

How to Install an OSCAR Cluster
Software Version 3.0
Documentation Version 3.0

<http://oscar.sourceforge.net/>
oscar-users@lists.sourceforge.net

The Open Cluster Group
<http://www.openclustergroup.org/>

November 25, 2003

Contents

1	Introduction	7
1.1	Latest Documentation	7
1.2	Terminology	7
1.3	Supported Distributions	7
1.4	Minimum System Requirements	8
1.5	Document Organization	8
2	Downloading an OSCAR Distribution Package	9
3	Release Notes	9
3.1	Notes for All Systems	10
3.2	Red Hat 8/9.0 Notes	11
3.3	Mandrake 9.0 Notes	12
3.4	IA64 and Other Bleeding Edge Systems Notes	13
4	Overview of System Installation Suite (SIS)	13
5	Detailed Cluster Installation Procedure	13
5.1	Server Installation and Configuration	14
5.1.1	Install Linux on the server machine	14
5.1.2	Disk space and directory considerations	14
5.1.3	Download a copy of OSCAR and unpack on the server	14
5.1.4	Configure and Install OSCAR	15
5.1.5	Configure the ethernet adapter for the cluster	16
5.1.6	Copy distribution installation RPMs to /tftpboot/rpm	17
5.2	Launching the OSCAR Installer	17
5.3	Downloading Additional OSCAR Packages	18
5.4	Selecting Packages to Install	19
5.5	Configuring OSCAR Packages	19
5.5.1	Selecting a Default MPI Implementation	21
5.6	Install OSCAR Server Packages	21
5.7	Build OSCAR Client Image	21
5.8	Define OSCAR Clients	23
5.9	Setup Networking	24
5.10	Client Installations	26
5.10.1	Network boot the client nodes	26
5.10.2	Check completion status of nodes	26
5.10.3	Reboot the client nodes	27
5.11	Complete the Cluster Setup	27
5.12	Test Cluster Setup	27
5.13	Congratulations!	28
5.14	Adding and Deleting client nodes	28
5.14.1	Adding OSCAR clients	28
5.14.2	Deleting clients	28
5.15	Install/Uninstall OSCAR Packages	29
5.15.1	Selecting the Right Package	29

5.15.2	Single Image Restriction	29
5.15.3	Failures to Install and Uninstall	30
5.15.4	More Debugging Info	30
5.16	Starting over – installing OSCAR again	30
6	Package-Specific Installation Notes	31
6.1	Disabling Services	31
6.2	HDF5	31
6.3	Selecting a Different Kernel with <code>kernel_picker</code>	32
6.4	Networking Package	32
6.5	Managing NTP for the OSCAR Server and Clients	33
6.5.1	Configuring NTP	33
6.5.2	Enabling/Disabling the NTP Service	34
6.6	The OSCAR Password Installer and User Management (OPIUM)	34
6.6.1	SSH Configuration	34
6.7	The Portable Batch System (PBS) and Maui Scheduler	34
6.7.1	Configuring PBS	35
6.7.2	PBS Resources	35
6.7.3	An FAQueue	35
6.8	Packet Filtering with <code>pfilter</code>	36
6.9	Managing machines and images in SIS	37
6.9.1	Multicast Installs with <code>SystemImager</code>	37
6.10	Which MPI do you want to use?	37
6.10.1	Setting the system-level default	37
6.10.2	Setting the user-level default	38
6.10.3	Use <code>switcher</code> with care!	39
A	Network Booting Client Nodes	40
B	What Happens During Client Installation	40
C	Installing without a DHCP server	41
C.1	Boot Floppies	41
C.2	Differences from Standard OSCAR Install	42
C.2.1	Cluster Definition	42
C.2.2	Adding and Deleting client nodes	42
D	Security	42
D.1	Security layers	42
D.2	Router packet filtering	42
D.3	Network stack protections	42
D.4	Host based packet filtering	42
D.5	TCP Wrappers	43
D.6	Service paring	43
D.7	Service configuration	43
D.8	Secure communications	44

E	Screen-by-Screen Walkthrough	45
E.1	Running <code>install_cluster</code>	45
E.2	Download Additional OSCAR Packages...	45
E.3	Select OSCAR Packages to Install	45
E.4	Configure Selected OSCAR Packages	45
E.5	Install OSCAR Server Packages	46
E.6	Build OSCAR Client Image	46
E.7	Define OSCAR Clients	53
E.8	Setup Networking	53
E.9	Step 5: Complete Cluster Setup	59
E.10	Test Cluster Setup	59
E.11	Delete Node Button	60
E.12	Add Node Button	60
E.13	Install/Uninstall OSCAR Packages	60

List of Tables

1	OSCAR supported distributions	8
2	OSCAR distribution package file and directory layout.	15

List of Figures

1	OSCAR Wizard.	18
2	OSCAR package selection.	20
3	OSCAR package configuration.	20
4	Build the image.	22
5	Define the Clients.	24
6	Collect client MAC addresses.	26
7	Setup cluster tests	27
8	Adding OSCAR clients.	28
9	Deleting OSCAR clients.	29
10	Unpacking OSCAR.	45
11	Configure & Install the OSCAR Toolkit.	46
12	Running the <code>install_cluster</code> script.	47
13	Enter an OSCAR Database (ODA) password.	47
14	The OSCAR Installation Wizard.	48
15	Downloading additional OSCAR packages.	49
16	Adding additional OPD Repositories.	50
17	Selecting which OSCAR packages to install.	50
18	Configuring selected OSCAR packages. For example, the Environment Switcher package has configuration options.	51
19	Successfully installed the OSCAR server packages.	51
20	Building the OSCAR client image.	52
21	Defining the OSCAR clients.	53
22	Setup networking: initial window.	54
23	Setup networking: building an autoinstall floppy.	54
24	Setup networking: booting the client.	55
25	Setup networking: scan for booting client (DHCP request).	55
26	Setup networking: client is broadcasting, allows capture of MAC address.	56
27	Setup networking: Assigning MACs to IPs	56
28	Booting the client a second time to download the image.	57
29	Client partitioning disk, setting up the disk tables, and starting to download the image.	57
30	Client downloading and installing the image.	58
31	A client has finished the install and is asking to be rebooted.	58
32	Complete Cluster Setup.	59
33	Test the cluster.	59
34	Delete OSCAR Clients.	60
35	Add OSCAR Clients.	60
36	Add node / defining the clients.	61
37	Add node / setup networking.	61
38	Add node / complete cluster setup.	62
39	Install/Uninstall OSCAR Packages.	62

40 Install/Uninstall OSCAR Package selector sample colors. 63

1 Introduction

The OSCAR (Open Source Cluster Application Resource) software package is intended to simplify the complex tasks required to install a cluster. While the usual intended use for OSCAR clusters is for high-performance computing (HPC), OSCAR clusters can be used for any cluster-enabled kinds of applications. Note that since OSCAR is aimed towards HPC, several HPC-related packages are installed by default, such as popular MPI implementations, PVM, PBS, etc.

This document provides a step-by-step installation guide for system administrators, as well as a detailed explanation of what is happening as you install. Note that this installation guide is specific to OSCAR version 3.0.

1.1 Latest Documentation

Please be sure that you have the latest version of this document. It is possible (and probable!) that newer versions of this document were released on the main OSCAR web site after the software was released. You are *strongly* encouraged to check <http://oscar.sourceforge.net/> for the latest version of these instructions before proceeding. Document versions can be compared by checking their version number and date on the cover page.

1.2 Terminology

A common term used in this document is *cluster*, which refers to a group of individual computers bundled together using hardware and software in order to make them work as a single machine.

Each individual machine of a cluster is referred to as a *node*. Within the OSCAR cluster to be installed, there are two types of nodes: *server* and *client*. A *server* node is responsible for servicing the requests of *client* nodes. A *client* node is dedicated to computation.

An OSCAR cluster consists of one server node and one or more client nodes, where all the client nodes [currently] must have homogeneous hardware. The software contained within OSCAR does support doing multiple cluster installs from the same server, but no documentation is provided on how to do so.

An *OSCAR package* is a set of files that is used to install a software package in an OSCAR cluster. An OSCAR package can be as simple as a single RPM file, or it can be more complex, perhaps including a mixture of RPM and other auxiliary configuration / installation files. OSCAR packages provide the majority of functionality in OSCAR clusters.

OSCAR packages fall into one of three categories:

- *Core packages* are required for the operation of OSCAR itself (mostly involved with the installer).
- *Included packages* are shipped in the official OSCAR distribution. These are usually authored and/or packaged by OSCAR developers, and have some degree of official testing before release.
- *Third party packages* are not included in the official OSCAR distribution; they are “add-ons” that can be unpacked in the OSCAR tree, and therefore installed using the OSCAR installation framework.

1.3 Supported Distributions

OSCAR has been tested to work with several distributions. Table 1 lists each distribution and version and specifies the level of support for each. In order to ensure a successful installation, most users should stick to a distribution that is listed as *Fully supported*.

Distribution and Release	Status
RedHat 9.0	Fully supported
RedHat 8.0	Fully supported
Mandrake 9.0	Fully supported

Table 1: OSCAR supported distributions

1.4 Minimum System Requirements

The following is a list of minimum system requirements for the OSCAR server node:

- CPU of i586 or above
- A network interface card that supports a TCP/IP stack
- If your OSCAR server node is going to be the router between a public network and the cluster nodes, you will need a second network interface card that supports a TCP/IP stack
- At least 4GB total free space – 2GB under / and 2GB under /var
- An installed version of Linux, preferably a *Fully supported* distribution from Table 1

The following is a list of minimum system requirements for the OSCAR client nodes:

- CPU of i586 or above
- A disk on each client node, at least 2GB in size (OSCAR will format the disks during the installation)
- A network interface card that supports a TCP/IP stack¹
- Same Linux distribution and version as the server node
- All clients must have the same architecture (e.g., ia32 vs. ia64)
- Monitors and keyboards may be helpful, but are not required
- Floppy or PXE enabled BIOS

1.5 Document Organization

Due to the complicated nature of putting together a high-performance cluster, it is strongly suggested that even experienced administrators read this document through, without skipping any sections, and then use the detailed installation procedure to install your OSCAR cluster. Novice users will be comforted to know that anyone who has installed and used Linux can successfully navigate through the OSCAR cluster install.

The rest of this document is organized as follows. First, Section 2 tells how to obtain an OSCAR version 3.0 distribution package. Next, the “Release Notes” section (Section 3) that applies to OSCAR version 3.0 contains some requirements and update issues that need to be resolved before the install. Section 4 provides an overview for the System Installation Suite software package used in OSCAR to perform the bulk of the

¹Beware of certain models of 3COM cards – not all models of 3COM cards are supported by the installation Linux kernel that is shipped with OSCAR. See the OSCAR web site for more information.

cluster installation. Section 5 details the cluster installation procedure (the level of detail lies somewhere between “the install will now update some files” and “the install will now replace the string ‘xyz’ with ‘abc’ in file `some_file`.”)

Finally, Section 6 contains system administration notes about several of the individual packages that are installed by OSCAR. *This section is a “must read” for all OSCAR system administrators.*

Appendix A covers the topic of network booting client nodes, which is so important that it deserved its own section. Appendix B provides curious users an overview of what really happens during a client install. Appendix C discusses how to install an OSCAR cluster without a DHCP server. Appendix D covers a primer of some security aspects of a Linux cluster. Although not intended to be a comprehensive description of cluster security, it is a good overview for those who know relatively little about system administration and security. Finally, Appendix E is a screen-by-screen walk through of a typical OSCAR installation.

More information is available on the OSCAR web site and archives of the various OSCAR mailing lists. If you have a question that cannot be answered by this document (including answers to common installation problems), be sure to visit the OSCAR web site:

<http://oscar.sourceforge.net/>

2 Downloading an OSCAR Distribution Package

The OSCAR distribution packages can be downloaded from:

<http://oscar.sourceforge.net/>

Note that there are actually three flavors of distribution packages for OSCAR, depending on your bandwidth and installation/development needs:

1. “Regular”: All the OSCAR installation material that most users need to install and operate an OSCAR cluster.
2. “Extra Crispy”: Same as Regular, except that the SRPMs for most of the RPMs in OSCAR are also included. SRPMs are *not* required for installation or normal operation of an OSCAR cluster. This distribution package is significantly larger than Regular, and is not necessary for most users – only those who are interested in source RPMs need the Extra Crispy distribution. The SRPMs can be found under `packages/*/SRPMS/` directories.
3. “Secret Sauce”: This distribution contains *only* the SRPMs for RPMs in OSCAR in the Regular and Extra Crispy distributions (it’s essentially (Extra Crispy - Regular)). It is intended only for those who initially downloaded the Regular distribution and later decided that they wanted the SRPMs as well. The Secret Sauce distribution is intended to be expanded over your Regular installation – it will create `packages/*/SRPMS/` directories and populate them with the relevant `.src.rpm` files.

All three distributions can be downloaded from the main OSCAR web page.

3 Release Notes

The following release notes apply to OSCAR version 3.0.

3.1 Notes for All Systems

- Each package in OSCAR has its own installation and release notes. See Section 6 for additional release notes.
- All nodes must have a hostname other than “localhost” that does not contain any underscores (“_”).
- A domain name must be specified for the client nodes when defining them.
- Due to some distribution portability issues, OSCAR currently installs a “compatibility” (`python2-compat-1.0-1`) RPM to resolve the Python2 prerequisite that is slightly different across different Linux distributions. Also see the file `packages/c3/RPMS/NOTE.python2`.
- In some cases, the test window that is opened from the OSCAR wizard may close suddenly when there is a test failure. If this happens, run the test script, `testing/test_cluster`, manually in a shell window to diagnose the problem.
- Although OSCAR can be installed on pre-existing server nodes, it is typically easiest to use a machine that has a new, fresh install of a distribution listed in Table 1 *with no updates installed*. If the updates are installed, there may be conflicts in RPM requirements. It is recommended to install RedHat updates *after* the initial OSCAR installation has completed. On the Mandrake systems the security updates must be added prior to the install.
- The following benign warning messages will appear multiple times during the OSCAR installation process:

```
rsync_stub_dir: no such variable at ...  
  
Use of uninitialized value in pattern match (m//) at  
/usr/lib/perl5/site_perl/oda.pm ...
```

It is safe to ignore these messages.

- The OSCAR installer will install the MySQL package on the server node if it is not already installed. You will be prompted to enter a password to be used for the MySQL database.
- The OSCAR installer GUI provides little protection for user mistakes. If the user executes steps out of order, or provides erroneous input, Bad Things may happen. Users are strongly encouraged to closely follow the instructions provided in this document.
- The OSCAR installer GUI currently does not support deleting a node and adding the same node back *in the same session*. If you wish to delete a node and then add it back, you must delete the node, close the OSCAR installer GUI, launch the OSCAR installer GUI again, and then add the node.
- During the **<Build OSCAR Client Image>** step, the “Successfully created image” notice will appear even though the status bar looks incomplete. This incomplete status bar can be ignored.
- If `ssh` produces warnings when logging into the compute nodes from the OSCAR head node, the C3 tools (e.g., `cexec`) may experience difficulties. For example, if you use `ssh` to login in to the OSCAR head node from a terminal that does not support X windows and then try to run `cexec`, you might see a warning message in the `cexec` output:

```
Warning: No xauth data; using fake authentication data for
X11 forwarding.
```

Although this is only a warning message from `ssh`, `cexec` may interpret it as a fatal error, and not run across all cluster nodes properly (e.g., the `<Install/Uninstall Packages>` button will likely not work properly).

Note that this is actually an `ssh` problem, not a C3 problem. As such, you need to eliminate any warning messages from `ssh` (more specifically, eliminate any output from `stderr`). In the example above, you can tell the C3 tools to use the “-x” switch to `ssh` in order to disable X forwarding:

```
# export C3_RSH=ssh -x
# cexec uptime
```

The warnings about `xauth` should no longer appear (and the `<Install/Uninstall Packages>` button should work properly).

- The `<Cancel>` button in the `<Install/Uninstall Package>` step does not work properly; if any packages are selected to be installed or uninstalled, clicking the `<Cancel>` button still triggers the execution of the package installer/uninstaller. This will be fixed in a future release. The same behavior occurs if you close the window via the window manager’s “close” functionality.

Note that if you do not select any additional packages to install/uninstall, nothing will run (as expected).

3.2 Red Hat 8/9.0 Notes

There are a few issues that may crop up when using OSCAR on Red Hat 9.0 and/or Red Hat 8.0 clusters. The following items highlight these issues.

