

Tru64 UNIX

Guide to Preparing Product Kits

Part Number: AA-RH9SC-TE

August 2000

Product Version: Tru64 UNIX Version 5.1 or higher

This manual describes the procedures for creating, delivering, and installing layered product kits for use on Compaq Tru64™ UNIX operating systems.

© 2000 Compaq Computer Corporation

COMPAQ, the Compaq logo, AlphaServer, and TruCluster Registered in the U.S. Patent and Trademark Office. Tru64 is a trademark of Compaq Information Technologies Group, L.P.

UNIX is a trademark of The Open Group. All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE, OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES AND WHETHER IN AN ACTION OF CONTRACT OR TORT, INCLUDING NEGLIGENCE.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

Contents

About This Manual

1 Introducing Product Kits

1.1	Overview	1-1
1.2	Product Types	1-2
1.3	Kit Formats	1-3
1.4	Kit-Building Process	1-3

2 Creating Kit Directories

2.1	Obtaining a Unique Product Code	2-1
2.2	Creating a Kit Building Directory Structure	2-1
2.3	Populating the Source Directory	2-3
2.3.1	Using Standard Directory Structures	2-3
2.3.2	Using Context-Dependent Symbolic Links	2-5
2.3.2.1	Knowing When to Use CDSLs	2-6
2.3.2.2	Identifying CDSLs	2-7
2.3.2.3	Creating CDSLs	2-7
2.3.2.4	Restrictions	2-8
2.3.3	Placing Files in the Kit Source Directory	2-9
2.3.4	Setting Up the Sample Product Kit Source Directory	2-10

3 Creating Subsets

3.1	Grouping Files into Subsets	3-1
3.2	Creating the Master Inventory File	3-3
3.3	Creating the Key File	3-6
3.4	Creating Subset Control Programs	3-9
3.4.1	Creating SCP Source Files	3-10
3.4.2	Setting Up Initial SCP Processing	3-11
3.4.2.1	Including Library Routines	3-11
3.4.2.2	Setting Global Variables	3-13
3.4.3	Working in a Cluster Environment	3-14
3.4.4	Working in a Dataless Environment	3-16
3.4.5	Associating SCP Tasks with setld Utility Phases	3-17
3.4.5.1	Displaying the Subset Menu (M Phase)	3-19

3.4.5.2	Before Loading the Subset (PRE_L Phase)	3-21
3.4.5.3	After Loading the Subset (POST_L Phase)	3-23
3.4.5.4	After Securing the Subset (C INSTALL Phase)	3-26
3.4.5.5	Before Deleting a Subset (C DELETE Phase)	3-28
3.4.5.6	Before Deleting a Subset (PRE_D Phase)	3-29
3.4.5.7	After Deleting a Subset (POST_D Phase)	3-31
3.4.5.8	Verifying the Subset (V Phase)	3-32
3.4.6	Stopping the Installation	3-32
3.4.7	Creating SCPs for Different Product Kit Types	3-33
3.4.7.1	User Product Kit SCPs	3-33
3.4.7.2	Kernel Product Kit SCPs	3-35
3.5	Producing Subsets	3-39
3.5.1	Compression Flag File	3-43
3.5.2	Image Data File	3-43
3.5.3	Subset Control Files	3-44
3.5.4	Subset Inventory File	3-46
3.6	Testing Subsets	3-48
3.6.1	Loading All Subsets	3-48
3.6.2	Repeating Subset Load	3-49
3.6.3	Removing All Subsets	3-50
3.6.4	Loading Mandatory Subsets Only	3-50
3.6.5	Testing in a Cluster	3-51
3.6.5.1	Loading the Kit onto a Cluster	3-52
3.6.5.2	Deleting the Kit from a Cluster	3-53
3.7	Updating Inventory After Creating Subsets	3-53

4 Producing User Product Kits

4.1	Overview	4-1
4.2	Producing Distribution Media	4-1
4.2.1	Editing the /etc/kitcap File	4-3
4.2.1.1	Disk Media Descriptor	4-4
4.2.1.2	Tape Media Descriptor	4-5
4.2.2	Building a User Product Kit on Disk Media	4-6
4.2.3	Building a User Product Kit on Magnetic Tape	4-8
4.3	Testing the Distribution Media	4-10

5 Producing Kernel Product Kits

5.1	Overview	5-1
5.2	Creating Additional Installation Files	5-2
5.2.1	The files File Fragment	5-5
5.2.2	The sysconfigtab File Fragment	5-6

5.2.3	The Object Module File	5-7
5.2.4	The Source and Header Files	5-7
5.2.5	The Method Files	5-7
5.3	Producing Distribution Media	5-8
5.3.1	Editing the /etc/kitcap File	5-10
5.3.1.1	Disk Media Descriptor	5-11
5.3.1.2	Tape Media Descriptor	5-12
5.3.2	Building a Kernel Product Kit on Disk Media	5-13
5.3.3	Building a Kernel Product Kit on Magnetic Tape	5-15
5.4	Testing the Distribution Media	5-16
5.4.1	Testing a Kernel Product Kit with the setld Utility	5-17
5.4.2	Testing a Kernel Product Kit in a RIS Area	5-20

A Creating a Consolidated Firmware CD-ROM

A.1	Build Instructions	A-1
A.1.1	Prepare for the Build Session	A-1
A.1.2	Build the Consolidated Firmware CD-ROM	A-2
A.2	Sample Build Session	A-4
A.2.1	Preparing for the Build Session	A-5
A.2.2	Building the Consolidated Firmware CD-ROM	A-5

Glossary

Index

Examples

2-1	Creating the Kit-Building Directories	2-3
3-1	Sample ODB Kit Master Inventory File	3-5
3-2	Sample ODB Kit Key File	3-6
3-3	Sample setld Installation Menu	3-20
3-4	Sample Test for Alpha Processor During M Phase	3-20
3-5	Sample Backup of Existing Files During PRE_L Phase	3-22
3-6	Sample Backward Link Creation During POST_L Phase	3-25
3-7	Sample Message Output During C INSTALL Phase	3-27
3-8	Sample C DELETE Phase	3-28
3-9	Sample PRE_D Phase Reversal of POST_L Phase Actions	3-30
3-10	Sample File Restoration During POST_D Phase	3-31
3-11	Sample ODB User Product SCP	3-33
3-12	Sample ODB Kernel Product SCP	3-36

3-13	Using the kits Utility to Build ODB Subsets	3-40
3-14	Sample Image Data File	3-44
3-15	Sample Subset Control File	3-45
3-16	Sample ODB Product Subset Inventory File	3-47
4-1	Sample Disk Media Descriptor for User Product	4-5
4-2	Sample Tape Media Descriptor for User Product	4-6
4-3	Sample gendisk Command for User Product	4-8
4-4	Sample gentapes Command for User Product	4-9
5-1	Sample Disk Media Descriptor for Kernel Product	5-12
5-2	Sample Tape Media Descriptor	5-13
5-3	Sample gendisk Command	5-15
5-4	Sample gentapes Command	5-16

Figures

1-1	Steps in the Kit-Building Process	1-4
1-2	Kit Directory Hierarchies	1-5
2-1	Kit Directory Hierarchies	2-2
2-2	Member-Specific Directories in Source Hierarchy	2-6
2-3	Sample Product Kit Source Directory	2-11
3-1	ODB Product Subsets and Files	3-2
3-2	Time Lines for setld Utility Phases	3-18
3-3	ODB output Directory	3-43
4-1	User Product Kit File Formats	4-3
5-1	Kernel Product Source Directory	5-3
5-2	Editing the files File Fragment	5-6
5-3	Editing the sysconfigtab File Fragment	5-7
5-4	Kernel Product Kit File Formats	5-9
5-5	Static Configuration of a Driver	5-19

Tables

1-1	Using This Manual	1-2
2-1	File Locations in Kit Directories	2-9
3-1	Master Inventory File	3-4
3-2	Key File Product Attributes	3-7
3-3	Key File Subset Descriptors	3-9
3-4	Available SCP Library Routines	3-12
3-5	Available Library Installation Routines	3-12
3-6	Available Installed Software Database Library Routines	3-12
3-7	STL_ScpInit Global Variables	3-13
3-8	SCP Operations on a Cluster	3-15
3-9	Installation Control Files in the instctrl Directory	3-41

3-10	Image Data File Field Descriptions	3-44
3-11	Subset Control File Field Descriptions	3-44
3-12	Subset Inventory File Field Descriptions	3-46

About This Manual

A product kit is the standard mechanism by which software products are delivered to and maintained on a Compaq Tru64 UNIX operating system. This manual describes the procedures for creating, installing, and managing the collections of files and directories that make up a product kit to be installed on a customer's system. Kits can be distributed on CD-ROM, diskette, or magnetic tape.

Audience

This manual is intended for software developers who are responsible for creating product kits. They are expected to have experience with UNIX based operating systems, shell script programming, and system administration.

New and Changed Features

The following list describes the major changes made to this manual:

- New electronic mail address to request product codes
- More information about using context-dependent symbolic links (CDSLs) and preparing product kits to run on clusters
- Updated examples with CDSLs
- Removed references to hardware product kits and new hardware delivery (NHD)

Previous versions of the manual are available on the World Wide Web at the following location: http://www.tru64unix.compaq.com/faqs/publications/pub_page/doc_list.html.

Organization

This manual is organized as follows:

<i>Chapter 1</i>	Introduces the kit-building process
<i>Chapter 2</i>	Describes how to create and populate kit directories
<i>Chapter 3</i>	Describes how to organize product files into subsets, create kit production files, and produce the subsets and related control files
<i>Chapter 4</i>	Describes how to create, test, and deliver user product kits

<i>Chapter 5</i>	Describes how to create, test, and deliver kernel product kits
<i>Appendix A</i>	Describes how to create a CD-ROM that lets you upgrade your processor firmware at the same time that you install the operating system
<i>Glossary</i>	Defines terms used in this manual

Related Documentation

Icons on Tru64 UNIX Printed Books

The printed version of the Tru64 UNIX documentation uses letter icons on the spines of the books to help specific audiences quickly find the books that meet their needs. (You can order the printed documentation from Compaq.) The following list describes this convention:

- G Books for general users
- S Books for system and network administrators
- P Books for programmers
- D Books for device driver writers
- R Books for reference page users

Some books in the documentation help meet the needs of several audiences. For example, the information in some system books is also used by programmers. Keep this in mind when searching for information on specific topics.

The *Documentation Overview* provides information on all of the books in the Tru64 UNIX documentation set.

The Tru64 UNIX documentation set is available on the World Wide Web at the following URL: **http://www.tru64unix.compaq.com/faqs/publications/pub_page/pubs_page.html**

You may find the following documents helpful when preparing product kits:

Sharing Software on a Local Area Network

This manual describes Remote Installation Services (RIS) and Dataless Management Services (DMS). RIS is used to install software across a network instead of using locally mounted media. DMS allows a server system to maintain the `root`, `/usr`, and `/var` file systems for client systems. Each client system has its own `root` file system on the server, but shares the `/usr` and `/var` file systems.

This manual may be helpful if you are preparing a product kit that will be installed in a RIS environment.

Writing Device Drivers

This manual provides information for systems engineers who write device drivers for hardware that runs the operating system. Systems engineers can find information on driver concepts, device driver interfaces, kernel interfaces used by device drivers, kernel data structures, configuration of device drivers, and header files related to device drivers.

This manual may be helpful if you are preparing product kits for a device driver.

Installation Guide

This manual describes the procedures to perform an Update Installation or a basic installation of the operating system on all supported processors and single-board computers. It explains how to prepare your system for installation, boot the processor, and perform the installation procedure.

Installation Guide — Advanced Topics

This manual describes the procedures for performing an advanced installation of the operating system on all supported processors and single-board computers.

System Administration

This manual describes how to configure, use, and maintain the operating system. It includes information on general day-to-day activities and tasks, changing your system configuration, and locating and eliminating sources of trouble. This manual is intended for the system administrators responsible for managing the operating system. It assumes a knowledge of operating system concepts, commands, and configurations.

Reference Pages Sections 8 and 1m

This section describes commands for system operation and maintenance. It is intended for system administrators. In printed format, this section is divided into two volumes.

Release Notes

The *Release Notes* describe known problems you might encounter when working with the operating system and provides possible solutions for those problems. The printed format also contains information about new and changed features of the operating system, as well as plans to

retire obsolete features of the operating system. Obsolete features are features that have been replaced by new technology or otherwise outdated and are no longer needed. The *Release Notes* are intended for anyone installing the operating system or for anyone using the operating system after it is installed.

Reader's Comments

Compaq welcomes any comments and suggestions you have on this and other Tru64 UNIX manuals.

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UBPG Publications, ZKO3-3/Y32
- Internet electronic mail: `readers_comment@zk3.dec.com`

A Reader's Comment form is located on your system in the following location:

```
/usr/doc/readers_comment.txt
```

- Mail:

Compaq Computer Corporation
UBPG Publications Manager
ZKO3-3/Y32
110 Spit Brook Road
Nashua, NH 03062-2698

A Reader's Comment form is located in the back of each printed manual. The form is postage paid if you mail it in the United States.

Please include the following information along with your comments:

- The full title of the book and the order number. (The order number is printed on the title page of this book and on its back cover.)
- The section numbers and page numbers of the information on which you are commenting.
- The version of Tru64 UNIX that you are using.
- If known, the type of processor that is running the Tru64 UNIX software.

The Tru64 UNIX Publications group cannot respond to system problems or technical support inquiries. Please address technical questions to your local system vendor or to the appropriate Compaq technical support office. Information provided with the software media explains how to send problem reports to Compaq.

Conventions

The following conventions are used in this manual:

<code>%</code>	
<code>\$</code>	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Bourne, Korn, and POSIX shells.
<code>#</code>	A number sign represents the superuser prompt.
<code>% cat</code>	Boldface type in interactive examples indicates typed user input.
<i>file</i>	Italic (slanted) type indicates variable values, placeholders, and function argument names.
[] { }	In syntax definitions, brackets indicate items that are optional and braces indicate items that are required. Vertical bars separating items inside brackets or braces indicate that you choose one item from among those listed.
...	In syntax definitions, a horizontal ellipsis indicates that the preceding item can be repeated one or more times.
<code>cat(1)</code>	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, <code>cat(1)</code> indicates that you can find information on the <code>cat</code> command in Section 1 of the reference pages.
<code>Return</code>	In an example, a key name enclosed in a box indicates that you press that key.
<code>Ctrl/x</code>	This symbol indicates that you hold down the first named key while pressing the key or mouse button that follows the slash. In examples, this key combination is enclosed in a box (for example, <code>Ctrl/C</code>).

Introducing Product Kits

This guide is intended to provide product and kit developers with the proper method to create a product kit. The following topics are discussed in this chapter:

- An overview of product kits (Section 1.1)
- Defining the different product types and describing the sample product used in this guide (Section 1.2)
- Describing the available formats for layered product kits (Section 1.3)
- Illustrating the kit-building process (Section 1.4)

1.1 Overview

A product kit is the collection of files and directories that represent new or upgraded software to be installed onto a customer's system. The kit contains not only the actual files and directories that compose the product, but also includes the supporting files that are required to install the product on the system. The product kit is the standard mechanism by which most products are delivered and maintained on a customer's system. Kits for user and kernel products can be distributed on a CD-ROM, diskette, or tape for installation onto the customer's system.

Note

A consolidated firmware CD-ROM lets you upgrade your processor firmware at the same time that you install the operating system. Appendix A describes how to create a consolidated firmware CD-ROM.

Before building a kit, consider the kind of product for which you are building a kit:

- Does it run in user space or kernel space?
- Is it used during the initial installation and bootstrap of the operating system?

The answers to these questions determine the type of format you choose, the type of medium you use to distribute the kit, and the installation procedures that your users run when they install the kit on their systems.

This chapter helps you answer these questions. It describes the product types supported by the kit-building process and the options for packaging and installing the kit on the customer's system. It leads you through the steps involved in building kits for the various kinds of products, and it describes the installation options that the operating system supports.

Once you determine the type of product kit that you are creating, you can use the specific chapters in this manual as shown in Table 1-1:

Table 1-1: Using This Manual

ALL product kits	
1. Introducing Product Kits	
2. Creating Kit Directories	
3. Creating Subsets	
ONLY user product kits	ONLY kernel product kits
4. Producing User Product Kits	5. Producing Kernel Product Kits

This manual uses the fictitious Orpheus Document Builder (ODB) product to demonstrate how to build kits for each product type. OAT is the code assigned to *Orpheus Authoring Tools, Inc.*, the fictitious company developing the ODB product, and 100 is the product version number. The same product is used for each type of product kit, but Chapter 4 and Chapter 5 describe the files specific to user and kernel product kits.

1.2 Product Types

The process described in this book lets you deliver layered products for a customer's system. A layered product is any software product that is not part of the base operating system. There are three kinds of layered products:

- **User product**
A user product runs in user space. Commands and utilities are in this category, as are applications such as text editors and database systems. Users interact directly with user products, for example, through commands or window interfaces.
- **Kernel product**
A kernel product runs in kernel space. Users do not directly run kernel products, but the operating system and utilities access them to perform their work. For example, a device driver is one common type of kernel

product. A user runs an application or utility, which generates system requests to perform operations such as opening a file or writing data to a disk. The system determines which device driver should service this request and then calls the required driver interface.

1.3 Kit Formats

Before being copied onto the distribution media (diskette, CD-ROM, or tape), the product files are gathered into subsets. A subset groups together related files and specifies whether they are required or optional for the product. You can copy the product files onto the distribution media in one of the following formats:

- Compressed `tar` format

In compressed `tar` format, the product files belonging to the same subset are written to the distribution media as a single file. Kits for user and kernel products should be produced in compressed `tar` format. During installation, the `setld` utility uncompresses the files and moves them onto the customer's system, preserving the files' original directory structure. The `gentapes` and `gendisk` utilities can create kits in compressed `tar` format.

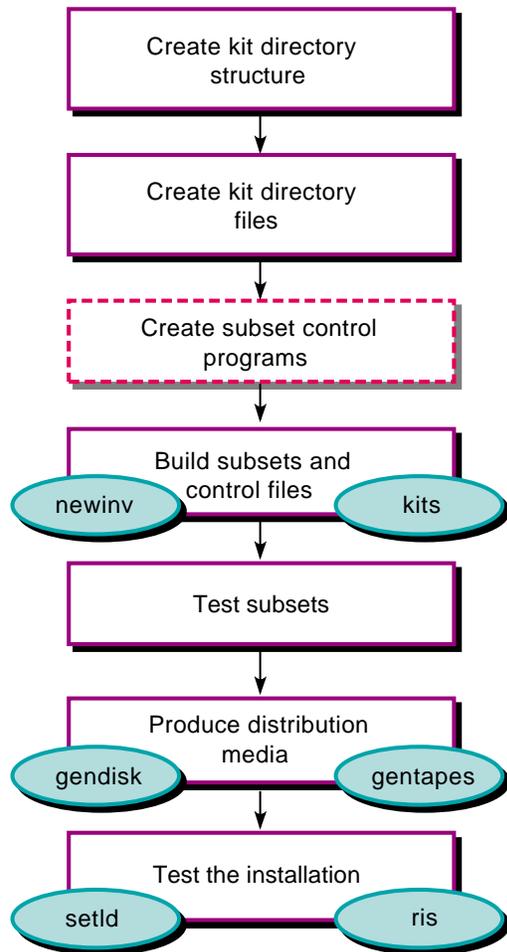
- Direct CD-ROM (DCD) format

In direct CD-ROM (DCD) format, the files are written to any disk media (CD-ROM, hard disk, or diskette) as a UNIX file system (UFS). Subsets distributed in DCD format cannot be compressed. The `gendisk` utility can create kits in DCD format.

1.4 Kit-Building Process

Figure 1-1 illustrates the process of creating and packaging a kit. In the figure, boxes drawn with dashed lines represent optional steps; for example, you do not have to create subset control programs if your kit requires no special handling when it is installed. In Figure 1-1, the commands enclosed in ovals perform the associated steps of the kit-building process.

Figure 1–1: Steps in the Kit-Building Process



ZK-0460U-AI

The kit-building process consists of the following steps:

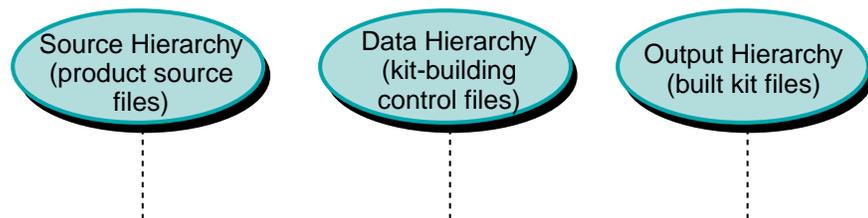
1. Create the kit directory structure that contains the source files.

On the development system, create the following directory structure for the kit you want to build:

- A source hierarchy, which contains all the files that make up the product
- A data hierarchy, which contains files needed to build the kit
- An output hierarchy, which holds the result of the kit-building process — one or more subsets that make up the product kit

Figure 1–2 illustrates these directory hierarchies.

Figure 1–2: Kit Directory Hierarchies



ZK-0461U-AI

This directory structure is the same for user product kits and kernel product kits. Only the contents of these directories differ among the product types. For example, a kernel product kit needs additional files that are unique to this specific kit type. Refer to Chapter 4 and Chapter 5 for additional information about the requirements for each product kit type.

2. Create kit production files.

This includes a master inventory file containing information about each file in the subset, a key file to define product attributes such as the product name, product version, subset definitions, and additional files for kernel product kits.

3. Create subset control programs (if needed).

The `setld` utility can call a subset control program (SCP) to perform installation steps specific to your kit. You supply an SCP for your kit only if the product requires special installation steps, such as those needed for kernel product kits. The SCP is optional for user products. Most layered products supply a subset control program, though the actions the programs perform differ for each product type. For example, the subset control program for a kernel product may call the `kreg` utility to maintain the system file that registers kernel layered products, while the subset control program for a user product would not.

4. Build subsets and control files.

Before transferring your kit onto distribution media, organize the product files into subsets. Subsets group together related files. For example, one subset could contain optional product files, while another subset could contain the files required to run the product. The `kits` utility creates subsets according to the specifications you define in the master inventory and key files. Invoke the `kits` utility from the same directory where the master inventory file is located. Refer to Chapter 3 for information about the master inventory and key files.

5. Test subsets.

You must test your subsets to ensure that they can be loaded onto a running system, that the product runs on the system, and that the subsets can be deleted. Subset testing includes loading all subsets onto a running system and deleting all subsets from a running system. If your kit includes optional subsets, you also should load the mandatory subsets onto a running system to determine if the product works as expected. If not, you may have to reclassify some optional subsets as mandatory.

6. Produce distribution media.

When you have created the subsets for the product, you are ready to package the kit. At this point, you must decide whether to create the kit in DCD format or in `tar` format. To do this, use the `gendisk` or `gentapes` utility. If you are creating a kit for a kernel product, you may need to modify the kit and add files to support your system's bootstrap link.

7. Test product kit media.

After you have successfully created the kit, you should test its installation from the new media. Chapter 4 and Chapter 5 tell you how to test the installation of each of the product kit types.

Creating Kit Directories

After you develop a software product, you package the product files to process them into a kit. First you must organize these files by function and use, then place them into a kit-building directory structure. When you design the kit-building directory structure, consider where you want to place the product files on the customer's system and then create a kit directory structure that closely mirrors that on the customer's system.

This chapter discusses the following topics:

- Obtaining a unique three-letter manufacturer's product code (Section 2.1)
- Creating the directory structure needed to build a product kit (Section 2.2)
- Populating the source directory on the kit-building system (Section 2.3)

2.1 Obtaining a Unique Product Code

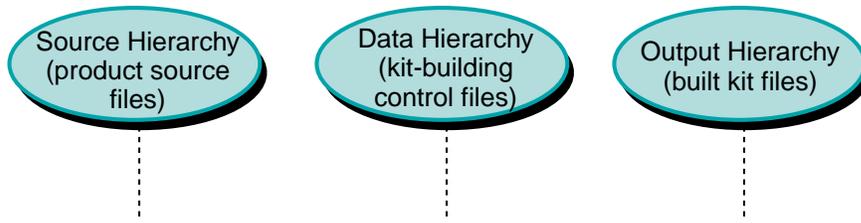
Before you can create a product kit, you must have a unique three-letter product code. To obtain this product code, send electronic mail to `product@dssr.sqp.zko.dec.com`. You use this product code and a product version number that you assign to name your product-specific subdirectories.

Examples in this book use `OAT` as the prefix as the unique three-letter product code for the Orpheus Document Builder (ODB) product kit. Assuming this is the first release of the product, the examples use `100` as the version number.

2.2 Creating a Kit Building Directory Structure

To create a kit, you need three separate directory hierarchies on the kit development system. Figure 2-1 illustrates these directory hierarchies.

Figure 2–1: Kit Directory Hierarchies



ZK-0461U-AI

The following definitions describe each directory hierarchy:

- **Source hierarchy**

The source hierarchy exactly mirrors the directory structure into which customers install your finished kit. You must place each file that is to become part of your kit into the required directory in the source hierarchy. You can create the source hierarchy under any directory you choose.

- **Data hierarchy**

The data hierarchy contains the following files to specify the contents of the kit and how it is organized:

- A master inventory file lists each of the files in the kit and defines which subset contains each file.
- A key file specifies the kit's attributes, such as the product name and version and whether the subsets are in compressed or uncompressed format.
- A subdirectory named `scps` contains any subset control programs that the product requires.
- Additional files may be required, depending on the kit type.

The `kits` utility is run from this data directory. There is no specific requirement for the location of the data hierarchy, but it is good practice to place it under the same directory as the source hierarchy.

- **Output hierarchy**

The output hierarchy contains the results of building the kit in the same format that the distribution media will contain when it is delivered to the customer. There is no specific requirement for the location of the output hierarchy, but it is good practice to place it under the same directory as the source and data hierarchies.

Example 2–1 shows how you can create the kit-building directories:

Example 2–1: Creating the Kit-Building Directories

```
# mkdir /mykit
# mkdir /mykit/src
# mkdir /mykit/data
# mkdir /mykit/output
```

2.3 Populating the Source Directory

The components of a kit can be installed in any directory on the customer's system. However, guidelines exist for kit file placement. The standard system directory structure separates files by function and use and is designed for efficient organization.

This section discusses the following topics:

- Using standard directory structures (Section 2.3.1)
- Using context-dependent symbolic links (Section 2.3.2)
- Placing files in the kit source directory (Section 2.3.3)
- Setting up the sample product kit source directory (Section 2.3.4)

You can choose any method for populating the source hierarchy. For example, you could create a `Makefile` to use with the `make` command, or you could copy files with the `cp` command.

2.3.1 Using Standard Directory Structures

File placement in a standard directory structure depends upon whether the system has a separate `var` file system.

- If the system has a separate `var` file system, use the `/var/opt` directory.
- If the system does not have a separate `var` file system, use the `/usr/var/opt` directory.

