

CaGrid Installer User's Guide

The caGrid 1.1 installer provides a graphical, wizard-like interface for installing caGrid dependencies, source, services, and applications. It is a complete overhaul of the caGrid 1.0 installer. New features include:

- Small download: < 6MB
- Component installers: install all caGrid services and applications
- Designed for re-use: can be used to re-install/re-configure previous installations

Downloading

The installer can be downloaded from http://gforge.nci.nih.gov/frs/?group_id=25 as [caGrid-1.1-installer.zip](#)

Save the installer zip file to your file system. Unzip the file into a directory of your choosing, for example /opt/cagrid/installer. The contents of that directory should look like this:

```
caGrid-1.1-installer.jar
caGrid-1.1-installer.zip
lib/
  caGrid-1.1-wizard.jar
  commons-logging.jar
  jce-jdk13-125.jar
  log4j-1.2.14.jar
  mysql-connector-java-3.0.16-ga-bin.jar
  tools.jar
  xmltask-v1.14.jar
scripts/
  build.xml
resources/
  spring-beans.dtd
  web-app_2_3.dtd
  web-facesconfig_1_1.dtd
```

Running

caGrid requires the Java 1.5 JDK. If the installer cannot locate this version of Java, an error message will be displayed when you try to run the installer. If you see this error message, read the [Environment Settings](#) section below to learn how to configure your environment on various platforms. You will also need to be connected to the Internet because the installer will download caGrid and its dependencies. If you are behind a firewall or need to install to hosts that have no route to the Internet, read the [Off-Line Installation](#) section below.

You will find a Java archive (Jar) file named caGrid-1.1-installer.jar in the directory where you unzipped the downloaded zip file. This is the installer program. On some operating systems (e.g. Microsoft Windows, Macintosh), you can double-click on this file.

NOTE: Some programs that unzip zip files will display the contents of the zip file in a folder-like view (e.g. Microsoft Windows XP and Vista). Be sure to first extract the zip file before double-clicking on the installer Jar file. Otherwise, the installer will fail to start.

On other systems (e.g. Linux, Unix), you must run this program from the command line. To run this program from the command line, change to the directory where you unzipped the the downloaded file. Then type the following:

```
java -jar caGrid-1.1-installer.jar
```

NOTE: Make sure you have changed to the same directory as the Jar file before running the installer.

When the installer starts, you will see a splash screen which displays the caGrid logo.

Usage

The installer has been designed for the following usage scenarios:

1. Installing the caGrid source distribution and dependencies.
2. Configuring a container into which caGrid services can be deployed.
3. Installing any of the caGrid services.
4. Installing caGrid applications (portal and browser).

If you are developing your own services that you wish to expose to an existing grid, then you will be interested in only scenarios 1 and 2. If you are creating your own grid or you would like to integrate your institution's security domain into other grids, then you would also be interested in scenario 3. If you would like to be able to use a web interface to discover and view metadata that is exposed by services which do not advertise to the NCICB's index services, then you would be interested in scenario 4.

Scenarios 1 and 2 require essentially no knowledge of how caGrid works. In a matter of minutes, you should be able to download and install caGrid, configure your container, and start using Introduce (and other core caGrid tools) to develop and deploy services. (Refer to the Introduce guide to get started developing services:

<http://www.cagrid.org/mwiki/index.php?title=Introduce>)

Scenarios 3 and 4 require knowledge of how caGrid and each of its components work - both on their own and together. You should refer to the [caGrid User's Guide](#) for each component. You can also refer to the deployment documentation that describes how the installer is used to deploy the NCICB production grid: http://www.cagrid.org/mwiki/index.php?title=CaGrid:How-To:Deploy_caGrid_1.1

The rest of this section will describe usage scenarios 1 and 2.

Installation Phases

Each run of the installer goes through these phases:

1. Select installation type and desired components
2. Download/install required components (optional)
3. Configure caGrid and container (optional)
4. Configure components
5. Install components

Installation Types

The first screen displayed by the installer (after the license screen) prompts the user to select one of the following installation types:

1. Install caGrid
2. Configure Container
3. Install caGrid Services
4. Install caGrid Portal
5. Install caGrid Browser

If you are interested in usage scenarios 1 and 2 (described above), then you will need to select only *Install caGrid* and/or *Configure Container*. If you are interested in usage scenarios 3 and 4, then you would select one or more of the remaining options.

NOTE: The installer will try to do the "right" thing, based on your selections. If you select only *Install caGrid Services*, but you haven't yet installed caGrid or configured a container, the installer will detect this and prompt you through the process of establishing a working system.

