*Chapter 2*

# Preparing for Installation

This chapter contains procedures and information for preparing yourself and your site for software installation. You should review all of the information in this chapter, but you do not need to perform all of the procedures described. The procedures that you do will depend on the type of installation that you plan and specific conditions at your site.

The chapter contains these sections:

- "Selecting a Distribution Source"
- "Setting Up an Installation Server"
- "Backing Up the Target Systems"
- "Planning the Order of Installation"
- "Collecting the Information That You Need"
- "Preparing the Remote CD-ROM Drive"

## Selecting a Distribution Source

When selecting a distribution source, consider the speed and reliability of your network, the frequency with which installations are performed, and the amount of software that will be installed. If disk space is available and the network is fast and reliable, consider creating a centralized distribution directory on an installation server. A centralized directory is particularly useful if users perform their own installations, or if the availability of the server's CD-ROM drive is subject to interruption.

When you are installing software on one or two target systems and the targets contain local CD_ROM drives, using a locally mounted distribution CD is often the most efficient distribution source, particularly if your network is slow and you plan to install a considerable amount of software. For any target that is not equipped with a CD-ROM drive, the distribution source must be a remote CD-ROM drive or distribution directory.

## Setting Up an Installation Server

You can create an installation server on almost any system in your network that is reliable and has adequate disk resources. The distribution source may be a local CD-ROM drive or a distribution directory.

**Note:** Do not create an installation server on "multi-homed" system (routers with more than one IP address). Packets sent to multi-homed systems are not necessarily returned to the sender by the same route. This can cause problems during miniroot installations.

Any system that you plan to use as an installation server must be accessible from the remote targets. This means that communications between the server and targets must support forwarding of boot files (for miniroot installations) and Trivial File Transfer Protocol (TFTP) file transfers. In addition, the installation server must contain a user account that is available to target systems.

The procedures in this section assume that you plan to set up a distribution directory on an installation server. If you plan to use a remote CD-ROM drive as the distribution source, complete all procedures in this section except "Creating a Distribution Directory" on page 19.

### Enabling BOOTP Forwarding on Routers

Inst uses the Internet Bootstrap Protocol (BOOTP) to obtain IP addresses during miniroot installations. For this reason, all routers between an installation server and a remote target must allow *bootp* (see the bootp 1M reference page) forwarding. On factory-shipped Silicon Graphics systems, however, *bootp* forwarding is disabled in the file */etc/inetd.conf*.

**Note:** The *inetd.conf* file is stored in */usr/etc* on systems running versions of IRIX that are earlier than 5.2.

The procedure below describes how to identify routers and change the */etc/inetd.conf* file to enable *bootp* forwarding on them. This procedure might not be appropriate for routers that are not Silicon Graphics systems.

1. Determine the names of routers.

   Enter this command from the target system:

**ping -R -c 1** *installation_server*

Look for an entry that begins with *RR:* in *ping* output. It shows the route of a packet from the target to the installation server and back. Each node listed, except the installation server and target, is a router.

2. Become the superuser on the router.

```
% su
Password:
#
```

3. Change the *bootp* line in */etc/inetd.conf* on each router, as needed.

   The default *bootp* line in */etc/inetd.conf* looks like this:

```
bootp  dgram   udp   wait  root  /usr/etc/bootp  bootp
```

   Add the **-f** flag to enable *bootp* forwarding. The modified line should look like this:

```
bootp  dgram   udp   wait  root  /usr/etc/bootp  bootp -f
```

4. Put your changes into effect:

```
# killall -v -HUP inetd
```

5. Repeat steps 2 and 3 for additional routers.

When the server is no longer needed for software installation, you can return the */etc/inetd.conf* file on the router to its original state, then enter the *killall* command to put the change into effect (see the killall(1M) reference page).

## Enabling TFTP Access on an Installation Server

During miniroot installations, Inst uses TFTP to transfer files from remote systems. For this reason, any system acting as the server for miniroot installations must allow *tftp* access. However, on factory-shipped Silicon Graphics systems, *tftpd*(1M) access is disabled in the */etc/inetd.conf* file.

