

AFS Quick Beginnings

Version 3.6

AFS Quick Beginnings: Version 3.6

Copyright © 2000 IBM Corporation. All Rights Reserved

This edition applies to:

IBM AFS for AIX, Version 3.6
IBM AFS for Digital Unix, Version 3.6
IBM AFS for HP-UX, Version 3.6
IBM AFS for Linux, Version 3.6
IBM AFS for SGI IRIX, Version 3.6
IBM AFS for Solaris, Version 3.6

and to all subsequent releases and modifications until otherwise indicated in new editions. This softcopy version is based on the printed edition of this book. Some formatting amendments have been made to make this information more suitable for softcopy.

Revision History

Revision 3.6 April 2000
First Edition

Table of Contents

About This Guide.....	i
Audience and Purpose.....	i
Organization of the Document.....	i
How to Use This Document.....	i
Related Documents.....	i
Typographical Conventions.....	ii
1. Installation Overview	1
The Procedures Described in this Guide.....	1
Required Initial Procedures.....	1
Incorporating AFS Into the Kernel.....	1
Installing the First AFS Machine.....	1
As-needed Procedures.....	2
Upgrading the Operating System.....	2
Installing Additional File Server Machines.....	2
Configuring or Decommissioning Database Server Machines.....	2
Installing Additional AFS Client Machines.....	2
Building AFS from Source Code.....	2
Recommended Reading List.....	2
Requirements.....	3
Login Identity.....	3
General Requirements.....	3
File Server Machine Requirements.....	4
Client Machine Requirements.....	4
Supported System Types.....	5
About Upgrading the Operating System.....	5
The AFS Binary Distribution.....	5
How to Continue.....	6
2. Installing the First AFS Machine	7
Requirements and Configuration Decisions.....	7
Overview: Installing Server Functionality.....	8
Choosing the First AFS Machine.....	8
Creating AFS Directories.....	9
Performing Platform-Specific Procedures.....	9
Getting Started on AIX Systems.....	10
Loading AFS into the AIX Kernel.....	10
Configuring Server Partitions on AIX Systems.....	11
Replacing the fsck Program Helper on AIX Systems.....	12
Enabling AFS Login on AIX Systems.....	12
Getting Started on Digital UNIX Systems.....	13
Loading AFS into the Digital UNIX Kernel.....	14
Building AFS into the Digital UNIX Kernel.....	14
Configuring Server Partitions on Digital UNIX Systems.....	16
Replacing the fsck Program on Digital UNIX Systems.....	17
Enabling AFS Login on Digital UNIX Systems.....	18
Getting Started on HP-UX Systems.....	19

Building AFS into the HP-UX Kernel.....	19
Configuring Server Partitions on HP-UX Systems.....	20
Configuring the AFS-modified fsck Program on HP-UX Systems	21
Enabling AFS Login on HP-UX Systems	22
Getting Started on IRIX Systems.....	24
Loading AFS into the IRIX Kernel	25
Building AFS into the IRIX Kernel.....	26
Configuring Server Partitions on IRIX Systems.....	27
Enabling AFS Login on IRIX Systems	28
Getting Started on Linux Systems	28
Loading AFS into the Linux Kernel	28
Configuring Server Partitions on Linux Systems	29
Enabling AFS Login on Linux Systems	30
Getting Started on Solaris Systems	32
Loading AFS into the Solaris Kernel	32
Configuring the AFS-modified fsck Program on Solaris Systems	33
Configuring Server Partitions on Solaris Systems.....	35
Enabling AFS Login and Editing the File Systems Clean-up Script on Solaris Systems	35
Starting the BOS Server	38
Defining Cell Name and Membership for Server Processes.....	40
Starting the Database Server Processes.....	40
Initializing Cell Security	41
Starting the File Server, Volume Server, and Salvager	44
Starting the Server Portion of the Update Server.....	45
Starting the Controller for NTPD.....	46
Overview: Installing Client Functionality.....	48
Copying Client Files to the Local Disk.....	48
Defining Cell Membership for Client Processes.....	48
Creating the Client CellServDB File	49
Configuring the Cache	51
Configuring a Disk Cache	52
Configuring a Memory Cache	52
Configuring the Cache Manager	53
Overview: Completing the Installation of the First AFS Machine	55
Verifying the AFS Initialization Script	55
Activating the AFS Initialization Script.....	58
Activating the Script on AIX Systems.....	58
Activating the Script on Digital UNIX Systems.....	58
Activating the Script on HP-UX Systems	59
Activating the Script on IRIX Systems	59
Activating the Script on Linux Systems	60
Activating the Script on Solaris Systems.....	60
Configuring the Top Levels of the AFS Filespace	61
Storing AFS Binaries in AFS.....	62
Storing AFS Documents in AFS.....	64
Storing System Binaries in AFS	66
Enabling Access to Foreign Cells	67
Improving Cell Security	69

Controlling root Access	69
Controlling System Administrator Access	69
Protecting Sensitive AFS Directories	70
Removing Client Functionality	70
3. Installing Additional Server Machines	73
Installing an Additional File Server Machine	73
Creating AFS Directories and Performing Platform-Specific Procedures	74
Getting Started on AIX Systems	75
Getting Started on Digital UNIX Systems	76
Getting Started on HP-UX Systems	79
Getting Started on IRIX Systems	81
Getting Started on Linux Systems	85
Getting Started on Solaris Systems	86
Starting Server Programs	88
Installing Client Functionality	91
Completing the Installation	93
Installing Database Server Functionality	97
Summary of Procedures	98
Instructions	99
Removing Database Server Functionality	101
Summary of Procedures	101
Instructions	102
4. Installing Additional Client Machines	105
Summary of Procedures	105
Creating AFS Directories on the Local Disk	105
Performing Platform-Specific Procedures	105
Getting Started on AIX Systems	106
Loading AFS into the AIX Kernel	106
Enabling AFS Login on AIX Systems	107
Getting Started on Digital UNIX Systems	108
Building AFS into the Digital UNIX Kernel	108
Enabling AFS Login on Digital UNIX Systems	110
Getting Started on HP-UX Systems	111
Building AFS into the HP-UX Kernel	111
Enabling AFS Login on HP-UX Systems	113
Getting Started on IRIX Systems	115
Loading AFS into the IRIX Kernel	116
Building AFS into the IRIX Kernel	117
Enabling AFS Login on IRIX Systems	118
Getting Started on Linux Systems	118
Loading AFS into the Linux Kernel	118
Enabling AFS Login on Linux Systems	119
Getting Started on Solaris Systems	124
Loading AFS into the Solaris Kernel	124
Enabling AFS Login on Solaris Systems	125
Loading and Creating Client Files	128
Configuring the Cache	129

Configuring a Disk Cache	130
Configuring a Memory Cache	131
Configuring the Cache Manager	131
Starting the Cache Manager and Installing the AFS Initialization Script.....	133
Running the Script on AIX Systems	134
Running the Script on Digital UNIX Systems	134
Running the Script on HP-UX Systems	135
Running the Script on IRIX Systems	135
Running the Script on Linux Systems	136
Running the Script on Solaris Systems	137
Setting Up Volumes and Loading Binaries into AFS.....	138
Linking /usr/afsws on an Existing System Type	138
Creating Binary Volumes for a New System Type	138
A. Appendix A. Building AFS from Source Code.....	143
Loading the Source Files.....	143
Compiling AFS Binaries Using the washtool Program	143
Index.....	147

About This Guide

This section describes the purpose, organization, and conventions of this document.

Audience and Purpose

This guide explains how to install and configure AFS^(R) server and client machines. It assumes that the reader is familiar with UNIX^(R) system administration, but not AFS.

The instructions explain how to issue AFS commands in the context of specific tasks, but do not describe a command's function or arguments in detail. Refer to the *IBM AFS Administration Reference* as necessary.

Organization of the Document

See "The Procedures Described in this Guide" on page 1.

How to Use This Document

See "The Procedures Described in this Guide" on page 1 and "How to Continue" on page 6.

Related Documents

The AFS documentation set also includes the following documents.

IBM AFS Administration Guide

This guide describes the concepts and procedures that a system administrator must know to manage an AFS cell. It assumes familiarity with UNIX, but requires no previous knowledge of AFS.

The first chapters of the *IBM AFS Administration Guide* present basic concepts and guidelines. Understanding them is crucial to successful administration of an AFS cell. The remaining chapters in the guide provide step-by-step instructions for specific administrative tasks, along with discussions of the concepts important to that particular task.

IBM AFS Administration Reference

This reference manual details the syntax and effect of each AFS command. It is intended for the experienced AFS administrator, programmer, or user.

The *IBM AFS Administration Reference* lists AFS files and commands in alphabetical order. The reference page for each command specifies its syntax, including the acceptable aliases and abbreviations. It then describes the command's function, arguments, and output if any. Examples and a list of related commands are provided, as are warnings where appropriate.

This manual complements the *IBM AFS Administration Guide*: it does not include procedural information, but describes commands in more detail than the *IBM AFS Administration Guide*.

IBM AFS User Guide

This guide presents the basic concepts and procedures necessary for using AFS effectively. It assumes that the reader has some experience with UNIX, but does not require familiarity with networking or AFS.

The guide explains how to perform basic functions, including authenticating, changing a password, protecting AFS data, creating groups, and troubleshooting. It provides illustrative examples for each function and describes some of the differences between the UNIX file system and AFS.

IBM AFS Release Notes

This document provides information specific to each release of AFS, such as a list of new features and commands, a list of requirements and limitations, and instructions for upgrading server and client machines.

Typographical Conventions

This document uses the following typographical conventions:

- Command and option names appear in **bold type** in syntax definitions, examples, and running text. Names of directories, files, machines, partitions, volumes, and users also appear in **bold type**.
- Variable information appears in *italic type*. This includes user-supplied information on command lines and the parts of prompts that differ depending on who issues the command. New terms also appear in *italic type*.
- Examples of screen output and file contents appear in `monospace type`.

In addition, the following symbols appear in command syntax definitions, both in the documentation and in AFS online help statements. When issuing a command, do not type these symbols.

- Square brackets [] surround optional items.
- Angle brackets < > surround user-supplied values in AFS commands.
- A superscripted plus sign + follows an argument that accepts more than one value.
- The percent sign % represents the regular command shell prompt. Some operating systems possibly use a different character for this prompt.
- The number sign # represents the command shell prompt for the local superuser **root**. Some operating systems possibly use a different character for this prompt.
- The pipe symbol | in a command syntax statement separates mutually exclusive values for an argument.

For additional information on AFS commands, including a description of command string components, acceptable abbreviations and aliases, and how to get online help for commands, see the appendix to the *IBM AFS Administration Guide*.

Chapter 1. Installation Overview

This chapter describes the type of instructions provided in this guide and the hardware and software requirements for installing AFS^(R).

Before beginning the installation of your cell's first machine, read this chapter and the material from the *IBM AFS Administration Guide* listed in "Recommended Reading List" on page 2. It is also best to read through "Installing the First AFS Machine" on page 7 before beginning the installation, so that you understand the overall scope of the installation procedure. Similarly, before installing additional server or client machines it is best to read through "Installing Additional Server Machines" on page 73 and "Installing Additional Client Machines" on page 105.

If you are already running a version of AFS, consult the upgrade instructions in the *IBM AFS Release Notes* or contact the AFS Product Support group before proceeding with the installation.

The Procedures Described in this Guide

This guide describes two types of installation procedures: initial procedures (such as installing the first AFS machine or incorporating AFS into the kernel) and as-needed procedures (such as installing additional server machines or client machines).

Required Initial Procedures

You must perform the following basic procedures to start using AFS.

Incorporating AFS Into the Kernel

You must incorporate AFS modifications into the kernel of every AFS file server and client machine. Depending on the operating system, you either use a program for dynamic kernel loading, build a new static kernel, or can choose between the two. For your convenience, the instructions for incorporating AFS into the kernel appear in full in every chapter where you need to use them.

Installing the First AFS Machine

You install the first AFS machine in your cell to function as both an AFS server and client machine. You can disable the client functionality after completing the installation, if you wish.

The first server machine in a cell performs several functions:

- It acts as the *system control machine* (if your AFS distribution includes the required encryption files), distributing certain configuration files to the other server machines in the cell
- It acts as the *binary distribution machine* for its system type, distributing AFS binaries to other server machines of its system type
- It acts as the first *database server machine*, running the server processes that maintain the AFS administrative databases

After you install server and client functionality, you complete other procedures specific to the first machine, including setting up the top levels of your cell's AFS filesystem.

As-needed Procedures

Upgrading the Operating System

Upgrading the operating system requires you to take several steps to protect data and AFS-modified binaries from being lost or overwritten. For guidelines, see "About Upgrading the Operating System" on page 5.

Installing Additional File Server Machines

See "Installing an Additional File Server Machine" on page 73.

Configuring or Decommissioning Database Server Machines

See "Installing Database Server Functionality" on page 97 and "Removing Database Server Functionality" on page 101.

Installing Additional AFS Client Machines

See "Installing Additional Client Machines" on page 105.

Building AFS from Source Code

See "Appendix A, Building AFS from Source Code" on page 143.

Recommended Reading List

To develop the best understanding of the overall scope of an installation procedure, read through the entire chapter or section that describes it before performing any actions.

In addition, familiarity with some basic AFS concepts can make the installation more efficient, because you understand better the purpose of the steps. The following is a prioritized list of material to read before installing the first AFS machine. At minimum, read the first chapter of the *IBM AFS Administration Guide*. Then continue your reading in the indicated order, as extensively as you can. It is more important at this point to read the conceptual material in each section than the instructions.

Selected Topics in the *IBM AFS Administration Guide*

- The chapter titled *An Overview of AFS Administration*

- Selected sections in the *Administering Server Machines* chapter: *Local Disk Files on a Server Machine, The Four Roles for a Server Machine, Maintaining the Server CellServDB File*
- Selected sections in the *Monitoring and Controlling Server Processes* chapter: *Controlling and Checking Process Status*
- Selected sections in the *Managing Server Encryption Keys* chapter: *About Server Encryption Keys*
- Selected sections in the *Managing Volumes* chapter: *About Volumes, Creating Read/write Volumes, Clones and Cloning, Mounting Volumes*
- Selected sections in the *Administering Client Machines and the Cache Manager* chapter: *Overview of Cache Manager Customization, Configuration and Cache-related Files on the Local Disk, Determining the Cache Type, Size, and Location*
- Selected sections in the *Managing Access Control Lists* chapter: *Protecting Data in AFS*

More Selected Topics in the *IBM AFS Administration Guide*

- Selected sections in the *Managing Volumes* chapter: *Creating and Releasing Read-only Volumes (Replication), Creating Backup Volumes*
- Selected sections in the *Administering the Protection Database* chapter: *About the Protection Database*
- Selected sections in the *Administering User Accounts* chapter: *The Components of an AFS User Account*
- Selected sections in the *Managing Administrative Privilege* chapter: *An Overview of Administrative Privilege*

Requirements

You must comply with the following requirements to install AFS successfully.

Login Identity

Log into the machine you are installing as the local superuser **root**. When instructed, also authenticate with AFS as the administrative user **admin**.

General Requirements

- You must have the AFS Binary Distribution for each system type you are installing. Unless otherwise noted, the Binary Distribution includes software for both client and server machines. If you are using the CD-ROM version of the distribution, the machine you are installing must be able to access the CD-ROMs, either through a local CD drive or via an NFS^(R) mount of a CD drive attached to a machine that is accessible by network.
- All AFS machines that belong to a cell must be able to access each other via the network.

- The machine must be running the standard, vendor-supplied version of the operating system supported by the current version of AFS. The operating system must already be installed on the machine's root partition.
- You must be familiar with the current operating system and disk configuration of the machine you are installing.
- All hardware and non-AFS software on the machine must be functioning normally.
- No critical processes can be running on the machine you are installing, because you must reboot it during the installation.

File Server Machine Requirements

- Cell configuration is simplest if the first machine you install has the lowest IP address of any database server machine you currently plan to install. If you later configure a machine with a lower IP address as a database server machine, you must update the **/usr/vice/etc/CellServDB** file on all of your cell's client machines before the installation. For further discussion, see "Installing Database Server Functionality" on page 97.

- The partition mounted on the **/usr** directory must have at least 18 MB of disk space available for storing the AFS server binaries (stored by convention in the **/usr/afs/bin** directory). If the machine is also a client, there must be additional local disk space available, as specified in "Client Machine Requirements" on page 4. The complete set of AFS binaries requires yet more space, but they are normally stored in an AFS volume rather than on a machine's local disk.

More significant amounts of space on the partition are required by the administrative databases stored in the **/usr/afs/db** directory and the server process log files stored in the **/usr/afs/logs** directory. The exact requirement depends on many factors, such as the size of your cell and how often you truncate the log files.

- There must be at least one partition (or logical volume, if the operating system and AFS support them) dedicated exclusively to storing AFS volumes. The total number and size of server partitions on all file server machines in the cell determines how much space is available for AFS files.

Client Machine Requirements

- The partition mounted on the **/usr** directory must have at least 4 MB of disk space available for storing the AFS client binaries and kernel library files (stored by convention in the **/usr/vice/etc** directory). The complete set of AFS binaries requires more space, but they are normally stored in an AFS volume rather than on a machine's local disk. For most system types, the instructions have you copy only the one kernel library file appropriate for the machine you are installing. If you choose to store all of the library files on the local disk, the space requirement can be significantly greater.
- On a client machine that uses a disk cache, there must be enough free space on the cache partition (by convention, mounted on the **/usr/vice/cache** directory) to accommodate the cache. The minimum recommended cache size is 10 MB, but larger caches generally perform better.

- On a client machine that uses a memory cache, there must be at least 5 MB of machine memory to devote to caching, but again more memory generally leads to better performance. For further discussion, see the sections in "Installing Additional Client Machines" on page 105 about configuring the cache.

Supported System Types

The *IBM AFS Release Notes* for each AFS release list the supported system types. Support for subsequent revisions of an operating system often becomes available between AFS releases. The AFS Product Support group can provide details.

It is the goal of the AFS Development and Product Support groups to support AFS on a wide range of popular system types. Furthermore, each time an operating system vendor releases a new general availability version of a supported operating system, it is a goal to certify and support AFS on it within a short time. Support can be delayed a bit longer if it is necessary to generate completely new binaries.

It is not always possible to support AFS on every intermediate version of an operating system or for certain processor types. In some cases, platform limitations make certain AFS functionality (such as file server or NFS/AFS translator functionality) unavailable on one or more platforms. For a list of limitations, see the *IBM AFS Release Notes* or ask the AFS Product Support group.

About Upgrading the Operating System

Whenever you upgrade an AFS machine to a different operating system, you must take several actions to maintain proper AFS functionality. These actions include, but are not necessarily limited to, the following.

- Unmount the AFS server partitions (mounted at **/vicep_{xx}** directories) on all file server machines, to prevent the vendor-supplied **fsck** program from running on them when you reboot the machine during installation of the new operating system. Before upgrading the operating system, it is prudent to comment out commands in the machine's initialization file that remount the server partitions, to prevent them from being remounted until you can replace the standard **fsck** program with the AFS-modified version. The instructions in this guide for installing AFS server machines explain how to replace the **fsck** program.
- Protect the AFS-modified versions of commands and configuration files from being overwritten by vendor-supplied versions. These include **vfsck** (the AFS version of **fsck**), binaries for the UNIX remote services such as **inetd**, and configuration files such as the one for the Pluggable Authentication Module (PAM). After you have successfully installed the operating system, remember to move the AFS-modified commands and files back to the locations where they are accessed during normal functioning.
- Reformat the server partitions to accommodate AFS-specific information, in certain cases. The upgrade instructions that accompany the new AFS binaries for an affected platform always detail the required procedure.

The AFS Binary Distribution

The AFS Binary Distribution includes a separate CD-ROM for each supported system type, containing all AFS binaries and files for both server and client machines. The instructions in this guide specify when to mount the CD-ROM and which files or directories to copy to the local disk or into an AFS volume.

How to Continue

If you are installing the first AFS machine in your cell, proceed to "Installing the First AFS Machine" on page 7.

If you are installing an additional file server machine, or configuring or decommissioning a database server machine, proceed to "Installing Additional Server Machines" on page 73.

If you are installing an additional client machine, proceed to "Installing Additional Client Machines" on page 105.

Chapter 2. Installing the First AFS Machine

This chapter describes how to install the first AFS machine in your cell, configuring it as both a file server machine and a client machine. After completing all procedures in this chapter, you can remove the client functionality if you wish, as described in "Removing Client Functionality" on page 70.

To install additional file server machines after completing this chapter, see "Installing Additional Server Machines" on page 73.

To install additional client machines after completing this chapter, see "Installing Additional Client Machines" on page 105.

Requirements and Configuration Decisions

The instructions in this chapter assume that you meet the following requirements.

- You are logged onto the machine's console as the local superuser **root**
- A standard version of one of the operating systems supported by the current version of AFS is running on the machine
- You can access the data on the AFS CD-ROMs, either through a local CD drive or via an NFS mount of a CD drive attached to a machine that is accessible by network

You must make the following configuration decisions while installing the first AFS machine. To speed the installation itself, it is best to make the decisions before beginning. See the chapter in the *IBM AFS Administration Guide* about issues in cell administration and configuration for detailed guidelines.

- Select the first AFS machine
- Select the cell name
- Decide which partitions or logical volumes to configure as AFS server partitions, and choose the directory names on which to mount them
- Decide whether to use the standard AFS authentication and authorization software or Kerberos as obtained from another source. On several system types, the decision determines how you incorporate AFS into the machine's authentication system. If you wish to use Kerberos, contact the AFS Product Support group now to learn about how you must modify the installation procedure.
- Decide how big to make the client cache
- Decide how to configure the top levels of your cell's AFS filesystem

This chapter is divided into three large sections corresponding to the three parts of installing the first AFS machine. Perform all of the steps in the order they appear. Each functional section begins with a summary of the procedures to perform. The sections are as follows:

- Installing server functionality (begins in "Overview: Installing Server Functionality" on page 8)
- Installing client functionality (begins in "Overview: Installing Client Functionality" on page 48)

- Configuring your cell's filesystem, establishing further security mechanisms, and enabling access to foreign cells (begins in "Overview: Completing the Installation of the First AFS Machine" on page 55)

Overview: Installing Server Functionality

In the first phase of installing your cell's first AFS machine, you install file server and database server functionality by performing the following procedures:

1. Choose which machine to install as the first AFS machine
2. Create AFS-related directories on the local disk
3. Incorporate AFS modifications into the machine's kernel
4. Configure partitions or logical volumes for storing AFS volumes
5. On some system types, install and configure an AFS-modified version of the **fsck** program
6. If the machine is to remain a client machine, incorporate AFS into its authentication system
7. Start the Basic OverSeer (BOS) Server
8. Define the cell name and the machine's cell membership
9. Start the database server processes: Authentication Server, Backup Server, Protection Server, and Volume Location (VL) Server
10. Configure initial security mechanisms
11. Start the **fs** process, which incorporates three component processes: the File Server, Volume Server, and Salvager
12. Start the server portion of the Update Server
13. Start the controller process (called **runntp**) for the Network Time Protocol Daemon, which synchronizes machine clocks

Choosing the First AFS Machine

The first AFS machine you install must have sufficient disk space to store AFS volumes. To take best advantage of AFS's capabilities, store client-side binaries as well as user files in volumes. When you later install additional file server machines in your cell, you can distribute these volumes among the different machines as you see fit.

These instructions configure the first AFS machine as a *database server machine*, the *binary distribution machine* for its system type, and the cell's *system control machine*. For a description of these roles, see the *IBM AFS Administration Guide*.

Installation of additional machines is simplest if the first machine has the lowest IP address of any database server machine you currently plan to install. If you later install database server functionality on a machine with a lower IP address, you must first update the `/usr/vice/etc/CellServDB` file on all of your cell's client machines. For more details, see "Installing Database Server Functionality" on page 97.

Creating AFS Directories

Create the `/usr/afs` and `/usr/vice/etc` directories on the local disk, to house server and client files respectively. Subsequent instructions copy files from the AFS CD-ROM into them. Create the `/cdrom` directory as a mount point for CD-ROMs, if it does not already exist.

```
# mkdir /usr/afs
# mkdir /usr/vice
# mkdir /usr/vice/etc
# mkdir /cdrom
```

Performing Platform-Specific Procedures

Several of the initial procedures for installing a file server machine differ for each system type. For convenience, the following sections group them together for each system type:

- Incorporate AFS modifications into the kernel.

The kernel on every AFS file server and client machine must incorporate AFS extensions. On machines that use a dynamic kernel module loader, it is conventional to alter the machine's initialization script to load the AFS extensions at each reboot.

- Configure server partitions or logical volumes to house AFS volumes.

Every AFS file server machine must have at least one partition or logical volume dedicated to storing AFS volumes (for convenience, the documentation hereafter refers to partitions only). Each server partition is mounted at a directory named `/vicepxx`, where `xx` is one or two lowercase letters. By convention, the first 26 partitions are mounted on the directories called `/vicepa` through `/vicepz`, the 27th one is mounted on the `/vicepaa` directory, and so on through `/vicepaz` and `/vicepba`, continuing up to the index corresponding to the maximum number of server partitions supported in the current version of AFS (which is specified in the *IBM AFS Release Notes*).

The `/vicepxx` directories must reside in the file server machine's root directory, not in one of its subdirectories (for example, `/usr/vicepa` is not an acceptable directory location).

You can also add or remove server partitions on an existing file server machine. For instructions, see the chapter in the *IBM AFS Administration Guide* about maintaining server machines.

Note: Not all file system types supported by an operating system are necessarily supported as AFS server partitions. For possible restrictions, see the *IBM AFS Release Notes*.

- On some system types, install and configure a modified `fsck` program which recognizes the structures that the File Server uses to organize volume data on AFS server partitions. The `fsck` program provided with the operating system does not understand the AFS data structures, and so removes them to the `lost+found` directory.

- If the machine is to remain an AFS client machine, modify the machine's authentication system so that users obtain an AFS token as they log into the local file system. Using AFS is simpler and more convenient for your users if you make the modifications on all client machines. Otherwise, users must perform a two-step login procedure (login to the local file system and then issue the **klog** command). For further discussion of AFS authentication, see the chapter in the *IBM AFS Administration Guide* about cell configuration and administration issues.

To continue, proceed to the appropriate section:

- "Getting Started on AIX Systems" on page 10
- "Getting Started on Digital UNIX Systems" on page 13
- "Getting Started on HP-UX Systems" on page 19
- "Getting Started on IRIX Systems" on page 24
- "Getting Started on Linux Systems" on page 28
- "Getting Started on Solaris Systems" on page 32

Getting Started on AIX Systems

Begin by running the AFS initialization script to call the AIX kernel extension facility, which dynamically loads AFS modifications into the kernel. Then use the **SMIT** program to configure partitions for storing AFS volumes, and replace the AIX **fsck** program helper with a version that correctly handles AFS volumes. If the machine is to remain an AFS client machine, incorporate AFS into the AIX secondary authentication system.

Loading AFS into the AIX Kernel

The AIX kernel extension facility is the dynamic kernel loader provided by IBM Corporation. AIX does not support incorporation of AFS modifications during a kernel build.

For AFS to function correctly, the kernel extension facility must run each time the machine reboots, so the AFS initialization script (included in the AFS distribution) invokes it automatically. In this section you copy the script to the conventional location and edit it to select the appropriate options depending on whether NFS is also to run.

After editing the script, you run it to incorporate AFS into the kernel. In later sections you verify that the script correctly initializes all AFS components, then configure the AIX **inittab** file so that the script runs automatically at reboot.

1. Mount the AFS CD-ROM for AIX on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your AIX documentation. Then change directory as indicated.

```
# cd /cdrom/rs_aix42/root.client/usr/vice/etc
```

- Copy the AFS kernel library files to the local `/usr/vice/etc/dkload` directory, and the AFS initialization script to the `/etc` directory.

```
# cp -rp dkload /usr/vice/etc
# cp -p rc.afs /etc/rc.afs
```

- Edit the `/etc/rc.afs` script, setting the `NFS` variable as indicated.

If the machine is not to function as an NFS/AFS Translator, set the `NFS` variable as follows.

```
NFS=$NFS_NONE
```

If the machine is to function as an NFS/AFS Translator and is running AIX 4.2.1 or higher, set the `NFS` variable as follows. Note that NFS must already be loaded into the kernel, which happens automatically on systems running AIX 4.1.1 and later, as long as the file `/etc/exports` exists.

```
NFS=$NFS_IAUTH
```

- Invoke the `/etc/rc.afs` script to load AFS modifications into the kernel. You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/rc.afs
```

Configuring Server Partitions on AIX Systems

Every AFS file server machine must have at least one partition or logical volume dedicated to storing AFS volumes. Each server partition is mounted at a directory named `/vicepxx`, where `xx` is one or two lowercase letters. The `/vicepxx` directories must reside in the file server machine's root directory, not in one of its subdirectories (for example, `/usr/vicepa` is not an acceptable directory location). For additional information, see "Performing Platform-Specific Procedures" on page 9.

To configure server partitions on an AIX system, perform the following procedures:

- Create a directory called `/vicepxx` for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

- Use the **SMIT** program to create a journaling file system on each partition to be configured as an AFS server partition.
- Mount each partition at one of the `/vicepxx` directories. Choose one of the following three methods:

- Use the **SMIT** program
- Use the **mount -a** command to mount all partitions at once
- Use the **mount** command on each partition in turn

Also configure the partitions so that they are mounted automatically at each reboot. For more information, refer to the AIX documentation.

Replacing the fsck Program Helper on AIX Systems

In this section, you make modifications to guarantee that the appropriate **fsck** program runs on AFS server partitions. The **fsck** program provided with the operating system must never run on AFS server partitions. Because it does not recognize the structures that the File Server uses to organize volume data, it removes all of the data. To repeat:

Never run the standard fsck program on AFS server partitions. It discards AFS volumes.

On AIX systems, you do not replace the **fsck** binary itself, but rather the *program helper* file included in the AIX distribution as `/sbin/helpers/v3fshelper`.

1. Move the AIX **fsck** program helper to a safe location and install the version from the AFS distribution in its place. The AFS CD-ROM must still be mounted at the `/cdrom` directory.

```
# cd /sbin/helpers
# mv v3fshelper v3fshelper.noafs
# cp -p /cdrom/rs_aix42/root.server/etc/v3fshelper v3fshelper
```

2. If you plan to retain client functionality on this machine after completing the installation, proceed to "Enabling AFS Login on AIX Systems" on page 12. Otherwise, proceed to "Starting the BOS Server" on page 38.

Enabling AFS Login on AIX Systems

Note: If you plan to remove client functionality from this machine after completing the installation, skip this section and proceed to "Starting the BOS Server" on page 38.

Follow the instructions in this section to incorporate AFS modifications into the AIX secondary authentication system.

1. Issue the **ls** command to verify that the **afs_dynamic_auth** and **afs_dynamic_kerbauth** programs are installed in the local `/usr/vice/etc` directory.

```
# ls /usr/vice/etc
```

If the files do not exist, mount the AFS CD-ROM for AIX (if it is not already), change directory as indicated, and copy them.

```
# cd /cdrom/rs_aix42/root.client/usr/vice/etc
# cp -p afs_dynamic* /usr/vice/etc
```

2. Edit the local `/etc/security/user` file, making changes to the indicated stanzas:

- In the default stanza, set the `registry` attribute to **DCE** (not to **AFS**), as follows:

```
registry = DCE
```

- In the default stanza, set the `SYSTEM` attribute as indicated.

If the machine is an AFS client only, set the following value:

```
SYSTEM = "AFS OR (AFS[UNAVAIL] AND compat[SUCCESS]) "
```

If the machine is both an AFS and a DCE client, set the following value (it must appear on a single line in the file):

```
SYSTEM = "DCE OR DCE[UNAVAIL] OR AFS OR (AFS[UNAVAIL] \
AND compat[SUCCESS]) "
```

- In the `root` stanza, set the `registry` attribute as follows. It enables the local superuser **root** to log into the local file system only, based on the password listed in the local password file.

```
root:
    registry = files
```

3. Edit the local `/etc/security/login.cfg` file, creating or editing the indicated stanzas:

- In the `DCE` stanza, set the `program` attribute as follows.

If you use the AFS Authentication Server (**kaserver** process):

```
DCE:
    program = /usr/vice/etc/afs_dynamic_auth
```

If you use a Kerberos implementation of AFS authentication:

```
DCE:
    program = /usr/vice/etc/afs_dynamic_kerbauth
```

- In the `AFS` stanza, set the `program` attribute as follows.

If you use the AFS Authentication Server (**kaserver** process):

```
AFS:
    program = /usr/vice/etc/afs_dynamic_auth
```

If you use a Kerberos implementation of AFS authentication:

```
AFS:
    program = /usr/vice/etc/afs_dynamic_kerbauth
```

4. Proceed to "Starting the BOS Server" on page 38 (or if referring to these instructions while installing an additional file server machine, return to "Starting Server Programs" on page 88).

Getting Started on Digital UNIX Systems

Begin by either building AFS modifications into a new static kernel or by setting up to dynamically load the AFS kernel module. Then create partitions for storing AFS volumes, and replace the Digital UNIX **fsck** program with a version that correctly handles AFS volumes. If the machine is to remain an AFS client machine, incorporate AFS into the machine's Security Integration Architecture (SIA) matrix.

Loading AFS into the Digital UNIX Kernel

The **sysconfig** program is the dynamic kernel loader provided for Digital UNIX systems.

For AFS to function correctly, the **sysconfig** program must run each time the machine reboots, so the AFS initialization script (included on the AFS CD-ROM) invokes it automatically. In this section you copy the appropriate AFS library file to the location where the **sysconfig** program accesses it and then run the script.

Mount the AFS CD-ROM for Digital UNIX on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your Digital UNIX documentation. Then change directory as indicated.

```
# cd /cdrom/alpha_dux40/root.client
```

Copy the AFS initialization script to the local directory for initialization files (by convention, **/sbin/init.d** on Digital UNIX machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp usr/vice/etc/afs.rc /sbin/init.d/afs
```

Copy the AFS kernel module to the local **/subsys** directory.

```
# cp bin/afs.mod /subsys/afs.mod
```

Set up the system to load the module at startup.

```
# /sbin/init.d/autosysconfig add afs
```

Reboot the machine to start using the new kernel, and login again as the superuser **root**.

```
# cd /
# shutdown -r now
login: root
Password: root_password
```

Building AFS into the Digital UNIX Kernel

Use the following instructions to build AFS modifications into the kernel on a Digital UNIX system.

1. Create a copy called **AFS** of the basic kernel configuration file included in the Digital UNIX distribution as **/usr/sys/conf/machine_name**, where *machine_name* is the machine's hostname in all uppercase letters.

```
# cd /usr/sys/conf
# cp machine_name AFS
```

2. Add AFS to the list of options in the configuration file you created in the previous step, so that the result looks like the following:

```

      .
      .
options  UFS
options  NFS
options  AFS
      .
      .

```

3. Add an entry for AFS to two places in the file `/usr/sys/conf/files`.

- Add a line for AFS to the list of `OPTIONS`, so that the result looks like the following:

```

      .
      .
OPTIONS/nfs      optional nfs
OPTIONS/afs      optional afs
OPTIONS/nfs_server optional nfs_server
      .
      .

```

- Add an entry for AFS to the list of `MODULES`, so that the result looks like the following:

```

      .
      .
      .
      .
#
MODULE/nfs_server optional nfs_server Binary
nfs/nfs_server.c  module nfs_server optimize -g3
nfs/nfs3_server.c module nfs_server optimize -g3
#
MODULE/afs        optional afs Binary
afs/libafs.c      module afs
#

```

4. Add an entry for AFS to two places in the file `/usr/sys/vfs/vfs_conf.c`.

- Add AFS to the list of defined file systems, so that the result looks like the following:

```

      .
      .
#include <afs.h>
#if defined(AFS) && AFS
    extern struct vfsops afs_vfsops;
#endif
      .
      .

```

- Put a declaration for AFS in the `vfssw[]` table's `MOUNT_ADDON` slot, so that the result looks like the following:

```

      .
      .
      .
      .
&fdfs_vfsops, "fdfs", /* 12 = MOUNT_FDFS */
#if defined(AFS)

```

```
        &afs_vfsops,          "afs",
    #else
        (struct vfsops *)0,   "",          /* 13 = MOUNT_ADDON */
    #endif
    #if NFS && INFS_DYNAMIC
        &nfs3_vfsops,         "nfsv3", /* 14 = MOUNT_NFS3 */
```

5. Mount the AFS CD-ROM for Digital UNIX on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your Digital UNIX documentation. Then change directory as indicated.

```
# cd /cdrom/alpha_dux40/root.client
```

6. Copy the AFS initialization script to the local directory for initialization files (by convention, **/sbin/init.d** on Digital UNIX machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp usr/vice/etc/afs.rc /sbin/init.d/afs
```

7. Copy the AFS kernel module to the local **/usr/sys/BINARY** directory.

If the machine's kernel supports NFS server functionality:

```
# cp bin/libafs.o /usr/sys/BINARY/afs.mod
```

If the machine's kernel does not support NFS server functionality:

```
# cp bin/libafs.nonfs.o /usr/sys/BINARY/afs.mod
```

8. Configure and build the kernel. Respond to any prompts by pressing **<Return>**. The resulting kernel resides in the file **/sys/AFS/vmunix**.

```
# doconfig -c AFS
```

9. Rename the existing kernel file and copy the new, AFS-modified file to the standard location.

```
# mv /vmunix /vmunix_noafs
# cp /sys/AFS/vmunix /vmunix
```

10. Reboot the machine to start using the new kernel, and login again as the superuser **root**.

```
# cd /
# shutdown -r now
login: root
Password: root_password
```

Configuring Server Partitions on Digital UNIX Systems

Every AFS file server machine must have at least one partition or logical volume dedicated to storing AFS volumes. Each server partition is mounted at a directory named **/vicep_{xx}**, where **xx** is one or two lowercase letters. The **/vicep_{xx}** directories must reside in the file server machine's root directory, not in

one of its subdirectories (for example, **/usr/vicepa** is not an acceptable directory location). For additional information, see "Performing Platform-Specific Procedures" on page 9.

1. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

2. Add a line with the following format to the file systems registry file, **/etc/fstab**, for each directory just created. The entry maps the directory name to the disk partition to be mounted on it.

```
/dev/disk /vicepxx ufs rw 0 2
```

The following is an example for the first partition being configured.

```
/dev/rz3a /vicepa ufs rw 0 2
```

3. Create a file system on each partition that is to be mounted at a **/vicep_{xx}** directory. The following command is probably appropriate, but consult the Digital UNIX documentation for more information.