Examples in this manual assume that `var` is not a separate file system, and show the `/usr/var/opt` directory.

A standard directory structure has the following advantages:

- Avoiding name conflicts

When a layered product installs a file that overwrites a file shipped by another product, it is known as a name space conflict. Shipping the files in the product-specific `opt` subdirectories of `root`, `usr` and `var` avoids this conflict because each three-letter product code is unique to a particular manufacturer. The product-specific directories for the

examples in this manual are `/opt/OAT100`, `/usr/opt/OAT100`, and `/usr/var/opt/OAT100`.

- **Locating kit components**

If disk partition restructuring or product maintenance becomes necessary, it is easier to find all of your kit if its components are in the `/opt` directories rather than scattered throughout the standard directories.

- **Serving multiple versions of the same product to different clients**

Exporting software to share across a network is simplified and more secure. You need to export only the specific directories under `/opt`, `/usr/opt`, and `/usr/var/opt` that contain the product you want, then create links on the importing system. You can set up a server with multiple versions of a given product, using the links created on the client systems to determine which version a given client uses. In this way, you can maintain software for multiple dissimilar hardware platforms on the same server.

Specific directory requirements exist for each type of product kit. In some cases, additional files are required for the kit to build successfully.

- You do not need any additional installation files for a user product.
- Section 5.2 describes the additional installation files you need for a kernel product. Figure 5-1 illustrates these files.

Install product files in product-specific subdirectories of the root (`/`), `/usr`, and `/usr/var` directories, as described in the following list:

- **Boot files reside under `/opt`**

Install files that are required at bootstrap time, such as device drivers, in a product-specific subdirectory of the `/opt` directory, such as `/opt/OAT100`. This also includes any files to be accessed before file systems other than root are mounted.

- **Read-only files reside under `/usr/opt`**

Install read-only files (such as commands), startup files (not modified by individual users), or data files in a product-specific subdirectory of the `/usr/opt` directory, such as `/usr/opt/OAT100`.

- **Read/write files reside under `/usr/var/opt`**

Install files that users can read and modify, such as data lists, in a product-specific subdirectory of the `/usr/var/opt` directory.

2.3.2 Using Context-Dependent Symbolic Links

If you are preparing a product kit that may run on a cluster, you need to create context-dependent symbolic link (CDSL) to a member-specific file instead of using a shared file in your product kit's inventory.

- A member-specific file is used by a specific cluster member. The contents of a member-specific file differ for each cluster member, and each member has its own copy of a member-specific file.

For example, the `sysconfigtab` file contains information that can be different for each cluster member, so the `/etc/sysconfigtab` file is a CDSL that points to the member's `sysconfigtab` file.

- A shared file is used by all members of a cluster. There is only one copy of a shared file.

For example, executable files may be shared by all cluster members.

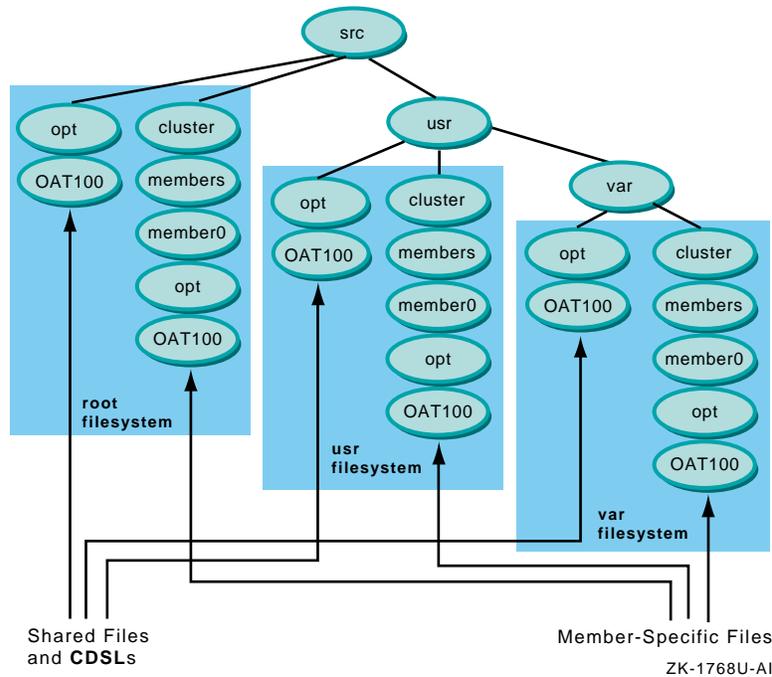
- A context-dependent symbolic link (CDSL) is a special kind of symbolic link that points to a member-specific file. The CDSL references the member-specific file for the member that accesses the CDSL — the member that accesses the CDSL determines its target.

This section provides the following information:

- When to use CDSLs in your product kit (Section 2.3.2.1)
- How to identify CDSLs (Section 2.3.2.2)
- How to create CDSLs for your product kit (Section 2.3.2.3)
- Restrictions on using CDSLs in your product kit (Section 2.3.2.4)

Figure 2–2 shows member-specific directories in the ODB product kit source hierarchy.

Figure 2–2: Member-Specific Directories in Source Hierarchy



Caution

Layered product kits always should refer to CDSLs and not to the corresponding member-specific files. Doing this isolates CDSL awareness from the layered product and keeps it in the file system hierarchy, making it easier for you to maintain and upgrade your layered product kit. CDSL references access the correct file whether the product is installed on a cluster or on a single system, and continue to work if file types change in a future version of the product.

For example, refer to the `/usr/var/X11/Xserver.conf` file rather than to the `/usr/var/cluster/members/member0/X11/Xserver.conf` file. These are the same if the `/usr/var/X11/Xserver.conf` file is a CDSL.

2.3.2.1 Knowing When to Use CDSLs

If your product kit might run on a cluster, you may need to create CDSLs.

- Use CDSLs if you need a file to be different on every machine running the product kit software.

- Use shared files if the file is always the same on every machine running the product kit software.

Do not make a directory into a CDSL. Once a directory is a CDSL you cannot change it back to a regular directory. If a directory contains all member-specific files, make each file a CDSL.

2.3.2.2 Identifying CDSLs

You can identify a CDSL by the presence of the `{memb}` variable in its pathname. For example:

```
lrwxrwxrwx 1 root system 57 May 19 10:54 /etc/sysconfigtab -> \
  ../cluster/members/{memb}/boot_partition/etc/sysconfigtab
```

Note

The backslash (`\`) character in this example indicates line continuation, and is not present in the actual output.

To resolve a CDSL's pathname, the kernel replaces the `{memb}` variable with the string `memberN`, where `N` is the member ID of the cluster member that is referencing the file. If a cluster member with member ID 2 is accessing the `/etc/sysconfigtab` file in this example, the pathname is resolved to `/cluster/members/member2/boot_partition/etc/sysconfigtab`.

2.3.2.3 Creating CDSLs

Follow these steps to create CDSLs:

1. Determine the file system where the CDSLs should reside: `root (/)`, `/usr`, or `/usr/var`.
2. Log in as `root` or use the `su` command to gain superuser privileges.
3. Use the `mkdir` command to create member-specific directories for each file system.

For example, if you are building the OAT100 kit in the `/mykit` directory, enter the following commands for the `root (/)`, `/usr`, and `/usr/var` file systems:

```
# mkdir -p /mykit/src/cluster/members/member0/opt/OAT100
# mkdir -p /mykit/src/usr/cluster/members/member0/opt/OAT100
# mkdir -p /mykit/src/usr/var/cluster/members/member0/opt/OAT100
```

4. Copy the `odb_log` file to the member-specific area, creating directories as needed. Refer to Figure 2-2 and to the sample OAT100 product kit source directory shown in Figure 2-3.

5. If you are working on a single system, use the `ln -s` command to create the CDSLs in the shared file directories. Specify relative pathnames, as shown in the following example:

```
# cd /mykit/src
# mkdir -p usr/var/opt/OAT100/log_files
# cd usr/var/opt/OAT100/log_files
# ln -s \
../../../../cluster/members/{memb}/opt/OAT100/log_files/odb_log
```

Note

In the previous step, you created member specific files in the `cluster/members/member0` path, but in this step you are creating symbolic links to the `cluster/members/{memb}` paths. When your product kit is installed on a cluster, this lets the system resolve the `{memb}` variable to `memberN`, where `N` is the member ID of the cluster member that is referencing the file.

The backslash (`\`) character in the sample user input shows line continuation. You do not need to use a backslash to continue your input line.

6. Use the `ls -l` command to verify the symbolic link, as shown in the following example:

```
# ls -l odb_log
lrwxrwxrwx 1 root system 36 May 19 15:46 odb_log@ -> \
../../../../cluster/members/member0/opt/OAT100/log_files/odb_log
```

Note

The backslash (`\`) character shows line continuation and is not present in the actual output.

If you are working on a cluster, you can use the `mkcdsl` command to create CDSLs. Refer to the `mkcdsl(8)` reference page for information about this command.

2.3.2.4 Restrictions

The following restrictions apply when you use CDSLs in product kits:

1. If you are creating a product kit to run on a version of the operating system prior to Version 5.0A, you cannot use CDSLs.

2. If you are creating a product kit on a system that is running a version of the operating system prior to Version 5.0A, you can use CDSLs in your product kit but you must test it on a system running Version 5.0A or higher of the operating system.
3. If you are creating a product kit on a cluster member running Version 5.0A or higher of the operating system, CDSLs cannot be resolved to the member-specific areas in the `cluster/members/member0` directories. The kernel resolves `{memb}` to the cluster ID of the member where you are creating the kit, `memberN`.

2.3.3 Placing Files in the Kit Source Directory

Files in the kit source directory hierarchy should match their intended installation location under the `root (/)`, `/usr`, and `/var` file systems. Plan to install all of your kit files in product-specific subdirectories in the `/opt`, `/usr/opt`, and `/usr/var/opt` directories to prevent name space conflicts with the base operating system and other layered products.

Table 2–1 shows where to place kit files that will be installed in the `root (/)`, `/usr`, and `/var` file systems:

Table 2–1: File Locations in Kit Directories

Location in Kit src Directory	Installed File System
<code>opt/PRODCODE</code>	<code>root (/)</code>
<code>usr/opt/PRODCODE</code>	<code>/usr</code>
<code>usr/var/opt/PRODCODE</code>	<code>/usr/var</code>

If you overwrite base operating system files, you may encounter the following problems:

- Your product can be corrupted during an Update Installation of the operating system. The Update Installation will overwrite any file that was shipped as part of the old version of the operating system with the version of the file shipped with the operating system.
- An Update Installation may not complete successfully if you overwrite a base operating system file. This can make the system unusable.
- Your product may have to be removed from the system to complete an Update Installation. Your product would have to be reinstalled after the Update Installation is completed.
- Removing your product corrupts the operating system.

2.3.4 Setting Up the Sample Product Kit Source Directory

Figure 2–3 shows how the Orpheus Document Builder (ODB) product is installed in the standard directory structure, under `/opt`, `/usr/opt`, and `/usr/var/opt`. The directories shown between the `src` and the `OAT*` directories are the existing directories on the customer's system. All directories and files created by the product are shipped under the `OAT*` directories. In this example, directory names begin with `OAT` because `OAT` is the three-letter product code assigned to *Orpheus Authoring Tools, Inc.*

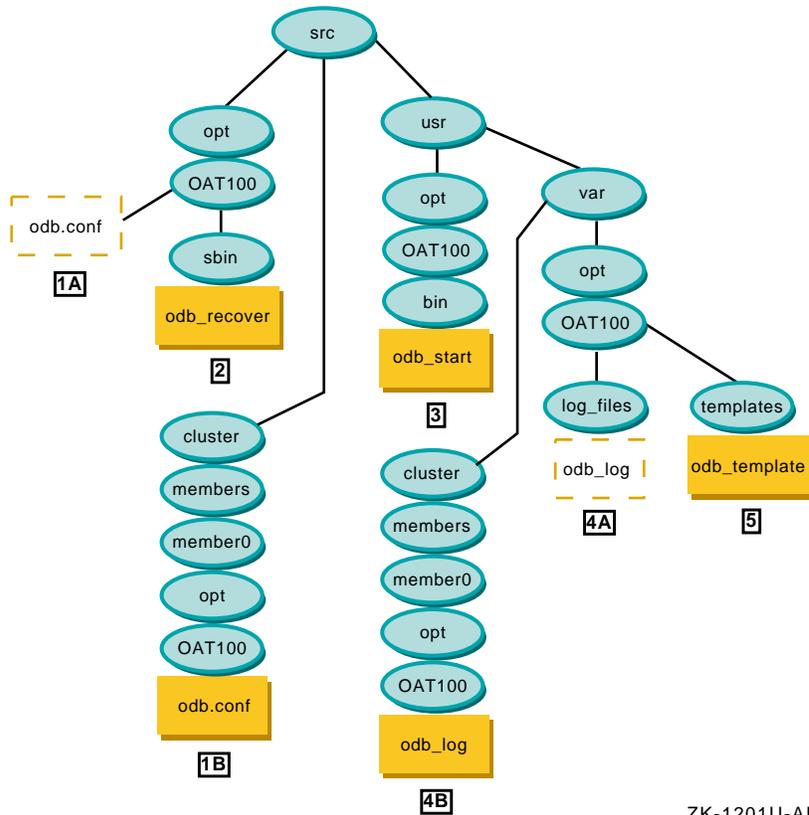
Caution

File attributes (ownership and permissions) for files and directories in the kit's source hierarchy must be set exactly as they should be on the customer's system. This means that you must be superuser when populating the source hierarchy so that you can change these file attributes.

Do not attempt to circumvent this requirement by setting file attributes in your subset control programs. If a superuser on the customer's system runs the `fverify` command on your subsets, attributes that have been modified by the subset control programs are reset to their original values in the kit's master inventory files.

Figure 2–3 shows a sample source directory for the `OAT` product.

Figure 2–3: Sample Product Kit Source Directory



ZK-1201U-AI

The following files have been provided by the ODB developers for the OAT kit. Each of these files has a path and an associated description, and must be placed correctly in the source hierarchy for the OAT product to build successfully.

- 1 /odb.conf — the ODB product configuration file

If the product kit can be installed on a cluster, each cluster member must have its own copy of the configuration file. To accommodate this requirement, create a context-dependent symbolic link (CDSL) targeted to the member-specific file.

- A. The context-dependent symbolic link (CDSL) for the odb.conf file is linked to the member-specific cluster/members/{memb}/opt/OAT100/odb.conf file. This CDSL is installed in the root file system and placed in the opt/OAT100 source directory.

- B. The member-specific file `odb.conf` can differ on each cluster member. This file is installed in the cluster member's root file system and is placed in the `cluster/members/member0/opt/OAT100` source directory.

Refer to Section 2.3.2 for information about CDSLs.

- 2 `/sbin/odb_recover` — a utility to recover corrupt ODB documents when the system boots. The `odb_recover` script executes when the system boots and the `/usr` file system may not be mounted.

This file is installed in the root file system and is placed in the `opt/OAT100/sbin` source directory.

- 3 `/usr/bin/odb_start` — the ODB product startup script. The `odb_start` script is a user command.

This file is installed in the `/usr` file system and is placed in the `usr/opt/OAT100/bin` source directory.

- 4 `/usr/var/log_files/odb_log` — the ODB product log file

If the product kit can be installed on a cluster, each cluster member must have its own copy of the log file. To accommodate this requirement, create a context-dependent symbolic link (CDSL) targeted to a member-specific file.

- A. The context-dependent symbolic link (CDSL) for the `odb_log` file is linked to the member-specific `usr/var/cluster/members/{memb}/opt/OAT100/log_files/odb.log` file. This CDSL is installed in the `/usr/var` file system and placed in the `usr/var/opt/OAT100/log_files` source directory.
- B. The member-specific file `odb_log` can differ on each cluster member. This file is installed in the cluster member's `/usr/var` file system and is placed in the `/usr/var/cluster/members/member0/opt/OAT100/log_files` source directory.

Refer to Section 2.3.2 for information about CDSLs.

- 5 `/usr/var/templates/odb_template` — a document template that can be modified by a user.

This file is installed in the `/var` file system and is placed in the `usr/var/opt/OAT100/templates` source directory.

Each of these files will be installed in the path provided by the kit developer. They reside in the same relative location in the kit source directory with `opt/OAT100` inserted into the pathname.

For users to make effective use of your product after it is installed, they should add the directories that contain your product commands to the search path in their `.profile` or `.login` files. For example, the Orpheus Document

Builder (ODB) product is installed in the standard directory structure under `/opt`, `/usr/opt` and `/usr/var/opt`. The `src` directory mirrors the root (`/`) directory on the customer's system. The commands for the product are located in the `/opt/OAT100/sbin` and `/usr/opt/OAT100/bin` directories. To use ODB commands without specifying the full path on the command line, the user can add these directory names to the `PATH` environment variable.

You can ship a symbolic link to make commands accessible through the standard directories. For example, the ODB kit contains the command `/usr/opt/OAT100/bin/odb_start`. A symbolic link can be created from `/usr/bin/odb_start` to `/usr/opt/OAT100/bin/odb_start`. This also makes the `odb_start` command available to users as a part of their normal search path, since `/usr/bin` is part of the standard path.

You can ship a symbolic link only if one of the following conditions apply:

- The symbolic link does not conflict with any base operating system file.

Using our example, it means that you could create the `/usr/bin/odb_start` link only if the operating system does not already contain a `/usr/bin/odb_start` file. If the operating system did contain a `/usr/bin/odb_start` file, shipping the symbolic link would overwrite an existing file on the customer's operating system.

- The command name does not conflict with any standard operating system command.

For example, if the `/usr/opt/OAT100/bin/odb_start` command is shipped in the ODB kit and was part of the standard operating system in `/bin/odb_start`, when a user entered the `odb_start` command there would be a command name conflict. Depending upon whether `/bin` or `/usr/bin` is first in the search path, the user could be accessing the operating system version or the symbolically linked ODB product version.

3

Creating Subsets

In a product kit, a subset is the smallest installable entity compatible with the `setld` utility. The kit developer specifies how many subsets are included in the kit and what files each subset contains.

Note

A file's physical location is not necessarily a factor in determining the subset to which it belongs.

A kit developer must perform the following tasks to build subsets and associated control files:

1. Organize product files into subsets. (Section 3.1)
2. Create a master inventory file containing information about each file in the subset. (Section 3.2)
3. Create a key file to define product attributes such as the product name, product version, and subset definitions. (Section 3.3)
4. Optionally, create subset control programs (SCPs). (Section 3.4)
5. Use the `kits` utility to produce the subsets and related control files. (Section 3.5)
6. Test your subsets to ensure that they can be loaded onto a running system, that the product runs on the system, and that the subsets can be deleted. (Section 3.6)
7. Optionally, update the master inventory file after you have created subsets. (Section 3.7)

3.1 Grouping Files into Subsets

Files that are required for the product to work should be grouped together by function. For example, if the product has two parts such as a user interface and underlying functional code, you should group them into two subsets.

Optional files also should be grouped together by function, but should be grouped separately from mandatory files. This prevents unnecessary files from being loaded when you install the mandatory subsets.

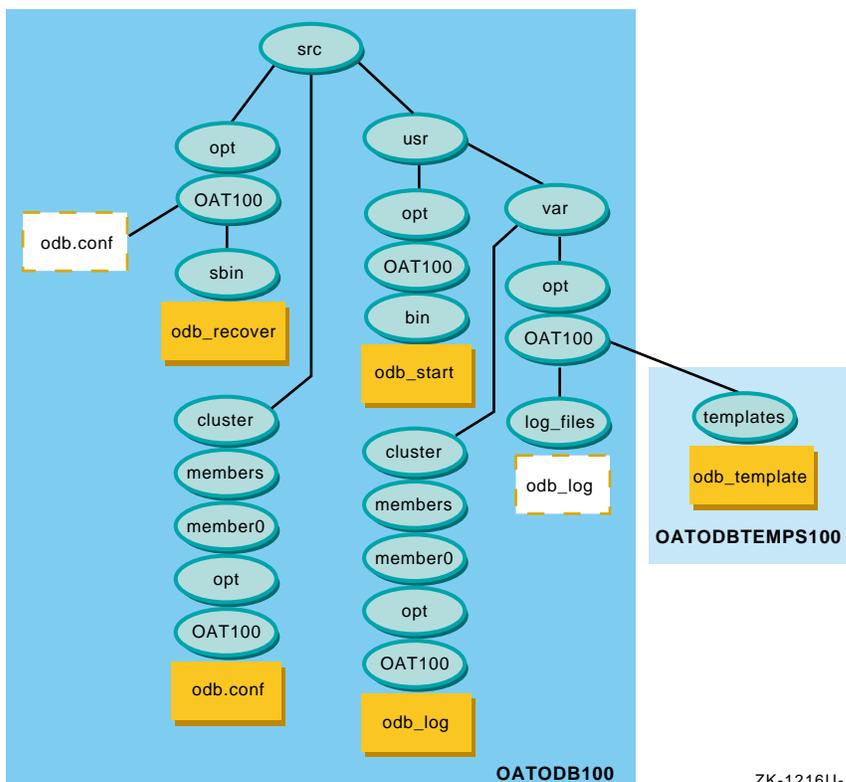
The fictitious ODB user product requires two subsets:

- OATODB100, a mandatory subset, contains the files needed to run the product. This includes all product files except the `odb_template` file.
- OATODBTEMPS100, an optional subset, contains documentation templates in the `odb_template` file.

By placing the documentation templates in a separate subset, the customer's system administrator can choose not to install them if storage space is limited.

Figure 3–1 shows how the files that make up the ODB product are grouped into subsets. The physical location of a file is not necessarily a factor in determining the subset to which it belongs.

Figure 3–1: ODB Product Subsets and Files



ZK-1216U-AI

3.2 Creating the Master Inventory File

After choosing subset names and deciding their contents, you have to specify in a master inventory file the subset names and the files that each subset contains.

You can create a master inventory file with any text editor you like, or create the file with the `touch` command. The master inventory file name must consist of the product code and version with the suffix `.mi`. The file should be located in the `data` directory of the kit. For example:

```
% cd /mykit/data
% touch OAT100.mi
```

The first time you process a kit, the master inventory file is empty. You must enter one record for each file that belongs on the kit. To get an initial list of these files, use the `newinv` command with the file name of the empty master inventory file and the pathname of the source hierarchy's top-level directory. For example, to invoke `newinv` on the master inventory file for the ODB product, specify the pathname to the source hierarchy as a relative path from the current directory (`data`), similar to the following:

```
% newinv OAT100.mi ../src
Scanning new baselevel files...done.

Sorting inventories...done.

Joining...done.

Awking...done.

*** THIS BUILD CONTAINS FILES THAT ARE NOT IN THE PREVIOUS BUILD ***

You will be placed in the editor with the file containing
the names of these new files.

If you wish these new files to become part of the product,
you must convert the line for the wanted files into
an inventory record.

Any records remaining in the file when you exit the editor
will become part of the new inventory.

Type <RETURN> when you are ready or CTRL/C to quit:
```

The `newinv` utility produces a list of files that are present in the source hierarchy and opens a working copy of the master inventory file in the `vi` editor (or the editor specified by your `$EDITOR` environment variable) to make the required changes.

Caution

Be extremely careful when you edit the master inventory file. Separate fields in this file with single Tab characters, not spaces. File names must not contain Space or Tab characters.

Use dot-relative pathnames for the files listed in the master inventory file; do not use absolute pathnames. By default, the `setld` utility operates from the system's root (`/`) directory unless you specify an alternate root with the `-D` option.

First, remove the entries for any files that should not appear on the kit. Second, add the flags and subset identifier to the entry for each file that should appear in the kit.

Caution

The first time that you run the `newinv` utility, the following files are created in the `/mykit/data` directory in addition to the `OAT100.mi` file:

```
OAT100.mi.bkp
OAT100.mi.dead
OAT100.mi.extra
OAT100.mi.join
OAT100.mi.tmp
```

Do not modify or delete these additional files. They are used during subsequent master inventory file updates with the `newinv` utility.

Table 3-1 describes the fields in the master inventory file.

Table 3-1: Master Inventory File

Field	Description
Flags	16-bit unsigned integer Bit 1 is the <code>v</code> (volatility) bit. When this bit is set, changes to an existing copy of the file can occur either during or after kit installation. The remaining bits are reserved, so valid values for this field are 0 (protected) or 2 (unprotected). The volatility bit usually is set for log files such as <code>/usr/spool/mqueue/syslog</code> .
Pathname	The dot-relative (<code>./path</code>) path to the file.
Subset identifier	The name of the subset that contains the file. Subset names consist of the product code, subset mnemonic, and version number. You must not include standard system directories in your subsets. In the ODB master inventory file, several records specify directories that are part of the standard system hierarchy. Instead of a subset identifier, these records specify <code>RESERVED</code> ; this keyword prevents <code>setld</code> from overwriting existing directories.

After you edit the file list and exit the editor, you see output similar to the following:

Merging...done.

Sorting...done.

Master inventory update complete.

Example 3–1 shows that the ODB kit has two subsets:

- The OATODB100 subset contains mandatory commands and utilities.
- The OATODBTMP100 subset contains optional document templates.

Example 3–1: Sample ODB Kit Master Inventory File

```
0 ./cluster RESERVED 1
0 ./cluster/members RESERVED 1
0 ./cluster/members/member0 RESERVED 1
0 ./cluster/members/member0/opt RESERVED 1
0 ./cluster/members/member0/opt/OAT100 OATODB100 2
2 ./cluster/members/member0/opt/OAT100/odb.conf OATODB100 23
0 ./opt RESERVED 1
0 ./opt/OAT100 OATODB100 2
2 ./opt/OAT100/odb.conf OATODB100 23
0 ./opt/OAT100/sbin OATODB100 2
0 ./opt/OAT100/sbin/odb_recover OATODB100 2
0 ./usr RESERVED 1
0 ./usr/opt RESERVED 1
0 ./usr/opt/OAT100 OATODB100 2
0 ./usr/opt/OAT100/bin OATODB100 2
0 ./usr/opt/OAT100/bin/odb_start OATODB100 2
0 ./usr/var RESERVED 1
0 ./usr/var/cluster RESERVED 1
0 ./usr/var/cluster/members RESERVED 1
0 ./usr/var/cluster/members/member0 RESERVED 1
0 ./usr/var/cluster/members/member0/opt RESERVED 1
0 ./usr/var/cluster/members/member0/opt/OAT100 OATODB100 2
0 ./usr/var/cluster/members/member0/opt/OAT100/odb_log OATODB100 24
0 ./usr/var/opt RESERVED 1
0 ./usr/var/opt/OAT100 OATODB100 2
0 ./usr/var/opt/OAT100/log_files OATODB100 2
0 ./usr/var/opt/OAT100/log_files/odb_log OATODB100 24
0 ./usr/var/opt/OAT100/templates OATODBTMP100 5
0 ./usr/var/opt/OAT100/templates/odb_template OATODBTMP100 5
```

1 Add the RESERVED identifier to the directories shown. This tells the setld utility not to overwrite the directory if it already exists on the customer's system.