Install caGrid

If you select *Install caGrid* and de-select all other options, the installer will do the following:

1. Download and extract Ant
2. Download and extract Globus
3. Download and extract caGrid
4. Compile caGrid
5. Configure caGrid for a target grid (e.g. NCICB Production Grid)

If your ANT_HOME environment variable already points to the correct version of Ant, the installer will give you the option to use this, or reinstall Ant to the location you choose. Similarly, if your GLOBUS_LOCATION environment variable points to the right version of Globus, you'll be given the same options. If you have already used the installer to install caGrid 1.1, you'll be

given the option re-install it as well. You'll also be given the option to reconfigure it for different target grid.

Configure Container

caGrid services can be deployed to either the Globus container or to the Globus web application that has been deployed to a Jakarta Tomcat servlet container. When you select the *Configure Container* option, you'll be prompted to select the container type: *Globus* or *Tomcat*.

If you select the Globus container, the installer will:

1. (Re-)Install Ant (optional)
2. (Re-)Install Globus (optional)
3. (Re-)Install caGrid (optional)

If you select the Tomcat container, the installer will also give you the option to (re-)install Tomcat.

Once the installer has installed (or identified) the appropriate container, it will ask you a series of questions to determine how the container should be configured.

Secure Container

The first question will ask whether the container should be secure. That means, do you want the container to be able to host secure services? Secure services are those that require clients to use one of the Globus Security Infrastructure (GSI) authentication mechanisms. The authentication mechanisms require that the container be provisioned with grid credentials (i.e. a X.509 certificate and key). Grid credentials that are used for a container that is hosting grid services are referred to as *host credentials*. The GAARDS infrastructure (caGrid's security infrastructure) provides mechanisms to obtain host credentials.

In caGrid 1.1, the primary way to obtain host credentials is to request them from the Dorian CA via the GAARDS security UI. That process is described [here](#). The NCICB Production grid includes a Dorian "Level 1" service that will automatically grant host certificates to any grid user, for any host. If you use this approach, your services will be trusted by any other caGrid services that trust the Dorian level 1 CA.

On the other hand, if you just want to experiment with secure grid services, and are not concerned about integrating with an existing trust fabric, you can use the installer to generate both your own CA and host certificates. (The installer just uses the tools found in the caGrid *gridca* project to generate these certificates.)

- If you choose to install a secure container, the installer will then ask for the host name of the machine. (You can't just make this up. The host name must be a DNS name that resolves to the same IP address as the IP address(es) to which the container is listening on the specified port.). Enter the host name and press *Next*.

- If you have chosen the Tomcat container, the next screen will prompt you for the Shutdown Port and HTTPS port. If you are installing multiple containers on the same machine, ensure that these are different than the ports that other containers are using.
 - **NOTE:** The installer will attempt to remove the existing HTTP Connector element in Tomcat's server.xml. But it will not remove any other Connector elements. If you are running multiple containers on the same machine, you will have to manually modify/remove other Connectors (e.g. the AJP/1.3 Connector) if there are port conflicts.
- If you have chosen the Globus container, you will not be prompted for port information. The port that the Globus container listens to is specified as a command line argument to the container's startup script. See [here](#) for further information.
- Next the installer will ask if the service certificate and key (i.e. the host credentials) are present.
 - If you have these, select the *Yes* check box, and click *Next*. On the next screen, you'll be prompted for the location of these credentials.
 - If you don't have host credentials, do **NOT** select the *Yes* check box. Then press *Next*.
 - The installer will then ask if you have a CA certificate and key. These are what the installer will need in order to generate host credentials for you.
 - If you have these, select the *Yes* check box. Then press *Next*. On the following screen, you'll be prompted for the locations of the CA certificate and key as well as the password for the key. Once you've entered these, press *Next*. You'll then be prompted for the location to write the host certificates.
 - If you don't have a CA certificate and key, do **NOT** select the *Yes* check box on the screen that asks if they are present. Then press *Next*.
- You'll then be prompted for locations to write the certificate and key as well as the CA password, the distinguished name (DN) to be used, and the number of days that the CA certificate will be valid. Once you've entered that information, press *Next*. The following screen will prompt you for the location to write the host certificates.

Non-secure Container

If you have not chosen to configure a secure container, the installer will just ask you for the host name and, if using Tomcat, the shutdown and HTTP ports.

Install caGrid services

If you choose the *Install caGrid Services* installation type, the installer will present the list of caGrid services that can be installed:

1. My Introduce Service
2. Dorian
3. GTS

4. SyncGTS
5. AuthenticationService
6. GridGrouper
7. Index Service
8. GME
9. EVS
10. caDSR
11. FQP
12. Workflow

For details on the configuration of each of these components, see the [caGrid User's Guide](#). You could also look [here](#) to see how these services have been deployed using the installer at NCICB.

The *My Introduce Service* option allows you to configure the standard properties of any Introduce-created service and deploy it.

If you choose to install the Workflow service, ActiveBPEL will be downloaded and installed as well, in the same manner as other dependencies such as Tomcat or Ant.