- Deselecting `Pfilter` causes the image creation to fail. This is due to a dependency with `IPtables` and when `Pfilter` is not selected the `IPtables` RPM is not listed in the node (image) `rpmlist`. The simple fix is to add “`iptables`” to the Red Hat 9.0 `rpmlist` if you are not installing `Pfilter` on the compute nodes.
- The RPM system has been updated with this Red Hat release. The OSCAR install process will likely display several warnings due to unsigned RPMS. These warnings can be ignored.
- In some OSCAR pre-release testing, RPM would hang during the building of a client image (Section 5.7). This is a documented bug in the version of RPM that ships with Redhat 8 and 9.0; it is not a problem with OSCAR. The procedure that was used to remedy this situation is outlined below (excerpts taken from <http://www.rpm.org/hintskinks/repairdb-2003-06/>):
 - If RPM hangs at any point (e.g., building the client image) – first ensure that it really has hung and just isn’t taking a long, long time to complete. Typical indications that it has genuinely hung include: the disk is not running and load goes down to 0 (or nearly 0) and stays there.
 - Then do a `ps` and find the PID of the `rpm` process:

```
# ps -eaf | grep rpm | grep -v grep
...output...
# kill <PID_of_RPM>
```

- This will probably not kill the process (it's likely to be in a state where it is ignoring signals), but it should be tried anyway – this would allow `rpm` to exit cleanly. If `rpm` does exit cleanly, jump down to the last step in this procedure.
- If `rpm` does not exit within a short period of time, `kill -9 <PID_OF_RPM>`. This guarantees that `rpm` will not exit cleanly, but in this case, it's ok. Now, do the following:
 1. Save a copy of the RPM database (just to be safe):


```
# cd /var/lib
# tar zcvf /tmp/rpmdb.tar.gz rpm
```
 2. Delete any existing RPM database locks:


```
# cd /var/lib/rpm
# rm -f __db*
```
 3. Rebuild the RPM database:


```
# rpm -vv --rebuilddb
```
- Now re-run the OSCAR step that hung. If RPM hangs again, repeat these steps to un-hang it. Testing has shown that it may be necessary to repeat these steps multiple times in order to get a successful RPM run.

3.3 Mandrake 9.0 Notes

The following may need to be run before attempting to install OSCAR on a Mandrake 9.0 cluster. If there are problems during the server preparation related to C3 or Python2 this may solve the issue.

- Install the Python 2 compatibility RPM. This RPM can be found in the OSCAR distribution package, under the `packages/c3/RPMS` directory:

```
# cd oscar-2.3/packages/c3/RPMS
# rpm -Uvh python2-compat-1.0-1.noarch.rpm
```

- The OSCAR wizard makes use of `xterm` so this package must be installed. You can check the RPM database for this package by typing `rpm -q xterm`. If `xterm` is not available, you must install the `xterm-165-3mdk.i586.rpm` RPM (you may need Mandrake CD's to run the `urpmi` command as mentioned above):

```
# urpmi xterm
```

- Lastly, `root`'s default shell configuration files hardcode the value for the `PATH` environment variable, regardless of what is added via `profile.d/` startup scripts. This effects various OSCAR installed components, which are installed into locations such as `/opt`. As such, the OSCAR installer appends the global path, e.g., `$PATH` or `${PATH}` depending on the shell, to the end of the path so OSCAR installed applications are in the search environment. This potentially makes changes to the following files:

- `/root/.bashrc`
- `/root/.cshrc`
- `/root/.tcshrc`

3.4 IA64 and Other Bleeding Edge Systems Notes

IA64 support was removed for the 3.0 release of OSCAR due to timing problems and the significant differences between current IA32 distributions and the freely available IA64 distributions. IA64 support is fully expected to return in the next release.

4 Overview of System Installation Suite (SIS)

The first question you may have is “what is SIS?” The System Installation Suite (SIS) is a cluster installation tool developed by the collaboration of the IBM Linux Technology Center and the SystemImager team. SIS was chosen to be the installation mechanism for OSCAR for multiple reasons:

- SIS is a high-quality, third party, open source product that works well in production environments
- SIS does not require the client nodes to already have Linux installed
- SIS maintains a database containing installation and configuration information about each node in the cluster
- SIS uses RPM as a standard for software installation
- SIS supports heterogenous hardware and software installation (although this feature is not [yet] used by OSCAR)

In order to understand some of the steps in the upcoming install, you will need knowledge of the main concepts used within SIS. The first concept is that of an *image*. In SIS, an *image* is defined for use by the cluster nodes. This image is a copy of the operating system files stored on the server. The client nodes install by replicating this image to their local disk partitions. Another important concept from SIS is the client definition. A SIS client is defined for each of your cluster nodes. These client definitions keep track of the pertinent information about each client. The server node is responsible for creating the cluster information database and for servicing client installation requests. The information that is stored for each client includes:

- IP information such as hostname, IP address, route.
- Image name.

Each of these pieces of information will be discussed further as part of the detailed install procedure.

For additional information on the concepts in SIS and how to use it, you should refer to the `SIS(1)` man page. In addition, you can visit the SIS web site at <http://sisuite.org/> for recent updates.

5 Detailed Cluster Installation Procedure

All actions specified below should be performed by the `root` user on the server node unless noted otherwise. Note that if you login as a regular user and use the `su` command to change to the `root` user, you *must* use “`su -`” to get the full `root` environment. Using “`su`” (with no arguments) may not be sufficient, and may cause obscure errors during an OSCAR installation.

Note that all the steps below are mandatory unless explicitly marked as optional.

5.1 Server Installation and Configuration

During this phase, you will prepare the machine to be used as the server node in the OSCAR cluster.

5.1.1 Install Linux on the server machine

If you have a machine you want to use that already has Linux installed, ensure that it meets the minimum requirements as listed in Section 1.4. If it does, you may skip ahead to Section 5.1.2.

It should be noted that OSCAR is only supported on the distributions listed in Table 1 (page 8). As such, use of distributions other than those listed will likely require some porting of OSCAR, as many of the scripts and software within OSCAR are dependent on those distributions.

When installing Linux, it is not necessary to perform a “custom” install since OSCAR will usually install all the software on which it depends. The main Linux installation requirement is that some X windowing environment such as GNOME or KDE must be installed. Typically, a “Workstation” install with the “Software Development Tools” group added yields a sufficient installation for OSCAR to install successfully.

It is best to *not* install distribution updates after you install Linux; doing so may disrupt some of OSCAR’s RPM dependencies. Instead, install OSCAR first, and then install the distribution updates.

5.1.2 Disk space and directory considerations

OSCAR has certain requirements for server disk space. Space will be needed to store the Linux RPMs and to store the images. The RPMs will be stored in `/tftpboot/rpm`. 2GB is usually enough to store the RPMs. The images are stored in `/var/lib/systemimager` and will need approximately 2GB per image. Although only one image is required for OSCAR, you may want to create more images in the future.

If you are installing a new server, it is suggested that you allow for 4GB in both the `/` (which contains `/tftpboot`) and `/var` filesystems when partitioning the disk on your server.

If you are using an existing server, you will need to verify that you have enough space on the disk partitions. Again 4GB of free space is recommended under each of `/` and `/var`.

You can check the amount of free space on your drive’s partitions by issuing the command `df -h` in a terminal. The result for each file system is located below the Available column heading. If your root (`/`) partition has enough free space, enter the following command in a terminal:

```
# mkdir -p /tftpboot/rpm
```

If your root partition does not have enough free space, create the directories on a different partition that does have enough free space and create symbolic links to them from the root (`/`) directory. For example, if the partition containing `/usr` contains enough space, you could do so by using the following commands:

```
# mkdir -p /usr/tftpboot/rpm
# ln -s /usr/tftpboot /tftpboot
```

The same procedure should be repeated for the `/var/lib/systemimager` subdirectory.

5.1.3 Download a copy of OSCAR and unpack on the server

If you are reading this, you probably already have a copy of an OSCAR distribution package. If not, go to <http://oscar.sourceforge.net/> and download the latest OSCAR Regular or Extra Crispy distribution package (see Section 2, page 9). Ensure that you have the latest documentation (later documentation may be available on the OSCAR web site than in the OSCAR distribution package).

Place the OSCAR distribution package in a directory such as `root`’s home directory on the server node. Although there is no required installation directory (note that you may not use the directory `/usr/local/oscar`, `/opt/oscar`, `/var/lib/oscar`, or `/var/cache/oscar` – they are reserved for use by OSCAR),

the rest of these instructions will assume that you downloaded the OSCAR distribution package to `root`'s home directory.

Do **not** unpack the tarball on a Windows-based machine and copy the directories over to the server, as this will convert all the scripts to "DOS" format and will render them useless under Linux.

Open a command terminal and issue the following commands to unpack the OSCAR distribution package:

```
# cd
# tar zxf <filename>
```

Where `<filename>` is either `oscar-3.0.tar.gz` (regular distribution) or `oscar-including-srpms-3.0.tar.gz` (extra crispy distribution).

Directory	Contents
<code>~/oscar-3.0/</code>	the base OSCAR directory
<code>~/oscar-3.0/COPYING</code>	GNU General Public License v2
<code>~/oscar-3.0/dist</code>	distribution scripts for configure/install
<code>~/oscar-3.0/doc</code>	OSCAR documentation directory
<code>~/oscar-3.0/images</code>	auxiliary images used in the GUI
<code>~/oscar-3.0/install_cluster</code>	main installation script
<code>~/oscar-3.0/lib</code>	auxiliary library routines
<code>~/oscar-3.0/oscarsamples</code>	sample configuration files
<code>~/oscar-3.0/packages</code>	RPM and installation files for the OSCAR packages
<code>~/oscar-3.0/README</code>	text README document
<code>~/oscar-3.0/scripts</code>	contains scripts that do most of the work
<code>~/oscar-3.0/share</code>	more auxiliary helper files
<code>~/oscar-3.0/testing</code>	contains OSCAR cluster test helper scripts
<code>~/oscar-3.0/VERSION</code>	file containing the OSCAR version number

Table 2: OSCAR distribution package file and directory layout.

5.1.4 Configure and Install OSCAR

Starting with OSCAR 2.3, after unpacking the tarball you will need to configure and install OSCAR. The configure portion sets up the release to be permanently installed on the system in `/opt/oscar` (default location). The `--prefix=ALT-DIR` flag can be used to configure and install OSCAR into an alternate directory.

The first step is to run the `configure` script, which is provided in the top-level directory of the OSCAR release.

```
# cd /root/oscar-3.0
# ./configure
```

Once the configure script successfully completes you are ready to actually install OSCAR on the server. At this point no changes have been made to the system beyond the `oscar-3.0` directory itself. Running the following will copy the files to the system. The files copied will include the base OSCAR toolkit as well as startup scripts for the `profile.d/` area. The startup scripts add OSCAR to your path and set environment variables like `OSCAR_HOME`.

```
# make install
```

At this point you will be ready to change to either the default `/opt/oscar` directory or whatever path was used with the `--prefix` flag to perform the installation steps discussed in Section 5.2. In the remainder of this document, the variable `$OSCAR_HOME` will be used in place of the directory you installed OSCAR to – by default this is `/opt/oscar`.

5.1.5 Configure the ethernet adapter for the cluster

Assuming you want your server to be connected to both a public network and the private cluster subnet, you will need to have two ethernet adapters installed in the server. This is the preferred OSCAR configuration because exposing your cluster may be a security risk and certain software used in OSCAR (such as DHCP) may conflict with your external network.

Once both adapters have been physically installed in the server node, you need to configure them.² Any network configurator is sufficient; popular applications include `neat`, `netcfg`, or a text editor.

The following major requirements need to be satisfied:

Hostname. Most Linux distributions default to the hostname “localhost” (or “localhost.localdomain”). This must be changed in order to successfully install OSCAR – choose another name that does not include any underscores (“_”).

Public adapter. This is the adapter that connects the server node to a public network. Although it is not required to have such an adapter, if you do have one, you must configure it as appropriate for the public network (you may need to consult with your network administrator).

Private adapter. This is the adapter connected to the TCP/IP network with the rest of the cluster nodes. This adapter must be configured as follows:

- Use a private IP address³
- Use an appropriate netmask⁴
- Ensure that the interface is activated at boot time
- Set the interface control protocol to “none”

Now reboot the server node to ensure that all the changes are propagated to the appropriate configuration files. To confirm that all ethernet adapters are in the “up” state, once the machine has rebooted, open another terminal window and enter the following command:

```
# ifconfig -a
```

You should see UP as the first word on the third line of output for each adapter. If not, there is a problem that you need to resolve before continuing. Typically, the problem is that the wrong module is specified for the given device. Try using the network configuration utility again to resolve the problem.

²Beware of certain models of 3COM network cards. See footnote 1 on page 8.

³There are three private IP address ranges: 10.0.0.0 to 10.255.255.255; 172.16.0.0 to 172.32.255.255; and 192.168.0.0 to 192.168.255.255. Additional information on private intranets is available in RFC 1918. You should not use the IP addresses 10.0.0.0 or 172.16.0.0 or 192.168.0.0 for the server. If you use one of these addresses, the network installs of the client nodes will fail.

⁴A class C netmask of 255.255.255.0 should be sufficient for most OSCAR clusters.

5.1.6 Copy distribution installation RPMs to `/tftpboot/rpm`

In this step, you need to copy the RPMs included with your Linux distribution into the `/tftpboot/rpm` directory. When each CD is inserted, Linux usually automatically makes the contents of the CD be available in the `/mnt/cdrom` directory (you may need to execute a command such as “`mount /mnt/cdrom`” if the CD does not mount automatically).

For each CD, locate the directory that contains the RPMs. In RedHat 7.x, the RPMs are located in the `RedHat/RPMS` directory (i.e., `/mnt/cdrom/RedHat/RPMS`). Copy the RPMS into the `/tftpboot/rpm` directory with a command such as ⁵:

```
# cp /mnt/cdrom/RedHat/RPMS/*.rpm /tftpboot/rpm
```

Be sure to repeat the above process for *all* CDs. After using each CD you will have to unmount it from the local file system and eject it by issuing these commands:

```
# cd
# eject cdrom
```

5.2 Launching the OSCAR Installer

Change directory to the top-level OSCAR directory and start the OSCAR install wizard:

```
# cd $OSCAR_HOME
# ./install_cluster <device>
```

In the above command, substitute the device name (e.g., *eth1*) in place of `<device>` for your server’s private network ethernet adapter. The script will execute some setup / configuration steps, including (but not limited to):

1. installs prerequisite packages on the server
2. copies OSCAR RPMs to `/tftpboot/rpm`
3. installs all OSCAR server RPMs
4. updates `/etc/hosts` with OSCAR aliases
5. updates `/etc/exports`
6. adds OSCAR paths to `/etc/profile`
7. updates system startup (`/etc/rc.d/init.d`) scripts
8. restarts affected services

A lot of output will be displayed in the console window where you invoked `install_cluster`. This reflects normal operational output from the various installation commands that OSCAR executes. The output is also saved in the file `oscarinstall.log` for later reference (particularly if something goes wrong during during the installation).

After the steps listed above have successfully finished, the OSCAR installation wizard GUI will automatically be launched.

⁵It may be helpful to use `rsync` for these mass copies to ensure the integrity of the data (at a cost of longer transfer times). For example replace the `cp` command in the example with `rsync -cav --stats`.

The wizard, as shown in Figure 1, is provided to guide you through the rest of the cluster installation. To use the wizard, you will complete a series of steps, with each step being initiated by the pressing of a button on the wizard. Do not go on to the next step until the instructions say to do so, as there are times when you may need to complete an action outside of the wizard before continuing on with the next step. For each step, there is also a <Help> button located directly to the right of the step button. When pressed, the <Help> button displays a message box describing the purpose of the step.

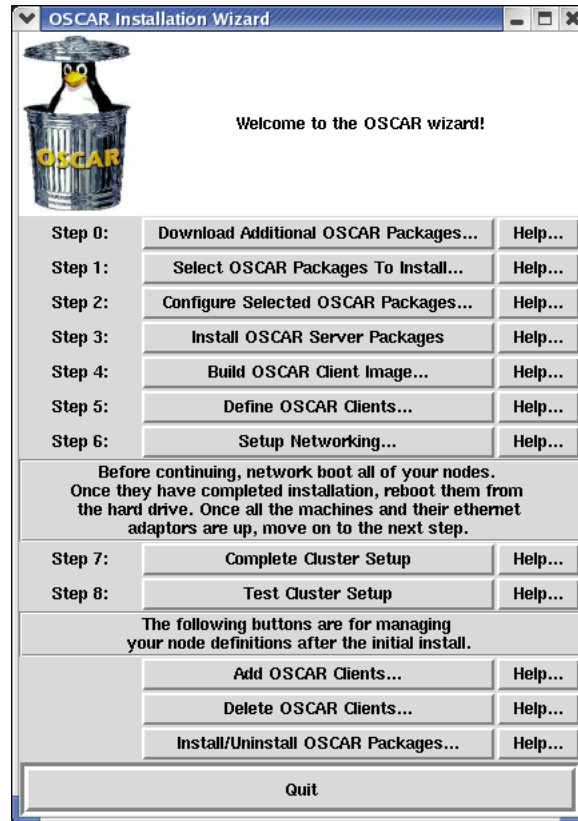


Figure 1: OSCAR Wizard.

5.3 Downloading Additional OSCAR Packages

Note: This step is optional.