```
# newfs -v /dev/disk
```

4. Mount each partition by issuing either the **mount -a** command to mount all partitions at once or the **mount** command to mount each partition in turn.

Replacing the fsck Program on Digital UNIX Systems

In this section, you make modifications to guarantee that the appropriate **fsck** program runs on AFS server partitions. The **fsck** program provided with the operating system must never run on AFS server partitions. Because it does not recognize the structures that the File Server uses to organize volume data, it removes all of the data. To repeat:

Never run the standard fsck program on AFS server partitions. It discards AFS volumes.

On Digital UNIX systems, the files **/sbin/fsck** and **/usr/sbin/fsck** are driver programs. Rather than replacing either of them, you replace the actual binary included in the Digital UNIX distribution as **/sbin/ufs_fsck** and **/usr/sbin/ufs_fsck**.

1. Install the **vfscck** binary to the **/sbin** and **/usr/sbin** directories. The AFS CD-ROM must still be mounted at the **/cdrom** directory.

```
# cd /cdrom/alpha_dux40/root.server/etc
# cp vfscck /sbin/vfscck
# cp vfscck /usr/sbin/vfscck
```

2. Rename the Digital UNIX **fsck** binaries and create symbolic links to the **vfscck** program.

```
# cd /sbin
# mv ufs_fsck ufs_fsck.noafs
# ln -s vfscck ufs_fsck
```

```
# cd /usr/sbin
# mv ufs_fsck ufs_fsck.noafs
# ln -s vfsck ufs_fsck
```

3. If you plan to retain client functionality on this machine after completing the installation, proceed to "Enabling AFS Login on Digital UNIX Systems" on page 18. Otherwise, proceed to "Starting the BOS Server" on page 38.

Enabling AFS Login on Digital UNIX Systems

Note: If you plan to remove client functionality from this machine after completing the installation, skip this section and proceed to "Starting the BOS Server" on page 38.

On Digital UNIX systems, the AFS initialization script automatically incorporates the AFS authentication library file into the Security Integration Architecture (SIA) matrix on the machine, so that users with AFS accounts obtain a token at login. In this section you copy the library file to the appropriate location.

For more information on SIA, see the Digital UNIX reference page for **matrix.conf**, or consult the section on security in your Digital UNIX documentation.

Note: If the machine runs both the DCE and AFS client software, AFS must start after DCE. Consult the AFS initialization script for suggested symbolic links to create for correct ordering. Also, the system startup script order must initialize SIA before any long-running process that uses authentication.

Perform the following steps to enable AFS login.

1. Mount the AFS CD-ROM for Digital UNIX on the local **/cdrom** directory, if it is not already. Change directory as indicated.

```
# cd /cdrom/alpha_dux40/lib/afs
```

2. Copy the appropriate AFS authentication library file to the local **/usr/shlib** directory.

If you use the AFS Authentication Server (**kaserver** process) in the cell:

```
# cp libafssiad.so /usr/shlib
```

If you use a Kerberos implementation of AFS authentication, rename the library file as you copy it:

```
# cp libafssiad.krb.so /usr/shlib/libafssiad.so
```

3. Proceed to "Starting the BOS Server" on page 38 (or if referring to these instructions while installing an additional file server machine, return to "Starting Server Programs" on page 88).

Getting Started on HP-UX Systems

Begin by building AFS modifications into a new kernel; HP-UX does not support dynamic loading. Then create partitions for storing AFS volumes, and install and configure the AFS-modified **fsck** program to run on AFS server partitions. If the machine is to remain an AFS client machine, incorporate AFS into the machine's Pluggable Authentication Module (PAM) scheme.

Building AFS into the HP-UX Kernel

Use the following instructions to build AFS modifications into the kernel on an HP-UX system.

1. Move the existing kernel-related files to a safe location.

```
# cp /stand/vmunix /stand/vmunix.noafs
# cp /stand/system /stand/system.noafs
```

2. Mount the AFS CD-ROM for HP-UX on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your HP-UX documentation. Then change directory as indicated.

```
# cd /cdrom/hp_ux110/root.client
```

3. Copy the AFS initialization file to the local directory for initialization files (by convention, **/sbin/init.d** on HP-UX machines). Note the removal of the **.rc** extension as you copy the file.

```
# cp usr/vice/etc/afs.rc /sbin/init.d/afs
```

4. Copy the file **afs.driver** to the local **/usr/conf/master.d** directory, changing its name to **afs** as you do.

```
# cp usr/vice/etc/afs.driver /usr/conf/master.d/afs
```

5. Copy the AFS kernel module to the local **/usr/conf/lib** directory.

If the machine's kernel supports NFS server functionality:

```
# cp bin/libafs.a /usr/conf/lib
```

If the machine's kernel does not support NFS server functionality, change the file's name as you copy it:

```
# cp bin/libafs.nonfs.a /usr/conf/lib/libafs.a
```

6. Incorporate the AFS driver into the kernel, either using the **SAM** program or a series of individual commands.

- To use the **SAM** program:

- a. Invoke the **SAM** program, specifying the hostname of the local machine as *local_hostname*. The **SAM** graphical user interface pops up.

```
# sam -display local_hostname:0
```

- b. Choose the **Kernel Configuration** icon, then the **Drivers** icon. From the list of drivers, select **afs**.
- c. Open the pull-down **Actions** menu and choose the **Add Driver to Kernel** option.
- d. Open the **Actions** menu again and choose the **Create a New Kernel** option.
- e. Confirm your choices by choosing **Yes** and **OK** when prompted by subsequent pop-up windows. The **SAM** program builds the kernel and reboots the system.
- f. Login again as the superuser **root**.

```
login: root
Password: root_password
```

- To use individual commands:

- a. Edit the file **/stand/system**, adding an entry for **afs** to the `Subsystems` section.
- b. Change to the **/stand/build** directory and issue the **mk_kernel** command to build the kernel.

```
# cd /stand/build
# mk_kernel
```

- c. Move the new kernel to the standard location (**/stand/vmunix**), reboot the machine to start using it, and login again as the superuser **root**.

```
# mv /stand/build/vmunix_test /stand/vmunix
# cd /
# shutdown -r now
login: root
Password: root_password
```

Configuring Server Partitions on HP-UX Systems

Every AFS file server machine must have at least one partition or logical volume dedicated to storing AFS volumes. Each server partition is mounted at a directory named **/vicep_{xx}**, where **xx** is one or two lowercase letters. The **/vicep_{xx}** directories must reside in the file server machine's root directory, not in one of its subdirectories (for example, **/usr/vicepa** is not an acceptable directory location). For additional information, see "Performing Platform-Specific Procedures" on page 9.

1. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

2. Use the **SAM** program to create a file system on each partition. For instructions, consult the HP-UX documentation.

3. On some HP-UX systems that use logical volumes, the **SAM** program automatically mounts the partitions. If it has not, mount each partition by issuing either the **mount -a** command to mount all partitions at once or the **mount** command to mount each partition in turn.

Configuring the AFS-modified fsck Program on HP-UX Systems

In this section, you make modifications to guarantee that the appropriate **fsck** program runs on AFS server partitions. The **fsck** program provided with the operating system must never run on AFS server partitions. Because it does not recognize the structures that the File Server uses to organize volume data, it removes all of the data. To repeat:

Never run the standard fsck program on AFS server partitions. It discards AFS volumes.

On HP-UX systems, there are several configuration files to install in addition to the AFS-modified **fsck** program (the **vfsc** binary).

1. Create the command configuration file `/sbin/lib/mfsconfig.d/afs`. Use a text editor to place the indicated two lines in it:

```
format_revision 1
fsck             0             m,P,p,d,f,b:c:y,n,Y,N,q,
```

2. Create and change directory to an AFS-specific command directory called `/sbin/fs/afs`.

```
# mkdir /sbin/fs/afs
# cd /sbin/fs/afs
```

3. Copy the AFS-modified version of the **fsck** program (the **vfsc** binary) and related files from the distribution directory to the new AFS-specific command directory.

```
# cp -p /cdrom/hp_ux110/root.server/etc/* .
```

4. Change the **vfsc** binary's name to **fsck** and set the mode bits appropriately on all of the files in the `/sbin/fs/afs` directory.

```
# mv vfsc fsck
# chmod 755 *
```

5. Edit the `/etc/fstab` file, changing the file system type for each AFS server partition from `hfs` to `afs`. This ensures that the AFS-modified **fsck** program runs on the appropriate partitions.

The sixth line in the following example of an edited file shows an AFS server partition, `/vicepa`.

```
/dev/vg00/lvol1 / hfs defaults 0 1
/dev/vg00/lvol4 /opt hfs defaults 0 2
/dev/vg00/lvol5 /tmp hfs defaults 0 2
/dev/vg00/lvol6 /usr hfs defaults 0 2
/dev/vg00/lvol8 /var hfs defaults 0 2
/dev/vg00/lvol9 /vicepa afs defaults 0 2
/dev/vg00/lvol7 /usr/vice/cache hfs defaults 0 2
```

6. If you plan to retain client functionality on this machine after completing the installation, proceed to "Enabling AFS Login on HP-UX Systems" on page 22. Otherwise, proceed to "Starting the BOS Server" on page 38.

Enabling AFS Login on HP-UX Systems

Note: If you plan to remove client functionality from this machine after completing the installation, skip this section and proceed to "Starting the BOS Server" on page 38.

At this point you incorporate AFS into the operating system's Pluggable Authentication Module (PAM) scheme. PAM integrates all authentication mechanisms on the machine, including login, to provide the security infrastructure for authenticated access to and from the machine.

Explaining PAM is beyond the scope of this document. It is assumed that you understand the syntax and meanings of settings in the PAM configuration file (for example, how the `other` entry works, the effect of marking an entry as `required`, `optional`, or `sufficient`, and so on).

The following instructions explain how to alter the entries in the PAM configuration file for each service for which you wish to use AFS authentication. Other configurations possibly also work, but the instructions specify the recommended and tested configuration.

Note: The instructions specify that you mark each entry as `optional`. However, marking some modules as optional can mean that they grant access to the corresponding service even when the user does not meet all of the module's requirements. In some operating system revisions, for example, if you mark as optional the module that controls login via a dial-up connection, it allows users to login without providing a password. See the *IBM AFS Release Notes* for a discussion of any limitations that apply to this operating system.

Also, with some operating system versions you must install patches for PAM to interact correctly with certain authentication programs. For details, see the *IBM AFS Release Notes*.

The recommended AFS-related entries in the PAM configuration file make use of one or more of the following three attributes.

try_first_pass

This is a standard PAM attribute that can be included on entries after the first one for a service; it directs the module to use the password that was provided to the first module. For the AFS module, it means that AFS authentication succeeds if the password provided to the module listed first is the user's correct AFS password. For further discussion of this attribute and its alternatives, see the operating system's PAM documentation.

ignore_root

This attribute, specific to the AFS PAM module, directs it to ignore not only the local superuser **root**, but also any user with UID 0 (zero).

setenv_password_expires

This attribute, specific to the AFS PAM module, sets the environment variable `PASSWORD_EXPIRES` to the expiration date of the user's AFS password, which is recorded in the Authentication Database.

Perform the following steps to enable AFS login.

1. Mount the AFS CD-ROM for HP-UX on the `/cdrom` directory, if it is not already. Then change directory as indicated.

```
# cd /usr/lib/security
```

2. Copy the AFS authentication library file to the `/usr/lib/security` directory. Then create a symbolic link to it whose name does not mention the version. Omitting the version eliminates the need to edit the PAM configuration file if you later update the library file.

If you use the AFS Authentication Server (**kaserver** process) in the cell:

```
# cp /cdrom/hp_ux110/lib/pam_afs.so.1 .
# ln -s pam_afs.so.1 pam_afs.so
```

If you use a Kerberos implementation of AFS authentication:

```
# cp /cdrom/hp_ux110/lib/pam_afs.krb.so.1 .
# ln -s pam_afs.krb.so.1 pam_afs.so
```

3. Edit the `Authentication` management section of the HP-UX PAM configuration file, `/etc/pam.conf` by convention. The entries in this section have the value `auth` in their second field.

First edit the standard entries, which refer to the HP-UX PAM module (usually, the file `/usr/lib/security/libpam_unix.1`) in their fourth field. For each service for which you want to use AFS authentication, edit the third field of its entry to read `optional`. The `pam.conf` file in the HP-UX distribution usually includes standard entries for the **login** and **ftp** services, for instance.

If there are services for which you want to use AFS authentication, but for which the `pam.conf` file does not already include a standard entry, you must create that entry and place the value `optional` in its third field. For instance, the HP-UX `pam.conf` file does not usually include standard entries for the **remsh** or **telnet** services.

Then create an AFS-related entry for each service, placing it immediately below the standard entry. The following example shows what the `Authentication` Management section looks like after you have edited or created entries for the services mentioned previously. Note that the example AFS entries appear on two lines only for legibility.

```
login  auth  optional  /usr/lib/security/libpam_unix.1
login  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root setenv_password_expires
ftp    auth  optional  /usr/lib/security/libpam_unix.1
ftp    auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
remsh  auth  optional  /usr/lib/security/libpam_unix.1
remsh  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
```

```
telnet  auth  optional  /usr/lib/security/libpam_unix.1
telnet  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass  ignore_root  setenv_password_expires
```

4. If you use the Common Desktop Environment (CDE) on the machine and want users to obtain an AFS token as they log in, also add or edit the following four entries in the `Authentication` management section. Note that the AFS-related entries appear on two lines here only for legibility.

```
dtlogin  auth  optional  /usr/lib/security/libpam_unix.1
dtlogin  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass  ignore_root
dtaction  auth  optional  /usr/lib/security/libpam_unix.1
dtaction  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass  ignore_root
```

5. Proceed to "Starting the BOS Server" on page 38 (or if referring to these instructions while installing an additional file server machine, return to "Starting Server Programs" on page 88).

Getting Started on IRIX Systems

To incorporate AFS into the kernel on IRIX systems, choose one of two methods:

- Run the AFS initialization script to invoke the `ml` program distributed by Silicon Graphics, Incorporated (SGI), which dynamically loads AFS modifications into the kernel
- Build a new static kernel

Then create partitions for storing AFS volumes. You do not need to replace the IRIX `fsck` program because SGI has already modified it to handle AFS volumes properly. If the machine is to remain an AFS client machine, verify that the IRIX login utility installed on the machine grants an AFS token.

In preparation for either dynamic loading or kernel building, perform the following procedures:

1. Mount the AFS CD-ROM for IRIX on the `/cdrom` directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your IRIX documentation. Then change directory as indicated.

```
# cd /cdrom/sgi_65/root.client
```

2. Copy the AFS initialization script to the local directory for initialization files (by convention, `/etc/init.d` on IRIX machines). Note the removal of the `.rc` extension as you copy the script.

```
# cp -p usr/vice/etc/afs.rc /etc/init.d/afs
```

3. Issue the `uname -m` command to determine the machine's CPU board type. The `IPxx` value in the output must match one of the supported CPU board types listed in the *IBM AFS Release Notes* for the current version of AFS.


```
# uname -m
```

4. Proceed to either "Loading AFS into the IRIX Kernel" on page 25 or "Building AFS into the IRIX Kernel" on page 26.

Loading AFS into the IRIX Kernel

The **ml** program is the dynamic kernel loader provided by SGI for IRIX systems. If you use it rather than building AFS modifications into a static kernel, then for AFS to function correctly the **ml** program must run each time the machine reboots. Therefore, the AFS initialization script (included on the AFS CD-ROM) invokes it automatically when the **afsm1** configuration variable is activated. In this section you activate the variable and run the script.

In later sections you verify that the script correctly initializes all AFS components, then create the links that incorporate AFS into the IRIX startup and shutdown sequence.

1. Create the local **/usr/vice/etc/sgiload** directory to house the AFS kernel library file.

```
# mkdir /usr/vice/etc/sgiload
```

2. Copy the appropriate AFS kernel library file to the **/usr/vice/etc/sgiload** directory. The **IPxx** portion of the library file name must match the value previously returned by the **uname -m** command. Also choose the file appropriate to whether the machine's kernel supports NFS server functionality (NFS must be supported for the machine to act as an NFS/AFS Translator). Single- and multiprocessor machines use the same library file.

(You can choose to copy all of the kernel library files into the **/usr/vice/etc/sgiload** directory, but they require a significant amount of space.)

If the machine's kernel supports NFS server functionality:

```
# cp -p usr/vice/etc/sgiload/libafs.IPxx.o /usr/vice/etc/sgiload
```

If the machine's kernel does not support NFS server functionality:

```
# cp -p usr/vice/etc/sgiload/libafs.IPxx.nonfs.o \
    /usr/vice/etc/sgiload
```

3. Issue the **chkconfig** command to activate the **afsm1** configuration variable.

```
# /etc/chkconfig -f afsm1 on
```

If the machine is to function as an NFS/AFS Translator and the kernel supports NFS server functionality, activate the **afsxnfs** variable.

```
# /etc/chkconfig -f afsxnfs on
```

4. Run the **/etc/init.d/afs** script to load AFS extensions into the kernel. The script invokes the **ml** command, automatically determining which kernel library file to use based on this machine's CPU type and the activation state of the **afsxnfs** variable.

You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/init.d/afs start
```

5. Proceed to "Configuring Server Partitions on IRIX Systems" on page 27.

Building AFS into the IRIX Kernel

Use the following instructions to build AFS modifications into the kernel on an IRIX system.

1. Copy the kernel initialization file **afs.sm** to the local **/var/sysgen/system** directory, and the kernel master file **afs** to the local **/var/sysgen/master.d** directory.

```
# cp -p bin/afs.sm /var/sysgen/system
# cp -p bin/afs /var/sysgen/master.d
```

2. Copy the appropriate AFS kernel library file to the local file **/var/sysgen/boot/afs.a**; the **IP_{xx}** portion of the library file name must match the value previously returned by the **uname -m** command. Also choose the file appropriate to whether the machine's kernel supports NFS server functionality (NFS must be supported for the machine to act as an NFS/AFS Translator). Single- and multiprocessor machines use the same library file.

If the machine's kernel supports NFS server functionality:

```
# cp -p bin/libafs.IPxx.a /var/sysgen/boot/afs.a
```

If the machine's kernel does not support NFS server functionality:

```
# cp -p bin/libafs.IPxx.nonfs.a /var/sysgen/boot/afs.a
```

3. Issue the **chkconfig** command to deactivate the **afsm1** configuration variable.

```
# /etc/chkconfig -f afsm1 off
```

If the machine is to function as an NFS/AFS Translator and the kernel supports NFS server functionality, activate the **afsnfs** variable.

```
# /etc/chkconfig -f afsnfs on
```

4. Copy the existing kernel file, **/unix**, to a safe location. Compile the new kernel, which is created in the file **/unix.install**. It overwrites the existing **/unix** file when the machine reboots in the next step.

```
# cp /unix /unix_noafs
# autoconfig
```

5. Reboot the machine to start using the new kernel, and login again as the superuser **root**.

```
# cd /
# shutdown -i6 -g0 -y
login: root
```

```
Password: root_password
```

Configuring Server Partitions on IRIX Systems

Every AFS file server machine must have at least one partition or logical volume dedicated to storing AFS volumes. Each server partition is mounted at a directory named **/vicep_{xx}**, where **xx** is one or two lowercase letters. The **/vicep_{xx}** directories must reside in the file server machine's root directory, not in one of its subdirectories (for example, **/usr/vicepa** is not an acceptable directory location). For additional information, see "Performing Platform-Specific Procedures" on page 9.

AFS supports use of both EFS and XFS partitions for housing AFS volumes. SGI encourages use of XFS partitions.

1. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

2. Add a line with the following format to the file systems registry file, **/etc/fstab**, for each partition (or logical volume created with the XLV volume manager) to be mounted on one of the directories created in the previous step.

For an XFS partition or logical volume:

```
/dev/dsk/disk /vicepxx xfs rw,raw=/dev/rdisk/disk 0 0
```

For an EFS partition:

```
/dev/dsk/disk /vicepxx efs rw,raw=/dev/rdisk/disk 0 0
```

The following are examples of an entry for each file system type:

```
/dev/dsk/dks0d2s6 /vicepa xfs rw,raw=/dev/rdisk/dks0d2s6 0 0
/dev/dsk/dks0d3s1 /vicepb efs rw,raw=/dev/rdisk/dks0d3s1 0 0
```

3. Create a file system on each partition that is to be mounted on a **/vicep_{xx}** directory. The following commands are probably appropriate, but consult the IRIX documentation for more information. In both cases, *raw_device* is a raw device name like **/dev/rdisk/dks0d0s0** for a single disk partition or **/dev/rxlv/xlv0** for a logical volume.

For XFS file systems, include the indicated options to configure the partition or logical volume with inodes large enough to accommodate AFS-specific information:

```
# mkfs -t xfs -i size=512 -l size=4000b raw_device
```

For EFS file systems:

```
# mkfs -t efs raw_device
```

4. Mount each partition by issuing either the **mount -a** command to mount all partitions at once or the **mount** command to mount each partition in turn.

5. **(Optional)** If you have configured partitions or logical volumes to use XFS, issue the following command to verify that the inodes are configured properly (are large enough to accommodate AFS-specific information). If the configuration is correct, the command returns no output. Otherwise, it specifies the command to run in order to configure each partition or logical volume properly.

```
# /usr/afs/bin/xfs_size_check
```

6. If you plan to retain client functionality on this machine after completing the installation, proceed to "Enabling AFS Login on IRIX Systems" on page 28. Otherwise, proceed to "Starting the BOS Server" on page 38.

Enabling AFS Login on IRIX Systems

Note: If you plan to remove client functionality from this machine after completing the installation, skip this section and proceed to "Starting the BOS Server" on page 38.

The standard IRIX command-line **login** program and the graphical **xdm** login program both automatically grant an AFS token when AFS is incorporated into the machine's kernel. However, some IRIX distributions use another login utility by default, and it does not necessarily incorporate the required AFS modifications. If that is the case, you must disable the default utility if you want AFS users to obtain AFS tokens at login. For further discussion, see the *IBM AFS Release Notes*.

If you configure the machine to use an AFS-modified login utility, then the **afsauthlib.so** and **afskauthlib.so** files (included in the AFS distribution) must reside in the **/usr/vice/etc** directory. Issue the **ls** command to verify.

```
# ls /usr/vice/etc
```

If the files do not exist, mount the AFS CD-ROM for IRIX (if it is not already), change directory as indicated, and copy them.

```
# cd /cdrom/sgi_65/root.client/usr/vice/etc
# cp -p *authlib* /usr/vice/etc
```

After taking any necessary action, proceed to "Starting the BOS Server" on page 38.

Getting Started on Linux Systems

Begin by running the AFS initialization script to call the **insmod** program, which dynamically loads AFS modifications into the kernel. Then create partitions for storing AFS volumes. You do not need to replace the Linux **fsck** program. If the machine is to remain an AFS client machine, incorporate AFS into the machine's Pluggable Authentication Module (PAM) scheme.

Loading AFS into the Linux Kernel

The **insmod** program is the dynamic kernel loader for Linux. Linux does not support incorporation of AFS modifications during a kernel build.

For AFS to function correctly, the **insmod** program must run each time the machine reboots, so the AFS initialization script (included on the AFS CD-ROM) invokes it automatically. The script also includes commands that select the appropriate AFS library file automatically. In this section you run the script.

In later sections you verify that the script correctly initializes all AFS components, then activate a configuration variable, which results in the script being incorporated into the Linux startup and shutdown sequence.

1. Mount the AFS CD-ROM for Linux on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your Linux documentation. Then change directory as indicated.

```
# cd /cdrom/i386_linux22/root.client/usr/vice/etc
```

2. Copy the AFS kernel library files to the local **/usr/vice/etc/modload** directory. The filenames for the libraries have the format **libafs-version.o**, where *version* indicates the kernel build level. The string **.mp** in the *version* indicates that the file is appropriate for machines running a multiprocessor kernel.

```
# cp -rp modload /usr/vice/etc
```

3. Copy the AFS initialization script to the local directory for initialization files (by convention, **/etc/rc.d/init.d** on Linux machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp -p afs.rc /etc/rc.d/init.d/afs
```

4. Run the AFS initialization script to load AFS extensions into the kernel. You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/rc.d/init.d/afs start
```

Configuring Server Partitions on Linux Systems

Every AFS file server machine must have at least one partition or logical volume dedicated to storing AFS volumes. Each server partition is mounted at a directory named **/vicep_{xx}**, where *xx* is one or two lowercase letters. The **/vicep_{xx}** directories must reside in the file server machine's root directory, not in one of its subdirectories (for example, **/usr/vicepa** is not an acceptable directory location). For additional information, see "Performing Platform-Specific Procedures" on page 9.

1. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

2. Add a line with the following format to the file systems registry file, **/etc/fstab**, for each directory just created. The entry maps the directory name to the disk partition to be mounted on it.

```
/dev/disk /vicepxx ext2 defaults 0 2
```

The following is an example for the first partition being configured.

```
/dev/sda8 /vicepa ext2 defaults 0 2
```

3. Create a file system on each partition that is to be mounted at a **/vicep_{xx}** directory. The following command is probably appropriate, but consult the Linux documentation for more information.

```
# mkfs -v /dev/disk
```

4. Mount each partition by issuing either the **mount -a** command to mount all partitions at once or the **mount** command to mount each partition in turn.
5. If you plan to retain client functionality on this machine after completing the installation, proceed to "Enabling AFS Login on Linux Systems" on page 30. Otherwise, proceed to "Starting the BOS Server" on page 38.

Enabling AFS Login on Linux Systems

Note: If you plan to remove client functionality from this machine after completing the installation, skip this section and proceed to "Starting the BOS Server" on page 38.

At this point you incorporate AFS into the operating system's Pluggable Authentication Module (PAM) scheme. PAM integrates all authentication mechanisms on the machine, including login, to provide the security infrastructure for authenticated access to and from the machine.

Explaining PAM is beyond the scope of this document. It is assumed that you understand the syntax and meanings of settings in the PAM configuration file (for example, how the `other` entry works, the effect of marking an entry as `required`, `optional`, or `sufficient`, and so on).

The following instructions explain how to alter the entries in the PAM configuration file for each service for which you wish to use AFS authentication. Other configurations possibly also work, but the instructions specify the recommended and tested configuration.

The recommended AFS-related entries in the PAM configuration file make use of one or more of the following three attributes.

try_first_pass

This is a standard PAM attribute that can be included on entries after the first one for a service; it directs the module to use the password that was provided to the first module. For the AFS module, it means that AFS authentication succeeds if the password provided to the module listed first is the user's correct AFS password. For further discussion of this attribute and its alternatives, see the operating system's PAM documentation.

ignore_root

This attribute, specific to the AFS PAM module, directs it to ignore not only the local superuser **root**, but also any user with UID 0 (zero).

setenv_password_expires

This attribute, specific to the AFS PAM module, sets the environment variable **PASSWORD_EXPIRES** to the expiration date of the user's AFS password, which is recorded in the Authentication Database.

Perform the following steps to enable AFS login.

1. Mount the AFS CD-ROM for Linux on the **/cdrom** directory, if it is not already. Then change to the directory for PAM modules, which depends on which Linux distribution you are using.

If you are using a Linux distribution from Red Hat Software:

```
# cd /lib/security
```

If you are using another Linux distribution:

```
# cd /usr/lib/security
```

2. Copy the appropriate AFS authentication library file to the directory to which you changed in the previous step. Create a symbolic link whose name does not mention the version. Omitting the version eliminates the need to edit the PAM configuration file if you later update the library file.

If you use the AFS Authentication Server (**kaserver** process):

```
# cp /cdrom/i386_linux22/lib/pam_afs.so.1 .
# ln -s pam_afs.so.1 pam_afs.so
```

If you use a Kerberos implementation of AFS authentication:

```
# cp /cdrom/i386_linux22/lib/pam_afs.krb.so.1 .
# ln -s pam_afs.krb.so.1 pam_afs.so
```

3. For each service with which you want to use AFS authentication, insert an entry for the AFS PAM module into the **auth** section of the service's PAM configuration file. (Linux uses a separate configuration file for each service, unlike some other operating systems which list all services in a single file.) Mark the entry as **sufficient** in the second field.

Place the AFS entry below any entries that impose conditions under which you want the service to fail for a user who does not meet the entry's requirements. Mark these entries **required**. Place the AFS entry above any entries that need to execute only if AFS authentication fails.

Insert the following AFS entry if using the Red Hat distribution:

```
auth sufficient /lib/security/pam_afs.so try_first_pass ignore_root
```

Insert the following AFS entry if using another distribution:

```
auth sufficient /usr/lib/security/pam_afs.so try_first_pass ignore_root
```

The following example illustrates the recommended configuration of the configuration file for the **login** service (**/etc/pam.d/login**) on a machine using the Red Hat distribution.

```
##PAM-1.0
auth      required    /lib/security/pam_securetty.so
auth      required    /lib/security/pam_nologin.so
auth      sufficient  /lib/security/pam_afs.so try_first_pass ignore_root
auth      required    /lib/security/pam_pwdb.so shadow nullok
account   required    /lib/security/pam_pwdb.so
password  required    /lib/security/pam_cracklib.so
password  required    /lib/security/pam_pwdb.so shadow nullok use_authok
session   required    /lib/security/pam_pwdb.so
```

4. Proceed to "Starting the BOS Server" on page 38 (or if referring to these instructions while installing an additional file server machine, return to "Starting Server Programs" on page 88).

Getting Started on Solaris Systems

Begin by running the AFS initialization script to call the **modload** program distributed by Sun Microsystems, which dynamically loads AFS modifications into the kernel. Then create partitions for storing AFS volumes, and install and configure the AFS-modified **fsck** program to run on AFS server partitions. If the machine is to remain an AFS client machine, incorporate AFS into the machine's Pluggable Authentication Module (PAM) scheme.

Loading AFS into the Solaris Kernel

The **modload** program is the dynamic kernel loader provided by Sun Microsystems for Solaris systems. Solaris does not support incorporation of AFS modifications during a kernel build.

For AFS to function correctly, the **modload** program must run each time the machine reboots, so the AFS initialization script (included on the AFS CD-ROM) invokes it automatically. In this section you copy the appropriate AFS library file to the location where the **modload** program accesses it and then run the script.

In later sections you verify that the script correctly initializes all AFS components, then create the links that incorporate AFS into the Solaris startup and shutdown sequence.

1. Mount the AFS CD-ROM for Solaris on the **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your Solaris documentation. Then change directory as indicated.

```
# cd /cdrom/sun4x_56/root.client/usr/vice/etc
```

2. Copy the AFS initialization script to the local directory for initialization files (by convention, **/etc/init.d** on Solaris machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp -p afs.rc /etc/init.d/afs
```


- Copy the appropriate AFS kernel library file to the local file `/kernel/fs/afs`.

If the machine is running Solaris 2.6 or the 32-bit version of Solaris 7, its kernel supports NFS server functionality, and the `nfsd` process is running:

```
# cp -p modload/libafs.o /kernel/fs/afs
```

If the machine is running Solaris 2.6 or the 32-bit version of Solaris 7, and its kernel does not support NFS server functionality or the `nfsd` process is not running:

```
# cp -p modload/libafs.nonfs.o /kernel/fs/afs
```

If the machine is running the 64-bit version of Solaris 7, its kernel supports NFS server functionality, and the `nfsd` process is running:

```
# cp -p modload/libafs64.o /kernel/fs/sparcv9/afs
```

If the machine is running the 64-bit version of Solaris 7, and its kernel does not support NFS server functionality or the `nfsd` process is not running:

```
# cp -p modload/libafs64.nonfs.o /kernel/fs/sparcv9/afs
```

- Run the AFS initialization script to load AFS modifications into the kernel. You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/init.d/afs start
```

When an entry called `afs` does not already exist in the local `/etc/name_to_sysnum` file, the script automatically creates it and reboots the machine to start using the new version of the file. If this happens, log in again as the superuser `root` after the reboot and run the initialization script again. This time the required entry exists in the `/etc/name_to_sysnum` file, and the `modload` program runs.

```
login: root
Password: root_password
# /etc/init.d/afs start
```

Configuring the AFS-modified fsck Program on Solaris Systems

In this section, you make modifications to guarantee that the appropriate `fsck` program runs on AFS server partitions. The `fsck` program provided with the operating system must never run on AFS server partitions. Because it does not recognize the structures that the File Server uses to organize volume data, it removes all of the data. To repeat:

Never run the standard `fsck` program on AFS server partitions. It discards AFS volumes.

- Create the `/usr/lib/fs/afs` directory to house the AFS-modified `fsck` program and related files.

```
# mkdir /usr/lib/fs/afs
# cd /usr/lib/fs/afs
```

- Copy the `vfsc` binary to the newly created directory, changing the name as you do so.

```
# cp /cdrom/sun4x_56/root.server/etc/vfsc fsck
```

3. Working in the `/usr/lib/fs/afs` directory, create the following links to Solaris libraries:

```
# ln -s /usr/lib/fs/ufs/clri
# ln -s /usr/lib/fs/ufs/df
# ln -s /usr/lib/fs/ufs/edquota
# ln -s /usr/lib/fs/ufs/ff
# ln -s /usr/lib/fs/ufs/fsdb
# ln -s /usr/lib/fs/ufs/fsirand
# ln -s /usr/lib/fs/ufs/fstyp
# ln -s /usr/lib/fs/ufs/labelit
# ln -s /usr/lib/fs/ufs/lockfs
# ln -s /usr/lib/fs/ufs/mkfs
# ln -s /usr/lib/fs/ufs/mount
# ln -s /usr/lib/fs/ufs/ncheck
# ln -s /usr/lib/fs/ufs/newfs
# ln -s /usr/lib/fs/ufs/quot
# ln -s /usr/lib/fs/ufs/quota
# ln -s /usr/lib/fs/ufs/quotaoff
# ln -s /usr/lib/fs/ufs/quotaon
# ln -s /usr/lib/fs/ufs/repquota
# ln -s /usr/lib/fs/ufs/tunefs
# ln -s /usr/lib/fs/ufs/ufsdump
# ln -s /usr/lib/fs/ufs/ufsrestore
# ln -s /usr/lib/fs/ufs/volcopy
```

4. Append the following line to the end of the file `/etc/dfs/fstypes`.

```
afs AFS Utilities
```

5. Edit the `/sbin/mountall` file, making two changes.

- Add an entry for AFS to the `case` statement for option 2, so that it reads as follows:

```
case "$2" in
ufs)    foptions="-o p"
        ;;
afs)    foptions="-o p"
        ;;
s5)    foptions="-y -t /var/tmp/tmp$$ -D"
        ;;
*)     foptions="-y"
        ;;
```

- Edit the file so that all AFS and UFS partitions are checked in parallel. Replace the following section of code:

```
# For fsck purposes, we make a distinction between ufs and
# other file systems
#
if [ "$fstype" = "ufs" ]; then
    ufs_fscklist="$ufs_fscklist $fsckdev"
    saveentry $fstype "$OPTIONS" $special $mountp
    continue
fi
```

with the following section of code:

```

# For fsck purposes, we make a distinction between ufs/afs
# and other file systems.
#
if [ "$fstype" = "ufs" -o "$fstype" = "afs" ]; then
    ufs_fscklist="$ufs_fscklist $fsckdev"
    saveentry $fstype "$OPTIONS" $special $mountp
    continue
fi

```

Configuring Server Partitions on Solaris Systems

Every AFS file server machine must have at least one partition or logical volume dedicated to storing AFS volumes. Each server partition is mounted at a directory named **/vicep_{xx}**, where **xx** is one or two lowercase letters. The **/vicep_{xx}** directories must reside in the file server machine's root directory, not in one of its subdirectories (for example, **/usr/vicepa** is not an acceptable directory location). For additional information, see "Performing Platform-Specific Procedures" on page 9.

1. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

2. Add a line with the following format to the file systems registry file, **/etc/vfstab**, for each partition to be mounted on a directory created in the previous step. Note the value **afs** in the fourth field, which tells Solaris to use the AFS-modified **fsck** program on this partition.

```
/dev/dsk/disk /dev/rdisk/disk /vicepxx afs boot_order yes
```

The following is an example for the first partition being configured.

```
/dev/dsk/c0t6d0s1 /dev/rdisk/c0t6d0s1 /vicepa afs 3 yes
```

3. Create a file system on each partition that is to be mounted at a **/vicep_{xx}** directory. The following command is probably appropriate, but consult the Solaris documentation for more information.

```
# newfs -v /dev/rdisk/disk
```

4. Issue the **mountall** command to mount all partitions at once.
5. If you plan to retain client functionality on this machine after completing the installation, proceed to "Enabling AFS Login and Editing the File Systems Clean-up Script on Solaris Systems" on page 35. Otherwise, proceed to "Starting the BOS Server" on page 38.

Enabling AFS Login and Editing the File Systems Clean-up Script on Solaris Systems

Note: If you plan to remove client functionality from this machine after completing the installation, skip this section and proceed to "Starting the BOS Server" on page 38.

At this point you incorporate AFS into the operating system's Pluggable Authentication Module (PAM) scheme. PAM integrates all authentication mechanisms on the machine, including login, to provide the security infrastructure for authenticated access to and from the machine.

Explaining PAM is beyond the scope of this document. It is assumed that you understand the syntax and meanings of settings in the PAM configuration file (for example, how the `other` entry works, the effect of marking an entry as `required`, `optional`, or `sufficient`, and so on).

The following instructions explain how to alter the entries in the PAM configuration file for each service for which you wish to use AFS authentication. Other configurations possibly also work, but the instructions specify the recommended and tested configuration.

Note: The instructions specify that you mark each entry as `optional`. However, marking some modules as `optional` can mean that they grant access to the corresponding service even when the user does not meet all of the module's requirements. In some operating system revisions, for example, if you mark as `optional` the module that controls login via a dial-up connection, it allows users to login without providing a password. See the *IBM AFS Release Notes* for a discussion of any limitations that apply to this operating system.

Also, with some operating system versions you must install patches for PAM to interact correctly with certain authentication programs. For details, see the *IBM AFS Release Notes*.

The recommended AFS-related entries in the PAM configuration file make use of one or more of the following three attributes.

try_first_pass

This is a standard PAM attribute that can be included on entries after the first one for a service; it directs the module to use the password that was provided to the first module. For the AFS module, it means that AFS authentication succeeds if the password provided to the module listed first is the user's correct AFS password. For further discussion of this attribute and its alternatives, see the operating system's PAM documentation.

ignore_root

This attribute, specific to the AFS PAM module, directs it to ignore not only the local superuser `root`, but also any user with UID 0 (zero).

setenv_password_expires

This attribute, specific to the AFS PAM module, sets the environment variable `PASSWORD_EXPIRES` to the expiration date of the user's AFS password, which is recorded in the Authentication Database.

Perform the following steps to enable AFS login.

1. Mount the AFS CD-ROM for Solaris on the `/cdrom` directory, if it is not already. Then change directory as indicated.

```
# cd /usr/lib/security
```

2. Copy the AFS authentication library file to the `/usr/lib/security` directory. Then create a symbolic link to it whose name does not mention the version. Omitting the version eliminates the need to edit the PAM configuration file if you later update the library file.