Install caGrid Portal and Browser

See [here](#) for details on install the caGrid portal and browser.

Step-by-Step: Install Secure Container

This section will show you, step-by-step, how to create a secure container using the caGrid Installer. The overall process invokes the following basic steps:

- Download, extract, run the installer to install just caGrid.
- Use the caGrid security UI to request host credentials.
- Run the installer again to install the container and point it at the credentials created in the previous step.

First, [download](#) and [run](#) the installer. After agreeing to the license, select the *Install caGrid* option and de-select all other options.

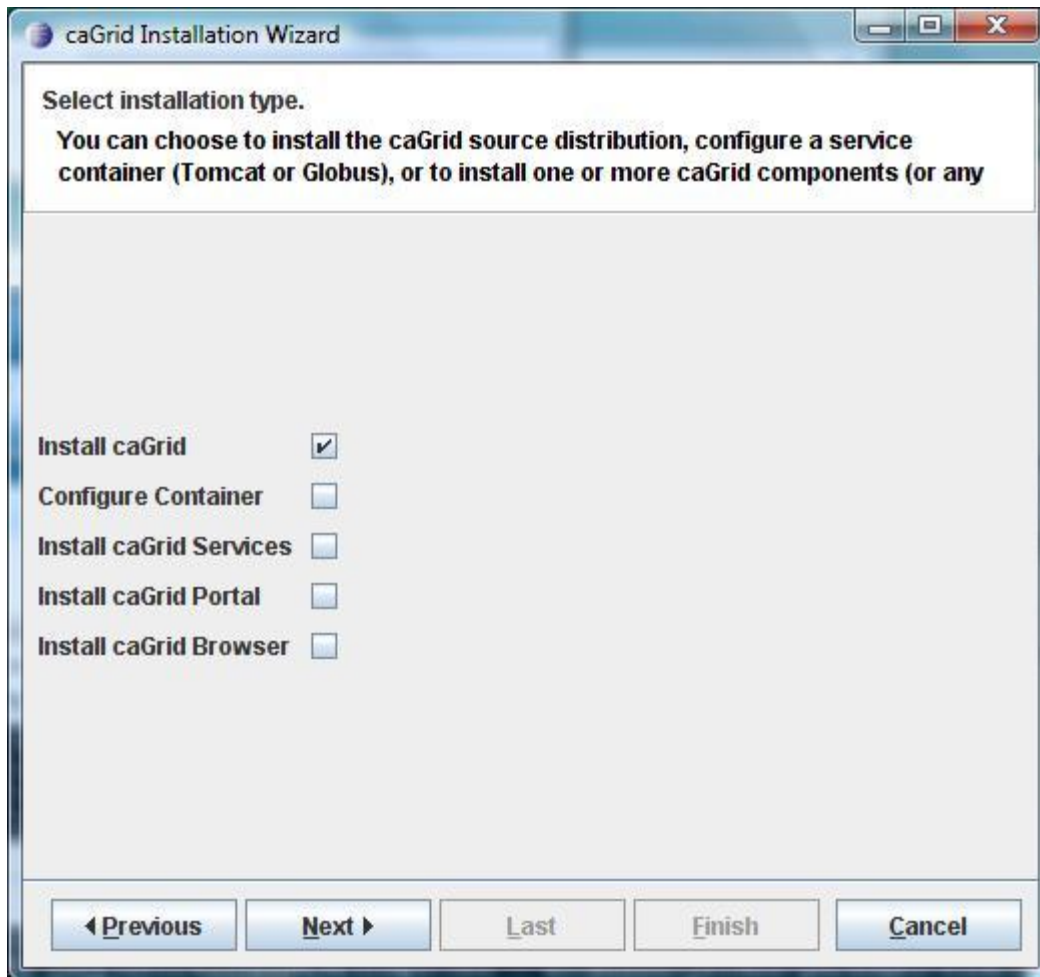


Figure 1: Select caGrid

The press *Next*. In the following screens, indicate where Ant, Globus and caGrid should be installed. On the *Select Target Grid* screen, select the grid that you intend to integrate with. Then press *Next*.

