

# ECCE 4.5 Installation and Administration

## Overview

Before installing ECCE we highly recommend reviewing the [platform requirements](#) in order to save unnecessary aggravation should the system you wish to install on not meet these specifications for operating system version, graphics hardware, memory, disk space, etc.

When running ECCE three distinct types of compute resources are used:

- Desktop workstations running ECCE client graphical user interface applications
- Data and messaging servers
- Compute resources running chemistry codes

ECCE client application software such as the Molecule Builder, Calculation Editor, and Calculation Viewer run on Linux and UNIX workstations under the X Window System. These machines can only be Intel Pentium and Itanium, AMD Athlon and Opteron processor Linux PCs, and SGI workstations, all constrained by the minimum platform requirements referenced above. The application software can be installed on each individual workstation running ECCE, or preferably, on a shared file system such as NFS or AFS that each workstation accesses.

The ECCE data and messaging server are packaged together to run on the same host. The data server is an Apache2 web server compiled with the supplied mod\_dav module to support distributed data access. The mod\_dav module implements the Distributed Authoring and Versioning (DAV) protocol (<http://www.webdav.org/>) for storing and retrieving data on a web server. The messaging server relies on an open source client/server implementation of the Java standard for interprocess communication (Java Messaging Service), named ActiveMQ (<http://activemq.apache.org>), to manage communications between the separate applications running in an ECCE session. ECCE has applied this technology to C++ applications by embedding a Java virtual machine inside the main Gateway application. In addition to the data and messaging server software, ECCE reference data libraries (basis set and structure libraries) and a small number of ECCE server administration scripts are provided. Either a single data and messaging server (referred to collectively as an ECCE server) can be installed at a site serving all workstations running the ECCE application software, or multiple servers can be installed with user access to each being configured by the ECCE site administrator.

The compute resources where the chemistry codes are run, sometimes referred to as compute servers, range from single processor desktop workstations to massively parallel supercomputers and clusters running a variety of queue management systems. They are not constrained by the platforms on which the ECCE application and server software must be run, although they must run UNIX, Linux, or Mac OS (OS X minimum). The NWChem computational code is bundled with the ECCE distribution and may be used to run chemistry jobs on the same host as ECCE client applications. However, for shared file system installations of ECCE client application software, NWChem performance is severely degraded to the

point that this may prove impractical. More than likely you will want to run NWChem and other chemistry codes on higher performance dedicated compute resources rather than where ECCE client applications are installed. This will require separate installation of the chemistry code from ECCE as well as registration of the compute resource in ECCE as described later in this document.

It is important to understand that each of these three compute resource types is distinct within ECCE. It is possible for a single desktop workstation to function as all three of these for an ECCE installation and represent the extent of the ECCE configuration at a site. Conversely, there may be dozens of workstations running the application software (Linux and SGI), accessing several ECCE servers, and farming out chemistry jobs to compute resources both at the site and around the world.

The installation procedures assume a basic familiarity with UNIX/Linux system administration. Commands given are for csh and should also apply to tcsh; if you are using another shell such as Bourne sh or ksh, you will need to adjust the syntax accordingly. The example installations described in this document are for a Linux workstation. If you are installing on an SGI workstation instead, you would replace the obvious Linux name references, such as "i686\_Linux\_glibc2.3", with the appropriate IRIX names.

## Install Third-Party Software

These PREREQUISITE software packages should be installed by a system administrator on all hosts where ECCE will be run. The information below for each package specifies whether they are needed on ECCE application software workstations, ECCE servers, and/or compute resources. Please make sure you have obtained and properly installed all these third-party applications before installing ECCE. The required packages are:

- **Perl** (needed by application software workstations, ECCE servers, and compute resources): Locate or install the Perl script language for all platforms where ECCE will be run. Users must have Perl in their PATH on both the host where ECCE applications are run and all compute resources where chemistry codes are run. Additionally, ECCE cannot be installed without Perl being found. Establish whether perl is on your machine using the command:

```
prompt% which perl
```

- **Secure Shell (ssh)** (needed by application software workstations and compute resources): Locate or install ssh and scp on your machine and make sure it is version 1.2.25 or newer. ECCE supports both ssh protocols version 1 and version 2 (only openssh has been tested for version 2). Because it provides secure user authentication and data transfer, ssh is the recommended remote communications shell for job launching and other remote communication in ECCE. Other shells, including telnet and rsh, are also supported so ssh is not strictly required, but strongly recommended. There is also a mechanism within ECCE for sites to define remote shells other than those with built-in support, which has been used to support shells like Kerberized rsh. Ssh, and scp as the remote file transfer command used in conjunction with ssh, must be in the users' PATH on both the ECCE

application workstations and compute resources where it is used. To find out if ssh and scp are available and the version, use the commands:

```
prompt% which ssh  
prompt% which scp  
prompt% ssh -V
```

- **Netscape, Mozilla, or Firefox** (needed by application software workstations): Locate or install netscape, mozilla, or firefox on your machine and make sure it is in a directory in the PATH for each user. Online help and the user support request web page rely on one of these web browsers; which one is determined by installation settings.
- **OpenGL** (needed by application software workstations):

Linux: ECCE works with desktops and laptops using either Mesa OpenGL software rendering or hardware OpenGL rendering using a graphics card along with the associated Linux driver software. We suggest using Linux PCs with hardware graphics cards supporting OpenGL whenever possible, with NVIDIA and ATI cards outfitted with the most on-board memory available, being our top recommendations when purchasing new systems for use with ECCE. If you need to use software graphics rendering, the Mesa OpenGL library will most likely be included with your version of Linux depending upon the vendor and version.