This procedure explains how to change the */etc/inetd.conf* to allow *tftp* access.

**Note:** If your site has strict security requirements, step 2 of this procedure offers an option that limits *tftp* access to the distribution directory.

1. Become the superuser on the installation server.

```
% su
Password:
#
```

2. Modify the *tftp* line in */etc/inetd.conf* on the server.

   You can modify */etc/inetd.conf* to allow *tftp* access in a way that suits the needs of your site. The default *tftp* entry in */etc/inetd.conf* looks similar to this (your entry might contain additional arguments to the **-s** option):

```
tftp dgram udp wait guest /usr/etc/tftpd tftpd –s /usr/local/boot
```

   Take any of these actions to modify the *tftp* entry:

   - Remove the *–s /usr/local/boot*, and any additional directories that are specified by the *-s* option, from the entry to allow *tftp* access to all publicly readable directories. (This is not recommended for sites on the Internet.) The modified line should look like this:

```
tftp dgram udp wait guest /usr/etc/tftpd tftpd
```

   - Specify access to the distribution directory only (suggested for secure sites). The modified line should look like this:

```
tftp dgram udp wait guest /usr/etc/tftpd tftpd –s
                            /usr/local/boot <distdir>
```

   - Specify access to the CD-ROM drive. The modified line should look like this:

```
tftp dgram udp wait guest /usr/etc/tftpd tftpd –s
                           /usr/local/boot <CDdir>/dist
```

   **Note:** If an account other than *guest* is used for installation, replace the *guest* entry on the *tftp* line with the alternate user ID.

3. Put your changes into effect.

```
# killall –v –HUP inetd
```

When the remote system is no longer needed for software installation, you can return the */etc/inetd.conf* file to its original state, then enter the *killall* command to put your changes into effect.

## Configuring an Installation Account

During an installation, Inst defaults to using the *guest* account on the server to accept a connection from the target systems; *guest* must not be password protected. If the *guest* account on the server is either unavailable or password protected (to limit access to the target by *rsh*, for example), you must provide an alternate means for accessing the server. You can allow access to the server in any of these ways:

- Remove the password from *guest* while installations are taking place.
- Use an account other than *guest* on the server (the alternate account must not be password protected) and specify the alternate account when you start Inst (see "Specifying the Source on the Command Line" on page 84 for details on invoking Inst with the *-f* flag):

  **inst -f** *alternate_user***@***installation_server:distdir*

- Use a password-protected account on the server for installations and create an *.rhosts* file for the installation account. The installation account must have read permissions on the distribution source.

  The *.rhosts* file that you create must contain an entry for each target system that will access the server (see the hosts.equiv(4) reference page for information on *.rhosts*). For example, assume that the installation account on the server is *instuser*. The file */usr/people/guest/.rhosts* on the installation server contains these lines to permit installations on targets called *joesbox* and *lab1*:

  ```
  joesbox.engr.xxx.com   instuser
  lab1.engr.xxx.com      instuser
  ```

  Either of the following commands can be used to install software on *joesbox* or *lab1* (see "Step 2: Specifying the Source" on page 42 and "Specifying the Source on the Command Line" on page 84 for details on these commands):

  ```
  Inst> from instuser@installation_server:path
  # inst -f instuser@installation_server:path
  ```

## Creating a Distribution Directory

Follow the procedure below to create a distribution directory on the installation server. The CD-ROM drive from which you copy the distribution software may be either a local or remote drive. You can copy several CDs to the distribution directory if you wish; however, do not include more than one release of a given product in the directory—different distribution directories must be created for different releases of a product.