The first step of the Wizard, “Step 0”, enables you to download additional packages.⁶ The OSCAR Package Downloader (OPD) is a unified method of downloading OSCAR packages and inserting them in the OSCAR installation hierarchy so that they will be installed during the main OSCAR installation process. The Wizard uses a GUI frontend to OPD affectionately known as *OPDer*. The addition of this frontend limits the need for direct access to the command-line OPD tool directly.⁷ If you would like to add additional repository URLs for testing purposes, you can do this by accessing the <File> menu and

⁶Ganglia is a notable package that is not included in the main OSCAR distribution. The process mentioned in this section will enable you to download Ganglia and install it with the rest of the OSCAR components.

⁷Note, if using OPD directly, the packages must be downloaded before the <Select OSCAR Packages To Install> step (see Section 5.4).

then choosing **<Additional Repositories...>**. The remainder of this sub-section describes the underlying OPD tool.

The command-line OPD can be executed outside of the GUI Wizard from the top-level OSCAR directory with the following command:

```
# cd $OSCAR_HOME
# ./scripts/opd
```

OPD can use either the `wget` command or the built-in Perl LWP for downloading files. By default, `wget` will be used if it can be found. However, the use of LWP may be forced if the argument “`--lwp`” is given on the command line. LWP may be used in order to utilize a proxy, for example. See the `LWP::UserAgent` documentation for details on how to use proxies through LWP.

OPD requires some Perl modules to be installed before running. If they are not installed, you will receive a detailed error message listing which modules need to be installed. You have two options to install these modules:

1. Use a tool such as CPAN to download and install the required modules.
2. Launch the OSCAR installer (see Section 5.2) and then quit immediately when the GUI window appears. This is the preferred method, since the OSCAR installer will install of its own prerequisites (i.e., the process is automated).

Upon launching OPD, select a repository for packages from the choices listed. OPD will then connect to that repository and query it for a list of available packages. Select all the packages that you would like to download, and then use the “download” command to actually download them.

Use the “help” command in OPD for listings of additional commands and help.

5.4 Selecting Packages to Install

Note: This step is optional.

If you wish to change the list of packages that are installed, click on the **<Select OSCAR Packages To Install>** button. This step is optional – by default all packages directly included in OSCAR are selected and installed. However, if you downloaded any additional packages, e.g., via OPD/OPDer, they will not be selected for installation by default. Therefore you will need to click this button and select the appropriate OSCAR Packages to install on the cluster.

When you click on the button, a window similar to the one shown in Figure 2 appears. Each of the packages that the OSCAR installer has found are listed in the main frame. Core packages must be installed and cannot be unselected. Included packages can be unselected if desired.

Note that this window only shows *OSCAR packages* – it does not show individual RPMs.

Once you have a selected a set of OSCAR packages to install, click on the **<Exit>** button to save your selections and return to the main OSCAR window. Note that closing the window yields the same result and there is no way of ‘defaulting’ to the original settings, so make sure your package list is complete before proceeding to the next step.

5.5 Configuring OSCAR Packages

Note: This step is optional.

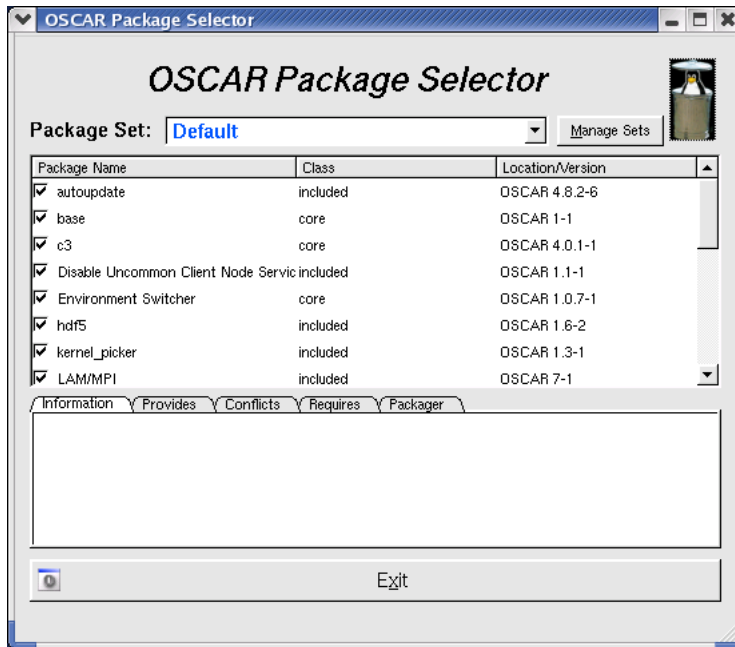


Figure 2: OSCAR package selection.

Some OSCAR packages allow themselves to be configured. Clicking on the **<Configure Selected OSCAR Packages>** button will bring up a window listing all the packages that can be configured. Figure 3 shows a sample with only the Environment Switcher package listed.

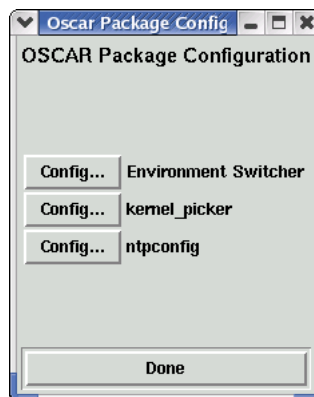


Figure 3: OSCAR package configuration.

Clicking on any of the packages' **<Config>** button will bring up a panel for configuring that package. Select whatever options are appropriate for that package, and then click on the **<Save>** button to save your selections, or the **<Cancel>** button to cancel all of your selections and leave the original settings. If you have saved your changes but want to go back to the default settings, simply click on the **<Default Configuration>** button and then the **<Save>** button to revert to the original settings.

This step is optional. If you do not click on the **<Configure Selected OSCAR Packages>** button, defaults for all packages will be used.

5.5.1 Selecting a Default MPI Implementation

Although multiple MPI implementations can be installed, only one can be “active” for each user at a time. Specifically, each user’s path needs to be set to refer to a “default” MPI that will be used for all commands. The Environment Switcher package provides a convenient mechanism for switching between multiple MPI implementations. Section 6.10 contains more details about this package (page 37).

The Environment Switcher package is mentioned now, however, because its configuration panel allows you to select which MPI implementation will be the initial “default” for all users. OSCAR currently includes two MPI implementations: LAM/MPI and MPICH. Using Environment Switcher’s configuration panel, you can select one of these two to be the cluster’s default MPI.

You can change this default setting later – see Section 6.10 for more details.

When you close the main Configuration window, the following benign warning may appear in the shell window (it is safe to ignore):

```
Tag "mpi" does not seem to exist yet.  Skipping.
```

5.6 Install OSCAR Server Packages

Press the <Install OSCAR Server Packages> button. This will invoke the installation of various RPMs and auxiliary configuration on the server node. Execution may take several minutes; text output and status messages will appear in the shell window.

A popup will appear indicating the success or failure of this step. Click on the <Close> button to dismiss it.

5.7 Build OSCAR Client Image

Before pressing the <Build OSCAR Client Image>, ensure that the following conditions on the server are true:

- Ensure that the SSH daemon’s configuration file (`/etc/ssh/sshd_config`) on the headnode has `PermitRootLogin` set to `yes`. After the OSCAR installation, you may set this back to `no` (if you want), but it needs to be `yes` during the install because the config file is copied to the client nodes, and `root` *must* be able to login to the client nodes remotely.
- By the same token, ensure that TCP wrappers settings are not “too tight”. The `/etc/hosts.allow` and `/etc/hosts.deny` files should allow all traffic from the entire private subnet.
- Also, beware of firewall software that restricts traffic in the private subnet.

If these conditions are not met, the installation may fail during this step or later steps.

Press the <Build OSCAR Client Image> button. A dialog will be displayed. In most cases, the defaults will be sufficient. You should verify that the disk partition file is the proper type for your client nodes. The sample files have the disk type as the last part of the filename. You may also want to change the post installation action and the IP assignment methods. **It is important to note that if you wish to use automatic reboot, you should make sure the BIOS on each client is set to boot from the local hard drive before attempting a network boot by default. If you have to change the boot order to do a network boot before a disk boot to install your client machines, you should not use automatic reboot.**

Building the image may take several minutes; the red progress bar on the bottom of the window will indicate how far along the process is.

There is a lot of output in the console window during the build. It is normal to see some warning messages in the console. You can safely ignore these messages and wait for the final popup window announcing the success or failure of the overall image build.

A sample dialog is shown in Figure 4.

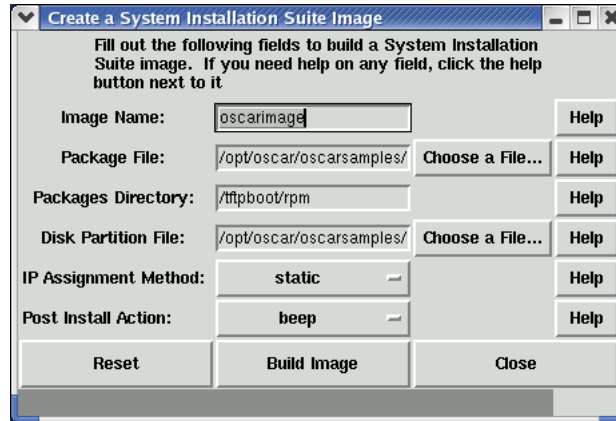


Figure 4: Build the image.

Customizing your image. The defaults of this panel use the sample disk partition and RPM package files that can be found in the `oscarsamples` directory. You may want to customize these files to make the image suit your particular requirements.

Disk partitioning. The disk partition file contains a line for each partition desired, where each line is in the following format:

```
<partition> <size in megabytes> <type> <mount point> <options>
```

Here is a sample (for a SCSI disk):

```
/dev/sda1      24          ext2      /boot    defaults
/dev/sda5     128         swap
/dev/sda6     *           ext2     /         defaults
nfs_oscar:/home -          nfs      /home    rw
```

An `*` in the size column causes that partition to grow to fill the entire disk. You can create your own partition files, but make sure that you do not exceed the physical capacity of your client hardware. Also be careful to not specify duplicate filesystems as this will cause problems during the installation. The sample listed above, and some others, are in the `oscarsamples` directory.

Package lists. The package list is simply a list of RPM file names (one per line). Be sure to include all prerequisites that any packages you might add. You do not need to specify the version, architecture, or extension of the RPM filename. For example, `bash-2.05-8.i386.rpm` need only be listed as “`bash`”.

Custom kernels. If you want to use a customized kernel, you can add it to the image after it is built (after installing the server OSCAR packages, but before building the client image). See the `kernel_picker` application description in Section 6.3 on page 32.

Build the Image. Once you are satisfied with the input, click the **<Build Image>** button. When the image completes, a popup window will appear indicating whether the build succeeded or failed. If successful, click the **<Close>** button to close the popup, and then press the **<Close>** button on the build image window. You will be back at the main OSCAR wizard menu.

If the build fails, look through the console output⁸ for some indication as to what happened to cause the failure. Common causes include: prerequisite failure, ran out of disk space, and missing package files. Also see the Release Notes for this version of OSCAR in Section 3 (page 9).

5.8 Define OSCAR Clients

Press the **<Define OSCAR Clients>** button. In the dialog box that is displayed, enter the appropriate information. Although the defaults will be sufficient for most cases, you will need to enter a value in the *Number of Hosts* field to specify how many clients you want to create.

1. The *Image Name* field should specify the image name that was used to create the image in the previous step.
2. The *Domain Name* field should be used to specify the client's IP domain name. It should contain the server node's domain (if it has one); if the server does not have a domain name, the default name `oscardomain` will be put in the field (although you may change it). **This field must have a value – it cannot be blank.** Note that especially for compute nodes on a private network, the domain name does not necessarily matter much. The domain name supplied in this field is used to form the fully-qualified name of each host in the OSCAR cluster. For example: `oscardnode1.oscardomain`, `oscardnode2.oscardomain`, etc. If your compute nodes are on a public network, you may want to use the “real” domain name that is part of their fully-qualified domain names.
3. The *Base name* field is used to specify the first part of the client name and hostname. It will have an index appended to the end of it. This name *cannot* contain an underscore character “_”.
4. The *Number of Hosts* field specifies how many clients to create. **This number must be greater than 0.**
5. The *Starting Number* specifies the index to append to the *Base Name* to derive the first client name. It will be incremented for each subsequent client.
6. The *Padding* specifies the number of digits to pad the client names, e.g., 3 digits would yield `oscardnode001`. The default is 0 to have no padding between base name and number (index).
7. The *Starting IP* specifies the IP address of the first client. It will be incremented for each subsequent client. See Footnote 3 on page 16 for more information on how to pick a starting IP address.

Clients will be given IP addresses starting with this IP address, and incrementing by 1 for each successive client. Ensure that the range of $[starting_ip, (starting_ip + num_clients)]$ does not conflict with the IP addresses of any other nodes on your network.

⁸Note that all console output is also spooled into the file `oscardinstall.log`.

IMPORTANT NOTE: Be sure that the resulting range of IP addresses does *not* include typical broadcast addresses such as *X.Y.Z.255*! If you have more hosts than will fit in a single address range, see the note at the end of this section about how to make multiple IP address ranges.

8. The *Subnet Mask* specifies the IP netmask for all clients. See Footnote 4 on page 16 for more information on how to select a netmask for your cluster.
9. The *Default Gateway* specifies the default route for all clients.

When finished entering information, press the **<Addclients>** button. When those clients have been created in the database, a popup will appear indicating the completion status. A sample dialog is shown in Figure 5.

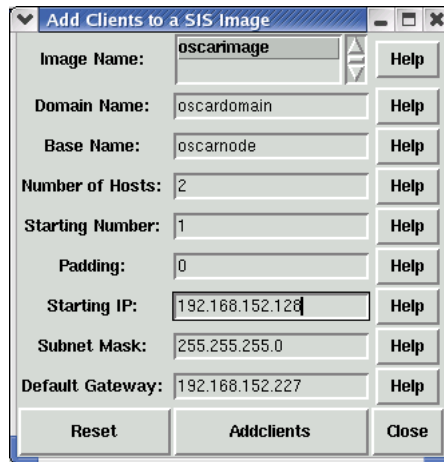


Image Name:	oscarimage	Help
Domain Name:	oscardomain	Help
Base Name:	oscardnode	Help
Number of Hosts:	2	Help
Starting Number:	1	Help
Padding:	0	Help
Starting IP:	192.168.152.128	Help
Subnet Mask:	255.255.255.0	Help
Default Gateway:	192.168.152.227	Help
Reset		Addclients
		Close

Figure 5: Define the Clients.

Note that this step can be executed multiple times. The GUI panel that is presented has limited flexibility in IP address numbering – the starting IP address will only increment the least significant byte by one for each successive client. Hence, if you need to define more than 254 clients (beware of broadcast addresses!), you will need to run this step multiple times and change the starting IP address. There is no need to close the panel and return to the main OSCAR menu before executing it again; simply edit the information and click on the **<Addclients>** button as many times as is required.

Additionally, you can run this step multiple times to use more structured IP addressing schemes. With a larger cluster, for example, it may be desirable to assign IP addresses based on the top-level switch that they are connected to. For example, the 32 clients connected to switch 1 should have an address of the form 192.168.1.x. The next 32 clients will be connected to switch 2, and should therefore have an address of the form 192.168.2.x. And so on.

After all clients have been created, you may press the **<Close>** button in the build clients dialog and continue with the next step.

5.9 Setup Networking

The MAC address of a client is a twelve hex-digit hardware address embedded in the client's ethernet adapter. For example, "00:0A:CC:01:02:03", as opposed to the familiar format of IP addresses. These

MAC addresses uniquely identify client machines on a network before they are assigned IP addresses. DHCP uses the MAC address to assign IP addresses to the clients.

In order to collect the MAC addresses, press the **<Setup Networking>** button. The OSCAR network utility dialog box will be displayed. To use this tool, you will need to know how to network boot your client nodes, or have a file that lists all the MACs from your cluster. For instructions on doing network booting, see Appendix A. A sample dialog is shown in Figure 6.

If you need to collect the MACs in your cluster, start the collection by pressing the **<Collect MAC Address>** button and then network boot the first client. As the clients boot up, their MAC addresses will show up in the left hand window. You have multiple options for assigning MACs to nodes; you can either:

- manually select MAC address and the appropriate client in the right side window. Click **<Assign MAC to Node>** to associate that MAC address with that node.
- click **<Assign all MACs>** button to assign all the MACs in the left hand window to all the open nodes in the right hand window.

Some notes that are relevant to collecting MAC addresses from the network:

- The **<Dynamic DHCP Update>** checkbox at the bottom right of the window controls refreshing the DHCP server. If it is selected (the default), the DHCP server configuration will be refreshed each time a MAC is assigned to a node. Note that if the DHCP reconfiguration takes place quick enough, you may not need to reboot the nodes a second time (i.e., if the DHCP server answers the request quick enough, the node may start downloading its image immediately).