SGI: Normally all IRIX workstations have OpenGL installed along with another graphics library needed by ECCE, OpenInventor, as part of the window system software. Some IRIX compute platforms such as the Power Challenge will not have the libraries installed although they are available. On a related note, the Power Challenge platform normally does not come with the graphics hardware required for ECCE so you will need to verify whether you have this hardware as well as the OpenGL and OpenInventor libraries, if you plan to run ECCE applications on a Power Challenge. You must obtain the OpenGL and OpenInventor libraries from SGI and install them using the SGI "inst" command. You can proceed to install ECCE and then determine whether you have the required graphics hardware based on whether an ECCE "visualization" application like the Molecule Builder displays chemical systems in the workspace area or not. If the OpenGL and/or OpenInventor libraries have not been installed you will get missing shared library error messages trying to start visualization applications.

- **Chemistry Codes** (needed by compute resources): If you want to run NWChem, Gaussian 03, Gaussian 98, or Amica jobs on machines where it is not currently available, please install them using the documentation provided with those packages. Note that ECCE requires NWChem version 4.0 or newer. A compatible version of NWChem is bundled, installed, and automatically registered for use on the host where ECCE client application software is installed. For installations of application software to local disk or for evaluating ECCE prior to any large-scale deployment, this distribution of NWChem may prove adequate.

## Register and Download the ECCE Distribution

You must provide the information asked for in the EMSL User System and then submit a software user agreement as directed in the EMSL User System before you can download the ECCE distribution. See the [ECCE registration page](#) for details. Once approved at EMSL, you will be provided with download instructions. Using a web browser, download the appropriate ECCE distribution for your platform to a local disk directory (or download multiple ECCE distributions if you will be running ECCE from more than one platform). File size for downloading is displayed on the web page although having sufficient additional free disk space during installation is critical since uncompress and untar steps will multiply the size requirements temporarily. The download web page provides information on which distribution is appropriate for your platform. Download the ECCE software distribution as a single file, e.g., `install_ecce.v4.5.i686_Linux_glibc2.3.csh`. Download distributions for each of the different platforms users will run the ECCE graphical user interfaces on. One ECCE server can be shared for all the different application software platforms with the server being installed on the platform of your choice. The ECCE distribution for each platform contains both the application software and server, although installation options allow you to selectively install only applications or the server, as well as both.

The application and server software is compressed using GNU gzip. For SGI, gunzip is not included with the operating system. If you do not already have gunzip, a link on the ECCE download web page points to miscellaneous download files, including the gunzip executable in standard UNIX “.Z” compressed format. Download, uncompress, and make sure gunzip is in your path when running the install script. The gunzip utility is already provided with all Linux operating systems.

## Install ECCE

In order to maintain the integrity of the installation, we recommend creating an account named `ecceadm`, for “ECCE Administrator”, or something similar and installing as that user. There are many configuration files, along with executables and libraries distributed with ECCE that if removed or improperly modified will corrupt the installation.

The ECCE distribution itself is a self-extracting C shell installation script along with the application software and server compressed tar files. The installation script has a main menu allowing you to select the type of install to be performed. Normally a “full” installation of both the application and server software is done. This option is appropriate for either standalone or networked hosts. Further options allow the application software and server to be installed independently. These options accommodate application software being installed on all desired platforms (by downloading and running the distribution script on each platform), with the server installed on just a single platform, among other scenarios. The main menu allows both new installations and upgrades of older releases of ECCE to the current version.

Depending upon the policies at your site, you may need to do the server side of the installation as another user rather than `ecceadm`. If you install as root the Apache `httpd` daemon will automatically run as the unprivileged “nobody” user, an additional security benefit when your network, firewall, and web server configuration allows external access. If you wish to do a server install as root, you may either install all of ECCE as root (“full install”), or you can install the application software first as `ecceadm` (or another

user) selecting “Application software only” from the main menu, and then doing a second install as root selecting “Server only” from the main menu.

The following steps document sample installations for version 4.5 of ECCE on a Linux host for two common scenarios--a full install and a full upgrade. If you are doing other than a full ECCE install or upgrade, then the prompts below will vary somewhat. The host these sample installations are performed on runs a version of the Linux operating system using version 2.3.x of the GNU C Library such as Red Hat Enterprise Linux Workstation 4 or Fedora Core 4. The ECCE download page directs you to which ECCE distribution to download based on your operating system. For other platforms, remember to modify the commands accordingly.

- **Full ECCE Install**

Run the `install_ecce.v4.5.*.csh` script in the directory where it was downloaded. You may need to add execute permission to the file first. Note that values in square brackets are the default and you may simply hit return to use the default. For clarity, values are always explicitly entered for the prompts in these examples even when the default value is used. Here is the first sample invocation of `install_ecce.v4.5.*.csh`, a full install, run as `ecceadm` with links to notes describing how to determine appropriate values for each of the configuration settings:

```
prompt% cd /myfiles
prompt% chmod +x ./install_ecce.v4.5.i686_Linux_glibc2.3.csh
prompt% ./install_ecce.v4.5.i686_Linux_glibc2.3.csh
```

Extracting ECCE distribution from `./install_ecce.v4.5.i686_Linux_glib2.3.csh...`

Main ECCE installation menu

- ```
=====
0) Help on main menu options
1) Full install
2) Full upgrade
3) Application software install
4) Application software upgrade
5) Server install
6) Server upgrade
```

IMPORTANT: If you are uncertain about any aspect of installing or running ECCE at your site, please refer to the detailed ECCE Installation and Administration Guide at <http://ecce.emsl.pnl.gov/docs/installation/2864B-Installation.pdf>

Hit <return> at prompts to accept the default value in brackets.

