SDSC
SAN DIEGO SUPERCOMPUTER CENTER

gx-map, a system for maintaining grid-mapfiles and CRLs

Keith Thompson <kst@sdsc.edu>
San Diego Supercomputer Center

GRIDS Center Community Workshop
June 24, 2005, Chicago, IL

grid-mapfile overview

- The Globus grid-mapfile is a plain text file mapping DNs (GSI distinguished names) to Unix user names.
- The default location is /etc/grid-security/
- Protecting the grid-mapfile from unauthorized updates is critical.
- Keeping it up to date can be tedious and time-consuming.

gx-map

- gx-map allows users to request grid-mapfile updates without administrative intervention. Updates are typically applied within a few minutes.
- Requests can be automatically propagated to multiple systems.
- The actual updates are performed from cron jobs running under a privileged account (typically “globus”).

gx-map

- The system is implemented in about 5000 lines of Perl. It's currently deployed at SDSC and on several TeraGrid systems.
- gx-map 0.4.5 is included in NMI 7.0.
- Why the funny name? An earlier version was part of "globus-extras", a set of auxiliary tools for Globus users at SDSC. gx-map is the sole survivor.

The “gx-request” command

- This is the user’s interface to the gx-map system. It can be run with many confusing command-line options or in interactive mode.
- Formerly called “gx-map”.
- Just type “`gx-request –interactive`” and follow the prompts.
- The result is a plain text request file, written to a world-writable directory.

Sample request file

```
comment            Just testing
dn                 "/O=Earth/CN=Keith Thompson"
email              kst@sdsc.edu
hostname           uffda.sdsc.edu
map_to_name        kst
map_to_uid         500
operation          add
requested_by_name  kst
requested_by_uid   500
timestamp          1119588482 Fri 2005-06-24 04:48:02 UTC
```

gx-map and the SDSC CA

- Another cron job (not part of the gx-map system) checks for new certificates issued by the SDSC CA and automatically invokes the gx-map command.
- Thus a user can obtain a certificate and have the DN added to multiple grid-mapfiles, all without administrative intervention.

The “gx-check-requests” command

- The “gx-check-requests” command is run from a cron job under a privileged account (typically “globus”).
- It checks for new request files generated by gx-request.
- Each new request is validated, annotated, and logged.

Sample annotated request

```
{
        NAMESPACE          KST
        OWNER_NAME         kst
        OWNER_UID          500
        PROCESSED          1119588600 Fri 2005-06-24 04:50:00 UTC
        REQUEST_FILE       1119588482-uffda.sdsc.edu-kst-12259.request
        SOURCE             user
        comment            Just testing
        dn                 "/O=Earth/CN=Keith Thompson"
        email              kst@sdsc.edu
        hostname           uffda.sdsc.edu
        map_to_name        kst
        map_to_uid         500
        operation          add
        requested_by_name  kst
        requested_by_uid   500
        timestamp          1119588482 Fri 2005-06-24 04:48:02 UTC
}
```

The “gx-gen-mapfile” command

- The “gx-gen-mapfile” command is run from a cron job under a privileged account on each host that needs a grid-mapfile.
- If the request log has been updated, it reads it, sorts it by timestamp, and traverses it, generating a new grid-mapfile from scratch.
- Multiple request logs can be read via http or ftp.

Sample cron jobs

```
#
# Every 5 minutes, check for new requests
#
4,9,14,19,24,29,34,39,44,49,54,59 * * * * \
   /usr/local/apps/gx-map-0.4.1/sbin/gx-check-requests \
   -namespace SDSC

#
# Every 5 minutes, update the grid-mapfile (if needed)
#
0,5,10,15,20,25,30,35,40,45,50,55 * * * * \
   /usr/local/apps/gx-map-0.4.1/sbin/gx-gen-mapfile \
        -req default \
        -req ftp://ftp.sdsc.edu/pub/sdsc/globus/software/gx-
   map/sdsc-data-0.4.1/requests.log \
        /usr/local/apps/grid-security/grid-mapfile
# (/etc/grid-security/grid-mapfile is a symlink to
# /usr/local/apps/grid-security/grid-mapfile)
```

Installation

- Unpack the tarball.
- Write a config file.
- Run “./configure-gx-map foo.conf”.
- Run “make install”.
- Sample config file:

```
PERL            /usr/bin/perl
PATH            /bin:/usr/bin
NAMESPACE       SAMPLE
INSTALL_DIR     /INSTALL/DIR/gx-map-0.4.5
DATA_DIR        /DATA/DIR/gx-map-0.4.5-data
REQUESTS_LOG_PERMISSIONS 444
GLOBUS_ADMINS   globus
ADMIN_EMAIL     foo@sample.edu
```

Namespaces

- A gx-map “namespace” is a consistent mapping of Unix user names and numeric UIDs to people.
- The “John Smith” problem: How do I know whether “jsmith@site1” and “jsmith@site2” are the same person?
- I’m currently working on a mechanism to support propagation of information across different namespaces, via an external user database. (This is for TeraGrid, but it should be reasonably extensible.)