If you use the AFS Authentication Server (**kaserver** process):

```
# cp /cdrom/sun4x_56/lib/pam_afs.so.1 .
# ln -s pam_afs.so.1 pam_afs.so
```

If you use a Kerberos implementation of AFS authentication:

```
# cp /cdrom/sun4x_56/lib/pam_afs.krb.so.1 .
# ln -s pam_afs.krb.so.1 pam_afs.so
```

3. Edit the `Authentication` management section of the Solaris PAM configuration file, `/etc/pam.conf` by convention. The entries in this section have the value `auth` in their second field.

First edit the standard entries, which refer to the Solaris PAM module (usually, the file `/usr/lib/security/pam_unix.so.1`) in their fourth field. For each service for which you want to use AFS authentication, edit the third field of its entry to read `optional`. The `pam.conf` file in the Solaris distribution usually includes standard entries for the **login**, **rlogin**, and **rsh** services, for instance.

If there are services for which you want to use AFS authentication, but for which the `pam.conf` file does not already include a standard entry, you must create that entry and place the value `optional` in its third field. For instance, the Solaris `pam.conf` file does not usually include standard entries for the **ftp** or **telnet** services.

Then create an AFS-related entry for each service, placing it immediately below the standard entry. The following example shows what the `Authentication` Management section looks like after you have edited or created entries for the services mentioned previously. Note that the example AFS entries appear on two lines only for legibility.

```
login  auth  optional  /usr/lib/security/pam_unix.so.1
login  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root setenv_password_expires
rlogin auth  optional  /usr/lib/security/pam_unix.so.1
rlogin auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root setenv_password_expires
rsh    auth  optional  /usr/lib/security/pam_unix.so.1
rsh    auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
ftp    auth  optional  /usr/lib/security/pam_unix.so.1
ftp    auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
telnet auth  optional  /usr/lib/security/pam_unix.so.1
```

```
telnet  auth  optional  /usr/lib/security/pam_afs.so      \  
      try_first_pass  ignore_root  setenv_password_expires
```

4. If you use the Common Desktop Environment (CDE) on the machine and want users to obtain an AFS token as they log in, also add or edit the following four entries in the `Authentication` management section. Note that the AFS-related entries appear on two lines here only for legibility.

```
dtlogin  auth  optional  /usr/lib/security/pam_unix.so.1  
dtlogin  auth  optional  /usr/lib/security/pam_afs.so      \  
      try_first_pass  ignore_root  
dtsession  auth  optional  /usr/lib/security/pam_unix.so.1  
dtsession  auth  optional  /usr/lib/security/pam_afs.so      \  
      try_first_pass  ignore_root
```

5. Some Solaris distributions include a script that locates and removes unneeded files from various file systems. Its conventional location is `/usr/lib/fs/nfs/nfsfind`. The script generally uses an argument to the `find` command to define which file systems to search. In this step you modify the command to exclude the `/afs` directory. Otherwise, the command traverses the AFS filesystem of every cell that is accessible from the machine, which can take many hours. The following alterations are possibilities, but you must verify that they are appropriate for your cell.

The first possible alteration is to add the `-local` flag to the existing command, so that it looks like the following:

```
find $dir -local -name .nfs\* -mtime +7 -mount -exec rm -f {} \;
```

Another alternative is to exclude any directories whose names begin with the lowercase letter `a` or a non-alphabetic character.

```
find /[A-Zb-z]* remainder of existing command
```

Do not use the following command, which still searches under the `/afs` directory, looking for a subdirectory of type `4.2`.

```
find / -fstype 4.2 /* do not use */
```

6. Proceed to "Starting the BOS Server" on page 38 (or if referring to these instructions while installing an additional file server machine, return to "Starting Server Programs" on page 88).

Starting the BOS Server

You are now ready to start the AFS server processes on this machine. Begin by copying the AFS server binaries from the CD-ROM to the conventional local disk location, the `/usr/afs/bin` directory. The following instructions also create files in other subdirectories of the `/usr/afs` directory.

Then issue the `bosserv` command to initialize the Basic OverSeer (BOS) Server, which monitors and controls other AFS server processes on its server machine. Include the `-noauth` flag to disable authorization checking. Because you have not yet configured your cell's AFS authentication and authorization mechanisms, the BOS Server cannot perform authorization checking as it does during

normal operation. In no-authorization mode, it does not verify the identity or privilege of the issuer of a **bos** command, and so performs any operation for anyone.

Disabling authorization checking gravely compromises cell security. You must complete all subsequent steps in one uninterrupted pass and must not leave the machine unattended until you restart the BOS Server with authorization checking enabled, in "Verifying the AFS Initialization Script" on page 55.

As it initializes for the first time, the BOS Server creates the following directories and files, setting the owner to the local superuser **root** and the mode bits to limit the ability to write (and in some cases, read) them. For a description of the contents and function of these directories and files, see the chapter in the *IBM AFS Administration Guide* about administering server machines. For further discussion of the mode bit settings, see "Protecting Sensitive AFS Directories" on page 70.

- **/usr/afs/db**
- **/usr/afs/etc/CellServDB**
- **/usr/afs/etc/ThisCell**
- **/usr/afs/local**
- **/usr/afs/logs**

The BOS Server also creates symbolic links called **/usr/vice/etc/ThisCell** and **/usr/vice/etc/CellServDB** to the corresponding files in the **/usr/afs/etc** directory. The AFS command interpreters consult the **CellServDB** and **ThisCell** files in the **/usr/vice/etc** directory because they generally run on client machines. On machines that are AFS servers only (as this machine currently is), the files reside only in the **/usr/afs/etc** directory; the links enable the command interpreters to retrieve the information they need. Later instructions for installing the client functionality replace the links with actual files.

1. On the local **/cdrom** directory, mount the AFS CD-ROM for this machine's system type, if it is not already. For instructions on mounting CD-ROMs (either locally or remotely via NFS), consult the operating system documentation.
2. Copy files from the CD-ROM to the local **/usr/afs** directory.

```
# cd /cdrom/sysname/root.server/usr/afs
# cp -rp * /usr/afs
```

3. Issue the **bosserver** command. Include the **-noauth** flag to disable authorization checking.

```
# /usr/afs/bin/bosserver -noauth &
```

4. Verify that the BOS Server created **/usr/vice/etc/ThisCell** and **/usr/vice/etc/CellServDB** as symbolic links to the corresponding files in the **/usr/afs/etc** directory.

```
# ls -l /usr/vice/etc
```

If either or both of **/usr/vice/etc/ThisCell** and **/usr/vice/etc/CellServDB** do not exist, or are not links, issue the following commands.

```
# cd /usr/vice/etc
# ln -s /usr/afs/etc/ThisCell
# ln -s /usr/afs/etc/CellServDB
```

Defining Cell Name and Membership for Server Processes

Now assign your cell's name. The chapter in the *IBM AFS Administration Guide* about cell configuration and administration issues discusses the important considerations, explains why changing the name is difficult, and outlines the restrictions on name format. Two of the most important restrictions are that the name cannot include uppercase letters or more than 64 characters.

Use the **bos setcellname** command to assign the cell name. It creates two files:

- **/usr/afs/etc/ThisCell**, which defines this machine's cell membership
- **/usr/afs/etc/CellServDB**, which lists the cell's database server machines; the machine named on the command line is placed on the list automatically

Note: In the following and every instruction in this guide, for the *machine name* argument substitute the fully-qualified hostname (such as **fs1.abc.com**) of the machine you are installing. For the *cell name* argument substitute your cell's complete name (such as **abc.com**).

1. Issue the **bos setcellname** command to set the cell name.

```
# cd /usr/afs/bin
# ./bos setcellname <machine name> <cell name> -noauth
```

Because you are not authenticated and authorization checking is disabled, the **bos** command interpreter possibly produces error messages about being unable to obtain tickets and running unauthenticated. You can safely ignore the messages.

2. Issue the **bos listhosts** command to verify that the machine you are installing is now registered as the cell's first database server machine.

```
# ./bos listhosts <machine name> -noauth
Cell name is cell_name
Host 1 is machine_name
```

Starting the Database Server Processes

Next use the **bos create** command to create entries for the four database server processes in the **/usr/afs/local/BosConfig** file and start them running. The four processes run on database server machines only:

- The Authentication Server (the **kaserver** process) maintains the Authentication Database

- The Backup Server (the **buserver** process) maintains the Backup Database
- The Protection Server (the **ptserver** process) maintains the Protection Database
- The Volume Location (VL) Server (the **vlserver** process) maintains the Volume Location Database (VLDB)

Note: AFS's authentication and authorization software is based on algorithms and other procedures known as *Kerberos*, as originally developed by Project Athena at the Massachusetts Institute of Technology. Some cells choose to replace the AFS Authentication Server and other security-related protocols with Kerberos as obtained directly from Project Athena or other sources. If you wish to do this, contact the AFS Product Support group now to learn about necessary modifications to the installation.

The remaining instructions in this chapter include the **-cell** argument on all applicable commands. Provide the cell name you assigned in "Defining Cell Name and Membership for Server Processes" on page 40. If a command appears on multiple lines, it is only for legibility.

1. Issue the **bos create** command to start the Authentication Server. The current working directory is still **/usr/afs/bin**.

```
# ./bos create <machine name> kserver simple /usr/afs/bin/kserver \
    -cell <cell name> -noauth
```

You can safely ignore the messages that tell you to add Kerberos to the **/etc/services** file; AFS uses a default value that makes the addition unnecessary. You can also ignore messages about the failure of authentication.

2. Issue the **bos create** command to start the Backup Server.

```
# ./bos create <machine name> buserver simple /usr/afs/bin/buserver \
    -cell <cell name> -noauth
```

3. Issue the **bos create** command to start the Protection Server.

```
# ./bos create <machine name> ptserver simple /usr/afs/bin/ptserver \
    -cell <cell name> -noauth
```

4. Issue the **bos create** command to start the VL Server.

```
# ./bos create <machine name> vlserver simple /usr/afs/bin/vlserver \
    -cell <cell name> -noauth
```

Initializing Cell Security

Now initialize the cell's security mechanisms. Begin by creating the following two initial entries in the Authentication Database:

- A generic administrative account, called **admin** by convention. If you choose to assign a different name, substitute it throughout the remainder of this document.

After you complete the installation of the first machine, you can continue to have all administrators use the **admin** account, or you can create a separate administrative account for each of them. The latter scheme implies somewhat more overhead, but provides a more informative audit trail for administrative operations.

- The entry for AFS server processes, called **afs**. No user logs in under this identity, but the Authentication Server's Ticket Granting Service (TGS) module uses the associated key to encrypt the server tickets that it grants to AFS clients for presentation to server processes during mutual authentication. (The chapter in the *IBM AFS Administration Guide* about cell configuration and administration describes the role of server encryption keys in mutual authentication.)

In Step "7" on page 43, you also place the initial AFS server encryption key into the **/usr/afs/etc/KeyFile** file. The AFS server processes refer to this file to learn the server encryption key when they need to decrypt server tickets.

You also issue several commands that enable the new **admin** user to issue privileged commands in all of the AFS suites.

The following instructions do not configure all of the security mechanisms related to the AFS Backup System. See the chapter in the *IBM AFS Administration Guide* about configuring the Backup System.

1. Enter **kas** interactive mode. Because the machine is in no-authorization checking mode, include the **-noauth** flag to suppress the Authentication Server's usual prompt for a password.

```
# kas -cell <cell name> -noauth
ka>
```

2. Issue the **kas create** command to create Authentication Database entries called **admin** and **afs**.

Do not provide passwords on the command line. Instead provide them as *afs_passwd* and *admin_passwd* in response to the **kas** command interpreter's prompts as shown, so that they do not appear on the standard output stream.

You need to enter the *afs_passwd* string only in this step and in Step "7" on page 43, so provide a value that is as long and complex as possible, preferably including numerals, punctuation characters, and both uppercase and lowercase letters. Also make the *admin_passwd* as long and complex as possible, but keep in mind that administrators need to enter it often. Both passwords must be at least six characters long.

```
ka> create afs
initial_password: afs_passwd
Verifying, please re-enter initial_password: afs_passwd
```

```
ka> create admin
initial_password: admin_passwd
Verifying, please re-enter initial_password: admin_passwd
```

3. Issue the **kas examine** command to display the **afs** entry. The output includes a checksum generated by encrypting a constant with the server encryption key derived from the *afs_passwd* string. In Step "8" on page 43 you issue the **bos listkeys** command to verify that the checksum in its output matches the checksum in this output.

```
ka> examine afs
User data for afs
key (0) cksum is checksum . . .
```

4. Issue the **kas setfields** command to turn on the `ADMIN` flag in the **admin** entry. This enables the **admin** user to issue privileged **kas** commands. Then issue the **kas examine** command to verify that the `ADMIN` flag appears in parentheses on the first line of the output, as shown in the example.

```
ka> setfields admin -flags admin
ka> examine admin
User data for admin (ADMIN) . . .
```

5. Issue the **kas quit** command to leave **kas** interactive mode.

```
ka> quit
```

6. Issue the **bos adduser** command to add the **admin** user to the `/usr/afs/etc/UserList` file. This enables the **admin** user to issue privileged **bos** and **vos** commands.

```
# ./bos adduser <machine name> admin -cell <cell name> -noauth
```

7. Issue the **bos addkey** command to define the AFS server encryption key in the `/usr/afs/etc/KeyFile` file.

Do not provide the password on the command line. Instead provide it as *afs_passwd* in response to the **bos** command interpreter's prompts, as shown. Provide the same string as in Step "2" on page 42.

```
# ./bos addkey <machine name> -kvno 0 -cell <cell name> -noauth
Input key: afs_passwd
Retype input key: afs_passwd
```

8. Issue the **bos listkeys** command to verify that the checksum for the new key in the **KeyFile** file is the same as the checksum for the key in the Authentication Database's **afs** entry, which you displayed in Step "3" on page 43.

```
# ./bos listkeys <machine name> -cell <cell name> -noauth
key 0 has cksum checksum
```

You can safely ignore any error messages indicating that **bos** failed to get tickets or that authentication failed.

If the keys are different, issue the following commands, making sure that the *afs_passwd* string is the same in each case. The *checksum* strings reported by the **kas examine** and **bos listkeys**

commands must match; if they do not, repeat these instructions until they do, using the **-kvno** argument to increment the key version number each time.

```
# ./kas -cell <cell name> -noauth
ka> setpassword afs -kvno 1
new_password: afs_passwd
Verifying, please re-enter initial_password: afs_passwd
ka> examine afs
User data for afs
  key (1) cksum is checksum . . .
ka> quit
# ./bos addkey <machine name> -kvno 1 -cell <cell name> -noauth
Input key: afs_passwd
Retype input key: afs_passwd
# ./bos listkeys <machine name> -cell <cell name> -noauth
key 1 has cksum checksum
```

9. Issue the **pts createuser** command to create a Protection Database entry for the **admin** user.

By default, the Protection Server assigns AFS UID 1 (one) to the **admin** user, because it is the first user entry you are creating. If the local password file (**/etc/passwd** or equivalent) already has an entry for **admin** that assigns it a UNIX UID other than 1, it is best to use the **-id** argument to the **pts createuser** command to make the new AFS UID match the existing UNIX UID. Otherwise, it is best to accept the default.

```
# ./pts createuser -name admin -cell <cell name> [-id <AFS UID>] -noauth
User admin has id AFS UID
```

10. Issue the **pts adduser** command to make the **admin** user a member of the **system:administrators** group, and the **pts membership** command to verify the new membership. Membership in the group enables the **admin** user to issue privileged **pts** commands and some privileged **fs** commands.

```
# ./pts adduser admin system:administrators -cell <cell name> -noauth
# ./pts membership admin -cell <cell name> -noauth
Groups admin (id: 1) is a member of:
  system:administrators
```

11. Issue the **bos restart** command with the **-all** flag to restart the database server processes, so that they start using the new server encryption key.

```
# ./bos restart <machine name> -all -cell <cell name> -noauth
```

Starting the File Server, Volume Server, and Salvager

Start the **fs** process, which consists of the File Server, Volume Server, and Salvager (**fileserver**, **volserver** and **salvager** processes).

1. Issue the **bos create** command to start the **fs** process. The command appears here on multiple lines only for legibility.

```
# ./bos create <machine name> fs fs /usr/afs/bin/fileserver \
                /usr/afs/bin/volserver /usr/afs/bin/salvager \
                -cell <cell name> -noauth
```

Sometimes a message about Volume Location Database (VLDB) initialization appears, along with one or more instances of an error message similar to the following:

```
FSYNC_clientInit temporary failure (will retry)
```

This message appears when the **volserver** process tries to start before the **fileserver** process has completed its initialization. Wait a few minutes after the last such message before continuing, to guarantee that both processes have started successfully.

You can verify that the **fs** process has started successfully by issuing the **bos status** command. Its output mentions two `proc starts`.

```
# ./bos status <machine name> fs -long -noauth
```

2. Your next action depends on whether you have ever run AFS file server machines in the cell:
 - If you are installing the first AFS server machine ever in the cell (that is, you are not upgrading the AFS software from a previous version), create the first AFS volume, **root.afs**.

For the *partition name* argument, substitute the name of one of the machine's AFS server partitions (such as **/vicepa**).

```
# ./vos create <machine name> <partition name> root.afs \
                -cell <cell name> -noauth
```

The Volume Server produces a message confirming that it created the volume on the specified partition. You can ignore error messages indicating that tokens are missing, or that authentication failed.

- If there are existing AFS file server machines and volumes in the cell, issue the **vos syncvldb** and **vos syncserv** commands to synchronize the VLDB with the actual state of volumes on the local machine. To follow the progress of the synchronization operation, which can take several minutes, use the **-verbose** flag.

```
# ./vos syncvldb <machine name> -cell <cell name> -verbose -noauth
# ./vos syncserv <machine name> -cell <cell name> -verbose -noauth
```

You can ignore error messages indicating that tokens are missing, or that authentication failed.

Starting the Server Portion of the Update Server

Start the server portion of the Update Server (the **upserver** process), to distribute the contents of directories on this machine to other server machines in the cell. It becomes active when you configure the client portion of the Update Server on additional server machines.

Distributing the contents of its **/usr/afs/etc** directory makes this machine the cell's *system control machine*. The other server machines in the cell run the **upclientetc** process (an instance of the client portion of the Update Server) to retrieve the configuration files. Use the **-crypt** argument to the **upserver** initialization command to specify that the Update Server distributes the contents of the **/usr/afs/etc** directory only in encrypted form, as shown in the following instruction. Several of the files in the directory, particularly the **KeyFile** file, are crucial to cell security and so must never cross the network unencrypted.

(You can choose not to configure a system control machine, in which case you must update the configuration files in each server machine's **/usr/afs/etc** directory individually. The **bos** commands used for this purpose also encrypt data before sending it across the network.)

Distributing the contents of its **/usr/afs/bin** directory to other server machines of its system type makes this machine a *binary distribution machine*. The other server machines of its system type run the **upclientbin** process (an instance of the client portion of the Update Server) to retrieve the binaries.

The binaries in the **/usr/afs/bin** directory are not sensitive, so it is not necessary to encrypt them before transfer across the network. Include the **-clear** argument to the **upserver** initialization command to specify that the Update Server distributes the contents of the **/usr/afs/bin** directory in unencrypted form unless an **upclientbin** process requests encrypted transfer.

Note that the server and client portions of the Update Server always mutually authenticate with one another, regardless of whether you use the **-clear** or **-crypt** arguments. This protects their communications from eavesdropping to some degree.

For more information on the **upclient** and **upserver** processes, see their reference pages in the *IBM AFS Administration Reference*. The commands appear on multiple lines here only for legibility.

1. Issue the **bos create** command to start the **upserver** process.

```
# ./bos create <machine name> upserver simple \
    "/usr/afs/bin/upserver -crypt /usr/afs/etc \
    -clear /usr/afs/bin" -cell <cell name> -noauth
```

Starting the Controller for NTPD

Keeping the clocks on all server and client machines in your cell synchronized is crucial to several functions, and in particular to the correct operation of AFS's distributed database technology, Ubik. The chapter in the *IBM AFS Administration Guide* about administering server machines explains how time skew can disturb Ubik's performance and cause service outages in your cell.

The AFS distribution includes a version of the Network Time Protocol Daemon (NTPD) for synchronizing the clocks on server machines. If a time synchronization program is not already running on the machine, then in this section you start the **runntp** process to configure NTPD for use with AFS.

Note: Do not run the **runntp** process if NTPD or another time synchronization protocol is already running on the machine. Some versions of some operating systems run a time synchronization program by default, as detailed in the *IBM AFS Release Notes*.

Attempting to run multiple instances of the NTPD causes an error. Running NTPD together with another time synchronization protocol is unnecessary and can cause instability in the clock setting.

If you run the **runntp** process and your cell has reliable network connectivity to machines outside your cell, then it is conventional to configure the first AFS machine to refer to a time source outside the cell. When you later install the **runntp** program on other server machines in the cell, it configures NTPD to choose a time source at random from among the database server machines listed in the **/usr/afs/etc/CellServDB** file. Time synchronization therefore works in a chained manner: this database server machine refers to a time source outside the cell, the database server machines refer to the machine among them that has access to the most accurate time (NTPD itself includes code for determining this), and each non-database server machine refers to a local database server machine chosen at random from the **/usr/afs/etc/CellServDB** file. If you ever decide to remove database server functionality from this machine, it is best to transfer responsibility for consulting an external time source to a remaining database server machine.

If your cell does not have network connectivity to external machines, or if the connectivity is not reliable, include the **-localclock** flag to the **runntp** command as indicated in the following instructions. The flag tells NTPD to rely on the machine's internal clock when all external time sources are inaccessible. The **runntp** command has other arguments that are possibly useful given your cell configuration; see the *IBM AFS Administration Reference*.

Choosing an appropriate external time source is important, but involves more considerations than can be discussed here. If you need help in selecting a source, contact the AFS Product Support group.

As the **runntp** process initializes NTPD, trace messages sometimes appear on the standard output stream. You can ignore them, but they can be informative if you understand how NTPD works.

1. Issue the **bos create** command to start the **runntp** process. For the *host* argument, substitute the fully-qualified hostname or IP address of one or more machines outside the cell that are to serve as time sources. Separate each name with a space.
 - If your cell usually has reliable network connectivity to an external time source, use the following command:


```
# ./bos create <machine name> runntp simple \
  "/usr/afs/bin/runntp <host>+" -cell <cell name> -noauth
```
 - If your cell does not have network connectivity to an external time source, use the following command:


```
# ./bos create <machine name> runntp simple \
  "/usr/afs/bin/runntp -localclock" -cell <cell name> -noauth
```
 - If your cell has network connectivity to an external time source, but the network connection is frequently interrupted, use the following command:


```
# ./bos create <machine name> runntp simple \
  "/usr/afs/bin/runntp -localclock <host>+" \
```

```
-cell <cell name> -noauth
```

Overview: Installing Client Functionality

The machine you are installing is now an AFS file server machine, database server machine, system control machine, and binary distribution machine. Now make it a client machine by completing the following tasks:

1. Define the machine's cell membership for client processes
2. Create the client version of the **CellServDB** file
3. Define cache location and size
4. Create the **/afs** directory and start the Cache Manager

Copying Client Files to the Local Disk

Before installing and configuring the AFS client, copy the necessary files from the AFS CD-ROM to the local **/usr/vice/etc** directory.

1. On the local **/cdrom** directory, mount the AFS CD-ROM for this machine's system type, if it is not already. For instructions on mounting CD-ROMs (either locally or remotely via NFS), consult the operating system documentation.
2. Copy files to the local **/usr/vice/etc** directory.

This step places a copy of the AFS initialization script (and related files, if applicable) into the **/usr/vice/etc** directory. In the preceding instructions for incorporating AFS into the kernel, you copied the script directly to the operating system's conventional location for initialization files. When you incorporate AFS into the machine's startup sequence in a later step, you can choose to link the two files.

On some system types that use a dynamic kernel loader program, you previously copied AFS library files into a subdirectory of the **/usr/vice/etc** directory. On other system types, you copied the appropriate AFS library file directly to the directory where the operating system accesses it. The following commands do not copy or recopy the AFS library files into the **/usr/vice/etc** directory, because on some system types the library files consume a large amount of space. If you want to copy them, add the **-r** flag to the first **cp** command and skip the second **cp** command.

```
# cd /cdrom/sysname/root.client/usr/vice/etc
# cp -p * /usr/vice/etc
# cp -rp C /usr/vice/etc
```


Defining Cell Membership for Client Processes

Every AFS client machine has a copy of the `/usr/vice/etc/ThisCell` file on its local disk to define the machine's cell membership for the AFS client programs that run on it. The **ThisCell** file you created in the `/usr/afs/etc` directory (in "Defining Cell Name and Membership for Server Processes" on page 40) is used only by server processes.

Among other functions, the **ThisCell** file on a client machine determines the following:

- The cell in which users authenticate when they log onto the machine, assuming it is using an AFS-modified login utility
- The cell in which users authenticate by default when they issue the **klog** command
- The cell membership of the AFS server processes that the AFS command interpreters on this machine contact by default

1. Change to the `/usr/vice/etc` directory and remove the symbolic link created in "Starting the BOS Server" on page 38.

```
# cd /usr/vice/etc
# rm ThisCell
```

2. Create the **ThisCell** file as a copy of the `/usr/afs/etc/ThisCell` file. Defining the same local cell for both server and client processes leads to the most consistent AFS performance.

```
# cp /usr/afs/etc/ThisCell ThisCell
```

Creating the Client CellServDB File

The `/usr/vice/etc/CellServDB` file on a client machine's local disk lists the database server machines for each cell that the local Cache Manager can contact. If there is no entry in the file for a cell, or if the list of database server machines is wrong, then users working on this machine cannot access the cell. The chapter in the *IBM AFS Administration Guide* about administering client machines explains how to maintain the file after creating it.

As the **afsd** program initializes the Cache Manager, it copies the contents of the **CellServDB** file into kernel memory. The Cache Manager always consults the list in kernel memory rather than the **CellServDB** file itself. Between reboots of the machine, you can use the **fs newcell** command to update the list in kernel memory directly; see the chapter in the *IBM AFS Administration Guide* about administering client machines.

The AFS distribution includes the file **CellServDB.sample**, and you have already copied it to the `/usr/vice/etc` directory. It includes an entry for all AFS cells that agreed to share their database server machine information at the time your AFS CD-ROM was created. The AFS Product Support group also maintains a copy of the file, updating it as necessary. If you are interested in participating in the global AFS namespace, it is a good policy to consult the file occasionally for updates. Ask the AFS Product Support group for a pointer to its location.

The **CellServDB.sample** file can be a good basis for the client **CellServDB** file, because all of the entries in it use the correct format. You can add or remove cell entries as you see fit. Later (in "Enabling

Access to Foreign Cells" on page 67) you perform additional steps that enable the Cache Manager actually to reach the cells.

In this section, you add an entry for the local cell to the local **CellServDB** file. The current working directory is still **/usr/vice/etc**.

1. Remove the symbolic link created in "Starting the BOS Server" on page 38 and rename the **CellServDB.sample** file to **CellServDB**.

```
# rm CellServDB
# mv CellServDB.sample CellServDB
```

2. Add an entry for the local cell to the **CellServDB** file. One easy method is to use the **cat** command to append the contents of the server **/usr/afs/etc/CellServDB** file to the client version.

```
# cat /usr/afs/etc/CellServDB >> CellServDB
```

Then open the file in a text editor to verify that there are no blank lines, and that all entries have the required format, which is described just following. The ordering of cells is not significant, but it can be convenient to have the client machine's home cell at the top; move it there now if you wish.

- The first line of a cell's entry has the following format:

```
>cell_name      #organization
```

where *cell_name* is the cell's complete Internet domain name (for example, **abc.com**) and *organization* is an optional field that follows any number of spaces and the number sign (#). By convention it names the organization to which the cell corresponds (for example, the ABC Corporation).

- After the first line comes a separate line for each database server machine. Each line has the following format:

```
IP_address      #machine_name
```

where *IP_address* is the machine's IP address in dotted decimal format (for example, 192.12.105.3). Following any number of spaces and the number sign (#) is *machine_name*, the machine's fully-qualified hostname (for example, **db1.abc.com**). In this case, the number sign does not indicate a comment; *machine_name* is a required field.

3. If the file includes cells that you do not wish users of this machine to access, remove their entries.

The following example shows entries for two cells, each of which has three database server machines:

```
>abc.com        #ABC Corporation (home cell)
192.12.105.3    #db1.abc.com
192.12.105.4    #db2.abc.com
192.12.105.55   #db3.abc.com
```

```
>stateu.edu      #State University cell
138.255.68.93    #serverA.stateu.edu
138.255.68.72    #serverB.stateu.edu
138.255.33.154   #serverC.stateu.edu
```

Configuring the Cache

The Cache Manager uses a cache on the local disk or in machine memory to store local copies of files fetched from file server machines. As the **afsd** program initializes the Cache Manager, it sets basic cache configuration parameters according to definitions in the local **/usr/vice/etc/cacheinfo** file. The file has three fields:

1. The first field names the local directory on which to mount the AFS filesystem. The conventional location is the **/afs** directory.
2. The second field defines the local disk directory to use for the disk cache. The conventional location is the **/usr/vice/cache** directory, but you can specify an alternate directory if another partition has more space available. There must always be a value in this field, but the Cache Manager ignores it if the machine uses a memory cache.
3. The third field specifies the number of kilobyte (1024 byte) blocks to allocate for the cache.

The values you define must meet the following requirements.

- On a machine using a disk cache, the Cache Manager expects always to be able to use the amount of space specified in the third field. Failure to meet this requirement can cause serious problems, some of which can be repaired only by rebooting. You must prevent non-AFS processes from filling up the cache partition. The simplest way is to devote a partition to the cache exclusively.
- The amount of space available in memory or on the partition housing the disk cache directory imposes an absolute limit on cache size.
- The maximum supported cache size can vary in each AFS release; see the *IBM AFS Release Notes* for the current version.
- For a disk cache, you cannot specify a value in the third field that exceeds 95% of the space available on the partition mounted at the directory named in the second field. If you violate this restriction, the **afsd** program exits without starting the Cache Manager and prints an appropriate message on the standard output stream. A value of 90% is more appropriate on most machines. Some operating systems (such as AIX) do not automatically reserve some space to prevent the partition from filling completely; for them, a smaller value (say, 80% to 85% of the space available) is more appropriate.
- For a memory cache, you must leave enough memory for other processes and applications to run. If you try to allocate more memory than is actually available, the **afsd** program exits without initializing the Cache Manager and produces the following message on the standard output stream.

```
afsd: memCache allocation failure at number KB
```

The *number* value is how many kilobytes were allocated just before the failure, and so indicates the approximate amount of memory available.

Within these hard limits, the factors that determine appropriate cache size include the number of users working on the machine, the size of the files with which they work, and (for a memory cache) the number of processes that run on the machine. The higher the demand from these factors, the larger the cache needs to be to maintain good performance.

Disk caches smaller than 10 MB do not generally perform well. Machines serving multiple users usually perform better with a cache of at least 60 to 70 MB. The point at which enlarging the cache further does not really improve performance depends on the factors mentioned previously and is difficult to predict.

Memory caches smaller than 1 MB are nonfunctional, and the performance of caches smaller than 5 MB is usually unsatisfactory. Suitable upper limits are similar to those for disk caches but are probably determined more by the demands on memory from other sources on the machine (number of users and processes). Machines running only a few processes possibly can use a smaller memory cache.

Configuring a Disk Cache

Note: Not all file system types that an operating system supports are necessarily supported for use as the cache partition. For possible restrictions, see the *IBM AFS Release Notes*.

To configure the disk cache, perform the following procedures:

1. Create the local directory to use for caching. The following instruction shows the conventional location, `/usr/vice/cache`. If you are devoting a partition exclusively to caching, as recommended, you must also configure it, make a file system on it, and mount it at the directory created in this step.

```
# mkdir /usr/vice/cache
```

2. Create the `cacheinfo` file to define the configuration parameters discussed previously. The following instruction shows the standard mount location, `/afs`, and the standard cache location, `/usr/vice/cache`.

```
# echo "/afs:/usr/vice/cache:#blocks" > /usr/vice/etc/cacheinfo
```

The following example defines the disk cache size as 50,000 KB:

```
# echo "/afs:/usr/vice/cache:50000" > /usr/vice/etc/cacheinfo
```

Configuring a Memory Cache

To configure a memory cache, create the `cacheinfo` file to define the configuration parameters discussed previously. The following instruction shows the standard mount location, `/afs`, and the standard cache location, `/usr/vice/cache` (though the exact value of the latter is irrelevant for a memory cache).

```
# echo "/afs:/usr/vice/cache:#blocks" > /usr/vice/etc/cacheinfo
```

The following example allocates 25,000 KB of memory for the cache.

```
# echo "/afs:/usr/vice/cache:25000" > /usr/vice/etc/cacheinfo
```

Configuring the Cache Manager

By convention, the Cache Manager mounts the AFS filesystem on the local `/afs` directory. In this section you create that directory.

The `afsd` program sets several cache configuration parameters as it initializes the Cache Manager, and starts daemons that improve performance. You can use the `afsd` command's arguments to override the parameters' default values and to change the number of some of the daemons. Depending on the machine's cache size, its amount of RAM, and how many people work on it, you can sometimes improve Cache Manager performance by overriding the default values. For a discussion of all of the `afsd` command's arguments, see its reference page in the *IBM AFS Administration Reference*.

The `afsd` command line in the AFS initialization script on each system type includes an `OPTIONS` variable. You can use it to set nondefault values for the command's arguments, in one of the following ways:

- You can create an `afsd options file` that sets values for arguments to the `afsd` command. If the file exists, its contents are automatically substituted for the `OPTIONS` variable in the AFS initialization script. The AFS distribution for some system types includes an options file; on other system types, you must create it.

You use two variables in the AFS initialization script to specify the path to the options file: `CONFIG` and `AFSDOPT`. On system types that define a conventional directory for configuration files, the `CONFIG` variable indicates it by default; otherwise, the variable indicates an appropriate location.

List the desired `afsd` options on a single line in the options file, separating each option with one or more spaces. The following example sets the `-stat` argument to 2500, the `-daemons` argument to 4, and the `-volumes` argument to 100.

```
-stat 2500 -daemons 4 -volumes 100
```

- On a machine that uses a disk cache, you can set the `OPTIONS` variable in the AFS initialization script to one of `SMALL`, `MEDIUM`, or `LARGE`. The AFS initialization script uses one of these settings if the `afsd options file` named by the `AFSDOPT` variable does not exist. In the script as distributed, the `OPTIONS` variable is set to the value `MEDIUM`.

Note: Do not set the `OPTIONS` variable to `SMALL`, `MEDIUM`, or `LARGE` on a machine that uses a memory cache. The arguments it sets are appropriate only on a machine that uses a disk cache.

The script (or on some system types the `afsd options file` named by the `AFSDOPT` variable) defines a value for each of `SMALL`, `MEDIUM`, and `LARGE` that sets `afsd` command arguments appropriately for client machines of different sizes:

- **SMALL** is suitable for a small machine that serves one or two users and has approximately 8 MB of RAM and a 20-MB cache
 - **MEDIUM** is suitable for a medium-sized machine that serves two to six users and has 16 MB of RAM and a 40-MB cache
 - **LARGE** is suitable for a large machine that serves five to ten users and has 32 MB of RAM and a 100-MB cache
- You can choose not to create an **afsd** options file and to set the **OPTIONS** variable in the initialization script to a null value rather than to the default **\$MEDIUM** value. You can then either set arguments directly on the **afsd** command line in the script, or set no arguments (and so accept default values for all Cache Manager parameters).

1. Create the local directory on which to mount the AFS filesystem, by convention **/afs**. If the directory already exists, verify that it is empty.

```
# mkdir /afs
```

2. On AIX systems, add the following line to the **/etc/vfs** file. It enables AIX to unmount AFS correctly during shutdown.

```
afs      4      none      none
```

3. On Linux systems, copy the **afsd** options file from the **/usr/vice/etc** directory to the **/etc/sysconfig** directory, removing the **.conf** extension as you do so.

```
# cp /usr/vice/etc/afs.conf /etc/sysconfig/afs
```

4. Edit the machine's AFS initialization script or **afsd** options file to set appropriate values for **afsd** command parameters. The script resides in the indicated location on each system type:

- On AIX systems, **/etc/rc.afs**
- On Digital UNIX systems, **/sbin/init.d/afs**
- On HP-UX systems, **/sbin/init.d/afs**
- On IRIX systems, **/etc/init.d/afs**
- On Linux systems, **/etc/sysconfig/afs** (the **afsd** options file)
- On Solaris systems, **/etc/init.d/afs**

Use one of the methods described in the introduction to this section to add the following flags to the **afsd** command line. If you intend for the machine to remain an AFS client, also set any performance-related arguments you wish.

- Add the **-nosetime** flag, because this is a file server machine that is also a client. The flag prevents the machine from picking a file server machine in the cell as its source for the correct time, which client machines normally do. File server machines instead use NTPD (as controlled by the **runntp** process) or another protocol to synchronize their clocks.
- Add the **-memcache** flag if the machine is to use a memory cache.

- Add the **-verbose** flag to display a trace of the Cache Manager's initialization on the standard output stream.

Overview: Completing the Installation of the First AFS Machine

The machine is now configured as an AFS file server and client machine. In this final phase of the installation, you initialize the Cache Manager and then create the upper levels of your AFS filesystem, among other procedures. The procedures are:

1. Verify that the initialization script works correctly, and incorporate it into the operating system's startup and shutdown sequence
2. Create and mount top-level volumes
3. Create and mount volumes to store system binaries in AFS
4. Enable access to foreign cells
5. Institute additional security measures
6. Remove client functionality if desired

Verifying the AFS Initialization Script

At this point you run the AFS initialization script to verify that it correctly invokes all of the necessary programs and AFS processes, and that they start correctly. The following are the relevant commands:

- The command that dynamically loads AFS modifications into the kernel, on some system types (not applicable if the kernel has AFS modifications built in)
- The **bosserv** command, which starts the BOS Server; it in turn starts the server processes for which you created entries in the **/usr/afs/local/BosConfig** file
- The **afsd** command, which initializes the Cache Manager

On system types that use a dynamic loader program, you must reboot the machine before running the initialization script, so that it can freshly load AFS modifications into the kernel.

If there are problems during the initialization, attempt to resolve them. The AFS Product Support group can provide assistance if necessary.