Selection: [0] **1**

Host name: [mymachine.mysite.mydomain] **mymachine.mysite.mydomain**  
Application installation directory: [/myfiles/ecce-v4.5/apps] **/sharednfs/ecce-v4.5/apps**  
Server installation directory: [/myfiles/ecce-v4.5/server] **/myfiles/ecce-v4.5/server**

ECCE v4.5 will be installed using the settings:

Installation type: [full install]  
Host name: [mymachine.mysite.mydomain]  
Application installation directory: [/sharednfs/ecce-v4.5/apps]  
Server installation directory: [/myfiles/ecce-v4.5/server]

Are these choices correct (yes/no/quit)? [yes] **yes**

Installing ECCE application software in /sharednfs/ecce-v4.5/apps...  
Extracting application distribution...  
Extracting NWChem distribution...  
Extracting Java JRE distribution...  
Extracting client WebHelp distribution...  
Configuring application software...  
Configuring NWChem...  
Changing application file permissions...

Installing ECCE server in /myfiles/ecce-v4.5/server...  
Extracting data server in /myfiles/ecce-v4.5/server/apache...  
Extracting data libraries in /myfiles/ecce-v4.5/server/data...  
Extracting messaging server in /myfiles/ecce-v4.5/server/activemq...  
Configuring ECCE server...

ECCE installation succeeded.

\*\*\*\*\*

!! You MUST perform the following steps in order to use ECCE !!

- Start the ECCE server as ecceadm by running  
  '/myfiles/ecce-v4.5/server/ecce-utils/start\_ecce\_server'
  
- To register machines to run computational codes, please see  
  the installation and compute resource registration manuals  
  at <http://ecce.emsl.pnl.gov/using/installguide.shtml>
  
- To run ECCE each user must source either the runtime\_setup  
  (csh/tcsh) or runtime\_setup.sh (sh/bash/ksh) script in the  
  directory /sharednfs/ecce-v4.5/apps/scripts  
  from their shell environment setup script. For example,

```
with csh or tcsh, add the following to ~/.cshrc:
if (-e /sharednfs/ecce-v4.5/apps/scripts/runtime_setup) then
    source /sharednfs/ecce-v4.5/apps/scripts/runtime_setup
endif
```

-- To run the NWChem distribution packaged with ECCE, each user must create a .nwchemrc file in their home directory copying the template created during this install. For example:

```
cp /sharednfs/ecce-v4.5/apps/nwchem/etc/default.nwchemrc ~/.nwchemrc
*****
prompt%
```

### • Full ECCE Upgrade

Here is the second sample invocation of install\_ecce.v4.5.\*.csh, a full upgrade from a release of ECCE v4.0.2, run as eceedm with links to notes describing how to determine appropriate values for each of the configuration settings:

```
prompt% cd /myfiles
prompt% chmod +x ./install_ecce.v4.5.i686_Linux_glibc2.3.csh
prompt% ./install_ecce.v4.5.i686_Linux_glibc2.3.csh
```

Extracting ECCE distribution from ./install\_ecce.v4.5.i686\_Linux\_glibc2.3.csh...

Main ECCE installation menu

- ```
=====
0) Help on main menu options
1) Full install
2) Full upgrade
3) Application software install
4) Application software upgrade
5) Server install
6) Server upgrade
```

IMPORTANT: If you are uncertain about any aspect of installing or running ECCE at your site, please refer to the detailed ECCE Installation and Administration Guide at <http://ecce.emsl.pnl.gov/docs/installation/2864B-Installation.pdf>

Hit <return> at prompts to accept the default value in brackets.

Selection: [0] **2**

Host name: [mymachine.mysite.mydomain] **mymachine.mysite.mydomain**

New application installation directory: [/myfiles/ecce-v4.5/apps] **/sharednfs/ecce-v4.5/apps**  
Existing application directory to upgrade: **/sharednfs/ecce-v4.0.2/apps**  
New server installation directory: [/myfiles/ecce-v4.5/server] **/myfiles/ecce-v4.5/server**  
Existing server directory to upgrade: **/myfiles/ecce-v4.0.2/server**  
Backup existing server user data (yes/no)? [yes] **yes**

ECCE v4.5 will be installed using the settings:

Installation type: [full upgrade]  
Host name: [mymachine.mysite.mydomain]  
Application installation directory: [/sharednfs/ecce-v4.5/apps]  
Application directory to upgrade: [/sharednfs/ecce-v4.0.2/apps]  
Server installation directory: [/myfiles/ecce-v4.5/server]  
Server directory to upgrade: [/myfiles/ecce-v4.0.2/server]  
Backup existing server user data: [yes]

Are these choices correct (yes/no/quit)? [yes] **yes**

Installing ECCE application software in /sharednfs/ecce-v4.5/apps...  
Extracting application distribution...  
Extracting NWChem distribution...  
Extracting Java JRE distribution...  
Extracting client WebHelp distribution...  
Configuring application software...  
Configuring NWChem...  
Changing application file permissions...

