Security aspects of the lcg-CE

Maarten Litmaath (CERN)

EGEE'08

How to start/stop the services

- Start the services

```
/etc/init.d/globus-gass-cache-marshal start
/etc/init.d/globus-job-manager-marshal start
/etc/init.d/globus-gatekeeper start
/etc/init.d/globus-gridftp start
/etc/init.d/bdii start
/etc/init.d/gLite start
```

- Stop the services

```
/etc/init.d/globus-gatekeeper stop
/etc/init.d/globus-gridftp stop
/etc/init.d/globus-gass-cache-marshal stop
/etc/init.d/globus-job-manager-marshal stop
/etc/init.d/bdii stop
/etc/init.d/gLite stop
```

Network usage

- globus-gatekeeper
  - Listens on port 2119
- globus-job-manager
  - Started by the gatekeeper
  - Listens in GLOBUS_TCP_PORT_RANGE
  - Connects back to RB/WMS/...
- grid_manager_monitor_agent
  - 1 per DN per RB/WMS/Condor-G submit node
  - globus-url-copy job state summaries back to RB/WMS/Condor-G
- globus-gridftp-server
  - Listens on port 2811 and in GLOBUS_TCP_PORT_RANGE
- glite-lb-logd + glite-lb-interlogd
  - glite-lb-logd listens on port 9002
    - Only WNs should connect
  - glite-lb-interlogd connects to LB servers
- bdii
  - Listens on port 2170

User mapping (1)

- Gatekeeper and gridftp-server have authZ callout
  - /etc/grid-security/gsi-authz.conf
    - /opt/glite/etc/lcas/lcas.db
    - /opt/glite/etc/lcmaps/lcmaps.db

- /opt/glite/etc/lcas/lcas.db → authorization
  - lcas_userban.mod
    - /opt/glite/etc/lcas/ban_users.db
  - lcas_voms.mod
    - VOMS proxy with valid (!) extensions
      - Connection is closed if extensions have expired
    - plain grid proxy

User mapping (2)

- /opt/glite/etc/lcmaps/lcmaps.db → mapping
  - VOMS proxy with valid extensions
    - lcmaps_voms_localgroup.mod
      - /etc/grid-security/groupmapfile → primary and secondary GIDs
    - lcmaps_voms_localaccount.mod
      - /etc/grid-security/grid-mapfile → static account UID
    - lcmaps_voms_poolaccount.mod
      - /etc/grid-security/grid-mapfile → pool account UID
      - /etc/grid-security/gridmapdir → remember mapping
  - Else
    - lcmaps_localaccount.mod
      - /etc/grid-security/grid-mapfile → static account
    - lcmaps_poolaccount.mod
      - /etc/grid-security/grid-mapfile → pool account
      - /etc/grid-security/gridmapdir → remember mapping

User mapping (3)

- /etc/grid-security/groupmapfile
  - derived from YAIM's groups.conf and users.conf
- /etc/grid-security/grid-mapfile
  - Concatenation of 2 parts, needed for LCAS
    - classic grid-mapfile
      - contents updated by edg-mkgridmap cron job
    - /etc/grid-security/voms-grid-mapfile
      - derived from YAIM's groups.conf and users.conf

User mapping (4)

- /etc/grid-security/gridmapdir
  - Should have mode 0770
  - Contains root-owned empty files named after pool accounts
  - The file for an unused account has a hard link count of 1

```
[root@my-CE gridmapdir]# ls -li ops034
2467593 -rw-r--r-- 1 root root 0 Dec 15 2006 ops034
```

  - The file for a used account has a hard link count of 2

```
[root@my-CE gridmapdir]# ls -li ops022
2467581 -rw-r--r-- 2 root root 0 Sep 14 03:57 ops022
```

  - The other link's name encodes the DN and VOMS UNIX groups

```
[root@my-CE gridmapdir]# ls -li | grep ^2467581
2467581 -rw-r--r-- 2 root root 0 Sep 14 03:57
  %2fdc%3dch%2fdc%3dcern%2fcn%3darthur%20dent:ops
2467581 -rw-r--r-- 2 root root 0 Sep 14 03:57 ops022
```