1. Issue the **bos shutdown** command to shut down the AFS server processes other than the BOS Server. Include the **-wait** flag to delay return of the command shell prompt until all processes shut down completely.

```
# /usr/afs/bin/bos shutdown <machine name> -wait
```

2. Issue the **ps** command to learn the **bossserver** process's process ID number (PID), and then the **kill** command to stop it.

```
# ps appropriate_ps_options | grep bossserver  
# kill bossserver_PID
```

3. Issue the appropriate commands to run the AFS initialization script for this system type.

On AIX systems:

- a. Reboot the machine and log in again as the local superuser **root**.

```
# cd /  
# shutdown -r now  
login: root  
Password: root_password
```

- b. Run the AFS initialization script.

```
# /etc/rc.afs
```

On Digital UNIX systems:

- a. Run the AFS initialization script.

```
# /sbin/init.d/afs start
```

On HP-UX systems:

- a. Run the AFS initialization script.

```
# /sbin/init.d/afs start
```

On IRIX systems:

- a. If you have configured the machine to use the **ml** dynamic loader program, reboot the machine and log in again as the local superuser **root**.

```
# cd /  
# shutdown -i6 -g0 -y  
login: root  
Password: root_password
```

- b. Issue the **chkconfig** command to activate the **afsserver** and **afsclient** configuration variables.

```
# /etc/chkconfig -f afsserver on  
# /etc/chkconfig -f afsclient on
```

- c. Run the AFS initialization script.

```
# /etc/init.d/afs start
```


On Linux systems:

- a. Reboot the machine and log in again as the local superuser **root**.

```
# cd /
# shutdown -r now
login: root
Password: root_password
```

- b. Run the AFS initialization script.

```
# /etc/rc.d/init.d/afs start
```

On Solaris systems:

- a. Reboot the machine and log in again as the local superuser **root**.

```
# cd /
# shutdown -i6 -g0 -y
login: root
Password: root_password
```

- b. Run the AFS initialization script.

```
# /etc/init.d/afs start
```

4. Wait for the message that confirms that Cache Manager initialization is complete.

On machines that use a disk cache, it can take a while to initialize the Cache Manager for the first time, because the **afsd** program must create all of the **V_n** files in the cache directory. Subsequent Cache Manager initializations do not take nearly as long, because the **V_n** files already exist.

As a basic test of correct AFS functioning, issue the **klog** command to authenticate as the **admin** user. Provide the password (*admin_passwd*) you defined in "Initializing Cell Security" on page 41.

```
# /usr/afs/bin/klog admin
Password: admin_passwd
```

5. Issue the **tokens** command to verify that the **klog** command worked correctly. If it did, the output looks similar to the following example for the **abc.com** cell, where **admin**'s AFS UID is 1. If the output does not seem correct, resolve the problem. Changes to the AFS initialization script are possibly necessary. The AFS Product Support group can provide assistance as necessary.

```
# /usr/afs/bin/tokens
Tokens held by the Cache Manager:
User's (AFS ID 1) tokens for afs@abc.com [Expires May 22 11:52]
--End of list--
```

6. Issue the **bos status** command to verify that the output for each process reads `Currently running normally`.

```
# /usr/afs/bin/bos status <machine name>
```

7. Change directory to the local file system root (`/`) and issue the **fs checkvolumes** command.

```
# cd /  
# /usr/afs/bin/fs checkvolumes
```

Activating the AFS Initialization Script

Now that you have confirmed that the AFS initialization script works correctly, take the action necessary to have it run automatically at each reboot. Proceed to the instructions for your system type:

- "Activating the Script on AIX Systems" on page 58
- "Activating the Script on Digital UNIX Systems" on page 58
- "Activating the Script on HP-UX Systems" on page 59
- "Activating the Script on IRIX Systems" on page 59
- "Activating the Script on Linux Systems" on page 60
- "Activating the Script on Solaris Systems" on page 60

Activating the Script on AIX Systems

1. Edit the AIX initialization file, `/etc/inittab`, adding the following line to invoke the AFS initialization script. Place it just after the line that starts NFS daemons.

```
rcafs:2:wait:/etc/rc.afs > /dev/console 2>&1 # Start AFS services
```

2. **(Optional)** There are now copies of the AFS initialization file in both the `/usr/vice/etc` and `/etc` directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc  
# rm rc.afs  
# ln -s /etc/rc.afs
```

3. Proceed to "Configuring the Top Levels of the AFS Filespace" on page 61.

Activating the Script on Digital UNIX Systems

1. Change to the `/sbin/init.d` directory and issue the `ln -s` command to create symbolic links that incorporate the AFS initialization script into the Digital UNIX startup and shutdown sequence.

```
# cd /sbin/init.d
# ln -s ../init.d/afs /sbin/rc3.d/S67afs
# ln -s ../init.d/afs /sbin/rc0.d/K66afs
```

2. **(Optional)** There are now copies of the AFS initialization file in both the `/usr/vice/etc` and `/sbin/init.d` directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /sbin/init.d/afs afs.rc
```

3. Proceed to "Configuring the Top Levels of the AFS Filespace" on page 61.

Activating the Script on HP-UX Systems

1. Change to the `/sbin/init.d` directory and issue the `ln -s` command to create symbolic links that incorporate the AFS initialization script into the HP-UX startup and shutdown sequence.

```
# cd /sbin/init.d
# ln -s ../init.d/afs /sbin/rc2.d/S460afs
# ln -s ../init.d/afs /sbin/rc2.d/K800afs
```

2. **(Optional)** There are now copies of the AFS initialization file in both the `/usr/vice/etc` and `/sbin/init.d` directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /sbin/init.d/afs afs.rc
```

3. Proceed to "Configuring the Top Levels of the AFS Filespace" on page 61.

Activating the Script on IRIX Systems

1. Change to the `/etc/init.d` directory and issue the `ln -s` command to create symbolic links that incorporate the AFS initialization script into the IRIX startup and shutdown sequence.

```
# cd /etc/init.d
# ln -s ../init.d/afs /etc/rc2.d/S35afs
# ln -s ../init.d/afs /etc/rc0.d/K35afs
```

2. **(Optional)** There are now copies of the AFS initialization file in both the `/usr/vice/etc` and `/etc/init.d` directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /etc/init.d/afs afs.rc
```

3. Proceed to "Configuring the Top Levels of the AFS Filespace" on page 61.

Activating the Script on Linux Systems

1. Issue the `chkconfig` command to activate the `afs` configuration variable. Based on the instruction in the AFS initialization file that begins with the string `#chkconfig`, the command automatically creates the symbolic links that incorporate the script into the Linux startup and shutdown sequence.

```
# /sbin/chkconfig --add afs
```

2. **(Optional)** There are now copies of the AFS initialization file in both the `/usr/vice/etc` and `/etc/rc.d/init.d` directories, and copies of the `afsd` options file in both the `/usr/vice/etc` and `/etc/sysconfig` directories. If you want to avoid potential confusion by guaranteeing that the two copies of each file are always the same, create a link between them. You can always retrieve the original script or options file from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc afs.conf
# ln -s /etc/rc.d/init.d/afs afs.rc
# ln -s /etc/sysconfig/afs afs.conf
```

3. Proceed to "Configuring the Top Levels of the AFS Filespace" on page 61.

Activating the Script on Solaris Systems

1. Change to the `/etc/init.d` directory and issue the `ln -s` command to create symbolic links that incorporate the AFS initialization script into the Solaris startup and shutdown sequence.

```
# cd /etc/init.d
# ln -s ../init.d/afs /etc/rc3.d/S99afs
# ln -s ../init.d/afs /etc/rc0.d/K66afs
```

2. **(Optional)** There are now copies of the AFS initialization file in both the `/usr/vice/etc` and `/etc/init.d` directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /etc/init.d/afs afs.rc
```

Configuring the Top Levels of the AFS Filespace

If you have not previously run AFS in your cell, you now configure the top levels of your cell's AFS filesystem. If you have run a previous version of AFS, the filesystem is already configured. Proceed to "Storing AFS Binaries in AFS" on page 62.

You created the **root.afs** volume in "Starting the File Server, Volume Server, and Salvager" on page 44, and the Cache Manager mounted it automatically on the local **/afs** directory when you ran the AFS initialization script in "Verifying the AFS Initialization Script" on page 55. You now set the access control list (ACL) on the **/afs** directory; creating, mounting, and setting the ACL are the three steps required when creating any volume.

After setting the ACL on the **root.afs** volume, you create your cell's **root.cell** volume, mount it as a subdirectory of the **/afs** directory, and set the ACL. Create both a read/write and a regular mount point for the **root.cell** volume. The read/write mount point enables you to access the read/write version of replicated volumes when necessary. Creating both mount points essentially creates separate read-only and read-write copies of your filesystem, and enables the Cache Manager to traverse the filesystem on a read-only path or read/write path as appropriate. For further discussion of these concepts, see the chapter in the *IBM AFS Administration Guide* about administering volumes.

Then replicate both the **root.afs** and **root.cell** volumes. This is required if you want to replicate any other volumes in your cell, because all volumes mounted above a replicated volume must themselves be replicated in order for the Cache Manager to access the replica.

When the **root.afs** volume is replicated, the Cache Manager is programmed to access its read-only version (**root.afs.readonly**) whenever possible. To make changes to the contents of the **root.afs** volume (when, for example, you mount another cell's **root.cell** volume at the second level in your filesystem), you must mount the **root.afs** volume temporarily, make the changes, release the volume and remove the temporary mount point. For instructions, see "Enabling Access to Foreign Cells" on page 67.

1. Issue the **fs setacl** command to edit the ACL on the **/afs** directory. Add an entry that grants the **l (lookup)** and **r (read)** permissions to the **system:anyuser** group, to enable all AFS users who can reach your cell to traverse through the directory. If you prefer to enable access only to locally authenticated users, substitute the **system:authuser** group.

Note that there is already an ACL entry that grants all seven access rights to the **system:administrators** group. It is a default entry that AFS places on every new volume's root directory.

```
# /usr/afs/bin/fs setacl /afs system:anyuser rl
```

2. Issue the **vos create** command to create the **root.cell** volume. Then issue the **fs mkmount** command to mount it as a subdirectory of the **/afs** directory, where it serves as the root of your cell's local AFS filesystem. Finally, issue the **fs setacl** command to create an ACL entry for the **system:anyuser** group (or **system:authuser** group).

For the *partition name* argument, substitute the name of one of the machine's AFS server partitions (such as **/vicepa**). For the *cellname* argument, substitute your cell's fully-qualified Internet domain name (such as **abc.com**).

```
# /usr/afs/bin/vos create <machine name> <partition name> root.cell
# /usr/afs/bin/fs mkmount /afs/cellname root.cell
# /usr/afs/bin/fs setacl /afs/cellname system:anyuser rl
```

3. **(Optional)** Create a symbolic link to a shortened cell name, to reduce the length of pathnames for users in the local cell. For example, in the **abc.com** cell, **/afs/abc** is a link to **/afs/abc.com**.

```
# cd /afs
# ln -s full_cellname short_cellname
```

4. Issue the **fs mkmount** command to create a read/write mount point for the **root.cell** volume (you created a regular mount point in Step "2" on page 61).

By convention, the name of a read/write mount point begins with a period, both to distinguish it from the regular mount point and to make it visible only when the **-a** flag is used on the **ls** command.

Change directory to **/usr/afs/bin** to make it easier to access the command binaries.

```
# cd /usr/afs/bin
# ./fs mkmount /afs/.cellname root.cell -rw
```

5. Issue the **vos addsite** command to define a replication site for both the **root.afs** and **root.cell** volumes. In each case, substitute for the *partition name* argument the partition where the volume's read/write version resides. When you install additional file server machines, it is a good idea to create replication sites on them as well.

```
# ./vos addsite <machine name> <partition name> root.afs
# ./vos addsite <machine name> <partition name> root.cell
```

6. Issue the **fs examine** command to verify that the Cache Manager can access both the **root.afs** and **root.cell** volumes, before you attempt to replicate them. The output lists each volume's name, volumeID number, quota, size, and the size of the partition that houses them. If you get an error message instead, do not continue before taking corrective action.

```
# ./fs examine /afs
# ./fs examine /afs/cellname
```

7. Issue the **vos release** command to release a replica of the **root.afs** and **root.cell** volumes to the sites you defined in Step "5" on page 62.

```
# ./vos release root.afs
# ./vos release root.cell
```

8. Issue the **fs checkvolumes** to force the Cache Manager to notice that you have released read-only versions of the volumes, then issue the **fs examine** command again. This time its output mentions the read-only version of the volumes (**root.afs.readonly** and **root.cell.readonly**) instead of the read/write versions, because of the Cache Manager's bias to access the read-only version of the **root.afs** volume if it exists.

```
# ./fs checkvolumes
# ./fs examine /afs
# ./fs examine /afs/cellname
```

Storing AFS Binaries in AFS

In the conventional configuration, you make AFS client binaries and configuration files available in the subdirectories of the `/usr/afsws` directory on client machines (**afsws** is an acronym for **AFS workstation**). You can conserve local disk space by creating `/usr/afsws` as a link to an AFS volume that houses the AFS client binaries and configuration files for this system type.

In this section you create the necessary volumes. The conventional location to which to link `/usr/afsws` is `/afs/cellname/sysname/usr/afsws`, where `sysname` is the appropriate system type name as specified in the *IBM AFS Release Notes*. The instructions in "Installing Additional Client Machines" on page 105 assume that you have followed the instructions in this section.

If you have previously run AFS in the cell, the volumes possibly already exist. If so, you need to perform Step "8" on page 64 only.

The current working directory is still `/usr/afs/bin`, which houses the **fs** and **vos** command suite binaries. In the following commands, it is possible you still need to specify the pathname to the commands, depending on how your `PATH` environment variable is set.

1. Issue the **vos create** command to create volumes for storing the AFS client binaries for this system type. The following example instruction creates volumes called `sysname`, `sysname.usr`, and `sysname.usr.afsws`. Refer to the *IBM AFS Release Notes* to learn the proper value of `sysname` for this system type.

```
# vos create <machine name> <partition name> sysname
# vos create <machine name> <partition name> sysname.usr
# vos create <machine name> <partition name> sysname.usr.afsws
```

2. Issue the **fs mkmount** command to mount the newly created volumes. Because the **root.cell** volume is replicated, you must precede the `cellname` part of the pathname with a period to specify the read/write mount point, as shown. Then issue the **vos release** command to release a new replica of the **root.cell** volume, and the **fs checkvolumes** command to force the local Cache Manager to access them.

```
# fs mkmount -dir /afs/.cellname/sysname -vol sysname
# fs mkmount -dir /afs/.cellname/sysname/usr -vol sysname.usr
# fs mkmount -dir /afs/.cellname/sysname/usr/afsws -vol sysname.usr.afsws
# vos release root.cell
# fs checkvolumes
```

3. Issue the **fs setacl** command to grant the **l** (**lookup**) and **r** (**read**) permissions to the **system:anyuser** group on each new directory's ACL.

```
# cd /afs/.cellname/sysname
# fs setacl -dir . usr usr/afsws -acl system:anyuser rl
```

4. Issue the **fs setquota** command to set an unlimited quota on the volume mounted at the `/afs/cellname/sysname/usr/afsws` directory. This enables you to copy all of the appropriate files from the CD-ROM into the volume without exceeding the volume's quota.

If you wish, you can set the volume's quota to a finite value after you complete the copying operation. At that point, use the **vos examine** command to determine how much space the volume is occupying. Then issue the **fs setquota** command to set a quota that is slightly larger.

```
# fs setquota /afs/.cellname/sysname/usr/afsws 0
```

5. Mount the AFS CD-ROM for this machine's system type on the local **/cdrom** directory, if it is not already. For instructions on mounting CD-ROMs (either locally or remotely via NFS), consult the operating system documentation.

6. Copy the contents of the indicated directories from the CD-ROM into the **/afs/cellname/sysname/usr/afsws** directory.

```
# cd /afs/.cellname/sysname/usr/afsws
# cp -rp /cdrom/sysname/bin .
# cp -rp /cdrom/sysname/etc .
# cp -rp /cdrom/sysname/include .
# cp -rp /cdrom/sysname/lib .
```

7. Issue the **fs setacl** command to set the ACL on each directory appropriately. To comply with the terms of your AFS License agreement, you must prevent unauthorized users from accessing AFS software. To enable access for locally authenticated users only, set the ACL on the **etc**, **include**, and **lib** subdirectories to grant the **l** and **r** permissions to the **system:authuser** group rather than the **system:anyuser** group. The **system:anyuser** group must retain the **l** and **r** permissions on the **bin** subdirectory to enable unauthenticated users to access the **klog** binary. To ensure that unauthorized users are not accessing AFS software, check periodically that the ACLs on these directories are set properly.

```
# cd /afs/.cellname/sysname/usr/afsws
# fs setacl -dir etc include lib -acl system:authuser rl \
          system:anyuser none
```

8. Create **/usr/afsws** on the local disk as a symbolic link to the directory **/afs/cellname/@sys/usr/afsws**. You can specify the actual system name instead of **@sys** if you wish, but the advantage of using **@sys** is that it remains valid if you upgrade this machine to a different system type.

```
# ln -s /afs/cellname/@sys/usr/afsws /usr/afsws
```

9. **(Optional)** To enable users to issue commands from the AFS suites (such as **fs**) without having to specify a pathname to their binaries, include the **/usr/afsws/bin** and **/usr/afsws/etc** directories in the **PATH** environment variable you define in each user's shell initialization file (such as **.cshrc**).

Storing AFS Documents in AFS

The AFS distribution includes the following documents:

- *IBM AFS Release Notes*
- *IBM AFS Quick Beginnings*
- *IBM AFS User Guide*
- *IBM AFS Administration Reference*
- *IBM AFS Administration Guide*

The AFS CD-ROM for each system type has a top-level **Documentation** directory, with a subdirectory for each document format provided. The different formats are suitable for online viewing, printing, or both.

This section explains how to create and mount a volume to house the documents, making them available to your users. The recommended mount point for the volume is `/afs/cellname/afsdoc`. If you wish, you can create a link to the mount point on each client machine's local disk, called `/usr/afsdoc`. Alternatively, you can create a link to the mount point in each user's home directory. You can also choose to permit users to access only certain documents (most probably, the *IBM AFS User Guide*) by creating different mount points or setting different ACLs on different document directories.

The current working directory is still `/usr/afs/bin`, which houses the **fs** and **vos** command suite binaries you use to create and mount volumes. In the following commands, it is possible you still need to specify the pathname to the commands, depending on how your `PATH` environment variable is set.

1. Issue the **vos create** command to create a volume for storing the AFS documentation. Include the **-maxquota** argument to set an unlimited quota on the volume. This enables you to copy all of the appropriate files from the CD-ROM into the volume without exceeding the volume's quota.

If you wish, you can set the volume's quota to a finite value after you complete the copying operations. At that point, use the **vos examine** command to determine how much space the volume is occupying. Then issue the **fs setquota** command to set a quota that is slightly larger.

```
# vos create <machine name> <partition name> afsdoc -maxquota 0
```

2. Issue the **fs mkmount** command to mount the new volume. Because the **root.cell** volume is replicated, you must precede the *cellname* with a period to specify the read/write mount point, as shown. Then issue the **vos release** command to release a new replica of the **root.cell** volume, and the **fs checkvolumes** command to force the local Cache Manager to access them.

```
# fs mkmount -dir /afs/.cellname/afsdoc -vol afsdoc
# vos release root.cell
# fs checkvolumes
```

3. Issue the **fs setacl** command to grant the **rl** permissions to the **system:anyuser** group on the new directory's ACL.

```
# cd /afs/.cellname/afsdoc
# fs setacl . system:anyuser rl
```

4. Mount the AFS CD-ROM for any system type on the local **/cdrom** directory, if one is not already. For instructions on mounting CD-ROMs (either locally or remotely via NFS), consult the operating system documentation.

5. Copy the AFS documents in one or more formats from the CD-ROM into subdirectories of the `/afs/cellname/afsdoc` directory. Repeat the commands for each format.

```
# mkdir format_name
# cd format_name
# cp -rp /cdrom/Documentation/format .
```

If you choose to store the HTML version of the documents in AFS, note that in addition to a subdirectory for each document there are several files with a **.gif** extension, which enable readers to move easily between sections of a document. The file called **index.htm** is an introductory HTML page that contains a hyperlink to each of the documents. For online viewing to work properly, these files must remain in the top-level HTML directory (the one named, for example, */afs/cellname/afsdoc/html*).

6. **(Optional)** If you believe it is helpful to your users to access the AFS documents in a certain format via a local disk directory, create **/usr/afsdoc** on the local disk as a symbolic link to the documentation directory in AFS (*/afs/cellname/afsdoc/format_name*).

```
# ln -s /afs/cellname/afsdoc/format_name /usr/afsdoc
```

An alternative is to create a link in each user's home directory to the */afs/cellname/afsdoc/format_name* directory.

Storing System Binaries in AFS

You can also choose to store other system binaries in AFS volumes, such as the standard UNIX programs conventionally located in local disk directories such as **/etc**, **/bin**, and **/lib**. Storing such binaries in an AFS volume not only frees local disk space, but makes it easier to update binaries on all client machines.

The following is a suggested scheme for storing system binaries in AFS. It does not include instructions, but you can use the instructions in "Storing AFS Binaries in AFS" on page 62 (which are for AFS-specific binaries) as a template.

Some files must remain on the local disk for use when AFS is inaccessible (during bootup and file server or network outages). The required binaries include the following:

- A text editor, network commands, and so on
- Files used during the boot sequence before the **afsd** program runs, such as initialization and configuration files, and binaries for commands that mount file systems
- Files used by dynamic kernel loader programs

In most cases, it is more secure to enable only locally authenticated users to access system binaries, by granting the **l (lookup)** and **r (read)** permissions to the **system:authuser** group on the ACLs of directories that contain the binaries. If users need to access a binary while unauthenticated, however, the ACL on its directory must grant those permissions to the **system:anyuser** group.

The following chart summarizes the suggested volume and mount point names for storing system binaries. It uses a separate volume for each directory. You already created a volume called *sysname* for this machine's system type when you followed the instructions in "Storing AFS Binaries in AFS" on page 62.

You can name volumes in any way you wish, and mount them at other locations than those suggested here. However, this scheme has several advantages:

- Volume names clearly identify volume contents
- Using the *sysname* prefix on every volume makes it is easy to back up all of the volumes together, because the AFS Backup System enables you to define sets of volumes based on a string included in all of their names
- It makes it easy to track related volumes, keeping them together on the same file server machine if desired
- There is a clear relationship between volume name and mount point name

Volume Name	Mount Point
<i>sysname</i>	<i>/afs/cellname/sysname</i>
<i>sysname.bin</i>	<i>/afs/cellname/sysname/bin</i>
<i>sysname.etc</i>	<i>/afs/cellname/sysname/etc</i>
<i>sysname.usr</i>	<i>/afs/cellname/sysname/usr</i>
<i>sysname.usr.afs</i>	<i>/afs/cellname/sysname/usr/afs</i>
<i>sysname.usr.bin</i>	<i>/afs/cellname/sysname/usr/bin</i>
<i>sysname.usr.etc</i>	<i>/afs/cellname/sysname/usr/etc</i>
<i>sysname.usr.inc</i>	<i>/afs/cellname/sysname/usr/include</i>
<i>sysname.usr.lib</i>	<i>/afs/cellname/sysname/usr/lib</i>
<i>sysname.usr.loc</i>	<i>/afs/cellname/sysname/usr/local</i>
<i>sysname.usr.man</i>	<i>/afs/cellname/sysname/usr/man</i>
<i>sysname.usr.sys</i>	<i>/afs/cellname/sysname/usr/sys</i>

Enabling Access to Foreign Cells

In this section you create a mount point in your AFS filespace for the **root.cell** volume of each foreign cell that you want to enable your users to access. For users working on a client machine to access the cell, there must in addition be an entry for it in the client machine's local **/usr/vice/etc/CellServDB** file. (The instructions in "Creating the Client CellServDB File" on page 49 suggest that you use the **CellServDB.sample** file included in the AFS distribution as the basis for your cell's client **CellServDB** file. The sample file lists all of the cells that had agreed to participate in the AFS global namespace at the time your AFS CD-ROM was created. As mentioned in that section, the AFS Product Support group also maintains a copy of the file, updating it as necessary.)

The chapter in the *IBM AFS Administration Guide* about cell administration and configuration issues discusses the implications of participating in the global AFS namespace. The chapter about administering client machines explains how to maintain knowledge of foreign cells on client machines, and includes suggestions for maintaining a central version of the file in AFS.

1. Issue the **fs mkmount** command to mount each foreign cell's **root.cell** volume on a directory called

/afs/foreign_cell. Because the **root.afs** volume is replicated, you must create a temporary mount point for its read/write version in a directory to which you have write access (such as your cell's */afs/cellname* directory). Create the mount points, issue the **vos release** command to release new replicas to the read-only sites for the **root.afs** volume, and issue the **fs checkvolumes** command to force the local Cache Manager to access the new replica.

Note: You need to issue the **fs mkmount** command only once for each foreign cell's **root.cell** volume. You do not need to repeat the command on each client machine.

Substitute your cell's name for *cellname*.

```
# cd /afs/.cellname
# /usr/afs/bin/fs mkmount temp root.afs
```

Repeat the **fs mkmount** command for each foreign cell you wish to mount at this time.

```
# /usr/afs/bin/fs mkmount temp/foreign_cell root.cell -c foreign_cell
```

Issue the following commands only once.

```
# /usr/afs/bin/fs rmmount temp
# /usr/afs/bin/vos release root.afs
# /usr/afs/bin/fs checkvolumes
```

2. If this machine is going to remain an AFS client after you complete the installation, verify that the local */usr/vice/etc/CellServDB* file includes an entry for each foreign cell.

For each cell that does not already have an entry, complete the following instructions:

- a. Create an entry in the **CellServDB** file. Be sure to comply with the formatting instructions in "Creating the Client CellServDB File" on page 49.
- b. Issue the **fs newcell** command to add an entry for the cell directly to the list that the Cache Manager maintains in kernel memory. Provide each database server machine's fully qualified hostname.

```
# /usr/afs/bin/fs newcell <foreign_cell> <dbserver1> \
    [<dbserver2>] [<dbserver3>]
```

- c. If you plan to maintain a central version of the **CellServDB** file (the conventional location is */afs/cellname/common/etc/CellServDB*), create it now as a copy of the local */usr/vice/etc/CellServDB* file. Verify that it includes an entry for each foreign cell you want your users to be able to access.

```
# mkdir common
# mkdir common/etc
# cp /usr/vice/etc/CellServDB common/etc
# /usr/afs/bin/vos release root.cell
```

3. Issue the **ls** command to verify that the new cell's mount point is visible in your filesystem. The output lists the directories at the top level of the new cell's AFS filesystem.

```
# ls /afs/foreign_cell
```

4. Please register your cell with the AFS Product Support group at this time. If you do not want to participate in the global AFS namespace, they list your cell in a private **CellServDB** file that is not available to other AFS cells.

Improving Cell Security

This section discusses ways to improve the security of AFS data in your cell. Also see the chapter in the *IBM AFS Administration Guide* about configuration and administration issues.

Controlling root Access

As on any machine, it is important to prevent unauthorized users from logging onto an AFS server or client machine as the local superuser **root**. Take care to keep the **root** password secret.

The local **root** superuser does not have special access to AFS data through the Cache Manager (as members of the **system:administrators** group do), but it does have the following privileges:

- On client machines, the ability to issue commands from the **fs** suite that affect AFS performance
- On server machines, the ability to disable authorization checking, or to install rogue process binaries

Controlling System Administrator Access

Following are suggestions for managing AFS administrative privilege:

- Create an administrative account for each administrator named something like *username.admin*. Administrators authenticate under these identities only when performing administrative tasks, and destroy the administrative tokens immediately after finishing the task (either by issuing the **unlog** command, or the **klog** command to adopt their regular identity).
- Set a short ticket lifetime for administrator accounts (for example, 20 minutes) by using the **-lifetime** argument to the **kas setfields** command, which is described in the *IBM AFS Administration Reference*. Do not however, use a short lifetime for users who issue long-running **backup** commands.
- Limit the number of system administrators in your cell, especially those who belong to the **system:administrators** group. By default they have all ACL rights on all directories in the local AFS filesystem, and therefore must be trusted not to examine private files.
- Limit the use of system administrator accounts on machines in public areas. It is especially important not to leave such machines unattended without first destroying the administrative tokens.
- Limit the use by administrators of standard UNIX commands that make connections to remote machines (such as the **telnet** utility). Many of these programs send passwords across the network without encrypting them.

Protecting Sensitive AFS Directories

Some subdirectories of the `/usr/afs` directory contain files crucial to cell security. Unauthorized users must not read or write to these files because of the potential for misuse of the information they contain.

As the BOS Server initializes for the first time on a server machine, it creates several files and directories (as mentioned in "Starting the BOS Server" on page 38). It sets their owner to the local superuser **root** and sets their mode bits to enable writing by the owner only; in some cases, it also restricts reading.

At each subsequent restart, the BOS Server checks that the owner and mode bits on these files are still set appropriately. If they are not, it writes the following message to the `/usr/afs/logs/BosLog` file:

```
Bosserver reports inappropriate access on server directories
```

The BOS Server does not reset the mode bits, which enables you to set alternate values if you wish.

The following chart lists the expected mode bit settings. A question mark indicates that the BOS Server does not check that mode bit.

<code>/usr/afs</code>	<code>drwxr?xr-x</code>
<code>/usr/afs/backup</code>	<code>drwx???---</code>
<code>/usr/afs/bin</code>	<code>drwxr?xr-x</code>
<code>/usr/afs/db</code>	<code>drwx???---</code>
<code>/usr/afs/etc</code>	<code>drwxr?xr-x</code>
<code>/usr/afs/etc/KeyFile</code>	<code>-rw????---</code>
<code>/usr/afs/etc/UserList</code>	<code>-rw?????--</code>
<code>/usr/afs/local</code>	<code>drwx???---</code>
<code>/usr/afs/logs</code>	<code>drwxr?xr-x</code>

Removing Client Functionality

Follow the instructions in this section only if you do not wish this machine to remain an AFS client. Removing client functionality means that you cannot use this machine to access AFS files.

1. Remove the files from the `/usr/vice/etc` directory. The command does not remove the directory for files used by the dynamic kernel loader program, if it exists on this system type. Those files are still needed on a server-only machine.

```
# cd /usr/vice/etc
# rm *
# rm -rf C
```

2. Create symbolic links to the **ThisCell** and **CellServDB** files in the `/usr/afs/etc` directory. This makes it possible to issue commands from the AFS command suites (such as **bos** and **fs**) on this machine.

```
# ln -s /usr/afs/etc/ThisCell ThisCell
# ln -s /usr/afs/etc/CellServDB CellServDB
```

3. On IRIX systems, issue the **chkconfig** command to deactivate the **afsclient** configuration variable.

```
# /etc/chkconfig -f afsclient off
```

4. Reboot the machine. Most system types use the **shutdown** command, but the appropriate options vary.

```
# cd /
# shutdown appropriate_options
```


Chapter 3. Installing Additional Server Machines

Instructions for the following procedures appear in the indicated section of this chapter.

- "Installing an Additional File Server Machine" on page 73
- "Installing Database Server Functionality" on page 97
- "Removing Database Server Functionality" on page 101

The instructions make the following assumptions.

- You have already installed your cell's first file server machine by following the instructions in "Installing the First AFS Machine" on page 7
- You are logged in as the local superuser **root**
- You are working at the console
- A standard version of one of the operating systems supported by the current version of AFS is running on the machine
- You can access the data on the AFS CD-ROMs, either through a local CD-ROM drive or via an NFS mount of a CD-ROM drive attached to a machine that is accessible by network

Installing an Additional File Server Machine

The procedure for installing a new file server machine is similar to installing the first file server machine in your cell. There are a few parts of the installation that differ depending on whether the machine is the same AFS system type as an existing file server machine or is the first file server machine of its system type in your cell. The differences mostly concern the source for the needed binaries and files, and what portions of the Update Server you install:

- On a new system type, you must load files and binaries from the AFS CD-ROM. You install the server portion of the Update Server to make this machine the binary distribution machine for its system type.
- On an existing system type, you can copy files and binaries from a previously installed file server machine, rather than from the CD-ROM. You install the client portion of the Update Server to accept updates of binaries, because a previously installed machine of this type was installed as the binary distribution machine.

These instructions are brief; for more detailed information, refer to the corresponding steps in "Installing the First AFS Machine" on page 7.

To install a new file server machine, perform the following procedures:

1. Copy needed binaries and files onto this machine's local disk

2. Incorporate AFS modifications into the kernel
3. Configure partitions for storing volumes
4. Replace the standard **fsck** utility with the AFS-modified version on some system types
5. Start the Basic OverSeer (BOS) Server
6. Start the appropriate portion of the Update Server
7. Start the **fs** process, which incorporates three component processes: the File Server, Volume Server, and Salvager
8. Start the controller process (called **runntp**) for the Network Time Protocol Daemon, which synchronizes clocks

After completing the instructions in this section, you can install database server functionality on the machine according to the instructions in "Installing Database Server Functionality" on page 97.

Creating AFS Directories and Performing Platform-Specific Procedures

Create the **/usr/afs** and **/usr/vice/etc** directories on the local disk. Subsequent instructions copy files from the AFS distribution CD-ROM into them, at the appropriate point for each system type.

```
# mkdir /usr/afs
# mkdir /usr/afs/bin
# mkdir /usr/vice
# mkdir /usr/vice/etc
# mkdir /cdrom
```

As on the first file server machine, the initial procedures in installing an additional file server machine vary a good deal from platform to platform. For convenience, the following sections group together all of the procedures for a system type. Most of the remaining procedures are the same on every system type, but differences are noted as appropriate. The initial procedures are the following.

- Incorporate AFS modifications into the kernel, either by using a dynamic kernel loader program or by building a new static kernel
- Configure server partitions to house AFS volumes
- Replace the operating system vendor's **fsck** program with a version that recognizes AFS data
- If the machine is to remain an AFS client machine, modify the machine's authentication system so that users obtain an AFS token as they log into the local file system. (For this procedure only, the instructions direct you to the platform-specific section in "Installing the First AFS Machine" on page 7.)

To continue, proceed to the section for this system type:

- "Getting Started on AIX Systems" on page 75
- "Getting Started on Digital UNIX Systems" on page 76

- "Getting Started on HP-UX Systems" on page 79
- "Getting Started on IRIX Systems" on page 81
- "Getting Started on Linux Systems" on page 85
- "Getting Started on Solaris Systems" on page 86

Getting Started on AIX Systems

Begin by running the AFS initialization script to call the AIX kernel extension facility, which dynamically loads AFS modifications into the kernel. Then configure partitions and replace the AIX **fsck** program with a version that correctly handles AFS volumes.

1. Mount the AFS CD-ROM for AIX on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your AIX documentation. Then change directory as indicated.

```
# cd /cdrom/rs_aix42/root.client/usr/vice/etc
```

2. Copy the AFS kernel library files to the local **/usr/vice/etc/dkload** directory, and the AFS initialization script to the **/etc** directory.

```
# cp -rp dkload /usr/vice/etc
# cp -p rc.afs /etc/rc.afs
```

3. Edit the **/etc/rc.afs** script, setting the **NFS** variable as indicated.

If the machine is not to function as an NFS/AFS Translator, set the **NFS** variable as follows.

```
NFS=$NFS_NONE
```

If the machine is to function as an NFS/AFS Translator and is running AIX 4.2.1 or higher, set the **NFS** variable as follows. Note that NFS must already be loaded into the kernel, which happens automatically on systems running AIX 4.1.1 and later, as long as the file **/etc/exports** exists.

```
NFS=$NFS_IAUTH
```

4. Invoke the **/etc/rc.afs** script to load AFS modifications into the kernel. You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/rc.afs
```

5. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

6. Use the **SMIT** program to create a journaling file system on each partition to be configured as an AFS server partition.

7. Mount each partition at one of the **/vicep_{xx}** directories. Choose one of the following three methods:

- Use the **SMIT** program

- Use the **mount -a** command to mount all partitions at once
- Use the **mount** command on each partition in turn

Also configure the partitions so that they are mounted automatically at each reboot. For more information, refer to the AIX documentation.

8. Move the AIX **fsck** program helper to a safe location and install the version from the AFS distribution in its place. The AFS CD-ROM must still be mounted at the **/cdrom** directory.

```
# cd /sbin/helpers
# mv v3fshelper v3fshelper.noafs
# cp -p /cdrom/rs_aix42/root.server/etc/v3fshelper v3fshelper
```

9. If the machine is to remain an AFS client, incorporate AFS into its authentication system, following the instructions in "Enabling AFS Login on AIX Systems" on page 12.
10. Proceed to "Starting Server Programs" on page 88.

Getting Started on Digital UNIX Systems

Begin by building AFS modifications into the kernel, then configure server partitions and replace the Digital UNIX **fsck** program with a version that correctly handles AFS volumes.

If the machine's hardware and software configuration exactly matches another Digital UNIX machine on which AFS is already built into the kernel, you can copy the kernel from that machine to this one. In general, however, it is better to build AFS modifications into the kernel on each machine according to the following instructions.

1. Create a copy called **AFS** of the basic kernel configuration file included in the Digital UNIX distribution as **/usr/sys/conf/machine_name**, where *machine_name* is the machine's hostname in all uppercase letters.

```
# cd /usr/sys/conf
# cp machine_name AFS
```

2. Add AFS to the list of options in the configuration file you created in the previous step, so that the result looks like the following:

```
      .
      .
options      UFS
options      NFS
options      AFS
      .
      .
```

3. Add an entry for AFS to two places in the file **/usr/sys/conf/files**.

- Add a line for AFS to the list of **OPTIONS**, so that the result looks like the following:

```
      .
      .
      .
```


6. Copy the AFS initialization script to the local directory for initialization files (by convention, `/sbin/init.d` on Digital UNIX machines). Note the removal of the `.rc` extension as you copy the script.

```
# cp usr/vice/etc/afs.rc /sbin/init.d/afs
```

7. Copy the AFS kernel module to the local `/usr/sys/BINARY` directory.

If the machine's kernel supports NFS server functionality:

```
# cp bin/libafs.o /usr/sys/BINARY/afs.mod
```

If the machine's kernel does not support NFS server functionality:

```
# cp bin/libafs.nonfs.o /usr/sys/BINARY/afs.mod
```

8. Configure and build the kernel. Respond to any prompts by pressing `<Return>`. The resulting kernel resides in the file `/sys/AFS/vmunix`.

```
# doconfig -c AFS
```

9. Rename the existing kernel file and copy the new, AFS-modified file to the standard location.

```
# mv /vmunix /vmunix_noafs
# cp /sys/AFS/vmunix /vmunix
```

10. Reboot the machine to start using the new kernel, and login again as the superuser `root`.

```
# cd /
# shutdown -r now
login: root
Password: root_password
```

11. Create a directory called `/vicepxx` for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

12. Add a line with the following format to the file systems registry file, `/etc/fstab`, for each directory just created. The entry maps the directory name to the disk partition to be mounted on it.

```
/dev/disk /vicepxx ufs rw 0 2
```

The following is an example for the first partition being configured.

```
/dev/rz3a /vicepa ufs rw 0 2
```

13. Create a file system on each partition that is to be mounted at a `/vicepxx` directory. The following command is probably appropriate, but consult the Digital UNIX documentation for more information.

```
# newfs -v /dev/disk
```

14. Mount each partition by issuing either the `mount -a` command to mount all partitions at once or the `mount` command to mount each partition in turn.

15. Install the **vfsc** binary to the **/sbin** and **/usr/sbin** directories. The AFS CD-ROM must still be mounted at the **/cdrom** directory.