If this option is off, you will need to click the **<Configure DHCP Server>** (at least once) to give it the associations between MACs and IP addresses.

- To remove extraneous MAC addresses from the left hand window (e.g., if the collector finds MACs that are not part of your cluster), select the address and click on the **<Remove>** button. Or click on the **<Remove All>** button to remove all of them.
- At any time, you may click on the **<Export MACs to file...>** button to save the MAC address list to a file. If you need to re-run the OSCAR installation, you can later click on **<Import MACs from file...>** to import this file rather than re-collecting all the MACs.
- When you have collected all of the MAC addresses, click the **<Stop Collecting MACs>** button. If you do not have **<Dynamic DHCP update>** selected, you need to click the **<Configure DHCP Server>** button to configure the DHCP server.
- You *must* click on the **<Stop Collecting MACs>** before closing the MAC Address Collection window!

You may also configure your remote boot method from this panel. The **<Build Autoinstall Floppy>** button will build a boot floppy for client nodes that do not support PXE booting. The **<Setup Network Boot>** button will configure the server to answer PXE boot requests if your client hardware supports it. See Appendix A for more details.

If your network switch supports multicasting, there is a new feature in OSCAR 3.0 which uses multi-cast to push files to the clients. To enable this feature simply click on the **<Enable Multicasting>** checkbox.

Once this feature is enabled, rsync will not be used for file distribution but instead by a program called Flamethrower which is bundled with SystemImager.

Since this new feature is still in its early stage, we recommend only the adventurous to try it as it may not work on all networking gear. Some have reported favorable results with the Force10 and HP 4000M network switches.

Incase you cannot get this feature working and want to switch back to using rsync, simply go back to the <Setup Networking> menu, make sure that the <Enable Multicasting> checkbox is disabled (should be disabled by default) and then click on <Configure DHCP Server> - this should revert back to the default settings.

When you have collected the addresses for all your client nodes and completed the networks setup, press the <Close> button.

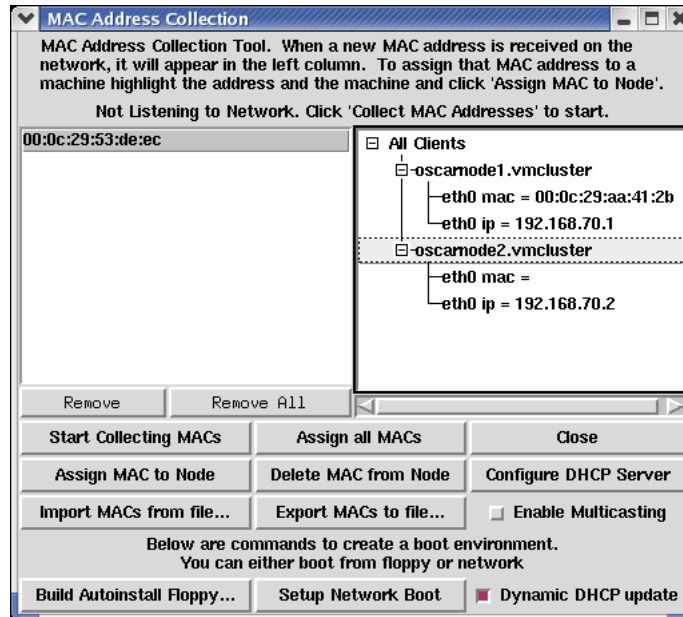


Figure 6: Collect client MAC addresses.

5.10 Client Installations

During this phase, you will network boot your client nodes and they will automatically be installed and configured. For a detailed explanation of what happens during client installation, see Appendix B.

5.10.1 Network boot the client nodes

See Appendix A for instructions on network booting clients.

Network boot all of your clients. As each machine boots, it will automatically start downloading and installing the OSCAR image from the server node.

5.10.2 Check completion status of nodes

After several minutes, the clients should complete the installation. You can watch the client consoles to monitor the progress. Depending on the Post Installation Action you selected when building the image, the clients will either halt, reboot, or beep incessantly when the installation is completed.

The time required for installation depends on the capabilities of your server, your clients, your network, and the number of simultaneous client installations. Generally, it should complete within several minutes.

5.10.3 Reboot the client nodes

After confirming that a client has completed its installation, you should reboot the node from its hard drive. If you chose to have your clients reboot after installation, they will do this on their own. If the clients are not set to reboot, you must manually reboot them. The filesystems will have been unmounted so it is safe to simply reset or power cycle them.

Note: If you had to change the BIOS boot order on the client to do a network boot before booting from the local disk, you will need to reset the order to prevent the node from trying to do another network install.

5.11 Complete the Cluster Setup

Ensure that all client nodes have fully booted before proceeding with this step.

Press the **<Complete Cluster Setup>** button. This will run the final installation configurations scripts from each OSCAR software package, and perform various cleanup and re-initialization functions.

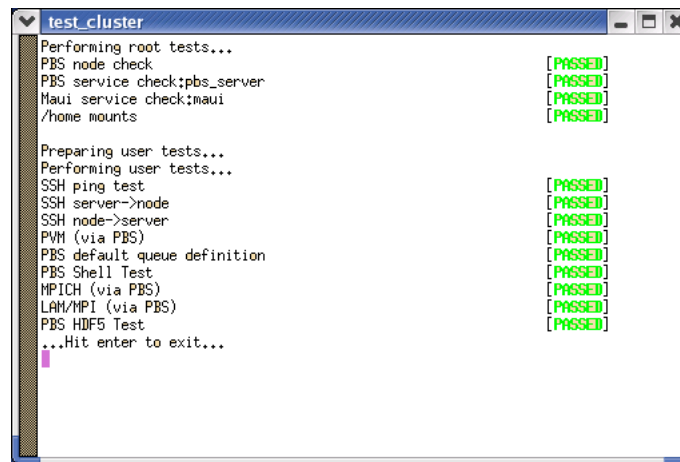
A popup window will indicate the success or failure of this step. Press the **<Close>** button to dismiss it.

5.12 Test Cluster Setup

A simplistic test suite is provided in OSCAR to ensure that the key cluster components (OpenSSH, PBS, MPI, PVM, etc.) are functioning properly.

Press the **<Test Cluster Setup>** button. This will open a separate window to run the tests in. The cluster's basic services are checked and then a set of `root` and user level tests are run.

A sample dialog is shown in Figure 7. If any of the test fail, then there may be problem with your installation.



```
test_cluster
Performing root tests...
PBS node check [ PASSED ]
PBS service check:pbs_server [ PASSED ]
Maui service check:maui [ PASSED ]
/home mounts [ PASSED ]

Preparing user tests...
Performing user tests...
SSH ping test [ PASSED ]
SSH server->node [ PASSED ]
SSH node->server [ PASSED ]
PVM (via PBS) [ PASSED ]
PBS default queue definition [ PASSED ]
PBS Shell Test [ PASSED ]
MPICH (via PBS) [ PASSED ]
LAN/MPI (via PBS) [ PASSED ]
PBS HDF5 Test [ PASSED ]
...Hit enter to exit...
```

Figure 7: Setup cluster tests

At the beginning of the test, when the PBS server is shut down, you may see a benign error message about the `pbsnodes` command not being able to connect to the PBS server. This is safe to ignore.

5.13 Congratulations!

Your cluster setup is now complete. Your cluster nodes should be ready for work.

Be sure to read Section 6 (starting on page 31) – it contains vital system administrator-level information on several of the individual packages that were installed as part of OSCAR.

5.14 Adding and Deleting client nodes

This section describes the steps need when it becomes necessary to add or delete client nodes. If you have already built your cluster successfully and would like to add or delete a client node, execute the following from the top-level OSCAR directory:

```
# ./install_cluster <device>
```

Like before, you must substitute the device name (e.g., eth1) for the server node's internal ethernet adapter in the above command. See Section 5.2 (page 17). Once the OSCAR wizard appears, you are ready to add or delete clients. Note that these steps will reuse the existing images made with the initial install, however, it will extend or contract the set of defined clients in the cluster.

5.14.1 Adding OSCAR clients

Press the button of the wizard entitled **<Add OSCAR Clients>**. A sample dialog is shown in Figure 8. These steps should seem familiar – they are same as the initial install steps. Refer to Sections 5.8, 5.9, and 5.11.

Note that when adding nodes, in the Defining OSCAR Clients step, you will typically need to change the following fields to suit your particular configuration:

- Number of hosts
- Starting number
- Starting IP

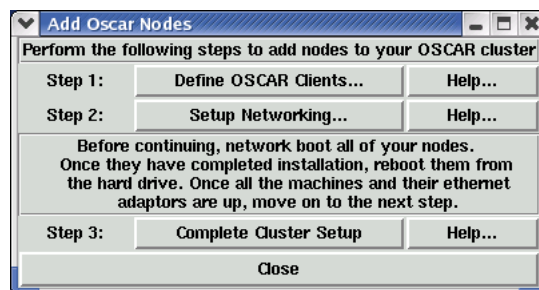


Figure 8: Adding OSCAR clients.

5.14.2 Deleting clients

Press the button of the wizard entitled **<Delete OSCAR Clients>**. A sample dialog is shown in Figure 9. Select the node(s) that you wish to delete and press the button **<Delete clients>**, then press **<Close>**.

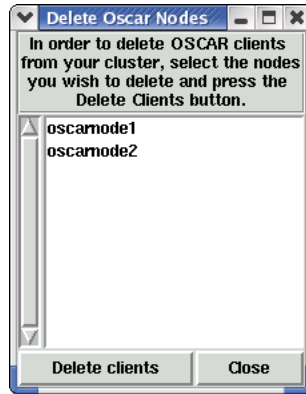


Figure 9: Deleting OSCAR clients.

5.15 Install/Uninstall OSCAR Packages

The package installation system is a new feature in OSCAR 3.0 and is designed to help simplify adding and removing OSCAR packages after the initial system installation. The system is fairly straightforward for the user as packages are selected for the appropriate action from the wizard. This is the first release to support package install and uninstall from the Wizard therefore some general background is provided to assist the user if problems are encountered.

5.15.1 Selecting the Right Package

Because of the way OSCAR handles packages, two packages of the same name cannot currently co-exist nicely on the system. Packages can be in one of two spots:

1. `OSCAR_HOME/packages` – OSCAR installation directory
2. `/var/lib/oscar/package` – OPD download area

If there are multiple packages on the system of the same name, the package located in the OPD download area is the package that the system recognizes. It is possible that a package placed in `/var/lib/oscar/package` via a download or manually could get loaded into the database, while the original package that is located in `OSCAR_HOME/packages` is the one that is actually installed. That is why this system does not support upgrades. That is planned for a future OSCAR release. This can be avoided by uninstalling the original package first before you download anything and then re-running the wizard. When the wizard first runs (`./install_cluster ethX`), the XML files are re-read for all packages and the database re-initialized.

5.15.2 Single Image Restriction

Currently, if more than one image is detected on your system, the package install/uninstall system will not run. Packages are installed and uninstalled to the clients (compute nodes), the server, and an image (singular). This release does not support the mapping of nodes to an image for package install/uninstall through the wizard. This single image restriction is instituted to help reduce the risk of error with this initial release, i.e., "keep it simpler". The system looks for images in `/var/lib/systemimager/images` and no other place.

5.15.3 Failures to Install and Uninstall

A package is not installed or uninstalled until it successfully installs or uninstalls on the compute nodes, the server, and a single image. We will run through an example problem as the best way of describing how to debug the system. As a note, the system does sanity checking to try to avoid the following situation.

For example, during the install process the clients may have been successfully installed, but the server installation died half way through for some unknown reason, and the install program exits with an error. If this happens, a second attempt to install will probably fail even if the server's problem is resolved.

The reason for this is that something happened on the clients and it succeeded, probably RPM's were installed. So the second time you go to install the package, the `rpm -Uvh` command on the nodes will fail. A `-f` flag is avoided.

So, what do you do?

Really, you have two options. The first is to look back through the log to see what commands were run, and undo those commands by hand. The second would be to manually push out the uninstall scripts to the compute nodes, and run them outside of the wizard environment. All scripts are re-runnable in OSCAR, and for several good reasons.

These options are not good ones, and we recognize this fact. We are working on better tools for general use. Generally, it is a difficult problem to judge error codes on remote systems and try to guess the appropriate actions to take. We have developed some additional software to start solving this problem, but there is still work that needs to be done in this area.

On the uninstall side of things, the picture is much simpler. The only action taken by the system is to run the package's uninstall script and judge the return code. So an uninstall failure is probably a faulty uninstall script.

Note if you hit the **<Cancel>** button or close down the window, some unexpected results may occur (post installation scripts will still run) - this is perfectly normal - please refer to the release notes in Section 3 for more info.

5.15.4 More Debugging Info

Some amount of trouble was taken to insure that decent debugging output was printed to the log – making best use of this output is your best bet in an error case.

5.16 Starting over – installing OSCAR again

If you feel that you want to start the cluster installation process over from scratch in order to recover from irresolvable errors, you can do so with the `start_over` script located in the `scripts` subdirectory.

It is important to note that `start_over` is *not* an uninstaller. That is, `start_over` does *not* guarantee to return the head node to the state that it was in before OSCAR was installed. It does a “best attempt” to do so, but the only guarantee that it provides is that the head node will be suitable for OSCAR re-installation. For example, the RedHat 7.x series ships with a LAM/MPI RPM. The OSCAR install process removes this RedHat-default RPM and installs a custom OSCAR-ized LAM/MPI RPM. The `start_over` script only removes the OSCAR-ized LAM/MPI RPM – it does not re-install the RedHat-default LAM/MPI RPM.

Another important fact to note is that because of the environment manipulation that was performed via `switcher` from the previous OSCAR install, *it is necessary to re-install OSCAR from a shell that was not tainted by the previous OSCAR installation*. Specifically, the `start_over` script can remove most files and packages that were installed by OSCAR, but it cannot chase down and patch up any currently-running user environments that were tainted by the OSCAR environment manipulation packages.

Ensuring to have an untainted environment can be done in one of two ways:

1. After running `start_over`, completely logout and log back in again before re-installing. *Simply launching a new shell may not be sufficient* (e.g., if the parent environment was tainted by the previous OSCAR install). This will completely erase the previous OSCAR installation's effect on the environment in all user shells, and establish a set of new, untainted user environments.
2. Use a shell that was established *before* the previous OSCAR installation was established. Although perhaps not entirely intuitive, this may include the shell was that initially used to install the previous OSCAR installation.

Note that the logout/login method is *strongly* encouraged, as it may be difficult to otherwise absolutely guarantee that a given shell/window has an untainted environment.

6 Package-Specific Installation Notes

The following sections provide package-specific notes regarding installation.

6.1 Disabling Services

The `disable-services` OSCAR package disables the following services (if they exist) on the client nodes:

- Incoming mail service: the `sendmail`, `exim`, and `postfix` daemons are disabled. This prevents incoming mail from being received on the nodes. Note that outgoing mail is not disabled; most outgoing mail is sent immediately. If there is some transient failure and the mail is not sent immediately, it will go to the mail service's queue. A `cron.hourly` crontab is installed that runs every hour to trigger the queue in case this happens.
- Kudzu: the Kudzu service looks for new hardware, usually upon boot up. At boot time, this process takes over 30 seconds. It is disabled in order to speed up the booting of individual nodes.
- `slocate`: the `slocate` service runs a top-level `find` command over all local filesystems periodically (some Linux distributions have it set to run daily, others have it set to run weekly) in order to index all filenames for quick lookup using the `locate` command. The top-level `find` command takes significant amounts of system resources to run, and is therefore disabled.
- `makewhatis`: the `makewhatis` command is run via crontab (sometimes daily, sometimes weekly – depending on the specific Linux distribution) to generate manual page indexes. As with `slocate`, this command takes significant amounts of system resources to run, and is therefore disabled.

Note that these services are not uninstalled – they are simply disabled. Administrators are free to re-enable them if they wish.

6.2 HDF5

HDF5 is a Hierarchical Data Format product consisting of a data format specification and a supporting library implementation.

HDF5 includes the following features which make it a widely-used scientific data format:

- supports a user-defined hierarchical grouping structure with varied datatypes and attributes to organize large and varied data in research, development, and production environments,

- runs on parallel computing platforms using MPI I/O (note that the compiled version of HDF5 supplied in OSCAR currently only supports serial operation),
- supports an all-encompassing variety of datatypes, including compound and user-defined datatypes,
- supports very large files (theoretically terabyte and larger, though the outer limits have not been tested)

HDF5 is fully documented – see `/opt/hdf5-oscar-1.4.4-post2/doc/hdf5/index.html` for the documentation that accompanied the release of HDF5 that is installed with this OSCAR distribution. For documentation of the current release of HDF5 as served from the HDF Group’s website, see <http://hdf.ncsa.uiuc.edu/HDF5/doc/>.

For further information regarding HDF5, such as lists of current users, supporting agencies, and applications that employ the library, see <http://hdf.ncsa.uiuc.edu/HDF5/>.