2 Add the OATODB100 subset identifier to the files and directories shown. This is the mandatory subset for the ODB product.

- ❸ The `odb.conf` file is the ODB product configuration file. The member-specific file resides in the `./cluster/members/member0/opt/OAT100` directory, and the context-dependent symbolic link (CDSL) resides in the `./opt/OAT100` directory. Change the `Flags` field to 2 to show that this file can change either during or after subset installation.
- ❹ The `odb_log` file is the ODB product log file. The member-specific file resides in the `./usr/var/cluster/members/member0/opt/OAT100/log_files` directory, and the CDSL resides in the `./usr/var/opt/OAT100/log_files` directory.
- ❺ Add the `OATODBTMP100` subset identifier to the files and directories shown.

3.3 Creating the Key File

The key file defines product attributes such as the product name, product version, and subset definitions, as well as the name of the kit's master inventory file. It consists of a product attributes section and a subset descriptor section. The key file name must consist of the product code and version followed by `.k`, so that `OAT100.k` is the key file for the ODB kit. Create the key file in the `data` directory of your kit-building directory structure, for example: `/mykit/data`.

Example 3–2 shows the ODB product kit key file with the two sections separated by two percent signs (`%%`) on their own line:

Example 3–2: Sample ODB Kit Key File

```

#
# Product-level attributes ❶
#
NAME='Orpheus Document Builder'
CODE=OAT
VERS=100
MI=/mykit/data/OAT100.mi ❷
COMPRESS=1 ❸
#
# Subset definitions ❹
#
%% ❺
OATODB100 . 0 'Document Builder Tools' ❻
OATODBTMP100 OATODB100|OSFDCMT??? 2 'Document Builder Templates' ❼

```

- ❶ The product attributes portion of the file describes the naming conventions for the kit and provides kit-level instructions for the `kits` command. This section of the key file consists of several lines of attribute-value pairs as described in Table 3–2. The order of these attribute-value pairs is not significant. Each attribute name is separated from its value by an equal sign, for example:

attribute=value. You can include comment lines, which begin with a pound sign, for example: *# Comment line in key file*.

- 2 The value of the `MI` attribute contains the path to the master inventory file. This may be either an absolute path or a relative path from the directory where the `kits` command is executed.
- 3 The `COMPRESS` attribute has a value of 0 for uncompressed subsets or 1 for compressed `tar` format subsets. User and kernel product kit subsets may be compressed or uncompressed.
- 4 The subset descriptor portion of the file describes each of the subsets in the kit and provides subset-level instructions for the `kits` command. This section contains one line for each subset in the kit. Each line consists of four fields, each separated by a single Tab character. You cannot include comments in this section of the key file. Table 3-3 describes the subset descriptor fields.
- 5 Separate the product-level attributes and the subset definitions with two percent signs (`%%`) on their own line. If this line is not present, the `kits` utility terminates with an error message.
- 6 In this entry, the Dependency list field value for `OATODB100` is `.` (dot), meaning that the subset has no dependencies.
Set the Flags field to 0 (zero), indicating that the subset is mandatory.
- 7 In this entry, the `OATODBTEMP100` subset is optional; its `FLAGS` field is set to 2 (two). This subset is dependent on both the `OATODB100` subset, part of the `ODB` kit, and the `OSFDCMT???` subset, part of the base operating system. The `???` notation is a wildcard to specify any version of the `OSFDCMT` subset.

Enclose the Subset Description field in single quotes, for example:
'Subset Description'.

The key file product attributes section describes the naming conventions for the kit and provides kit-level instructions for the `kits` command. This section consists of attribute-value pairs as described in Table 3-2. Each attribute name is separated from its value by an equal sign (`=`). Comment lines in this section begin with a pound sign (`#`).

Table 3-2: Key File Product Attributes

Attribute	Description
<code>NAME</code>	The product name; for example, <code>'Orpheus Document Builder'</code> . Enclose the product name in single quotation marks (<code>'</code>) if it contains spaces.

Table 3–2: Key File Product Attributes (cont.)

Attribute	Description
CODE	<p>A unique three-character product code, for example, OAT. The first character must be a letter. The first three letters of a subset name must be the same as the product code. In this guide, OAT is the three character code assigned to the fictional <i>Orpheus Authoring Tools, Inc.</i> company.</p> <p>Several product codes are reserved, including (but not limited to) the following: DNP, DNU, EPI, FOR, LSP, ORT, OSF, SNA, UDT, UDW, UDX, ULC, ULT, ULX, and UWS.</p> <p>Send electronic mail to <code>product@dssr.sgp.zko.dec.com</code> to request a product code.</p>
VERS	<p>A three-digit version code; for example, 100. The <code>setld</code> utility interprets this version code as 1.0.0. The first digit should reflect the product's major release number, the second the minor release number, and the third the upgrade level, if any. The version number cannot be lower than 100. The version number is assigned by the kit developer.</p>
MI	<p>The name of the master inventory file. If the master inventory file is not in the same directory where the <code>kits</code> utility is run, you must specify the explicit relative path from the directory where you are running the <code>kits</code> utility to the directory where the master inventory file resides. The file name of the product's master inventory file consists of the product code and version plus the <code>.mi</code> extension. You create and maintain the master inventory file with the <code>newinv</code> utility.</p>
ROOT	<p>Not illustrated in the example, the operating system has reserved this optional attribute for the base operating system. <code>ROOT</code> has a string value that names the root image file. Do not use this attribute for a layered product.</p>
COMPRESS	<p>An optional flag that is set to 1 if you want to create compressed subset files. For kits in Direct CD-ROM (DCD) format, you must set this flag to 0 (zero). Compressed files require less space on the distribution media (sometimes as little as 40% of the space required by uncompressed files), but they take longer to install than uncompressed files. If missing, this flag defaults to 0 (zero).</p>

The key file subset descriptor section describes each of the subsets in the kit and provides subset-level instructions for the `kits` command. This section contains one line for each subset in the kit and consists of four fields, each separated by a single Tab character. You cannot include comments in this section of the key file. Table 3–3 describes the subset descriptor fields.

Table 3–3: Key File Subset Descriptors

Field	Description
Subset identifier	A character string up to 80 characters in length, composed of the product code (for example, OAT), a mnemonic identifying the subset (for example, ODB), and the three-digit version code (for example, 100). In this example, the subset identifier is OATODB100. All letters in the subset identifier must be uppercase.
Dependency list	Either a list of subsets upon which this subset is dependent (OATODB100 OSFDCMT510), or a single period (.) indicating that there are no subset dependencies. Separate multiple subset dependencies with a pipe character ().
Flags	<p>A 16-bit unsigned integer; the operating system defines the use of the lower 8 bits.</p> <p>Set bit 0 to indicate whether the subset can be removed (0=removable, 1=protected).</p> <p>Set bit 1 to indicate whether the subset is optional (0=mandatory, 1=optional).</p> <p>Set bit 2 to indicate compression (0=compressed, 1=uncompressed).</p> <p>Bits 3-7 are reserved for future use. You can use bits 8-15 to relay special subset-related information to your subset control program.</p>
Subset description	A short description of the subset, delimited by single quotation marks ('), for example: 'Document Builder Tools'. The percent sign (%) is reserved in this field; do not use it for layered products.

3.4 Creating Subset Control Programs

This section describes common tasks required to write subset control programs for product kits. The following topics are discussed:

- Creating SCP source files (Section 3.4.1)
- Setting up initial SCP processing (Section 3.4.2)
- Working in a cluster environment (Section 3.4.3)
- Working in a DMS environment (Section 3.4.4)
- Associating SCP tasks with `setld` processing phases (Section 3.4.5)
- Stopping the installation (Section 3.4.6)
- Creating SCPs for different types of product kits (Section 3.4.7)

A subset control program (SCP) performs special tasks beyond the basic installation tasks managed by the `setld` utility. The following list includes some of the reasons why you might need to write a subset control program:

- Some of your kit's files have to be customized before the product will work properly.
- You want to offer the user the option to install some of the files in a nonstandard location.
- You want to register and statically or dynamically configure a device driver.
- Your kit depends on the presence of other products.
- You need to establish nonstandard permissions or ownership for certain files.
- Your kit requires changes in system files such as `/etc/passwd`.

A subset control program can perform all of these tasks.

Caution

Code your subset control program so that it can run more than once without causing operational problems. This does not mean that you must repeat the SCP tasks, but that multiple executions will not cause the SCP to fail or to corrupt existing files.

Remove any code from existing SCPs that refer to `subset-id.lk` or `subset-id.dw` files.

- Use the `DEPS` field in the subset control file (`subset_id.ctrl`) for subset dependency processing.
DO NOT use the `subset-id.lk` files to check for subset dependencies.
 - Use the library routines described in Table 3-6 to determine if a subset is installed.
 - Effective with this version of the operating system, `subset-id.dw` files are no longer created. Use the library routines described in Table 3-6 to determine if a subset is corrupt.
-

3.4.1 Creating SCP Source Files

Create one subset control program for each subset that requires special handling during installation. You can write the program in any programming language, but your subset control program must be executable on all

platforms on which the kit can be installed. If your product works on more than one hardware platform, you cannot write your subset control program in a compiled language. For this reason, it is recommended that you write your subset control program as a script for `/sbin/sh`. All of the examples in this chapter are written in this way.

Keep your subset control programs short. If written as a shell script, a subset control program should be under 100 lines in length. If your subset control program is lengthy, it is likely that you are trying to make up for a deficiency in the architecture or configuration of the product itself.

Caution

Subset control programs should not require any interactive responses, and should not generate errors when run repeatedly.

Place all subset control programs that you write in the `scps` directory, a subdirectory of the `data` directory. Each subset control program's file name must match the subset name to which it belongs, and it must end with the `scp` suffix. For example, the ODB product defines two subsets, named `OATODB100` and `OATODBTMP100`. If each of these subsets required a subset control program, the source file names would be `OATODB100.scp` and `OATODBTMP100.scp`.

When you create the subsets as described in Section 3.5, the `kits` utility copies the subset control programs from the `./data/scps` directory to the `./output/instctrl` directory. If a subset has no SCP, the `kits` utility creates an empty subset control program file for it in the `./output/instctrl` directory.

3.4.2 Setting Up Initial SCP Processing

Your subset control program should perform the following tasks within the program:

- Including library routines (Section 3.4.2.1)
- Setting global variables (Section 3.4.2.2)

The following sections describe the resources available to perform these tasks.

3.4.2.1 Including Library Routines

The operating system provides a set of routines in the form of Bourne shell script code located in the `/usr/share/lib/shell` directory. Do not copy these routines into your subset control program. This would prevent your

kit from receiving the benefit of enhancements or bug fixes made in future releases. Use the following syntax to include these library routines:

```
. /usr/share/lib/shell/lib_name
```

In the previous example, *lib_name* is one of the shell scripts specified in one or more of the following tables.

Table 3-4 lists the library routines available in the `libscp` shell script.

Table 3-4: Available SCP Library Routines

Purpose	Library Routine
Architecture checking	STL_ArchAssert
Dependency locking	STL_LockInit ^a STL_DepLock ^a STL_DepUnLock ^a
Dataless environment checking	STL_IsDataless STL_NoDataless
Forward symbolic linking	STL_LinkCreate STL_LinkRemove
Backward symbolic linking	STL_LinkInit STL_LinkBack
SCP initialization	STL_ScpInit

^a Do not use this library routine. It is provided for backward compatibility only. Use the `DEPS` field in the subset control file (*subset_id.ctrl*) for subset dependency processing.

Table 3-5 lists the library installation routines available in the `libinstall` shell script.

Table 3-5: Available Library Installation Routines

Purpose	Library Routine
Cluster member identification	INST_GetMemberID

Table 3-6 lists the library routines available in the `libswdb` shell script.

Table 3-6: Available Installed Software Database Library Routines

Purpose	Library Routine
Installed subsets list	SWDB_FindInstalledSubsets
3-digit version number	SWDB_FindInstalledVersions SWDB_FindLatestVersions

Table 3–6: Available Installed Software Database Library Routines (cont.)

Purpose	Library Routine
Product name in subset control file (<i>subset_id.ctrl</i>)	SWDB_GetProductName
Subset installation status	SWDB_IsCorrupt
	SWDB_IsInstalled
Subset dependencies	SWDB_IsLocked
	SWDB_ListLockingSubsets

3.4.2.2 Setting Global Variables

You can call the `STL_ScpInit` routine to define these variables and initialize them to their values for the current subset. This routine eliminates the need to hard code subset information in your subset control program.

Note

Use the `STL_ScpInit` routine to initialize global variables at the beginning of all `setld` utility processing phases in your SCP except the M phase. The control file is not read before the M phase.

All predefined global variable names begin with an underscore (`_`) for easier identification.

Table 3–7 lists global variables that the subset control program can use to access information about the current subset.

Table 3–7: STL_ScpInit Global Variables

Variable	Description
<code>_SUB</code>	Subset identifier, for example, OATODB100
<code>_DESC</code>	Subset description, for example, Document Builder Tools
<code>_PCODE</code>	Product code, for example, OAT
<code>_VCODE</code>	Version code, for example, 100
<code>_PVCODE</code>	Concatenation of product code and version code, for example, OAT100
<code>_PROD</code>	Product description, for example, Orpheus Document Builder
<code>_ROOT</code>	The root directory of the installation
<code>_SMDB</code>	The location of the subset control files, <code>./usr/.smdb.</code>
<code>_INV</code>	The inventory file, for example, OATODB100.inv

Table 3–7: STL_Scplnit Global Variables (cont.)

Variable	Description
<code>_CTRL</code>	The subset control file, for example, <code>OATODB100.ctrl</code>
<code>_OPT</code>	The directory specifier <code>/opt/</code>
<code>_ORGEXT</code>	File extension for files saved by the <code>STL_LinkCreate</code> routine, set to <code>pre\$_PVCODE</code>
<code>_OOPS</code>	The NULL string, for dependency checking ^a

^a Do not use this global variable. It is provided for backward compatibility only. Use the `DEPS` field in the subset control file (`subset_id.ctrl`) for subset dependency processing.

3.4.3 Working in a Cluster Environment

A cluster is a tightly coupled collection of otherwise independent servers that share storage and other resources, providing high availability of applications and data, increased scalability of services, and ease of service and system management.

A TruCluster Server cluster uses a common root file system and shares a single name space for files and directories. This cluster file system (CFS) provides the same view of directly attached file systems to all cluster members.

For additional information about clusters, refer to the TruCluster Server documentation set.

When you create your subset control programs, consider the following restrictions so that your SCP tasks do not cause operational problems:

- Any `setld` utility phase of your SCP must be able to run more than once without causing operational problems. This does not mean that you must repeat the SCP tasks, but that multiple executions will not cause the SCP to fail or corrupt existing files.
- In a cluster, some `setld` utility phases will run on each cluster member. Any changes made by the SCP must first check to see if the change already has been made to that cluster member. If so, the SCP should not attempt to make the change again.
- The SCP should never change the file type of an existing CDSL.

Do not try to change a CDSL to a file or directory. This would create a shared file where the cluster expects to find a link to a member-specific file and would cause cluster-wide problems.

- The SCP should never change base operating system or cluster inventory.
- The SCP should not decline a delete operation.

If a deconfiguration fails, report the error and continue the deletion, but do not exit with a nonzero status. The user must fix the problem after the software is removed.

Table 3–8 describes how `setld` utility phases behave when your SCP runs on a cluster.

Table 3–8: SCP Operations on a Cluster

setld Phase	Cluster Behavior
<i>all phases</i>	All phases of <code>setld</code> utility processing must be able to run more than once without causing operational problems. This does not mean that you must repeat the SCP tasks, but that multiple executions will not cause the SCP to fail.
M	Runs only on the cluster member where you run the SCP. The SCP invokes the <code>setld</code> utility on that member.
PRE_L	Runs only once for the entire cluster. If you must run member-specific operations, include SCP code that performs the operation on all cluster members in the C INSTALL phase. Do not decline software loading because of one cluster member.
POST_L	Runs only once for the entire cluster. If you must run member-specific operations, include SCP code that performs the operation on all cluster members in the C INSTALL phase. Do not decline software configuration because of one cluster member.
C INSTALL	Runs multiple times, once on each cluster member. If your SCP needs to access member-specific files, perform those operations during the C INSTALL phase.
C DELETE	Runs multiple times, once on each cluster member. Always return a zero exit status from the C DELETE phase. A nonzero status tells the <code>setld</code> utility not to delete the software, but if the <code>setld</code> utility has run the C DELETE phase on other cluster members then the software already may be marked as corrupt. If the operation fails, report the error and continue processing. The user must fix the problem after the software is removed. ^a
PRE_D	Runs only once for the entire cluster. If you must run member-specific operations, include SCP code that performs the operation on all cluster members in the C DELETE phase. Always return a zero exit status from the PRE_D phase. A nonzero status tells the <code>setld</code> utility not to delete the software, but since the <code>setld</code> utility has run the C DELETE phase on other cluster members then the software already may be marked as corrupt. If the operation fails, report the error and continue processing. The user must fix the problem after the software is removed. ^a

Table 3–8: SCP Operations on a Cluster (cont.)

setld Phase	Cluster Behavior
POST_D	<p>Runs only once for the entire cluster. If you must run member-specific operations, include SCP code that performs the operation on all cluster members in the C DELETE phase.</p> <p>Always return a zero exit status from the POST_D phase. If the operation fails, report the error and continue processing. The user must fix the problem after the software is removed.^a</p>
V	No cluster-specific restrictions.

^a Refer to the *Installation Guide* and the `setld(8)` reference page for recovery information.

3.4.4 Working in a Dataless Environment

In a Dataless Management Services (DMS) environment, one computer acts as a server by storing the operating system software on its disk. Other computers, called clients, access this software across the Local Area Network (LAN) rather than from their local disks. Refer to *Sharing Software on a Local Area Network* for more information about DMS.

The `setld` utility uses an alternate root directory in a Dataless Management Services (DMS) environment. To make your subset control program DMS compliant, use dot-relative pathnames for file names and full absolute pathnames starting from `root (/)` for commands in your subset control program. This ensures that the proper command is executed when running on either the server or the client in the dataless environment.

The following example shows the default path for SCP processing commands to be run from the server in a DMS environment:

```
/sbin:/usr/sbin:/usr/sbin:/usr/bin:.
```

A subset control program may need to perform differently in a dataless environment or disallow installation of the subset on such a system. If the product will be installed onto a DMS server, use relative pathnames in your SCP. The dataless environment root is the DMS area rather than the DMS server's root file system.

Caution

When running on a dataless client, the `/usr` area is not writable. You cannot install the product kit if any files reside in the `/usr` directory. Make sure the subset control program does not attempt to write to any files located in the `/usr` directory.

You can use the following routines to perform SCP processing in dataless environments:

`STL_IsDataless`

Checks to see if a subset is being installed into a dataless environment.

`STL_NoDataless`

Declines installation of a subset into a dataless environment.

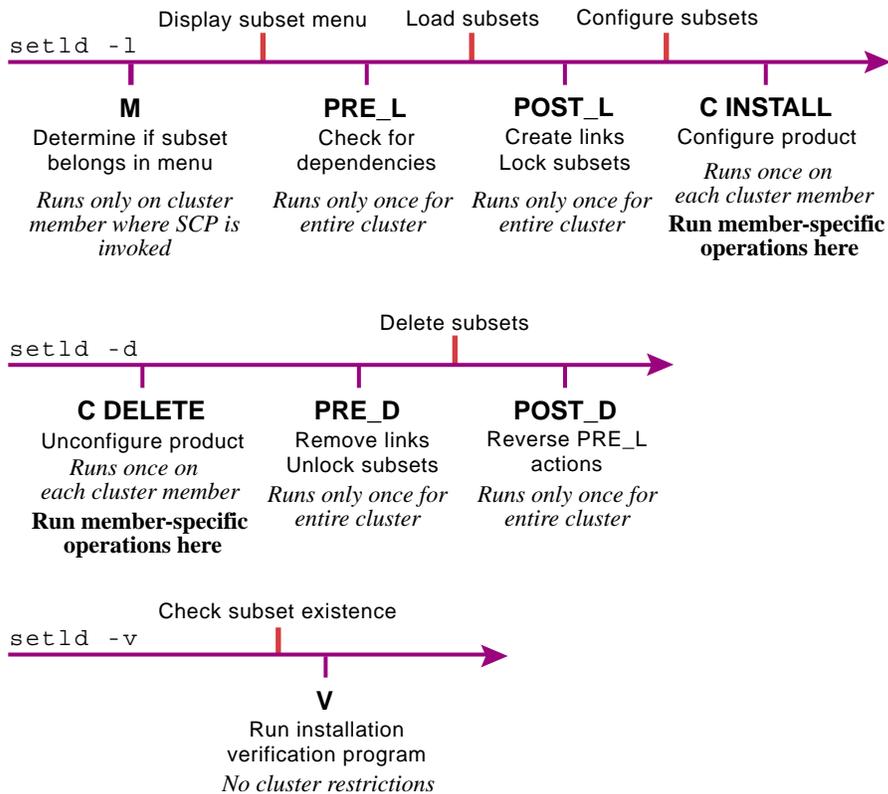
3.4.5 Associating SCP Tasks with `setld` Utility Phases

The `setld` utility invokes the subset control program during different phases of its processing. The SCP can perform certain tasks during any of these phases, such as creating or deleting a file or displaying messages. Other tasks that may be required, such as performing dependency checks or creating links, should be performed only during specific phases.

Some tasks must take place during specific phases. For example, checking dependency relationships between subsets occurs during the `PRE_L` phase; creating links between product files and the standard directory structure occurs during the `POST_L` phase.

Figure 3-2 shows `setld` utility time lines for the `-l`, `-d`, and `-v` options.

Figure 3–2: Time Lines for setld Utility Phases



ZK-1220U-AI

The actions taken by the `setld` utility are shown above the time lines. The SCP actions taken during each `setld` processing phase are shown below the time lines, along with any restrictions when the SCP is run on a cluster.

When the `setld` utility enters a new phase, it first sets the `ACT` environment variable to a corresponding value, then invokes the subset control program. The SCP checks the value of the `ACT` environment variable and any command line arguments to determine the required action.

Caution

Do not include wildcard characters in your subset control program's option-parsing routine. Write code only for the cases the subset control program actually handles. For example, the subset control programs in this chapter provide no code for several conditions under which they could be invoked, for example, the `V` phase.

The following sections describe the tasks that a subset control program may perform in each `setld` processing phase:

- Displaying the subset menu: `M` phase (Section 3.4.5.1)
- Before loading the subset: `PRE_L` phase (Section 3.4.5.2)
- After loading the subset: `POST_L` phase (Section 3.4.5.3)
- After securing the subset: `C INSTALL` phase (Section 3.4.5.4)
- Before deleting a subset: `C DELETE` phase (Section 3.4.5.5)
- Before deleting a subset: `PRE_D` phase (Section 3.4.5.6)
- After deleting a subset: `POST_D` phase (Section 3.4.5.7)
- Verifying the subset: `V` phase (Section 3.4.5.8)

Caution

Any `setld` utility phase of your SCP must be able to run more than once without causing operational problems. This does not mean that you must repeat the SCP tasks, but that multiple executions will not cause the SCP to fail.

Refer to the `setld(8)` and `stl_scp(4)` reference pages for more information about the `setld` utility and conventions for subset control programs.

3.4.5.1 Displaying the Subset Menu (M Phase)

Whenever it performs an operation, the `setld` utility uses the `M` phase to determine if the subset should be included in that operation. Before displaying the menu, `setld` sets the `ACT` environment variable to `M` and calls the subset control program for each subset. At this time, the subset control program can determine whether to include its subset in the menu. The subset control program should return a value of 0 (zero) if the subset can be included in the menu.

Note

In a cluster environment, the `M` phase runs only on the cluster member where the `setld` utility is invoked by the SCP.

Example 3–3 shows a sample `setld` installation menu, listing the subsets available for installation.

Example 3–3: Sample setld Installation Menu

```
1) Kit One Name: Subset Description
2) Kit Two Name: Subset Description
3) Kit Three Name: Subset Description
4) Kit Four Name: Subset Description
5) ALL of the above
6) CANCEL selections and redisplay menus
7) EXIT without installing any subsets
```

Enter your choices or press RETURN to redisplay menus.

Choices (for example, 1 2 4-6):

When it calls the subset control program during this phase, the `setld` utility passes one argument, which can have one of two values:

- The `-l` argument indicates that the operation is a subset load.
- The `-x` argument is reserved for extraction of the subset into a RIS server's product area.

When `setld` extracts a subset into a RIS server's product area, the server also executes the subset control program to make use of the program's code for the `M` phase of installation. You should code the `M` phase to detect the difference between extraction of the subset into a RIS area and loading of the subset for use of its contents. To make this determination, check the value of the `$1` command argument (either `-x` for RIS extraction or `-l` for loading). For RIS extraction, the subset control program should take no action during the `M` phase. When loading subsets, the SCP should perform a machine test. The following Bourne shell example illustrates one way to code the `M` phase. In Example 3–4, the subset control program is checking to determine the type of processor on which it is running. In this example, there is no special code for the RIS extract case.

Example 3–4: Sample Test for Alpha Processor During M Phase

```
#
# The ACT variable is set by setld and determines which
# phase of the SCP should be executed.
#
case $ACT in
#
# This is the menu phase of the SCP
#
M)
#
# Setld invokes the M phase with an argument and if
# the argument is "-l" it means that a software load
# is occurring.
#
case $1 in
-l)
#
```

Example 3–4: Sample Test for Alpha Processor During M Phase (cont.)

```
# Examine the machine architecture to be sure
# that this software is being installed on an
# alpha machine.  If it is not, exit with an
# error status so that setld will not display
# this subset on the menu of subsets to load.
#
ARCH=`./bin/machine`
[ "$ARCH" = alpha ] || exit 1
;;
esac
;;
:
:
```

In Example 3–4, the SCP returns the following codes to the `setld` utility:

- 0 - offer the subset on the menu
- 1 - do not offer the subset on the menu

Note

Installation for a dataless client uses the client's local copy of the `machine` shell script even though the installation is performed in a DMS area on the server. Refer to Section 3.4.4 and *Sharing Software on a Local Area Network* for more information about DMS.

3.4.5.2 Before Loading the Subset (PRE_L Phase)

After presenting the menu and before loading the subset, the `setld` utility sets the `ACT` environment variable to `PRE_L` and calls the subset control program. At this time, the subset control program can take any action required to prepare the system for subset installation, such as protecting existing files.

Note

In a cluster environment, the `PRE_L` phase is run only once for the whole cluster. If you must run member-specific operations when your subset is loaded, include code in your SCP that performs the operation on all cluster members in the `C INSTALL` phase.

Do not decline software loading because of one cluster member.

Use the `DEPS` field in the subset control file (`subset_id.ctrl`) for subset dependency processing.