```
# cd /cdrom/alpha_dux40/root.server/etc
# cp vfsc /sbin/vfsc
# cp vfsc /usr/sbin/vfsc
```

16. Rename the Digital UNIX **fsck** binaries and create symbolic links to the **vfsc** program.

```
# cd /sbin
# mv ufs_fsck ufs_fsck.noafs
# ln -s vfsc ufs_fsck
# cd /usr/sbin
# mv ufs_fsck ufs_fsck.noafs
# ln -s vfsc ufs_fsck
```

17. If the machine is to remain an AFS client, incorporate AFS into its authentication system, following the instructions in "Enabling AFS Login on Digital UNIX Systems" on page 18.
18. Proceed to "Starting Server Programs" on page 88.

Getting Started on HP-UX Systems

Begin by building AFS modifications into the kernel, then configure server partitions and replace the HP-UX **fsck** program with a version that correctly handles AFS volumes.

If the machine's hardware and software configuration exactly matches another HP-UX machine on which AFS is already built into the kernel, you can copy the kernel from that machine to this one. In general, however, it is better to build AFS modifications into the kernel on each machine according to the following instructions.

1. Move the existing kernel-related files to a safe location.

```
# cp /stand/vmunix /stand/vmunix.noafs
# cp /stand/system /stand/system.noafs
```

2. Mount the AFS CD-ROM for HP-UX on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your HP-UX documentation. Then change directory as indicated.

```
# cd /cdrom/hp_ux110/root.client
```

3. Copy the AFS initialization file to the local directory for initialization files (by convention, **/sbin/init.d** on HP-UX machines). Note the removal of the **.rc** extension as you copy the file.

```
# cp usr/vice/etc/afs.rc /sbin/init.d/afs
```

4. Copy the file **afs.driver** to the local **/usr/conf/master.d** directory, changing its name to **afs** as you do.

```
# cp usr/vice/etc/afs.driver /usr/conf/master.d/afs
```

5. Copy the AFS kernel module to the local **/usr/conf/lib** directory.

If the machine's kernel supports NFS server functionality:

```
# cp bin/libafs.a /usr/conf/lib
```

If the machine's kernel does not support NFS server functionality, change the file's name as you copy it:

```
# cp bin/libafs.nonfs.a /usr/conf/lib/libafs.a
```

6. Incorporate the AFS driver into the kernel, either using the **SAM** program or a series of individual commands.

- To use the **SAM** program:

- a. Invoke the **SAM** program, specifying the hostname of the local machine as *local_hostname*. The **SAM** graphical user interface pops up.

```
# sam -display local_hostname:0
```

- b. Choose the **Kernel Configuration** icon, then the **Drivers** icon. From the list of drivers, select **afs**.

- c. Open the pull-down **Actions** menu and choose the **Add Driver to Kernel** option.

- d. Open the **Actions** menu again and choose the **Create a New Kernel** option.

- e. Confirm your choices by choosing **Yes** and **OK** when prompted by subsequent pop-up windows. The **SAM** program builds the kernel and reboots the system.

- f. Login again as the superuser **root**.

```
login: root
Password: root_password
```

- To use individual commands:

- a. Edit the file **/stand/system**, adding an entry for **afs** to the `Subsystems` section.

- b. Change to the **/stand/build** directory and issue the **mk_kernel** command to build the kernel.

```
# cd /stand/build
# mk_kernel
```

- c. Move the new kernel to the standard location (**/stand/vmunix**), reboot the machine to start using it, and login again as the superuser **root**.

```
# mv /stand/build/vmunix_test /stand/vmunix
# cd /
# shutdown -r now
login: root
Password: root_password
```

7. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```


8. Use the **SAM** program to create a file system on each partition. For instructions, consult the HP-UX documentation.
9. On some HP-UX systems that use logical volumes, the **SAM** program automatically mounts the partitions. If it has not, mount each partition by issuing either the **mount -a** command to mount all partitions at once or the **mount** command to mount each partition in turn.
10. Create the command configuration file `/sbin/lib/mfsconfig.d/afs`. Use a text editor to place the indicated two lines in it:

```
format_revision 1
fsck             0             m,P,p,d,f,b:c:y,n,Y,N,q,
```

11. Create and change directory to an AFS-specific command directory called `/sbin/fs/afs`.

```
# mkdir /sbin/fs/afs
# cd /sbin/fs/afs
```

12. Copy the AFS-modified version of the **fsck** program (the **vfsc** binary) and related files from the distribution directory to the new AFS-specific command directory.

```
# cp -p /cdrom/hp_ux110/root.server/etc/* .
```

13. Change the **vfsc** binary's name to **fsck** and set the mode bits appropriately on all of the files in the `/sbin/fs/afs` directory.

```
# mv vfsc fsck
# chmod 755 *
```

14. Edit the `/etc/fstab` file, changing the file system type for each AFS server partition from `hfs` to `afs`. This ensures that the AFS-modified **fsck** program runs on the appropriate partitions.

The sixth line in the following example of an edited file shows an AFS server partition, `/vicepa`.

```
/dev/vg00/lvol1 / hfs defaults 0 1
/dev/vg00/lvol4 /opt hfs defaults 0 2
/dev/vg00/lvol5 /tmp hfs defaults 0 2
/dev/vg00/lvol6 /usr hfs defaults 0 2
/dev/vg00/lvol8 /var hfs defaults 0 2
/dev/vg00/lvol9 /vicepa afs defaults 0 2
/dev/vg00/lvol7 /usr/vice/cache hfs defaults 0 2
```

15. If the machine is to remain an AFS client, incorporate AFS into its authentication system, following the instructions in "Enabling AFS Login on HP-UX Systems" on page 22.
16. Proceed to "Starting Server Programs" on page 88.

Getting Started on IRIX Systems

Begin by incorporating AFS modifications into the kernel. Either use the **ml** dynamic loader program, or build a static kernel. Then configure partitions to house AFS volumes. AFS supports use of both EFS and XFS partitions for housing AFS volumes. SGI encourages use of XFS partitions.

You do not need to replace IRIX **fsck** program, because the version that SGI distributes handles AFS volumes properly.

1. Prepare for incorporating AFS into the kernel by performing the following procedures.

- a. Mount the AFS CD-ROM for IRIX on the **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your IRIX documentation. Then change directory as indicated.

```
# cd /cdrom/sgi_65/root.client
```

- b. Copy the AFS initialization script to the local directory for initialization files (by convention, **/etc/init.d** on IRIX machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp -p usr/vice/etc/afs.rc /etc/init.d/afs
```

- c. Issue the **uname -m** command to determine the machine's CPU board type. The **IP_{xx}** value in the output must match one of the supported CPU board types listed in the *IBM AFS Release Notes* for the current version of AFS.

```
# uname -m
```

2. Incorporate AFS into the kernel, either using the **ml** program or by building AFS modifications into a static kernel.

- To use the **ml** program:

- a. Create the local **/usr/vice/etc/sgiload** directory to house the AFS kernel library file.

```
# mkdir /usr/vice/etc/sgiload
```

- b. Copy the appropriate AFS kernel library file to the **/usr/vice/etc/sgiload** directory. The **IP_{xx}** portion of the library file name must match the value previously returned by the **uname -m** command. Also choose the file appropriate to whether the machine's kernel supports NFS server functionality (NFS must be supported for the machine to act as an NFS/AFS Translator). Single- and multiprocessor machines use the same library file.

(You can choose to copy all of the kernel library files into the **/usr/vice/etc/sgiload** directory, but they require a significant amount of space.)

If the machine's kernel supports NFS server functionality:

```
# cp -p usr/vice/etc/sgiload/libafs.IPxx.o /usr/vice/etc/sgiload
```

If the machine's kernel does not support NFS server functionality:

```
# cp -p usr/vice/etc/sgiload/libafs.IPxx.nonfs.o \
    /usr/vice/etc/sgiload
```

- c. Issue the **chkconfig** command to activate the **afsm1** configuration variable.

```
# /etc/chkconfig -f afsm1 on
```

If the machine is to function as an NFS/AFS Translator and the kernel supports NFS server functionality, activate the **afsxnfs** variable.

```
# /etc/chkconfig -f afsxnfs on
```

- d. Run the `/etc/init.d/afs` script to load AFS extensions into the kernel. The script invokes the `ml` command, automatically determining which kernel library file to use based on this machine's CPU type and the activation state of the `afsnfs` variable.

You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/init.d/afs start
```

- e. Proceed to Step "3" on page 84.

- If you prefer to build a kernel, and the machine's hardware and software configuration exactly matches another IRIX machine on which AFS is already built into the kernel, you can copy the kernel from that machine to this one. In general, however, it is better to build AFS modifications into the kernel on each machine according to the following instructions.

- a. Copy the kernel initialization file `afs.sm` to the local `/var/sysgen/system` directory, and the kernel master file `afs` to the local `/var/sysgen/master.d` directory.

```
# cp -p bin/afs.sm /var/sysgen/system
# cp -p bin/afs /var/sysgen/master.d
```

- b. Copy the appropriate AFS kernel library file to the local file `/var/sysgen/boot/afs.a`; the `IPxx` portion of the library file name must match the value previously returned by the `uname -m` command. Also choose the file appropriate to whether the machine's kernel supports NFS server functionality (NFS must be supported for the machine to act as an NFS/AFS Translator). Single- and multiprocessor machines use the same library file.

If the machine's kernel supports NFS server functionality:

```
# cp -p bin/libafs.IPxx.a /var/sysgen/boot/afs.a
```

If the machine's kernel does not support NFS server functionality:

```
# cp -p bin/libafs.IPxx.nonfs.a /var/sysgen/boot/afs.a
```

- c. Issue the `chkconfig` command to deactivate the `afsm1` configuration variable.

```
# /etc/chkconfig -f afsm1 off
```

If the machine is to function as an NFS/AFS Translator and the kernel supports NFS server functionality, activate the `afsnfs` variable.

```
# /etc/chkconfig -f afsxnfs on
```

- d. Copy the existing kernel file, `/unix`, to a safe location. Compile the new kernel, which is created in the file `/unix.install`. It overwrites the existing `/unix` file when the machine reboots in the next step.

```
# cp /unix /unix_noafs
# autoconfig
```

- e. Reboot the machine to start using the new kernel, and login again as the superuser **root**.

```
# cd /
# shutdown -i6 -g0 -y
login: root
Password: root_password
```

3. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

4. Add a line with the following format to the file systems registry file, **/etc/fstab**, for each partition (or logical volume created with the XLV volume manager) to be mounted on one of the directories created in the previous step.

For an XFS partition or logical volume:

```
/dev/dsk/disk /vicepxx xfs rw,raw=/dev/rdisk/disk 0 0
```

For an EFS partition:

```
/dev/dsk/disk /vicepxx efs rw,raw=/dev/rdisk/disk 0 0
```

The following are examples of an entry for each file system type:

```
/dev/dsk/dks0d2s6 /vicepa xfs rw,raw=/dev/rdisk/dks0d2s6 0 0
/dev/dsk/dks0d3s1 /vicepb efs rw,raw=/dev/rdisk/dks0d3s1 0 0
```

5. Create a file system on each partition that is to be mounted on a **/vicep_{xx}** directory. The following commands are probably appropriate, but consult the IRIX documentation for more information. In both cases, *raw_device* is a raw device name like **/dev/rdisk/dks0d0s0** for a single disk partition or **/dev/rxlv/xlv0** for a logical volume.

For XFS file systems, include the indicated options to configure the partition or logical volume with inodes large enough to accommodate AFS-specific information:

```
# mkfs -t xfs -i size=512 -l size=4000b raw_device
```

For EFS file systems:

```
# mkfs -t efs raw_device
```

6. Mount each partition by issuing either the **mount -a** command to mount all partitions at once or the **mount** command to mount each partition in turn.
7. **(Optional)** If you have configured partitions or logical volumes to use XFS, issue the following command to verify that the inodes are configured properly (are large enough to accommodate AFS-specific information). If the configuration is correct, the command returns no output. Otherwise, it specifies the command to run in order to configure each partition or logical volume properly.

```
# /usr/afs/bin/xfs_size_check
```

8. If the machine is to remain an AFS client, incorporate AFS into its authentication system, following the instructions in "Enabling AFS Login on IRIX Systems" on page 28.

9. Proceed to "Starting Server Programs" on page 88.

Getting Started on Linux Systems

Begin by running the AFS initialization script to call the **insmod** program, which dynamically loads AFS modifications into the kernel. Then create partitions for storing AFS volumes. You do not need to replace the Linux **fsck** program.

1. Mount the AFS CD-ROM for Linux on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your Linux documentation. Then change directory as indicated.

```
# cd /cdrom/i386_linux22/root.client/usr/vice/etc
```

2. Copy the AFS kernel library files to the local **/usr/vice/etc/modload** directory. The filenames for the libraries have the format **libafs-version.o**, where *version* indicates the kernel build level. The string **.mp** in the *version* indicates that the file is appropriate for machines running a multiprocessor kernel.

```
# cp -rp modload /usr/vice/etc
```

3. Copy the AFS initialization script to the local directory for initialization files (by convention, **/etc/rc.d/init.d** on Linux machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp -p afs.rc /etc/rc.d/init.d/afs
```

4. Run the AFS initialization script to load AFS extensions into the kernel. You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/rc.d/init.d/afs start
```

5. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

6. Add a line with the following format to the file systems registry file, **/etc/fstab**, for each directory just created. The entry maps the directory name to the disk partition to be mounted on it.

```
/dev/disk /vicepxx ext2 defaults 0 2
```

The following is an example for the first partition being configured.

```
/dev/sda8 /vicepa ext2 defaults 0 2
```

7. Create a file system on each partition that is to be mounted at a **/vicep_{xx}** directory. The following command is probably appropriate, but consult the Linux documentation for more information.

```
# mkfs -v /dev/disk
```

8. Mount each partition by issuing either the **mount -a** command to mount all partitions at once or the **mount** command to mount each partition in turn.

9. If the machine is to remain an AFS client, incorporate AFS into its authentication system, following the instructions in "Enabling AFS Login on Linux Systems" on page 30.
10. Proceed to "Starting Server Programs" on page 88.

Getting Started on Solaris Systems

Begin by running the AFS initialization script to call the **modload** program, which dynamically loads AFS modifications into the kernel. Then configure partitions and replace the Solaris **fsck** program with a version that correctly handles AFS volumes.

1. Mount the AFS CD-ROM for Solaris on the **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your Solaris documentation. Then change directory as indicated.

```
# cd /cdrom/sun4x_56/root.client/usr/vice/etc
```

2. Copy the AFS initialization script to the local directory for initialization files (by convention, **/etc/init.d** on Solaris machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp -p afs.rc /etc/init.d/afs
```

3. Copy the appropriate AFS kernel library file to the local file **/kernel/fs/afs**.

If the machine is running Solaris 2.6 or the 32-bit version of Solaris 7, its kernel supports NFS server functionality, and the **nfsd** process is running:

```
# cp -p modload/libafs.o /kernel/fs/afs
```

If the machine is running Solaris 2.6 or the 32-bit version of Solaris 7, and its kernel does not support NFS server functionality or the **nfsd** process is not running:

```
# cp -p modload/libafs.nonfs.o /kernel/fs/afs
```

If the machine is running the 64-bit version of Solaris 7, its kernel supports NFS server functionality, and the **nfsd** process is running:

```
# cp -p modload/libafs64.o /kernel/fs/sparcv9/afs
```

If the machine is running the 64-bit version of Solaris 7, and its kernel does not support NFS server functionality or the **nfsd** process is not running:

```
# cp -p modload/libafs64.nonfs.o /kernel/fs/sparcv9/afs
```

4. Run the AFS initialization script to load AFS modifications into the kernel. You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/init.d/afs start
```

When an entry called **afs** does not already exist in the local **/etc/name_to_sysnum** file, the script automatically creates it and reboots the machine to start using the new version of the file. If this happens, log in again as the superuser **root** after the reboot and run the initialization script again. This time the required entry exists in the **/etc/name_to_sysnum** file, and the **modload** program runs.

```
login: root
Password: root_password
# /etc/init.d/afs start
```

5. Create the `/usr/lib/fs/afs` directory to house the AFS-modified `fsck` program and related files.

```
# mkdir /usr/lib/fs/afs
# cd /usr/lib/fs/afs
```

6. Copy the `vfsc` binary to the newly created directory, changing the name as you do so.

```
# cp /cdrom/sun4x_56/root.server/etc/vfsc fsck
```

7. Working in the `/usr/lib/fs/afs` directory, create the following links to Solaris libraries:

```
# ln -s /usr/lib/fs/ufs/clri
# ln -s /usr/lib/fs/ufs/df
# ln -s /usr/lib/fs/ufs/edquota
# ln -s /usr/lib/fs/ufs/ff
# ln -s /usr/lib/fs/ufs/fsdb
# ln -s /usr/lib/fs/ufs/fsirand
# ln -s /usr/lib/fs/ufs/fstyp
# ln -s /usr/lib/fs/ufs/labelit
# ln -s /usr/lib/fs/ufs/lockfs
# ln -s /usr/lib/fs/ufs/mkfs
# ln -s /usr/lib/fs/ufs/mount
# ln -s /usr/lib/fs/ufs/ncheck
# ln -s /usr/lib/fs/ufs/newfs
# ln -s /usr/lib/fs/ufs/quot
# ln -s /usr/lib/fs/ufs/quota
# ln -s /usr/lib/fs/ufs/quotaoff
# ln -s /usr/lib/fs/ufs/quotaon
# ln -s /usr/lib/fs/ufs/repquota
# ln -s /usr/lib/fs/ufs/tunefs
# ln -s /usr/lib/fs/ufs/ufsdump
# ln -s /usr/lib/fs/ufs/ufsrestore
# ln -s /usr/lib/fs/ufs/volcopy
```

8. Append the following line to the end of the file `/etc/dfs/fstypes`.

```
afs AFS Utilities
```

9. Edit the `/sbin/mountall` file, making two changes.

- Add an entry for AFS to the `case` statement for option 2, so that it reads as follows:

```
case "$2" in
ufs)    foptions="-o p"
        ;;
afs)    foptions="-o p"
        ;;
s5)    foptions="-y -t /var/tmp/tmp$$ -D"
        ;;
*)    foptions="-y"
        ;;
```

- Edit the file so that all AFS and UFS partitions are checked in parallel. Replace the following section of code:

```
# For fsck purposes, we make a distinction between ufs and
# other file systems
#
if [ "$fstype" = "ufs" ]; then
    ufs_fscklist="$ufs_fscklist $fsckdev"
    saveentry $fstype "$OPTIONS" $special $mountp
    continue
fi
```

with the following section of code:

```
# For fsck purposes, we make a distinction between ufs/afs
# and other file systems.
#
if [ "$fstype" = "ufs" -o "$fstype" = "afs" ]; then
    ufs_fscklist="$ufs_fscklist $fsckdev"
    saveentry $fstype "$OPTIONS" $special $mountp
    continue
fi
```

10. Create a directory called **/vicep_{xx}** for each AFS server partition you are configuring (there must be at least one). Repeat the command for each partition.

```
# mkdir /vicepxx
```

11. Add a line with the following format to the file systems registry file, **/etc/vfstab**, for each partition to be mounted on a directory created in the previous step. Note the value **afs** in the fourth field, which tells Solaris to use the AFS-modified **fsck** program on this partition.

```
/dev/dsk/disk /dev/rdisk/disk /vicepxx afs boot_order yes
```

The following is an example for the first partition being configured.

```
/dev/dsk/c0t6d0s1 /dev/rdisk/c0t6d0s1 /vicepa afs 3 yes
```

12. Create a file system on each partition that is to be mounted at a **/vicep_{xx}** directory. The following command is probably appropriate, but consult the Solaris documentation for more information.

```
# newfs -v /dev/rdisk/disk
```

13. Issue the **mountall** command to mount all partitions at once.
14. If the machine is to remain an AFS client, incorporate AFS into its authentication system, following the instructions in "Enabling AFS Login and Editing the File Systems Clean-up Script on Solaris Systems" on page 35.
15. Proceed to "Starting Server Programs" on page 88.

Starting Server Programs

In this section you initialize the BOS Server, the Update Server, the controller process for NTPD, and the **fs** process. You begin by copying the necessary server files to the local disk.

1. Copy file server binaries to the local **/usr/afs/bin** directory.

- On a machine of an existing system type, you can either load files from the AFS CD-ROM or use a remote file transfer protocol to copy files from an existing server machine of the same system type. To load from the CD-ROM, see the instructions just following for a machine of a new system type. If using a remote file transfer protocol, copy the complete contents of the existing server machine's **/usr/afs/bin** directory.
- On a machine of a new system type, you must use the following instructions to copy files from the AFS CD-ROM.
 - a. On the local **/cdrom** directory, mount the AFS CD-ROM for this machine's system type, if it is not already. For instructions on mounting CD-ROMs (either locally or remotely via NFS), consult the operating system documentation.
 - b. Copy files from the CD-ROM to the local **/usr/afs** directory.

```
# cd /cdrom/sysname/root.server/usr/afs
# cp -rp * /usr/afs
```

2. Copy the contents of the **/usr/afs/etc** directory from an existing file server machine, using a remote file transfer protocol such as **ftp** or NFS. If you use a system control machine, it is best to copy the contents of its **/usr/afs/etc** directory. If you choose not to run a system control machine, copy the directory's contents from any existing file server machine.

3. Change to the **/usr/afs/bin** directory and start the BOS Server (**bosserv** process). Include the **-noauth** flag to prevent the AFS processes from performing authorization checking. This is a grave compromise of security; finish the remaining instructions in this section in an uninterrupted pass.

```
# cd /usr/afs/bin
# ./bosserv -noauth &
```

4. If you run a system control machine, create the **upclientetc** process as an instance of the client portion of the Update Server. It accepts updates of the common configuration files stored in the system control machine's **/usr/afs/etc** directory from the **upserver** process (server portion of the Update Server) running on that machine. The cell's first file server machine was installed as the system control machine in "Starting the Server Portion of the Update Server" on page 45. (If you do not run a system control machine, you must update the contents of the **/usr/afs/etc** directory on each file server machine, using the appropriate **bos** commands.)

By default, the Update Server performs updates every 300 seconds (five minutes). Use the **-t** argument to specify a different number of seconds. For the *machine name* argument, substitute the name of the machine you are installing. The command appears on multiple lines here only for legibility reasons.

```
# ./bos create <machine name> upclientetc simple \
  "/usr/afs/bin/upclient <system control machine> \
  [-t <time>] /usr/afs/etc" -cell <cell name> -noauth
```

5. Create an instance of the Update Server to handle distribution of the file server binaries stored in the **/usr/afs/bin** directory.

- If this is the first file server machine of its AFS system type, create the **upserver** process as an instance of the server portion of the Update Server. It distributes its copy of the file server process binaries to the other file server machines of this system type that you install in future. Creating this process makes this machine the binary distribution machine for its type.

```
# ./bos create <machine name> upserver simple \  
    "/usr/afs/bin/upserver -clear /usr/afs/bin" \  
    -cell <cell name> -noauth
```

- If this machine is an existing system type, create the **upclientbin** process as an instance of the client portion of the Update Server. It accepts updates of the AFS binaries from the **upserver** process running on the binary distribution machine for its system type. For distribution to work properly, the **upserver** process must already be running on that machine.

Use the **-clear** argument to specify that the **upclientbin** process requests unencrypted transfer of the binaries in the **/usr/afs/bin** directory. Binaries are not sensitive and encrypting them is time-consuming.

By default, the Update Server performs updates every 300 seconds (five minutes). Use the **-t** argument to specify a different number of seconds.

```
# ./bos create <machine name> upclientbin simple \  
    "/usr/afs/bin/upclient <binary distribution machine> \  
    [-t <time>] -clear /usr/afs/bin" -cell <cell name> -noauth
```

6. Start the **runntp** process, which configures the Network Time Protocol Daemon (NTPD) to choose a database server machine chosen randomly from the local **/usr/afs/etc/CellServDB** file as its time source. In the standard configuration, the first database server machine installed in your cell refers to a time source outside the cell, and serves as the basis for clock synchronization on all server machines.

```
# ./bos create <machine name> runntp simple \  
    /usr/afs/bin/runntp -cell <cell name> -noauth
```

Note: Do not run the **runntp** process if NTPD or another time synchronization protocol is already running on the machine. Some versions of some operating systems run a time synchronization program by default, as detailed in the *IBM AFS Release Notes*.

Attempting to run multiple instances of the NTPD causes an error. Running NTPD together with another time synchronization protocol is unnecessary and can cause instability in the clock setting.

7. Start the **fs** process, which binds together the File Server, Volume Server, and Salvager.

```
# ./bos create <machine name> fs fs \  
    /usr/afs/bin/fileserver /usr/afs/bin/volserver \  
    /usr/afs/bin/salvager
```

```
/usr/afs/bin/salvager -cell <cell name> -noauth
```

Installing Client Functionality

If you want this machine to be a client as well as a server, follow the instructions in this section. Otherwise, skip to "Completing the Installation" on page 93.

Begin by loading the necessary client files to the local disk. Then create the necessary configuration files and start the Cache Manager. For more detailed explanation of the procedures involved, see the corresponding instructions in "Installing the First AFS Machine" on page 7 (in the sections following "Overview: Installing Client Functionality" on page 48).

If another AFS machine of this machine's system type exists, the AFS binaries are probably already accessible in your AFS filesystem (the conventional location is `/afs/cellname/sysname/usr/afsws`). If not, or if this is the first AFS machine of its type, copy the AFS binaries for this system type into an AFS volume by following the instructions in "Storing AFS Binaries in AFS" on page 62. Because this machine is not yet an AFS client, you must perform the procedure on an existing AFS machine. However, remember to perform the final step (linking the local directory `/usr/afsws` to the appropriate location in the AFS file tree) on this machine itself. If you also want to create AFS volumes to house UNIX system binaries for the new system type, see "Storing System Binaries in AFS" on page 66.

1. Copy client binaries and files to the local disk.

- On a machine of an existing system type, you can either load files from the AFS CD-ROM or use a remote file transfer protocol to copy files from an existing server machine of the same system type. To load from the CD-ROM, see the instructions just following for a machine of a new system type. If using a remote file transfer protocol, copy the complete contents of the existing client machine's `/usr/vice/etc` directory.
- On a machine of a new system type, you must use the following instructions to copy files from the AFS CD-ROM.
 - a. On the local `/cdrom` directory, mount the AFS CD-ROM for this machine's system type, if it is not already. For instructions on mounting CD-ROMs (either locally or remotely via NFS), consult the operating system documentation.
 - b. Copy files to the local `/usr/vice/etc` directory.

This step places a copy of the AFS initialization script (and related files, if applicable) into the `/usr/vice/etc` directory. In the preceding instructions for incorporating AFS into the kernel, you copied the script directly to the operating system's conventional location for initialization files. When you incorporate AFS into the machine's startup sequence in a later step, you can choose to link the two files.

On some system types that use a dynamic kernel loader program, you previously copied AFS library files into a subdirectory of the `/usr/vice/etc` directory. On other system types, you copied the appropriate AFS library file directly to the directory where the operating system accesses it. The following commands do not copy or recopy the AFS library files into the `/usr/vice/etc` directory, because on some system types the library files consume a large

amount of space. If you want to copy them, add the **-r** flag to the first **cp** command and skip the second **cp** command.

```
# cd /cdrom/sysname/root.client/usr/vice/etc
# cp -p * /usr/vice/etc
# cp -rp C /usr/vice/etc
```

2. Change to the **/usr/vice/etc** directory and create the **ThisCell** file as a copy of the **/usr/afs/etc/ThisCell** file. You must first remove the symbolic link to the **/usr/afs/etc/ThisCell** file that the BOS Server created automatically in "Starting Server Programs" on page 88.

```
# cd /usr/vice/etc
# rm ThisCell
# cp /usr/afs/etc/ThisCell ThisCell
```

3. Remove the symbolic link to the **/usr/afs/etc/CellServDB** file.

```
# rm CellServDB
```

4. Create the **/usr/vice/etc/CellServDB** file. Use a network file transfer program such as **ftp** or NFS to copy it from one of the following sources, which are listed in decreasing order of preference:

- Your cell's central **CellServDB** source file (the conventional location is **/afs/cellname/common/etc/CellServDB**)
- The global **CellServDB** file maintained by the AFS Product Support group
- An existing client machine in your cell
- The **CellServDB.sample** file included in the ***sysname*/root.client/usr/vice/etc** directory of each AFS CD-ROM; add an entry for the local cell by following the instructions in "Creating the Client CellServDB File" on page 49

5. Create the **cacheinfo** file for either a disk cache or a memory cache. For a discussion of the appropriate values to record in the file, see "Configuring the Cache" on page 51.

To configure a disk cache, issue the following commands. If you are devoting a partition exclusively to caching, as recommended, you must also configure it, make a file system on it, and mount it at the directory created in this step.

```
# mkdir /usr/vice/cache
# echo "/afs:/usr/vice/cache:#blocks" > cacheinfo
```

To configure a memory cache:

```
# echo "/afs:/usr/vice/cache:#blocks" > cacheinfo
```

6. Create the local directory on which to mount the AFS filespace, by convention **/afs**. If the directory already exists, verify that it is empty.

```
# mkdir /afs
```

7. On AIX systems, add the following line to the **/etc/vfs** file. It enables AIX to unmount AFS correctly during shutdown.

```
afs      4      none      none
```

- On Linux systems, copy the **afsd** options file from the **/usr/vice/etc** directory to the **/etc/sysconfig** directory, removing the **.conf** extension as you do so.

```
# cp /usr/vice/etc/afsd.conf /etc/sysconfig/afsd
```

- Edit the machine's AFS initialization script or **afsd** options file to set appropriate values for **afsd** command parameters. The script resides in the indicated location on each system type:
 - On AIX systems, **/etc/rc.afs**
 - On Digital UNIX systems, **/sbin/init.d/afsd**
 - On HP-UX systems, **/sbin/init.d/afsd**
 - On IRIX systems, **/etc/init.d/afsd**
 - On Linux systems, **/etc/sysconfig/afsd** (the **afsd** options file)
 - On Solaris systems, **/etc/init.d/afsd**

Use one of the methods described in "Configuring the Cache Manager" on page 53 to add the following flags to the **afsd** command line. If you intend for the machine to remain an AFS client, also set any performance-related arguments you wish.

- Add the **-nosettime** flag, because this is a file server machine that is also a client.
- Add the **-memcache** flag if the machine is to use a memory cache.
- Add the **-verbose** flag to display a trace of the Cache Manager's initialization on the standard output stream.

- If appropriate, follow the instructions in "Storing AFS Binaries in AFS" on page 62 to copy the AFS binaries for this system type into an AFS volume. See the introduction to this section for further discussion.

Completing the Installation

At this point you run the machine's AFS initialization script to verify that it correctly loads AFS modifications into the kernel and starts the BOS Server, which starts the other server processes. If you have installed client files, the script also starts the Cache Manager. If the script works correctly, perform the steps that incorporate it into the machine's startup and shutdown sequence. If there are problems during the initialization, attempt to resolve them. The AFS Product Support group can provide assistance if necessary.

If the machine is configured as a client using a disk cache, it can take a while for the **afsd** program to create all of the **Vn** files in the cache directory. Messages on the console trace the initialization process.

- Issue the **bos shutdown** command to shut down the AFS server processes other than the BOS Server. Include the **-wait** flag to delay return of the command shell prompt until all processes shut down completely.

```
# /usr/afs/bin/bos shutdown <machine name> -wait
```

2. Issue the **ps** command to learn the BOS Server's process ID number (PID), and then the **kill** command to stop the **bossver** process.

```
# ps appropriate_ps_options | grep bossver
# kill bossver_PID
```

3. Run the AFS initialization script by issuing the appropriate commands for this system type.

On AIX systems:

- a. Reboot the machine and log in again as the local superuser **root**.

```
# cd /
# shutdown -r now
login: root
Password: root_password
```

- b. Run the AFS initialization script.

```
# /etc/rc.afs
```

- c. Edit the AIX initialization file, **/etc/inittab**, adding the following line to invoke the AFS initialization script. Place it just after the line that starts NFS daemons.

```
rcafs:2:wait:/etc/rc.afs > /dev/console 2>&1 # Start AFS services
```

- d. **(Optional)** There are now copies of the AFS initialization file in both the **/usr/vice/etc** and **/etc** directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm rc.afs
# ln -s /etc/rc.afs
```

- e. Proceed to Step "4" on page 97.

On Digital UNIX systems:

- a. Run the AFS initialization script.

```
# /sbin/init.d/afs start
```

- b. Change to the **/sbin/init.d** directory and issue the **ln -s** command to create symbolic links that incorporate the AFS initialization script into the Digital UNIX startup and shutdown sequence.

```
# cd /sbin/init.d
# ln -s ../init.d/afs /sbin/rc3.d/S67afs
# ln -s ../init.d/afs /sbin/rc0.d/K66afs
```

- c. **(Optional)** There are now copies of the AFS initialization file in both the **/usr/vice/etc** and **/sbin/init.d** directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /sbin/init.d/afs afs.rc
```

d. Proceed to Step "4" on page 97.

On HP-UX systems:

a. Run the AFS initialization script.

```
# /sbin/init.d/afs start
```

b. Change to the `/sbin/init.d` directory and issue the `ln -s` command to create symbolic links that incorporate the AFS initialization script into the HP-UX startup and shutdown sequence.

```
# cd /sbin/init.d
# ln -s ../init.d/afs /sbin/rc2.d/S460afs
# ln -s ../init.d/afs /sbin/rc2.d/K800afs
```

c. **(Optional)** There are now copies of the AFS initialization file in both the `/usr/vice/etc` and `/sbin/init.d` directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /sbin/init.d/afs afs.rc
```

d. Proceed to Step "4" on page 97.

On IRIX systems:

a. If you have configured the machine to use the `ml` dynamic loader program, reboot the machine and log in again as the local superuser `root`.

```
# cd /
# shutdown -i6 -g0 -y
login: root
Password: root_password
```

b. Issue the `chkconfig` command to activate the `afsserver` configuration variable.

```
# /etc/chkconfig -f afsserver on
```

If you have configured this machine as an AFS client and want to it remain one, also issue the `chkconfig` command to activate the `afsclient` configuration variable.

```
# /etc/chkconfig -f afsclient on
```

c. Run the AFS initialization script.

```
# /etc/init.d/afs start
```

- d. Change to the **/etc/init.d** directory and issue the **ln -s** command to create symbolic links that incorporate the AFS initialization script into the IRIX startup and shutdown sequence.

```
# cd /etc/init.d
# ln -s ../init.d/afs /etc/rc2.d/S35afs
# ln -s ../init.d/afs /etc/rc0.d/K35afs
```

- e. **(Optional)** There are now copies of the AFS initialization file in both the **/usr/vice/etc** and **/etc/init.d** directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /etc/init.d/afs afs.rc
```

- f. Proceed to Step "4" on page 97.

On Linux systems:

- a. Reboot the machine and log in again as the local superuser **root**.

```
# cd /
# shutdown -r now
login: root
Password: root_password
```

- b. Run the AFS initialization script.

```
# /etc/rc.d/init.d/afs start
```

- c. Issue the **chkconfig** command to activate the **afs** configuration variable. Based on the instruction in the AFS initialization file that begins with the string **#chkconfig**, the command automatically creates the symbolic links that incorporate the script into the Linux startup and shutdown sequence.

```
# /sbin/chkconfig --add afs
```

- d. **(Optional)** There are now copies of the AFS initialization file in both the **/usr/vice/etc** and **/etc/rc.d/init.d** directories, and copies of the **afsd** options file in both the **/usr/vice/etc** and **/etc/sysconfig** directories. If you want to avoid potential confusion by guaranteeing that the two copies of each file are always the same, create a link between them. You can always retrieve the original script or options file from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc afs.conf
# ln -s /etc/rc.d/init.d/afs afs.rc
# ln -s /etc/sysconfig/afs afs.conf
```

- e. Proceed to Step "4" on page 97.

On Solaris systems:

- a. Reboot the machine and log in again as the local superuser **root**.

```
# cd /
# shutdown -i6 -g0 -y
login: root
Password: root_password
```

- b. Run the AFS initialization script.

```
# /etc/init.d/afs start
```

- c. Change to the **/etc/init.d** directory and issue the **ln -s** command to create symbolic links that incorporate the AFS initialization script into the Solaris startup and shutdown sequence.

```
# cd /etc/init.d
# ln -s ../init.d/afs /etc/rc3.d/S99afs
# ln -s ../init.d/afs /etc/rc0.d/K66afs
```

- d. **(Optional)** There are now copies of the AFS initialization file in both the **/usr/vice/etc** and **/etc/init.d** directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /etc/init.d/afs afs.rc
```

4. Verify that **/usr/afs** and its subdirectories on the new file server machine meet the ownership and mode bit requirements outlined in "Protecting Sensitive AFS Directories" on page 70. If necessary, use the **chmod** command to correct the mode bits.
5. To configure this machine as a database server machine, proceed to "Installing Database Server Functionality" on page 97.

Installing Database Server Functionality

This section explains how to install database server functionality. Database server machines have two defining characteristics. First, they run the Authentication Server, Protection Server, and Volume Location (VL) Server processes. They also run the Backup Server if the cell uses the AFS Backup System, as is assumed in these instructions. Second, they appear in the **CellServDB** file of every machine in the cell (and of client machines in foreign cells, if they are to access files in this cell).

Note the following requirements for database server machines.

- In the conventional configuration, database server machines also serve as file server machines (run the File Server, Volume Server and Salvager processes). If you choose not to run file server functionality on a database server machine, then the kernel does not have to incorporate AFS modifications, but the local **/usr/afs** directory must house most of the standard files and subdirectories. In particular, the

/usr/afs/etc/KeyFile file must contain the same keys as all other server machines in the cell. If you run a system control machine, run the **upclientetc** process on every database server machine other than the system control machine; if you do not run a system control machine, use the **bos addkey** command as instructed in the chapter in the *IBM AFS Administration Guide* about maintaining server encryption keys.

The instructions in this section assume that the machine on which you are installing database server functionality is already a file server machine. Contact the AFS Product Support group to learn how to install database server functionality on a non-file server machine.

- During the installation of database server functionality, you must restart all of the database server machines to force the election of a new Ubik coordinator (synchronization site) for each database server process. This can cause a system outage, which usually lasts less than 5 minutes.
- Updating the kernel memory list of database server machines on each client machine is generally the most time-consuming part of installing a new database server machine. It is, however, crucial for correct functioning in your cell. Incorrect knowledge of your cell's database server machines can prevent your users from authenticating, accessing files, and issuing AFS commands.

You update a client's kernel memory list by changing the **/usr/vice/etc/CellServDB** file and then either rebooting or issuing the **fs newcell** command. For instructions, see the chapter in the *IBM AFS Administration Guide* about administering client machines.

The point at which you update your clients' knowledge of database server machines depends on which of the database server machines has the lowest IP address. The following instructions indicate the appropriate place to update your client machines in either case.

- If the new database server machine has a lower IP address than any existing database server machine, update the **CellServDB** file on every client machine before restarting the database server processes. If you do not, users can become unable to update (write to) any of the AFS databases. This is because the machine with the lowest IP address is usually elected as the Ubik coordinator, and only the Coordinator accepts database writes. On client machines that do not have the new list of database server machines, the Cache Manager cannot locate the new coordinator. (Be aware that if clients contact the new coordinator before it is actually in service, they experience a timeout before contacting another database server machine. This is a minor, and temporary, problem compared to being unable to write to the database.)
- If the new database server machine does not have the lowest IP address of any database server machine, then it is better to update clients after restarting the database server processes. Client machines do not start using the new database server machine until you update their kernel memory list, but that does not usually cause timeouts or update problems (because the new machine is not likely to become the coordinator).

Summary of Procedures

To install a database server machine, perform the following procedures.

1. Install the **bos** suite of commands locally, as a precaution
2. Add the new machine to the **/usr/afs/etc/CellServDB** file on existing file server machines
3. Update your cell's central **CellServDB** source file and the file you make available to foreign cells
4. Update every client machine's **/usr/vice/etc/CellServDB** file and kernel memory list of database server machines
5. Start the database server processes (Authentication Server, Backup Server, Protection Server, and Volume Location Server)
6. Restart the database server processes on every database server machine
7. Notify the AFS Product Support group that you have installed a new database server machine

Instructions

Note: It is assumed that your PATH environment variable includes the directory that houses the AFS command binaries. If not, you possibly need to precede the command names with the appropriate pathname.

1. You can perform the following instructions on either a server or client machine. Login as an AFS administrator who is listed in the **/usr/afs/etc/UserList** file on all server machines.