OSCAR installs the HDF5 library at `/opt/hdf5-oscar-1.4.4-post2/lib/`. Please note that the HDF5 library included in OSCAR is presently only available for serial operation. Parallel operation will be enabled in a future version of OSCAR.

When compiling HDF5 in parallel mode, two compiler issues must be considered:

- there is a bug in the `gcc 2.96` compiler which affects HDF5. To work properly, both the MPI implementation and HDF5 must be compiled with a compiler other than 2.96. The HDF5 Group suggests using `gcc 2.95.3` for both.
- when compiling HDF5 in parallel mode with an MPI implementation on a Linux system with kernels 2.4 or greater, you must specify that it should support >2GB file sizes. For example, configure the installation of MPICH with the following configuration command-line option:

```
-cflags="-D_LARGEFILE_SOURCE -D_LARGEFILE64_SOURCE -D_FILE_OFFSET_BITS=64"
```

6.3 Selecting a Different Kernel with `kernel_picker`

`kernel_picker` is a Perl script which allows a user to install a given kernel into an OSCAR image different from the one which is installed by default. After step 1, but before step 2, you can run `kernel_picker` to substitute a given kernel into your OSCAR (SIS) image. If executed with no command line options, you will be prompted for all information. If you use any command line options, the program will assume that you know what you are doing and prompt you *only* for information which is required for correct execution.

The `kernel_picker` program assumes that the optional OSCAR image files you wish to use reside in a subdirectory in the `/var/lib/systemimager/images` directory. By default, the original OSCAR image is in a subdirectory named `oscarimage`.

`kernel_picker` is installed in the `/opt/kernel_picker/bin` directory. Documentation is available in HTML, PostScript, PDF, plain text, and manpage formats. To see the manpage documentation, type the following at a Unix command prompt:

```
$ man /opt/kernel_picker/man/man1/kernel_picker.1
```

6.4 Networking Package

The networking package was created to set up the OSCAR server as a caching nameserver.

By default, this package will set up a caching nameserver on the OSCAR server.

6.5 Managing NTP for the OSCAR Server and Clients

NTP is the Network Time Protocol which is used to synchronize the computer clock to external sources of time. The `ntpd` daemon can operate as a client (by connecting to other NTP servers to get the current time) and as a server (by providing the current time to other NTP clients).

OSCAR uses NTP to synchronize:

- the server to an external source of time (if the server is connected to the internet), and
- the clients to the OSCAR server (so that the clients do not need to be connected to the internet).

For the OSCAR server, NTP is configured to contact several public NTP servers for time synchronization. If none of these external servers can be contacted (because the server is not connected to the internet for example), it will still be used by the OSCAR clients as a source for the current time. Thus, your OSCAR cluster will be locally synchronized even if it is not in-sync with the rest of the world.

For each OSCAR client, NTP is configured to contact only the OSCAR server. If you want to have your clients connect to other public NTP servers, you will have to edit the NTP configuration files.

6.5.1 Configuring NTP

By default, the OSCAR server is set up to use several public NTP servers for its time source, and the OSCAR clients are set up to use the OSCAR server for their time sources. If you want to change the servers used for time synchronization, you will need to edit two configuration files: `/etc/ntp.conf` and `/etc/ntp/step-tickers`.

For the `/etc/ntp.conf` file, place entries of the form “`server SERVER_NAME_OR_IP`” at the top of the file, one entry per line. You can have as many `server` lines as you want. However, ALL of the listed servers are consulted for time synchronization, so you may want to limit the number of servers to three or less. Here is an example of the first few lines of a typical `/etc/ntp.conf` file:

```
# These are some servers for use by the ntpd daemon.
server 130.126.24.24
server ntp0.cornell.edu
server ntp.cmr.gov
```

Notice that you may use either FQDNs (fully qualified domain names) or IP addresses. Most NTP servers prefer that you use FQDNs in case they change the server IP address, but typically these IP addresses are fairly static.

For the `/etc/ntp/step-tickers` file, place these NTP servers on a single line separated by spaces. Here is an example of the `/etc/ntp/step-tickers` file using the same servers:

```
130.126.24.24 ntp0.cornell.edu ntp.cmr.gov
```

The servers used in the `/etc/ntp.conf` file and the `/etc/ntp/step-tickers` file do not have to be the same. The `/etc/ntp/step-tickers` file is used to force the clock to be set correctly at boot time, while the `/etc/ntp.conf` file is used to adjust the clock in small increments while the system is running.

6.5.2 Enabling/Disabling the NTP Service

By default, the `ntpd` daemon is configured to start at boot time in run levels 2 through 5. If for some reason you want to disable NTP without actually uninstalling it, execute the following commands:

```
# /etc/init.d/ntpd stop
# /sbin/chkconfig --level 2345 ntpd off
```

This will not only stop any currently running `ntpd` daemon, but also prevent NTP from starting up at boot time.

NOTE: You must be `root` to execute these commands.

To restart NTP and make NTP start up at boot time, execute the following commands:

```
# /etc/init.d/ntpd restart
# /sbin/chkconfig --level 2345 ntpd on
```

For more information on NTP, see the (rather lengthy) documentation at <http://www.ntp.org/>.

6.6 The OSCAR Password Installer and User Management (OPIUM)

Currently, OPIUM manages users and passwords by replicating files on a triggered basis. The user account files are copied to each node on a regular interval. It is possible to add your own files to this list, by editing the configuration file `/opt/opium/etc/sync_users.conf`.

6.6.1 SSH Configuration

The OPIUM package also handles the SSH key setup. This is done by placing startup scripts in `/etc/profile.d` which generate SSH user key files if they don't already exist. It also generates `authorized_keys` files which enable users to traverse the cluster securely without entering passwords. Since these files reside in the users' home directories, it is necessary for the home directory to be mounted on a node in order for a user to `ssh` to it. Also, a user must log on to the head node in order for the keys to be generated, as the `/etc/profile.d` scripts are not installed on the compute nodes in OSCAR 3.0. SSH1 and SSH2 key files are supported.

6.7 The Portable Batch System (PBS) and Maui Scheduler

PBS serves as the job launcher and batch queueing system for OSCAR. OSCAR uses the open source version of PBS, OpenPBS. A commercial version (PBSPro) exists as well. OpenPBS includes a basic FIFO scheduler with it, but is disabled by default in OSCAR. A more robust, open source scheduler named "Maui" is used instead. In OSCAR 3.0, we are using OpenPBS version 2.3.16 and Maui 3.2.5p2. PBS and Maui will continue to be updated in future OSCAR releases.

Basic PBS functionality is tested by OSCAR's test suite, and is also used to launch jobs when testing other software included in OSCAR. If the PBS test passes, PBS and Maui are up and working.

If the users of your OSCAR cluster have not used PBS before, you can expect somewhat of a learning curve. The OSCAR user's documentation contains some useful information to get them started. The user's document instructs users to ask the system administrator to provide them with sample PBS scripts used in the OSCAR test suite. Once the OSCAR test suite has been run (step 6 in the install process), these scripts can be found in the home directory of the `oscartst` user if the tests have been run previously.

6.7.1 Configuring PBS

By default, PBS installs without any queues or cluster specific parameters defined. OSCAR configures PBS with sensible defaults based on what it finds in the SIS database. When the “Complete Cluster Setup” step is executed from the wizard, the `post_install` script from the PBS package in OSCAR is called. The `post_install` configures only PBS parameters that are non-existent, so as not to overwrite local customizations by the system administrator. However, if you wish to force the default values back in place over any local customizations, the `post_install` script can be invoked manually with a `--default` option. This will revert all values to the original OSCAR settings.

`qmgr` can be used to configure queues and PBS server parameters. The OSCAR PBS `post_install` script (located off the top-level OSCAR installation directory in `packages/pbs/scripts`) uses `qmgr` behind the scenes. There are man pages available, but reading the PBS admin guide is the best way to learn how to use it. It is available on OpenPBS’s homepage, listed below. You will have to create an account on their site in order to download the admin guide.

6.7.2 PBS Resources

Arbitrary node properties can be set by the administrator. PBS calls these properties “resources”. These resources can be specified on the `qsub` command line when a user submits a job. This allows a user to restrict their jobs to run only on nodes exhibiting certain properties. If some nodes of a cluster have more memory, a different network, faster processors, etc., jobs can be submitted so they only run a specific subset. These properties are stored in plain text in `/usr/spool/PBS/server_priv/nodes`. However, if adjusted in the plain text file, the PBS server must be restarted in order for changes to take effect. The more elaborate method is to use the `qmgr` command to modify node properties via the PBS API. OSCAR gives each node a starting property of “all”.

6.7.3 An FAQueue

A popular misconception about PBS queues is that they are bound to a group of nodes. This is false. If you have a four node queue defined, it is not associated with any specific nodes. You can think of a queue as a multidimensional box that a job must fit in in order to allow submission. That is, the submitted parameters must fit within certain max and min values for nodes, ppn (procs per node), walltime, etc. If specific nodes are desired to run on, then resource attributes must be defined.

If you would like to get a full dump of your PBS server and queue configuration, you can issue this command:

```
# qmgr -c "print server"
```

The `qmgr` interface can be used to define additional queues and their parameters. You can also change the parameters on the default OSCAR queue, “workq”. For example, to show the configuration of the `workq`, execute the following:

```
# qmgr -c "list queue workq"
```

To change any of the values listed, use the following:

```
# qmgr -c "set queue workq PARAMETER = VALUE"
```

where `PARAMETER` is a parameter from the “`list queue`” command, and `VALUE` is a valid value for that parameter. You can use the “`print server`” and/or “`list queue`” commands to verify your changes.

Be aware that if you call the `post_install` command with the `--default` option, you will lose your customizations. Also note that OSCAR’s default wallclock limit on `workq` is 10,000 hours. Depending on the application mix that will run on your cluster, you may wish to adjust this value.

Some useful links:

- OpenPBS: <http://www.openpbs.org/>
- PBSPro: <http://www.pbspro.com/>
- Maui Scheduler: <http://www.supercluster.org/>

6.8 Packet Filtering with `pfilter`

`pfilter` is a firewall compiler. The `pfilter` package is used to control the packet filtering capabilities available in the Linux kernel. It takes in high-level firewall directives, and produces a complete firewall output commands file that can be turned on or off like other Linux services. Like other compilers, `pfilter` adds appropriate “glue” code to the compiled output. `pfilter`’s added “glue” code consists of common things that are done by any good firewall, including turning on TCP networking protective features.

When OSCAR was installed, it merged any pre-existing server `pfilter` installation configuration into a new configuration, and created new client machine `pfilter` configurations. The resulting `pfilter` configurations do the following:

- any network connections that were specifically allowed in the former server `pfilter` installation configuration are still allowed to the main OSCAR server node
- the main OSCAR server node and all client OSCAR nodes allow `ssh` logins from anywhere
- the main OSCAR server node has `http` access enabled from anywhere
- any remaining network connections from outside the cluster are blocked
- the logging of bad network packets to `syslog` is turned off
- all network communication of any kind between nodes in the cluster is enabled

If the main OSCAR server has two or more network interfaces, `pfilter` will attempt to determine which network interface is the public interface, and then enable packet forwarding and network address translation for the remaining interfaces if it can.

`pfilter` is turned on by default – this is considered good “defense in depth” security for a cluster (see Section D for more information on cluster security). If for some reason you need to disable `pfilter` packet filtering (perhaps for debugging – disabling it permanently is not recommended), you can disable `pfilter` on subsequent system boots with the following command:

```
# chkconfig --level=2345 pfilter off
```

To turn on `pfilter` packet filtering immediately, execute the following command:

```
# service pfilter start
```

For more information on how to use and configure `pfilter`, see the `pfilter(8)` and `pfilter.conf(5)` and `pfilter.rulesets(5)` man pages.

6.9 Managing machines and images in SIS

During the life of your cluster, you may want to delete unused images, create new images, or change the image that a client uses. Currently OSCAR doesn't have a direct interface to do this, but you can use the SIS commands directly. Here are some useful examples:

- To list all defined machines, run:

```
mksimachine --List
```

- To list all defined images, run:

```
mksiimage --List
```

- To delete an image, run:

```
mksiimage --Delete --name <imagename>
```

- To change which image a machine will install, run:

```
mksimachine --Update --name <machinename> --image <imagename>
```

There is also a SIS GUI that is available. Start it by running `tksis`.

More details on these commands can be obtained from their respective man pages.

6.9.1 Multicast Installs with SystemImager

SIS now includes multicast install capability with SystemImager v3.2.x and Flamethrower v1.0.x. Working with multicast can prove very beneficial, especially for large sites. However, multicast can be a tricky thing to get working reliably, based on networking equipment, multicast tuning parameters, machine speed, etc. If you are interested in giving multicast a try, see the "HOWTO Use Flamethrower for Multicast Installs" section in the SystemImager manual (<http://www.systemimager.org/documentation/>), and please provide us with feedback on your success!

6.10 Which MPI do you want to use?

OSCAR has a generalized mechanism to both set a system-level default MPI implementation, and also to allow users to override the system-level default with their own choice of MPI implementation.

This allows multiple MPI implementations to be installed on an OSCAR cluster (e.g., LAM/MPI and MPICH), yet still provide unambiguous MPI implementation selection for each user such that "`mpicc foo.c -o foo`" will give deterministic results.

6.10.1 Setting the system-level default

The system-level default MPI implementation can be set in two different (yet equivalent) ways:

1. During the OSCAR installation, the GUI will prompt asking which MPI should be the system-level default. This will set the default for all users on the system who do not provide their own individual MPI settings.

2. As root, execute the command:

```
# switcher mpi --list
```

This will list all the MPI implementations available. To set the system-level default, execute the command:

```
# switcher mpi = name --system
```

where “name” is one of the names from the output of the `--list` command.

NOTE: System-level defaults for `switcher` are currently propagated to the nodes on a periodic basis. If you set the system-level MPI default, you will either need to wait until the next automatic “push” of configuration information, or manually execute the `/opt/opium/bin/sync_users` command to push the changes to the compute nodes.

NOTE: Using the `switcher` command to change the default MPI implementation will modify the `PATH` and `MANPATH` for all *future* shell invocations – it does *not* change the environment of the shell in which it was invoked. For example:

```
# which mpicc
/opt/lam-1.2.3/bin/mpicc
# switcher mpi = mpich-4.5.6 --system
# which mpicc
/opt/lam-1.2.3/bin/mpicc
# bash
# which mpicc
/opt/mpich-4.5.6/bin/mpicc
```

6.10.2 Setting the user-level default

Setting a user-level default is essentially the same as setting the system-level default, except without the `--system` argument. This will set the user-level default instead of the system-level default:

```
$ switcher mpi = lam-1.2.3
```

Using the special name `none` will indicate that no module should be loaded for the `mpi` tag. It is most often used by users to specify that they do not want a particular software package loaded.

```
$ switcher mpi = none
```

Removing a user default (and therefore reverting to the system-level default) is done by removing the default attribute:

```
$ switcher mpi --show
user:default=mpich-1.2.4
system:exists=true
$ switcher mpi --rm-attr default
$ switcher mpi --show
system:default=lam-6.5.6
system:exists=true
```

6.10.3 Use `switcher` with care!

`switcher` immediately affects the environment of all future shell invocations (including the environment of scripts). To get a full list of options available, read the `switcher(1)` man page, and/or run `switcher --help`.

A Network Booting Client Nodes

There are two methods available for network booting your client nodes. The first is to use the Preboot eXecution Environment (PXE) network boot option in the client's BIOS, if available. If the option is not available, you will need to create a network boot floppy disk using the SystemImager boot package. Each method is described below.

1. **Network booting using PXE.** To use this method, the BIOS and network adapter on each of the client nodes will need to support PXE version 2.0 or later. The PXE specification is available at <http://developer.intel.com/ial/wfm/tools/pxepdk20/>. Earlier versions may work, but experience has shown that versions earlier than 2.0 are unreliable. As BIOS designs vary, there is not a standard procedure for network booting client nodes using PXE. More often than not, the option is presented in one of two ways.
 - (a) The first is that the option can be specified in the BIOS boot order list. If presented in the boot order list, you will need to set the client to have network boot as the first boot device. In addition, when you have completed the client installation, remember to reset the BIOS and remove network boot from the boot list so that the client will boot from its local hard drive and will not attempt to do the installation again.
 - (b) The second is that the user must watch the output of the client node while booting and press a specified key such as "N" at the appropriate time. In this case, you will need to do so for each client as it boots.
2. **Network booting using a SystemImager boot floppy.** The SystemImager boot package is provided with OSCAR just in case your machines do not have a BIOS network boot option. You can create a boot floppy through the OSCAR GUI installation wizard on the `<Setup Networking>` panel or by using the `mkautoinstalldiskette` command.

Once you have created the SystemImager boot floppy, set your client's BIOS to boot from the floppy drive. Insert the floppy and boot the machine to start the network boot. Check the output for errors to make sure your network boot floppy is working properly. Remember to remove the floppy when you reboot the clients after installation.