If you overwrite base operating system files, you may encounter the following problems:

- Your product can be corrupted during an Update Installation of the operating system. The Update Installation will overwrite any file that was shipped as part of the old version of the operating system with the version of the file currently shipped with the operating system.
- An Update Installation may not complete successfully if you overwrite a base operating system file. This can make the system unusable.
- Your product may have to be removed from the system to complete an Update Installation. Your product would have to be reinstalled after the Update Installation is completed.
- Removing your product corrupts the operating system.

If your subset control program is designed to overwrite existing files, it first should make a backup copy of the original file during the `PRE_L` phase and restore the copy in the `POST_D` phase described in Section 3.4.5.7.

In Example 3–5, the subset control program checks a list of files to be backed up if they already exist on the system. If it finds any, it creates a backup copy with an extension of `.OLD`.

Example 3–5: Sample Backup of Existing Files During `PRE_L` Phase

```

    :
#
# Here is a list of files to back up if found on
# the installed system.
#
BACKUP_FILES="\
./usr/var/opt/$_PVCODE/templates/odb_template \ 1
./usr/var/opt/$_PVCODE/log_files/odb_log" 1
    :
#
# The ACT variable is set by setld and determines which
# phase of the SCP should be executed.
#
case $ACT in
    :
#
# This is the pre-load phase of the SCP
#
PRE_L)
#
# Loop through the list of backup files and create
# backup copies for any file that is found on the
# system.
#
for FILE in $BACKUP_FILES
do
#
# If the file to be backed up exists, create
# a backup copy with a .OLD extension.
```

Example 3–5: Sample Backup of Existing Files During PRE_L Phase (cont.)

```
#
if [ -f $FILE -a ! -f $FILE.OLD ]
then
  cp $FILE $FILE.OLD [2]
fi
done
;;
:
```

- [1] The `STL_ScpInit` routine sets the value of the `$_PVCODE` global variable to `OAT100`. Using this variable allows the SCP to be used for the next version of the product, `OAT200`, without changing the pathnames.
- [2] A backup copy is made if the specified file exists and if a backup copy does not already exist. This will be restored in the `POST_D` phase.

In Example 3–5, the SCP returns the following codes to the `setld` utility:

```
0 - load the subset
1 - do not load the subset
```

3.4.5.3 After Loading the Subset (POST_L Phase)

After loading the subset, the `setld` utility sets the `ACT` environment variable to `POST_L` and calls the subset control program for each subset. At this time the subset control program can make any modifications required to subset files that usually are protected from modification when the installation is complete. The subset control program should create backward links and perform dependency locking at this time

Note

In a cluster environment, the `POST_L` phase is run only once for the whole cluster. If you must run member-specific operations when your subset is loaded, include code in your SCP that performs the operation on all cluster members in the `C INSTALL` phase.

Do not decline software configuration because of one cluster member.

Sometimes you may need to create links within your product-specific directories that refer to files in the standard hierarchy. Such backward links must be created carefully because the layered product directories themselves can be symbolic links. This means that you cannot rely on

knowing in advance the correct number of directory levels (. . /) to include when you create your backward links. For example, `/var` is frequently a link to `/usr/var`.

When a kit is installed on a Network File System (NFS) server, the SCP should make the backward links in the server's kit area. When the server's kit area is exported to clients, the links are already in place and you do not need to create any backward links in the client area. This is done so that installation on an NFS client cannot overwrite any existing backward links in the server's kit areas. You do not run the subset control program on an NFS client. Your subset control program should create and remove backward links in the `POST_L` and `PRE_D` phases, respectively.

Caution

NFS clients importing products with backward links must have directory hierarchies that exactly match those on the server. Otherwise, the backward links fail.

Use the `STL_LinkInit` and `STL_LinkBack` routines to create backward links as follows, and use the `rm` command to remove them:

`STL_LinkInit`

Used in the `POST_L` phase to establish internal variables for the `STL_LinkBack` routine. Before you use `STL_LinkBack` to create a link, you must execute `STL_LinkInit` once. This routine has no arguments and returns no status.

`STL_LinkBack link_file file_path link_path`

Creates a valid symbolic link from your product area (under `/usr/opt` or `/usr/var/opt`) to a directory within the standard UNIX directory structure. In this example, `link_file` is the file to link, `file_path` is the dot-relative path of the directory where the file actually resides, and `link_path` is the dot-relative path of the directory where you should place the link. You can use `STL_LinkBack` repeatedly to create as many links as required. This routine returns no status.

Example 3-6 uses `STL_LinkBack` in the `POST_L` phase to create a link named `/opt/OAT100/sbin/ls` that refers to the real file `/sbin/ls`, and removes the link in the `PRE_D` phase.

Example 3–6: Sample Backward Link Creation During POST_L Phase

```
case $ACT in
    :
    :
#
# This is the post-load phase of the SCP
#
POST_L)
#
# Initializes the variables so that the STL_LinkBack
# routine can be executed
#
STL_LinkInit
#
# Create a symbolic link in the ./opt/$_PVCODE/sbin
# directory that points to the ./sbin/ls file.
#
STL_LinkBack ls ./sbin ./opt/$_PVCODE/sbin 1
;;
PRE_D)
#
# Remove the links created in the POST_L phase
#
rm -f ./opt/$_PVCODE/sbin/ls 2
;;
    :
    :
```

1 The `STL_LinkBack` routine creates a backward link in the product-specific area. If you used the `STL_LinkCreate` routine, it would create an unacceptable link in the OSF name space. The `STL_ScpInit` routine sets the value of the `$_PVCODE` global variable to `OAT100`. Using this variable allows the SCP to be used for the next version of the product, `OAT200`, without changing the pathnames.

2 The SCP uses the `rm` command to remove the links created in the `POST_L` phase. The `STL_LinkRemove` routine is used only to remove links created by the `STL_LinkCreate` routine.

In Example 3–6, the SCP returns the following codes to the `setld` utility:

- 0 - continue subset configuration
- 1 - terminate subset configuration; leave the subset corrupt

The `setld` utility creates an empty `subsetID.lk` lock file when it installs a subset. After successful installation, that subset is then available for dependency checks and locking is performed when other subsets are installed later. A subset's lock file can then contain any number of records, each naming a single dependent subset.

For example, the ODB kit requires that some version of the Orpheus Document Builder base product must be installed for the ODB product to work properly. Suppose that the `OATBASE200` subset is present. When

the `setld` utility installs the `OATODBTEMPS100` subset from the ODB kit, it inserts a record that contains the subset identifier `OATODBTEMPS100` into the `OATBASE200.lk` file. When the system administrator uses the `setld` utility to remove the `OATBASE200` subset, the `setld` utility checks `OATBASE200.lk` and finds a record that indicates that `OATODBTEMPS100` depends on `OATBASE200`, displays a warning message, and requires confirmation that the user really intends to remove the `OATBASE200` subset.

If the administrator removes the `OATODBTEMPS100` subset, the `setld` utility removes the corresponding record from the `OATBASE200.lk` file. Thereafter, the administrator can remove the `OATBASE200` subset without causing a dependency warning.

The SCP uses the `DEPS` field in the subset control file (`subset-id.ctrl`) to perform dependency locking.

3.4.5.4 After Securing the Subset (C INSTALL Phase)

After securing the subset, the `setld` utility sets the `ACT` environment variable to `C` (configuration) and calls the subset control program for each subset, passing `INSTALL` as an argument. At this time, the subset control program can perform any configuration operations required for product-specific tailoring. For example, a kernel kit can statically or dynamically configure a device driver at this point.

The `setld` utility enters the `C INSTALL` phase when `setld` is invoked with the `-l` (load) option.

Note

In a cluster environment, the `C INSTALL` phase runs multiple times, once on each cluster member. You must be able to run any SCP operations in the `C INSTALL` phase more than once without causing a problem.

If your SCP needs to access member-specific files, perform those operations during the `C INSTALL` phase.

The subset control program cannot create a layered product's symbolic links during the `C INSTALL` phase.

Example 3-7 shows the `C INSTALL` portion of the SCP that issues a message upon successful subset installation.

Example 3–7: Sample Message Output During C INSTALL Phase

```
# The ACT variable is set by setld and determines which
# phase of the SCP should be executed.
#
case $ACT in
:
:
#
# This is the configuration phase of the SCP
#
C)
#
# Setld invokes the C phase with an argument that is
# either INSTALL or DELETE. The INSTALL argument is
# used on a setld load, while the DELETE argument is
# used on a setld delete.
#
case $1 in
INSTALL) ❶
#
# Output a message letting the user know
# that they should check the odb.conf file
# before using the product.
#
echo "
The installation of the $_DESC ($_SUB) ❷
software subset is complete.
Please check the /opt/$_PVCODE/odb.conf file before ❸
using the $_DESC product." ❷
;;
:
:
esac
;;
:
:
:
```

❶ During the C phase, the SCP checks to see if the first argument passed by the setld utility has the value of INSTALL. If so, the program displays a message indicating that the installation is complete.

❷ The STL_ScpInit routine sets the value of the \$_DESC global variable to Orpheus Document Builder and the \$_SUB global variable to OATODB100, resulting in the following message:

```
The installation of the Orpheus Document Builder (OATODB100)
software subset is complete.
Please check the /opt/OAT100/odb.conf file before
using the Orpheus Document Builder product."
```

❸ The STL_ScpInit routine sets the value of the \$_PVCODE global variable to OAT100. Using this variable allows the SCP to be used for the next version of the product, OAT200, without changing the pathnames.

In Example 3–7, the SCP returns the following codes to the `setld` utility:

- 0 - successful load and configure
- 1 - unsuccessful load and configure; leave the subset corrupt

3.4.5.5 Before Deleting a Subset (C DELETE Phase)

When the user invokes the `setld` utility with the `-d` option, the utility sets the `ACT` environment variable to `C` and calls the subset control program for each subset, passing `DELETE` as an argument. At this time, the subset control program can make configuration modifications to remove evidence of the subset's existence from the system. For example, a kernel kit would deconfigure a statically or dynamically configured driver during this phase. The `C DELETE` phase should reverse any changes made during the `C INSTALL` phase.

Note

In a cluster environment, the `C DELETE` phase runs multiple times, once on each cluster member. You must be able to run any SCP operations in the `C DELETE` phase more than once without causing a problem.

The SCP always should return a zero exit status in the `C DELETE` phase. A nonzero return status tells the `setld` utility not to delete the software, but if the SCP has run the `C DELETE` phase on other cluster members the software already may be marked as corrupt.

If an operation fails, report the error and continue processing. The user must fix the problem after the software is removed.

The subset control program cannot remove a layered product's links during the `C DELETE` phase.

Example 3–8 shows the `C DELETE` portion of the SCP that would reverse any changes made during the `C INSTALL` phase.

Example 3–8: Sample C DELETE Phase

```
#
# The ACT variable is set by setld and determines which
# phase of the SCP should be executed.
#
case $ACT in
:
:
#
# This is the configuration phase of the SCP
#
C)
```

Example 3–8: Sample C DELETE Phase (cont.)

```
#
# Setld invokes the C phase with an argument that is
# either INSTALL or DELETE. The INSTALL argument is
# used on a setld load, while the DELETE argument is
# used on a setld delete.
#
case $1 in
INSTALL)
#
# Output a message letting the user know
# that they should check the odb.conf file
# before using the product.
#
echo "
The installation of the $_DESC ($_SUB)
software subset is complete.
Please check the /opt/$_PVCODE/odb.conf file before
using the $_DESC product."
;;
DELETE)
;; 1
esac
;;
:
:
```

- 1 This phase should reverse any changes made during the C INSTALL phase. Since no changes were made in Example 3–7, no action is taken in the C DELETE phase.

In Example 3–8, the SCP returns the following codes to the setld utility:

- 0 - continue with the deletion
- 1 - terminate the deletion

3.4.5.6 Before Deleting a Subset (PRE_D Phase)

When the user invokes the setld utility with the -d option, the utility sets the ACT environment variable to PRE_D and calls the subset control program for each subset. At this time, the subset control program can reverse modifications made during the POST_L phase of installation, such as removing links and dependency locks.

Note

In a cluster environment, the PRE_D phase runs only once for the entire cluster. If you must run member-specific operations when your subset is deleted, include code in your SCP that performs the operation on all cluster members in the C DELETE phase.

The SCP always should return a zero exit status in the `PRE_D` phase. A nonzero return status tells the `setld` utility not to delete the software, but since the SCP has run the `C DELETE` phase the software is already marked as corrupt.

If an operation fails, report the error and continue processing. The user must fix the problem after the software is removed.

You can call the following routines to remove links and unlock subsets:

`STL_LinkRemove`

Removes links created by `STL_LinkCreate` and restores any original files that `STL_LinkCreate` saved. Call `STL_ScpInit` first to initialize required global variables. The `STL_LinkRemove` routine cannot remove modified links.

The SCP uses the `DEPS` field in the subset control file (`subset-id.ctrl`) to perform dependency unlocking.

In Example 3–6, the SCP used `STL_LinkBack` in the `POST_L` phase to create the `/opt/OAT100/sbin/ls` link, referring to the `/sbin/ls` file. Example 3–9 shows the SCP removing this link in the `PRE_D` phase.

Example 3–9: Sample `PRE_D` Phase Reversal of `POST_L` Phase Actions

```
case $ACT in
    :
#
# This is the pre-deletion phase of the SCP
#
PRE_D)
#
# Remove the links created in the POST_L phase
#
    rm -f ./opt/${PVCODE}/sbin/ls 1
;;
    :

```

- 1** The SCP uses the `rm` command to remove the links created in the `POST_L` phase. The `STL_LinkRemove` routine is used only to remove links created by the `STL_LinkCreate` routine.

In Example 3–9, the SCP returns the following codes to the `setld` utility:

- 0 - continue with the deletion
- 1 - terminate the deletion

3.4.5.7 After Deleting a Subset (POST_D Phase)

During the `POST_D` phase, after deleting a subset, the `setld` utility sets the `ACT` environment variable to `POST_D` and calls the subset control program for each subset. At this time the subset control program can reverse any modifications made during the `PRE_L` phase of installation.

Note

In a cluster environment, the `POST_D` phase is run only once for the whole cluster. If you must run member-specific operations when your subset is deleted, include code in your SCP that performs the operation on all cluster members in the `C DELETE` phase.

The SCP always should return a zero exit status in the `POST_D` phase. A nonzero return status tells the `setld` utility not to delete the software, but the subset already has been removed. This causes cluster corruption.

If an operation fails, report the error and continue processing. The user must fix the problem after the software is removed.

In Example 3–10, the subset control program checks a list of files to be backed up if they already exist on the system. If it finds any, it restores the backup copy.

Example 3–10: Sample File Restoration During `POST_D` Phase

```

:
#
# Here is a list of files to back up if found on
# the installed system.
#
BACKUP_FILES="\
./usr/var/opt/$_PVCODE/templates/odb_template \ 1
./usr/var/opt/$_PVCODE/log_files/odb_log"
#
# The ACT variable is set by setld and determines which
# phase of the SCP should be executed.
#
case $ACT in
:
#
# This is the post-deletion phase of the SCP
#
POST_D)
```

Example 3–10: Sample File Restoration During POST_D Phase (cont.)

```
#
# Restore the backup copies created during the PRE_L phase
#

for FILE in $BACKUP_FILES
do
  [ -f $FILE.OLD ] &&
  mv $FILE.OLD $FILE [2]
done
;;
:
esac
```

- [1] The `STL_ScpInit` routine sets the value of the `$_PVCODE` global variable to `OAT100`. Using this variable allows the SCP to be used for the next version of the product, `OAT200`, without changing the pathnames.
- [2] Restores any files backed up in the `PRE_L` phase, as shown in Example 3–5.

In Example 3–10, the SCP returns the following codes to the `setld` utility:

- 0 - continue with the deletion
- 1 - terminate the deletion

3.4.5.8 Verifying the Subset (V Phase)

When the user invokes the `setld` utility with the `-v` option, the utility sets the `ACT` environment variable to `V` and calls the subset control program for each subset. Any `V` phase processing included in the subset control program is executed at this time.

The `setld` utility checks for the existence of the installed subset and if the subset exists, the `setld` utility verifies the size and checksum information for each file in the subset. The `setld` utility does not execute subset control program `V` phase processing during the installation process.

3.4.6 Stopping the Installation

Depending on the tests performed, your subset control program could decide to stop the installation or deletion of its subset. For example, if it finds a later version of the product already installed, the subset control program can stop the process.

To stop the installation or deletion of the subset, the subset control program must return a nonzero status to the `setld` utility upon exiting from the particular phase for which it was called. If the subset control program returns a status of 0 (zero), the `setld` utility assumes that the subset control program is satisfied that the `setld` process should continue.

3.4.7 Creating SCPs for Different Product Kit Types

This section provides examples of subset control programs for each of the different product kit types. The following topics are discussed:

- Creating SCPs for user product kits (Section 3.4.7.1)
- Creating SCPs for kernel product kits (Section 3.4.7.2)

3.4.7.1 User Product Kit SCPs

User product kits do not require subset control programs. You may need to provide one if your user product requires special installation tasks.

Example 3–11 shows a subset control program for the ODB user product, illustrating the types of operations that can be performed during different `setld` phases and illustrating one method of using the value of the `ACT` environment variable to determine what actions to perform.

Example 3–11: Sample ODB User Product SCP

```
#!/sbin/sh
#
# Load all of the standard SCP library routines
#
. /usr/share/lib/shell/libscp
#
# Initialize the global variables, except in the M phase
#
if [ "$ACT" != "M" ]
then
    STL_ScpInit
fi
#
# Here is a list of files to back up if found on
# the installed system.
#
BACKUP_FILES="\
./usr/var/opt/${PVCODE}/templates/odb_template \
./usr/var/opt/${PVCODE}/log_files/odb_log"
#
# The ACT variable is set by setld and determines which
# phase of the SCP should be executed.
#
case $ACT in
#
# This is the menu phase of the SCP
#
M)
#
# Setld invokes the M phase with an argument and if
```

Example 3–11: Sample ODB User Product SCP (cont.)

```
# the argument is "-l" it means that a software load
# is occurring.
#
case $1 in
-l)
#
# Examine the machine architecture to be sure
# that this software is being installed on an
# alpha machine. If it is not, exit with an
# error status so that setld will not display
# this subset on the menu of subsets to load.
#
ARCH=`./bin/machine`
[ "$ARCH" = alpha ] || exit 1
;;
esac
;;
PRE_L)
#
# Loop through the list of backup files and create
# backup copies for any file that is found on the
# system.
#
for FILE in $BACKUP_FILES
do
#
# If the file to be backed up exists, create
# a backup copy with a .OLD extension.
#
if [ -f $FILE ]
then
cp $FILE $FILE.OLD
fi
done
;;
POST_L)
#
# Initializes the variables so that the STL_LinkBack
# routine can be executed
#
STL_LinkInit
#
# Create a symbolic link in the ./opt/$_PVCODE/sbin
# directory that points to the ./sbin/ls file.
#
STL_LinkBack ls ./sbin ./opt/$_PVCODE/sbin
;;
PRE_D)
#
# Remove the links created in the POST_L phase
#
rm -f ./opt/$_PVCODE/sbin/ls
;;
POST_D)
#
# Restore the backup copies created during the PRE_L phase
#
for FILE in $BACKUP_FILES
do
[ -f $FILE.OLD ] &&
mv $FILE.OLD $FILE
done

```

Example 3–11: Sample ODB User Product SCP (cont.)

```
done
;;
C)
#
# Setld invokes the C phase with an argument that is
# either INSTALL or DELETE. The INSTALL argument is
# used on a setld load, while the DELETE argument is
# used on a setld delete.
#
case $1 in
INSTALL)
#
# Output a message letting the user know
# that they should check the odb.conf file
# before using the product.
#
echo "
The installation of the $_DESC ($_SUB)
software subset is complete.

Please check the /opt/$_PVCODE/odb.conf file before
using the $_DESC product."
;;
DELETE)
;;
esac
;;
esac
exit 0
```

3.4.7.2 Kernel Product Kit SCPs

In addition to the optional processing described in Section 3.4.5, a subset control program for a kernel product such as a device driver also must configure the driver into the kernel. When building subset control programs for a kernel product, you can choose one of the following configuration strategies:

- Write one subset control program for a kit that contains the software subset associated with the single binary module for a statically configured driver.
- Write one subset control program for a kit that contains the software subset associated with the single binary module for a dynamically configured driver.
- Write one subset control program for a kit that contains the software subsets associated with the device driver that can be statically or dynamically configured.

Example 3–12 shows the subset control program for the single binary module associated with the `odb_kernel` driver. The user can choose to configure

this single binary module into the kernel either statically or dynamically. The subset control program runs the doconfig utility to configure the driver into the kernel.

Example 3–12: Sample ODB Kernel Product SCP

```
#!/sbin/sh
#
# Load all of the standard SCP library routines
#
. /usr/share/lib/shell/libscp
#
# Load the standard Error library routines
# (location of the Error routine)
#
. /usr/share/lib/shell/Error
#
# Load the standard String library routines
# (location of the ToUpper routine)
#
. /usr/share/lib/shell/Strings
#
# This routine rebuilds the static kernel
#
Rebuild_Static_Kernel()
{
  HNAME=`/sbin/hostname -s`
  HOSTNAME=`ToUpper $HNAME`
  if doconfig -c $HOSTNAME
  then
    echo "\nThe /sys/${HOSTNAME}/vmunix kernel has been"
    echo "moved to /vmunix and the changes will take effect"
    echo "the next time the system is rebooted."
    return 0
  else
    Error "
An error occurred while building the static kernel."
    return 1
  fi
}
KERNEL=/cluster/members/{memb}/boot_partition/vmunix
#
# Initialize the global variables, except in M phase
#
if [ "$ACT" != "M" ]
then
  STL_ScpInit
fi
#
# The ACT variable is set by setld and determines which
# phase of the SCP should be executed.
#
case $ACT in
C)
  #
  # The kreg database file where all of the kernel
  # layered products are registered.
  #
  KREGFILE=./usr/sys/conf/.product.list
  case $1 in
  INSTALL)
    #
    # Merge the graphics support into the existing
```

Example 3–12: Sample ODB Kernel Product SCP (cont.)

```
# /etc/sysconfigtab file
#
sysconfigdb -m -f ./opt/$_PVCODE/etc/sysconfigtab odb_graphics
echo "*** $_DESC Product Installation Menu ***\n"
echo "1. Statically configure the graphics support"
echo "2. Dynamically configure the graphics support"
echo "\nType the number of your choice []: \c"
read answer
case ${answer} in
1)
#
# Determine if the product is already registered
# with the kreg database, and if it is, skip
# registering it.
#
grep -q $_SUB $KREGFILE
if [ "$?" != "0" ]
then
#
# Register the product with the
# kernel using kreg
#
/sbin/kreg -l $_PCODE $_SUB \
./opt/$_PVCODE/sys/BINARY
fi
#
# Rebuild the static kernel
#
Rebuild_Static_Kernel
#
# Successful rebuild, so back up the existing
# kernel and move the new one into place.
#
if [ "$?" = "0" ]
then
#
# Make a backup copy of the kernel
# as it existed prior to installing
# this subset. Since a subset can
# be installed more than once (due to
# load/configuration failures or even
# because the user removed files) make
# sure that the backup does not already
# exist.
#
if [ ! -f $KERNEL.pre_${_SUB} ]
then
mv $KERNEL $KERNEL/pre_${_SUB}
fi
#
# Move the new kernel into place
#
mv /sys/${HOSTNAME}/vmunix /vmunix
#
# Place a marker on the system so that
# upon subset removal the SCP can
# determine if it needs to remove a
# static or dynamic configuration.
#
touch ./opt/$_PVCODE/sys/BINARY/odb_graphics_static
```

Example 3–12: Sample ODB Kernel Product SCP (cont.)

```
fi
;;
2)
#
# Dynamically load the odb_graphics subsystem
# into the kernel
#
sysconfig -c odb_graphics
;;
esac
;;
DELETE)
#
# If the marker is present then the kernel option
# was added statically.
#
if [ -f ./opt/$_PVCODE/sys/BINARY/odb_graphics_static ]
then
#
# Clean-up the marker
#
rm -f ./opt/$_PVCODE/sys/BINARY/odb_graphics_static
#
# Deregister the product using kreg
#
/sbin/kreg -d $_SUB
#
# Rebuild the static kernel
#
Rebuild_Static_Kernel
#
# Successful rebuild, remove the old backup
# copy that was created when we installed.
#
if [ "$?" = "0" ]
then
mv /sys/${HOSTNAME}/vmunix /vmunix
rm -f $KERNEL.pre_${_SUB}
fi
else
#
# Unload the dynamic kernel module
#
sysconfig -u odb_graphics
fi
#
# Remove the entry from the /etc/sysconfigtab file
#
sysconfigdb -d odb_graphics
;;
esac
;;
esac
exit 0
```

3.5 Producing Subsets

After you create the master inventory and key files, run the `kits` utility to produce subsets and control files. The `kits` utility creates the following files:

- The compression flag file (Section 3.5.1)
- The image data file (Section 3.5.2)
- The subset control files (Section 3.5.3)
- The subset inventory files (Section 3.5.4)

Caution

Do not create these files before you run the `kits` utility.

Use the following syntax for the `kits` command:

kits *key-file* *input-path* *output-path* [*subset*]...

key-file

This mandatory parameter is the pathname of the key file created in Section 3.3.

input-path

This mandatory parameter specifies the top of the file hierarchy that contains the source files.

output-path

This mandatory parameter specifies the directory used to store the subset images and data files produced.

subset

This optional parameter specifies the name of an individual subset to be built. You may specify multiple subsets in a space-separated list. If you use the *subset* argument, the `kits` utility assumes the following:

- Only the subsets named as arguments to this parameter are to be built.
- The *key-file* contains descriptors for each of the named subsets.
- All other subsets in the product have been built already.
- The *output-path* directory contains images of the previously built subsets.

If you do not use the *subset* argument, the `kits` utility builds all subsets listed in the key file.

Refer to the `kits(1)` reference page for more information.

Caution

The master inventory file (`*.mi`) and the key file (`*.k`) are typically in the same directory. If they are not, the `MI=` attribute in the key file must contain the explicit relative path from the directory where you are running the `kits` utility to the directory where the master inventory file resides. The `scps` directory that contains any subset control programs must be in the same directory where the `kits` utility is invoked.