Installing ECCE server in /myfiles/ecce-v4.5/server...  
Extracting data server in /myfiles/ecce-v4.5/server/apache...  
Extracting data libraries in /myfiles/ecce-v4.5/server/data...  
Extracting messaging server in /myfiles/ecce-v4.5/server/activemq...  
Configuring ECCE server...  
Copying user data from server to be upgraded...  
Copying share data from server to be upgraded...

ECCE installation succeeded.

\*\*\*\*\*

**!! You MUST complete the following steps in order to use ECCE !!**

-- Start the ECCE server as ecceadm by running  
    '/myfiles/ecce-v4.5/server/ecce-utils/start\_ecce\_server'

-- To register machines to run computational codes, please see

the installation and compute resource registration manuals  
at <http://ecce.emsl.pnl.gov/using/installguide.shtml>

```
-- To run ECCE each user must source either the runtime_setup
(csh/tcsh) or runtime_setup.sh (sh/bash/ksh) script in the
directory /sharednfs/ecce-v4.5/apps/scripts
from their shell environment setup script. For example,
with csh or tcsh, add the following to ~/.cshrc:
    if (-e /sharednfs/ecce-v4.5/apps/scripts/runtime_setup) then
        source /sharednfs/ecce-v4.5/apps/scripts/runtime_setup
    endif
```

```
-- To run the NWChem distribution packaged with ECCE, each user
must create a .nwchemrc file in their home directory copying
the template created during this install. For example:
    cp /sharednfs/ecce-v4.5/apps/nwchem/etc/default.nwchemrc ~/.nwchemrc
*****
prompt%
```

After the `install_ecce.v4.5.*.csh` script has completed with both the application software and server being installed as desired, you may delete the distribution script although we recommend waiting until ECCE is completely tested at your site.

The next several paragraphs describe how to choose appropriate values for the ECCE installation prompts as shown above. If you have successfully run the installation script and understand the prompts, you may skip to [Post-Install Configuration](#) to continue with the installation.

**Main menu selection:** There are three basic types of ECCE installations where each type can either be a from-scratch install or an upgrade of an existing installation. An upgrade is done rather than a from-scratch install, referred to as simply an “install”, when you have an existing version of ECCE where you wish to incorporate the machine registrations and/or calculation data for that version into the new version. An install is done rather than an upgrade when you have either not previously installed ECCE or do not wish to incorporate data from an existing version. The following describes each of the six main menu options (a seventh option gives help on the other main menu options):

1. Full install: Install both the ECCE application software and server without incorporating data created from a previous ECCE install, if any. This is the recommended selection when special needs, as described below, do not dictate separate application and server software installations. A full install can be done either for use from a single machine on or off the network, or for use from multiple machines. For use from a single machine the host name is often entered as “localhost”, if defined properly in the `/etc/hosts` file, to facilitate taking machines such as laptops on and off networks with ECCE working in both configurations. Application software is installed to local disk for use from a single machine or to a shared file system, such as NFS or AFS, for use from multiple machines. The ECCE server is always installed to local disk, but

the server will be accessible externally to the extent allowed by the network and firewall configuration of the server host. Security and performance of the ECCE server would both be compromised if it was installed to a shared file system and there are no advantages in doing so. Even when installed for use from a single machine, computational jobs may still be run to other machines with a full install, assuming a network connection exists of course.

2. Full upgrade: Install a new version of the ECCE application software and server on a machine incorporating the machine registrations and calculation data from a previous install. Upgrades can only be done when the original installation was a v3.2.x install. The large number of changes between each successive version of ECCE prevents us from supporting upgrades of all but the most recent previous releases.
3. Application software install: Install only the application software and not the server. Used when ECCE will be run from multiple platforms, but the server has been already been installed on a different platform with a full install, or will be installed with a server only install.
4. Application software upgrade: Install a new version of the application software incorporating machine registrations and other application software configuration data from a previous installation. Used for installing patches to application software when there are no server enhancements, or upgrading multiple platforms for running applications. Also can be used to only incorporate machine registrations from a previous version of ECCE without incorporating user calculation data from the previous version.
5. Server install: Install only the server and not the application software. Used when the server needs to run on a different machine or platform than where the application software is installed.
6. Server upgrade: Install a new version of the server incorporating user calculation data from a previous installation. Used for installing server-side only patches or when user calculation data from a previous version needs to be upgraded without application software being upgraded.

**Host name:** This is the full domain name of the machine where you are currently running the `install_ecce.v4.5.*.csh` script. You cannot install the ECCE server on a different machine than the one you run the script on. If you do not wish to install the server on the machine you are currently on, you must copy the `install_ecce.v4.5.*.csh` script to the desired machine and run there. The host name prompt confirms the name detected by the installation script is correct. In the majority of cases you will simply accept the default server name as shown in the square brackets, provided it is the correct full domain name. In the instance you wish to refer to the server name by an alias, use the IP address rather than the name, or use "localhost" as the name for a machine that is taken off the network, you may enter the desired name at this prompt. Note, of course, that if you set the name as "localhost", you will only be able to run ECCE with this server from the local machine regardless of whether it is currently on or off the network.

**Application installation directory:** ECCE applications are designed for installation under a single shared directory at a site, with all Linux and SGI hosts accessing the common application installation.