B What Happens During Client Installation

Once the client is network booted, it either boots off the autoinstall Diskette that you created or uses PXE to network boot, and loads the install kernel. It then broadcasts a BOOTP/DHCP request to obtain the IP address associated with its MAC address. The DHCP server provides the IP information and the client looks for its auto-install script in `/var/lib/systemimager/scripts/`. The script is named `<nodename>.sh` and is a symbolic link to the script for the desired image. The auto-install script is the installation workhorse, and does the following:

1. partitions the disk as specified in the image in `<imagedir>/etc/systemimager/partitionschemes`.
2. mounts the newly created partitions on `/a`.
3. chroots to `/a` and uses `rsync` to bring over all the files in the image.

4. invokes `systemconfigurator` to customize the image to the client's particular hardware and configuration.
5. unmounts `/a`.

Once clone completes, the client will either reboot, halt, or beep as specified when defining the image.

C Installing without a DHCP server

This section provides information on installing OSCAR without the use of a DHCP server on the server node. For most people, this information is not applicable. In addition, this section only presents information where the installation differs from the standard OSCAR install - information presented elsewhere in this document is applicable unless stated otherwise.

Please note that installing an OSCAR cluster with these instructions is not generally supported by the OSCAR developers. While at least one group in the OSCAR team regularly installs and maintains a cluster configured as described here, that doesn't mean it will work in every release. In other words, continue at your own risk.

C.1 Boot Floppies

PXE boot and installation only works correctly if the server running the tftpboot server is also running the DHCP server. As this is obviously not the case if the OSCAR server node can not run a DHCP server, client nodes must be installed using a boot floppy. Further, the boot floppies created with the "Make Boot Floppy" option in the OSCAR wizard assumes there is a DHCP server running on the server node, so we must use a special boot floppy.

Unfortunately, we have to make a boot floppy for every node (since the boot floppy will have IP information encoded in it)⁹. First, a configuration file that provides SIS with the needed boot information must be created. The sample below is used to create a machine `thumb1.osl.iu.edu`. The fields are self-explanatory. The `IMAGESERVER` should be set to the IP address of the OSCAR server node (where you are running the OSCAR wizard). The other networking information must match the information given when building client image during the installation, otherwise your system might not be installable.

```
HOSTNAME=thumb1
DOMAINNAME=osl.iu.edu
DEVICE=eth0
IPADDR=129.79.247.11
NETMASK=255.255.252.0
NETWORK=129.79.244.0
BROADCAST=129.79.247.255
GATEWAY=129.79.247.254
IMAGESERVER=129.79.247.10
```

The boot floppy can be created by running `mkautoinstalldiskette` with the `-config FILE` option specifying the location of the configuration file created above.

⁹This process can obviously be automated, but we do not go into detail on that process here.

C.2 Differences from Standard OSCAR Install

C.2.1 Cluster Definition

When setting up the networking in the OSCAR installation wizard, it is not necessary to collect the MAC addresses of your client nodes. However, to make sure all the proper steps of the OSCAR installation are performed, you should ensure that you open the Wizard pane and click on the “Done” option.

C.2.2 Adding and Deleting client nodes

The Add / Delete client node functionality has not been well tested in this configuration. While there should not be any problems in adding a node to the cluster using boot floppies ¹⁰, please be cautious when adding a new node to the cluster.

D Security

D.1 Security layers

Linux cluster security should, ideally, consist of multiple layers. The main security layers are router packet filtering, network stack protections, host based packet filtering, TCP wrappers, service paring, service configuration, and secure communications. Oscar installs a host based packet filtering package called `pfilter`.

D.2 Router packet filtering

This involves adding packet filtering rules to your border network router. Normally, this is not done because of the difficulty in modifying router tables. Packet filtering involves looking at each network packet, and deciding whether each packet should be allowed, dropped, or rejected, based on tables of rules.

D.3 Network stack protections

Linux kernels have security features built in that can help prevent outsiders from pretending that they are part of your internal network. These features are enabled through `/proc` filesystem entries. A good firewall package will turn these on appropriately (`pfilter` does this when enabled).

D.4 Host based packet filtering

Like router packet filtering, host based network packet filtering involves examining each packet and deciding what do do with it. But with host based filtering, each machine individually filters the network packets going to, from, or through it. Linux kernels from 2.4 on include support for connection tracking and “statefull” packet filtering, which keeps track of ongoing network connections, allowing better filtering decisions to be made based on whether packets are part of an already allowed connection.

The problem with packet filtering is that it requires generating filtering “rulesets” that the `iptables` or `ipchains` programs interpret and store in the running kernel. Creating these rulesets is similar to writing software in assembly language. There are now higher level “languages” and compilers that can be used to generate the rulesets and provide firewalls. OSCAR installs a ruleset compiler/firewall package called `pfilter`. For more information, see the <http://pfilter.sourceforge.net/>.

¹⁰Just remember to create another boot floppy

D.5 TCP Wrappers

TCP wrappers are an access control system that allows control over which network addresses or address ranges can access particular network services on a computer host. This is controlled by the `/etc/hosts.allow` and `/etc/hosts.deny` files. This allows certain services to be only accessible from your local domain, for instance. A common use of this would be to limit exported NFS filesystems to only be accessible from your local domain, while allowing security logins through `ssh` to come in from anywhere. This would be done with a `/etc/hosts.deny` file that looks like this:

```
ALL:    ALL
```

and a `/etc/hosts.allow` file that looks like this:

```
# allow NFS service to domain.net only
portmap:    .domain.net
rpc.mountd: .domain.net
# allow ssh logins from anywhere
sshd:      ALL
```

D.6 Service paring

This is probably the most used of all the security layers, since turning off unneeded network services gets rid of opportunities for network breakins. To hunt down and turn off unwanted services, the `lsof`, `chkconfig`, and `service` system commands can be used. To display which network services are currently listening on a system, do this:

```
# lsof -i | grep LISTEN | awk '{print $1,$(NF-2),$(NF-1)}' | sort | uniq
```

To list the services that will be started by default at the current run level do this:

```
# chkconfig --list | grep 'grep :initdefault: /etc/inittab | \
awk -F: '{print $2}':on | awk '{print $1}' | sort | column
```

To find services started by `xinetd` do this:

```
# chkconfig --list | awk 'NF==2&&$2=='off''{print}' | \
awk -F: '{print $1}' | sort | column
```

The `nmap` port scanning command is also useful to get a hackers-eye view of your systems. The `chkconfig` and `service` commands can be used to turn on and off system services.

D.7 Service configuration

Some network services have their own configuration files. These should be edited to tighten down outside access. For example, the NFS filesystem uses the `/etc/exports` to determine which network addresses can access individual file systems, and which have read-write or read-only access. Indeed, OSCAR sets `/etc/exports` to only NFS export the `/home` directory to the OSCAR cluster nodes.

D.8 Secure communications

The OSCAR team *strongly* recommends using `ssh` for all network logins, and `scp` to copy files either between OSCAR nodes, or to remote hosts. The use of `telnet`, `rsh`, `ftp`, and other legacy network applications are strongly discouraged because of inherent security risks (passwords are transmitted “in the clear”).

`ssh` and `scp` are fully functional replacements for these legacy commands – OSCAR takes care of setting up all the authentication and key management issues for all users.

E Screen-by-Screen Walkthrough

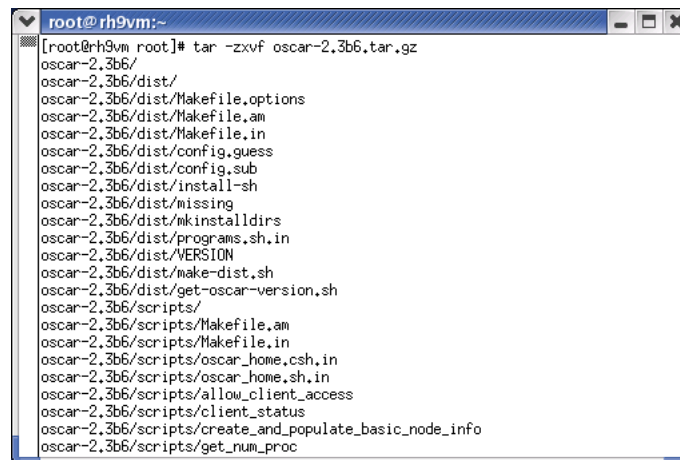
The following is a screen-by-screen walkthrough of a simple installation. It is intended as supplementary material to aid in providing a better feel for the general progression of the installation. For a detailed discussion of the steps, please refer to the Detailed Cluster Installation Procedure.

Note the example screen shots were based in a Red Hat 9.0 using a pre-release version of OSCAR 3.0 in the GNOME graphical environment. Since this section is intended as a supplementary source of information, it is judged to be “close enough” to the real 3.0 release.

Also note that these images have been scaled down to fit within the document. As such, although they are readable, the images may not render nicely on a screen. The images tend to be much more readable on an actual printout.

E.1 Running `install_cluster`

These Figures 10 – 13 walk through the steps needed to prepare and run the `install_cluster` script with the network interface name, which will begin the actual cluster installation. See details in Section 5.2, page 17.



```
root@rh9vm:~  
[root@rh9vm root]# tar -zxvf oscar-2.3b6.tar.gz  
oscar-2.3b6/  
oscar-2.3b6/dist/  
oscar-2.3b6/dist/Makefile.options  
oscar-2.3b6/dist/Makefile.am  
oscar-2.3b6/dist/Makefile.in  
oscar-2.3b6/dist/config.guess  
oscar-2.3b6/dist/config.sub  
oscar-2.3b6/dist/install-sh  
oscar-2.3b6/dist/missing  
oscar-2.3b6/dist/mkinstalldirs  
oscar-2.3b6/dist/programs.sh.in  
oscar-2.3b6/dist/VERSION  
oscar-2.3b6/dist/make-dist.sh  
oscar-2.3b6/dist/get-oscar-version.sh  
oscar-2.3b6/scripts/  
oscar-2.3b6/scripts/Makefile.am  
oscar-2.3b6/scripts/Makefile.in  
oscar-2.3b6/scripts/oscar_home.csh.in  
oscar-2.3b6/scripts/oscar_home.sh.in  
oscar-2.3b6/scripts/allow_client_access  
oscar-2.3b6/scripts/client_status  
oscar-2.3b6/scripts/create_and_populate_basic_node_info  
oscar-2.3b6/scripts/get_num_proc
```

Figure 10: Unpacking OSCAR.

E.2 Download Additional OSCAR Packages...

An optional step. See details in Section 5.3, page 18.

E.3 Select OSCAR Packages to Install

An optional step. See details in Section 5.4, page 19.

E.4 Configure Selected OSCAR Packages

An optional step. See details in Section 5.5, page 19.

```

root@rh9vm:~/oscar-2.3b6
oscar-2.3b6/testing/Makefile.in
oscar-2.3b6/testing/pbs_test
oscar-2.3b6/testing/testprint
oscar-2.3b6/testing/test_cluster
oscar-2.3b6/testing/ssh_user_tests
oscar-2.3b6/testing/README.pbs_test
[root@rh9vm root]#
[root@rh9vm root]# cd oscar-2.3b6
[root@rh9vm oscar-2.3b6]# ./configure
Configuring OSCAR version 2.3b6
checking for a BSD-compatible install... /usr/bin/install -c
checking whether build environment is sane... yes
checking for gawk... gawk
checking whether make sets $(MAKE)... yes
checking build system type... i686-pc-linux-gnu
checking host system type... i686-pc-linux-gnu
checking for documentation LaTeX source... not found
checking for location of OSCAR_HOME... /opt/oscar
configure: creating ./config.status
config.status: creating Makefile
config.status: creating VERSION
config.status: creating dist/Makefile
config.status: creating dist/programs.sh

root@rh9vm:~/oscar-2.3b6
config.status: creating packages/autoupdate/RPMS/Makefile
config.status: creating packages/autoupdate/Makefile
config.status: creating packages/Makefile
[root@rh9vm oscar-2.3b6]#
[root@rh9vm oscar-2.3b6]# make install
Making install in dist
make[1]: Entering directory `/root/oscar-2.3b6/dist'
make[2]: Entering directory `/root/oscar-2.3b6/dist'
make[2]: Nothing to be done for `install-exec-am'.
make[2]: Nothing to be done for `install-data-am'.
make[2]: Leaving directory `/root/oscar-2.3b6/dist'
make[1]: Leaving directory `/root/oscar-2.3b6/dist'
Making install in doc
make[1]: Entering directory `/root/oscar-2.3b6/doc'
make[2]: Entering directory `/root/oscar-2.3b6/doc'
make[2]: Nothing to be done for `install-exec-am'.
make[3]: Entering directory `/root/oscar-2.3b6/doc'
make[4]: Entering directory `/root/oscar-2.3b6/doc'
make[4]: Nothing to be done for `install-exec-am'.
/bin/sh ../dist/mkinstalldirs /opt/oscar/doc
mkdir -p -- /opt/oscar/doc
/usr/bin/install -c -m 644 install.pdf /opt/oscar/doc/install.pdf
make[4]: Leaving directory `/root/oscar-2.3b6/doc'
make[3]: Leaving directory `/root/oscar-2.3b6/doc'

root@rh9vm:~/oscar-2.3b6
make[3]: Entering directory `/root/oscar-2.3b6'
-----
OSCAR has been successfully installed to /opt/oscar.

You should change to that directory and start the OSCAR
installation process:

  shell# cd /opt/oscar
  shell# ./install_cluster <network_device>

Consult the OSCAR installation guide for more information.
-----
make[3]: Leaving directory `/root/oscar-2.3b6'
/bin/sh ../dist/mkinstalldirs /opt/oscar/.
/usr/bin/install -c -m 644 COPYING /opt/oscar/./COPYING
/usr/bin/install -c -m 644 README /opt/oscar/./README
/usr/bin/install -c -m 644 VERSION /opt/oscar/./VERSION
/bin/sh ../dist/mkinstalldirs /opt/oscar/.
/usr/bin/install -c install_cluster /opt/oscar/./install_cluster
make[2]: Leaving directory `/root/oscar-2.3b6'
make[1]: Leaving directory `/root/oscar-2.3b6'
[root@rh9vm oscar-2.3b6]#

```

Figure 11: Configure & Install the OSCAR Toolkit.

E.5 Install OSCAR Server Packages

This step is used to setup the server for the OSCAR cluster. See details in Section 5.6, page 21.

E.6 Build OSCAR Client Image

This step builds a disk image for the clients to download and install onto their local disks. See the details in Section 5.7, page 21.

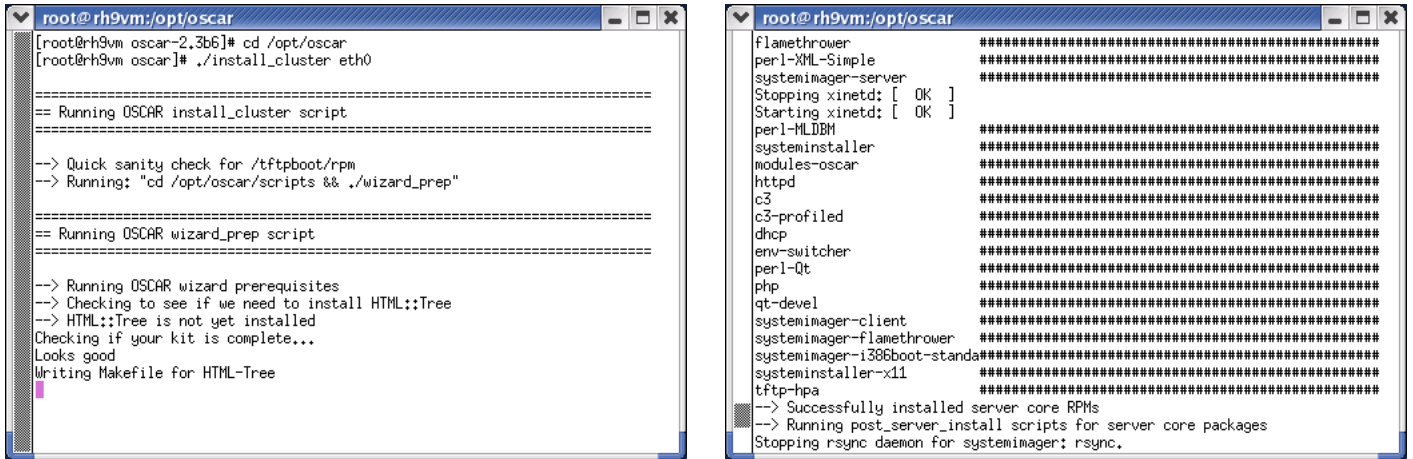


Figure 12: Running the install_cluster script.




Figure 13: Enter an OSCAR Database (ODA) password.