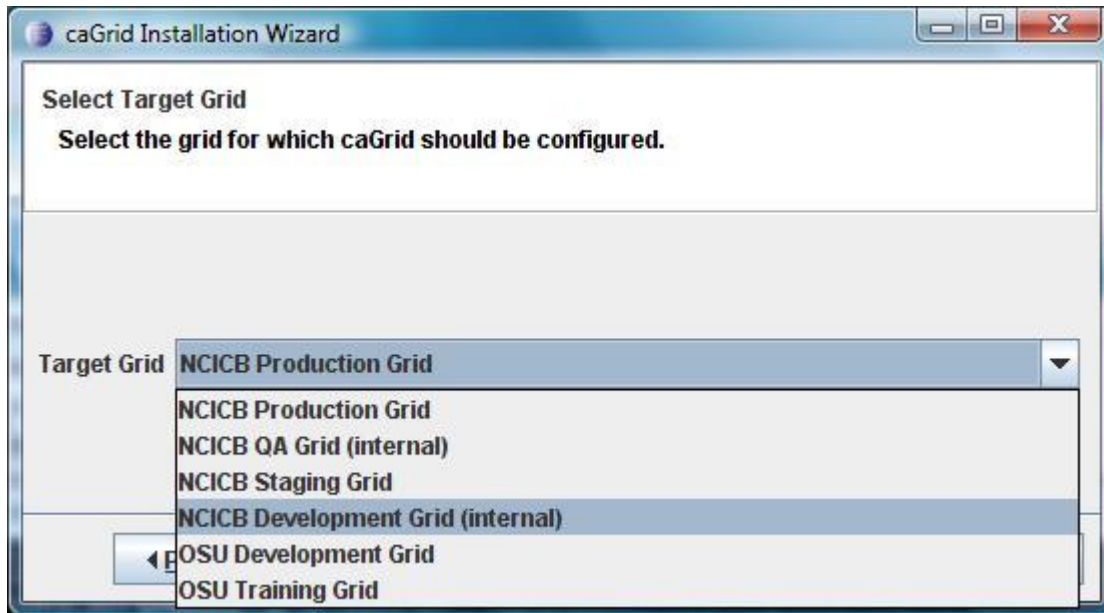


Figure 2: Select target grid

On the *Install Components* screen, press the *Start* button and wait while the installer downloads and extracts the required components. Once this step is finished, press the *Next* button.

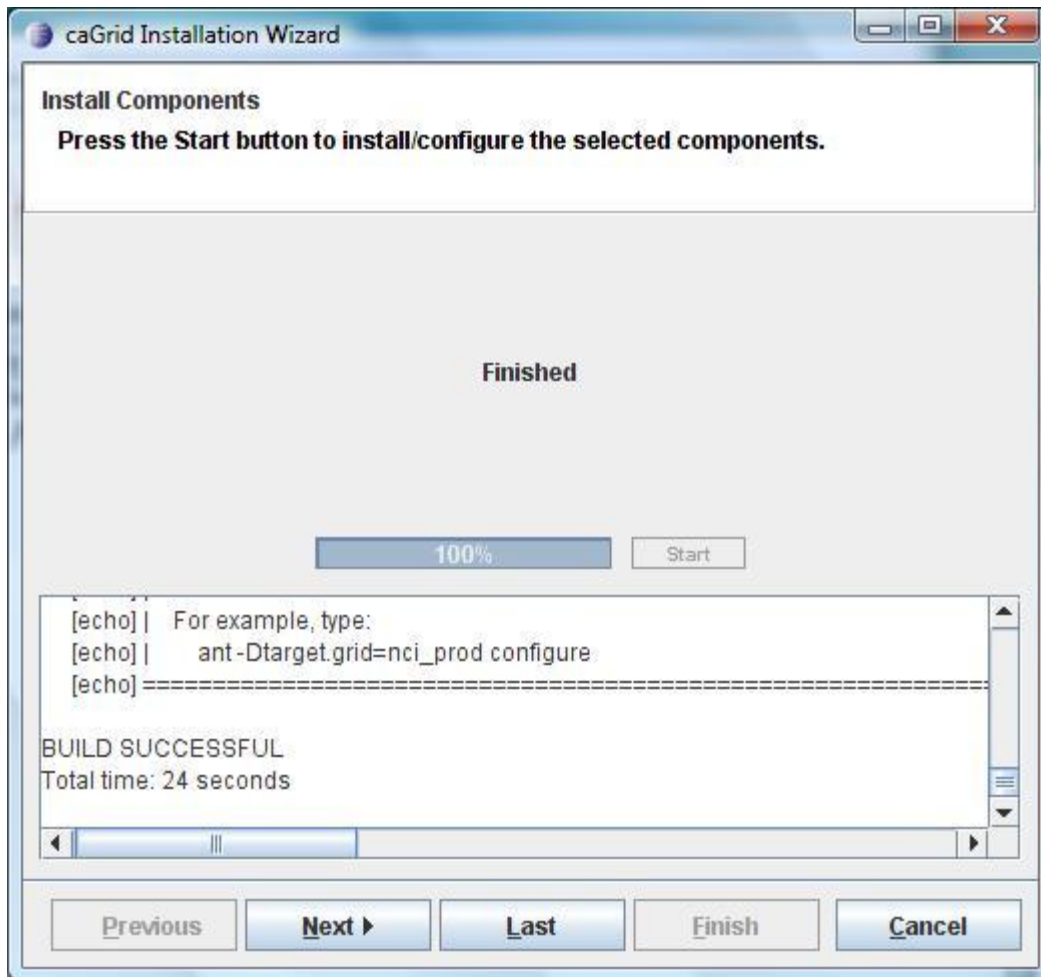


Figure 3: Required components installed

The next panel will indicate that caGrid has been successfully installed. It will also list the environment variables that should be set.