```
% klog admin_user
Password: admin_password
```

2. If you are working on a client machine configured in the conventional manner, the **bos** command suite resides in the **/usr/afsws/bin** directory, a symbolic link to an AFS directory. An error during installation can potentially block access to AFS, in which case it is helpful to have a copy of the **bos** binary on the local disk. This step is not necessary if you are working on a server machine, where the binary resides in the local **/usr/afs/bin** directory.

```
% cp /usr/afsws/bin/bos /tmp
```

3. Issue the **bos addhost** command to add the new database server machine to the **/usr/afs/etc/CellServDB** file on existing server machines (as well as the new database server machine itself).

Substitute the new database server machine's fully-qualified hostname for the *host name* argument. If you run a system control machine, substitute its fully-qualified hostname for the *machine name* argument. If you do not run a system control machine, repeat the **bos addhost** command once for each server machine in your cell (including the new database server machine

itself), by substituting each one's fully-qualified hostname for the *machine name* argument in turn.

```
% bos addhost <machine name> <host name>
```

If you run a system control machine, wait for the Update Server to distribute the new **CellServDB** file, which takes up to five minutes by default. If you are issuing individual **bos addhost** commands, attempt to issue all of them within five minutes.

Note: It is best to maintain a one-to-one mapping between hostnames and IP addresses on a multihomed database server machine (the conventional configuration for any AFS machine). The BOS Server uses the **gethostbyname()** routine to obtain the IP address associated with the *host name* argument. If there is more than one address, the BOS Server records in the **CellServDB** entry the one that appears first in the list of addresses returned by the routine. The routine possibly returns addresses in a different order on different machines, which can create inconsistency.

4. (Optional) Issue the **bos listhosts** command on each server machine to verify that the new database server machine appears in its **CellServDB** file.

```
% bos listhosts <machine name>
```

5. Add the new database server machine to your cell's central **CellServDB** source file, if you use one. The standard location is **/afs/cellname/common/etc/CellServDB**.

If you are willing to make your cell accessible to users in foreign cells, add the new database server machine to the file that lists your cell's database server machines. The conventional location is **/afs/cellname/service/etc/CellServDB.local**.

6. If this machine's IP address is lower than any existing database server machine's, update every client machine's **/usr/vice/etc/CellServDB** file and kernel memory list to include this machine. (If this machine's IP address is not the lowest, it is acceptable to wait until Step "12" on page 101.)

There are several ways to update the **CellServDB** file on client machines, as detailed in the chapter of the *IBM AFS Administration Guide* about administering client machines. One option is to copy over the central update source (which you updated in Step "5" on page 100), with or without using the **package** program. To update the machine's kernel memory list, you can either reboot after changing the **CellServDB** file or issue the **fs newcell** command.

7. Start the Authentication Server (the **kaserver** process).

```
% bos create <machine name> kaserver simple /usr/afs/bin/kaserver
```

8. Start the Backup Server (the **buserver** process). You must perform other configuration procedures before actually using the AFS Backup System, as detailed in the *IBM AFS Administration Guide*.

```
% bos create <machine name> buserver simple /usr/afs/bin/buserver
```

9. Start the Protection Server (the **ptserver** process).

```
% bos create <machine name> ptserver simple /usr/afs/bin/ptserver
```

10. Start the Volume Location (VL) Server (the **vlserver** process).

```
% bos create <machine name> vlserver simple /usr/afs/bin/vlserver
```

11. Issue the **bos restart** command on every database server machine in the cell, including the new machine. The command restarts the Authentication, Backup, Protection, and VL Servers, which forces an election of a new Ubik coordinator for each process. The new machine votes in the election and is considered as a potential new coordinator.

A cell-wide service outage is possible during the election of a new coordinator for the VL Server, but it normally lasts less than five minutes. Such an outage is particularly likely if you are installing your cell's second database server machine. Messages tracing the progress of the election appear on the console.

Repeat this command on each of your cell's database server machines in quick succession. Begin with the machine with the lowest IP address.

```
% bos restart <machine name> kaserver buserver ptserver vlserver
```

If an error occurs, restart all server processes on the database server machines again by using one of the following methods:

- Issue the **bos restart** command with the **-bosserver** flag for each database server machine
- Reboot each database server machine, either using the **bos exec** command or at its console

12. If you did not update the **CellServDB** file on client machines in Step "6" on page 100, do so now.

13. Send the new database server machine's name and IP address to the AFS Product Support group.

If you wish to participate in the AFS global name space, your cell's entry appear in a **CellServDB** file that the AFS Product Support group makes available to all AFS sites. Otherwise, they list your cell in a private file that they do not share with other AFS sites.

Removing Database Server Functionality

Removing database server machine functionality is nearly the reverse of installing it.

Summary of Procedures

To decommission a database server machine, perform the following procedures.

1. Install the **bos** suite of commands locally, as a precaution
2. Notify the AFS Product Support group that you are decommissioning a database server machine
3. Update your cell's central **CellServDB** source file and the file you make available to foreign cells

4. Update every client machine's **/usr/vice/etc/CellServDB** file and kernel memory list of database server machines
5. Remove the machine from the **/usr/afs/etc/CellServDB** file on file server machines
6. Stop the database server processes and remove them from the **/usr/afs/local/BosConfig** file if desired
7. Restart the database server processes on the remaining database server machines

Instructions

Note: It is assumed that your PATH environment variable includes the directory that houses the AFS command binaries. If not, you possibly need to precede the command names with the appropriate pathname.

1. You can perform the following instructions on either a server or client machine. Login as an AFS administrator who is listed in the **/usr/afs/etc/UserList** file on all server machines.

```
% klog admin_user
Password: admin_password
```

2. If you are working on a client machine configured in the conventional manner, the **bos** command suite resides in the **/usr/afsws/bin** directory, a symbolic link to an AFS directory. An error during installation can potentially block access to AFS, in which case it is helpful to have a copy of the **bos** binary on the local disk. This step is not necessary if you are working on a server machine, where the binary resides in the local **/usr/afs/bin** directory.

```
% cp /usr/afsws/bin/bos /tmp
```

3. Send the revised list of your cell's database server machines to the AFS Product Support group.

This step is particularly important if your cell is included in the global **CellServDB** file. If the administrators in foreign cells do not learn about the change in your cell, they cannot update the **CellServDB** file on their client machines. Users in foreign cells continue to send database requests to the decommissioned machine, which creates needless network traffic and activity on the machine. Also, the users experience time-out delays while their request is forwarded to a valid database server machine.

4. Remove the decommissioned machine from your cell's central **CellServDB** source file, if you use one. The conventional location is **/afs/cellname/common/etc/CellServDB**.

If you maintain a file that users in foreign cells can access to learn about your cell's database server machines, update it also. The conventional location is **/afs/cellname/service/etc/CellServDB.local**.

- Update every client machine's `/usr/vice/etc/CellServDB` file and kernel memory list to exclude this machine. Altering the `CellServDB` file and kernel memory list before stopping the actual database server processes avoids possible time-out delays that result when users send requests to a decommissioned database server machine that is still listed in the file.

There are several ways to update the `CellServDB` file on client machines, as detailed in the chapter of the *IBM AFS Administration Guide* about administering client machines. One option is to copy over the central update source (which you updated in Step "5" on page 100), with or without using the `package` program. To update the machine's kernel memory list, you can either reboot after changing the `CellServDB` file or issue the `fs newcell` command.

- Issue the `bos removehost` command to remove the decommissioned database server machine from the `/usr/afs/etc/CellServDB` file on server machines.

Substitute the decommissioned database server machine's fully-qualified hostname for the `host name` argument. If you run a system control machine, substitute its fully-qualified hostname for the `machine name` argument. If you do not run a system control machine, repeat the `bos removehost` command once for each server machine in your cell (including the decommissioned database server machine itself), by substituting each one's fully-qualified hostname for the `machine name` argument in turn.

```
% bos removehost <machine name> <host name>
```

If you run a system control machine, wait for the Update Server to distribute the new `CellServDB` file, which takes up to five minutes by default. If issuing individual `bos removehost` commands, attempt to issue all of them within five minutes.

- (Optional) Issue the `bos listhosts` command on each server machine to verify that the decommissioned database server machine no longer appears in its `CellServDB` file.

```
% bos listhosts <machine name>
```

- Issue the `bos stop` command to stop the database server processes on the machine, by substituting its fully-qualified hostname for the `machine name` argument. The command changes each process's status in the `/usr/afs/local/BosConfig` file to `NotRun`, but does not remove its entry from the file.

```
% bos stop <machine name> kaserver buserver ptserver vlserver
```

- (Optional) Issue the `bos delete` command to remove the entries for database server processes from the `BosConfig` file. This step is unnecessary if you plan to restart the database server functionality on this machine in future.

```
% bos delete <machine name> kaserver buserver ptserver vlserver
```

- Issue the `bos restart` command on every database server machine in the cell, to restart the Authentication, Backup, Protection, and VL Servers. This forces the election of a Ubik coordinator for each process, ensuring that the remaining database server processes recognize that the machine is no longer a database server.

A cell-wide service outage is possible during the election of a new coordinator for the VL Server, but it normally lasts less than five minutes. Messages tracing the progress of the election appear on the console.

Chapter 3. Installing Additional Server Machines

Repeat this command on each of your cell's database server machines in quick succession. Begin with the machine with the lowest IP address.

```
% bos restart <machine name> kaserver buserver ptserver vlserver
```

If an error occurs, restart all server processes on the database server machines again by using one of the following methods:

- Issue the **bos restart** command with the **-bossver** flag for each database server machine
- Reboot each database server machine, either using the **bos exec** command or at its console

Chapter 4. Installing Additional Client Machines

This chapter describes how to install AFS client machines after you have installed the first AFS machine. Some parts of the installation differ depending on whether or not the new client is of the same AFS system type (uses the same AFS binaries) as a previously installed client machine.

Summary of Procedures

1. Incorporate AFS into the machine's kernel
2. Define the machine's cell membership
3. Define cache location and size
4. Create the `/usr/vice/etc/CellServDB` file, which determines which foreign cells the client can access in addition to the local cell
5. Create the `/afs` directory and start the Cache Manager
6. Create and mount volumes for housing AFS client binaries (necessary only for clients of a new system type)
7. Create a link from the local `/usr/afsws` directory to the AFS directory housing the AFS client binaries
8. Modify the machine's authentication system to enable AFS users to obtain tokens at login

Creating AFS Directories on the Local Disk

Create the `/usr/vice/etc` directory on the local disk, to house client binaries and configuration files. Subsequent instructions copy files from the AFS CD-ROM into them. Create the `/cdrom` directory as a mount point for the CD-ROM, if it does not already exist.

```
# mkdir /usr/vice
# mkdir /usr/vice/etc
# mkdir /cdrom
```

Performing Platform-Specific Procedures

Every AFS client machine's kernel must incorporate AFS modifications. Some system types use a dynamic kernel loader program, whereas on other system types you build AFS modifications into a static kernel. Some system types support both methods.

Also modify the machine's authentication system so that users obtain an AFS token as they log into the local file system. Using AFS is simpler and more convenient for your users if you make the modifications on all client machines. Otherwise, users must perform a two-step login procedure (login to the local file system and then issue the `klog` command). For further discussion of AFS authentication, see the chapter in the *IBM AFS Administration Guide* about cell configuration and administration issues.

For convenience, the following sections group the two procedures by system type. Proceed to the appropriate section.

- "Getting Started on AIX Systems" on page 106
- "Getting Started on Digital UNIX Systems" on page 108
- "Getting Started on HP-UX Systems" on page 111
- "Getting Started on IRIX Systems" on page 115
- "Getting Started on Linux Systems" on page 118
- "Getting Started on Solaris Systems" on page 124

Getting Started on AIX Systems

In this section you load AFS into the AIX kernel. Then incorporate AFS modifications into the machine's secondary authentication system, if you wish to enable AFS login.

Loading AFS into the AIX Kernel

The AIX kernel extension facility is the dynamic kernel loader provided by IBM Corporation. AIX does not support incorporation of AFS modifications during a kernel build.

For AFS to function correctly, the kernel extension facility must run each time the machine reboots, so the AFS initialization script (included in the AFS distribution) invokes it automatically. In this section you copy the script to the conventional location and edit it to select the appropriate options depending on whether NFS is also to run.

After editing the script, you run it to incorporate AFS into the kernel. In a later section you verify that the script correctly initializes the Cache Manager, then configure the AIX **inittab** file so that the script runs automatically at reboot.

1. Mount the AFS CD-ROM for AIX on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your AIX documentation. Then change directory as indicated.

```
# cd /cdrom/rs_aix42/root.client/usr/vice/etc
```

2. Copy the AFS kernel library files to the local **/usr/vice/etc/dkload** directory, and the AFS initialization script to the **/etc** directory.

```
# cp -rp dkload /usr/vice/etc
# cp -p rc.afs /etc/rc.afs
```

3. Edit the **/etc/rc.afs** script, setting the **NFS** variable as indicated.

If the machine is not to function as an NFS/AFS Translator, set the **NFS** variable as follows.

```
NFS=$NFS_NONE
```

If the machine is to function as an NFS/AFS Translator and is running AIX 4.2.1 or higher, set the `NFS` variable as follows. Note that NFS must already be loaded into the kernel, which happens automatically on systems running AIX 4.1.1 and later, as long as the file `/etc/exports` exists.

```
NFS=$NFS_IAUTH
```

4. Invoke the `/etc/rc.afs` script to load AFS modifications into the kernel. You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/rc.afs
```

Enabling AFS Login on AIX Systems

Now incorporate AFS into the AIX secondary authentication system.

1. Issue the `ls` command to verify that the `afs_dynamic_auth` and `afs_dynamic_kerbauth` programs are installed in the local `/usr/vice/etc` directory.

```
# ls /usr/vice/etc
```

If the files do not exist, mount the AFS CD-ROM for AIX (if it is not already), change directory as indicated, and copy them.

```
# cd /cdrom/rs_aix42/root.client/usr/vice/etc
# cp -p afs_dynamic* /usr/vice/etc
```

2. Edit the local `/etc/security/user` file, making changes to the indicated stanzas:

- In the default stanza, set the `registry` attribute to **DCE** (not to **AFS**), as follows:

```
registry = DCE
```

- In the default stanza, set the `SYSTEM` attribute as indicated.

If the machine is an AFS client only, set the following value:

```
SYSTEM = "AFS OR (AFS[UNAVAIL] AND compat[SUCCESS])"
```

If the machine is both an AFS and a DCE client, set the following value (it must appear on a single line in the file):

```
SYSTEM = "DCE OR DCE[UNAVAIL] OR AFS OR (AFS[UNAVAIL] \
AND compat[SUCCESS])"
```

- In the `root` stanza, set the `registry` attribute as follows. It enables the local superuser **root** to log into the local file system only, based on the password listed in the local password file.

```
root:
    registry = files
```

3. Edit the local `/etc/security/login.cfg` file, creating or editing the indicated stanzas:

- In the DCE stanza, set the `program` attribute as follows.

If you use the AFS Authentication Server (**kaserver** process):

```
DCE:
    program = /usr/vice/etc/afs_dynamic_auth
```

If you use a Kerberos implementation of AFS authentication:

```
DCE:
    program = /usr/vice/etc/afs_dynamic_kerbauth
```

- In the AFS stanza, set the `program` attribute as follows.

If you use the AFS Authentication Server (**kaserver** process):

```
AFS:
    program = /usr/vice/etc/afs_dynamic_auth
```

If you use a Kerberos implementation of AFS authentication:

```
AFS:
    program = /usr/vice/etc/afs_dynamic_kerbauth
```

4. Proceed to "Loading and Creating Client Files" on page 128.

Getting Started on Digital UNIX Systems

In this section you build AFS into the Digital UNIX kernel. Then incorporate AFS modifications into the machine's Security Integration Architecture (SIA) matrix, if you wish to enable AFS login.

Building AFS into the Digital UNIX Kernel

On Digital UNIX systems, you must build AFS modifications into a new static kernel; Digital UNIX does not support dynamic loading. If the machine's hardware and software configuration exactly matches another Digital UNIX machine on which AFS is already built into the kernel, you can choose to copy the kernel from that machine to this one. In general, however, it is better to build AFS modifications into the kernel on each machine according to the following instructions.

1. Create a copy called **AFS** of the basic kernel configuration file included in the Digital UNIX distribution as `/usr/sys/conf/machine_name`, where `machine_name` is the machine's hostname in all uppercase letters.

```
# cd /usr/sys/conf
# cp machine_name AFS
```

2. Add AFS to the list of options in the configuration file you created in the previous step, so that the result looks like the following:

```

      .
options      UFS
options      NFS
options      AFS
      .
      .

```

3. Add an entry for AFS to two places in the file `/usr/sys/conf/files`.

- Add a line for AFS to the list of `OPTIONS`, so that the result looks like the following:

```

      .
      .
      .
OPTIONS/nfs      optional nfs
OPTIONS/afs      optional afs
OPTIONS/nfs_server optional nfs_server
      .
      .
      .

```

- Add an entry for AFS to the list of `MODULES`, so that the result looks like the following:

```

      .
      .
      .
      .
#
MODULE/nfs_server optional nfs_server Binary
nfs/nfs_server.c   module nfs_server optimize -g3
nfs/nfs3_server.c  module nfs_server optimize -g3
#
MODULE/afs         optional afs Binary
afs/libafs.c       module afs
#

```

4. Add an entry for AFS to two places in the file `/usr/sys/vfs/vfs_conf.c`.

- Add AFS to the list of defined file systems, so that the result looks like the following:

```

      .
      .
#include <afs.h>
#if defined(AFS) && AFS
    extern struct vfsops afs_vfsops;
#endif
      .
      .

```

- Put a declaration for AFS in the `vfssw[]` table's `MOUNT_ADDON` slot, so that the result looks like the following:

```

      .
      .
      .
      &fdfs_vfsops,      "fdfs",      /* 12 = MOUNT_FDFS */
#if defined(AFS)
      &afs_vfsops,      "afs",
#else
      (struct vfsops *)0, "",      /* 13 = MOUNT_ADDON */
#endif

```

```
#if NFS && INFS_DYNAMIC
    &nfs3_vfsops,      "nfsv3", /* 14 = MOUNT_NFS3 */
```

5. Mount the AFS CD-ROM for Digital UNIX on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your Digital UNIX documentation. Then change directory as indicated.

```
# cd /cdrom/alpha_dux40/root.client
```

6. Copy the AFS initialization script to the local directory for initialization files (by convention, **/sbin/init.d** on Digital UNIX machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp usr/vice/etc/afs.rc /sbin/init.d/afs
```

7. Copy the AFS kernel module to the local **/usr/sys/BINARY** directory.

If the machine's kernel supports NFS server functionality:

```
# cp bin/libafs.o /usr/sys/BINARY/afs.mod
```

If the machine's kernel does not support NFS server functionality:

```
# cp bin/libafs.nonfs.o /usr/sys/BINARY/afs.mod
```

8. Configure and build the kernel. Respond to any prompts by pressing **<Return>**. The resulting kernel resides in the file **/sys/AFS/vmunix**.

```
# doconfig -c AFS
```

9. Rename the existing kernel file and copy the new, AFS-modified file to the standard location.

```
# mv /vmunix /vmunix_noafs
# cp /sys/AFS/vmunix /vmunix
```

10. Reboot the machine to start using the new kernel, and login again as the superuser **root**.

```
# cd /
# shutdown -r now
login: root
Password: root_password
```

Enabling AFS Login on Digital UNIX Systems

On Digital UNIX systems, the AFS initialization script automatically incorporates the AFS authentication library file into the Security Integration Architecture (SIA) matrix on the machine, so that users with AFS accounts obtain a token at login. In this section you copy the library file to the appropriate location.

For more information on SIA, see the Digital UNIX reference page for **matrix.conf**, or consult the section on security in your Digital UNIX documentation.

Note: If the machine runs both the DCE and AFS client software, AFS must start after DCE. Consult the AFS initialization script for suggested symbolic links to create for correct ordering. Also, the system startup script order must initialize SIA before any long-running process that uses authentication.

Perform the following steps to enable AFS login.

1. Mount the AFS CD-ROM for Digital UNIX on the local **/cdrom** directory, if it is not already. Change directory as indicated.

```
# cd /cdrom/alpha_dux40/lib/afs
```

2. Copy the appropriate AFS authentication library file to the local **/usr/shlib** directory.

If you use the AFS Authentication Server (**kaserver** process) in the cell:

```
# cp libafssiad.so /usr/shlib
```

If you use a Kerberos implementation of AFS authentication, rename the library file as you copy it:

```
# cp libafssiad.krb.so /usr/shlib/libafssiad.so
```

3. Proceed to "Loading and Creating Client Files" on page 128.

Getting Started on HP-UX Systems

In this section you build AFS into the HP-UX kernel. Then incorporate AFS modifications into the machine's Pluggable Authentication Module (PAM) system, if you wish to enable AFS login.

Building AFS into the HP-UX Kernel

On HP-UX systems, you must build AFS modifications into a new static kernel; HP-UX does not support dynamic loading. If the machine's hardware and software configuration exactly matches another HP-UX machine on which AFS is already built into the kernel, you can choose to copy the kernel from that machine to this one. In general, however, it is better to build AFS modifications into the kernel on each machine according to the following instructions.

1. Move the existing kernel-related files to a safe location.

```
# cp /stand/vmunix /stand/vmunix.noafs  
# cp /stand/system /stand/system.noafs
```

2. Mount the AFS CD-ROM for HP-UX on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your HP-UX documentation. Then change directory as indicated.

```
# cd /cdrom/hp_ux110/root.client
```

3. Copy the AFS initialization file to the local directory for initialization files (by convention, `/sbin/init.d` on HP-UX machines). Note the removal of the `.rc` extension as you copy the file.

```
# cp usr/vice/etc/afs.rc /sbin/init.d/afs
```

4. Copy the file `afs.driver` to the local `/usr/conf/master.d` directory, changing its name to `afs` as you do.

```
# cp usr/vice/etc/afs.driver /usr/conf/master.d/afs
```

5. Copy the AFS kernel module to the local `/usr/conf/lib` directory.

If the machine's kernel supports NFS server functionality:

```
# cp bin/libafs.a /usr/conf/lib
```

If the machine's kernel does not support NFS server functionality, change the file's name as you copy it:

```
# cp bin/libafs.nonfs.a /usr/conf/lib/libafs.a
```

6. Incorporate the AFS driver into the kernel, either using the **SAM** program or a series of individual commands.

- To use the **SAM** program:

- a. Invoke the **SAM** program, specifying the hostname of the local machine as `local_hostname`. The **SAM** graphical user interface pops up.

```
# sam -display local_hostname:0
```

- b. Choose the **Kernel Configuration** icon, then the **Drivers** icon. From the list of drivers, select **afs**.
- c. Open the pull-down **Actions** menu and choose the **Add Driver to Kernel** option.
- d. Open the **Actions** menu again and choose the **Create a New Kernel** option.
- e. Confirm your choices by choosing **Yes** and **OK** when prompted by subsequent pop-up windows. The **SAM** program builds the kernel and reboots the system.
- f. Login again as the superuser **root**.

```
login: root
Password: root_password
```

- To use individual commands:

- a. Edit the file `/stand/system`, adding an entry for **afs** to the `Subsystems` section.

- b. Change to the `/stand/build` directory and issue the `mk_kernel` command to build the kernel.

```
# cd /stand/build
# mk_kernel
```

- c. Move the new kernel to the standard location (`/stand/vmunix`), reboot the machine to start using it, and login again as the superuser **root**.


```
# mv /stand/build/vmunix_test /stand/vmunix
# cd /
# shutdown -r now
login: root
Password: root_password
```

Enabling AFS Login on HP-UX Systems

At this point you incorporate AFS into the operating system's Pluggable Authentication Module (PAM) scheme. PAM integrates all authentication mechanisms on the machine, including login, to provide the security infrastructure for authenticated access to and from the machine.

Explaining PAM is beyond the scope of this document. It is assumed that you understand the syntax and meanings of settings in the PAM configuration file (for example, how the `other` entry works, the effect of marking an entry as `required`, `optional`, or `sufficient`, and so on).

The following instructions explain how to alter the entries in the PAM configuration file for each service for which you wish to use AFS authentication. Other configurations possibly also work, but the instructions specify the recommended and tested configuration.

Note: The instructions specify that you mark each entry as `optional`. However, marking some modules as `optional` can mean that they grant access to the corresponding service even when the user does not meet all of the module's requirements. In some operating system revisions, for example, if you mark as `optional` the module that controls login via a dial-up connection, it allows users to login without providing a password. See the *IBM AFS Release Notes* for a discussion of any limitations that apply to this operating system.

Also, with some operating system versions you must install patches for PAM to interact correctly with certain authentication programs. For details, see the *IBM AFS Release Notes*.

The recommended AFS-related entries in the PAM configuration file make use of one or more of the following three attributes.

try_first_pass

This is a standard PAM attribute that can be included on entries after the first one for a service; it directs the module to use the password that was provided to the first module. For the AFS module, it means that AFS authentication succeeds if the password provided to the module listed first is the user's correct AFS password. For further discussion of this attribute and its alternatives, see the operating system's PAM documentation.

ignore_root

This attribute, specific to the AFS PAM module, directs it to ignore not only the local superuser **root**, but also any user with UID 0 (zero).

setenv_password_expires

This attribute, specific to the AFS PAM module, sets the environment variable `PASSWORD_EXPIRES` to the expiration date of the user's AFS password, which is recorded in the Authentication Database.

Perform the following steps to enable AFS login.

1. Mount the AFS CD-ROM for HP-UX on the `/cdrom` directory, if it is not already. Then change directory as indicated.

```
# cd /usr/lib/security
```

2. Copy the AFS authentication library file to the `/usr/lib/security` directory. Then create a symbolic link to it whose name does not mention the version. Omitting the version eliminates the need to edit the PAM configuration file if you later update the library file.

If you use the AFS Authentication Server (**kaserver** process) in the cell:

```
# cp /cdrom/hp_ux110/lib/pam_afs.so.1 .
# ln -s pam_afs.so.1 pam_afs.so
```

If you use a Kerberos implementation of AFS authentication:

```
# cp /cdrom/hp_ux110/lib/pam_afs.krb.so.1 .
# ln -s pam_afs.krb.so.1 pam_afs.so
```

3. Edit the `Authentication` management section of the HP-UX PAM configuration file, `/etc/pam.conf` by convention. The entries in this section have the value `auth` in their second field.

First edit the standard entries, which refer to the HP-UX PAM module (usually, the file `/usr/lib/security/libpam_unix.1`) in their fourth field. For each service for which you want to use AFS authentication, edit the third field of its entry to read `optional`. The `pam.conf` file in the HP-UX distribution usually includes standard entries for the **login** and **ftp** services, for instance.

If there are services for which you want to use AFS authentication, but for which the `pam.conf` file does not already include a standard entry, you must create that entry and place the value `optional` in its third field. For instance, the HP-UX `pam.conf` file does not usually include standard entries for the **remsh** or **telnet** services.

Then create an AFS-related entry for each service, placing it immediately below the standard entry. The following example shows what the `Authentication` Management section looks like after you have edited or created entries for the services mentioned previously. Note that the example AFS entries appear on two lines only for legibility.

```
login  auth  optional  /usr/lib/security/libpam_unix.1
login  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root  setenv_password_expires
ftp    auth  optional  /usr/lib/security/libpam_unix.1
ftp    auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
remsh  auth  optional  /usr/lib/security/libpam_unix.1
remsh  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
```

```
telnet  auth  optional  /usr/lib/security/libpam_unix.1
telnet  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass  ignore_root  setenv_password_expires
```

4. If you use the Common Desktop Environment (CDE) on the machine and want users to obtain an AFS token as they log in, also add or edit the following four entries in the `Authentication` management section. Note that the AFS-related entries appear on two lines here only for legibility.

```
dtlogin  auth  optional  /usr/lib/security/libpam_unix.1
dtlogin  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass  ignore_root
dtaction  auth  optional  /usr/lib/security/libpam_unix.1
dtaction  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass  ignore_root
```

5. Proceed to "Loading and Creating Client Files" on page 128.

Getting Started on IRIX Systems

In this section you incorporate AFS into the IRIX kernel, choosing one of two methods:

- Dynamic loading using the **ml** program distributed by Silicon Graphics, Incorporated (SGI).
- Building a new static kernel.

Then see "Enabling AFS Login on IRIX Systems" on page 118 to read about integrated AFS login on IRIX systems.

In preparation for either dynamic loading or kernel building, perform the following procedures:

1. Mount the AFS CD-ROM for IRIX on the **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your IRIX documentation. Then change directory as indicated.

```
# cd /cdrom/sgi_65/root.client
```

2. Copy the AFS initialization script to the local directory for initialization files (by convention, **/etc/init.d** on IRIX machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp -p  usr/vice/etc/afs.rc  /etc/init.d/afs
```

3. Issue the **uname -m** command to determine the machine's CPU board type. The **IP_{xx}** value in the output must match one of the supported CPU board types listed in the *IBM AFS Release Notes* for the current version of AFS.

```
# uname -m
```

4. Proceed to either "Loading AFS into the IRIX Kernel" on page 116 or "Building AFS into the IRIX Kernel" on page 117.

Loading AFS into the IRIX Kernel

The **ml** program is the dynamic kernel loader provided by SGI for IRIX systems. If you use it rather than building AFS modifications into a static kernel, then for AFS to function correctly the **ml** program must run each time the machine reboots. Therefore, the AFS initialization script (included on the AFS CD-ROM) invokes it automatically when the **afsml** configuration variable is activated. In this section you activate the variable and run the script.

In a later section you verify that the script correctly initializes the Cache Manager, then create the links that incorporate AFS into the IRIX startup and shutdown sequence.

1. Create the local **/usr/vice/etc/sgiload** directory to house the AFS kernel library file.

```
# mkdir /usr/vice/etc/sgiload
```

2. Copy the appropriate AFS kernel library file to the **/usr/vice/etc/sgiload** directory. The **IP_{xx}** portion of the library file name must match the value previously returned by the **uname -m** command. Also choose the file appropriate to whether the machine's kernel supports NFS server functionality (NFS must be supported for the machine to act as an NFS/AFS Translator). Single- and multiprocessor machines use the same library file.

(You can choose to copy all of the kernel library files into the **/usr/vice/etc/sgiload** directory, but they require a significant amount of space.)

If the machine's kernel supports NFS server functionality:

```
# cp -p usr/vice/etc/sgiload/libafs.IPxx.o /usr/vice/etc/sgiload
```

If the machine's kernel does not support NFS server functionality:

```
# cp -p usr/vice/etc/sgiload/libafs.IPxx.nonfs.o \
    /usr/vice/etc/sgiload
```

3. Issue the **chkconfig** command to activate the **afsml** configuration variable.

```
# /etc/chkconfig -f afsml on
```

If the machine is to function as an NFS/AFS Translator and the kernel supports NFS server functionality, activate the **afsxnfs** variable.

```
# /etc/chkconfig -f afsxnfs on
```

4. Run the **/etc/init.d/afs** script to load AFS extensions into the kernel. The script invokes the **ml** command, automatically determining which kernel library file to use based on this machine's CPU type and the activation state of the **afsxnfs** variable.

You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/init.d/afs start
```

5. Proceed to "Enabling AFS Login on IRIX Systems" on page 118.

Building AFS into the IRIX Kernel

If you prefer to build a kernel, and the machine's hardware and software configuration exactly matches another IRIX machine on which AFS is already built into the kernel, you can choose to copy the kernel from that machine to this one. In general, however, it is better to build AFS modifications into the kernel on each machine according to the following instructions.

1. Copy the kernel initialization file **afs.sm** to the local **/var/sysgen/system** directory, and the kernel master file **afs** to the local **/var/sysgen/master.d** directory.

```
# cp -p bin/afs.sm /var/sysgen/system
# cp -p bin/afs /var/sysgen/master.d
```

2. Copy the appropriate AFS kernel library file to the local file **/var/sysgen/boot/afs.a**; the **IP_{xx}** portion of the library file name must match the value previously returned by the **uname -m** command. Also choose the file appropriate to whether the machine's kernel supports NFS server functionality (NFS must be supported for the machine to act as an NFS/AFS Translator). Single- and multiprocessor machines use the same library file.

If the machine's kernel supports NFS server functionality:

```
# cp -p bin/libafs.IPxx.a /var/sysgen/boot/afs.a
```

If the machine's kernel does not support NFS server functionality:

```
# cp -p bin/libafs.IPxx.nonfs.a /var/sysgen/boot/afs.a
```

3. Issue the **chkconfig** command to deactivate the **afsm1** configuration variable.

```
# /etc/chkconfig -f afsm1 off
```

If the machine is to function as an NFS/AFS Translator and the kernel supports NFS server functionality, activate the **afsxnfs** variable.

```
# /etc/chkconfig -f afsxnfs on
```

4. Copy the existing kernel file, **/unix**, to a safe location. Compile the new kernel, which is created in the file **/unix.install**. It overwrites the existing **/unix** file when the machine reboots in the next step.

```
# cp /unix /unix_noafs
# autoconfig
```

5. Reboot the machine to start using the new kernel, and login again as the superuser **root**.

```
# cd /
# shutdown -i6 -g0 -y
```

```
login: root
Password: root_password
```

6. Proceed to "Enabling AFS Login on IRIX Systems" on page 118.

Enabling AFS Login on IRIX Systems

The standard IRIX command-line **login** program and the graphical **xdm** login program both automatically grant an AFS token when AFS is incorporated into the machine's kernel. However, some IRIX distributions use another login utility by default, and it does not necessarily incorporate the required AFS modifications. If that is the case, you must disable the default utility if you want AFS users to obtain AFS tokens at login. For further discussion, see the *IBM AFS Release Notes*.

If you configure the machine to use an AFS-modified login utility, then the **afsauthlib.so** and **afskauthlib.so** files (included in the AFS distribution) must reside in the **/usr/vice/etc** directory. Issue the **ls** command to verify.

```
# ls /usr/vice/etc
```

If the files do not exist, mount the AFS CD-ROM for IRIX (if it is not already), change directory as indicated, and copy them.

```
# cd /cdrom/sgi_65/root.client/usr/vice/etc
# cp -p *authlib* /usr/vice/etc
```

After taking any necessary action, proceed to "Loading and Creating Client Files" on page 128.

Getting Started on Linux Systems

In this section you load AFS into the Linux kernel. Then incorporate AFS modifications into the machine's Pluggable Authentication Module (PAM) system, if you wish to enable AFS login.

Loading AFS into the Linux Kernel

The **insmod** program is the dynamic kernel loader for Linux. Linux does not support incorporation of AFS modifications during a kernel build.

For AFS to function correctly, the **insmod** program must run each time the machine reboots, so the AFS initialization script (included on the AFS CD-ROM) invokes it automatically. The script also includes commands that select the appropriate AFS library file automatically. In this section you run the script.

In a later section you also verify that the script correctly initializes the Cache Manager, then activate a configuration variable, which results in the script being incorporated into the Linux startup and shutdown sequence.

1. Mount the AFS CD-ROM for Linux on the local **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your Linux documentation. Then change directory as indicated.

```
# cd /cdrom/i386_linux22/root.client/usr/vice/etc
```

2. Copy the AFS kernel library files to the local **/usr/vice/etc/modload** directory. The filenames for the libraries have the format **libafs-version.o**, where *version* indicates the kernel build level. The string **.mp** in the *version* indicates that the file is appropriate for machines running a multiprocessor kernel.

