

INSTALLING THE NGS VDT-BASED SOFTWARE STACK

Revision: INSTALL.txt,v 1.6 2011/01/02 01:42:44 jason-lander Exp

Overview

The impatient can go straight to the Installation Guide below.

Background

Users and hosts of the NGS authenticate using X509 certificates.

On a compute service, these certificates are used to associate the user with one of a set of *pool* accounts under which any programs they wish to use will run.

There can be several sets of pool accounts in use. Additional information - in the form of Virtual Organisation (VO) attributes - attached to the certificate is used to associate the user with one set.

The NGS uses Globus middleware to provide access (<http://www.globus.org>). The VO support is provided by an extension to Globus called LCAS/LCMAPS.

LCAS/LCMAPS will pass the users details through a series of plugins until a final mapping to a username and group are made.

The plugins can use VO information or fall back to reading entries in `/etc/grid-security/grid-mapfile`.

The mapping between users' certificates and their pool accounts is recorded in a directory called `/etc/grid-security/gridmapdir`. This contains empty files for every pool account which are *hard-linked* to files representing the users as they are assigned.

Installing a suitably-configured set of software involves:

1. Obtaining and installing a key and certificate for your system.
2. Installing a base Globus installation.

This document uses a precompiled version provided by the version 1.8.1 of the Virtual Data Toolkit. (<http://vdt.cs.wisc.edu/>) which is itself installed via a tool called Pacman.

3. Building and installing LCAS/LCMAPS.

This is built from source code and requires a suitable software development environment.

4. Configuring LCAS/LCMAPS and creating a suitable set of pool accounts.
5. Enabling LCAS/LCMAPS support within Globus.

Where possible, scripts have been provided to do the installation and configuration.

Manifest

The following files and directories are required, The latest versions can be found at <http://sourceforge.net/projects/ukngi>

INSTALL.txt

This installation guide

ngs-prereq-check.sh

Script to check if system is capable of building and installing the software

ngs-install-host-cert.sh, ngs-pacman-installer.sh, ngs-vdt-installer.sh

Scripts to install and configure Globus from the VDT distribution

ngs-glite-build.sh, ngs-glite-patches

Script and patches required to download and compile LCAS/LCMAPS

`configure_lcas_lcmaps_ngs, ngs-voms-configure`

Scripts to create and configure a suitable environment

Installation Guide

Preparation

Set convenience variables:

```
Pacman=/usr/local/pacman  
Vdt=/usr/local/VDT
```

Run:

```
./ngs-prereq-check.sh
```

to check if all the prerequisites mentioned in this document are present.

Host certificates

Before deploying the software, you should to obtain a host key and certificate signed by the UK eScience certification authority (<http://ca.grid-support.ac.uk>) and install these as:

```
``/etc/grid-security/hostkey.pem``  
``/etc/grid-security/hostcert.pem``
```

The key should be readable only by root.

The key and certificate will usually be exported from a web-browser within a single PKCS12-format file. If you have such a file, the individual `hostkey.pem` and `hostcert.pem` files can be installed by running the following as root:

```
mkdir /etc/grid-security  
./ngs-install-host-cert.sh filename.p12 /etc/grid-security
```

Note

It is possible to install a host certificate *after* the software has been installed but the Globus gatekeeper must be reconfigured by running `$GLOBUS_LOCATION/setup/globus/setup-globus-gatekeeper`

Pacman

You will need *one* of

- curl
- wget
- lwp-request
- lynx

to install Pacman using the installer script.

As root, run:

```
mkdir $Pacman
./ngs-pacman-installer.sh $Pacman
```

Base VDT

Prerequisites:

- perl
- patch
- gcc (to build Globus-Core)

You will see warnings if 'wget' is missing.

As root:

```
. $Pacman/setup.sh

mkdir $Vdt

/usr/sbin/adduser -r globus           # Or local equivalent
/usr/sbin/addgroup -r gridmapdir     # Or local equivalent

./ngs-vdt-installer.sh              \
  --site-name=my-ngs-site           \
  --contact-email=my-contact-email@address \
  --city=Leeds                       \
  --latitude=53.809812               \
  --longitude=-1.55699               \
  $Vdt
```

Other supported options:

```
--accept-licenses           # Do not prompt user to accept
                             # license terms.
--pbs-path=<path>           # Enable PBS support
--lsf-path=<path>           # Enable LSF support
--sge-path=<path>           # Enable SGE support
--condor-path=path         # Enable Condor support
--condor-config=path       # Path to Condor configuration file
--gridmapdir-group=<name>  # Set name of group which owns the gridmapdir
                             # directory
```

The script can take a very long time to complete. If you wish to check on progress, you should look at the files

```
$Vdt/o..pacman..o/logs/pacman.log $Vdt/vdt-install.log
```

which are updated as the installation progresses.

The ngs-vdt-install.sh script can be *rerun* - in which case it will attempt to reconfigure the system to match the new settings.

LCAS/LCMAPS Installation

Prerequisites:

- cvs
- automake
- autoconf
- libtoolize
- make
- yacc
- flex
- gcc
- g++

NOTE: if g++ is missing you will see a cryptic messages about /lib/cpp failing a sanity check.