A local file system disk is selected for the application installation directory when ECCE will only be run from a single machine. If you wish to have multiple machines and/or platforms running ECCE, we strongly recommend installing the application software in a shared file system as it substantially reduces parallel administration. The directory where ECCE applications are installed is independent of the directory where the distribution is initially downloaded, although the default will be a subdirectory of the download directory. For access from multiple machines the distribution is typically downloaded to local disk to speed file transfer, and then installed to a shared file system. This file system for the application installation directory should have at least 300 megabytes free for each platform that will be installed. To install multiple platforms to a shared file system, specify the same application installation directory for each platform (requires downloading and running the installation script distribution for each platform). The ECCE installation script automatically extracts platform-dependent libraries and executables into separate subdirectories. The absolute path to the directory where ECCE applications are installed is stored in the ECCE runtime\_setup script as the environment variable \$ECCE\_HOME. The rest of this document uses the term “ECCE application installation directory” and slight variations interchangeably with the variable \$ECCE\_HOME, especially when referring to file paths.

You are not allowed to overwrite an existing ECCE application installation except in the case where you are installing different platforms under the same top-level application installation directory. The installation script will verify that the directory specified is not a previous installation for the same platform and prompt for a new directory if it is. This restriction prevents the inadvertent loss of a working installation should the new one have some kind of problem. If you do wish to install a new version of the ECCE applications to the same top level directory as an existing installation for the same platform, you must move the old installation to another directory prior to running the install\_ecce.v4.5.\*.csh script. We recommend that you only remove an existing installation after you have verified the operation of the new installation. Certain releases, such as quick turnaround patches done for individual sites, have not been tested to the extent as major releases so you may find problems severe enough that you wish to revert to the previous install.

**Server installation directory:** The Apache server binaries, Apache server data directory, ActiveMQ JMS server, Java runtime environment needed by ActiveMQ, and a directory named ecce-utils containing some scripts to facilitate ECCE server administration, will all be placed under this directory. This directory should always be on local disk. All data created within ECCE will reside under the “data” subdirectory of the server installation directory including both the original and ECCE-formatted input and output data from all jobs run with ECCE. For this reason, we highly recommend that you dedicate an entire disk (tens of gigabytes if not more) for the ECCE server, depending upon anticipated usage. If you are short on disk space, it is possible to configure multiple ECCE servers and divide users between the servers, but planning ahead will make administration easier.

**Are these choices correct:** Pressing return or entering “yes” results in the ECCE installation completing based on the summary of settings given before this prompt. Entering “no” results in all items being prompted for again including the main menu selection.

A list of steps of additional steps is printed at the end of the installation process. If you are installing an ECCE server, the installation script will attempt to update the ECCE application software files in the `$ECCE_HOME/siteconfig` directory that configure access to the new ECCE server. If this fails for some reason, such as the application software not being accessible from the machine where the ECCE server is currently being installed, the instructions will be given for manually completing this step. In the event you have multiple installations of the application software at your site (not recommended due to additional burden of maintaining installations in parallel) and you wish to share the same ECCE server, the ECCE installation script will only be able to update the application software `$ECCE_HOME/siteconfig` files for one of these installations. In this case, you will need to copy the `add_ecce_server` script from the `ecce-utils` directory under the server installation directory to the other hosts where the application software is installed and run it manually on each host.

**Existing application directory to upgrade:** Each site must customize their ECCE application installation in order to recognize the compute resources and ECCE servers available to the users at the site. This information is stored in the `$ECCE_HOME/siteconfig` directory. Performing an upgrade allows you to automatically copy the site-specific customizations for an existing installation into the new one. Enter the application installation `$ECCE_HOME` directory of the existing (old) installation at this prompt. If you specify an invalid directory for the existing installation for the given platform, a warning will be issued and you will be prompted again. As well as copying the contents of an existing `siteconfig` directory, this upgrade feature is also used to update files that have changed format in the new version of ECCE. Therefore, it is recommended that this feature be used instead of manually upgrading an existing installation.

**Existing server directory to upgrade:** Enter the path for the existing ECCE server installation at this prompt. User accounts and all data created under a previous server installation will be incorporated into the new server.

**Backup existing server user data:** We HIGHLY recommend making a backup copy of all user data stored under the server data directory before an upgrade between server versions is done. Hitting return or entering “yes” at this prompt will copy all data stored in the “users” and “share” subdirectories of the ECCE server data directory from the existing server installation into the new server installation. This can require considerable disk space and time, which is the reason this step is optional. If you do not make the backup copy, by entering “no” at this prompt, then the “users” and “share” data directories will simply be moved from the old to the new server installation with no possibility of continuing to use the existing server. If there is not enough disk space to make the backup, or for some other reason, you do not wish the ECCE installation script to automatically make the backup, you should make a backup manually (such as a tar file) to disk or another storage device before doing the upgrade.

## Post-install Configuration

- **Change the Default Web Browser**

ECCE uses a web browser for online help and the user support request web page. At this time, Netscape, Mozilla, and Firefox are compatible with how ECCE controls the display of web pages externally (KDE Konqueror, for instance, does not allow the control needed by ECCE). By default Mozilla is configured as the web browser within ECCE for both Linux and SGI platforms. This can be changed on a per-platform basis by editing the application software `$ECCE_HOME/siteconfig/site_runtime` file and scrolling down to the entries for web browsers. The only valid values for the browsers are “netscape”, “mozilla”, and “firefox”.