Security

- The worst-case scenario: Allowing you to map your DN to my Unix account.
- The gx-request command itself is unprivileged; anyone can easily create a fake request file.
- The gx-check-requests command validates the ownership of the request file. Some systems allow non-root chown, which would break the security model; gx-check-requests now detects this and refuses to run.

Security, Security, Security

- gx-map is a security-critical application.
- The author is not a security expert.
- Does this make you nervous? Good!
- gx-map has no known security bugs.
- Equivalently (and perhaps more accurately), all the security bugs are unknown ones.
- I think it’s fairly robust, but there are no guarantees.
- If you install it and it breaks your system, it’s *your* fault for trusting me. 8-)} (sort of)

Paranoid mode

- The command-line arguments to gx-gen-mapfile allow you to specify the location of the grid-mapfile. This doesn’t have to be “/etc/grid-security/grid-mapfile”.
- If you don’t quite trust gx-map, you can have it update a separate file; periodically, you can examine the separate file and manually copy it to /etc/grid-security if it looks ok.
- When/if you’ve decided to trust gx-map, you can modify the cron job so it writes directly to /etc/grid-security/grid-mapfile (or you can make /etc/grid-security/grid-mapfile a symlink).

Numeric UIDs?

- We assume that both user names and numeric UIDs are consistent within a namespace (typically a site or organization).
- Q: Why worry about UIDs? They don’t appear in the grid-mapfile.
- A: The system on which gx-check-requests runs may not have all user accounts in /etc/passwd. In this case, gx-check-requests records the UID; it doesn’t know the user name.

Numeric UIDs? (cont.)

- This is workable but ugly. Possible alternatives:
  - Assume/require that gx-check-requests runs on a system with all accounts, or make UID dependence configurable at installation time.
  - If a user doesn’t have an account on the system running gx-check-requests, require administrative intervention.
  - Get username/UID information from somewhere other than /etc/passwd (system-specific).

User interface

- The first version of gx-map had only a command-line interface, with a dozen or so options. It all seemed perfectly clear to me (there’s even a “-help” option) until I let someone else use it.
- The command-line interface is too complex, especially for a tool that most users will run only once.
- The command-line interface is still supported (mostly for use by automated tools), but the main user interface is now interactive, prompting the user for each required piece of information.

Command-line options

(See, I told you they were confusing)

```
% gx-request -long-help
Usage: gx-request [options]
Option names may be abbreviated.
    -help                   : Show a brief usage message and exit.
    -version                : Show version information and exit.
    -interactive            : Run interactively (recommended).
    -long-help              : Show this long usage message (recommended
                              only for Globus administrators and masochists).
    -add                    : Add the specified mapping.
    -remove                 : Remove the specified mapping.
    -remove-dn              : Remove all mappings for the specified
                              distinguished name.  For use only by Globus
                              administrators.
    -remove-user            : Remove all mappings for the specified user.
    -update                 : Request an update of all grid-mapfiles.
                              This normally isn't necessary, but it can be
                              useful if you already have a certificate and
                              get a new account on a machine.
Note: Exactly one of "-interactive", "-add", "-remove", "-remove-dn",
      "-remove-user", and "-update"
      (or "-help", "-usage", or -long-usage) must be specified.
------------------------------------------------------------------------
    -quiet                  : Work silently.
                              Implies -force.
    -force                  : Apply mapping without prompting.
                              Default is to ask for verification before
                              proceeding.
    -no-admin               : Assume the user is not a Globus administrator.
                              Intended for testing only; has no effect if
                              you're not already a Globus administrator.
    -dn "string"            : Distinguished name.
                              Default is extracted from ~/.globus/usercert.pem
    -certificate-file file  : Name of file from which to extract DN.
                              If neither "-dn" nor "-certificate-file" is
                              specified, extract DN from
                              $HOME/.globus/usercert.pem
    -force-dn               : Normally, gx-map (minimally) checks the DN for
                              proper syntax; this option overrides that check.
    -username name          : Unix user name to map.
                              This option is for use by Globus
                              administrators only.
    -secondary              : Request a secondary mapping.
                              See the documentation (not yet written) for
                              details.
    -directory dir          : Specify an alternate data directory.
                              This option is for use by Globus administrators
                              only.
                              The default data directory is
                              /usr/local/apps/gx-map-0.4.5/var .
    -email addr             : Your contact e-mail address (optional).
                              This may be used to contact you if there's
                              a problem with your certificate.
    -no-email               : Ignored (provided for compatibility with 0.3).
    -comment "string"       : Comment to be added to request log (optional)
    -source string          : Specify the source of the mapping.
                              Argument may consist only of letters, digits,
                              underscore, period, and hyphen characters
                              ([A-Za-z0-9_.-]).
                              This option is for use by Globus administrators
                              only.
    -debugging              : Enable debugging output.
Note: If this help message has scrolled off the top of your screen, try
    gx-request -long-help | less
```

User interface (cont.)