1. Become the superuser on the server with the CD-ROM drive.

   ```
   % su -
   Password:
   #
   ```

   **Note:** If the distribution directory is on a different system, become the superuser on that system also.

2. Get the controller and unit numbers for the CD-ROM drive.

   ```
   # hinv
   ```

   Look for a line like this in the output of *hinv* and note numbers that appear in parentheses beside the *scsi* and *cdrom* entries:

   ```
   SCSI CDROM: scsi (0) cdrom (4)
   ```

3. Eject the CD that is currently in the drive, if any.

   You can omit the argument from this command if there is only one CD-ROM drive on the workstation:

   ```
   # eject  /dev/scsi/sccntlrdunitl0
   ```

   **Note:** This command has a lowercase *l* and a zero at the end of it.

4. Create the directory for the distribution.

   ```
   # mkdir distdir
   ```

5. Stop the CD-ROM daemon, if it is running.

   Enter this command on the system with the CD-ROM drive to stop the mediad daemon (see the mediad1M) reference page):

   ```
   # mediad -k
   ```

6. Create a mount-point directory for the CD-ROM if none exists.

   ```
   # mkdir /CDROM
   ```

7. Insert the CD containing the distribution and mount it.

   ```
   # mount -o ro /dev/dsk/dkscntlrdunits7 /CDROM
   ```

8. Copy the distribution to the new directory.

   On the system containing the distribution directory, enter the *cp* command to copy the installable software from the CD to the distribution directory. The following sample commands illustrate how to copy all distribution software to the directory, creating the necessary subdirectories simultaneously. Use the first command if the CD-ROM is local and the second command if the CD-ROM is remote:

```
# cp -r /CDROM/dist/. distdir
# rcp -r guest@server:/CDROM/dist/. distdir
```

You can create a software distribution directory that contains fewer products than are in the CD-ROM distribution by copying the files for just the products that you want. Remember that distribution directories and CD-ROM distributions have an identical structure. For example, to copy just one product from a local CD-ROM drive, use this command:

```
% cp /CDROM/dist/product* distdir/product
```

9. Copy the release notes to the server (optional).

   Enter this command to create a directory for the release notes and copy them to the installation server:

   ```
   # cp -r /CDROM/relnotes/* relnotes_dir
   ```

10. Eject the distribution CD.

    Enter this command on the system with the CD-ROM drive. You can omit the argument if there is just one CD-ROM drive on the system:

    ```
    # eject /dev/scsi/sccntlrdunit10
    ```

11. Repeat steps 7 through 9 for additional CDs.

12. Restart the CD-ROM daemon, if you stopped it.

    If you stopped *mediad* in step 5, restart it by entering this command on the system with the CD-ROM drive:

    ```
    # mediad
    ```

    **Note:** If you have more than one CD-ROM drive, you might get an error from this command, but you can safely ignore it.

## Backing Up the Target Systems

Although backing up the target is not requirement for installation, it is strongly recommended. The *Personal System Administration Guide* describes general backup procedures and using System Manager to perform backups. Other programs you can use to make backups are *backup*, *bru*, *cpio*, and *tar*. For more information, see "IRIX Admin: Backup, Security, and Accounting" and section 1 of the IRIX reference pages for these programs.

Consider backing up these files:

- *User files* that are created or copied to the target. Any file on the target that was not put there during the software installation process is considered a user file.
- *Configuration files* that contain information unique to the target system or the site. These files are created during installation but are likely to be modified after they are installed. The unique information in these files is not destroyed during an installation. However, the pre-installation copy of these configuration files is helpful if you decide to go back to the earlier software release after installation. You can identify modified configuration files with this command:

  ```
  # showfiles -c -m -s
  ```

  Use this command to back up the configuration files onto tape:

  ```
  # showfiles -c -m -s | tar cv -
  ```

  Use this command to retrieve specified configuration files from tape and overwrite existing files:

  **Note:** If you follow these instructions, you will overwrite the installed configuration file. Refer to "Managing Configuration Files" on page 62 for more information on how configuration files are handled during installations.

  ```
  # tar xv files
  ```

  Use this command to retrieve all configuration files from tape and overwrite existing files:

  ```
  # tar xv
  ```

## Planning the Order of Installation

Inst automatically manages the installation order of the products on a single CD or in one distribution directory. However, if you expect to install software from more than one CD or distribution directory, you must plan the installation order, since some products require that other products be installed first.

If you are installing from multiple CDs, use the sequence numbers on their labels to put the CDs in order. Install the CDs starting with the lowest sequence number first. Use these guidelines to plan the order of your installation:

- For miniroot installations, be sure that the first CD or distribution directory that you install contains installation tools. CDs containing installation tools are clearly marked.