Figure 4: caGrid installation complete

Press *Finish*, then *Close*.

The ANT_HOME and GLOBUS_LOCATION variables should now be set as the installer indicated. If you haven't already set the JAVA_HOME variable, the installer will also indicate what that should be set to. See [here](#) for instructions on setting environment variables on various operating systems.

Open a command prompt and change directories to where caGrid has been installed. Then type:
ant security

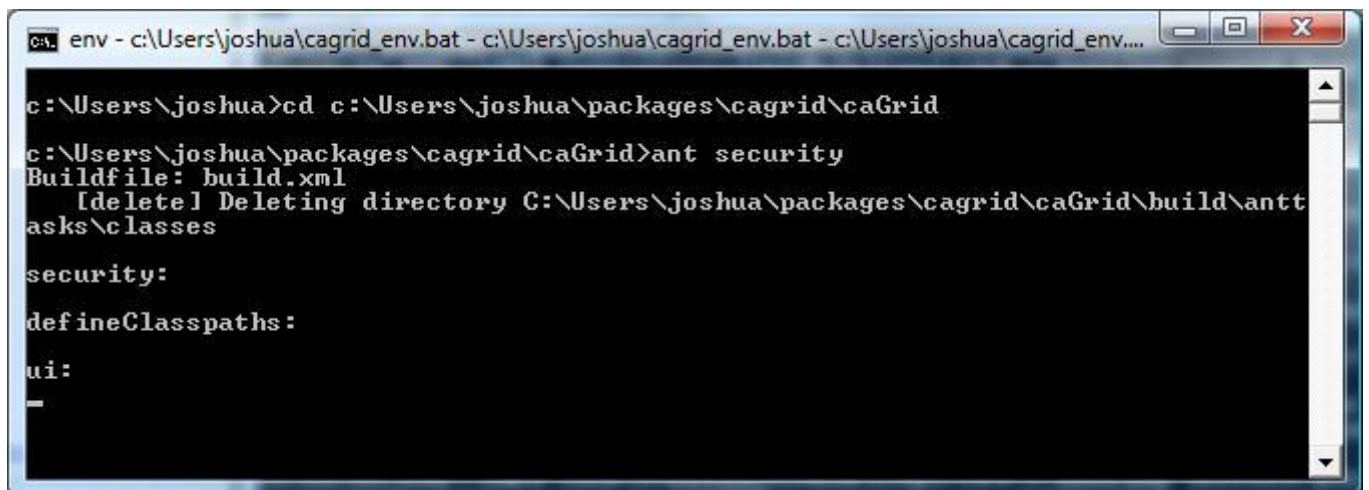


Figure 5: Start up the caGrid security UI

You will need to authenticate to the grid, which requires an account. Anyone can get an account automatically using the caGrid Dorian Level 1 service. If your institution already has an

AuthenticationService that is integrated with the NCICB Production grid, you will be able to use the same username and password that you use at your institution. See [here](#) for directions on requesting an account. See [here](#) for directions on authenticating through the caGrid security UI. Once you have successfully authenticated, follow [these](#) instructions for requesting host credentials.

Once you have obtained host certificates, [start up](#) the installer again.

1. Agree to the license; then select *Configure Container* and de-select all other options. Press *Next*.
2. Select the *Tomcat* container type. Press *Next*.
3. Do **NOT** select the *Yes* check box on the *Reinstall Ant?* screen. Press *Next*.
4. On the *Tomcat* screen, indicate where Tomcat should be installed. Press *Next*.
5. Press *Next*. Indicate where Tomcat should be installed. Press *Next*.
6. Do **NOT** select the *Yes* check box on the *Reinstall Globus?* screen. Press *Next*.
7. Do **NOT** select the *Yes* check box on the *Reinstall caGrid?* screen. Press *Next*.
8. Do **NOT** select the *Yes* check box on the *Reconfigure caGrid for target grid?* screen. Press *Next*.
9. Press the *Start* button on the *Install Components* screen. When it finishes, press *Next*.
10. On the *Secure Deployment* screen, select the *Yes* check box. Press *Next*.
11. On the *Service Hostname* screen, enter the host name of this machine. Press *Next*.
12. On the *Specify Tomcat Ports* screen, enter the *Shutdown Port* and *HTTPS* port. (In most cases, the defaults should be fine) Press *Next*.
13. On the *Service Credentials* screen, select the *Yes* check box. Press *Next*.
14. On the *Service Credentials Information* screen, navigate to and select the certificate and key files that you created through the caGrid security UI. Press *Next*.
15. Press the *Start* button.