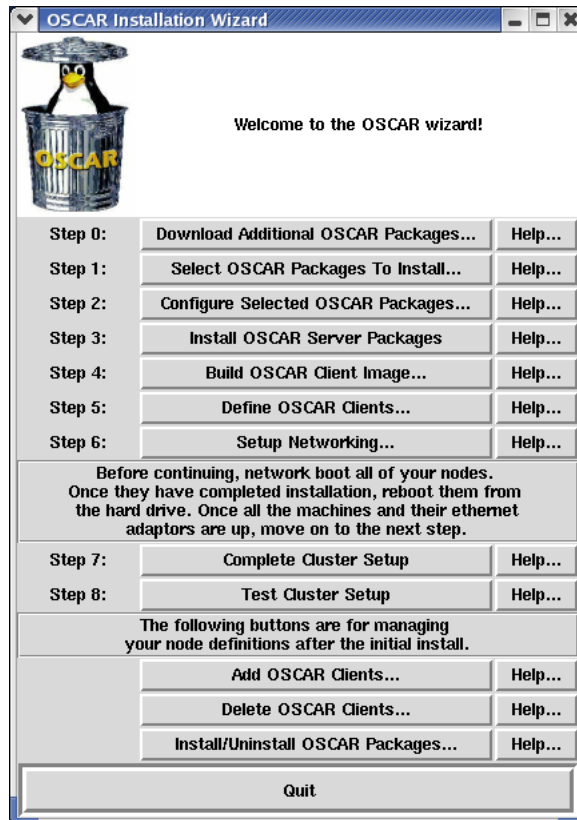


Figure 14: The OSCAR Installation Wizard.

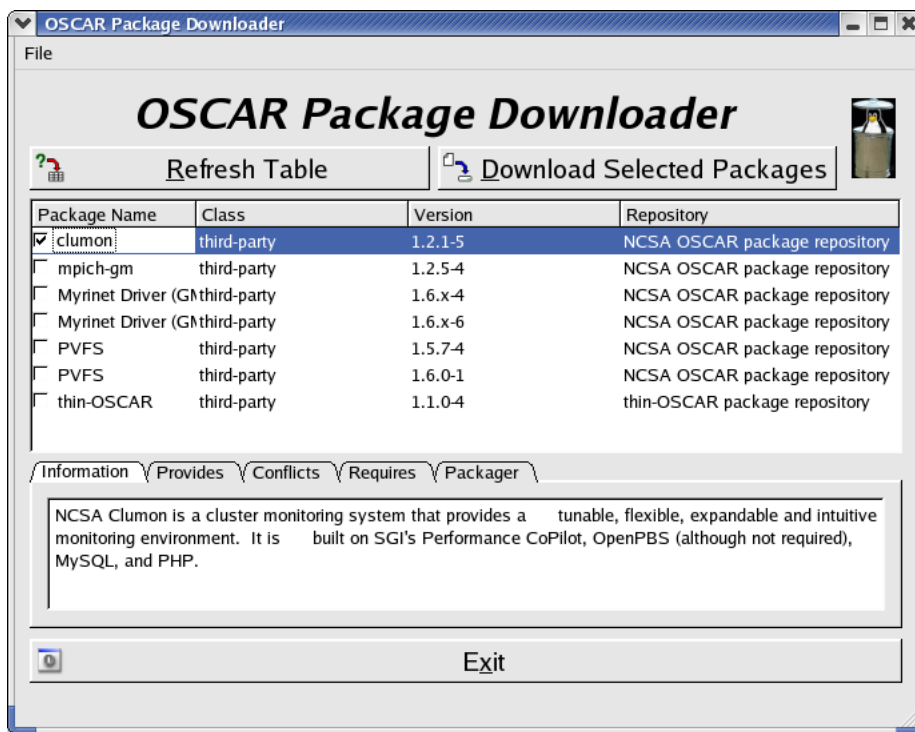


Figure 15: Selecting additional OSCAR packages to download using OPD/OPDer.

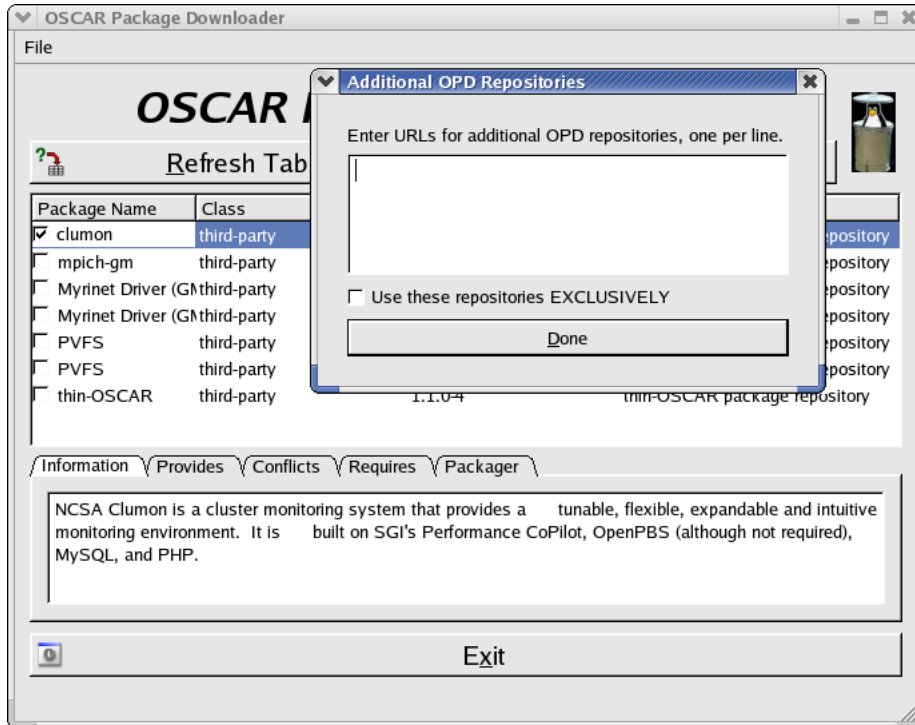


Figure 16: Adding Additional OPD Repositories.

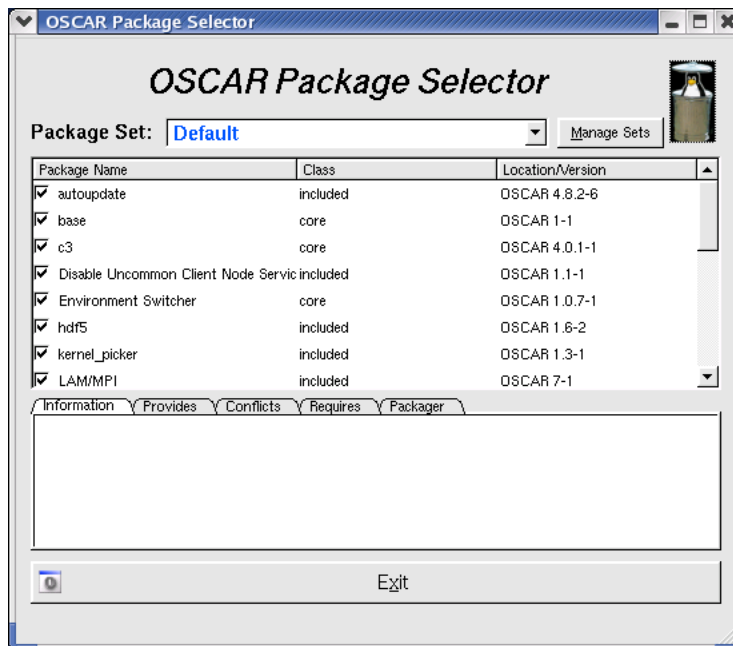


Figure 17: Selecting which OSCAR packages to install.

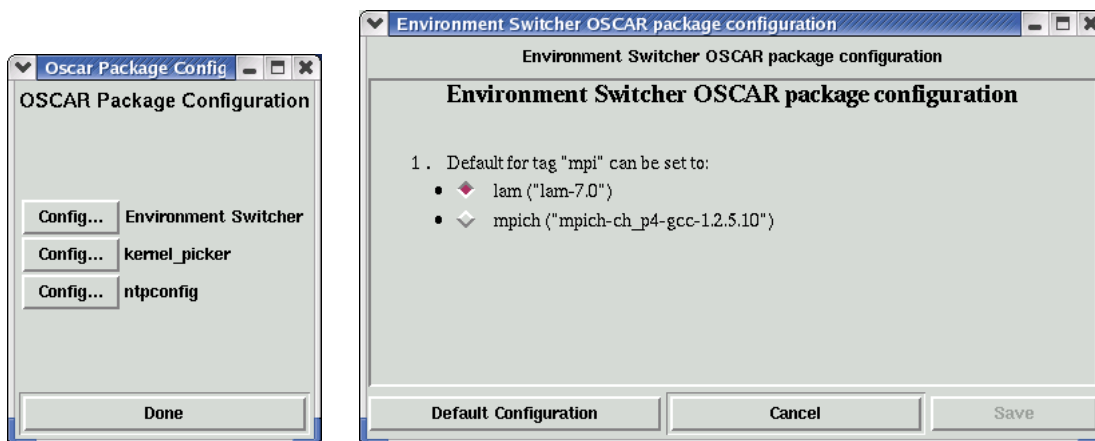


Figure 18: Configuring selected OSCAR packages. For example, the Environment Switcher package has configuration options.

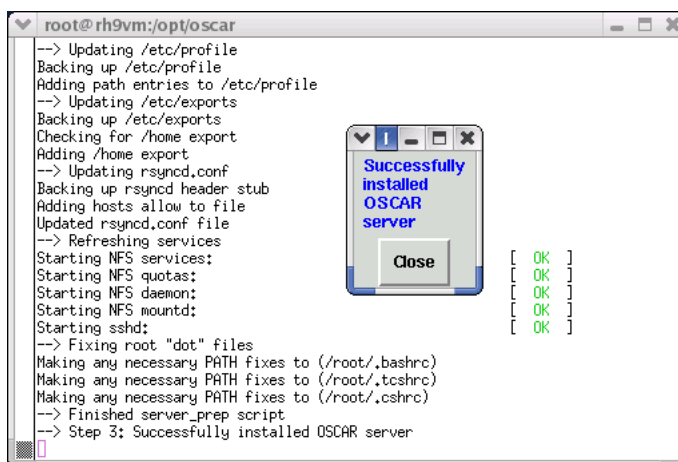


Figure 19: Successfully installed the OSCAR server packages.

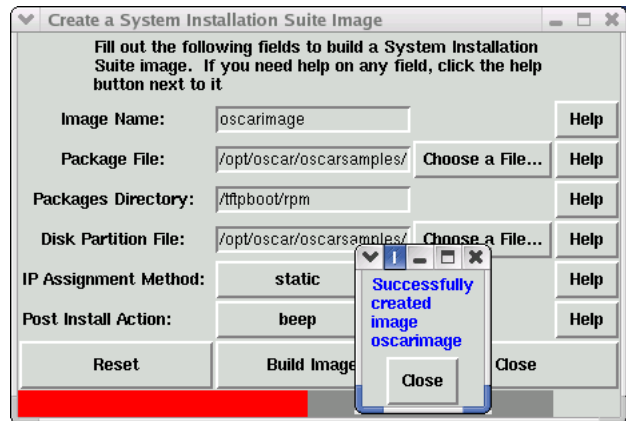
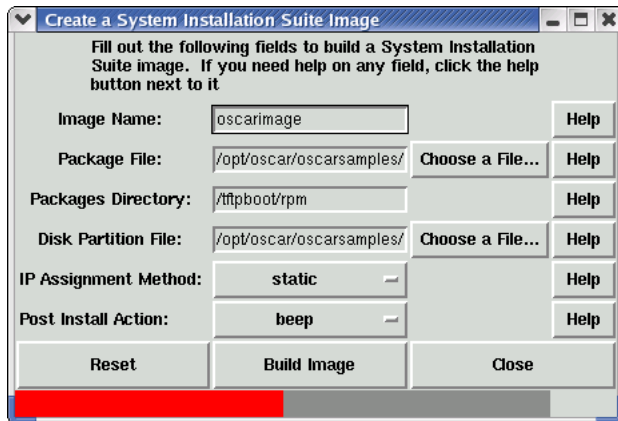
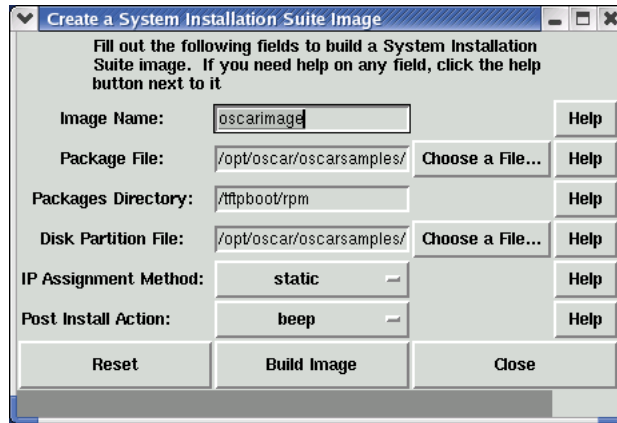


Figure 20: Building the OSCAR client image.

E.7 Define OSCAR Clients

Step 3 is used to specify how many clients there will be, and what their TCP/IP characteristics will be. See the details in Section 5.8, page 23.

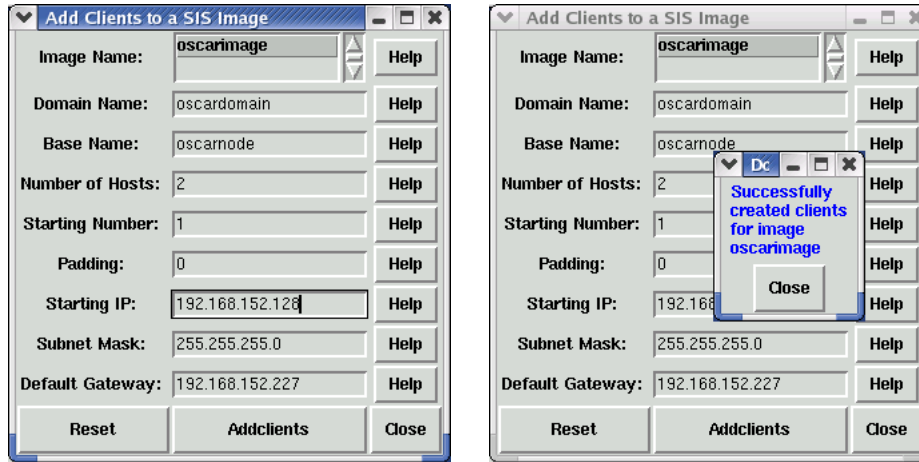


Figure 21: Defining the OSCAR clients.

E.8 Setup Networking

This step is used to collect the MAC addresses of the clients, and then download the disk images to the clients. See the details in Section 5.9, page 24.

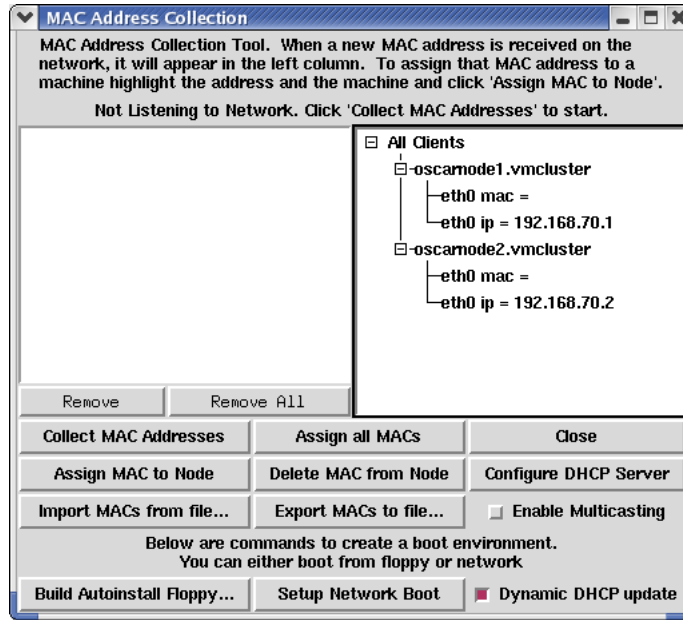


Figure 22: Setup networking: initial window.