Example 3–13 shows a sample of using the `kits` utility to build the subsets for the ODB product kit:

Example 3–13: Using the `kits` Utility to Build ODB Subsets

```
# cd /mykit/data
# kits OAT100.k ../src ../output
%#
Creating 2 Orpheus Document Builder subsets.
1 Subset OATODB100 [1]
  Generating media creation information...done
  Creating OATODB100 control file...done.
  Making tar image...done. [2]
  Compressing [3]
    OATODB100: Compression: 92.64%
  -- replaced with OATODB100.Z

  *** Finished creating media image for OATODB100. ***

2 Subset OATODBTEMP100 [1]
  Generating media creation information...done
  Creating OATODBTEMP100 control file...done.
  Making tar image...done.
  Compressing [3]
    OATODBTEMP100: Compression: 98.39%
  -- replaced with OATODBTEMP100.Z
  Null subset control program created for OATODBTEMP100.

  *** Finished creating media image for OATODBTEMP100. ***

Creating OAT.image [4]

Creating INSTCTRL [5]
a OAT.image 1 Blocks
a OAT100.comp 0 Blocks
a OATODB100.ctrl 1 Blocks
a OATODB100.inv 2 Blocks
a OATODB100.scp 7 Blocks
a OATODBTEMP100.ctrl 1 Blocks
a OATODBTEMP100.inv 0 Blocks
a OATODBTEMP100.scp 0 Blocks

Media image production complete.
```

In Example 3–13, the `kits` utility performs the following steps and reports its progress:

- ❶ Creates the subsets.
- ❷ If the subset is not in DCD format, creates a `tar` image of the subset.
- ❸ Compresses each subset if you specify the `COMPRESS` attribute in the key file.
- ❹ Creates the image data file `OAT.image`.
- ❺ Creates the `INSTCTRL` file, which contains a `tar` image of all the following installation control files:
 - Compression flag file `product-id.comp`
 - Image data file `product-code.image`
 - Subset control file `subset-id.ctrl`
 - Subset inventory file `subset-id.inv`
 - Subset control program file `subset-id.scp`

These files are described in Table 3–9.

The `INSTCTRL` file is placed in the output directory.

Table 3–9 shows the installation control files in the `instctrl` directory after you run the `kits` utility.

Table 3–9: Installation Control Files in the `instctrl` Directory

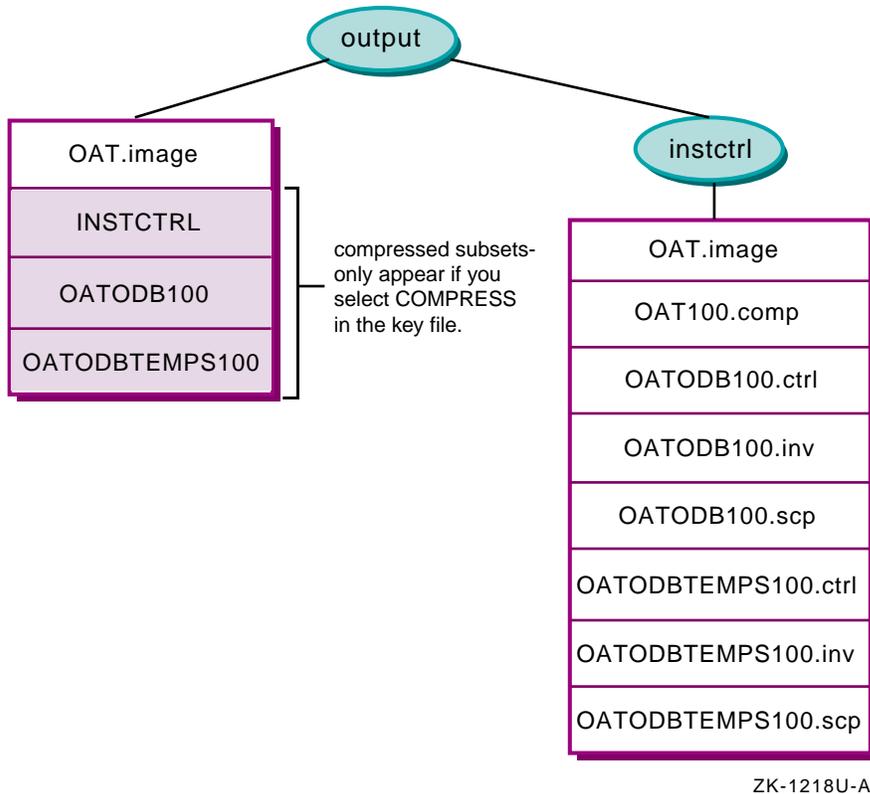
File	Description
<code>product-id.comp</code>	Compression flag file. This empty file is created only if you specified the <code>COMPRESS</code> attribute in the key file. Its presence signals to the <code>setld</code> utility that the subset files are compressed. The ODB kit's compression flag file is named <code>OAT100.comp</code> .
<code>product-code.image</code>	Image data file. This file contains size and checksum information for the subsets. The ODB kit's image data file is named <code>OAT.image</code> .
<code>subset-id.ctrl</code>	Subset control file. This file contains the <code>setld</code> utility control information. There is one subset control file for each subset. The ODB kit's subset control files are named <code>OATODB100.ctrl</code> and <code>OATODBTMP100.ctrl</code> .

Table 3–9: Installation Control Files in the instctrl Directory (cont.)

File	Description
<i>subset-id.inv</i>	Subset inventory file. This file contains an inventory of the files in the subset. Each record describes one file. There is one subset inventory file for each subset. The ODB kit's subset inventory files are named <i>OATODB100.inv</i> and <i>OATODBTEMPS100.inv</i> .
<i>subset-id.scp</i>	Subset control program. If you created subset control programs for your kit, these files are copied from the <i>scps</i> directory to the <i>instctrl</i> directory. There is one subset control program for each subset; if you have not created a subset control program for a subset, the <i>kits</i> utility creates a blank file. The ODB kit's subset control program files are named <i>OATODB100.scp</i> and <i>OATODBTEMPS100.scp</i> .

Figure 3–3 shows the contents of the output directory after you run the *kits* utility.

Figure 3–3: ODB output Directory



The subset files and the files in the `instctrl` directory are constituents of the final kit. The following sections describe the contents of the installation files.

3.5.1 Compression Flag File

The compression flag file is an empty file whose name consists of the product code and the version number with the string `comp` as a suffix; for example, `OAT100.comp`. If the compression flag file exists, the `setld` utility knows that the subset files are compressed.

3.5.2 Image Data File

The image data file is used by the `setld` utility to verify subset image integrity before starting the actual installation process. The image data file name consists of the product code with the string `image` as a suffix. The image data file contains one record for each subset in the kit, with three fields in each record.

Table 3–10 describes the image data file.

Table 3–10: Image Data File Field Descriptions

Field	Description
Checksum	The modulo-65536 (16-bit) checksum of the subset file, as provided by the <code>sum</code> utility. If the file is compressed, the checksum after compression. ^a
Size	The size of the subset file in kilobytes. If the file is compressed, the size after compression.
Subset identifier	The product code, subset mnemonic, and version number. For example, <code>OATODB100</code> .

^a Refer to the `sum(1)` reference page.

Example 3–14 shows `OAT.image`, the image data file for the ODB kit:

Example 3–14: Sample Image Data File

```
13601 10 OATODB100
12890 10 OATODBTMP100
```

3.5.3 Subset Control Files

The `setld` utility uses the subset control files as a source of descriptive information about subsets. Subset control file fields are described in Table 3–11.

Table 3–11: Subset Control File Field Descriptions

Field	Description
NAME	Specifies the product name. This value is from the Name field in the Key File.
DESC	Briefly describes the subset. This value is from the Subset Description field in the Subset Descriptor section of the Key File.
ROOTSIZE	Specifies (in bytes) the space the subset requires in the <code>root (/)</code> file system.
USRSIZE	Specifies (in bytes) the space the subset requires in the <code>usr</code> file system. This value is calculated by the <code>kits</code> utility.
VARSIZE	Specifies (in bytes) the space the subset requires in the <code>var</code> file system. This value is calculated by the <code>kits</code> utility.
NVOLS	Specifies disk volume identification information as two colon-separated integers (the volume number of the disk that contains the subset archive and the number of disks required to contain the subset archive). This value is calculated by the <code>kits</code> utility.

Table 3–11: Subset Control File Field Descriptions (cont.)

Field	Description
MTLOC	Specifies the tape volume number and subset's location on the tape as two colon-separated integers (the volume number of the tape that contains the subset archive and the file offset at which the subset archive begins). On tape volumes, the first three files are reserved for a bootable operating system image and are not used by the <code>setld</code> utility. An offset of 0 (zero) indicates the fourth file on the tape. The fourth file is a <code>tar</code> archive named <code>INSTCTRL</code> , which contains the kit's installation control files (listed in Table 3–9). This value is calculated by the <code>kits</code> utility.
DEPS	Specifies either a list of subsets upon which this subset is dependent (<code>DEPS="OATODB100 OSFDCMT510"</code>), or a single period (<code>DEPS="."</code>) indicating that there are no subset dependencies. If there is more than one subset dependency, each subset name is separated by a Space character. This value is from the Dependency List field in the Subset Descriptor section of the Key file. You may use the following wildcard characters when you specify subset names in the <code>DEPS</code> field: <ul style="list-style-type: none">• An asterisk (*) represents any number of characters. For example, <code>OAT*100</code> will match <code>OAT100</code>, <code>OATODB100</code>, <code>OATODBTEMPS100</code>, and so on.• A question mark (?) represents a single numeric character. For example, <code>OATODB1??</code> matches <code>OATODB100</code>, <code>OATODB101</code>, and so on up to <code>OATODB199</code>.
FLAGS	Specifies the value in the flags field of the subsets record in the key file. This value is from the Flags field in the Subset Descriptor section of the Key file. Bit 0 indicates whether the subset can be removed (0=removable, 1=protected). Bit 1 indicates whether the subset is mandatory (0=mandatory, 1=optional). Bit 2 indicates whether the subset is compressed (0=compressed, 1=uncompressed). Bits 3 to 7 are reserved; bits 8 to 15 are undefined.

Example 3–15 shows `OATODB100.ctrl`, the control file for the ODB kit's `OATODB100` subset:

Example 3–15: Sample Subset Control File

```
NAME='Orpheus Document Builder OATODB100'  
DESC='Document Builder Tools'  
ROOTSIZE=16668  
USRSIZE=16459  
VARSIZE=16384
```

Example 3–15: Sample Subset Control File (cont.)

```
NVOLS=1:0
MTLOC=1:1
DEPS="."
FLAGS=0
```

3.5.4 Subset Inventory File

The subset inventory file describes each file in the subset, listing its size, checksum, permissions, and other information. The `kits` utility generates this information, which reflects the exact state of the files in the source hierarchy from which the kit was built. The `setld` utility uses the information to duplicate that state, thus transferring an exact copy of the source hierarchy to the customer's system. Table 3–12 describes subset inventory file fields.

Each record of the inventory is composed of 12 fields, each separated by single Tab characters. Table 3–12 describes the contents of these fields.

Table 3–12: Subset Inventory File Field Descriptions

Name	Description
Flags	A 16-bit unsigned integer. Bit 1 is the v (volatility) bit. When this bit is set, changes to an existing copy of the file can occur either during or after kit installation. The volatility bit usually is set for log files such as <code>/usr/spool/mqueue/syslog</code> . Valid values for this field are 0 (protected) or 2 (unprotected).
Size	The actual number of bytes in the file.
Checksum	The modulo-65536 (16-bit) checksum of the file.
uid	The user ID of the file's owner.
gid	The group ID of the file's owner.
Mode	The six-digit octal representation of the file's mode.
Date	The file's last modification date.
Revision	The version code of the product that includes the file.
Type	A letter that describes the file: b – Block device. c – Character device. d – Directory containing one or more files.

Table 3–12: Subset Inventory File Field Descriptions (cont.)

Name	Description
	<p>f – Regular file. For regular files with a link count greater than one, see file type l.</p> <p>l – Hard link. Other files in the inventory have the same inode number. The first (in ASCII collating sequence) is listed in the referent field.</p> <p>p – Named pipe (FIFO).</p> <p>s – Symbolic link.</p>
Pathname	The dot-relative (./) pathname of the file.
Referent	For file types l and s, the path to which the file is linked; for types b and c, the major and minor numbers of the device; for all other types, none.
Subset identifier	The name of the subset that contains the file.

Example 3–16 shows the OATODB100.inv inventory file for the ODB kit's OATODB100 subset.

Example 3–16: Sample ODB Product Subset Inventory File

```

0 512 00000 0 0 040755 5/15/00 100 d\
./cluster/members/member0/opt/OAT100 none OATODB100
2 44 56771 0 0 100644 5/15/00 100 f\
./cluster/members/member0/opt/OAT100/odb.conf none OATODB100
0 512 56771 0 0 040755 5/15/00 100 d\
./opt/OAT100 none OATODB100
2 51 00000 0 0 120777 5/15/00 100 s\
./opt/OAT100/odb.conf\
../../../../cluster/members/{memb}/opt/OAT100/odb.conf OATODB100
0 512 00000 0 0 040755 5/15/00 100 d\
./opt/OAT100/sbin none OATODB100
0 28 06280 0 0 100644 5/15/00 100 f\
./opt/OAT100/sbin/odb_recover none OATODB100
0 512 06280 0 0 040755 5/15/00 100 d\
./usr/opt/OAT100 none OATODB100
0 512 06280 0 0 040755 5/15/00 100 d\
./usr/opt/OAT100/bin none OATODB100
0 27 33168 0 0 100644 5/15/00 100 f\
./usr/opt/OAT100/bin/odb_start none OATODB100
0 512 33168 0 0 040755 5/15/00 100 d\
./usr/var/cluster/members/member0/opt/OAT100 none OATODB100
0 23 43390 0 0 100644 5/15/00 100 f\
./usr/var/cluster/members/member0/opt/OAT100/odb_log none OATODB100
0 512 43390 0 0 040755 5/15/00 100 d\
./usr/var/opt/OAT100 none OATODB100
0 512 43390 0 0 040755 5/15/00 100 d\
./usr/var/opt/OAT100/log_files none OATODB100
0 58 00000 0 0 120777 5/15/00 100 s\
./usr/var/opt/OAT100/log_files/odb_log\
../../../../usr/var/cluster/members/{memb}/opt/OAT100/odb_log OATODB100

```

Note

The backslashes (\) in Example 3–16 indicate line continuation and are not present in the actual file.

Fields are separated by single Tab characters.

3.6 Testing Subsets

You must test your subsets to ensure that they can be loaded onto a running system, that the product runs on the system, and that the subsets can be deleted. The following sections describe this testing:

1. Loading all of the subsets onto a running system. (Section 3.6.1)
2. Repeating the subset loading test. (Section 3.6.2)
3. Deleting all of the subsets from a running system. (Section 3.6.3)
4. If your kit includes optional subsets, loading only the mandatory subsets onto a running system. (Section 3.6.4)
5. If your kit may run in a cluster environment, testing on a cluster. (Section 3.6.5)

Caution

It is important that you perform these tests in sequence.

3.6.1 Loading All Subsets

The examples in this section assume that your kit consists of the mandatory OATODB100 subset and the optional OATODBTMP100 subset, and that it resides in the `/mykit/output` directory.

Follow these steps to load all subsets:

1. Log in to the system as `root` or use the `su` command to gain superuser privileges.
2. Use the `setld` utility to load all of your subsets onto the system, as in the following example:

```
# setld -l /mykit/output
```
3. When prompted, select the option to install all subsets from the `setld` installation menu.
4. Verify that all files in your subsets were loaded. If any files are missing, check the master inventory file. Subset inventory files are created from master inventory file entries.

5. Verify each file's installed location, permissions, owner, and group. The `setld` utility uses the information in the subset inventory file to determine these attributes. If any are incorrect, modify the file in the source directory and rebuild the master inventory file and the subsets.
6. If you supplied SCP files, verify any actions that should have occurred in the `M`, `PRE_L`, `POST_L`, and `C INSTALL` phases. Refer to Section 3.4.5 for discussions of SCP tasks associated with these phases.
7. After successful installation, test all commands or utilities included with your product. Since file locations may have changed, especially if you installed in the `/opt`, `/usr/opt`, or `/usr/var/opt` directories, it is important that you test your product thoroughly to verify that everything works correctly.
8. Repeat the test to confirm that the SCP does not fail when it runs more than once.
 - a. Use the `setld -l` command to reload all of your subsets onto the system.
 - b. Verify that all files in your subsets were loaded. If any files are missing, check the master inventory file.
 - c. Verify each file's installed location, permissions, owner, and group. If any are incorrect, modify the file in the source directory and rebuild the master inventory file and the subsets.
 - d. If you supplied SCP files, verify any actions that should have occurred in the `M`, `PRE_L`, `POST_L`, and `C INSTALL` phases. Refer to Section 3.4.5 for discussions of SCP tasks associated with these phases.
 - e. After successful installation, test all commands or utilities included with your product.

3.6.2 Repeating Subset Load

Follow these steps to load all subsets:

1. Log in to the system as `root` or use the `su` command to gain superuser privileges.
2. Use the `setld` utility to load all of your subsets onto the system, as in the following example:

```
# setld -l /mykit/output
```
3. When prompted, select the option to install all subsets from the `setld` installation menu.

4. Verify that all files in your subsets were loaded. If any files are missing, check the master inventory file. Subset inventory files are created from master inventory file entries.
5. Verify each file's installed location, permissions, owner, and group. The `setld` utility uses the information in the subset inventory file to determine these attributes. If any are incorrect, modify the file in the source directory and rebuild the master inventory file and the subsets.
6. If you supplied SCP files, verify any actions that should have occurred in the `M`, `PRE_L`, `POST_L`, and `C INSTALL` phases. Refer to Section 3.4.5 for discussions of SCP tasks associated with these phases.
7. After successful installation, test all commands or utilities included with your product. Since file locations may have changed, especially if you installed in the `/opt`, `/usr/opt`, or `/usr/var/opt` directories, it is important that you test your product thoroughly to verify that everything works correctly.

3.6.3 Removing All Subsets

The examples in this section assume that your kit consists of the mandatory `OATODB100` subset and the optional `OATODBTMP100` subset, and that it resides in the `/mykit/output` directory.

Follow these steps to remove all subsets:

1. Log in to the system as `root` or use the `su` command to gain superuser privileges.
2. Use the `setld` utility to delete all of your subsets from the system, as in the following example:

```
# setld -d OATODB100 OATODBTMP100
```
3. Verify that all files loaded onto your system in Section 3.6.1 were deleted.
4. If you supplied SCP files, verify any actions that should have occurred in the `C DELETE`, `PRE_D`, and `POST_D` phases. Refer to Section 3.4.5 for discussions of SCP tasks associated with these phases.

3.6.4 Loading Mandatory Subsets Only

The examples in this section assume that your kit consists of the mandatory `OATODB100` subset and the optional `OATODBTMP100` subset, and that it resides in the `/mykit/output` directory.

Follow these steps to load only the mandatory subsets:

1. Log in to the system as `root`.
2. Use the `setld` utility to load all of your subsets onto the system, as in the following example:

```
# setld -l /mykit/output
```
3. When prompted, select the option to install only mandatory subsets from the `setld` installation menu.
4. Verify that all mandatory files in your subsets were loaded. If any files are missing, check the master inventory file. Subset inventory files are created from master inventory file entries.
5. Verify each file's installed location, permissions, owner, and group. The `setld` utility uses the information in the subset inventory file to determine these attributes. If any are incorrect, modify the file in the source directory and rebuild the master inventory file and the subsets.
6. If you supplied SCP files, verify any actions that should have occurred in the `M`, `PRE_L`, `POST_L`, and `C INSTALL` phases. Refer to Section 3.4.5 for discussions of SCP tasks associated with these phases.
7. After successful installation, test all commands or utilities included with your product. Since file locations may have changed, especially if you installed in the `/opt`, `/usr/opt`, or `/usr/var/opt` directories, it is important that you test your product thoroughly to verify that everything works correctly.

If your product does not work correctly, some of the files in your optional subsets may need to be moved to mandatory subsets.

3.6.5 Testing in a Cluster

To test your product kit in a cluster, you must ensure that your subsets can be loaded onto a running cluster, that the product runs on the cluster, and that the subsets can be deleted from the cluster. The following sections describe this testing:

1. Loading all of the subsets onto a cluster. (Section 3.6.5.1)
2. Deleting all of the subsets from a cluster. (Section 3.6.5.2)

Caution

It is important that you perform these tests in sequence.

The examples in this section assume that your kit consists of the mandatory OATODB100 subset and the optional OATODBTMP100 subset, and that it resides in the `/mykit/output` directory.

3.6.5.1 Loading the Kit onto a Cluster

Follow these steps to load the product kit onto a cluster:

1. Log in to the cluster as `root`.
2. Use the `setld` utility to load all of your subsets onto the cluster, as in the following example:

```
# setld -l /mykit/output
```
3. When prompted, select the option to install all subsets from the `setld` installation menu.
4. Verify that all files in your subsets were loaded. If any files are missing, check the master inventory file. Subset inventory files are created from master inventory file entries.
5. Verify each file's installed location, permissions, owner, and group. The `setld` utility uses the information in the subset inventory file to determine these attributes. If any are incorrect, modify the file in the source directory and rebuild the master inventory file and the subsets.
6. Perform the following checks on each cluster member:
 - a. Verify each member-specific file's location, permissions, owner, and group. The `setld` utility uses the information in the subset inventory file to determine these attributes. If any are incorrect, modify the file in the source directory and rebuild the master inventory file and the subsets.
 - b. Verify that each CDSL can be accessed and that it contains the correct information for each member.
 - c. If you supplied SCP files, verify any actions that should have occurred in the `M`, `PRE_L`, `POST_L`, and `C INSTALL` phases. Refer to Section 3.4.5 for discussions of SCP tasks associated with these phases.
 - d. After successful installation, test all commands or utilities included with your product. Since file locations may have changed, especially if you installed in the `/opt`, `/usr/opt`, or `/usr/var/opt` directories, it is important that you test your product thoroughly to verify that everything works correctly.

3.6.5.2 Deleting the Kit from a Cluster

Follow these steps to delete the kit from a cluster:

1. Log in to the cluster as `root` or use the `su` command to gain superuser privileges.
2. Use the `setld` utility to delete all of your subsets from the cluster, as in the following example:

```
# setld -d OATODB100 OATODTEMP100
```

3. Verify that all files loaded onto your system in Section 3.6.1 were deleted.
4. Perform the following checks on each cluster member:
 - a. Verify that each member-specific file was removed.
 - b. If you supplied SCP files, verify any actions that should have occurred in the `C DELETE`, `PRE_D`, and `POST_D` phases. Refer to Section 3.4.5 for discussions of SCP tasks associated with these phases.
 - c. Verify that all files not in the inventory are deleted. This includes any files created when your kit was installed.

3.7 Updating Inventory After Creating Subsets

You may have to update the master inventory file after you have created subsets. For example, kernel product kits require additional files, some of which must be added to your kit's inventory.

After you update the kit source directory, run the `newinv` utility to update the master inventory file, using the existing master inventory file as input. The `newinv` utility performs the following additional steps:

1. Creates a backup file, *inventory-file.bkp*.
2. Finds all the file and directory names in the source hierarchy.
3. Produces the following sorted groups of records:
 - Records that contain pathnames only, representing files now present that were not in the previous inventory (new records).
 - Records that represent files now present that were also present in the previous inventory. This list is empty the first time you create the inventory.
 - Records that were in the previous inventory but are no longer present (defunct records). This list is also empty the first time you create the inventory.

4. Lets you edit the group of defunct records, deleting records for files that no longer belong in the kit.
5. Lets you edit the group of new records by adding the flags and subset identification fields (see Table 3-1).
6. Merges the three groups of records and sorts the result to produce a finished master inventory file that matches the source hierarchy.

Run the `newinv` utility to update the master inventory file any time that you add, modify, or remove files in the kit's source directory. After you update the master inventory file, run the `kits` utility as described in Section 3.5 to produce updated subsets and control files.

4

Producing User Product Kits

This chapter tells you how to produce a user product kit. A user product runs in user space. This includes commands and utilities as well as applications such as text editors and database systems. Users interact directly with user products through commands or window interfaces.

Note

The information in this chapter describes how to produce user product kits. If you want to create a kernel product kit, go to Chapter 5.

Follow these steps to create and test a user product kit:

1. Read Chapter 1 for an overview of product kits.
2. Design the kit directory structure as described in Chapter 2.
3. Create subsets as described in Chapter 3. Subset control programs are optional for user product kits (Section 3.4).
4. Create the kit distribution media as described in Section 4.2.
5. Test the distribution media as described in Section 4.3.

No additional installation files are required for user product kits.

4.1 Overview

A user product is a layered product that contains software run directly by users. Commands and utilities are in this category, as are applications such as text editors and database systems. Users interact directly with user products through such means as commands or graphical interfaces.

4.2 Producing Distribution Media

After you have tested the subsets as described in Section 3.6, you can produce the distribution media.

Distribution media production consists of the following tasks:

1. Edit the `/etc/kitcap` file. (Section 4.2.1)
2. Build the user product kit on the distribution media:
 - Use the `gendisk` utility to build a kit on disk media (Section 4.2.2)
 - Use the `gentapes` utility to build a `tar` format kit on magnetic tape (Section 4.2.3)

Produce user product kits in `tar` format. You can use direct CD-ROM (DCD) format if you require access to kit files before or during installation, but installation time for DCD format kits is slower than for `tar` format kits.

- `tar` format

In `tar` format, the product files in each subset are written to the distribution media as a single file. During installation, the `setld` utility uncompresses the files and moves them onto the target system, preserving the files' original directory structure. Kits distributed in `tar` format install more quickly and consume less space on the distribution media.

- direct CD-ROM (DCD) format

In DCD format, the files are written to the distribution media as a UNIX file system where the product files are organized into a directory structure that mirrors the target system. Subsets distributed in DCD format cannot be compressed.

You can distribute user product kits on diskette, CD-ROM, or magnetic tape, as follows:

- Diskette

Diskettes are a good media for testing purposes or for small products. The product must fit on a single diskette; it cannot span multiple diskettes. Use the `gendisk` utility to produce kits for diskette media.

- CD-ROM

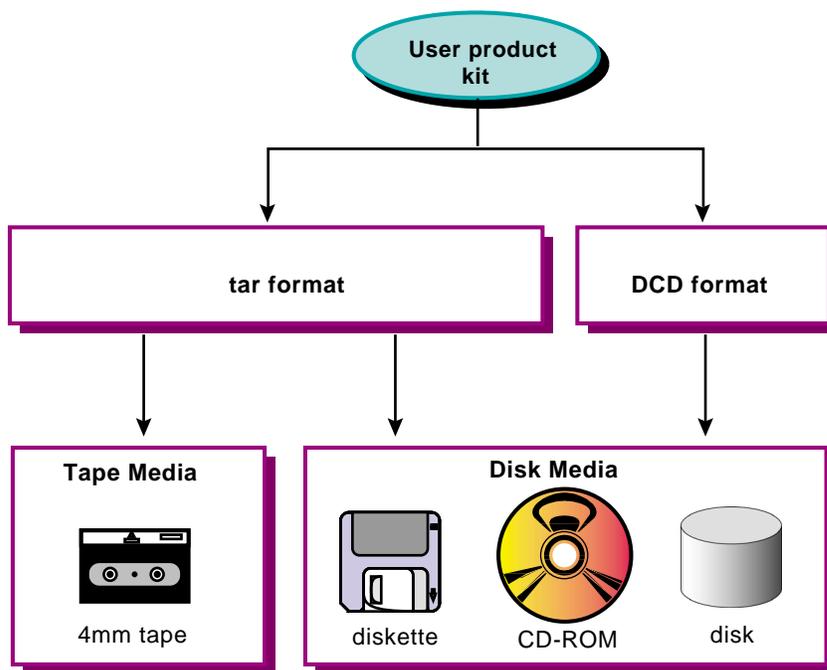
CD-ROM media can support large kits or multiple kits on a single media. The kit is first produced on the hard disk, then written onto the CD-ROM. Use the `gendisk` utility to produce the master kit on hard disk. Follow the CD-ROM manufacturer's instructions for writing the kit onto the CD-ROM media.