The installer will deploy Globus to Tomcat and configure Tomcat and the Globus web application accordingly. Once the installer is finished, press *Next*.

The installer will indicate that the CATALINA_HOME environment variable must now also be [set](#).

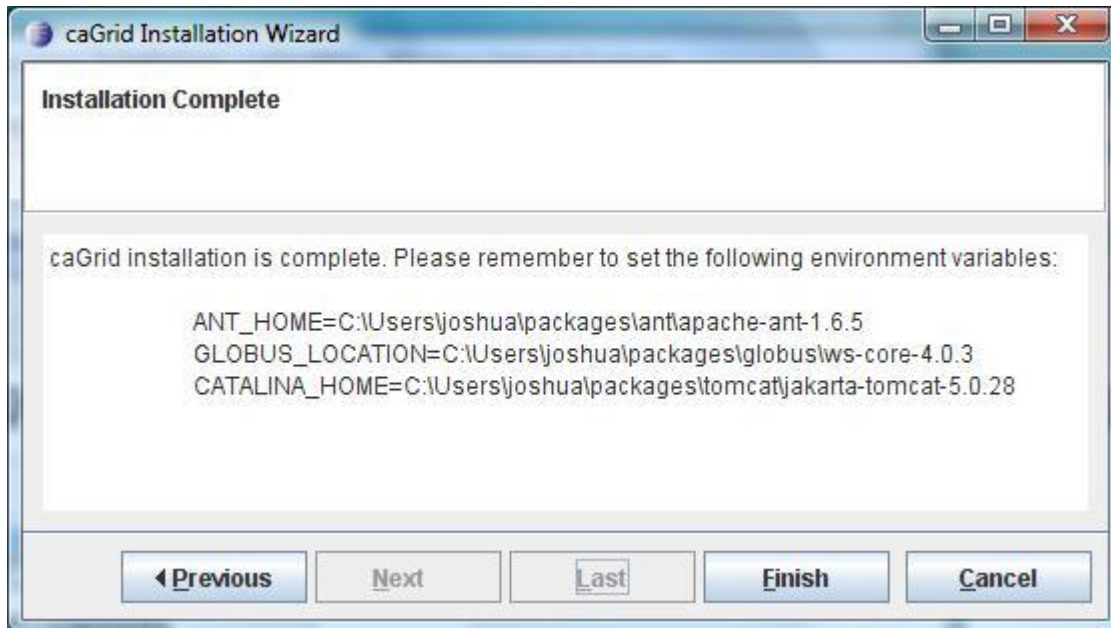


Figure 6: Container configuration complete

On some systems, you will also need to modify the file permissions of the components that have been installed. For example, on Mac and *nix type systems, you could running following commands:

```
chmod u+x $CATALINA_HOME/bin/*.sh
chmod u+x $ANT_HOME/bin/*.sh
chmod u+x $GLOBUS_LOCATION/bin/*
```

Environment Settings

In order for the caGrid development environment and deployed caGrid services to work properly, certain environment variables must be set. If these environment variables have not been set, the installer will indicate how they should be set, based on the input the user provides during the installation process. Also, since the installer is designed to work on systems where caGrid (and some or all of its dependencies) have already been installed, the installer will examine the environment and give the user options to re-use or re-install these components (possibly to a new location on the file system).

The environment variables that the installer looks for are as follows:

- **JAVA_HOME:** This should point to the Java 1.5 Java Develop Kit (JDK), not the Java Runtime Environment (JRE).
- **ANT_HOME:** Should point to the Ant 1.6.5 installation directory. The caGrid development environment requires this variable, but running caGrid services do not.
- **GLOBUS_LOCATION:** Should point to the caGrid Globus 4.0.3 distribution installation directory. This is required by the caGrid development environment and caGrid services that are running in the Globus container.

- **CATALINA_HOME:** Should point to the Jakarta Tomcat 5.0.28 installation directory. This is required by the caGrid development environment and is necessary to run the Tomcat container.

Following are examples of how to set these variables of various systems:

Unix/Linux and Mac

Suppose you have the following installation:

- Java: /usr/local/jdk1.5.0_12
- Ant: /usr/local/apache-ant-1.6.5
- Globus: /usr/local/ws-core-4.0.3
- Tomcat: /usr/local/jakarta-tomcat-5.0.28

If you are using the Bash shell, then you should add the following to your ~/.bash_profile file:

```
export JAVA_HOME=/usr/local/jdk1.5.0_12
export ANT_HOME=/usr/local/apache-ant-1.6.5
export GLOBUS_LOCATION=/usr/local/ws-core-4.0.3
export CATALINA_HOME=/usr/local/jakarta-tomcat-5.0.28
export PATH=$JAVA_HOME/bin:$ANT_HOME/bin:$PATH
```