The LCAS/LCMAPS build script downloads and compiles code and installs it at a user specified place.

We recommend that you do a test install into a empty directory, before rerunning it to install within the directory structure provided by VDT.

As a non-root user with sudoers rights:

```
Flavour=gcc64dbg          # or gcc32dbg on 32-bit systems
Testinst=/tmp/test-install # Test installation tree

. $Vdt/setup.sh

./ngs-glite-build.sh \
  --prefix=$Testinst \
  --with-glite-location=$GLITE_LOCATION \
  --with-edg-location=$EDG_LOCATION \
  --with-globus-location=$GLOBUS_LOCATION \
  --with-globus-flavour=$Flavour \
  --with-globus-thr-flavour=${Flavour}pthr \
  --install-gt4-callout \
  --disable-ldap-requirements \
  --voms-location=$GLITE_LOCATION \
  --install-lcmaps-accounting \
  --install-lcmaps-accounting-edg-tools \
  --with-expat-location=$VDT_LOCATION/expat
```

To do the 'final' install, repeat but using the:

```
--with-sudo-install
--prefix=$GLITE_LOCATION
```

options.

The builder script creates and uses a directory called:

```
ngs-glite-build
```

as a work space.

Each of the gLite packages required is compiled and installed separately. A logs for each build can be found within ngs-glite-build/log.

LCAS/LCMAPS configuration

As root:

```
. $Vdt/setup.sh
perl ./configure_lcas_lcmaps_ngs
```

This creates a skeleton LCAS/LCMAPS configuration - with all the user editable files within `lcas` and `lcmaps` directories within `/etc/grid-security`.

Copy:

```
./ngs-voms-configure
```

to `$GLITE_LOCATION/bin/`, make executable and run via cron using a wrapper like:

```
./path-to-vdt/setup.sh
ngs-voms-configure
```

This script keeps VOMS information up-to-date.

After it has been configured, enable LCAS/LCMAPS support with:

```
ln -s ./gsi-authz-lcas-lcmaps.conf /etc/grid-security/gsi-authz.conf
```

The `/etc/grid-security/gsi-authz-lcas-lcmaps.conf` file should have been created by `configure_lcas_lcmaps_ngs`.

Pool accounts

As root:

1. Create pool accounts. These will consist of a fixed prefix followed by a number of digits
By default NGS users are assigned `ngs<digits>` accounts.
2. Create an entry for each account under `/etc/grid-security/gridmapdir` using a series of command such as:

```
touch /etc/grid-security/gridmapdir/ngs0001
```

The `edg-mkgridpool` script can automate this step, but be aware that it assumes *any* username containing a number is a pool account.

3. For accounts which depend on VO credentials, add entries to `/etc/grid-security/lcmaps/gridmapfile` of the form:

```
"/monitoring.ngs.ac.uk/lcas_lcmaps/*" .ngsmon
```

and entries in `/etc/grid-security/lcmaps/groupmapfile` of the form:

```
"/monitoring.ngs.ac.uk/lcas_lcmaps/*" ngsmon
```

These will map accounts from the `monitoring.ngs.ac.uk` to a pool accounts started in `ngsmon` under the `ngsmon` group.

The ngs-voms-configure can automate this stage on some systems if you run it as:

```
ngs-voms-configure -setupaccounts
```

This assumes that the `useradd` command is sufficient to create a user and that you wish to support *all* current NGS VOs.

Understanding the Installation Scripts

This section is for those who care what the various installation scripts do. It is not needed to do an installation.

ngs-pacman-installer.sh

1. Downloads the latest version of Pacman from http://physics.bu.edu/pacman/sample_cache/tarballs
2. Cleans up...

ngs-vdt-installer.sh

1. Configures Pacman to download packages from http://vdt.cs.wisc.edu/vdt_181_cache
2. Sets environment variables to automatically answer questions for which the user otherwise be prompted.
3. Downloads the following packages and their dependencies
 - Globus
 - EDG-Make-Gridmap
 - GSIOpenSSH
 - MyProxy
 - OpenLDAP
 - GLUE-Schema
 - Generic-Information-Provider
 - KX509
 - VOMS-Client
 - Globus-Base-SDK
 - CA-Certificates-Updater

plus the Globus-*⁻Setup packages for the PBS, LSF or SGE batch systems

5. Unpacks new files embedded in the script.
6. Patches the `configure_gip` and `osg-info-dynamic-pbs` scripts
7. Downloads a version of the 1.3 Glue Schema that works with the Grid information service
8. Adds a new GIP plugin `ngs-uee-gip-plugin` that reports on the contents of the `/usr/ngs/` directory
9. Enables all services. GSISsh is enabled outside the VDT install mechanisms.
10. Populates the `grid-mapfile` using `edg-mkgridmap`
11. Reconfigures the `vdt-update-certs` service to download from the NGS's certificate repository

ngs-glite-build.sh

1. Downloads glite packages from the CVS repository
2. Applies patches
3. Downloads the NGS's own accounting plugin

4. Compiles and installs