```
# cp -rp modload /usr/vice/etc
```

3. Copy the AFS initialization script to the local directory for initialization files (by convention, **/etc/rc.d/init.d** on Linux machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp -p afs.rc /etc/rc.d/init.d/afs
```

4. Run the AFS initialization script to load AFS extensions into the kernel. You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/rc.d/init.d/afs start
```

Enabling AFS Login on Linux Systems

At this point you incorporate AFS into the operating system's Pluggable Authentication Module (PAM) scheme. PAM integrates all authentication mechanisms on the machine, including login, to provide the security infrastructure for authenticated access to and from the machine.

Explaining PAM is beyond the scope of this document. It is assumed that you understand the syntax and meanings of settings in the PAM configuration file (for example, how the `other` entry works, the effect of marking an entry as `required`, `optional`, or `sufficient`, and so on).

The following instructions explain how to alter the entries in the PAM configuration file for each service for which you wish to use AFS authentication. Other configurations possibly also work, but the instructions specify the recommended and tested configuration.

The recommended AFS-related entries in the PAM configuration file make use of one or more of the following three attributes.

Authentication Management

try_first_pass

This is a standard PAM attribute that can be included on entries after the first one for a service; it directs the module to use the password that was provided to the first module. For the AFS module, it means that AFS authentication succeeds if the password provided to the module listed first is the user's correct AFS password. For further discussion of this attribute and its alternatives, see the operating system's PAM documentation.

ignore_root

This attribute, specific to the AFS PAM module, directs it to ignore not only the local superuser **root**, but also any user with UID 0 (zero).

ignore_uid uid

This option is an extension of the "ignore_root" switch. The additional parameter is a limit. Users with a uid up to the given parameter are ignored by *pam_afs.so*. Thus, a system administrator still has the opportunity to add local user accounts to his system by choosing between "low" and "high" user ids. An example */etc/passwd* file for "ignore_uid 100" may have entries like these:

```
.
.
afsuserone:x:99:100::/afs/afscell/u/afsuserone:/bin/bash
afsusertwo:x:100:100::/afs/afscell/u/afsusertwo:/bin/bash
localuserone:x:101:100::/home/localuserone:/bin/bash
localusertwo:x:102:100::/home/localusertwo:/bin/bash
.
.
```

AFS accounts should be locked in the file */etc/shadow* like this:

```
.
.
afsuserone:!!:11500:0:99999:7:::
afsusertwo:!!:11500:0:99999:7:::
localuserone:<thelocaluserone'skey>:11500:0:99999:7:::
localusertwo:<thelocalusertwo'skey>:11500:0:99999:7:::
.
.
```

There is no need to store a local key in this file since the AFS password is sent and verified at the AFS cell server!

setenv_password_expires

This attribute, specific to the AFS PAM module, sets the environment variable **PASSWORD_EXPIRES** to the expiration date of the user's AFS password, which is recorded in the Authentication Database.

set_token

Some applications don't call *pam_setcred()* in order to retrieve the appropriate credentials (here the AFS token) for their session. This switch sets the credentials already in *pam_sm_authenticate()* obsoleting a call to *pam_setcred()*. **Caution: Don't use this switch for applications which do call *pam_setcred()*!** One example for an application not calling *pam_setcred()* are older versions of the samba server. Nevertheless, using applications with working pam session management is recommended as this setup conforms better with the PAM definitions.

refresh_token

This options is identical to "set_token" except that no new PAG is generated. This is necessary to handle processes like xlock or xscreensaver. It is not enough to give the screen and the keyboard free for the user who reactivated his screen typing in the correct AFS password, but one may also

need fresh tokens with full lifetime in order to work on, and the new token must be refreshed in the already existing PAG for the processes that have been started. This is achieved using this option.

use_klog

Activating this switch the authentication is done by calling the external program "klog". One program requiring this is for example *kdm* of KDE 2.x.

dont_fork

Usually, the password verification and the establishment of the token is performed in a sub process. Using this option *pam_afs* does not fork and performs all actions in a single process. **Only use this options in case you notice serious problems caused by the sub process.** This option has been developed in respect to the "mod_auth_pam"-project (see also *mod_auth_pam* (http://pam.sourceforge.net/mod_auth_pam/)). The *mod_auth_pam* module enables PAM authentication for the apache http server package.

Session Management

no_unlog

Normally the tokens are deleted (in memory) after the session ends. Using this options the tokens are left untouched. **This behaviour has been the default in *pam_afs* until *openafs-1.1.1!***

remainlifetime sec

The tokens are kept active for *sec* seconds before they are deleted. X display managers i.e. are used to inform the applications started in the X session before the logout and then end themselves. If the token was deleted immediately the applications would have no chance to write back their settings to i.e. the user's AFS home space. This option may help to avoid the problem.

Perform the following steps to enable AFS login.

1. Mount the AFS CD-ROM for Linux on the **/cdrom** directory, if it is not already. Then change to the directory for PAM modules, which depends on which Linux distribution you are using.

If you are using a Linux distribution from Red Hat Software:

```
# cd /lib/security
```

If you are using another Linux distribution:

```
# cd /usr/lib/security
```

2. Copy the appropriate AFS authentication library file to the directory to which you changed in the previous step. Create a symbolic link whose name does not mention the version. Omitting the version eliminates the need to edit the PAM configuration file if you later update the library file.

If you use the AFS Authentication Server (**kaserver** process):

```
# cp /cdrom/i386_linux22/lib/pam_afs.so.1 .
# ln -s pam_afs.so.1 pam_afs.so
```

If you use a Kerberos implementation of AFS authentication:

```
# cp /cdrom/i386_linux22/lib/pam_afs.krb.so.1 .
# ln -s pam_afs.krb.so.1 pam_afs.so
```

3. For each service with which you want to use AFS authentication, insert an entry for the AFS PAM module into the `auth` section of the service's PAM configuration file. (Linux uses a separate configuration file for each service, unlike some other operating systems which list all services in a single file.) Mark the entry as `sufficient` in the second field.

Place the AFS entry below any entries that impose conditions under which you want the service to fail for a user who does not meet the entry's requirements. Mark these entries `required`. Place the AFS entry above any entries that need to execute only if AFS authentication fails.

Insert the following AFS entry if using the Red Hat distribution:

```
auth sufficient /lib/security/pam_afs.so try_first_pass ignore_root
```

Insert the following AFS entry if using another distribution:

```
auth sufficient /usr/lib/security/pam_afs.so try_first_pass ignore_root
```

Check the PAM config files also for "session" entries. If there are lines beginning with "session" then please insert this line too:

```
session optional /lib/security/pam_afs.so
```

or

```
session optional /usr/lib/security/pam_afs.so
```

This guarantees that the user's tokens are deleted from memory after his session ends so that no other user coincidentally gets those tokens without authorization! The following examples illustrate the recommended configuration of the configuration file for several services:

Authentication Management

(`/etc/pam.d/login`)

```
##PAM-1.0
auth required /lib/security/pam_securetty.so
auth required /lib/security/pam_nologin.so
auth sufficient /lib/security/pam_afs.so try_first_pass ignore_root
#
#This enables AFS authentication for every user but root
auth required /lib/security/pam_pwdb.so shadow nullok
account required /lib/security/pam_pwdb.so
password required /lib/security/pam_cracklib.so
password required /lib/security/pam_pwdb.so shadow nullok use_authtok
session optional /lib/security/pam_afs.so
#Make sure tokens are deleted after the user logs out
session required /lib/security/pam_pwdb.so
```

(/etc/pam.d/samba)

```

auth      required      /lib/security/pam_afs.so ignore_uid 100 set_token
#
#Here, users with uid>100 are considered to belong to the AFS and users
#with uid<=100 are ignored by pam_afs. The token is retrieved already in
#pam_sm_authenticate() (this is an example pam config for a samba version
#that does not call pam_setcred(), it also does no sense to include session
#entries here since they would be ignored by this version of samba ).
account   required      /lib/security/pam_pwdb.so

```

(/etc/pam.d/xscreensaver)

```

auth      sufficient    /lib/security/pam_afs.so ignore_uid 100 refresh_token
#
#Avoid generating a new PAG for the new tokens, use the already existing PAG and
#establish a fresh token in it.
auth      required      /lib/security/pam_pwdb.so try_first_pass

```

(/etc/pam.d/httpd)

```

auth      required      /lib/security/pam_afs.so ignore_uid 100 dont_fork
#
#Don't fork for the verification of the password.

```

Session Management**(/etc/pam.d/su)**

```

auth      sufficient    /lib/security/pam_afs.so ignore_uid 100
auth      required      /lib/security/pam_pwdb.so try_first_pass
account   required      /lib/security/pam_pwdb.so
password  required      /lib/security/pam_cracklib.so
password  required      /lib/security/pam_pwdb.so use_authtok
session   required      /lib/security/pam_pwdb.so
session   optional     /lib/security/pam_afs.so no_unlog
#
#Don't delete the token in this case, since the user may still
#need it (for example if somebody logs in and changes to root
#afterwards he may still want to access his home space in AFS).
session   required      /lib/security/pam_login_access.so
session   optional     /lib/security/pam_xauth.so

```

(/etc/pam.d/xdm)

```
auth      required      /lib/security/pam_nologin.so
auth      required      /lib/security/pam_login_access.so
auth      sufficient    /lib/security/pam_afs.so ignore_uid 100 use_klog
auth      required      /lib/security/pam_pwdb.so try_first_pass
account   required      /lib/security/pam_pwdb.so
password  required      /lib/security/pam_cracklib.so
password  required      /lib/security/pam_pwdb.so shadow nullok use_authok
session   optional      /lib/security/pam_afs.so remainlifetime 10
#
#Wait 10 seconds before deleting the AFS tokens in order to give
#the programs of the X session some time to save their settings
#to AFS.
session   required      /lib/security/pam_pwdb.so
```

4. Proceed to "Loading and Creating Client Files" on page 128.

Getting Started on Solaris Systems

In this section you load AFS into the Solaris kernel. Then incorporate AFS modifications into the machine's Pluggable Authentication Module (PAM) system, if you wish to enable AFS login.

Loading AFS into the Solaris Kernel

The **modload** program is the dynamic kernel loader provided by Sun Microsystems for Solaris systems. Solaris does not support incorporation of AFS modifications during a kernel build.

For AFS to function correctly, the **modload** program must run each time the machine reboots, so the AFS initialization script (included on the AFS CD-ROM) invokes it automatically. In this section you copy the appropriate AFS library file to the location where the **modload** program accesses it and then run the script.

In a later section you verify that the script correctly initializes the Cache Manager, then create the links that incorporate AFS into the Solaris startup and shutdown sequence.

1. Mount the AFS CD-ROM for Solaris on the **/cdrom** directory. For instructions on mounting CD-ROMs (either locally or remotely via NFS), see your Solaris documentation. Then change directory as indicated.

```
# cd /cdrom/sun4x_56/root.client/usr/vice/etc
```

2. Copy the AFS initialization script to the local directory for initialization files (by convention, **/etc/init.d** on Solaris machines). Note the removal of the **.rc** extension as you copy the script.

```
# cp -p afs.rc /etc/init.d/afs
```

- Copy the appropriate AFS kernel library file to the local file `/kernel/fs/afs`.

If the machine is running Solaris 2.6 or the 32-bit version of Solaris 7, its kernel supports NFS server functionality, and the `nfsd` process is running:

```
# cp -p modload/libafs.o /kernel/fs/afs
```

If the machine is running Solaris 2.6 or the 32-bit version of Solaris 7, and its kernel does not support NFS server functionality or the `nfsd` process is not running:

```
# cp -p modload/libafs.nonfs.o /kernel/fs/afs
```

If the machine is running the 64-bit version of Solaris 7, its kernel supports NFS server functionality, and the `nfsd` process is running:

```
# cp -p modload/libafs64.o /kernel/fs/sparcv9/afs
```

If the machine is running the 64-bit version of Solaris 7, and its kernel does not support NFS server functionality or the `nfsd` process is not running:

```
# cp -p modload/libafs64.nonfs.o /kernel/fs/sparcv9/afs
```

- Run the AFS initialization script to load AFS modifications into the kernel. You can ignore any error messages about the inability to start the BOS Server or the Cache Manager or AFS client.

```
# /etc/init.d/afs start
```

When an entry called `afs` does not already exist in the local `/etc/name_to_sysnum` file, the script automatically creates it and reboots the machine to start using the new version of the file. If this happens, log in again as the superuser `root` after the reboot and run the initialization script again. This time the required entry exists in the `/etc/name_to_sysnum` file, and the `modload` program runs.

```
login: root
Password: root_password
# /etc/init.d/afs start
```

Enabling AFS Login on Solaris Systems

At this point you incorporate AFS into the operating system's Pluggable Authentication Module (PAM) scheme. PAM integrates all authentication mechanisms on the machine, including login, to provide the security infrastructure for authenticated access to and from the machine.

Explaining PAM is beyond the scope of this document. It is assumed that you understand the syntax and meanings of settings in the PAM configuration file (for example, how the `other` entry works, the effect of marking an entry as `required`, `optional`, or `sufficient`, and so on).

The following instructions explain how to alter the entries in the PAM configuration file for each service for which you wish to use AFS authentication. Other configurations possibly also work, but the instructions specify the recommended and tested configuration.

Note: The instructions specify that you mark each entry as `optional`. However, marking some modules as optional can mean that they grant access to the corresponding service even when the user does not meet all of the module's requirements. In some operating system revisions, for example, if you mark as optional the module that controls login via a dial-up connection, it allows users to login without providing a password. See the *IBM AFS Release Notes* for a discussion of any limitations that apply to this operating system.

Also, with some operating system versions you must install patches for PAM to interact correctly with certain authentication programs. For details, see the *IBM AFS Release Notes*.

The recommended AFS-related entries in the PAM configuration file make use of one or more of the following three attributes.

Authentication Management

`try_first_pass`

This is a standard PAM attribute that can be included on entries after the first one for a service; it directs the module to use the password that was provided to the first module. For the AFS module, it means that AFS authentication succeeds if the password provided to the module listed first is the user's correct AFS password. For further discussion of this attribute and its alternatives, see the operating system's PAM documentation.

`ignore_root`

This attribute, specific to the AFS PAM module, directs it to ignore not only the local superuser `root`, but also any user with UID 0 (zero).

`setenv_password_expires`

This attribute, specific to the AFS PAM module, sets the environment variable `PASSWORD_EXPIRES` to the expiration date of the user's AFS password, which is recorded in the Authentication Database.

Perform the following steps to enable AFS login.

1. Mount the AFS CD-ROM for Solaris on the `/cdrom` directory, if it is not already. Then change directory as indicated.

```
# cd /usr/lib/security
```

2. Copy the AFS authentication library file to the `/usr/lib/security` directory. Then create a symbolic link to it whose name does not mention the version. Omitting the version eliminates the need to edit the PAM configuration file if you later update the library file.

If you use the AFS Authentication Server (`kaserver` process):

```
# cp /cdrom/sun4x_56/lib/pam_afs.so.1 .
# ln -s pam_afs.so.1 pam_afs.so
```

If you use a Kerberos implementation of AFS authentication:

```
# cp /cdrom/sun4x_56/lib/pam_afs.krb.so.1 .
```

```
# ln -s pam_afs.krb.so.1 pam_afs.so
```

3. Edit the `Authentication` management section of the Solaris PAM configuration file, `/etc/pam.conf` by convention. The entries in this section have the value `auth` in their second field.

First edit the standard entries, which refer to the Solaris PAM module (usually, the file `/usr/lib/security/pam_unix.so.1`) in their fourth field. For each service for which you want to use AFS authentication, edit the third field of its entry to read `optional`. The `pam.conf` file in the Solaris distribution usually includes standard entries for the **login**, **rlogin**, and **rsh** services, for instance.

If there are services for which you want to use AFS authentication, but for which the `pam.conf` file does not already include a standard entry, you must create that entry and place the value `optional` in its third field. For instance, the Solaris `pam.conf` file does not usually include standard entries for the **ftp** or **telnet** services.

Then create an AFS-related entry for each service, placing it immediately below the standard entry. The following example shows what the `Authentication` Management section looks like after you have edited or created entries for the services mentioned previously. Note that the example AFS entries appear on two lines only for legibility.

```
login  auth  optional  /usr/lib/security/pam_unix.so.1
login  auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root setenv_password_expires
rlogin auth  optional  /usr/lib/security/pam_unix.so.1
rlogin auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root setenv_password_expires
rsh    auth  optional  /usr/lib/security/pam_unix.so.1
rsh    auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
ftp    auth  optional  /usr/lib/security/pam_unix.so.1
ftp    auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
telnet auth  optional  /usr/lib/security/pam_unix.so.1
telnet auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root setenv_password_expires
```

4. If you use the Common Desktop Environment (CDE) on the machine and want users to obtain an AFS token as they log in, also add or edit the following four entries in the `Authentication` management section. Note that the AFS-related entries appear on two lines here only for legibility.

```
dtlogin auth  optional  /usr/lib/security/pam_unix.so.1
dtlogin auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
dtsession auth  optional  /usr/lib/security/pam_unix.so.1
dtsession auth  optional  /usr/lib/security/pam_afs.so      \
      try_first_pass ignore_root
```

5. Some Solaris distributions include a script that locates and removes unneeded files from various file systems. Its conventional location is `/usr/lib/fs/nfs/nfsfind`. The script generally uses an argument to the **find** command to define which file systems to search. In this step you modify the command to exclude the `/afs` directory. Otherwise, the command traverses the AFS filespace of every cell that is

accessible from the machine, which can take many hours. The following alterations are possibilities, but you must verify that they are appropriate for your cell.

The first possible alteration is to add the **-local** flag to the existing command, so that it looks like the following:

```
find $dir -local -name .nfs\* -mtime +7 -mount -exec rm -f {} \;
```

Another alternative is to exclude any directories whose names begin with the lowercase letter **a** or a non-alphabetic character.

```
find /[A-Zb-z]* remainder of existing command
```

Do not use the following command, which still searches under the **/afs** directory, looking for a subdirectory of type **4.2**.

```
find / -fstype 4.2 /* do not use */
```

6. Proceed to "Loading and Creating Client Files" on page 128.

Loading and Creating Client Files

Now copy files from the AFS CD-ROM to the **/usr/vice/etc** directory. On some platforms that use a dynamic loader program to incorporate AFS modifications into the kernel, you have already copied over some the files. Copying them again does no harm.

Every AFS client machine has a copy of the **/usr/vice/etc/ThisCell** file on its local disk to define the machine's cell membership for the AFS client programs that run on it. Among other functions, this file determines the following:

- The cell in which users authenticate when they log onto the machine, assuming it is using an AFS-modified login utility
- The cell in which users authenticate by default when they issue the **klog** command
- The cell membership of the AFS server processes that the AFS command interpreters on this machine contact by default

Similarly, the **/usr/vice/etc/CellServDB** file on a client machine's local disk lists the database server machines in each cell that the local Cache Manager can contact. If there is no entry in the file for a cell, or the list of database server machines is wrong, then users working on this machine cannot access the cell. The chapter in the *IBM AFS Administration Guide* about administering client machines explains how to maintain the file after creating it. A version of the client **CellServDB** file was created during the installation of your cell's first machine (in "Creating the Client CellServDB File" on page 49). It is probably also appropriate for use on this machine.

Remember that the Cache Manager consults the **/usr/vice/etc/CellServDB** file only at reboot, when it copies the information into the kernel. For the Cache Manager to perform properly, the **CellServDB** file must be accurate at all times. Refer to the chapter in the *IBM AFS Administration Guide* about administering client machines for instructions on updating this file, with or without rebooting.

1. On the local **/cdrom** directory, mount the AFS CD-ROM for this machine's system type, if it is not already. For instructions on mounting CD-ROMs (either locally or remotely via NFS), consult the operating system documentation.
2. Copy files to the local **/usr/vice/etc** directory.

This step places a copy of the AFS initialization script (and related files, if applicable) into the **/usr/vice/etc** directory. In the preceding instructions for incorporating AFS into the kernel, you copied the script directly to the operating system's conventional location for initialization files. When you incorporate AFS into the machine's startup sequence in a later step, you can choose to link the two files.

On some system types that use a dynamic kernel loader program, you previously copied AFS library files into a subdirectory of the **/usr/vice/etc** directory. On other system types, you copied the appropriate AFS library file directly to the directory where the operating system accesses it. The following commands do not copy or recopy the AFS library files into the **/usr/vice/etc** directory, because on some system types the library files consume a large amount of space. If you want to copy them, add the **-r** flag to the first **cp** command and skip the second **cp** command.

```
# cd /cdrom/sysname/root.client/usr/vice/etc
# cp -p * /usr/vice/etc
# cp -rp C /usr/vice/etc
```

3. Create the **/usr/vice/etc/ThisCell** file.

```
# echo "cellname" > /usr/vice/etc/ThisCell
```

4. Create the **/usr/vice/etc/CellServDB** file. Use a network file transfer program such as **ftp** or NFS to copy it from one of the following sources, which are listed in decreasing order of preference:
 - Your cell's central **CellServDB** source file (the conventional location is **/afs/*cellname*/common/etc/CellServDB**)
 - The global **CellServDB** file maintained by the AFS Product Support group
 - An existing client machine in your cell
 - The **CellServDB.sample** file included in the ***sysname*/root.client/usr/vice/etc** directory of each AFS CD-ROM; add an entry for the local cell by following the instructions in "Creating the Client CellServDB File" on page 49

Configuring the Cache

The Cache Manager uses a cache on the local disk or in machine memory to store local copies of files fetched from file server machines. As the **afsd** program initializes the Cache Manager, it sets basic cache configuration parameters according to definitions in the local **/usr/vice/etc/cacheinfo** file. The file has three fields:

1. The first field names the local directory on which to mount the AFS filespace. The conventional location is the **/afs** directory.

2. The second field defines the local disk directory to use for the disk cache. The conventional location is the **/usr/vice/cache** directory, but you can specify an alternate directory if another partition has more space available. There must always be a value in this field, but the Cache Manager ignores it if the machine uses a memory cache.
3. The third field specifies the number of kilobyte (1024 byte) blocks to allocate for the cache.

The values you define must meet the following requirements.

- On a machine using a disk cache, the Cache Manager expects always to be able to use the amount of space specified in the third field. Failure to meet this requirement can cause serious problems, some of which can be repaired only by rebooting. You must prevent non-AFS processes from filling up the cache partition. The simplest way is to devote a partition to the cache exclusively.
- The amount of space available in memory or on the partition housing the disk cache directory imposes an absolute limit on cache size.
- The maximum supported cache size can vary in each AFS release; see the *IBM AFS Release Notes* for the current version.
- For a disk cache, you cannot specify a value in the third field that exceeds 95% of the space available on the partition mounted at the directory named in the second field. If you violate this restriction, the **afsd** program exits without starting the Cache Manager and prints an appropriate message on the standard output stream. A value of 90% is more appropriate on most machines. Some operating systems (such as AIX) do not automatically reserve some space to prevent the partition from filling completely; for them, a smaller value (say, 80% to 85% of the space available) is more appropriate.
- For a memory cache, you must leave enough memory for other processes and applications to run. If you try to allocate more memory than is actually available, the **afsd** program exits without initializing the Cache Manager and produces the following message on the standard output stream.

```
afsd: memCache allocation failure at number KB
```

The *number* value is how many kilobytes were allocated just before the failure, and so indicates the approximate amount of memory available.

Within these hard limits, the factors that determine appropriate cache size include the number of users working on the machine, the size of the files with which they work, and (for a memory cache) the number of processes that run on the machine. The higher the demand from these factors, the larger the cache needs to be to maintain good performance.

Disk caches smaller than 10 MB do not generally perform well. Machines serving multiple users usually perform better with a cache of at least 60 to 70 MB. The point at which enlarging the cache further does not really improve performance depends on the factors mentioned previously and is difficult to predict.

Memory caches smaller than 1 MB are nonfunctional, and the performance of caches smaller than 5 MB is usually unsatisfactory. Suitable upper limits are similar to those for disk caches but are probably determined more by the demands on memory from other sources on the machine (number of users and processes). Machines running only a few processes possibly can use a smaller memory cache.

Configuring a Disk Cache

Note: Not all file system types that an operating system supports are necessarily supported for use as the cache partition. For possible restrictions, see the *IBM AFS Release Notes*.

To configure the disk cache, perform the following procedures:

1. Create the local directory to use for caching. The following instruction shows the conventional location, **/usr/vice/cache**. If you are devoting a partition exclusively to caching, as recommended, you must also configure it, make a file system on it, and mount it at the directory created in this step.

```
# mkdir /usr/vice/cache
```

2. Create the **cacheinfo** file to define the configuration parameters discussed previously. The following instruction shows the standard mount location, **/afs**, and the standard cache location, **/usr/vice/cache**.

```
# echo "/afs:/usr/vice/cache:#blocks" > /usr/vice/etc/cacheinfo
```

The following example defines the disk cache size as 50,000 KB:

```
# echo "/afs:/usr/vice/cache:50000" > /usr/vice/etc/cacheinfo
```

Configuring a Memory Cache

To configure a memory cache, create the **cacheinfo** file to define the configuration parameters discussed previously. The following instruction shows the standard mount location, **/afs**, and the standard cache location, **/usr/vice/cache** (though the exact value of the latter is irrelevant for a memory cache).

```
# echo "/afs:/usr/vice/cache:#blocks" > /usr/vice/etc/cacheinfo
```

The following example allocates 25,000 KB of memory for the cache.

```
# echo "/afs:/usr/vice/cache:25000" > /usr/vice/etc/cacheinfo
```

Configuring the Cache Manager

By convention, the Cache Manager mounts the AFS filesystem on the local **/afs** directory. In this section you create that directory.

The **afsd** program sets several cache configuration parameters as it initializes the Cache Manager, and starts daemons that improve performance. You can use the **afsd** command's arguments to override the parameters' default values and to change the number of some of the daemons. Depending on the machine's cache size, its amount of RAM, and how many people work on it, you can sometimes improve

Cache Manager performance by overriding the default values. For a discussion of all of the **afsd** command's arguments, see its reference page in the *IBM AFS Administration Reference*.

The **afsd** command line in the AFS initialization script on each system type includes an `OPTIONS` variable. You can use it to set nondefault values for the command's arguments, in one of the following ways:

- You can create an **afsd options file** that sets values for arguments to the **afsd** command. If the file exists, its contents are automatically substituted for the `OPTIONS` variable in the AFS initialization script. The AFS distribution for some system types includes an options file; on other system types, you must create it.

You use two variables in the AFS initialization script to specify the path to the options file: `CONFIG` and `AFSDOPT`. On system types that define a conventional directory for configuration files, the `CONFIG` variable indicates it by default; otherwise, the variable indicates an appropriate location.

List the desired **afsd** options on a single line in the options file, separating each option with one or more spaces. The following example sets the **-stat** argument to 2500, the **-daemons** argument to 4, and the **-volumes** argument to 100.

```
-stat 2500 -daemons 4 -volumes 100
```

- On a machine that uses a disk cache, you can set the `OPTIONS` variable in the AFS initialization script to one of `$SMALL`, `$MEDIUM`, or `$LARGE`. The AFS initialization script uses one of these settings if the **afsd** options file named by the `AFSDOPT` variable does not exist. In the script as distributed, the `OPTIONS` variable is set to the value `$MEDIUM`.

Note: Do not set the `OPTIONS` variable to `$SMALL`, `$MEDIUM`, or `$LARGE` on a machine that uses a memory cache. The arguments it sets are appropriate only on a machine that uses a disk cache.

The script (or on some system types the **afsd** options file named by the `AFSDOPT` variable) defines a value for each of `SMALL`, `MEDIUM`, and `LARGE` that sets **afsd** command arguments appropriately for client machines of different sizes:

- `SMALL` is suitable for a small machine that serves one or two users and has approximately 8 MB of RAM and a 20-MB cache
 - `MEDIUM` is suitable for a medium-sized machine that serves two to six users and has 16 MB of RAM and a 40-MB cache
 - `LARGE` is suitable for a large machine that serves five to ten users and has 32 MB of RAM and a 100-MB cache
- You can choose not to create an **afsd** options file and to set the `OPTIONS` variable in the initialization script to a null value rather than to the default `$MEDIUM` value. You can then either set arguments

directly on the **afsd** command line in the script, or set no arguments (and so accept default values for all Cache Manager parameters).

1. Create the local directory on which to mount the AFS filesystem, by convention **/afs**. If the directory already exists, verify that it is empty.

```
# mkdir /afs
```

2. On AIX systems, add the following line to the **/etc/vfs** file. It enables AIX to unmount AFS correctly during shutdown.

```
afs      4      none      none
```

3. On Linux systems, copy the **afsd** options file from the **/usr/vice/etc** directory to the **/etc/sysconfig** directory, removing the **.conf** extension as you do so.

```
# cp /usr/vice/etc/afsd.conf /etc/sysconfig/afsd
```

4. Edit the machine's AFS initialization script or **afsd** options file to set appropriate values for **afsd** command parameters. The appropriate file for each system type is as follows:

- On AIX systems, **/etc/rc.afs**
- On Digital UNIX systems, **/sbin/init.d/afsd**
- On HP-UX systems, **/sbin/init.d/afsd**
- On IRIX systems, **/etc/init.d/afsd**
- On Linux systems, **/etc/sysconfig/afsd** (the **afsd** options file)
- On Solaris systems, **/etc/init.d/afsd**

Use one of the methods described in the introduction to this section to add the following flags to the **afsd** command line. Also set any performance-related arguments you wish.

- Add the **-memcache** flag if the machine is to use a memory cache.
- Add the **-verbose** flag to display a trace of the Cache Manager's initialization on the standard output stream.

Starting the Cache Manager and Installing the AFS Initialization Script

In this section you run the AFS initialization script to start the Cache Manager. If the script works correctly, perform the steps that incorporate it into the machine's startup and shutdown sequence. If there are problems during the initialization, attempt to resolve them. The AFS Product Support group can provide assistance if necessary.

On machines that use a disk cache, it can take a while for the **afsd** program to run the first time on a machine, because it must create all of the **Vn** files in the cache directory. Subsequent Cache Manager initializations do not take nearly as long, because the **Vn** files already exist.

On system types that use a dynamic loader program, you must reboot the machine before running the initialization script, so that it can freshly load AFS modifications into the kernel.

Proceed to the instructions for your system type:

- "Running the Script on AIX Systems" on page 134
- "Running the Script on Digital UNIX Systems" on page 134
- "Running the Script on HP-UX Systems" on page 135
- "Running the Script on IRIX Systems" on page 135
- "Running the Script on Linux Systems" on page 136
- "Running the Script on Solaris Systems" on page 137

Running the Script on AIX Systems

1. Reboot the machine and log in again as the local superuser **root**.

```
# cd /
# shutdown -r now
login: root
Password: root_password
```

2. Run the AFS initialization script.

```
# /etc/rc.afs
```

3. Edit the AIX initialization file, **/etc/inittab**, adding the following line to invoke the AFS initialization script. Place it just after the line that starts NFS daemons.

```
rcafs:2:wait:/etc/rc.afs > /dev/console 2>&1 # Start AFS services
```

4. **(Optional)** There are now copies of the AFS initialization file in both the **/usr/vice/etc** and **/etc** directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm rc.afs
# ln -s /etc/rc.afs
```

5. If a volume for housing AFS binaries for this machine's system type does not already exist, proceed to "Setting Up Volumes and Loading Binaries into AFS" on page 138. Otherwise, the installation is complete.

Running the Script on Digital UNIX Systems

1. Run the AFS initialization script.

```
# /sbin/init.d/afs start
```

2. Change to the `/sbin/init.d` directory and issue the `ln -s` command to create symbolic links that incorporate the AFS initialization script into the Digital UNIX startup and shutdown sequence.

```
# cd /sbin/init.d
# ln -s ../init.d/afs /sbin/rc3.d/S67afs
# ln -s ../init.d/afs /sbin/rc0.d/K66afs
```

3. **(Optional)** There are now copies of the AFS initialization file in both the `/usr/vice/etc` and `/sbin/init.d` directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /sbin/init.d/afs afs.rc
```

4. If a volume for housing AFS binaries for this machine's system type does not already exist, proceed to "Setting Up Volumes and Loading Binaries into AFS" on page 138. Otherwise, the installation is complete.

Running the Script on HP-UX Systems

1. Run the AFS initialization script.

```
# /sbin/init.d/afs start
```

2. Change to the `/sbin/init.d` directory and issue the `ln -s` command to create symbolic links that incorporate the AFS initialization script into the HP-UX startup and shutdown sequence.

```
# cd /sbin/init.d
# ln -s ../init.d/afs /sbin/rc2.d/S460afs
# ln -s ../init.d/afs /sbin/rc2.d/K800afs
```

3. **(Optional)** There are now copies of the AFS initialization file in both the `/usr/vice/etc` and `/sbin/init.d` directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /sbin/init.d/afs afs.rc
```

4. If a volume for housing AFS binaries for this machine's system type does not already exist, proceed to "Setting Up Volumes and Loading Binaries into AFS" on page 138. Otherwise, the installation is complete.

Running the Script on IRIX Systems

1. If you have configured the machine to use the **ml** dynamic loader program, reboot the machine and log in again as the local superuser **root**.

```
# cd /
# shutdown -i6 -g0 -y
login: root
Password: root_password
```

2. Issue the **chkconfig** command to activate the **afsclient** configuration variable.

```
# /etc/chkconfig -f afsclient on
```

3. Run the AFS initialization script.

```
# /etc/init.d/afs start
```

4. Change to the **/etc/init.d** directory and issue the **ln -s** command to create symbolic links that incorporate the AFS initialization script into the IRIX startup and shutdown sequence.

```
# cd /etc/init.d
# ln -s ../init.d/afs /etc/rc2.d/S35afs
# ln -s ../init.d/afs /etc/rc0.d/K35afs
```

5. **(Optional)** There are now copies of the AFS initialization file in both the **/usr/vice/etc** and **/etc/init.d** directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /etc/init.d/afs afs.rc
```

6. If a volume for housing AFS binaries for this machine's system type does not already exist, proceed to "Setting Up Volumes and Loading Binaries into AFS" on page 138. Otherwise, the installation is complete.

Running the Script on Linux Systems

1. Reboot the machine and log in again as the local superuser **root**.

```
# cd /
# shutdown -r now
login: root
Password: root_password
```

2. Run the AFS initialization script.

```
# /etc/rc.d/init.d/afs start
```


- Issue the **chkconfig** command to activate the **afs** configuration variable. Based on the instruction in the AFS initialization file that begins with the string `#chkconfig`, the command automatically creates the symbolic links that incorporate the script into the Linux startup and shutdown sequence.

```
# /sbin/chkconfig --add afs
```

- (Optional)** There are now copies of the AFS initialization file in both the **/usr/vice/etc** and **/etc/rc.d/init.d** directories, and copies of the **afsd** options file in both the **/usr/vice/etc** and **/etc/sysconfig** directories. If you want to avoid potential confusion by guaranteeing that the two copies of each file are always the same, create a link between them. You can always retrieve the original script or options file from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc afs.conf
# ln -s /etc/rc.d/init.d/afs afs.rc
# ln -s /etc/sysconfig/afs afs.conf
```

- If a volume for housing AFS binaries for this machine's system type does not already exist, proceed to "Setting Up Volumes and Loading Binaries into AFS" on page 138. Otherwise, the installation is complete.

Running the Script on Solaris Systems

- Reboot the machine and log in again as the local superuser **root**.

```
# cd /
# shutdown -i6 -g0 -y
login: root
Password: root_password
```

- Run the AFS initialization script.

```
# /etc/init.d/afs start
```

- Change to the **/etc/init.d** directory and issue the **ln -s** command to create symbolic links that incorporate the AFS initialization script into the Solaris startup and shutdown sequence.

```
# cd /etc/init.d
# ln -s ../init.d/afs /etc/rc3.d/S99afs
# ln -s ../init.d/afs /etc/rc0.d/K66afs
```

- (Optional)** There are now copies of the AFS initialization file in both the **/usr/vice/etc** and **/etc/init.d** directories. If you want to avoid potential confusion by guaranteeing that they are always the same, create a link between them. You can always retrieve the original script from the AFS CD-ROM if necessary.

```
# cd /usr/vice/etc
# rm afs.rc
# ln -s /etc/init.d/afs afs.rc
```

- If a volume for housing AFS binaries for this machine's system type does not already exist, proceed to "Setting Up Volumes and Loading Binaries into AFS" on page 138. Otherwise, the installation is complete.

Setting Up Volumes and Loading Binaries into AFS

In this section, you link `/usr/afsws` on the local disk to the directory in AFS that houses AFS binaries for this system type. The conventional name for the AFS directory is `/afs/cellname/sysname/usr/afsws`.

If this machine is an existing system type, the AFS directory presumably already exists. You can simply create a link from the local `/usr/afsws` directory to it. Follow the instructions in "Linking `/usr/afsws` on an Existing System Type" on page 138.

If this machine is a new system type (there are no AFS machines of this type in your cell), you must first create and mount volumes to store its AFS binaries, and then create the link from `/usr/afsws` to the new directory. See "Creating Binary Volumes for a New System Type" on page 138.

You can also store UNIX system binaries (the files normally stored in local disk directories such as `/bin`, `/etc`, and `/lib`) in volumes mounted under `/afs/cellname/sysname`. See "Storing System Binaries in AFS" on page 66 .

Linking `/usr/afsws` on an Existing System Type

If this client machine is an existing system type, there is already a volume mounted in the AFS filespace that houses AFS client binaries for it.

1. Create `/usr/afsws` on the local disk as a symbolic link to the directory `/afs/cellname/@sys/usr/afsws`. You can specify the actual system name instead of `@sys` if you wish, but the advantage of using `@sys` is that it remains valid if you upgrade this machine to a different system type.

```
# ln -s /afs/cellname/@sys/usr/afsws /usr/afsws
```

2. **(Optional)** If you believe it is helpful to your users to access the AFS documents in a certain format via a local disk directory, create `/usr/afsdoc` on the local disk as a symbolic link to the documentation directory in AFS (`/afs/cellname/afsdoc/format_name`).

```
# ln -s /afs/cellname/afsdoc/format_name /usr/afsdoc
```

An alternative is to create a link in each user's home directory to the `/afs/cellname/afsdoc/format_name` directory.

Creating Binary Volumes for a New System Type

If this client machine is a new system type, you must create and mount volumes for its binaries before you can link the local `/usr/afsws` directory to an AFS directory.

To create and mount the volumes, you use the `klog` command to authenticate as an administrator and then issue commands from the `vos` and `fs` command suites. However, the command binaries are not yet available on this machine (by convention, they are accessible via the `/usr/afsws` link that you are about to create). You have two choices:

- Perform all steps except the last one (Step "10" on page 140) on an existing AFS machine. On a file server machine, the **klog**, **fs** and **vos** binaries reside in the **/usr/afs/bin** directory. On client machines, the **klog** and **fs** binaries reside in the **/usr/afsws/bin** directory and the **vos** binary in the **/usr/afsws/etc** directory. Depending on how your **PATH** environment variable is set, you possibly need to precede the command names with a pathname.

If you work on another AFS machine, be sure to substitute the new system type name for the *sysname* argument in the following commands, not the system type of the machine on which you are issuing the commands.

- Copy the necessary command binaries to a temporary location on the local disk, which enables you to perform the steps on the local machine. The following procedure installs them in the **/tmp** directory and removes them at the end. Depending on how your **PATH** environment variable is set, you possibly need to precede the command names with a pathname.

Perform the following steps to create a volume for housing AFS binaries.

1. Working either on the local machine or another AFS machine, mount the AFS CD-ROM for the new system type on the **/cdrom** directory, if it is not already. For instructions on mounting CD-ROMs (either locally or remotely via NFS), consult the operating system documentation.
2. If working on the local machine, copy the necessary binaries to a temporary location on the local disk. Substitute a different directory name for **/tmp** if you wish.

```
# cd /cdrom/new_sysname/root.server/usr/afs/bin
# cp -p klog /tmp
# cp -p fs /tmp
# cp -p vos /tmp
```

3. Authenticate as the user **admin**.

```
# klog admin
Password: admin_password
```

4. Issue the **vos create** command to create volumes for storing the AFS client binaries for this system type. The following example instruction creates volumes called *sysname*, *sysname.usr*, and *sysname.usr.afsws*. Refer to the *IBM AFS Release Notes* to learn the proper value of *sysname* for this system type.

```
# vos create <machine name> <partition name> sysname
# vos create <machine name> <partition name> sysname.usr
# vos create <machine name> <partition name> sysname.usr.afsws
```

5. Issue the **fs mkmount** command to mount the newly created volumes. Because the **root.cell** volume is replicated, you must precede the *cellname* part of the pathname with a period to specify the read/write mount point, as shown. Then issue the **vos release** command to release a new replica of the **root.cell** volume, and the **fs checkvolumes** command to force the local Cache Manager to access them.

```
# fs mkmount -dir /afs/.cellname/sysname -vol sysname
# fs mkmount -dir /afs/.cellname/sysname/usr -vol sysname.usr
# fs mkmount -dir /afs/.cellname/sysname/usr/afsws -vol sysname.usr.afsws
```

```
# vos release root.cell
# fs checkvolumes
```

6. Issue the **fs setacl** command to grant the **l** (**lookup**) and **r** (**read**) permissions to the **system:anyuser** group on each new directory's ACL.

```
# cd /afs/.cellname/sysname
# fs setacl -dir . usr usr/afsws -acl system:anyuser rl
```

7. Issue the **fs setquota** command to set an unlimited quota on the volume mounted at the **/afs/cellname/sysname/usr/afsws** directory. This enables you to copy all of the appropriate files from the CD-ROM into the volume without exceeding the volume's quota.

