 30253 in template `wms_lemon`) with automated recovery at 200 files and considering WMS node unhealthy (`wms_refuse_new_job` exception) at 300. See also recovery procedure in `WmsOperationsAndTroubleshooting`.

Nordugrid CEs crashing Condor's Gridmanager continuously -- https://ggus.eu/ws/ticket_info.php?ticket=80215

Two workarounds are set up, both addressing the same issue. The first is a mitigation, the second is a fix:

1. By default there is only one Gridmanager instance, but we tell Condor to use a separate Gridmanager instance "nordugrid_gm" for all nordugrid CEs and another instance "generic_gm" for other CEs. This way not all LCG CEs will be affected by Nordugrid crashing the Gridmanager instance.

We add a `GRIDMANAGER_SELECTION_EXPR` setting in `/opt/condor-c/etc/condor_config` via a yaim post-action for `config_condor_wms` (defined in `glite.tpl`)

Manual action in order to undo when this is fixed: remove `GRIDMANAGER_SELECTION_EXPR` setting from `/opt/condor-c/etc/condor_config`

BUT we may want to keep this even now that a fix is available, as there are regular problems with nordugrid and this will improve overall isolation at little cost. So this might be kept as final configuration, not a workaround.

2. custom RPMs built by Maarten fix the problem, but not yet included in official distribution.

`~maart/public/www/wms-patch/condor-7.4.2-1.lcg.x86_64.rpm` (instead of "condor","7.4.2-1")

`~maart/public/www/wms-patch/glite-wms-helper-3.3.2-1.sl5.x86_64.rpm` (instead of "glite-wms-helper","3.3.2-0.sl5.")

They should get replaced automatically by EMI update process deployment procedure when fixed in EMI

Cron job deletes `/var/proxycache` when no activity -- https://ggus.eu/ws/ticket_info.php?ticket=78295

As a workaround we touch `/var/proxycache/DO_NOT_DELETE_THIS_FOLDER` in the yaim post-action `/opt/glite/yaim/functions/post/config_glite_wms` deployed via filecopy (this post function is used for other things as well)

The file should be deleted if this gets fixed.

Condor 7.4.2 not truncating Matchlog log file

This is a known problem in condor (cf.

<http://comments.gmane.org/gmane.comp.distributed.condor.user/20533>). The file can quickly grow to GB's. More recent versions of condor solve the problem but it would have to be seen with the developers what versions are actually supported (configuration files may change etc.). Version 7.4.2 is the one found in the EMI 3rd-party repository http://emisoft.web.cern.ch/emisoft/dist/EMI/1/sl5/x86_64/third-party/

As a workaround we set up a daily logrotate of the Matchlog file and keep only 1 rotated log (via `ncm-altlogrotate`)

Missing job to clean up LBProxy database entries -- https://ggus.eu/ws/ticket_info.php?ticket=75881

Fixed in EMI2: <https://savannah.cern.ch/bugs/index.php?88630>

Without this the mysql database keeps growing forever.

A cron job to run glite-lb-purge with -x option is added in glite.tpl. We must specify a value for --done otherwise all jobs in the Done state are cleaned up (while not in the "regular" LBServer cleanup = glite-lb-purge without -x)

See: http://egee.cesnet.cz/mediawiki/index.php/LB_purge_and_export_to_JP#glite-lb-purge

Note that jobs cleaned with -x are moved to the zombie_jobs table (querying their status results in "410 Gone")

LCMAPS logging problems -- https://ggus.eu/ws/ticket_info.php?ticket=80065

Several issues in this area:

1. LCMAPS is still logging excessively in /var/log/messages resulting in large messages log files. Supposed to be fixed in Update 11 but it's not (<https://savannah.cern.ch/bugs/?88569>)

As a workaround, in `glite.tpl` we use the `config_glite_wms` post-action to redirect LCMAPS/LCAS logging to `/var/log/glite/[lcaslcmaps]-gridftp.log` and set up `altlogrotate` to rotate the log files (12 weekly rotations i.e. 90 days)

2. We also fix the incorrect log rotation of `var/log/glite/lcmaps.log` reported in the same GGUS ticket with `altlogrotate`

Log file generation/rotation is not compliant with grid policies -- https://ggus.eu/tech/ticket_show.php?ticket=77161

As a workaround we archive logs with TSM, see TSM configuration below. When this is improved we can adjust the TSM configuration accordingly.

But if critical security information missing from /var/log/messages (cf. https://ggus.eu/tech/ticket_show.php?ticket=79938) is finally included, then security team collection of the important data via syslog will allow us to stop using TSM at all.

High memory usage of in glite-wms-workload daemon -- https://ggus.eu/ws/ticket_info.php?ticket=76149

High memory usage in glite-wms-workload daemon. cf. <https://savannah.cern.ch/task/?20342>: this affects WMS nodes dedicated to CMS.

WMS developers recommended (in https://ggus.eu/ws/ticket_info.php?ticket=76149) to install the `jemalloc` alternative `malloc()` to avoid high memory usage. Without it, we end up with the workload manager filling all the memory. But tests concluded that `jemalloc` performs not visibly better than default `malloc()`. On the other hand `libtcmalloc` from `google-perftools` reduces memory usage by 50%. A `google-perftools` build (and dependency `libunwind`) for `slc5/x86_64` was done by Linux support.

Missing job to clean up LBProxy database entries -- https://ggus.eu/ws/ticket_info.php?ticket=75881

We install the google-perftools package and set in /etc/glite-wms/glite_wms.conf (via customized YAIM actions):

```
RuntimeMalloc = "/usr/lib64/libtcmalloc_minimal.so.0";
```

With tcmalloc, the workload_manager process grows to 12.5GB of "RSS" memory usage then gets stable there on CMS nodes. This is OK as we have 24GB of RAM on these nodes and other processes use no more than 5-6GB in total.

As a safeguard we use the lemon monitoring to monitor memory usage of the workload_manager daemon and restart it in case it uses >14GB of memory. For more details see WmsMonitoring.

This topic: LCG > EMIWMSissues

Topic revision: r1 - 2012-07-14 - MaartenLitmaath



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)