Microsoft Windows

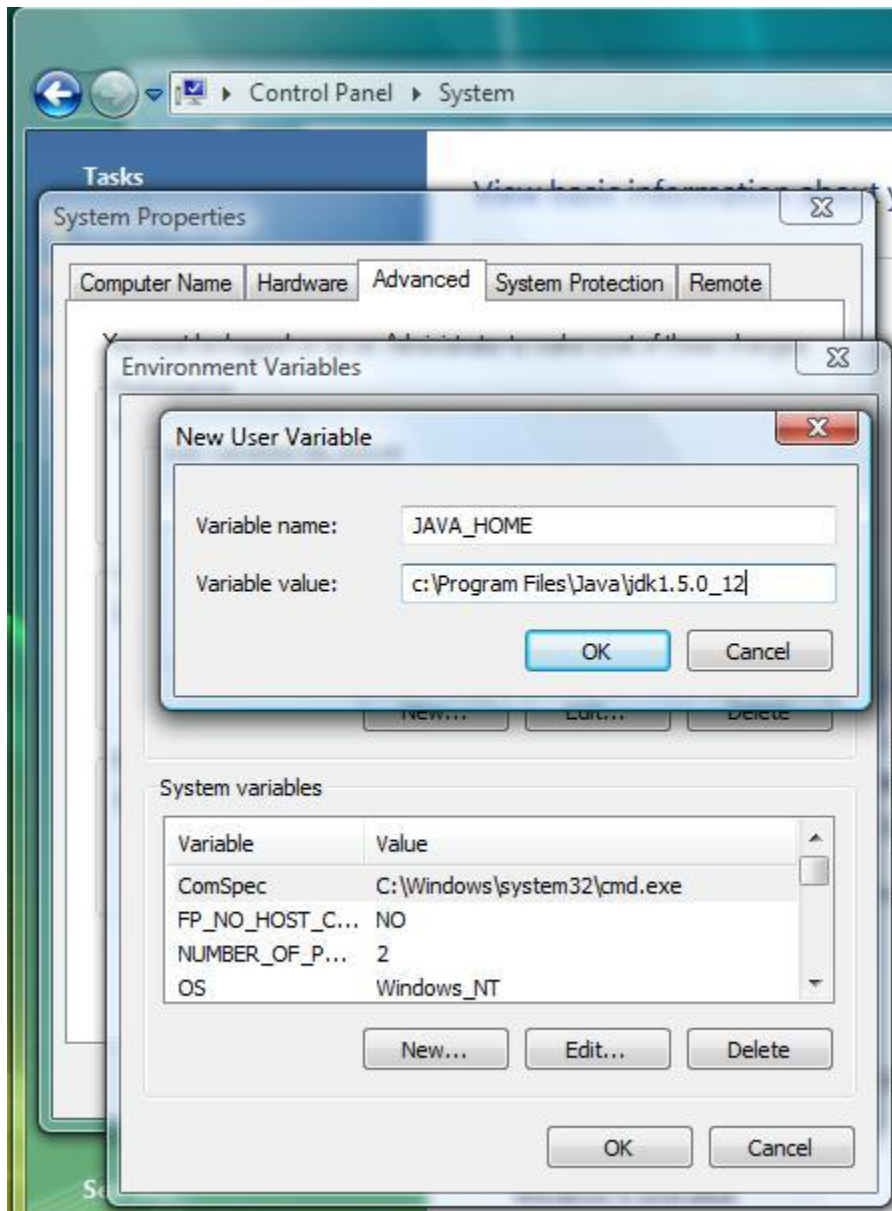


Figure 7: Setting MS Windows environment variables.

Suppose you have the following installation:

- Java: C:\Program Files\Java\jdk1.5.0_12
- Ant: C:\Users\me\apache-ant-1.6.5
- Globus: C:\Users\me\ws-core-4.0.3
- Tomcat: C:\Users\me\jakarta-tomcat-5.0.28

In the Microsoft Windows systems, you can set environment variables in the registry using the System administration application that is found in the Control Panel. From you desktop...

- Click the *Start* button.

- Select *Control Panel*.
- Double-click the *System* application.
- Select *Advanced System Settings* (or the *Advanced* tab, depending on your version of MS Windows).
- Click the *Environment Variables* button.
- Create environment variables as in Figure 7. These can be either *User* or *System* variables.

You could also run a batch file that contains the following from the command prompt:

```
set JAVA_HOME=C:\Program Files\Java\jdk1.5.0_12
set ANT_HOME=C:\Users\me\apache-ant-1.6.5
set GLOBUS_LOCATION=C:\Users\me\ws-core-4.0.3
set CATALINA_HOME=C:\Users\me\jakarta-tomcat-5.0.28
set PATH=%JAVA_HOME%\bin;%ANT_HOME%\bin;%PATH%
cmd
```

NOTE: Environment variables set in this way will only be noticed by programs that are launched from the same command prompt window.