- GUI? No.
- Web interface? No.
- Two reasons:
  1. I haven't had much practice implementing GUIs or web interfaces.
  2. I don't know how to integrate the gx-map security model into a fancy interface.
- gx-map has been tested only on Unix-like systems; it should run on anything that supports Perl and cron.

Levels of complexity

- The simplest case is a single system.
- The next level is a set of systems sharing a common filesystem.
- If several systems share the same account namespace but have no shared filesystem, information can be propagated by ftp or http (SDSC does this).
- Systems without a shared account namespace (i.e., a grid) are more of a challenge.

Multiple mappings

- The grid-mapfile format allows multiple user names per DN. (Some but not all Globus tools can use this.)
- **`"/O=Foobar/CN=John Smith" user1,user2`**
- This is supported via “secondary” mappings, but it’s clumsy; I’ll probably simplify the feature. (So far it’s been used accidentally more often than deliberately.)

Mapping somebody else’s DN

- gx-map doesn’t let you map to somebody else’s account.
- It does let you map somebody else’s DN to *your* account.
- If you want to give somebody else access to your account, that’s your problem.
- It might be possible to require a valid proxy before mapping a DN, but I haven’t done this.

gx-ca-update

- The gx-ca-update tool installs and maintains CA certificates, signing_policy files, and CRLs (Certificate Revocation Lists).
- *.cadesc files distributed with gx-map (more than 80 of them) describe the attributes of certificate authorities.
- A cron job specifies a list of CAs to be accepted. CRLs are automatically downloaded and installed as needed.
- If a CRL expires or becomes unavailable, warnings are sent by e-mail (mostly to me).

gx-ca-update (cont.)

Sample *.cadesc file:

```
# $Id: 3deda549.sdsc.cadesc,v 1.13 2005/04/17 09:43:23 kst Exp $
# $Source: /projects/globus/kst/CVS/tools/gx-map/ca/3deda549.sdsc.cadesc,v $

CA_NAME             SDSC CA
HOMEPAGE            http://www.sdsc.edu/CA/
CONTACT             Bill Link <bill@sdsc.edu>
HASH                3deda549
SUBJECT             /C=US/O=SDSC/OU=SDSC-CA/CN=Certificate Authority/UID=certman
MAY_SIGN            /C=US/O=SDSC/*
CERTIFICATE_MD5     07:83:1A:81:1F:2C:DD:AF:1E:BD:92:03:B5:F8:A9:C9
CERTIFICATE_SHA1    D9:90:69:8B:BE:C7:85:65:8A:EE:2D:3C:9E:F2:E2:6A:CE:C1:7D:F4
CERTIFICATE_URL     http://www.sdsc.edu/CA/3deda549.0
SIGNING_POLICY_URL  http://www.sdsc.edu/CA/3deda549.signing_policy
CRL_URL             http://www.sdsc.edu/CA/3deda549.r0
INDEX               http://www.sdsc.edu/CA/ca.db.index
CERT_EXPIRES        Sep  9 02:42:29 2014 GMT

# MD5 and SHA1 fingerprints from 3deda549.0 downloaded 2005-03-29
```

gx-ca-update (cont.)

- Sample cron job:

```
1,31 * * * *   /usr/local/apps/gx-map-0.4.5/sbin/gx-ca-update \
                 -target-dir /usr/local/apps/grid-security/certificates \
                 -ca 01621954 -ca 1c3f2ca8 -ca 2ca73e82 -ca 34a5e0db \
                 -ca 3deda549 -ca 4a6cd8b1 -ca 67e8acfa -ca 85ca9edc \
                 -ca 95009ddc -ca 9a1da9f9 -ca aa99c057 -ca b89793e4 \
                 -ca d1b603c3 -ca fa9c3452
```

Bugs

- Yes, there are bugs. (It’s software, after all.)
- None of the *known* bugs create security holes; at worst, there might be a denial of service.
- Design principle: When in doubt, bail out.
- Recent releases include a “bugs” directory (25 entries as of 0.4.5, 31 today).
- Most are fixed; the rest are either suggested enhancements or very minor.

Availability

- The gx-map home page is http://users.sdsc.edu/~kst/gx-map/.
- SDSC and several TeraGrid sites are running 0.4.1. The latest release is 0.4.5 (in NMI 7.0).
- Any questions: contact me, Keith Thompson, <kst@sdsc.edu>.
- If you find a security hole, *please* let me know ASAP.
- Released as open source under a BSD-like license.