- If any CD has two sequence numbers, that CD is used twice during the installation. If you find no intervening sequence numbers, you need to insert the CD only once during the installation.
- If you have two or more CDs with the same sequence number, the order of those CDs relative to each other does not matter. For example, assume that you have four CDs with sequence numbers 400, 500, 600, and 600. The CD labeled 400 is installed first, followed by the CD labeled 500. The order of the two CDs labeled 600 is irrelevant, as long as they are installed last.

If you are installing from several distribution directories, check to see what products are in each directory. Plan to install the products in this order:

- installation tools (for miniroot installations)
- operating system software
- communications software
- compilers
- optional software

## Collecting the Information That You Need

When you install software from a remote distribution source, you are required to supply information about the installation server and target. If your distribution source is a local CD-ROM drive, you may have to supply the controller and unit number of the drive. Jot down the information that you get as you complete these instructions and have it available when you start the installation.

### Getting CD-ROM Device Numbers

If you are planning a miniroot installation from the local CD-ROM drive on the target, you must be prepared to specify the controller and unit number of the drive during the installation. Use this command to determine the controller and unit number:

```
# hinv
```

Look for a line like the one below in the output of *hinv*. The numbers that you need appear in parentheses beside the *scsi* and *cdrom* fields:

```
SCSI CDROM: scsi (0) cdrom (4)
```

In the previous example, the controller number is *0* and the unit number is *4*.

## Getting the Target's Name and Address

The name and address of the target system are required for miniroot installations when the distribution source is on a remote installation server. Enter the *hostname* command from the target system to get its name:

```
% hostname
```

Enter the *ping* command from the target to get its Internet Protocol (IP) address. Use the output of the *hostname* command as an argument to *ping*:

```
% /usr/etc/ping -c 1 target_name
```

## Getting the Server's Name and Address

The name of the remote installation server is specified to identify the distribution source in miniroot and live installations. If you are already on the installation server, enter the *hostname* command to get its name:

```
% hostname
```

If you are on the target system, enter this *ping* command to get the Internet Protocol (IP) address of the installation server:

```
% /usr/etc/ping -c 1 installation_server
PING server (IPaddress): 2048 data bytes
----server PING Statistics----
100 packets transmitted, 100 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 0/2/7
```

Note the numbers in the *IPaddress* field of the output. If *ping* output is not similar to this example, the network connection might be faulty. You should resolve the problem before continuing with your installation (see Appendix B, "Troubleshooting Installations" for help).

### Getting the Installation Account Name

By default, Inst uses the *guest* account on the installation server for installations. If you plan to use an account other than *guest*, be prepared to specify the name of the alternate account (see "Configuring an Installation Account" on page 19).

## Preparing the Remote CD-ROM Drive

You can use a utility called *cdinstmgr* to manage the mounting and unmounting of CDs on a remote CD-ROM drive during installations. This utility prevents other users from ejecting distribution CDs during an installation and informs you of error conditions. You must start *cdinstmgr* before you begin the installation.

Use this procedure to start *cdinstmgr*:

1. Become root and enter the *cdinstmgr* command on the installation server.

   ```
   % su
   # cdinstmgr
   ```

2. Respond to *cdinstmgr* prompts.

   Follow *cdinstmgr* directions as it prompts for the name of the target and guides you through the mechanics of loading the first CD. Be sure to insert CDs with the label up. If mediad is not already running on the installation server, *cdinstmgr* prompts you for the superuser password to the server.

3. Note the mount directory for the CD-ROM drive.

   When you see the message below, note the value of *CDdir*, the mount point directory; you will need it later.

   ```
   CD on CDdir for host hostame: type the word "done" when you are
   finished with this CD, "quit" if you are completely done:
   ```

   You will respond to this prompt after you have finished using this CD.

4. Verify that the target can access the distribution source.

   Enter this command from the target to verify that it has access to the distribution CD:

   ```
   % rsh server -l user 'ls CDdir/dist'
   ```

   If the output to this command is a list of products contained on the distribution CD, the target has access to the distribution CD. If the output is an error message, you should diagnose and correct the error before beginning the installation.