User mapping (5)

- lcg-expiregridmapdir cron job recycles pool accounts as needed
  - When usage exceeds a threshold (80%)
  - Decided per set of accounts having the same prefix
  - The oldest accounts are recycled until the usage falls again below the threshold
  - An account can only be recycled when it has been idle for a certain period (was 2 days, 10 days in latest YAIM)

- This should happen only rarely → create sufficient accounts!
  - /var/log/lcg-expiregridmapdir.log shows current situation

```
VO dtmsgm: inuse / total = 13 / 99 = 0.13, thr = 0.8
VO ops: inuse / total = 24 / 50 = 0.48, thr = 0.8
VO dteam: inuse / total = 76 / 99 = 0.77, thr = 0.8
```

Logs for mapped clients

- "JMA" records in gatekeeper log and /var/log/messages
  - DN, client IP address, RB/WMS job ID, local account, batch system job ID
  - For "lcg" job managers only /var/log/messages has batch system job ID
- /opt/edg/var/gatekeeper/grid-jobmap_* summarize job records
  - Also the user's VOMS attributes
  - Directory is defined in
    - /etc/sysconfig/globus
    - /etc/sysconfig/dgas-add-record.conf
- /var/log/gridftp-session.log
  - DN, client hostname, local account, file transferred
- /var/log/globus-gridftp.log
  - Client IP address, local account, file transferred

Job traces (1)

- /var/log/globus-gatekeeper.log

```
JMA 2008/09/17 20:15:04 GATEKEEPER_JM_ID 2008-09-17.20:15:04.0000024170.0000000000
 has EDG_WL_JOBID 'https://rb113.cern.ch:9000/JgjCc4Vj_tuVj7dIonB3Aw'
JMA 2008/09/17 20:15:12 GATEKEEPER_JM_ID 2008-09-17.20:15:04.0000024170.0000000000
 for /DC=ch/DC=cern/..... on 128.142.173.75
JMA 2008/09/17 20:15:12 GATEKEEPER_JM_ID 2008-09-17.20:15:04.0000024170.0000000000
 mapped to opssgm (8892, 2688)
JMA 2008/09/17 20:15:12 GATEKEEPER_JM_ID 2008-09-17.20:15:04.0000024170.0000000000
 has GRAM_SCRIPT_JOB_ID 1221675312:lcglsf:internal_2807484216:24172.1221675304
 manager type lcglsf
JMA 2008/09/17 20:15:13 GATEKEEPER_JM_ID 2008-09-17.20:15:04.0000024170.0000000000
 JM exiting
```

- /var/log/messages

```
Sep 17 20:15:12 ce111 gridinfo[24172]: JMA 2008/09/17 20:15:12 GATEKEEPER_JM_ID
 2008-09-17.20:15:04.0000024170.0000000000 has GRAM_SCRIPT_JOB_ID
 1221675312:lcglsf:internal_2807484216:24172.1221675304 manager type lcglsf
Sep 17 20:15:48 ce111 gridinfo: [31283-24874] Submitted job
 1221675312:lcglsf:internal_2807484216:24172.1221675304 to batch system lcglsf
 with ID 9427948
Sep 17 20:20:48 ce111 gridinfo: [31283-31283] Job
 1221675312:lcglsf:internal_2807484216:24172.1221675304 (ID 9427948) has finished
```

Job traces (2)

- /opt/edg/var/gatekeeper/grid-jobmap_20080917

"localUser=**8892**"
"userDN=**/DC=ch/DC=cern/.....**"
"userFQAN=/ops/Role=lcgadmin/Capability=NULL"
"userFQAN=/ops/Role=NULL/Capability=NULL"
"jobID=**https://rb113.cern.ch:9000/JgjCc4Vj_tuVj7dIonB3Aw**"
"ceID=ce111.cern.ch:2119/jobmanager-lcglsf-grid_ops"
"lrmsID=**9427948**"
"timestamp=2008-09-17 18:15:48"