- Magnetic tape

You can distribute kits for user products on magnetic tape. Tape media does not support DCD format. Use the `gentapes` utility to produce kits for magnetic tape media.

Figure 4–1 shows the types of file formats and distribution media that are available for user product kits.

Figure 4–1: User Product Kit File Formats



ZK-1215U-AI

4.2.1 Editing the /etc/kitcap File

The `gendisk` and `gentapes` utilities refer to the `/etc/kitcap` file, a database containing information about the kits to be built on the system. Each record contains a product code and the names of the directories, files, and subsets that make up the product kit. Before you can build your kit, you must add a media descriptor record to the `/etc/kitcap` database.

Note

If you use the `gendisk` utility to produce your kit on disk distribution media, you can specify an alternate kit descriptor database. Refer to the `gendisk(1)` reference page for additional information.

Use the following conventions when you add a record to the `/etc/kitcap` file:

- Separate the first field from the rest of the record by a colon (:) for disk media descriptors and by a pipe character (|) for tape media descriptors.
- Separate all other fields with colons (:).
- Indicate continuation with a backslash (\) at the end of the line.
- Lines starting with a pound sign (#) are comments and are ignored.
- Comments within the record start with pound sign (#) and end with a colon (:). Use this feature sparingly.

The contents of a `kitcap` record differ depending on whether you are producing disk or tape media. You must add one record for each media type on which you plan to distribute your kit.

The contents of the record also can depend on the product type you are delivering. Refer to the `kitcap(4)` reference page for more information about the contents of the `/etc/kitcap` file.

4.2.1.1 Disk Media Descriptor

Create a disk media `kitcap` record when you produce kits for distribution on diskette or CD-ROM. The `kitcap` record for disk media contains the following elements:

- The kit name, consisting of two parts:
 - The product code, consisting of the product code and version number specified in the `CODE` and `VERS` fields of the kit's key file. Refer to Section 3.3 for information about the key file.
 - The media code `HD` to indicate disk media. This element is followed by a colon (:).
- The partition on the disk media where the product should be placed. The partition is a letter between `a` and `h`. Partition `c` is used most often, as it spans the entire disk.
- The destination directory for the subsets on the disk media. This allows a hierarchical structure so you can put multiple products on one disk, or put parts of one product on different areas of the same disk. You can use multiple destination directories in a `kitcap` record.
- The product description. This entry is taken from key file `NAME` field. Replace any spaces with an underscore (`_`) character, for example: `Product Description` becomes `Product_Description`.
- The name of the output directory where you created the kit, where the `gendisk` utility can find the product subsets.

- The `instctrl` directory, relative to the output directory specification.
- The names of the subsets that make up the kit.

Refer to the `kitcap(4)` reference page for more detailed information about the disk media record format.

Refer to Section 3.3 for information about the key file.

Example 4–1 shows the record to be added to the `/etc/kitcap` file to produce the ODB kit on disk media:

Example 4–1: Sample Disk Media Descriptor for User Product

```
OAT100HD:c:/:\
dd=/OAT100:Orpheus_Document_Builder:/mykit/output:\
instctrl:OATODB100:OATODBTEMPS100
```

Based on the information shown in Example 4–1, the `gendisk` utility places the kit on the `c` partition in the `/` (root) directory of the disk media. The product description is `Orpheus_Document_Builder` and the output directory where you created the kit is `/mykit/output`. The kit consists of two subsets: `OATODB100` and `OATODBTEMPS100`.

4.2.1.2 Tape Media Descriptor

The `kitcap` record for tape media contains the following elements:

- The kit name, consisting of two parts:
 - Product code, consisting of the product code and version number specified in the `CODE` and `VERS` fields of the kit's key file. Refer to Section 3.3 for information about the key file.
 - The media code, either `TK` for TK50 tapes or `MT` for 9-track magnetic tape. This element is followed by a pipe character (`|`).
- Product description. This entry is taken from the `NAME` field of the key file.
- Name of the output directory where you created the kit, where the `gentapes` utility can find the subsets.

Since the `gentapes` utility can take subsets from multiple products and merge them on tape as a combined product, you can specify multiple directories where the `gentapes` utility can find the subsets. There must be one directory entry for each `kitcap` descriptor.

- Three empty `SPACE` files to ensure compatibility with operating system kits. To create the `SPACE` file in the output area of the kit directory structure, use comments similar to the following:

```
# cd /mykit/output
# touch space
# tar -cf SPACE space
```

- The `INSTCTRL` image in the output directory containing `setld` control information.
- The names of the subsets that make up the kit. Each subset listed must be stored in one of the specified directories.
- Optional volume identifiers `%%N`, followed by the names of the subsets to be placed on that volume. You can use multiple tapes.

Refer to the `kitcap(4)` reference page for more detailed information about the tape media record format.

Example 4–2 shows the record to be added to the `/etc/kitcap` file to produce the ODB kit on TK50 tapes:

Example 4–2: Sample Tape Media Descriptor for User Product

```
OAT100TK|Orpheus Document Builder: \
/mykit/output:SPACE:SPACE:SPACE: \
INSTCTRL:OATODB100:OATODBTMP100
```

The product name, `OAT100`, is the same name that appears in the key file. The product description, `Orpheus Document Builder`, also appears in the key file. The name of the output directory is `/mykit/output`, and three `SPACE` files are included for compatibility with operating system kits. The last line of the record contains the `INSTCTRL` image in the output directory and the names of the subsets that make up the kit: `OATODB100` and `OATODBTMP100`.

4.2.2 Building a User Product Kit on Disk Media

After the product subsets are located in the output area of the kit directory structure, use the `gendisk` utility to produce the kit on a disk.

Note

The `gendisk` utility supports diskettes but does not let you create a chained diskette kit. A kit written to diskette must fit on a single diskette or be packaged as a set of kits on separate diskettes.

Use the following syntax for the `gendisk` command:

gendisk [-d] [-i] [-k *filename*] [-w] [-v] [*hostname*:] *prodID devname*

Note

If you do not use either the `-w` or `-v` options, the `gendisk` utility writes and then verifies the product media.

`-d`

Creates a distribution disk in direct CD format. This means that the distribution disk contains uncompressed file systems that are laid out just as the software is installed on the system.

Note

We recommend that you do not use the `-d` option when you use the `gendisk` utility to produce user product kits.

`-i`

Creates a distribution disk in ISO 9660 format. This means that the distribution disk contains an ISO 9660-compliant CD-ROM file system (CDFS).

`-k filename`

Uses an alternate kit descriptor database, *filename*, on the local system. You may use either a full absolute pathname or a relative pathname from the directory where you run the `gendisk` utility. The file does not have to be named `kitcap`.

`-w`

Writes the product media without verification, if used without the `-v` option. If used with the `-v` option, the `gendisk` utility writes and then verifies the product media.

`-v`

Verifies the product media without writing it first, if used without the `-w` option. This assumes that you already have written kit files to the distribution media. If used with the `-w` option, the `gendisk` utility writes and then verifies the product media.

hostname:

The optional *hostname*: operand is the name of a remote machine that contains the kit descriptor database. The `gendisk` utility searches the kit descriptor database on the remote machine for the kit identifier (*prodIDHD*) and uses it to create the distribution media. The colon (:) is a required delimiter for TCP/IP networks, and space is permitted between the colon and the *prodID*. For example, if the product code is OAT100 and you are using the kit descriptor database on node `mynode`, use `mynode:OAT100` for this option.

prodID

The mandatory *prodID* operand is a kit identifier consisting of the product code and version number specified in the `CODE` and `VERS` fields of the kit's key file. Refer to Section 3.3 for information about the key file.

devname

The mandatory *devname* operand specifies the device special file name for a raw or character disk device such as `/dev/rdisk/dsk1`. The `gendisk` utility uses the disk partition specified in the kit descriptor and ignores any partition specified on the command line.

The command shown in Example 4–3 creates a `tar` format user product kit for OAT100 on `dsk0`:

Example 4–3: Sample `gendisk` Command for User Product

```
# gendisk OAT100 /dev/rdisk/dsk0
```

Refer to the `gendisk(1)` reference page for more information about this utility.

4.2.3 Building a User Product Kit on Magnetic Tape

After the product subsets are located in the output area of the kit directory structure, use the `gentapes` utility to build the kit on magnetic tape.

Use the following syntax for the `gentapes` command:

```
/usr/bin/gentapes [-w | -v] [hostname:] prodID devname
```

`-w`

Writes the product media without verification. Do not use the `-w` option with the `-v` option.

-v

Verifies the product media without writing it first. Do not use the -v option with the -w option.

hostname:

The optional *hostname:* argument is the name of a remote network machine that contains the kit descriptor database. The `gentapes` utility searches the kit descriptor database on the remote machine for the kit identifier (*prodID*[TK|MT]) and uses it to create the media. The colon (:) is a required delimiter for TCP/IP networks, and space is permitted between the colon and the *prodID*. For example, if the product code is OAT100, and the kitcap file to be used is on node *mynode*, use `mynode:OAT100` for this option.

prodID

The mandatory *prodID* operand is a kit identifier consisting of the product code and version number specified in the `CODE` and `VERS` fields of the kit's key file. Refer to Section 3.3 for information about the key file.

devname

The mandatory *devname* operand specifies the device special file name for a no-rewind tape device such as `/dev/ntape/tape01`. The `gentapes` utility uses the default tape density for the device and ignores any suffix specified on the command line.

Note

If you do not use either the -w or -v option, the `gentapes` utility writes the tape, rewinds it, and then verifies the files in the kit descriptor.

The command shown in Example 4-4 creates a `tar` format user product kit for OAT100 on the magnetic tape in `/dev/ntape/dat`:

Example 4-4: Sample `gentapes` Command for User Product

```
# gentapes OAT100 /dev/ntape/dat
```

Refer to the `gentapes(1)` reference page for more information about this utility.

4.3 Testing the Distribution Media

Before shipping a user product kit to customers, you should test the kit with the same procedures that your customers will use on configurations that resemble your customers' systems.

Use the `setld` utility to test a user product kit as described in the following procedure for the OAT100 kit:

1. Log in to the system as `root` or use the `su` command to gain superuser privileges.

2. Place the CD-ROM in the drive.

3. Create a directory to be the media mount point, such as `/cdrom`:

```
# mkdir /cdrom
```

4. Mount the CD-ROM on `/cdrom`. For example, if the CD-ROM device is located on the `c` partition of `cdrom0`, enter the following command:

```
# mount -r /dev/disk/cdrom0c /cdrom
```

After mounting the CD-ROM, you can change to the `/cdrom` directory and view the directories on the CD-ROM.

5. Install the user product subsets:

```
# setld -l /cdrom/OAT100/kit
```

```
*** Enter subset selections ***
```

```
The following subsets are mandatory and will be installed automatically unless you choose to exit without installing any subsets:
```

```
    * Document Builder Tools
```

```
The subsets listed below are optional:
```

```
- Other:
```

```
    1) Document Builder Templates
```

```
Or you may choose one of the following options:
```

```
    2) ALL mandatory and all optional subsets
```

```
    3) MANDATORY subsets only
```

```
    4) CANCEL selections and redisplay menus
```

```
    5) EXIT without installing any subsets
```

```
Estimated free disk space (MB) in root:54.5 usr:347.0
```

```
Enter your choices or press RETURN to redisplay menus.
```

```
Choices (for example, 1 2 4-6): 2
```

```
You are installing the following mandatory subsets:
```

```
    Document Builder Kernel Support
    Document Builder Tools
```

```
You are installing the following optional subsets:
```

```

- Other:
  Document Builder Templates

Estimated free disk space (MB) in root:54.5 usr:347.0

Is this correct? (y/n): y

Checking file system space required to install selected subsets:

File system space checked OK.

3 subset(s) will be installed.

Loading subset 1 of 2 ...

Document Builder Tools
  Copying from /mykit/output (disk)
  Verifying

Loading subset 2 of 2 ...

Document Builder Templates
  Copying from /mykit/output (disk)
  Verifying

2 of 2 subset(s) installed successfully.

Configuring "Document Builder Tools" (OATODB100)

The installation of the Document Builder Tools (OATODB100)
software subset is complete.

Please check the /opt/OAT100/odb.conf file before
using the Document Builder Tools product.

Configuring "Document Builder Templates" (OATODBTEMPS100)

#

```

The `setld` utility displays prompts and messages to guide you through the process of selecting the subsets you want to install. As each subset is loaded, the `setld` utility calls the subset control program as needed.

6. After the installation finishes, unmount the CD-ROM:

```
# umount /cdrom
```

7. Verify that the installed product functions correctly.

Refer to the *Installation Guide* and the `setld(8)` reference page for more information about using the `setld` utility to install layered products.

5

Producing Kernel Product Kits

This chapter tells you how to produce a kernel product kit, what additional files are required, and how to test the kit installation.

Note

The information in this chapter describes how to produce kernel product kits. If you want to create a user product kit, go to Chapter 4.

Follow these steps to create and test a kernel product kit:

1. Read Chapter 1 for an overview of product kits.
2. Design the kit directory structure as described in Chapter 2.
3. Create subsets as described in Chapter 3. Kernel product kits require a subset control program (Section 3.4).
4. Create additional installation files as described in Section 5.2.
5. Rerun the `newinv` utility as described in Section 3.7 to update the master inventory file with the additional installation files.
6. Rerun the `kits` utility as described in Section 3.5 to produce updated subsets and control files.
7. Create the kit distribution media as described in Section 5.3.
8. Test the distribution media as described in Section 5.4.

5.1 Overview

A kernel product is a layered product that contains kernel support. Users do not run kernel products directly; the operating system and utilities access kernel products to perform their work. For example, a device driver is one common type of kernel product. A user runs an application or utility, which generates system requests to perform operations such as opening a file or writing data to a disk. The system determines which device driver should service this request and then calls the required driver interface.

The kernel modules and the kit support files are distributed as part of the kernel product kit, and can be installed either directly from the distribution media or loaded onto a Remote Installation Services (RIS) area for installation by RIS clients over a local area network (LAN).

5.2 Creating Additional Installation Files

The files needed to build a kernel product kit depend on whether the kit will be configured statically or dynamically on the customer's system. For example:

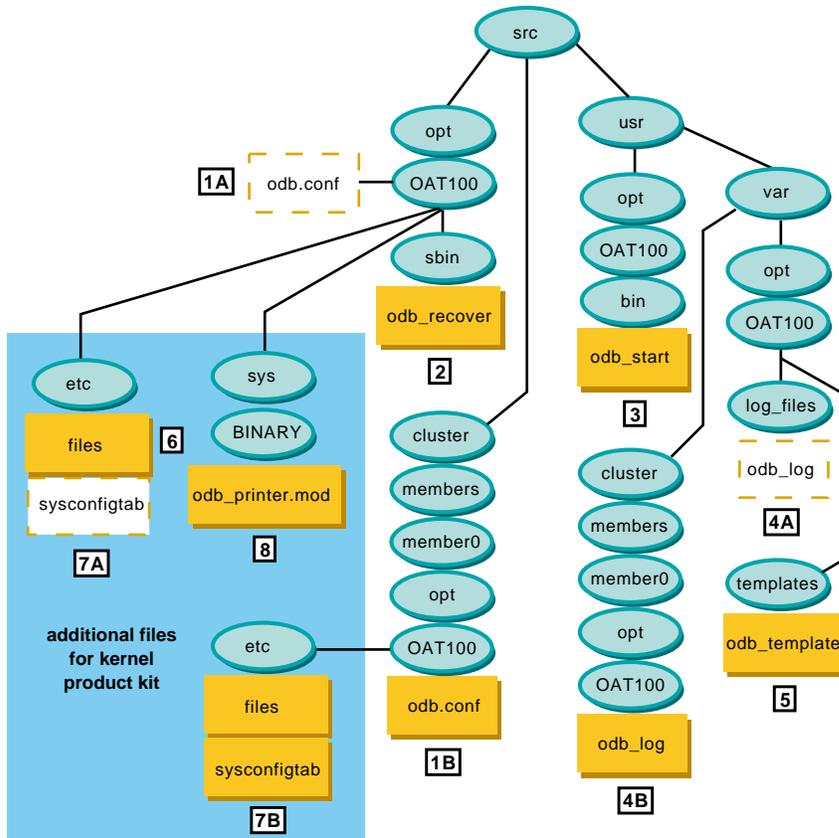
- A statically configured product is linked statically into the kernel at build (or bootlink) time and configured at boot time. A static product can be built from source files, binary objects, modules, or all three.
- A dynamically configured product is loaded into a running kernel after it has been booted. It is not part of the permanent kernel and is configured when it is loaded. It must be reloaded after each boot of the system. A dynamic product can be built from source files, binary objects, modules, or all three.

Note:

A module that can be loaded dynamically also can be linked statically. The only difference is the call to configure the product. For more information on static and dynamic drivers, refer to *Writing Device Drivers*.

Figure 5-1 shows the files that make up the ODB product kit source directory. Additional kernel product files are highlighted.

Figure 5–1: Kernel Product Source Directory



ZK-1555U-AI

1 /oddb.conf — the ODB product configuration file

If your product kit may run on a cluster, each cluster member must have its own configuration file. To accommodate this requirement, create a context-dependent symbolic link (CDSL) targeted to the member-specific file.

- A. The context-dependent symbolic link (CDSL) for the `oddb.conf` file is linked to the member-specific `cluster/members/{memb}/opt/OAT100/oddb.conf` file. This CDSL is installed in the root file system and placed in the `opt/OAT100` source directory.
- B. The member-specific file `oddb.conf` file can differ on each cluster member. This file is installed in the cluster member's root file system and is placed in the `cluster/members/member0/opt/OAT100` source directory.

- 2 /sbin/odb_recover — a utility to recover corrupt ODB documents when the system boots. The odb_recover script executes when the system boots and the /usr file system may not be mounted.

This file is installed in the root file system and is placed in the opt/OAT100/sbin source directory.

- 3 /usr/bin/odb_start — the ODB product startup script. The odb_start script is a user command.

This file is installed in the /usr file system and is placed in the usr/opt/OAT100/bin source directory.

- 4 /usr/var/log_files/odb_log — the ODB product log file.

If the product kit can be installed on a cluster, each cluster member must have its own copy of the log file. To accommodate this requirement, create a context-dependent symbolic link (CDSL) targeted to a member-specific file.

- A. The context-dependent symbolic link (CDSL) for the odb_log file is linked to the member-specific usr/var/cluster/members/{memb}/opt/OAT100/log_files/odb_log file. This CDSL is installed in the /usr/var file system and placed in the usr/var/opt/OAT100/log_files source directory.
- B. The member-specific file odb_log file can differ on each cluster member. This file is installed in the cluster member's /usr/var file system and is placed in the /usr/var/cluster/members/member0/opt/OAT100/log_files source directory.

Refer to Section 2.3.2 for information about CDSLs.

- 5 /usr/var/templates/odb_template — a document template that can be modified by a user.

This file is installed in the /var file system and is placed in the usr/var/opt/OAT100/templates source directory.

- 6 /etc/files — the files file fragment contains information about the location of the source code and modules associated with the driver, tags indicating when the driver is loaded into the kernel, and whether the source or binary form of the driver is supplied to the customer.

- 7 /etc/sysconfigtab — the sysconfigtab file fragment contains device special file information, bus option data, and information on contiguous memory usage for statically and dynamically configured drivers. The sysconfigtab file fragment gets appended to the /etc/sysconfigtab database when the kernel product kit is installed. For more information, refer to the sysconfigtab(4) reference page.

If the product kit can be installed on a cluster, each cluster member must have its own copy of the `sysconfigtab` file fragment. To accommodate this requirement, create a context-dependent symbolic link (CDSL) targeted to the member-specific file.

- A. The context-dependent symbolic link (CDSL) for the `sysconfigtab` file fragment is linked to the member-specific `cluster/members/{memb}/opt/OAT100/etc/sysconfigtab` file. This CDSL is installed in the root file system and placed in the `opt/OAT100/etc` source directory.
- B. The member-specific `sysconfigtab` file can differ on each cluster member. This file is installed in the cluster member's root file system and is placed in the `cluster/members/member0/opt/OAT100/etc` source directory.

Refer to Section 2.3.2 for information about CDSLs.

8 `/sys/BINARY/odb_printer.mod` — the object module file containing the single binary module for the ODB kernel product.

A kernel product kit requires you to include the following files on the distribution media to make the kernel product accessible during installation:

- The `files` file fragment (Section 5.2.1)
- The `sysconfigtab` file fragment (Section 5.2.2)
- The `driver.mod` object module file (Section 5.2.3)
- The `*.c` (source) and `*.h` (header) files (Section 5.2.4)
- The `device.mth` method files (Section 5.2.5)

5.2.1 The files File Fragment

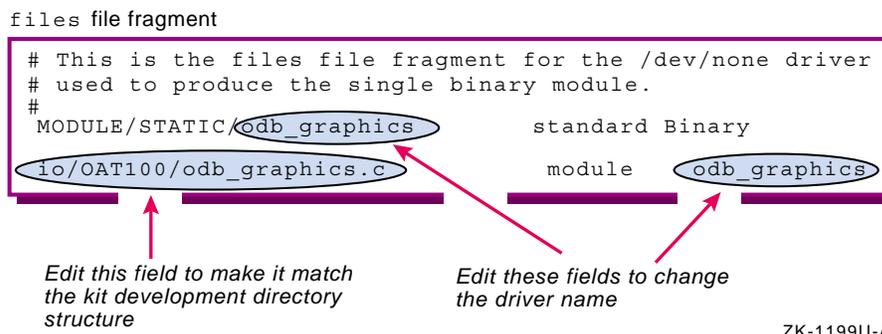
The `files` file fragment contains information about the location of the source code and modules associated with the driver, tags indicating when the driver is loaded into the kernel, and whether the source or binary form of the driver is supplied to the customer. You need to edit this file if the kit development directory structure differs from the driver development directory structure or if you must change the driver name for any reason.

Caution

The `files` file fragment must be in the same directory as the kernel modules or the `kreg` and `doconfig` utilities will not work properly.

Figure 5–2 shows which fields within the `files` file fragment need to change.

Figure 5–2: Editing the files File Fragment



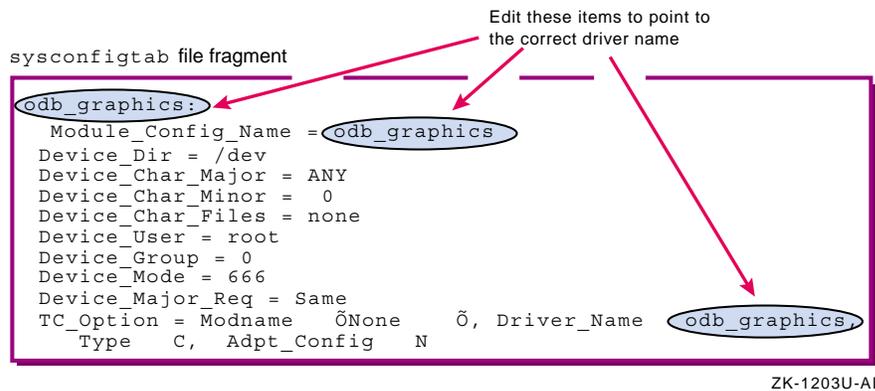
ZK-1199U-AI

5.2.2 The `sysconfigtab` File Fragment

The `sysconfigtab` file fragment contains device special file information, bus option data information, and information on contiguous memory usage for statically and dynamically configured drivers. When the user installs a kernel product kit, the driver's `sysconfigtab` file fragment gets appended to the `/etc/sysconfigtab` database. You should place this file fragment in a product directory, such as `/opt/OAT100/etc`.

You do not need to change the `sysconfigtab` file fragment unless you change the driver (subsystem) name. The driver name appears in three places within the file, as shown in Figure 5–3. In the example, the driver runs on a `TURBOchannel` bus (indicated by the `TC_Option` entry), but a similar set of bus options would be specified for other bus types.

Figure 5–3: Editing the sysconfigtab File Fragment



This file is described in the `sysconfigtab(4)` reference page.

5.2.3 The Object Module File

The `driver.mod` object module file contains the single binary module for both statically and dynamically configured drivers. Include this file in a product directory, such as `/opt/OAT100/sys/BINARY`, in the root (`/`) file system.

Caution

You cannot use RIS to link compressed modules. Use the `file` command to determine if a file is compressed. You see output similar to the following:

```
# file odb_graphics.mod
odb_graphics.mod:
  alpha compressed COFF executable or object module not stripped
```

5.2.4 The Source and Header Files

The `*.c` (source) and `*.h` (header) files contain the source code for the device driver. Include these files in a product directory, such as `/usr/opt/OAT100/src`, when the driver is statically configured and distributed in source form.

5.2.5 The Method Files

The `device.mth` method files contain driver methods that are called during automatic configuration to create device special files for dynamically configured drivers. The subset control program creates links to these device special files in the customer's `subsys` directory when the driver is

installed. The driver method files are on the distribution media and are not installed onto the customer's system. The device driver developer can tell you which method files the subset control program should link to, typically `/subsys/device.mth`. Link the method in a device driver kernel kit only if the driver needs to have device special files created for its devices.

5.3 Producing Distribution Media

After you have tested the subsets, you can produce the distribution media. Distribution media production consists of the following tasks:

1. Rerunning the `newinv` utility to update the master inventory file with the additional installation files. (Section 3.7)
2. Rerunning the `kits` utility to produce updated subsets and control files. (Section 3.5)
3. Editing the `/etc/kitcap` file. (Section 5.3.1)
4. Building the kernel product kit on the distribution media.
 - Use the `gendisk` utility to build a kit on disk media. (Section 5.3.2)
 - Use the `gentapes` utility to build a `tar` format kit on magnetic tape. (Section 5.3.3)

You can produce the kernel product kit in either `tar` format or direct CD-ROM (DCD) format. Installation time for `tar` format kits is faster than for DCD format kits.

- If your product kit does not access kernel modules during boot, use the `tar` format to compress your kit and save space on the media.

In `tar` format, the product files in each subset are written to the distribution media as a single file. During installation, the `setld` utility uncompresses the files and moves them onto the target system, preserving the files' original directory structure. Kits distributed in `tar` format install more quickly and consume less space on the distribution media.

- If your product kit accesses kernel modules during the boot process, you must use the DCD format and cannot produce the kit on magnetic tape media.