Advanced Usage

Off-Line Installation

The installer can be configured to pull dependencies from the file system (or local network).

Step 1: Create a file named `cagrid.installer.properties` and enter the following content into it

```
ant.download.url=file:/home/me/software/apache-ant-1.6.5-bin.zip
ant.temp.file.name=ant.zip
ant.dir.name=apache-ant-1.6.5
tomcat.download.url=file:/home/me/software/jakarta-tomcat-5.0.28.zip
tomcat.temp.file.name=tomcat.zip
tomcat.dir.name=jakarta-tomcat-5.0.28
```

```
globus.download.url=file:/home/me/software/ws-core-enum-4.0.3.zip
globus.temp.file.name=globus.zip
globus.dir.name=ws-core-4.0.3
cagrid.download.url=file:/home/me/software/caGrid-1.1-src.zip
cagrid.temp.file.name=cagrid.zip
cagrid.dir.name=caGrid
activebpel.download.url=file:/home/me/software/activebpel-3.0.3-all.zip
activebpel.dir.name=activebpel-3.0.3
activebpel.temp.file.name=activebpel-3.0.3-all.tar.gz
browser.download.url=file:/home/me/software/caGrid-1.1-browser.zip
browser.temp.file.name=browser.zip
browser.dir.name=cagrid-browser
available.target.grids=nci_prod,nci_qa,nci_stage,nci_dev,osu_dev,training
jdbc.driver.path=lib/mysql-connector-java-3.0.16-ga-bin.jar
jdbc.driver.classname=com.mysql.jdbc.Driver
java.version.pattern=1.5.*
```

Modify the properties that end in "download.url" to point to the location on your file system where you have saved copies of the caGrid dependencies.

NOTE: For Globus and ActiveBPEL, you have to use the caGrid-specific, patched versions:

- <http://gforge.nci.nih.gov/frs/download.php/1334/ws-core-enum-4.0.3.zip>
- <http://64.119.177.125/download/files/3.0/patches/activebpel-3.0.3-all.zip>

Step 2: Create a file named download.properties and enter the following content:

```
download.url=file:/home/me/cagrid.installer.properties
```

Modify the value of *download.url* to be the path to the 'cagrid.installer.properties' file that you created previously.

Step 3: Point the installer at your modified properties file

```
java -Ddownload.properties=/home/me/download.properties -jar caGrid-1.1-
installer.jar
```

Custom Target Grid

If you are setting up your own grid, you can configure the installer to provide the option to configure caGrid for your custom target grid.

Step 1: Following [these](#) directions to create a custom caGrid distribution.

Step 2: Configure the installer to present custom target grid option

Follow [these](#) instructions to create a custom cagrid.installer.properties file. In that file, edit the *available.target.grids* property to include the name of the custom target grid directory that you created in Step 1. For example, if you created...


```
caGrid/share/resources/target_grids/custom_grid
```

...then the property should look like this:

```
available.target.grids=nci_prod,nci_qa,nci_stage,nci_dev,osu_dev,training,cus  
tom_grid
```

Step 3: Add a label for your custom target grid.

First...

```
jar xf caGrid-1.1-installer.jar CaGridInstallerMessages_en_US.properties
```

...then add this new property to that file:

```
target.grid.custom_grid.label=My Custom Grid
```

...then,

```
jar uf caGrid-1.1-installer.jar CaGridInstallerMessages_en_US.properties
```

Step 4: Follow *Step 3* of [these](#) instructions to point the installer at you custom *cagrid.installer.properties* file.

Troubleshooting

Problems:

- **Error loading default properties:...**
 - This error message will be displayed at startup if the host on which the installer is running has no connection to the Internet. Check your network settings.
 - You may need to specify proxy settings if you are behind a firewall. For example:

```
java -Dhttp.proxyHost=<your proxy host> -Dhttp.proxyPort=<your proxy host> -  
jar caGrid-1.1-installer.jar
```

Consult the [java network settings](#) for more details.

- **Error initializing: The version of Java found at ... is not correct. ...**
 - This error message will be displayed if the installer could not locate the correct version of Java. You can point it to the correct version by setting the JAVA_HOME environment variable to point to the Java 1.5 JDK. See [here](#) for directions on setting environment variables.
- **Could not connect to database.**

- Several of the caGrid services require a database. During the configuration of these services, the installer will need to interact with the database (i.e. to validate the connection, create/destroy the database, pre-populate the database, etc.) If you are seeing a message like this, check that 1) the database services is running; 2) you have specified the correct connection and user account information; and 3) the user account has privileges to create and drop database, and query/insert/update/delete data.

Retrieved from "http://www.cagrid.org/mwiki/index.php?title=CaGrid:How-To:UseCaGridInstaller_1-1"