- For job submission failures
  - lrmsID may be “FAILED”
  - Local account home directory may contain relevant gram_job_mgr_*.log file which may provide a clue

How to fix errors

- Keep your services up to date w.r.t. gLite updates
- GOC Wiki trouble shooting pages describe most common errors
  - http://goc.grid.sinica.edu.tw/gocwiki/SiteProblemsFollowUpFaq
  - Some entries are out of date
  - Recently updated entries usually fairly correct
- Regional grids also have Wiki pages and mailing lists
- Your ROC will assist you if needed
  - Open a GGUS ticket
- The LCG-Rollout list may be of help

How to ban a user or a VO

- In /opt/glite/etc/lcas/ban_users.db
  - Add a line for each DN that is banned
  - Each line must start with the DN surrounded by double quotes
    - One can simply copy the corresponding line from the grid-mapfile
    - Trailing fields are ignored
  - Nothing needs to be restarted

- To ban a VO reconfigure the service without that VO
  - Will also adapt the information system

How to map suspect account to DN

- /opt/edg/var/gatekeeper/grid-jobmap_*
  - Only for jobs submitted to the batch system
  - Does not catch “fork” jobs running on the lcg-CE itself
    - Used by RB/WMS/Condor-G for “grid_monitor” processes
    - Useful for debugging or small-scale tests via globus-job-run
    - Can be abused...
- /var/log/globus-gatekeeper.log, /var/log/messages
  - “JMA” records
  - Time stamps may disambiguate multiple DNs mapped to the same static account
    - Which DN started this dubious “xyzsgm” process on my CE?
- /etc/grid-security/gridmapdir
  - Only works for pool accounts

Script to map pool account to DN

- http://litmaath.home.cern.ch/litmaath/user-to-dn.pl
- Essence:

```
my $gmd = "/etc/grid-security/gridmapdir";
my $gmf = "/etc/grid-security/grid-mapfile";

open(GMD, "ls -lai $gmd |");

while (<GMD>) {
    my @f = split;

    if ($f[-1] eq $user) {
        (my $dn = $map{$f[0]}) =~ s/%(..)/chr(hex($1))/eg;    # decode...
        $dn =~ s/:.*//;    # remove VOMS UNIX groups
        if ($dn eq "") {
            print "Not in use: $user\n";
            exit 0;
        }
        open(GMF, "$gmf") or die "$gmf: $!\n";
        while (<GMF>) {
            if (/^\s*("$dn")\s/i) {
                print "$1\n";
                exit 0;
            }
        }
        print "Not in grid-mapfile: \"$dn\"\n";    # user probably left the VO
        exit 0;
    }

    $map{$f[0]} = $f[-1];
}

print STDERR "Not found: $user\n";
exit 1;
```

How to handle suspicious jobs

- Pause or stop the batch system queues
- Suspend all active jobs, if the batch system supports it
- Stop gatekeeper and gridftp-server while suspected DNs not yet identified
- Ban suspected DNs or VO
- Keep the active jobs submitted by the suspected accounts suspended if possible, to facilitate forensic investigations
  - Otherwise kill the jobs
- Follow the EGEE Incident Response Procedure
  - http://osct.web.cern.ch/osct/incident-reporting.html

Software manager compromise

- If a software manager (“sgm”) account is suspected
  - Prevent the start of new jobs for that VO
  - Suspend or kill running jobs of that VO
  - The software area for the affected VO must be reinstalled from scratch before new jobs can be accepted for that VO
    - Avoid risk of Trojan horses

Security recommendations

- Use pool accounts instead of static accounts where possible
  - See https://twiki.cern.ch/twiki/bin/view/LCG/SgmPrdPoolAccounts
- Configure enough pool accounts so that recycling occurs rarely
  - Check /var/log/lcg-expiregridmapdir.log
- Ensure that each VO software area is only writable by the software managers for that VO
  - Group-writable only for the “sgm” accounts group, not the VO
- Do not mount the VO software areas on the lcg-CE, but only on the WNs and on the VOBOXes
  - Reduced exposure, reduced risk