- **Change the User Support Email Address**

Email submitted via the ECCE support request web page available from the Help menu in all applications should be sent to someone at your site if you support multiple users. However, the default configuration sends these requests back to [ecce-support@emsl.pnl.gov](mailto:ecce-support@emsl.pnl.gov), the support queue maintained by the ECCE development team in EMSL. Should the [ecce-support@emsl.pnl.gov](mailto:ecce-support@emsl.pnl.gov) queue start receiving several support requests from a site regarding issues that should be handled by the site contact you will be asked to change your support address to redirect requests. The person doing the installation of ECCE at the site is one possible choice if they are also an ECCE user so that requests pertaining to both using and the administration of ECCE can be handled. This address may be changed at any time by editing the `site_runtime` file in the `$ECCE_HOME/siteconfig` directory and changing the value for the `ECCE_SUPPORT` entry. The site point of contact for user support may pass along problems, questions, and requests that they don't know how to handle to [ecce-support@emsl.pnl.gov](mailto:ecce-support@emsl.pnl.gov), either acting as an intermediary or judiciously setting up direct communication between the ECCE team and the user.

- **Change the Apache Data or ActiveMQ Messaging Server Ports**

ECCE uses a default port number of 8080 for the Apache data server. This is a “high port” and is usually a safe value that closes off the web server from outside a firewall. However, if this value is inappropriate for your site based on firewall configuration or external access to ECCE you wish to allow or deny, it can be changed manually. First, change directory to `apache/conf` under the server installation directory. Next, edit `httpd.conf`, search for “ECCE PORT”, and follow the instructions provided there to change the default. Next, update all application software installations using this server by editing the `$ECCE_HOME/siteconfig/DataServers` file URL for the server.

The ActiveMQ messaging server uses a default port number of 8088 as configured in ECCE. To change this value, change directory to `activemq/conf` under the server installation directory. Edit the file `activemq.xml` and search for “openwire”. Change the port number at the end of the value given for the “uri” keyword in the same line. Finally, update all application software installations using this server with the new ActiveMQ port in the `$ECCE_HOME/siteconfig/jndi.properties` file for each installation. Near the top of this file you will see the “`java.naming.provider.url`” keyword that needs to be updated with the new port value.

- **Access Multiple ECCE Servers**

To configure access to more than a single ECCE server for an application software installation, you must edit the DataServers file in the application software \$ECCE\_HOME/siteconfig directory. Follow the documentation in this file for configuring access to multiple servers.

- **Add ECCE Server Users**

The default configuration of ECCE will automatically create an ECCE server account for each user the first time they run the ECCE application software, eliminating this extra chore of creating these accounts from the ECCE administrator. The ECCE administrator can disable this feature if there are concerns about allowing any user with access to the application software to create their own area with write permission on the server. Note that users do not automatically get read or write access to other users' data. The environment variable \$ECCE\_AUTO\_ACCOUNTS defined in the \$ECCE\_HOME/siteconfig/site\_runtime file (which should be editable only by the ECCE administrator like all other files under \$ECCE\_HOME) controls whether this feature is enabled. The automatic account creation feature is implemented with an Apache cgi-bin script with no critical data being passed unencrypted over the Internet. For security reasons, server passwords will need to be manually reset by the ECCE administrator using the ecce\_htpasswd script, described below, for those users who forget their login password and must start over by clearing it in the password dialog. The confirmation dialog shown to users when the login password is cleared tells them to contact the ECCE administrator to have this done. If you plan to use the automatic server account creation feature, you may skip to the next section describing the ecce\_htpasswd script. [Goto next section](#)

If the automatic server account creation feature is disabled accounts must be created by the ECCE administrator. The ECCE server includes a script named add\_ecce\_user for this purpose. This script is located in the ecce-utils directory, under the server installation directory and must be run as the same user the installation was done as (ecceadm, if you are following the recommended install procedure). It performs the steps necessary to add a single new user account to the ECCE server per invocation. The add\_ecce\_user script will prompt for all input needed to make a user account and home directory. An example run, done as the ecceadm user (i.e., the ECCE server owner), is shown below. The prompt for "User name" is the Linux/UNIX machine login name corresponding to the \$USER environment variable for the user on the host(s) where they are running ECCE application software. Here is the add\_ecce\_user run:

```
prompt% /myfiles/ecce-v4.5/server/ecce-utils/add_ecce_user
```

```
First name? Joe  
Last name? Smith  
User name? jsmith
```

```
Entry confirmation:
```

```
First name: Joe  
Last name: Smith  
User name: jsmith
```

Is this information correct? [yes/no] **yes**

New password: **(text not echoed)**

Re-type new password: **(text not echoed)**

Adding password for user jsmith

ECCE uses basic web server authentication, built on standard UNIX crypt functionality. ECCE server passwords entered through `add_ecce_user` or as described below are independent of machine login passwords for users. We recommend passwords be chosen by some random means or using public domain tools for generating “nonsensical” words for such purposes.