In DCD format, the files are written to the distribution media as a UNIX file system where the product files are organized into a directory structure that mirrors the target system. Subsets distributed in DCD format cannot be compressed.

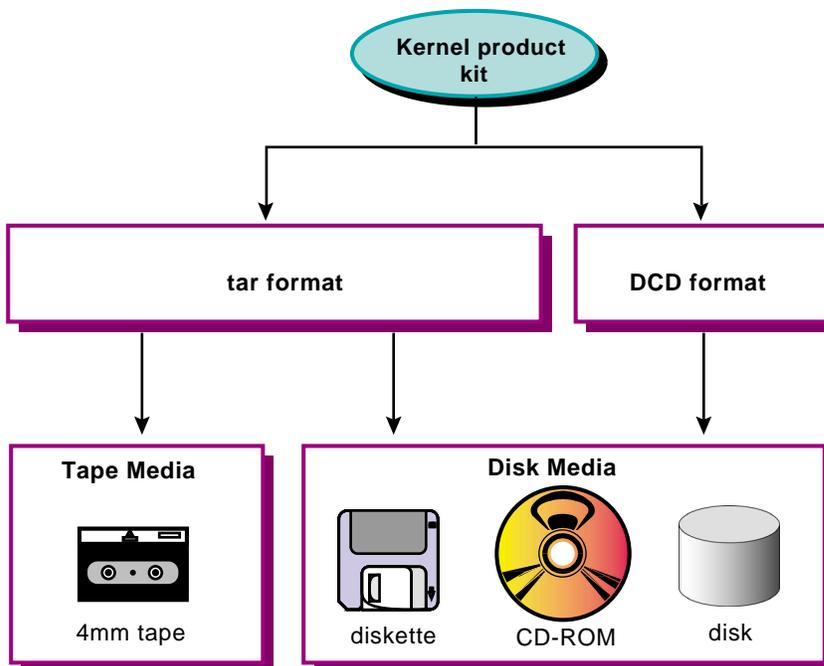
You can distribute user product kits on diskette, CD-ROM, or magnetic tape, as follows:

- Diskettes are good for testing purposes or for small products. The product must fit on a single diskette; it cannot span multiple diskettes. Use the `gendisk` utility to produce kits for diskette media.
- CD-ROMs can support large kits or multiple kits on a single media. The kit is first produced on the hard disk, then written onto the CD-ROM. Use the `gendisk` utility to produce the master kit on hard disk. Follow the CD-ROM manufacturer's instructions for writing the kit onto the CD-ROM media.
- Magnetic tape

Magnetic tape can be used for kernel product distribution if you do not need to access the kernel modules during the boot process. You must use `tar` format; magnetic tape does not support DCD format. Use the `gentapes` utility to produce kits for magnetic tape media.

Figure 5-4 shows the file formats and distribution media available for kernel product kits.

Figure 5-4: Kernel Product Kit File Formats



ZK-1552U-A1

5.3.1 Editing the `/etc/kitcap` File

The `gendisk` and `gentapes` utilities refer to the `/etc/kitcap` file, a database containing information about the kits to be built on the system. Each record contains a product code and the names of the directories, files, and subsets that make up the product kit. Before you can build your kit, you must add a media descriptor record to the `/etc/kitcap` database.

Note

If you use the `gendisk` utility to produce your kit on disk distribution media, you can specify an alternate kit descriptor database. Refer to the `gendisk(1)` reference page for additional information.

Use the following conventions when you add a record to the `/etc/kitcap` file:

- Separate the first field from the rest of the record by a colon (:) for disk media descriptors and by a pipe character (|) for tape media descriptors.
- Separate all other fields with colons (:).
- Indicate continuation with a backslash (\) at the end of the line.
- Lines starting with a pound sign (#) are comments and are ignored.
- Comments within the record start with pound sign (#) and end with a colon (:). Use this feature sparingly.

The contents of a `kitcap` record differ depending on whether you are producing disk or tape media. You must add one record for each media type on which you plan to distribute your kit.

The contents of the record also depend on the product type you are delivering. For example, the `kitcap` record for a kernel product may require the `kk=true` parameter and the `rootdd=` option. Refer to the `kitcap(4)` reference page for more information about the contents of the `/etc/kitcap` file.

5.3.1.1 Disk Media Descriptor

Create a disk media `kitcap` record when you produce kits for distribution on diskette or CD-ROM. The `kitcap` record for disk media contains the following elements:

- The kit name, consisting of two parts:
 - The product code, consisting of the product code and version number specified in the `CODE` and `VERS` fields of the kit's key file (`prodcode.k`).
 - The media code `HD` to indicate disk media. This element is followed by a colon (`:`).
- The partition on the disk media where the product should be placed. The partition is a letter between `a` and `h`. Partition `c` is used most often, as it spans the entire disk.
- The destination directory for the subsets on the disk media. This allows a hierarchical structure so you can put multiple products on one disk, or put parts of one product on different areas of the same disk. You can use multiple destination directories in a `kitcap` record.

The destination directory field also may include these parameters:

- `kk=true`
Indicates that the kit is needed during the boot process. When the `gendisk` utility finds this option, it automatically generates a `kitname.kk` file.
- `rootdd=dirname`
Specifies kit file placement on the distribution media, relative to the kit-specific directory such as `/OAT100/kit`. For example, `rootdd=..` would place the kit's root under the `/OAT100` directory.
- The product description. This entry is taken from key file `NAME` field. Replace any spaces with an underscore (`_`) character, for example: `Product Description` becomes `Product_Description`.
- The name of the output directory where you created the kit, where the `gendisk` utility can find the product subsets.
- The `instctrl` directory, relative to the output directory specification.
- The names of the subsets that make up the kit.

Refer to the `kitcap(4)` reference page for more information about the disk media record format.

Refer to Section 3.3 for information about the key file.

Example 5–1 shows the record to be added to the `/etc/kitcap` file to produce the ODB kit on disk media:

Example 5–1: Sample Disk Media Descriptor for Kernel Product

```
OAT100HD:c:\
  dd=/OAT100:Orpheus_Document_Builder:/mykit/output:\
  instctrl:OATODB100:OATODBTMP100:OATODBKERNEL100
```

Based on the information shown in Example 5–1, the `gendisk` utility places the kit on the `c` partition in the `/` (root) directory of the disk media. The product description is `Orpheus_Document_Builder` and the kit output directory is named `/mykit/output`. The kit consists of three subsets: `OATODB100`, `OATODBTMP100`, and `OATODBKERNEL100`.

5.3.1.2 Tape Media Descriptor

The `kitcap` record for tape media contains the following elements:

- The kit name, consisting of two parts:
 - Product code, consisting of the product code and version number specified in the `CODE` and `VERS` fields of the kit's key file (`prodcode.k`).
 - The media code, either `TK` for TK50 tapes or `MT` for 9-track magnetic tape. This element is followed by a pipe character (`|`).
- Product description. This entry is taken from the `NAME` field of the key file.
- Name of the output directory where you created the kit, where the `gentapes` utility can find the subsets.

Since the `gentapes` utility can take subsets from multiple products and merge them on tape as a combined product, you can specify multiple directories where the `gentapes` utility can find the subsets. There must be one directory entry for each `kitcap` descriptor.

- Three empty `SPACE` files to ensure compatibility with operating system kits. To create the `SPACE` file in the output area of the kit directory structure, issue the following commands:

```
# cd /mykit/output
# touch space
# tar -cf SPACE space
```

- The `INSTCTRL` image in the output directory containing `setld` control information.

- The names of the subsets that make up the kit. Each subset listed must be stored in one of the specified directories.
- Optional volume identifiers %%N, followed by the names of the subsets to be placed on that volume. You can use multiple tapes.

Refer to the `kitcap(4)` reference page for more detailed information about the tape media record format.

Example 5–2 shows the record to be added to the `/etc/kitcap` file to produce the ODB kit on TK50 tapes:

Example 5–2: Sample Tape Media Descriptor

```
OAT100TK|Orpheus Document Builder:\
/mykit/output:SPACE:SPACE:SPACE:\
INSTCTRL:OATODB100:OATODBTMP100:OATODBKERNEL100
```

The product name, `OAT100`, is the same name that appears in the key file. The product description, `Orpheus Document Builder` also appears in the key file. The name of the output directory is specified as `/mykit/output`, and three `SPACE` files are included for compatibility with operating system kits. The last line of the record contains the `INSTCTRL` image in the output directory and the names of the subsets that make up the kit: `OATODB100`, `OATODBTMP100`, and `OATODBKERNEL100`.

5.3.2 Building a Kernel Product Kit on Disk Media

When the product subsets are located in the output area of the kit directory structure, use the `gendisk` utility to create the kit on a disk.

Note

The `gendisk` utility supports diskettes but does not allow you to create a chained diskette kit. A kit written to diskette must fit on a single diskette or be packaged as a set of kits on separate diskettes.

Use the following syntax for the `gendisk` command:

```
gendisk [-d] [-i] [-k filename] [-w] [-v] [hostname:] prodID devname
-d
```

Creates a distribution disk in direct CD (DCD) format. This means that the distribution disk contains uncompressed file systems that are arranged in the same way as the software will be installed on the system.

-i

Creates a distribution disk in ISO 9660 format. This means that the distribution disk contains an ISO 9660-compliant CD-ROM file system (CDFS).

-k *filename*

Uses an alternate kit descriptor database, *filename*, on the local system. You may use either a full absolute pathname or a relative pathname from the directory where you run the `gendisk` utility. You do not have to name the file `kitcap`.

-w

If used without the `-v` option, writes the product media without verification. If used with the `-w` option, the `gendisk` utility writes and then verifies the product media.

-v

If used without the `-w` option, verifies the product media without writing it first. This assumes that you already have written kit files to the distribution media. If used with the `-w` option, the `gendisk` utility writes and then verifies the product media.

hostname:

The optional *hostname:* operand is the name of a remote machine that contains the `kitcap` file. The utility searches the `/etc/kitcap` file on the remote machine for the *prodID* and uses it for creating the media. The colon (`:`) is a required delimiter for TCP/IP networks, and space is permitted between the colon and the *prodID*. For example, if the product code is `OAT100` and you are using the kit descriptor database on node `mynode`, use `mynode:OAT100` for this option.

prodID

The mandatory *prodID* operand is a kit identifier consisting of the product code and version number specified in the `CODE` and `VERS` fields of the kit's key file. Refer to Section 3.3 for information about the key file.

devname

The mandatory *devname* operand specifies the device special file name for a raw or character disk device such as `/dev/rdisk/dsk1`. The `gendisk` utility uses the disk partition specified in the kit descriptor and ignores any partition specified on the command line.

Note

If you do not use either the `-w` or `-v` options, the `gendisk` utility writes and then verifies the product media.

The command illustrated in Example 5–3 creates a `tar` format kernel product kit for `OAT100` on `disk0`:

Example 5–3: Sample `gendisk` Command

```
# gendisk OAT100 /dev/rdisk/dsk0
```

Refer to the `gendisk(1)` reference page for more information about this utility.

5.3.3 Building a Kernel Product Kit on Magnetic Tape

When the product subsets are located in the output area of the kit directory structure, use the `gentapes` utility to create the kit on magnetic tape. Use the following syntax for the `gentapes` command:

gentapes [`-w` | `-v`] [*hostname*:] *prodID devname*

`-w`

Writes the product media without verification. Do not use the `-w` option with the `-v` option.

`-v`

Verifies the product media without writing it first. Do not use the `-v` option with the `-w` option.

hostname:

The optional *hostname*: argument is the name of a remote network machine that contains the kit descriptor database. The `gentapes` utility searches the kit descriptor database on the remote machine for the kit identifier (*prodID*[`TK`|`MT`]) and uses it to create the media. The colon (`:`) is a required delimiter for TCP/IP networks, and space is permitted between the colon and the *prodID*. For example, if the product code is `OAT100` and you are using the kit descriptor database on node `mynode`, use `mynode:OAT100` for this option.

prodID

The mandatory *prodID* operand is a kit identifier consisting of the product code and version number specified in the `CODE` and `VERS`

fields of the kit's key file. Refer to Section 3.3 for information about the key file.

devname

The mandatory *devname* operand specifies the device special file name for a no-rewind tape device such as `/dev/ntape/tape01`. The `gentapes` utility uses the default tape density for the device and ignores any suffix specified on the command line.

Note

If you do not use either the `-w` or `-v` option, the `gentapes` utility writes the tape, rewinds it, and then verifies the files in the kit descriptor.

The command illustrated in Example 5–4 creates a `tar` format user product kit for `OAT100` on the magnetic tape in `/dev/ntape/dat`:

Example 5–4: Sample `gentapes` Command

```
# gentapes OAT100 /dev/ntape/dat
```

Refer to the `gentapes(1)` reference page for more information about this utility.

5.4 Testing the Distribution Media

Before shipping a kernel product kit to customers, you should test the kit with the same procedures that your customers will use on configurations that resemble your customers' systems.

Run the following tests to test a kernel product kit:

1. Use the `setld` utility to verify that the subsets have been built correctly and that the files get installed into the correct locations on the target system. (Section 5.4.1)
2. Use the `ris` utility to add a kernel product kit into a RIS area and verify that the correct files are present on the kit. Then, register the client system to the RIS area and start a Full Installation on the client system. (Section 5.4.2)

5.4.1 Testing a Kernel Product Kit with the setld Utility

Use the `setld` utility to test a kernel product kit as described in the following procedure for the OAT100 kit:

1. Log in to the system as `root` or use the `su` command to gain superuser privileges.
2. Place the CD-ROM in the drive.
3. Create a directory to be the media mount point, such as `/cdrom`:

```
# mkdir /cdrom
```
4. Mount the CD-ROM on `/cdrom`. For example, if the CD-ROM device is located on the `c` partition of `cdrom0`, enter the following command:

```
# mount -r /dev/disk/cdrom0c /cdrom
```
5. Use the `setld` utility to install the kernel product subsets:

```
# setld -l /cdrom/OAT100/kit
```

Your session looks similar to the following:

```
*** Enter subset selections ***
```

```
The following subsets are mandatory and will be installed automatically unless you choose to exit without installing any subsets:
```

```
    * Document Builder Kernel Support
    * Document Builder Tools
```

```
The subsets listed below are optional:
```

```
- Other:
  1) Document Builder Templates
```

```
Or you may choose one of the following options:
```

```
  2) ALL mandatory and all optional subsets
  3) MANDATORY subsets only
  4) CANCEL selections and redisplay menus
  5) EXIT without installing any subsets
```

```
Estimated free disk space (MB) in root:54.5 usr:347.0
```

```
Enter your choices or press RETURN to redisplay menus.
```

```
Choices (for example, 1 2 4-6): 2
```

```
You are installing the following mandatory subsets:
```

```
    Document Builder Kernel Support
    Document Builder Tools
```

```
You are installing the following optional subsets:
```

```
- Other:
  Document Builder Templates
```

```
Estimated free disk space (MB) in root:54.5 usr:347.0
```

```

Is this correct? (y/n): y
Checking file system space required to install selected subsets:

File system space checked OK.

3 subset(s) will be installed.

Loading subset 1 of 3 ...

Document Builder Tools
  Copying from /cdrom/OAT100/kit (disk)
  Verifying

Loading subset 2 of 3 ...

Document Builder Templates
  Copying from /cdrom/OAT100/kit (disk)
  Verifying

Loading subset 3 of 3 ...

Document Builder Kernel Support
  Copying from /cdrom/OAT100/kit (disk)
  Verifying

3 of 3 subset(s) installed successfully.

Configuring "Document Builder Tools" (OATODB100)

The installation of the Document Builder Tools (OATODB100)
software subset is complete.

Please read the /opt/OAT100/README.odb file before
using the Document Builder Tools product.

Configuring "Document Builder Templates" (OATODBTEMPS100)

Configuring "Document Builder Kernel Support" (OATODBKERNEL100)

*** Document Builder Kernel Support Product Installation Menu ***

1. Statically configure the graphics support
2. Dynamically configure the graphics support

Type the number of your choice []: 1

*** KERNEL CONFIGURATION AND BUILD PROCEDURE ***

Saving /sys/conf/TEST01 as /sys/conf/TEST01.bck

Do you want to edit the configuration file? (y/n) [n]: n

*** PERFORMING KERNEL BUILD ***
Working...Fri May 19 14:49:25 EDT 2000

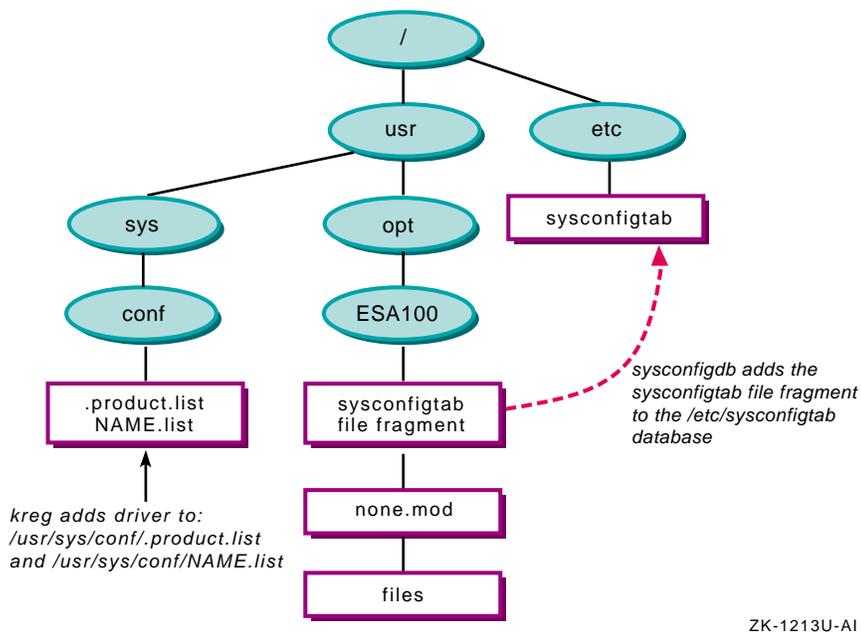
The new kernel is /sys/TEST01/vmunix

The /sys/TEST01/vmunix kernel has been
moved to /vmunix and the changes will take effect
the next time the system is rebooted.
#

```

The `setld` utility displays prompts and messages to guide you through the process of selecting the subsets you want to install. As each subset is loaded, the `setld` utility calls the subset control program as needed, including static or dynamic driver configuration. Figure 5-5 shows the steps the subset control program takes to statically configure the driver.

Figure 5-5: Static Configuration of a Driver



ZK-1213U-AI

- When the installation is complete, unmount the CD-ROM:

```
# umount /cdrom
```

- If the product was configured statically, restart the system with the new kernel:

```
# /usr/sbin/shutdown -r now
```

When the system restarts, the device driver is available on the system.

- Verify that the installed product functions correctly.
- Delete the ODB subsets with the `setld -d` command.

```
# setld -d OATODB100 OATODBTMP100 OATODBKERNEL100
```

You see output similar to the following:

```
Deleting "Document Builder Templates" (OATODBTMP100).
```

```
*** KERNEL CONFIGURATION AND BUILD PROCEDURE ***
```

```
Saving /sys/conf/TEST01 as /sys/conf/TEST01.bck
```

```

Do you want to edit the configuration file? (y/n) [n]: n

*** PERFORMING KERNEL BUILD ***
      Working...Fri May 19 14:55:31 EDT 2000

The new kernel is /sys/TEST01/vmunix

The /sys/TEST01/vmunix kernel has been
moved to /vmunix and the changes will take effect
the next time the system is rebooted.

Deleting "Document Builder Kernel Support" (OATODBKERNEL100).

Deleting "Document Builder Tools" (OATODB100).
#

```

10. Use the `shutdown` command to restart the system, ensuring that it reboots with the new kernel after the product is removed:

```
# /usr/sbin/shutdown -r now
```

When the system restarts, the product should not be available on the system.

Refer to the *Installation Guide* and the `setld(8)` reference page for more information about using the `setld` utility to install layered products.

5.4.2 Testing a Kernel Product Kit in a RIS Area

Use the `ris` utility to test a kernel product kit on a RIS server as described in the following procedure for the OAT100 kit. Refer to the *Sharing Software on a Local Area Network* manual for additional information about Remote Installation Services (RIS).

1. Log in to the RIS server as `root` or use the `su` command to gain superuser privileges.
2. Place the CD-ROM in the drive.
3. Create a directory to be the media mount point, such as `/cdrom`:

```
# mkdir /cdrom
```

4. Mount the CD-ROM on `/cdrom`. For example, if the CD-ROM device were located on the `c` partition of `cdrom0`, enter the following command:

```
# mount -r /dev/disk/cdrom0c /cdrom
```

5. Enter `/usr/sbin/ris` to start the `ris` utility.

You see the RIS Utility Main Menu:

```
*** RIS Utility Main Menu ***
```

Choices without key letters are not available.

```
) ADD a client
) DELETE software products
i) INSTALL software products
) LIST registered clients
) MODIFY a client
) REMOVE a client
) SHOW software products in remote installation environments
x) EXIT
```

Enter your choice:

6. Enter `i` to select Install software products. You see the RIS Software Installation Menu:

```
RIS Software Installation Menu:
```

```
1) Install software into a new area
2) Add software into an existing area
3) Return to previous menu
```

Enter your choice:

7. Depending on your test environment, enter `1` to select Install software into a new area or `2` to Add software into an existing area.
8. Install the software as described in *Sharing Software on a Local Area Network*.

To install the product kit from the RIS server onto the client system, first register the client system with the RIS server and then use the `setld` utility as described in the following procedure:

1. Log in to the RIS server as `root` or use the `su` command to gain superuser privileges.
2. Enter `/usr/sbin/ris` to start the `ris` utility. You see the RIS Utility Main Menu:

```
*** RIS Utility Main Menu ***
```

Choices without key letters are not available.

```
a) ADD a client
d) DELETE software products
i) INSTALL software products
) LIST registered clients
) MODIFY a client
) REMOVE a client
s) SHOW software products in remote installation environments
x) EXIT
```

Enter your choice:

3. Enter **a** to select ADD a client.
4. Enter the client information as described in *Sharing Software on a Local Area Network*.
5. Log in to the RIS client as `root` or use the `su` command to gain superuser privileges.
6. Use the `setld -l` command to load the product subsets from the RIS area. For example, if the RIS server is named `test01`, enter the following command:

```
# setld -l test01:
```

The `setld` utility displays prompts and messages to guide you through the installation process as described in *Sharing Software on a Local Area Network*.

Refer to the *Installation Guide* and the `setld(8)` reference page for more information on using the `setld` utility to install layered products.

A

Creating a Consolidated Firmware CD-ROM

A consolidated firmware CD-ROM lets you upgrade your processor firmware at the same time that you install the operating system. This appendix describes how to create a consolidated firmware CD-ROM in ISO 9660-compliant format (CDFS), using the operating system CD-ROM and the applicable *Alpha Systems Firmware CD-ROM*. The following information is included in this appendix:

- How to prepare for the build (Section A.1.1)
- How to build the consolidated firmware CD-ROM (Section A.1.2)
- Sample sessions for both the build preparation (Section A.2.1) and the actual building (Section A.2.2) of the consolidated firmware CD-ROM

This release includes the documentation and utilities that you need to build a consolidated firmware CD-ROM in ISO 9660-compliant format (CDFS). This documentation includes the `disklabel(8)` and `mkisofs(8)` reference pages.

A.1 Build Instructions

The following sections describe how to prepare and build a consolidated firmware CD-ROM.

Note

The examples in this appendix use the C shell.

A.1.1 Prepare for the Build Session

After you receive a new kit, follow these steps to move the necessary files from the CD-ROM to working directories on the build machine:

1. Log in as `root` or use the `su` command to gain superuser privileges.
2. Use the `disklabel` utility to set up a 635 Mb partition on a spare disk, starting at block 0, with a size of 1300480 512-byte blocks and a file system type of `unused`. Create a mount point for this partition (such as `/cdimage`) to use later.

For example, to set partition `a` of `dsk6` starting at offset 0 with a size of 1300480, and create mount point `/cdimage`:

```
% disklabel -F -r -e /dev/disk/dsk6
write new label [y] y
% mkdir /cdimage
```

3. Mount the operating system CD-ROM to a temporary mount point (such as `/mnt`) and use the `tar` command to copy the contents of the CD-ROM onto a suitably large directory on the system (at least 1.5 Gb). After this is done, unmount the CD-ROM.

Note

This step may take as long as 60 minutes to complete.

For example, using `/spare` as the target directory and `/dev/disk/cdrom4` as the CD-ROM drive:

```
% mount -r /dev/disk/cdrom4a /mnt
% cd /mnt
% tar cf /spare/os_copy.tar .
% cd /
% umount /mnt
```

A.1.2 Build the Consolidated Firmware CD-ROM

After you have completed the steps in Section A.1.1, follow these steps to consolidate the necessary data to a single CD-ROM in ISO 9660-compliant format (CDFS):

1. Log in as `root` or use the `su` command to gain superuser privileges.
2. Use the `newfs` command to initialize a file system on the partition reserved in Step 2 of Section A.1.1 and mount it to the mount point `/cdimage`. If you are prompted for confirmation, enter `y`. Use the `tar` utility to copy the base operating system image created in Step 3 of Section A.1.1 to `/cdimage`.

Note

This step may take as long as 60 minutes to complete.

For example, using `/spare` as the source and `dsk6c` as the target partition:

```
% newfs /dev/disk/dsk6c
% mount /dev/disk/dsk6c /cdimage
% cd /cdimage
% tar xpf /spare/os_copy.tar
% cd /
```

3. Optionally use the following multiple step operation to copy the firmware images to the target directory:

- a. Mount the *Alpha Systems Firmware CD-ROM* to a temporary mount point such as /mnt. For example, using /dev/disk/cdrom4a as the CD-ROM drive:

```
% mount -t cdfs -r /dev/disk/cdrom4a /mnt
```

- b. Copy the System Marketing Model (SMM) table from the *Alpha Systems Firmware CD-ROM* to the target directory. The SSM table maps system models to firmware image files.

```
% cp /mnt/SMMTABLE.TXT\;1 /cdimage/smmtable.txt
```

Caution

The target file name must be in lower case with the ;1 removed from the end.

- c. Look in the SMM table to find the name and locations of the firmware images to be copied by entering the following command:

```
% more /cdimage/smmtable.txt
```

As an example, the entry for the EV5 AlphaServer 1000A platform is similar to the following example (the actual table entry is on one line):

```
27 5 1270,1311,1558,1580-1581\  
[ALPHA1000A]AS1000A_E5_V5_1.EXE;1\  
6 5.1 ! AlphaServer 1000A 5/xxx
```

In this example, the firmware file on the CD-ROM is AS1000A_E5_V5_1.EXE;1.

- d. Create the required firmware directories in the target directory, and copy each of the platform firmware images that you want from the *Alpha Systems Firmware CD-ROM*.

Caution

The target file name must be in lower case with the “;1” removed from the end. Otherwise, the fwupgrade program cannot locate the firmware images. If the source file is AS1000A_E5_V5_1.EXE;1, the target file is as1000a_e5_v5_1.exe.