If you wish, you can set the volume's quota to a finite value after you complete the copying operation. At that point, use the **vos examine** command to determine how much space the volume is occupying. Then issue the **fs setquota** command to set a quota that is slightly larger.

```
# fs setquota /afs/.cellname/sysname/usr/afsws 0
```

8. Copy the contents of the indicated directories from the CD-ROM into the **/afs/cellname/sysname/usr/afsws** directory.

```
# cd /afs/.cellname/sysname/usr/afsws
# cp -rp /cdrom/sysname/bin .
# cp -rp /cdrom/sysname/etc .
# cp -rp /cdrom/sysname/include .
# cp -rp /cdrom/sysname/lib .
```

9. Issue the **fs setacl** command to set the ACL on each directory appropriately. To comply with the terms of your AFS License agreement, you must prevent unauthorized users from accessing AFS software. To enable access for locally authenticated users only, set the ACL on the **etc**, **include**, and **lib** subdirectories to grant the **l** and **r** permissions to the **system:authuser** group rather than the **system:anyuser** group. The **system:anyuser** group must retain the **l** and **r** permissions on the **bin** subdirectory to enable unauthenticated users to access the **klog** binary. To ensure that unauthorized users are not accessing AFS software, check periodically that the ACLs on these directories are set properly.

```
# cd /afs/.cellname/sysname/usr/afsws
# fs setacl -dir etc include lib -acl system:authuser rl \
    system:anyuser none
```

10. Perform this step on the new client machine even if you have performed the previous steps on another machine. Create **/usr/afsws** on the local disk as a symbolic link to the directory **/afs/cellname/@sys/usr/afsws**. You can specify the actual system name instead of **@sys** if you wish, but the advantage of using **@sys** is that it remains valid if you upgrade this machine to a different system type.

```
# ln -s /afs/cellname/@sys/usr/afsws /usr/afsws
```

11. (Optional) To enable users to issue commands from the AFS suites (such as **fs**) without having to specify a pathname to their binaries, include the **/usr/afsws/bin** and **/usr/afsws/etc** directories in the **PATH** environment variable you define in each user's shell initialization file (such as **.cshrc**).

12. **(Optional)** If you believe it is helpful to your users to access the AFS documents in a certain format via a local disk directory, create **/usr/afsdoc** on the local disk as a symbolic link to the documentation directory in AFS (**/afs/cellname/afsdoc/format_name**).

```
# ln -s /afs/cellname/afsdoc/format_name /usr/afsdoc
```

An alternative is to create a link in each user's home directory to the **/afs/cellname/afsdoc/format_name** directory.

13. **(Optional)** If working on the local machine, remove the AFS binaries from the temporary location. They are now accessible in the **/usr/afsws** directory.

```
# cd /tmp
# rm klog fs vos
```


Appendix A. Appendix A. Building AFS from Source Code

This chapter describes how to build AFS from source code.

Loading the Source Files

Working on an AFS client machine, perform these steps to load the AFS source tree from the AFS Source Distribution.

1. Create and mount a volume for housing the AFS source tree. These instructions name the volume **src.afs** and mount it at the **/afs/cellname/afs/src** directory.

Setting the **-maxquota** argument to **0** (zero) sets an unlimited quota on the volume, which enables you to copy all of the files into the volume without exceeding its quota. If you wish, you can set the volume's quota to a finite value after you complete the copying operation. At that point, use the **vos examine** command to determine how much space the volume is occupying. Then issue the **fs setquota** command to set a quota that is slightly larger.

```
# vos create <machine name> <partition name> src.afs -maxquota 0
# cd /afs/.cellname
# mkdir afs
# fs mkmount afs/src src.afs
# vos release root.cell
# fs checkvolumes
```

2. On the local **/cdrom** directory, mount the CD-ROM that contains the AFS source files. For instructions on mounting CD-ROMs (either locally or remotely via NFS), consult the operating system documentation.
3. Copy the source files from the CD-ROM into the newly created volume.

```
# cd /cdrom/src
# cp -rp * /afs/.cellname/afs/src
```

Compiling AFS Binaries Using the washtool Program

The AFS distribution includes the **washtool** program for managing a hierarchy of software development projects. The program builds project trees for program editing, compilation, and installation.

1. Create a subdirectory under the **/afs/.cellname/afs** directory for each system type for which you will build AFS binaries. Creating and mounting a volume for each system type is recommended, but you can also simply use the **mkdir** command. If you create a new volume, grant it an unlimited quota to avoid running out of space during the build process.

```
# cd /afs/.cellname/afs
```

If creating a new volume:

```
# vos create <machine name> <partition name> sysname -maxquota 0
# fs mkmount sysname sysname
```

If not creating a new volume:

```
# mkdir sysname
```

2. In the directory for each system type, create subdirectories called **dest**, **dest/bin**, and **obj**. If you plan to use the `@sys` variable in pathnames that refer to these directories, then you must use the conventional system names listed in the *IBM AFS Release Notes*.

```
# cd sysname
# mkdir dest
# mkdir dest/bin
# mkdir obj
```

3. Create the indicated directories and symbolic links in the `/afs/.cellname/afs` directory.

```
# cd /afs/.cellname/afs
# ln -s @sys/dest dest
# ln -s @sys/obj obj
# ln -s . PARENT
# ln -s src/Makefile Makefile
```

The following is an example directory listing for the `/afs/.cellname/afs` directory after completing the preceding steps. It includes two example system types.

```
lrwxr-xr-x admin 12 Jun 18 11:26 Makefile->src/Makefile
lrwxr-xr-x admin 1 Jun 18 11:26 PARENT -> .
lrwxr-xr-x admin 9 Jun 18 11:25 dest -> @sys/dest
lrwxr-xr-x admin 8 Jun 18 11:25 obj -> @sys/obj
drwxrwxrwx admin 4096 Jun 18 11:24 rcs
drwxrwxrwx admin 2048 Jun 18 11:27 rs_aix42
drwxrwxrwx admin 2048 Jun 18 11:10 src
drwxrwxrwx admin 2048 Jun 18 11:27 sun4x_56
```

4. (Optional) By default, the build procedure writes its results into a destination directory for each system type called `/afs/.cellname/afs/sysname/dest`. To write the results to a different destination directory, create a link from the **dest** directory to it.

```
# cd /afs/.cellname/afs/sysname
# ln -s full_path_of_alternate_directory dest
```

5. For each system type you plan to build, copy the binary for the **washtool** program to the directory specified in the AFS **Makefile**, which is `/afs/cellname/afs/sysname/dest/bin`. If you prefer to store the program in a different directory, you can use the `WASHTOOL` variable on the **make** command line as described in Step "6" on page 145.

If there is a volume that houses the AFS binaries for each system type (as recommended), the conventional location for the **washtool** binary is the `/afs/cellname/sysname/usr/afsws/bin` directory. Use the following instruction to copy it.

```
# cd /afs/cellname/sysname/usr/afsws/bin
# cp washtool /afs/.cellname/afs/sysname/dest/bin
```


Otherwise, mount the (binary) AFS CD-ROM for this system type on the local **/cdrom** directory, and copy the **washtool** binary directly from it.

```
# cd /cdrom/sysname/bin
# cp washtool /afs/.cellname/afs/sysname/dest/bin
```

6. Working in the **/afs/.cellname/afs** directory on a machine of the system type for which you are building AFS, issue the **make install** command. Set the **SYS_NAME** variable to the appropriate system type name.

If the **washtool** binary is not in the conventional directory (**/afs/cellname/afs/sysname/dest/bin**), set the **WASHTOOL** variable to the alternate full pathname of the binary.

```
# cd /afs/.cellname/afs
# make SYS_NAME=sysname [WASHTOOL=alternate_washtool_directory] install
```


Index

Symbols

- / as start to file and directory names
 - see alphabetized entries without initial slash, 9

A

- access
 - to local and foreign cells, 67
 - to root and admin accounts, 69
- access control list (ACL), setting, 61
- activating AFS init. script
 - (see installing)
- adding
 - entries to BosConfig file
 - database server machine, 100
 - first AFS machine, 40
 - server machine after first, 89
 - new db-server machine to CellServDB files, 99
- admin account
 - adding
 - to system:administrators group, 44
 - to UserList file, 43
 - controlling access to, 69
 - creating, 41
 - setting ADMIN flag on Auth. DB entry, 43
- afs (/afs) directory
 - as root of AFS filespace, 131
 - creating
 - client machine, 131
 - first AFS machine, 53
 - server machine after first, 92
- AFS Binary Distribution, 5
- AFS cache
 - (see cache)
- afs entry in Authentication Database, 41
- afs file
 - AFS initialization file, 134, 135, 135, 136, 137
 - afsd options file (Linux), 131
- AFS filespace
 - configuring top levels, 60
 - controlling access by root superuser, 69

- deciding how to configure, 7
- enabling access to foreign cells, 67
- root at /afs directory, 131
- AFS initialization script
 - adding to machine startup sequence
 - client machine, 133
 - first AFS machine, 58
 - server machine after first, 94
 - running
 - client machine, 133
 - first AFS machine, 55
 - server machine after first, 94
 - setting afsd parameters
 - client machine, 131
 - first AFS machine, 53
 - server machine after first, 92
 - verifying on first AFS machine, 55
- AFS kernel extensions
 - on client machine
 - AIX, 106
 - Digital UNIX, 108
 - HP-UX, 111
 - IRIX, 115
 - Linux, 118
 - Solaris, 124
 - on first AFS machine
 - AIX, 10
 - Digital UNIX, 14
 - HP-UX, 19
 - IRIX, 24
 - Linux, 28
 - Solaris, 32
 - on server machine after first
 - AIX, 75
 - Digital UNIX, 76
 - HP-UX, 79
 - IRIX, 82
 - Linux, 85
 - Solaris, 86
- AFS login
 - on client machine
 - AIX, 106
 - Digital UNIX, 108
 - HP-UX, 111
 - IRIX, 118
 - Linux, 118
 - Solaris, 124
 - on file server machine

- AIX, 12
- Digital UNIX, 18
- HP-UX, 22
- IRIX, 28
- Linux, 30
- Solaris, 35
- AFS server partition
 - configuring on first AFS machine
 - AIX, 11
 - Digital UNIX, 16
 - HP-UX, 20
 - IRIX, 27
 - Linux, 29
 - Solaris, 35
 - configuring on server machine after first
 - AIX, 75
 - Digital UNIX, 78
 - HP-UX, 80
 - IRIX, 84
 - Linux, 85
 - Solaris, 88
 - mounted on /vicep directory, 9
 - protecting during operating system upgrade, 5
- afscient variable (IRIX)
 - client machine, 135
 - first AFS machine, 56
 - server machine after first, 95
- afsd
 - command in AFS init. script, 131
 - options file (Linux), 131
- afsm1 variable (IRIX)
 - client machine, 116
 - first AFS machine, 25
 - server machine after first, 82
- afsserver variable (IRIX)
 - first AFS machine, 56
 - server machine after first, 95
- afsxnfs variable (IRIX)
 - client machine, 116
 - first AFS machine, 25
 - server machine after first, 82
- AIX
 - AFS initialization script
 - on add'l server machine, 94
 - on client machine, 134
 - on first AFS machine, 56, 58
 - AFS kernel extensions

B

- on add'l server machine, 75
- on client machine, 106
- on first AFS machine, 10
- AFS login
 - on client machine, 106
 - on file server machine, 12
- AFS server partition
 - on add'l server machine, 75
 - on first AFS machine, 11
- editing /etc/vfs file, 131
- fsck program
 - on add'l server machine, 76
 - on first AFS machine, 12
- Authentication Database, 41
- Authentication Server
 - starting
 - first AFS machine, 40
 - new db-server machine, 100
- authorization checking (disabling)
 - first AFS machine, 38
 - server machine after first, 89
- background reading list, 2
- Backup Server
 - starting
 - first AFS machine, 40
 - new db-server machine, 100
 - stopping, 103
- Basic OverSeer Server
 - (see BOS Server)
- binaries
 - storing AFS in volume, 62, 137
 - storing system in volumes, 66
- Binary Distribution (AFS), 5
- binary distribution machine, 45
- bos commands
 - addhost, 99
 - addkey, 43
 - adduser, 43
 - create, 41
 - delete, 103
 - listhosts, 40
 - listkeys, 43
 - removehost, 103
 - restart

- on first AFS machine, 44
 - on new db-server machine, 101
 - on removed db-server machine, 103
- setcellname, 40
- shutdown, 55
- status, 45
- stop, 103
- BOS Server
 - checking mode bits on AFS directories, 70
 - starting
 - first AFS machine, 38
 - server machine after first, 89
- BosConfig file
 - adding entries
 - database server machine, 100
 - first AFS machine, 40
 - server machine after first, 89
 - removing entries, 103
- bosserv command, 39
- building
 - AFS extensions into kernel
 - (see incorporating AFS kernel extensions)
 - AFS from source, 143
- buserv process
 - (see Backup Server)

C

- cache
 - choosing size, 129
 - configuring
 - client machine, 129
 - first AFS machine, 51
 - server machine after first, 92
 - requirements, 129
- Cache Manager
 - client machine, 131
 - first AFS machine, 53
 - server machine after first, 92
- cacheinfo file, 129
- CD-ROM
 - copying AFS binaries into volume, 64
 - copying AFS documentation from, 65
 - copying client files from
 - client machine, 128
 - first AFS machine, 48

- server machine after first, 91
- copying server files from
 - first AFS machine, 39
 - server machine after first, 89
- copying source files from, 143
- creating /cdrom directory
 - client machine, 105
 - first AFS machine, 9
 - server machine after first, 74
- packaging of AFS Binary Distribution, 5
- cdrom directory
 - client machine, 105
 - first AFS machine, 9
 - server machine after first, 74
- cell
 - enabling access to foreign, 67
 - improving security, 69
 - initializing security mechanisms, 41
- cell name
 - choosing, 7
 - defining during installation of first machine, 40
 - setting in client ThisCell file
 - client machine, 128
 - first AFS machine, 48
 - server machine after first, 92
 - setting in server ThisCell file
 - first AFS machine, 40
 - server machine after first, 92
 - symbolic link for abbreviated, 61
- CellServDB file (client)
 - adding entry
 - for foreign cell, 67
 - for new db-server machine, 100
 - creating
 - on client machine, 128
 - on first AFS machine, 49
 - on server machine after first, 92
 - removing entry, 102
 - required format, 49
- CellServDB file (server)
 - adding entry for new db-server machine, 99
 - creating
 - on first AFS machine, 40
 - on server machine after first, 89
 - displaying entries, 40
 - removing entry, 103
- client cache

- (see cache)
- client machine
 - /cdrom directory, 105
 - /usr/vice/etc directory, 105
 - AFS initialization script, 133
 - AFS kernel extensions
 - on AIX, 106
 - on Digital UNIX, 108
 - on HP-UX, 111
 - on IRIX, 115
 - on Linux, 118
 - on Solaris, 124
 - AFS login
 - on AIX, 106
 - on Digital UNIX, 108
 - on HP-UX, 111
 - on IRIX, 118
 - on Linux, 118
 - on Solaris, 124
 - afsd command parameters, 131
 - afsd options file (Linux), 131
 - Cache Manager, 131
 - cache size and location, 129
 - cell membership, 128
 - CellServDB file
 - adding entry, 100
 - creating during initial installation, 128
 - removing entry, 102
 - copying client files to local disk, 128
 - requirements for installation, 4
 - ThisCell file, 128
 - vfs file (AIX), 131
- clock synchronization, 46
- commands
 - afsd, 131
 - bos addhost, 99
 - bos addkey, 43
 - bos adduser, 43
 - bos create, 41
 - bos delete, 103
 - bos listhosts, 40
 - bos listkeys, 43
 - bos removehost, 103
 - bos restart
 - on first AFS machine, 44
 - on new db-server machine, 101
 - on removed db-server machine, 103
 - bos setcellname, 40
 - bos shutdown, 55
 - bos status, 45
 - bos stop, 103
 - bosserv, 39
 - fs checkvolumes, 58, 62
 - fs examine, 62
 - fs mkmount, 61
 - fs newcell, 68
 - fs setacl, 61
 - fs setquota, 63
 - kas (interactive), 42
 - kas create, 42
 - kas examine, 43
 - kas quit, 43
 - kas setfields, 43
 - klog, 57
 - make, 145
 - pts adduser, 44
 - pts createuser, 44
 - tokens, 57
 - vos addsite, 62
 - vos create
 - root.afs volume, 45
 - root.cell volume, 61
 - src.afs volume, 143
 - volume for AFS binaries, 63
 - volume for AFS documentation, 65
 - vos release, 62
 - vos syncserv, 45
 - vos syncvldb, 45
 - washtool, 144
- compiling AFS from source, 143
- configuring
 - AFS filespace (top levels), 60
 - AFS server partition on first AFS machine
 - AIX, 11
 - Digital UNIX, 16
 - HP-UX, 20
 - IRIX, 27
 - Linux, 29
 - Solaris, 35
 - AFS server partition on server machine
 - after first
 - AIX, 75
 - Digital UNIX, 78
 - HP-UX, 80
 - IRIX, 84
 - Linux, 85

- Solaris, 88
- cache
 - client machine, 129
 - first AFS machine, 51
 - server machine after first, 92
- Cache Manager
 - client machine, 131
 - first AFS machine, 53
 - server machine after first, 92
- copying
 - AFS binaries into volume, 64
 - AFS documentation from CD-ROM, 65
 - client files to local disk
 - client machine, 128
 - first AFS machine, 48
 - server machine after first, 91
 - server files to local disk
 - first AFS machine, 39
 - server machine after first, 89
 - source files from CD-ROM, 143
- creating
 - /cdrom directory
 - client machine, 105
 - first AFS machine, 9
 - server machine after first, 74
 - /usr/afs directory
 - first AFS machine, 9
 - server machine after first, 74
 - /usr/afs/bin directory
 - first AFS machine, 39
 - server machine after first, 74
 - /usr/afs/etc directory
 - first AFS machine, 39
 - server machine after first, 89
 - /usr/vice/etc directory
 - client machine, 105
 - first AFS machine, 9
 - server machine after first, 74
- admin account in Authentication Database, 41
- afs entry in Authentication Database, 41
- CellServDB file (client)
 - client machine, 128
 - first AFS machine, 49
 - server machine after first, 92
- CellServDB file (server)
 - first AFS machine, 40
 - server machine after first, 89

- mount point, 61
- read/write mount point, 62
- root.afs volume, 45
- root.cell volume, 61
- server encryption key
 - Authentication Database, 42
 - KeyFile file, 43
- src.afs volume, 143
- symbolic link
 - for abbreviated cell name, 61
 - to AFS binaries, 64
- UserList file entry, 43
- volume
 - for AFS binaries, 62, 137
 - for AFS documentation, 64
 - for AFS source, 143
 - for system binaries, 66

D

- database server machine
 - entry in client CellServDB file
 - for foreign cell, 67
 - for new db-server machine, 100
 - on client machine, 128
 - on first AFS machine, 49
 - on server machine after first, 92
 - removing, 102
 - entry in server CellServDB file
 - for new db-server machine, 99
 - on first AFS machine, 40
 - on server machine after first, 89
 - removing, 103
- installing
 - additional, 99
 - first, 40
- removing db-server processes from BosConfig file, 103
- removing from service, 101
- requirements for installation, 97
- starting database server processes, 100
- stopping database server processes, 103
- defining
 - cell name during installation of first machine, 40
 - first AFS machine as database server, 40
 - replication site for volume, 62

- Digital UNIX
 - AFS initialization script
 - on add'l server machine, 94
 - on client machine, 134
 - on first AFS machine, 56, 58
 - AFS login
 - on client machine, 108
 - on file server machine, 18
 - AFS server partition
 - on add'l server machine, 78
 - on first AFS machine, 16
 - AFS-modified kernel
 - on add'l server machine, 76
 - on client machine, 108
 - on first AFS machine, 14
 - fsck program
 - on add'l server machine, 78
 - on first AFS machine, 17
- directories
 - /afs, 131
 - /usr/afsdock, 64
 - /usr/afsws, 62, 137
 - /usr/vice/cache, 129
 - /vicepax
 - (see AFS server partition)
- disabling authorization checking
 - first AFS machine, 38
 - server machine after first, 89
- disk cache
 - (see cache)
- displaying
 - CellServDB file (server) entries, 40
 - server encryption key
 - Authentication Database, 43
 - KeyFile file, 43
- documentation, creating volume for AFS, 64

E

- enabling AFS login
 - client machine
 - AIX, 106
 - Digital UNIX, 108
 - HP-UX, 111
 - IRIX, 118
 - Linux, 118
 - Solaris, 124

- file server machine
 - AIX, 12
 - Digital UNIX, 18
 - HP-UX, 22
 - IRIX, 28
 - Linux, 30
 - Solaris, 35
- encryption files
 - in AFS Binary Distribution, 5
- encryption key
 - (see server encryption key)
- environment variables
 - (see variables)
- etc/init.d/afs
 - (see afs file)
- etc/rc.afs
 - (see rc.afs file)
- etc/rc.d/init.d/afs
 - (see afs file)
- etc/sysconfig/afs
 - (see afs file)
- etc/vfs file, 131

F

- File Server
 - first AFS machine, 44
 - server machine after first, 90
- file server machine, 7
 - (see also file server machine, additional)
- requirements for installation, 4
- file server machine, additional
 - /cdrom directory, 74
 - /usr/afs directory, 74
 - /usr/afs/bin directory, 74
 - /usr/afs/etc directory, 89
 - /usr/vice/etc directory, 74
 - AFS initialization script, 94
 - AFS kernel extensions
 - on AIX, 75
 - on Digital UNIX, 76
 - on HP-UX, 79
 - on IRIX, 82
 - on Linux, 85
 - Solaris, 86
 - AFS login
 - (see first AFS machine)

- AFS server partition
 - on AIX, 75
 - on Digital UNIX, 78
 - on HP-UX, 80
 - on IRIX, 84
 - on Linux, 85
 - on Solaris, 88
- afsd command parameters, 92
- authorization checking (disabling), 89
- BOS Server, 89
- Cache Manager, 92
- cache size and location, 92
- cell membership, defining
 - for client processes, 92
 - for server processes, 89
- client functionality, 91
- copying
 - client files to local disk, 91
 - server files to local disk, 89
- File Server, 90
- fs process, 90
- fsck program
 - on AIX, 76
 - on Digital UNIX, 78
 - on HP-UX, 81
 - on IRIX, 81
 - on Linux, 85
 - on Solaris, 87
- runntp process, 90
- server functionality, 88
- ThisCell file (client), 92
- ThisCell file (server), 89
- Update Server client portion, 89
- Update Server server portion, 89
- Volume Server, 90
- file systems clean-up script (Solaris)
 - client machine, 124
 - file server machine, 35
- files
 - afs
 - AFS initialization file, 134, 135, 135, 136, 137
 - afsd options file (Linux), 131
 - AFS initialization
 - (see AFS initialization script)
 - AFS source, 143
 - afsd options file (Linux), 131
 - BosConfig, 40
 - cacheinfo, 129
 - CellServDB (client), 49
 - CellServDB (server), 40
 - index.htm, 65
 - KeyFile, 41
 - protecting during operating system upgrade, 5
 - rc.afs, 134
 - ThisCell (client), 48
 - ThisCell (server), 40
 - UserList, 43
 - vfs (AIX), 131
- fileserver process
 - (see File Server)
- filesystem
 - (see AFS filesystem)
- first AFS machine
 - /cdrom directory, 9
 - /usr/afs directory, 9
 - /usr/vice/etc directory, 9
- AFS initialization script
 - activating, 58
 - running/verifying, 55
- AFS kernel extensions
 - on AIX, 10
 - on Digital UNIX, 14
 - on HP-UX, 19
 - on IRIX, 24
 - on Linux, 28
 - on Solaris, 32
- AFS login
 - on AIX, 12
 - on Digital UNIX, 18
 - on HP-UX, 22
 - on IRIX, 28
 - on Linux, 30
 - on Solaris, 35
- AFS server partition
 - on AIX, 11
 - on Digital UNIX, 16
 - on HP-UX, 20
 - on IRIX, 27
 - on Linux, 29
 - on Solaris, 35
- afsd command parameters, 53
- Authentication Server, 40
- authorization checking (disabling), 38
- Backup Server, 40

- BOS Server, 38
- Cache Manager, 53
- cache size and location, 51
- cell membership, defining
 - for client processes, 48
 - for server processes, 40
- CellServDB file (client), 49
- CellServDB file (server), 40
- client functionality
 - installing, 48
 - removing, 70
- completion of installation, 55
- copying
 - AFS binaries into volume, 64
 - AFS documentation from CD-ROM, 65
 - client files to local disk, 48
 - server files to local disk, 39
- defining
 - as binary distribution machine, 45
 - as database server, 40
 - as system control machine, 45
- File Server, fs process, 44
- fsck program
 - on AIX, 12
 - on Digital UNIX, 17
 - on HP-UX, 21
 - on IRIX, 24
 - on Linux, 28
 - on Solaris, 33
- Protection Server, 40
- roles, 1
- runntp process, 46
- Salvager, 44
- server functionality, 8
- subdirectories of /usr/afs, 39
- ThisCell file (client), 48
- ThisCell file (server), 40
- Update Server server portion, 45
- VL Server, 40
- Volume Server, 44
- foreign cell, enabling access, 67
- fs commands
 - checkvolumes, 58, 62
 - examine, 62
 - mkmount, 61
 - newcell, 68
 - setacl, 61
 - setquota, 63

- fs process
 - first AFS machine, 44
 - server machine after first, 90
- fsck program
 - on first AFS machine
 - AIX, 12
 - Digital UNIX, 17
 - HP-UX, 21
 - IRIX, 24
 - Linux, 28
 - Solaris, 33
 - on server machine after first
 - AIX, 76
 - Digital UNIX, 78
 - HP-UX, 81
 - IRIX, 81
 - Linux, 85
 - Solaris, 87

H

HP-UX

- AFS initialization script
 - on add'l server machine, 95
 - on client machine, 135
 - on first AFS machine, 56, 59
- AFS login
 - on client machine, 111
 - on file server machine, 22
- AFS server partition
 - on add'l server machine, 80
 - on first AFS machine, 20
- AFS-modified kernel
 - on add'l server machine, 79
 - on client machine, 111
 - on first AFS machine, 19
- fsck program
 - on add'l server machine, 81
 - on first AFS machine, 21

I

incorporating AFS kernel extensions

client machine

- AIX, 106
- Digital UNIX, 108
- HP-UX, 111
- IRIX, 115
- Linux, 118
- Solaris, 124

first AFS machine

- AIX, 10
- Digital UNIX, 14
- HP-UX, 19
- IRIX, 24
- Linux, 28
- Solaris, 32

server machine after first

- AIX, 75
- Digital UNIX, 76
- HP-UX, 79
- IRIX, 82
- Linux, 85
- Solaris, 86

index.htm file, 65

initializing

- cell security mechanisms, 41
- server process
(see starting)

installing

AFS initialization script

- client machine, 133
- first AFS machine, 58
- server machine after first, 94

client functionality

- client machine, 105
- first AFS machine, 48
- server machine after first, 91

database server machine

- additional, 99
- first, 40

file server machine after first, 73

first AFS machine, 7

server functionality

- first AFS machine, 8
- server machine after first, 88

instructions

- client machine, 105

database server machine, installing
additional, 99

database server machine, installing first, 40

database server machine, removing, 101

file server machine after first, 73

first AFS machine, 7

interactive mode for kas

entering, 42

quitting, 43

invoking AFS init. script

(see running)

IRIX

AFS initialization script

- on add'l server machine, 95
- on client machine, 135
- on first AFS machine, 56, 59

AFS kernel extensions

- on client machine, 116
- on first AFS machine, 25
- on server machine after first, 82

AFS login, 28

AFS server partition

- on add'l server machine, 84
- on first AFS machine, 27

AFS-modified kernel

- on add'l server machine, 83
- on client machine, 117
- on first AFS machine, 26

afsclient variable

- client machine, 135
- first AFS machine, 56
- server machine after first, 95

afsml variable

- client machine, 116
- first AFS machine, 25
- server machine after first, 82

afsserver variable

- first AFS machine, 56
- server machine after first, 95

afsxnfs variable

- client machine, 116
- first AFS machine, 25
- server machine after first, 82

fsck program replacement not necessary, 24

K

- kas commands
 - create, 42
 - examine, 43
 - interactive mode, entering, 42
 - quit, 43
 - setfields, 43
- kaserver process
 - (see Authentication Server)
- Kerberos, 41
- kernel extensions
 - (see AFS kernel extensions)
- key
 - (see server encryption key)
- KeyFile file
 - first AFS machine, 41
 - server machine after first, 89
- klog command, 57

L

- licensing requirements, 64
- Linux
 - AFS initialization script
 - on add'l server machine, 96
 - on client machine, 136
 - on first AFS machine, 57, 60
 - AFS kernel extensions
 - on add'l server machine, 85
 - on client machine, 118
 - on first AFS machine, 28
 - AFS login
 - on client machine, 118
 - on file server machine, 30
 - AFS server partition
 - on add'l server machine, 85
 - on first AFS machine, 29
 - afsd options file, 131
 - fsck program replacement not necessary, 28
- loading AFS kernel extensions
 - (see incorporating)
- logical volume
 - (see AFS server partition)

M

- make command, 145
- memory cache
 - (see cache)
- mode bits on local AFS directories, 70
- mount point, 61

N

- naming conventions for AFS server partition, 9
- NTPD
 - first AFS machine, 46
 - server machine after first, 90

O

- operating system upgrades, 5
- OPTIONS variable in AFS initialization file, 131
- overview
 - completing installation of first machine, 55
 - general installation requirements, 3
 - installing additional database server machine, 98
 - installing client functionality on first machine, 48
 - installing client machine, 105
 - installing server functionality on first AFS machine, 8
 - installing server machine after first, 73
 - removing database server machine, 101

P

- PAM
 - on HP-UX
 - client machine, 111
 - file server machine, 22
 - on Linux
 - client machine, 118
 - file server machine, 30
 - on Solaris
 - client machine, 124
 - file server machine, 35
- partition

- (see AFS server partition)
- PATH environment variable for users, 64
- Pluggable Authentication Module
 - (see PAM)
- Protection Database, 44
- Protection Server
 - starting
 - first AFS machine, 40
 - new db-server machine, 100
 - stopping, 103
- pts commands
 - adduser, 44
 - createuser, 44
- ptserver process
 - (see Protection Server)

Q

- quota for volume, 63

R

- rc.afs file (AFS init. file for AIX), 134
- read/write mount point for root.afs volume, 62
- reading list for background information, 2
- releasing replicated volume, 62
- removing
 - client functionality from first AFS machine, 70
 - database server machine from service, 101
 - entries from BosConfig File, 103
 - entry from CellServDB file, 102
- replacing fsck program
 - first AFS machine
 - AIX, 12
 - Digital UNIX, 17
 - HP-UX, 21
 - Solaris, 33
 - not necessary on IRIX, 24
 - not necessary on Linux, 28
 - server machine after first
 - AIX, 76
 - Digital UNIX, 78
 - HP-UX, 81
 - Solaris, 87
- replicating volumes, 61
- requirements

- AFS server partition name and location, 9
- cache, 129
- CellServDB file format (client version), 49
- client machine, 4
- database server machine, 97
- file server machine (additional), 73
- file server machine (general), 4
- first AFS machine, 7
- general, 3
- protecting binaries per AFS license, 64
- restarting server process
 - on first AFS machine, 44
 - on new db-server machine, 101
 - on removed db-server machine, 103
- roles for first AFS machine, 1
- root superuser
 - as installer's login identity, 3
 - controlling access, 69
- root.afs volume
 - creating, 45
 - read/write mount point, 62
 - replicating, 61
- root.cell volume
 - creating and replicating, 61
 - mounting for foreign cells in local filespace, 67
- running AFS init. script
 - client machine, 133
 - first AFS machine, 55
 - server machine after first, 94
- runntp process
 - first AFS machine, 46
 - server machine after first, 90

S

- Salvager (salvager process)
 - first AFS machine, 44
 - server machine after first, 90
- sbin/init.d/afs
 - (see afs file)
- scripts
 - AFS initialization
 - (see AFS initialization script)
 - file systems clean-up (Solaris)
 - client machine, 124
 - file server machine, 35

- secondary authentication system (AIX)
 - client machine, 106
 - server machine, 12
- security
 - improving, 69
 - initializing cell-wide, 41
- Security Integration Architecture
 - (see SIA)
- server encryption key
 - in Authentication Database, 42
 - in KeyFile file, 43
- server machine after first
 - (see file server machine, additional)
- server process
 - restarting
 - on first AFS machine, 44
 - on new db-server machine, 101
 - on removed db-server machine, 103
 - see also entry for each server's name, 40
- setting
 - ACL, 61
 - cache size and location
 - client machine, 129
 - first AFS machine, 51
 - server machine after first, 92
 - cell name in client ThisCell file
 - client machine, 128
 - first AFS machine, 48
 - server machine after first, 92
 - cell name in server ThisCell file
 - first AFS machine, 40
 - server machine after first, 89
 - volume quota, 63
- SIA (Digital UNIX)
 - client machine, 108
 - file server machine, 18
- Solaris
 - AFS initialization script
 - on add'l server machine, 96
 - on client machine, 137
 - on first AFS machine, 57, 60
 - AFS kernel extensions
 - on add'l server machine, 86
 - on client machine, 124
 - on first AFS machine, 32
 - AFS login
 - on client machine, 124
 - on file server machine, 35
 - AFS server partition
 - on add'l server machine, 88
 - on first AFS machine, 35
 - file systems clean-up script
 - on client machine, 124
 - on file server machine, 35
 - fsck program
 - on add'l server machine, 87
 - on first AFS machine, 33
 - source (AFS)
 - compiling, 143
 - storing in AFS volume, 143
 - src.afs volume, 143
 - starting
 - Authentication Server
 - first AFS machine, 40
 - new db-server machine, 100
 - Backup Server
 - first AFS machine, 40
 - new db-server machine, 100
 - BOS Server
 - first AFS machine, 38
 - server machine after first, 89
 - File Server
 - first AFS machine, 44
 - server machine after first, 90
 - fs process
 - first AFS machine, 44
 - server machine after first, 90
 - Protection Server
 - first AFS machine, 40
 - new db-server machine, 100
 - runntp process
 - first AFS machine, 46
 - server machine after first, 90
 - Update Server client portion, 89
 - Update Server server portion
 - first AFS machine, 45
 - server machine after first, 89
 - VL Server
 - first AFS machine, 40
 - new db-server machine, 101
 - Volume Server
 - first AFS machine, 44
 - server machine after first, 90
 - stopping
 - database server processes, 103
 - storing

- AFS binaries in volumes, 62, 137
- AFS documentation in volumes, 64
- AFS source in volume, 143
- system binaries in volumes, 66
- supported system types, 5
- symbolic link
 - for abbreviated cell name, 61
 - to AFS binaries from local disk, 64
- system control machine, 45
- system types supported, 5
- system:administrators group, 44
- SYS_NAME variable for washtool command, 145

T

- ThisCell file (client)
 - client machine, 128
 - first AFS machine, 48
 - server machine after first, 92
- ThisCell file (server)
 - first AFS machine, 40
 - server machine after first, 89
- time synchronization, 46
- tokens command, 57

U

- UNIX mode bits on local AFS directories, 70
- upclient process, 89
- Update Server
 - starting client portion, 89
 - starting server portion
 - first AFS machine, 45
 - server machine after first, 89
- upgrading the operating system, 5
- upserver process
 - (see Update Server)
- UserList file
 - first AFS machine, 43
 - server machine after first, 89
- usr/afs directory
 - first AFS machine, 9
 - server machine after first, 74
- usr/afs/bin directory
 - first AFS machine, 39
 - server machine after first, 74

- usr/afs/db directory, 39
- usr/afs/etc directory
 - first AFS machine, 39
 - server machine after first, 89
- usr/afs/etc/CellServDB file
 - (see CellServDB file (server))
- usr/afs/etc/KeyFile
 - (see KeyFile file)
- usr/afs/etc/ThisCell
 - (see ThisCell file (server))
- usr/afs/etc/UserList
 - (see UserList file)
- usr/afs/local directory, 39
- usr/afs/local/BosConfig
 - (see BosConfig file)
- usr/afs/logs directory, 39
- usr/afsdoc directory, 64
- usr/afsws directory, 62, 137
- usr/vice/cache directory, 129
- usr/vice/etc directory
 - client machine, 105
 - first AFS machine, 9
 - server machine after first, 74
- usr/vice/etc/cacheinfo
 - (see cacheinfo file)
- usr/vice/etc/CellServDB
 - (see CellServDB file (client))
- usr/vice/etc/ThisCell
 - (see ThisCell file (client))

V

- variables
 - afsclient (IRIX)
 - client machine, 135
 - first AFS machine, 56
 - server machine after first, 95
 - afsm1 (IRIX)
 - client machine, 116
 - first AFS machine, 25
 - server machine after first, 82
 - afsserver (IRIX)
 - first AFS machine, 56
 - server machine after first, 95
 - afsxnfs (IRIX)
 - client machine, 116
 - first AFS machine, 25

- server machine after first, 82
- OPTIONS (in AFS initialization file), 131
- PATH, setting for users, 64
- SYS_NAME for washtool command, 145
- WASHTOOL, 145
- vfs file, 131
- vicepxx directory
 - (see AFS server partition)
- VL Server (vlserver process)
 - starting
 - first AFS machine, 40
 - new db-server machine, 101
 - stopping, 103
- volserver process
 - (see Volume Server)
- volume
 - creating
 - root.afs, 45
 - root.cell, 61
 - src.afs, 143
 - defining replication site, 62
 - for AFS binaries, 62, 137
 - for AFS documentation, 64
 - for AFS source, 143
 - for system binaries, 66
 - mounting, 61
 - releasing replicated, 62
 - replicating root.afs and root.cell, 61
 - setting quota, 63
- Volume Location Server
 - (see VL Server)
- Volume Server
 - first AFS machine, 44
 - server machine after first, 90
- vos commands
 - addsite, 62
 - create
 - root.afs volume, 45
 - root.cell volume, 61
 - src.afs volume, 143
 - volume for AFS binaries, 63
 - volume for AFS documentation, 65
 - release, 62
 - syncserv, 45
 - syncvldb, 45

W

- washtool command, 144
- WASHTOOL variable, 145