It is also possible to setup users completely manually, bypassing the `add_ecce_user` script. The script merely simplifies the most common administrative duty of adding a new user by performing a number of small steps on your behalf. If you decide to setup access differently than is supported by the script or want to understand the process behind it, the following instructions lead you through manual user account creation. If you wish to use `add_ecce_user` without knowing the details you may skip to the next section describing the `ecce_htpasswd` script, which must be used to change passwords because the `add_ecce_user` script only supports adding new users. [Goto next section](#)

- **Add ECCE Server Users Manually**

The recommended configuration of ECCE for access control to data stored on the web server is to create a user name and password entry in the Apache server “users” file for each user. A default “home” directory should also be created (owned by the server data directory owner--either the ECCE server owner or “nobody” if root owns the ECCE server) under the ECCE/users directory of the data directory under the server installation. The home directory is named the same as the user name and contains an `.htaccess` file. This file directs the web server to require the user to authenticate with their user name and password to view and create data under their home directory.

To add a user named Joe Smith with a login name of jsmith first create a directory named jsmith in the ECCE/users directory under the server data directory. Within that directory create an `.htaccess` file with the following five lines:

```
AuthName "ECCE-Joe.Smith"  
<Limit OPTIONS HEAD GET PUT DELETE PROPFIND PROPPATCH MKCOL COPY MOVE  
LOCK UNLOCK POST>  
Allow from all  
require user jsmith ecceadm  
</Limit>
```

It is possible to be much more sophisticated with access control than is shown with this `.htaccess` file. You are referred to the numerous reference sources for Apache web server administration to configure

more complex access control. Besides creating a home directory and .htaccess file for each new user, a user name and password must be defined, as described next.

- **Run `ecce_htpasswd`**

The `ecce_htpasswd` script included in the server `ecce-utils` directory is used to add user name and password pairs to the server, or change the password for an existing user. This script is a simple wrapper for the standard Apache `htpasswd` program that sets up the proper `$LD_LIBRARY_PATH` and passes the correct argument for the user password file. An example invocation of `ecce_htpasswd` for adding the user `jsmith`, again run as `ecceadm`, follows:

```
prompt% /myfiles/ecce-v4.5/server/ecce-utils/ecce_htpasswd jsmith
New password: (text not echoed)
Re-type new password: (text not echoed)
Adding password for user jsmith
```

Changing a password for an existing user is done in the same way by passing the user name on the command line. Prompts and output from `ecce_htpasswd` are also similar in this case. The most common way to use `ecce_htpasswd` is with the user name as the only command line argument. However, other arguments supported by the standard `htpasswd` program can also be used with `ecce_htpasswd`. To see the full usage of the `htpasswd` program via `ecce_htpasswd`, type “`ecce_htpasswd`” with no arguments in the `ecce-utils` directory. Remember that the “`passwordfile`” argument is passed by `ecce_htpasswd` even though the usage description (which is for the standard `htpasswd` script) shows it. You may also run `htpasswd` directly if you set the `$LD_LIBRARY_PATH` to include the `apache/lib` directory under the server installation directory. ECCE is not compatible with MD5 or SHA password encryption when automatic account creation or password synchronization is used (`$ECCE_AUTO_ACCOUNTS` set to something besides “`false`”). Encryption is done on the application side for security reasons (to avoid sending passwords in the clear over the Internet) when `$ECCE_AUTO_ACCOUNTS` is enabled and only `CRYPT` format is supported. Finally, it is not necessary to restart the web server to recognize changes when running `add_ecce_user`, `ecce_htpasswd`, `htpasswd`, or manually editing `.htaccess` files.

- **Remove ECCE Server Users**

Due to the risk of accidentally losing valuable data, removing users from the ECCE server is done manually. Simply edit and delete the appropriate line from the Apache server “`users`” file (located in the “`apache`” directory under the server installation directory) and then remove their home directory (located in the `data/ECCE/users` directory under the server installation directory). Backing up this home directory before deletion is recommended if there is any chance this data will be needed in the future. As well as ECCE-specific files, all input and output log files from jobs run under ECCE are stored with user calculations, so they may later be used outside ECCE.

- **Start the ECCE Server**

At this point you should be able to start the Apache data and ActiveMQ messaging servers that make up the ECCE server. This should be done as the ECCE server installation owner (ecceadm if you are following the recommended install procedure) in order for the httpd web server daemon to have write access to files under the server data directory. Commands to start and stop the server are located in the ecce-utils directory under the ECCE server installation directory:

To start the server:

```
prompt% /myfiles/ecce-v4.5/server/ecce-utils/start_ecce_server
```

After starting the server, you may verify that the Apache httpd daemon and child httpd processes are running, as well as the ActiveMQ server:

```
prompt% ls /myfiles/ecce-v4.5/server/apache/logs/httpd.pid
prompt% ps -ef | grep "java" | grep "Dorg.apache"
```

To stop the server:

```
prompt% /myfiles/ecce-v4.5/server/ecce-utils/stop_ecce_server
```

Most likely, you will want the ECCE server automatically started during the machine boot sequence. Contact your system administrator to set this up.

## Register Compute Resources to Run Codes

The ECCE administrator must register each queued machine that serves as an ECCE compute resource where NWChem, GAMESS-UK, Gaussian 03, Gaussian 98, and Amica jobs will be run. Non-queued shared access machines are also typically registered by the ECCE administrator. Users may register their own personal workstations provided they don't use a queue management system. Also, compute resources shared by small groups that don't use queue managers can be registered by each user desiring access as an alternative to making them globally visible within ECCE user interfaces as is the case when the ECCE administrator performs the registration. Compute resource registration is done with a graphical interface application provided in the ECCE distribution. To use this application, you must first set up your environment to run as an ECCE user. This is done by changing the .cshrc for the ecceadm account (or whatever account owns the ECCE installation) as described in [Setup Users' Environment](#) and sourcing the modified .cshrc file. You will also need to give yourself an ECCE server user account as described above. For security reasons, there is no automatic server account creation for ecceadm. You will need to run the add\_ecce\_user script to create an account if you are running as ecceadm. Finally, enter the following command to start the machine registration graphical interface:

```
prompt% ecce -admin
```

Write permissions in the ECCE application software installation are required to save the compute resource registration information. We do not recommend opening up permissions on the "siteconfig"

directory where the registration information is stored, as improper changes can effectively disable ECCE for all users at the site.