For example, using the file names on the *Alpha Systems Firmware CD-ROM*:

```
% mkdir /cdimage/alpha800  
% mkdir /cdimage/alpha1000a  
% mkdir /cdimage/as4x00
```

```
% cp /mnt/ALPHA800/AS800_V5_1.EXE\;1 \
/cdimage/alpha800/as800_v5_1.exe
% cp /mnt/ALPHA1000A/AS1000A_E5_V5_1.EXE\;1 \
/cdimage/alpha1000a/as1000a_e5_v5_1.exe
% cp /mnt/AS4X00/AS4X00_IMAGE.EXE\;1 \
/cdimage/as4x00/as4x00_image.exe
```

- e. Unmount and remove the firmware CD-ROM.

```
% umount /mnt
```

Note

You cannot repackage firmware or software unless you have a specific licensing agreement with the manufacturer that allows you to do so.

4. Use the `mkisofs` program to build the target CDFS file image of the directory structure in `/cdimage`. For example, using `/spare` as the target location for the image:

```
% /usr/sbin/mkisofs -D -R -a -d -o \
/spare/cons_oper_sys.cdfs /cdimage/
```

Refer to the `mkisofs(8)` reference page for more information.

5. Use the `disklabel` command to insert a label into the file generated in Step 4.

```
% disklabel -r -w -t cdfs -f \
/spare/cons_oper_sys.cdfs \
/mdec/rzboot.cdfs /mdec/bootrz.cdfs
```

Refer to the `disklabel(8)` reference page for more information.

The CD image file `/spare/cons_oper_sys.cdfs` is ready to be written to a CD-ROM.

A.2 Sample Build Session

The following assumptions are made for the examples in this section:

- The target partition is on `/dev/disk/dsk6c`.
- The `/spare` directory has at least 1.5 Gb of free space.
- The CD-ROM drive is `/dev/disk/cdrom4`.

Note

The examples in this appendix use the C shell.

A.2.1 Preparing for the Build Session

Follow these steps to prepare for the CD-ROM build session:

1. Log in as `root` or use the `su` command to gain superuser privileges.
2. Enter the following commands:

```
% cd /
% disklabel -F -r -e /dev/disk/dsk6
write new label? [y] y
% mkdir /cdimage
```

3. Place the operating system CD-ROM into the CD-ROM drive, and enter the following commands:

```
% mount -r /dev/disk/cdrom4a /mnt
% cd /mnt
% tar cf /spare/os_copy.tar .
% cd /
% umount /mnt
```

4. Remove the operating system CD-ROM from the drive. The preparatory steps are complete.

A.2.2 Building the Consolidated Firmware CD-ROM

Follow these steps to build a consolidated firmware CD-ROM:

1. Log in as `root` or use the `su` command to gain superuser privileges.
2. Enter the following commands:

```
% cd /
% newfs /dev/disk/dsk6c
% mount /dev/disk/dsk6c /cdimage
% cd /cdimage
% tar xpf /spare/os_copy.tar
% cd /
```

3. Place the *Alpha Systems Firmware CD-ROM* into the CD-ROM drive, and enter the following command:

```
% mount -t cdfs -r /dev/disk/cdrom4a /mnt
% cp /mnt/SMMTABLE.TXT;1 /cdimage/smmtable.txt
% more /cdimage/smmtable.txt
```

4. Review the output to determine the required directory and file names for the firmware images that you want.

The following example uses the same firmware images as Step 3d of Section A.1.2:

```
% mkdir /cdimage/alpha800
% mkdir /cdimage/alpha1000a
% mkdir /cdimage/as4x00
% cp /mnt/ALPHA800/AS800_V5_1.EXE;1 \
/cdimage/alpha800/as800_v5_1.exe
% cp /mnt/ALPHA1000A/AS1000A_E5_V5_1.EXE;1 \
/cdimage/alpha1000a/as1000a_e5_v5_1.exe
% cp /mnt/AS4X00/AS4X00_IMAGE.EXE;1 \
```

```
    /cdimage/as4x00/as4x00_image.exe
% umount /mnt
```

5. Remove the *Alpha Systems Firmware CD-ROM* and enter the following commands:

```
% /usr/sbin/mkisofs -D -R -a -d -o \  
/spare/cons_oper_sys.cdfs /cdimage/
```

Output is similar to the following:

```
Using OSFMANWO.000;1 for \  
/cdimage/ALPHA/BASE/instctrl/OSFMANWOS510.scp \  
(OSFMANWOP510.scp)  
Using OSFMANWO.001;1 for \  
/cdimage/ALPHA/BASE/instctrl/OSFMANWOS510.inv \  
(OSFMANWOP510.inv)  
Using OSFMANWO.002;1 for \  
/cdimage/ALPHA/BASE/instctrl/OSFMANWOS510.ctrl \  
(OSFMANWOP510.ctrl)  
.  
.  
.  
Using PROCFS_V.000;1 for \  
/cdimage/usr/sys/procfs/procfs_vnops_stubs.c \  
(procfs_vfsops_stubs.c)  
 3.92% done, estimate finish Fri May 19 15:36:59  
 5.87% done, estimate finish Fri May 19 15:39:24  
.  
.  
.  
 99.74% done, estimate finish Fri May 19 15:41:52  
Total extents actually written = 255673  
Total translation table size: 0  
Total rockridge attributes bytes: 2066594  
Total directory bytes: 4239360  
Path table size(bytes): 10130  
Max brk space used b9ec60  
255673 extents written (499 Mb)
```

Note

The backslashes (\) in the previous example indicate line continuation and are not present in the output.

6. Enter the following commands:

```
% disklabel -r -w -t cdfs -f \  
/spare/cons_oper_sys.cdfs \  
/mdec/rzboot.cdfs /mdec/bootrz.cdfs
```

The information is consolidated, and the file can be written onto a CD-ROM.

Glossary

This glossary defines terms used in this manual.

A

attribute-value pair

In a product kit's key file, attribute-value pairs specify the names and values of the attributes of the kit, such as the name and version of the product. Attribute-value pairs control how the `kits` utility builds the kit and how the `setld` utility installs it.

B

Backus-Naur form

A conventional notation for describing context-free grammars, commonly used for defining syntax in computer languages. It is named for John Backus, developer of FORTRAN, and Peter Naur, developer of ALGOL. The term BNF is often used to refer to grammar specifications based on this form.

See also *postfix*

backward link

A backward link is a symbolic link from the directories in a layered product area to files in the standard hierarchy. The subset control program for a product creates backward links during installation.

boot utility

The `boot` utility performs the initial installation and bootstrap of the operating system. You invoke the `boot` utility from the `>>>` console prompt. Refer to your hardware documentation for information about valid parameters for the `boot` utility on your system.

C

CDFS

A CD-ROM file system (CDFS) is formatted to be compliant with ISO 9660. This lets you read a CD-ROM from multiple computer platforms.

CDSL

A context-dependent symbolic link (CDSL) is a special form of symbolic link that dynamically resolves to a member-specific file, depending upon the cluster member accessing the file. CDSLs make it possible to maintain

system-specific configuration and data files on file systems shared by the cluster.

See also *cluster*, *cluster member*, *member-specific file*, *shared file*

cluster

A cluster is a tightly coupled collection of one or more otherwise independent servers that share storage and other resources, providing high availability of applications and data, increased scalability of services, and ease of service and system management.

A TruCluster Server cluster uses a common root file system and shares a single name space for files and directories. This cluster file system (CFS) provides the same view of directly attached file systems to all cluster members.

See also *cluster alias*, *cluster member*

cluster alias

An IP address used to address all or a subset of the cluster members. A cluster alias makes some or all of the systems in a cluster look like a single system when viewed from outside the cluster.

See also *cluster*, *cluster member*

cluster member

A system configured with TruCluster Server software that is capable of joining a cluster. A cluster member must be physically connected to both a private physical bus for intracluster communications and to at least one shared SCSI bus.

See also *cluster*

compression flag file

The compression flag file is an empty file whose name consists of the product code and the version number with the string `comp` as a suffix; for example, `OAT100.comp`. If the compression flag file exists, the `setld` utility knows that the subset files are compressed.

context-dependent symbolic link

See *CDSL*

control file

One of a collection of files that the `kits` utility places in the `instctrl` directory. These files include the compression flag file, image data file, subset control file, subset inventory file, and subset control programs.

D

Dataless Management Services

See *DMS*

DMS

Dataless Management Services. A service where a server maintains the root (/), /usr, and /var file systems for client computer systems connected to the server by a local area network (LAN).

data hierarchy

In the kit-building directory structure, the data hierarchy contains the files that direct the `setld` utility in making subsets for the kit, such as the master inventory and key files. An `scps` subdirectory contains subset control programs written by the kit developer.

DCD format

A disk media format where files are written to any disk media (CD-ROM, hard disk, or diskette) as a UNIX file system (UFS). Subsets distributed in DCD format cannot be compressed.

See also *tar format*

dependency expression

A dependency expression is a postfix logical expression consisting of subset identifiers and relational operators to describe the current subset's relationship to the named subsets. Subset control programs evaluate dependency expressions under control of the `setld` utility.

See also *Backus-Naur form, locking, postfix, subset dependency*

direct CD-ROM format

See *DCD format*

distribution media

The distribution media for a product kit may be diskette, CD-ROM, or tape. A hard disk is sometimes referred to as a distribution media because it is used as the master copy for a CD-ROM kit.

E

/etc/sysconfigtab

See *sysconfigtab database*

/etc/kitcap

See *kitcap database*

F

forward link

A forward link is a symbolic link that connects a product file in the `/opt`, `/usr/opt`, or `/var/opt` directory to a standard UNIX directory, such as `/usr/bin`. Forward links allow layered products to be installed in a central location (the `opt` directories) and still be accessible to users through the standard directory structure.

G

gendisk utility

The `gendisk` utility is used to produce disk distribution media for a product kit. Refer to the `gendisk(1)` reference page.

See also *kitcap database*

gentapes utility

The `gentapes` utility is used to produce magnetic tape distribution media for a product kit. Refer to the `gentapes(1)` reference page.

See also *kitcap database*

I

image data file

The image data file is used by the `setld` utility to verify subset image integrity before starting the actual installation process, and contains one record for each subset in the kit.

See also *setld utility*

ISO 9660

ISO 9660 is an international file system standard adopted by major operating system manufacturers. A file system in this format can be read by most of the standard operating systems. Multiple specification levels allow different file naming conventions. ISO 9660-compliant file systems are usually on CD-ROM media.

K

kernel

The kernel is a software entity that runs in supervisor mode and does not communicate with a device except through calls to a device driver.

kernel product

A kernel product is a layered product that runs in kernel space. Users do not directly run kernel products, but the operating system and utilities access them to perform their work.

See also *layered product*, *user product*

key file

A key file identifies the product that the kit represents. You create this file in the `data` directory before running the `kits` utility.

kit

A kit is a collection of files and directories that represent one or more layered products. It is the standard mechanism by which layered product modifications are delivered and maintained on the operating system.

See also *layered product*

kitcap database

The `kitcap` file (located in `/etc/kitcap`) is a kit descriptor database for the `gentapes` and `gendisk` utilities. This database contains product codes, media codes, and the names of the directories, files, and subsets that make up a product description used by these utilities to create distribution media.

The `gentapes` and `gendisk` utilities can specify substitute `kitcap` files in alternate locations.

See also *gendisk utility*, *gentapes utility*

kit descriptor database

See *kitcap database*

kits utility

The `kits` utility creates subsets according to the specifications you define in the master inventory file and key file.

See also *key file*, *master inventory file*, *subset*

L**layered product**

A layered product is an optional software product designed to be installed as an added feature of the operating system.

See also *kernel product*, *user product*

locking

In products installed by the `setld` utility, locking inserts a subset name in the lock file of another subset. Any attempt to remove the latter subset

warns the user of the dependency. The user can choose whether to remove the subset in spite of the dependency.

See also *dependency expression*, *subset dependency*

M

master inventory file

A master inventory file lists all the product files and the subsets in which they belong. You create this file in the `data` directory by running the `newinv` utility. The file must exist before you can create the product subsets.

See also *newinv utility*, *subset*

{memb}

A system variable used to support context-dependent symbolic links (CDSLs). The kernel resolves the `{memb}` variable in a CDSL pathname to the string `memberN`, where `N` is the member ID of the cluster member that is referencing the link. If a cluster member with member ID 2 is accessing a CDSL, the kernel resolves the `{memb}` variable in the pathname to `member2`.

See also *CDSL*, *cluster*, *cluster member*

member-specific file

A file used by a specific cluster member. The contents of a member-specific file can differ for each cluster member, and each member has its own copy of a member-specific file.

See also *cluster*, *cluster member*, *shared file*

N

Network File System

See *NFS*

newinv utility

The `newinv` utility creates the master inventory file from the list of files in the current working directory. The list does not contain all the information needed in the master inventory file. You must edit this file to include information about the subsets to which the files belong.

See also *master inventory file*

NFS

Network File System, an open operating system that allows all network users to access shared files stored on computers of different types. Users can manipulate shared files as if they were stored locally on the user's own hard disk.

O

output hierarchy

The output hierarchy contains the result of the kit-building process, including the subsets that make up the kit and installation control files to direct the `setld` utility during the installation of the product.

P

postfix

A form of logical expression where the operators follow the operands, rather than being placed between them. Also known as reverse Polish notation, or RPN.

See also *Backus-Naur form*

product code

A unique three-letter code that identifies the manufacturer of a product kit. The examples in this manual use the `OAT` product code for the fictional *Orpheus Authoring Tools, Inc.* You request a product code by electronic mail to `product@dssr.sqp.zko.dec.com`.

product kit

See *kit*

R

Remote Installation Services

See *RIS*

RIS

Remote Installation Services. A remote software distribution method where a server is set up to allow installation of software products over a local area network (LAN). RIS clients are registered on the RIS server to allow them access to specific software products. Using a RIS server makes installation of layered products faster and easier for all the clients on the network.

S

SCP

Subset control program. A program written by the kit developer to perform installation operations that the `setld` utility would not otherwise perform. The `setld` utility invokes the subset control program several times during the installation of the kit.

setld utility

The `setld` utility is the standard software management utility. It allows you to load, delete, inventory, configure, and extract software subsets. Refer to the `setld(8)` reference page.

shared file

A file used by all members of a cluster. There is only one copy of a shared file.

See also *cluster*, *cluster member*, *member-specific file*

source hierarchy

In the kit-building directory structure, the source hierarchy contains the files that make up the product. These files are grouped into subsets by the `kits` utility.

subset

The smallest installable software kit module that is compatible with the operating system's `setld` software installation utility. It contains files of any type, usually related in some way.

subset control program

See *SCP*

subset dependency

A subset dependency is the condition under which a given subset requires the presence (or absence) of other subsets in order to function properly.

See also *dependency expression*, *locking*

subset inventory file

The subset inventory file, generated by the `kits` utility, describes each file in the subset to reflect the exact state of the files in the source hierarchy from which the kit was built. The `setld` utility uses this file to duplicate that state, transferring an exact copy of the source hierarchy to the customer's system.

See also *kits utility*, *setld utility*, *source hierarchy*, *subset*

sysconfigdb utility

The `sysconfigdb` utility is a system management tool that maintains the `sysconfigtab` database.

See also *sysconfigtab database*

sysconfigtab database

The `sysconfigtab` database (located in the `/etc/sysconfigtab` file) contains information about the attributes of subsystems, such as device drivers. Device drivers supply attributes in `sysconfigtab` file fragments,

which get appended to the `sysconfigtab` database when the subset control program calls the `sysconfigdb` utility during the installation of a kit.

See also *sysconfigdb utility*

T

tar format

A media format where the product files belonging to the same subset are written to the distribution media as a single file by the `tar` command. During installation, the `setld` utility uncompresses the files, then moves them onto the customer's system, preserving the files' original directory structure. Refer to the `tar(1)` reference page for information about using the `tar` command.

See also *DCD format*

U

user product

A user product is a layered product that runs in user space. Commands, utilities, and user applications fall into this category.

See also *kernel product*, *layered product*

Index

A

ACT environment variable, 3-17

B

backup file
 for master inventory file, 3-53
backward link, 3-23
 (*See also* link)
 creating, 3-23
bootstrap files, 2-4

C

C DELETE phase, 3-28
.c (source) files , 5-7
C INSTALL phase, 3-26
CD-ROM file system
 (*See* CDFS)
CD-ROM
 consolidated firmware, A-1
 (*See also* consolidated
 firmware CD-ROM)
 layered product distribution, 4-2,
 5-9
CDSL, 2-5
cluster-related files, 2-5
compression flag file, 3-43
consolidated firmware CD-ROM
 build instructions, A-1
 definition, A-1
 sample build session, A-4
 preparation, A-5
 procedure, A-5
context-dependent symbolic link

(*See* CDSL)

.ctrl installation control file, 3-41

D

data hierarchy, 2-2
dataless environment
 defined, 3-16
 SCP for, 3-16
 scp routines for, 3-17
DCD format
 defined, 1-3
 layered product files, 4-2, 5-8
dependency expression, 3-21
dependency list, 3-45t
dependency lock
 creating, 3-21
Direct CD-ROM format
 (*See* DCD format)
directory structure, 2-1
 kernel product kit, 5-2
 kit-building, 2-1
 standard, 2-3
disk media
 building a kit on, 4-6, 5-13
 kitcap record, 4-4, 5-11
diskette
 layered product distribution, 4-2,
 5-9
distribution format
 for user products, 4-2, 5-8
distribution media
 producing, 1-6
dot-relative pathnames
 in master inventory records, 3-4t
 in subset inventory records, 3-47t

dynamic configuration, 5-2

E

/etc/kitcap file, 4-3, 5-10

/etc/sysconfigtab database, 5-6

F

file

lock, 3-25

file permissions, 2-10n

files

cluster-related, 2-5

files file fragment, 5-5

G

gendisk utility

syntax, 4-6, 5-13

gentapes utility

preparing a kit on magnetic tape,
4-8, 5-15

global variables

setting in SCP, 3-13

H

.h (header) files, 5-7

I

image data file, 3-43

instctrl file, 3-41

instctrl subdirectory, 2-2

moving files into, 3-41

.inv installation control file, 3-41

ISO 9660

(*See* CDFS)

K

.k file

(*See* key file)

kernel

dynamic configuration, 5-2

static configuration, 5-2

kernel product

defined, 1-2

SCP, 3-35

kernel product kit, 5-1

additional files required for, 5-2

building on disk, 5-13

creating, 5-1

kit directory structure, 5-2

producing distribution media, 5-8

testing, 5-1, 5-16

testing in a RIS area, 5-20

key file

attribute descriptions, 3-7

contents, 3-6

defined, 3-6

in kit-building directory structure,
2-2

product attributes, 3-7

product attributes section, 3-6

sample, 3-6

subset descriptor section, 3-7

subset descriptors, 3-8

kit building process, 1-3

kit formats, 1-3

kit production files

creating, 1-5

kit structure

creating, 1-4

file permissions, 2-10n

kitcap record, 4-3, 5-10

disk media, 4-4, 4-5, 5-11

for tape media, 4-5, 5-12

syntax, 4-3, 5-10

kits utility, 3-39

L

layered product, 1-1

(*See also* kernel product; user
product)

- assigning product version number, 2-1
- defined, 1-1
- obtaining product code, 2-1
- physical location of files, 2-3
- types of products, 1-2
- layered product files
 - in DCD format, 4-2, 5-8
 - in tar format, 4-2, 5-8
- library routines in SCPs, 3-11
- link
 - creating backward, 3-23
 - removing, 3-29
- lock file, 3-25
 - removing, 3-29

M

- M phase, 3-19
- master inventory file
 - creating, 3-3
 - defined, 3-3
 - field descriptions, 3-4t
 - in kit-building directory structure, 2-2
 - sample, 3-5
 - updating, 3-53
- media
 - distribution
 - producing, 1-6
 - tape
 - building a kit on, 4-8, 5-15
- {memb} variable, 2-7
- member-specific file, 2-5
- method file
 - kernel product, 5-7
- .mi file
 - (See master inventory file)
- .mod object module file
 - compressing with objZ utility, 5-7
- .mth method file, 5-7

N

- newinv utility, 3-3, 3-53

O

- object module file
 - kernel product kit, 5-7
- ODB sample product, 1-2
- ODB user product
 - SCP, 3-33
- /opt directory, 2-4
- output hierarchy, 2-2

P

- POST_D phase, 3-31
- POST_L phase, 3-23
- PRE_D phase, 3-29
- PRE_L phase, 3-21
- product code
 - obtaining, 2-1
- product kit
 - kernel, 5-1
 - testing, 1-6
 - user, 4-1
- product subdirectories
 - naming, 2-1
- product version number
 - assigning, 2-1

R

- RIS
 - considerations in SCP, 3-19
 - installing a kernel product, 5-20

S

- sample product
 - for illustration, 1-2
 - SCP, 3-9, 3-16

- checking machine architecture, 3-19
- creating, 1-5
- creating source files, 3-10
- for /dev/none device driver, 3-35
- including library routines, 3-11
- invoking, 3-17
- kernel product, 3-35
- managing subset dependencies, 3-21
- ODB user product, 3-33
- RIS support, 3-19
- setld phase
 - C DELETE, 3-28
 - C INSTALL, 3-26
 - POST_D, 3-31
 - POST_L, 3-23
 - PRE_D, 3-29
 - PRE_L, 3-21
 - V, 3-32
- setld tasks
 - M phase, 3-19
- setting global variables, 3-13
- stopping the program, 3-32
- user product, 3-33
- .scp installation control file, 3-41
- scps subdirectory
 - in kit-building directory structure, 2-2
 - location of subset control files, 3-10
- setld utility
 - ACT environment variable, 3-17
 - C DELETE phase, 3-28
 - C INSTALL phase, 3-26
 - invoking SCP, 3-17
 - lock files, 3-25
 - M phase, 3-19
 - POST_D phase, 3-31
 - POST_L phase, 3-23
 - PRE_D phase, 3-29
 - PRE_L phase, 3-21
 - testing a kernel product, 5-17
 - testing subsets, 3-48
 - testing user product kit, 4-10
 - V phase, 3-32
- shared file, 2-5
- SMM table, A-3
- software subsets for kits, 3-1
- source file
 - kernel product, 5-7
 - SCP, 3-10
- source hierarchy, 2-2
 - file permissions, 2-10n
- SPACE file, 4-6, 5-12
- standard directory structure, 2-3
- static configuration, 5-2
- STL_IsDataless shell, 3-17
- STL_LinkBack shell, 3-24
- STL_LinkInit shell, 3-24
- STL_LinkRemove shell, 3-30
- STL_NoDataless shell, 3-17
- STL_ScpInit shell, 3-13
- subset
 - compressing, 3-41
 - creating, 1-5
 - creating with kits utility, 3-39
 - dependencies, 3-45t
 - dependency, 3-21
 - locking, 3-21, 3-25
 - moving onto distribution media, 4-1, 5-8
 - tar image, 3-41
 - testing, 1-6
- subset control file
 - field descriptions, 3-44
- subset control files for kits, 3-1
- subset control program
 - (*See* SCP)
 - for dataless environment, 3-16
- subset inventory file, 3-46
- subsets
 - testing with setld utility, 3-48
- sysconfigtab file fragment
 - kernel product, 5-6

T

tape media

- building a kit on, 4-8, 5-15
- kitcap record, 4-5, 5-12
- layered product distribution, 4-2, 5-9
- tar format
 - layered product files, 4-2, 5-8
 - producing kits in, 1-3
- test subsets, 1-6
- testing
 - kernel product in RIS area, 5-20
 - kernel product kit, 5-16
 - user product kit, 4-10

U

- user product
 - defined, 1-2

- SCP, 3-33
- user product kit, 4-1
 - building on disk, 4-6
 - creating, 4-1
 - producing distribution media, 4-1
 - testing, 4-1, 4-10
- /usr/opt directory, 2-4
- /usr/share/lib/shell/libscp library, 3-11
- /usr/var/opt directory, 2-4

V

- V phase, 3-32
- verification
 - subset installation, 3-32

How to Order Tru64 UNIX Documentation

To order Tru64 UNIX documentation in the United States and Canada, call **800-344-4825**. In other countries, contact your local Compaq subsidiary.

If you have access to Compaq's intranet, you can place an order at the following Web site:

<http://asmorder.nqo.dec.com/>

If you need help deciding which documentation best meets your needs, see the Tru64 UNIX *Documentation Overview*, which describes the structure and organization of the Tru64 UNIX documentation and provides brief overviews of each document.

The following table provides the order numbers for the Tru64 UNIX operating system documentation kits. For additional information about ordering this and related documentation, see the *Documentation Overview* or contact Compaq.

Name	Order Number
Tru64 UNIX Documentation CD-ROM	QA-6ADAA-G8
Tru64 UNIX Documentation Kit	QA-6ADAA-GZ
End User Documentation Kit	QA-6ADAB-GZ
Startup Documentation Kit	QA-6ADAC-GZ
General User Documentation Kit	QA-6ADAD-GZ
System and Network Management Documentation Kit	QA-6ADAE-GZ
Developer's Documentation Kit	QA-6ADAF-GZ
Reference Pages Documentation Kit	QA-6ADAG-GZ
TruCluster Server Documentation Kit	QA-6BRAA-GZ

Reader's Comments

Tru64 UNIX

Guide to Preparing Product Kits
AA-RH9SC-TE

Compaq welcomes your comments and suggestions on this manual. Your input will help us to write documentation that meets your needs. Please send your suggestions using one of the following methods:

- This postage-paid form
- Internet electronic mail: readers_comment@zk3.dec.com
- Fax: (603) 884-0120, Attn: UBPG Publications, ZKO3-3/Y32

If you are not using this form, please be sure you include the name of the document, the page number, and the product name and version.

Please rate this manual:

	Excellent	Good	Fair	Poor
Accuracy (software works as manual says)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Clarity (easy to understand)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Organization (structure of subject matter)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Figures (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Examples (useful)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Index (ability to find topic)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Usability (ability to access information quickly)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please list errors you have found in this manual:

Page	Description
_____	_____
_____	_____
_____	_____
_____	_____

Additional comments or suggestions to improve this manual:

What version of the software described by this manual are you using? _____

Name, title, department _____

Mailing address _____

Electronic mail _____

Telephone _____

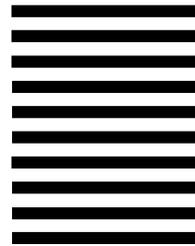
Date _____

----- Do Not Cut or Tear - Fold Here and Tape -----

COMPAQ



NO POSTAGE
NECESSARY IF
MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 33 MAYNARD MA

POSTAGE WILL BE PAID BY ADDRESSEE

COMPAQ COMPUTER CORPORATION
UBPG PUBLICATIONS MANAGER
ZKO3-3/Y32
110 SPIT BROOK RD
NASHUA NH 03062-2698



----- Do Not Cut or Tear - Fold Here -----

Cut on This Line