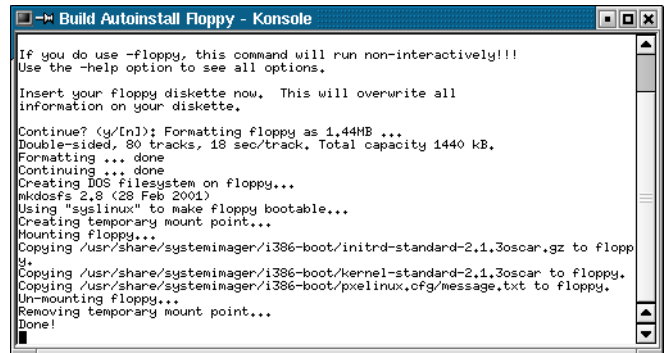
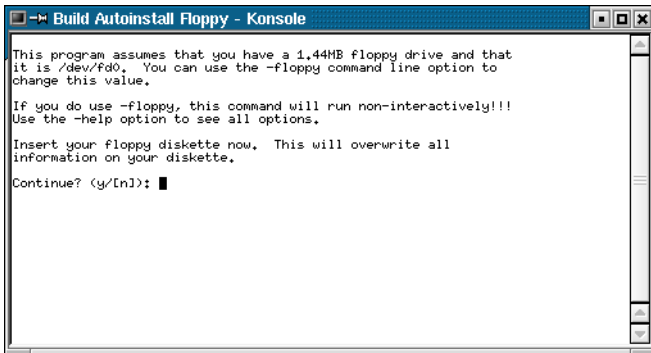
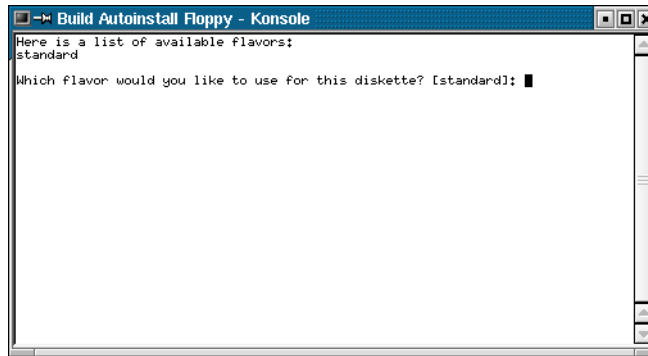


Figure 23: Setup networking: building an autoinstall floppy.

```
SYSLINUX 1.63 2001-08-06 Copyright (C) 1994-2001 H. Peter Anvin
Loading initrd-s.....
Loading kernel-s....._
```

Figure 24: Setup networking: booting the client.

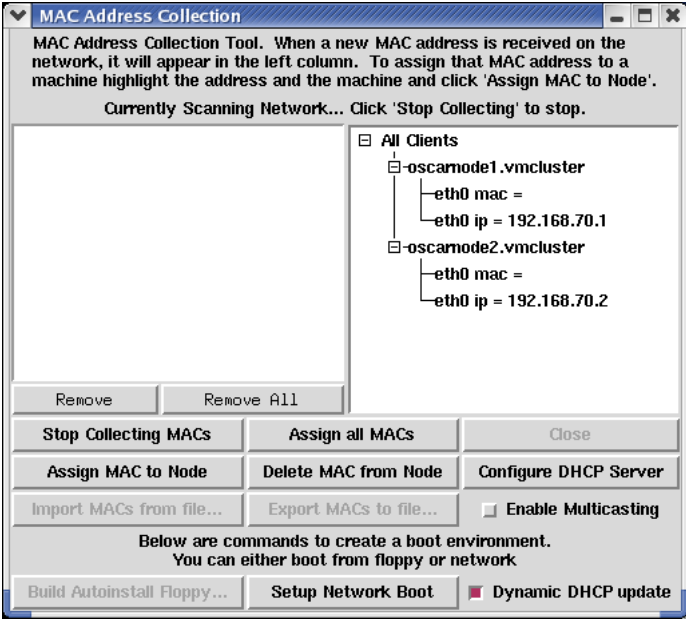


Figure 25: Setup networking: scan for booting client (DHCP request).

```

OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
2048 inodes, 8192 blocks
409 blocks (4.89%) reserved for the super user
First data block=1
1 block group
8192 blocks per group, 8192 fragments per group
2048 inodes per group
Writing inode tables: done
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 26 mounts or
180 days, whichever comes first. Use tune2fs -c or -l to override.
END Creating ram disk for /tmp to hold extra binaries
Checking for configuration file on floppy...
UFS: Can't find an ext2 filesystem on dev fd(2,28).
reiserfs_read_super: can't find a reiserfs filesystem on dev 02:1c.
reiserfs_read_super: try to find super block in old location
reiserfs_read_super: can't find a reiserfs filesystem on dev 02:1c.
No /local.cfg on floppy drive.
IP Address not set by local.cfg. I will use DHCP.
sleep 35: This is to give your switch (if you're using one) time to
recognize your ethernet card before we try the network.
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30
31 32 33 34 35
dhclient
Internet Software Consortium DHCP Client 2.0p15
Copyright 1995, 1996, 1997, 1998, 1999 The Internet Software Consortium.
All rights reserved.
Please contribute if you find this software useful.
For info, please visit http://www.isc.org/dhcp-contrib.html
cat: /floppy/local.cfg: No such file or directory
cat: /floppy/local.cfg: No such file or directory
Listening on LPP/eth0/00:50:56:40:43:78
Sending on LPP/eth0/00:50:56:40:43:78
Listening on LPP/lo/<null>
Sending on LPP/lo/<null>
Sending on Socket/fallback/fallback-net
DHCPDISCOVER on eth0 to 255.255.255.255 port 67 interval 8
DHCPDISCOVER on lo to 255.255.255.255 port 67 interval 7

```

Figure 26: Setup networking: client is broadcasting, allows capture of MAC address.

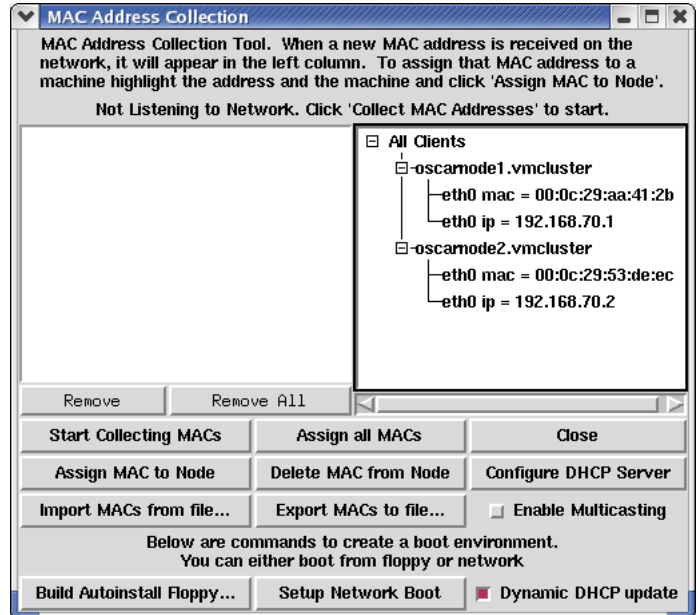
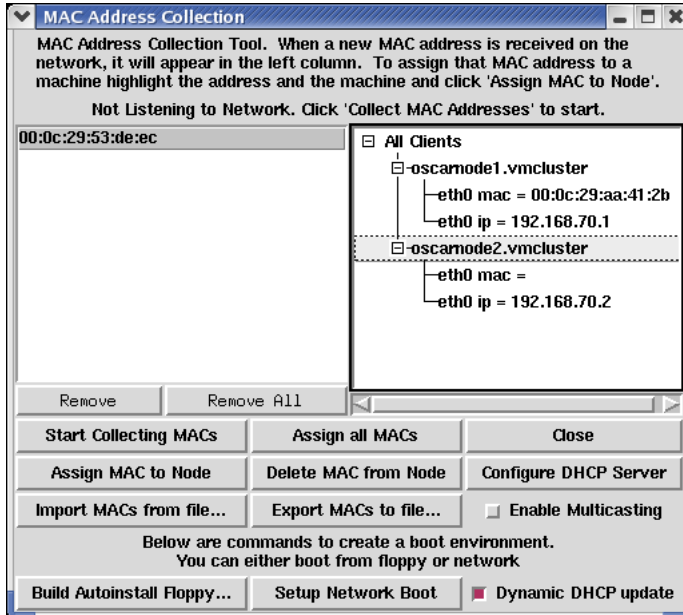


Figure 27: Setup networking: Scanning network, found first MAC address, then later assigned all MAC addresses.


```

Loading initrd-s..... ready.
Loading kernel-s.....Ok..... booting the kernel.
Linux version 2.2.20RAID (root@ia32.ncsa.uiuc.edu) (gcc version 2.95.2 20000220
(Debian GNU/Linux) #1 Thu Apr 4 20:01:46 UTC 2002
BIOS-provided physical RAM map:
 BIOS-e820: 0009f000 @ 00000000 (usable)
 BIOS-e820: 05f00000 @ 00100000 (usable)
Detected 903000 kHz processor.
Console: colour UGA+ 80x50
Calibrating delay loop... 1900.54 BogoMIPS
Memory: 94332k/98304k available (1384k kernel code, 412k reserved, 1228k data, 8
24k init)
Bentry hash table entries: 16384 (order 5, 128k)
Buffer cache hash table entries: 131072 (order 7, 512k)
Page cache hash table entries: 32768 (order 5, 128k)
Intel machine check architecture supported
Intel machine check reporting enabled on CPU#0.
256K L2 cache (8 way)
CPU: L2 cache: 256K
CPU: Intel Pentium III (Coppermine) stepping 0a
Checking 386/387 coupling... OK, FPU using exception 16 error reporting.
Checking halt instruction... OK
POSIX conformance testing by UNIFIX
PCI: PCI BIOS revision 2.10 entry at 0xfd980
PCI: Using configuration type 1
PCI: Probing PCI hardware
PCI: Enabling memory for device 00:80
Linux NET4.0 for Linux 2.2
Based upon Swansea University Computer Society NET3.039
NET4: Unix domain sockets 1.0 for Linux NET4.0.
NET4: Linux TCP/IP 1.0 for NET4.0
IP Protocols: ICMP, UDP, TCP
TCP: Hash tables configured (ehash 131072 bhash 65536)
Starting v 1.5
Serial driver version 4.27 with no serial options enabled
ttyS00 at 0x03e8 (irq = 4) is a 16550A
ttyS01 at 0x03f8 (irq = 3) is a 16550A
ttyS02 at 0x03e8 (irq = 4) is a 16550A
ttyS03 at 0x02e8 (irq = 3) is a 16550A
tty: Stopped by console configuration
Real Time Clock Driver v1.08
RAM disk driver initialized: 16 RAM disks of 8192K size
loop: registered at major 7
PIIX4: IDE controller on PCI bus 00 dev 39
PIIX4: not 100% native mode: will probe iras later
ide0: BM-DMA at 0x1058-0x105f, BIOS settings: hda:DMA, hdb:pio
ide1: BM-DMA at 0x1058-0x105f, BIOS settings: hdc:DMA, hdd:pio
hda: VMware Virtual IDE Hard Drive, ATA DISK drive

```

Figure 28: Booting the client a second time to download the image.

```

/dev/hda4 0 - 0 0 Empty swap
/dev/hda5 157+ 3996- 3840- 3931168+ 83 Linux swap
/dev/hda6 157+ 3996- 3840- 3931168+ 83 Linux
Warning: no primary partition is marked bootable (active)
This does not matter for LILO, but the DOS MBR will not boot this disk.
Successfully wrote the new partition table

Re-reading the partition table ...
 hda: hda1 hda2 < hda5 hda6 >

If you created or changed a DOS partition, /dev/foo7, say, then use dd(1)
to zero the first 512 bytes: dd if=/dev/zero of=/dev/foo7 bs=512 count=1
(See fdisk(8).)
Setting up swap-space version 1, size = 136212480 bytes
Adding Swap: 133016k swap-space (priority -1)
mke2fs 1.25 (20-Sep-2001)
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
7072 inodes, 28192 blocks
1409 blocks (5.00%) reserved for the super user
First data block=1
4 block groups
8192 blocks per group, 8192 fragments per group
1768 inodes per group
Superblock backups stored on blocks:
 8193, 24577

Writing inode tables: done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 26 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
mke2fs 1.25 (20-Sep-2001)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
431520 inodes, 982792 blocks
43139 blocks (5.00%) reserved for the super user
First data block=0
38 block groups
32768 blocks per group, 32768 fragments per group
16384 inodes per group
Superblock backups stored on blocks:
 32768, 98304, 163840, 229376, 294912, 819200, 884736

Writing inode tables: done
Writing superblocks and filesystem accounting information: -

```

Figure 29: Client partitioning disk, setting up the disk tables, and starting to download the image.

E.9 Step 5: Complete Cluster Setup

This step is used to unify the server and client installations into a single cluster. See the details in Section 5.11, page 27.

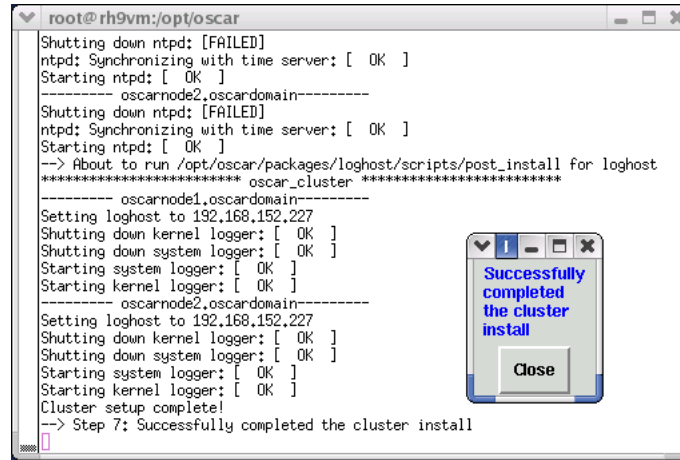


Figure 32: Complete Cluster Setup.

E.10 Test Cluster Setup

This step is used to test the cluster setup. It can either be run from within the wizard, or, as shown here, from manually launching a shell script at a root command prompt. See the details in Section 5.12, page 27.

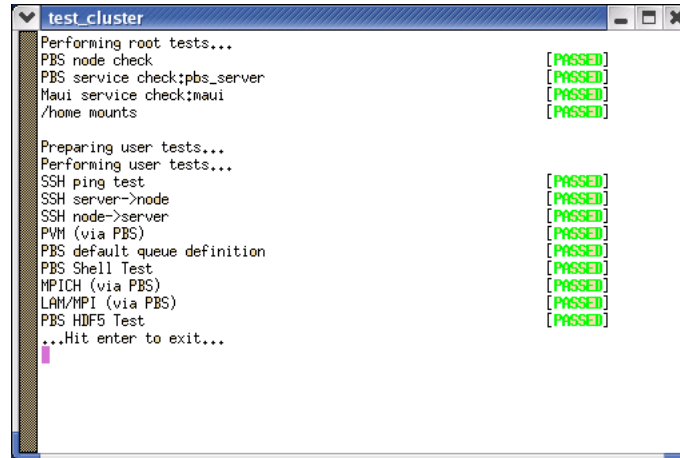


Figure 33: Test the cluster.

E.11 Delete Node Button

The Delete Node button can be used to delete clients from an OSCAR cluster. Note that this button only deletes OSCAR's knowledge of the clients – what physically happens to that client is not OSCAR's concern. This example shows deleting one of the clients setup in the previous sections – `oscardnode2`. The next section (Section E.12) will show adding it back. See the details on deleting clients in Section 5.14.2, page 28.

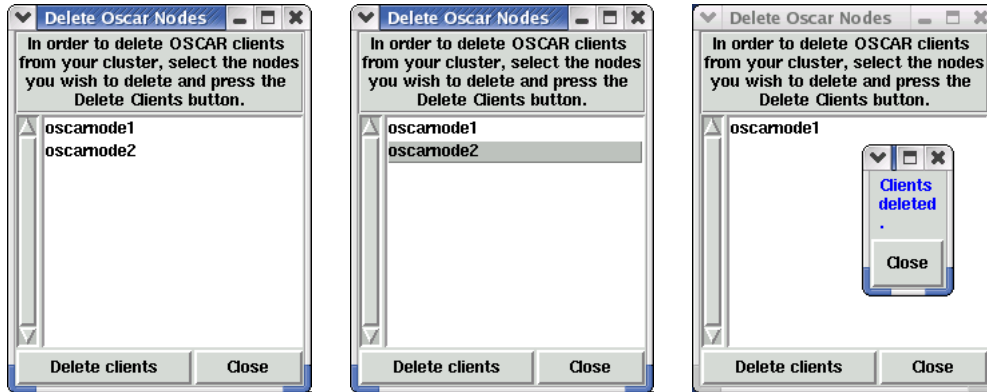


Figure 34: Delete OSCAR clients. First image is the initial window, second image is with a client selected, and third image is when the action has completed.

E.12 Add Node Button

The Add Node button will add clients into an existing OSCAR cluster. In this example, we will add back `oscardnode2` into the cluster. See the details of adding a client in Section 5.14.1, page 28.




Figure 35: Add OSCAR Clients.

E.13 Install/Uninstall OSCAR Packages

The Install/Uninstall OSCAR Packages button will let you add/remove packages from and an OSCAR cluster after the initial installation. See the details and more in-depth look at the underlying concepts of this new feature in Section 5.15, page 29.




Figure 36: Add node / defining the clients.




Figure 37: Add node / setup networking.




Figure 38: Add node / complete cluster setup.




Figure 39: Install/Uninstall packages from an existing OSCAR cluster.




Figure 40: Individual OSCAR packages are color-coded to show their current installation state.