Use the “Help” button to get information on using the interface, what all the input fields mean, and how to manually customize some queue management options that are not yet supported via the interface. Machines within EMSL accessible externally by EMSL collaborators are initially registered in the siteconfig directory as part of the ECCE distribution. There are examples for registering machines for each queue management system that ECCE currently supports in the CONFIG-Examples subdirectory of siteconfig. The configuration and site usage policies of most queued machines is too complex for them to be completely registered with the “ecce –admin” application. For these machines the “ecce -admin” application can be used to create template files based on the queue management system. These must be hand edited to account for the configuration of the individual machine. A PDF document describing compute resource registration is maintained on the ECCE web site and should be consulted whenever queued machines need to be registered.

## Setup Users’ Environment

- For csh and tcsh users, add the following lines to the user’s .cshrc (or .mycshrc within EMSL) file, substituting the correct path under which ECCE was installed on your host:

```
# setup to run ECCE
if ( -e /sharednfs/ecce/scripts/runtime_setup) then
  source /sharednfs/ecce/scripts/runtime_setup
endif
```

There is an equivalent script named runtime\_setup.sh in the same directory that can be sourced for sh, bash, and ksh shell users.

Hint: Within EMSL we maintain a symbolic link named scripts in a top-level shared ECCE directory that contains multiple releases of ECCE. The scripts link points to the scripts directory for the current production version of ECCE. This way the ECCE administrator only needs to update this link to change the version all ECCE users are running. Users wishing to run other than the default version the symbolic link points to can change their .cshrc file to reference the runtime\_setup script for the desired version. The path to runtime\_setup under /sharednfs/ecce/scripts in the example above is based upon this mechanism. In the example installation, ECCE was installed under /sharednfs/ecce-v4.5/apps. A symbolic link in the /sharednfs/ecce directory named scripts pointing to /sharednfs/ecce-v4.5/apps/scripts will make it easier for the ECCE administrator to upgrade users to new versions of ECCE. The commands to do this for the example installation are:

```
prompt% mkdir /sharednfs/ecce
prompt% cd /sharednfs/ecce
prompt% ln -s ../ecce-v4.5/apps/scripts scripts
```

- Logout then log back in again to make sure the environment is properly setup. Enter the command “which ecce” to make sure paths are correct.

## Run ECCE

- Start ECCE by typing...

```
prompt% ecce
```

Users will first be prompted for a password in a window containing the ECCE “wave” logo centered on their screen. The password, the single-point login for authenticating to ECCE, is set by the user the first time they start ECCE and can be changed at their will by hitting the “Change” button on the login password window or the “Clear” button if they have forgotten their password.

If automatic server account creation has been disabled, the first time a user runs ECCE there will be window prompting for the ECCE server user name and password after the login password window is dismissed. This window asks for the user name and password that was previously set with the `add_ecce_user` or `ecce_htpasswd` scripts on the ECCE server. Any time a user attempts to access data that is controlled by Apache server authentication, this same prompt will be displayed. Whenever a successful user name and password are entered, ECCE saves them for the user in an encrypted format so the user doesn’t need to re-enter them, including between ECCE sessions.

If automatic server account creation is enabled (the default), the user will not normally see these prompts for the ECCE server user name and password. Their ECCE server password will be automatically generated and saved as a pseudo-random sequence of characters with no user prompting required. Only when accessing data that is owned by another user, or their own data on an ECCE server that isn’t the default server (the default is listed first in the `SECCE_HOME/siteconfig/DataServers` file) are these prompts displayed.

If the login password for gaining access to ECCE is changed by a user, then all saved ECCE server passwords will be converted over to work with the newly selected login password. In the event that a user forgets their password, they must do a “clear” operation from the login password dialog to set a new one. In this instance, all saved ECCE server user names and machine passwords will be lost, and thus the prompt windows will be displayed again when trying to access these servers and compute resources. If automatic server accounts are being used (randomly generated passwords), the user will need to have their server password reset by the ECCE administrator as there is no secure way to allow users access to data when they have forgotten their login password. Clearing the login password also clears passwords for configured machines for running jobs. Users must select “Configure Machine...” from the Job Launcher or Machine Browser application to re-enter their password for each previously configured machine so ECCE will save them again. This design guarantees the security of login passwords as not even the ECCE administrator has access to them. It also emphasizes that users should not rely on ECCE to remember their passwords for them--ECCE stores passwords strictly as a convenience for users.

- **Where to go for help**

Check out the help pages on the web at <http://ecce.emsl.pnl.gov/help>, which are also accessible from within ECCE applications under the menubar Help menu. Also visit the user [FAQ](#) and [Release Notes](#) for version 4.5 of ECCE. The FAQ is out of date with respect to adding new items the last several releases of ECCE, but items have been removed that are no longer pertinent.