IBM Distributed Computing Environment Version 3.1
for AIX:

# High Availability Cluster Multi-Processing Guide for DCE and DFS

IBM Distributed Computing Environment Version 3.1
for AIX:

# High Availability Cluster Multi-Processing Guide for DCE and DFS

**Note**

Before using this document, be sure to read the general information under "Appendix. Notices" on page 19.

**First Edition (August, 1999)**

This edition applies to Version 3.1 of *IBM Distributed Computing Environment for AIX* and to all subsequent releases and modifications until other wise indicated in new editions or technical newsletters.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

IBM welcomes your comments. Send your comments to the following address:

International Business Machines Corporation
Department VLXA
11400 Burnet Road
Austin, Texas
78758

# Contents

# Figures

# About This Book

This book describes how to integrate the Distributed Computing Environment (DCE) for AIX® and Distributed File System (DFS™) for AIX with an IBM® High Availability Cluster Multi-Processing for AIX* (HACMP) cluster environment.

HACMP provides a computing configuration that survives multiple points of failure and uses reliable, recoverable shared disk resources. HACMP also provides higher availability of application services for applications such as DCE and DFS.

## Who Should Use This Book

This book is intended for system administrators who work with DCE or DFS and for others who want to use DCE or DFS with an HACMP cluster environment. The book assumes you have a working knowledge of HACMP and its requirements. You should also be familiar with:

- RISC System/6000® hardware
- AIX 4.2.1 or AIX 4.3
- DCE for AIX requirements
- DFS for AIX requirements

## Purpose of This Book

After reading this manual, you will know which HACMP configurations are supported for DCE and DFS. You will also understand planning considerations and requirements for integrating DCE and DFS with an HACMP cluster environment.

## How This Book is Organized

The book is organized as follows:

- Chapter 1. Introduction to HACMP for DCE and DFS, gives you an overview of HACMP services relating to DCE and DFS, as well as supported HACMP configurations.
- Chapter 2. Integration Planning for DCE and DFS with HACMP, provides information on planning, setup, and requirements for integrating DCE and DFS with HACMP.
- Chapter 3. Integration of DCE and DFS with HACMP, describes how to integrate DCE and DFS with HACMP. Sample application server start and stop scripts are also provided.

An index is also included.

## Related Publications

For information about DCE and DFS, refer to the following documents:

- *IBM DCE Version 3.1 for AIX: Quick Beginnings*
- *IBM DCE Version 3.1 for AIX: Administration Guide—Core Components*
- *IBM DCE Version 3.1 for AIX: Administration Commands Reference*

For information about HACMP for AIX, see the following documents:

- *High Availability Cluster Multi-Processing for AIX Concepts and Facilities Guide*
- *AIX High Availability Cluster Multi-Processing for AIX Installation Guide*
- *AIX High Availability Cluster Multi-Processing for AIX Planning Guide*
- *AIX High Availability Cluster Multi-Processing for AIX Administration Guide*
- *AIX High Availability Cluster Multi-Processing for AIX Troubleshooting Guide*

## Conventions Used in This Book

This book uses the following typographic conventions:

**Bold** **Bold** words or characters represent system elements that you must enter into the system literally, such as commands, filenames, directory names, and path names.

*Italic* *Italicized* words or characters represent variables that you must supply values for, or the first instance of a new term, if the definition is supplied.

**`Example Font`**
Examples and information displayed by the system are printed using an example font that is a `constant width typeface`. The same font is used for an instruction that you are to type.

**[ ]** Brackets enclose optional items found in format and syntax descriptions are enclosed in brackets.

**{ }** Braces enclose a list from which you must choose an item found in format and syntax descriptions are enclosed by braces.

**|** A vertical bar separates items in a list of choices.

**< >** Angle brackets enclose the name of a key on a keyboard.

**...** Horizontal ellipsis points indicate that you can repeat the preceding item one or more times. Vertical ellipsis points indicate that you can repeat the preceding item one or more times.

This document uses the following keying conventions :

**<Ctrl-*x*>** or ˆ*x*
The notation <Ctrl-*x*> or ˆ*x* followed by the name of a key indicates a control character sequence. For example, **<Ctrl-c>** means hold down the control key while pressing the **c** key.

**<Return>**
The notation **<Return>** refers to the key on your terminal or workstation that is labeled with the word “Return”, with “Enter”, or with a left arrow.

### Entering commands

When instructed to *enter* a command, type the command name and then press the **<Return>** key. For example, the instruction to “Enter the **ls** command” means that you type the **ls** command and then press the **<Return> key**.

# Chapter 1. Introduction to HACMP for DCE and DFS

This section provides an overview of High Availability Cluster Multi-Processing (HACMP) for AIX as used by applications such as Distributed Computing Environment (DCE) and Distributed File System (DFS). The section also lists HACMP release 4.2.1 cluster environments that are supported by DCE and DFS.

## HACMP Overview

HACMP is an environment of loosely coupled, clustered RISC System / 6000 machines executing the HACMP software on AIX to support high availability through redundancy and shared resource access. As a software solution, HACMP does not require any special hardware and works with base RISC System / 6000 machines.

HACMP provides a set of services that guarantees quick recovery upon system failure. Although HACMP does not provide complete fault tolerance and continuous operation, it does provide a minimal recovery time after a system failure. HACMP provides high availability for mission-critical database and transaction-processing applications.

Resources such as disk drives, adapters, and network links are monitored by the HACMP software for failures. If failures are detected, the HACMP software can automatically start administrator defined recovery actions.

## DCE in an HACMP Cluster Environment

DCE is a set of services and tools that supports creation, use, and maintenance of distributed applications in a heterogeneous computing environment. DCE provides the following services for distributed applications: Remote Procedure Call (RPC), Security Services, Cell Directory Service (CDS), and Distributed Time Service (DTS).

The DCE cell is a grouping of machines that share the same Security and CDS servers in a network. Users and nodes within the DCE cell rely on the CDS to register services and applications and on the Security Service to authorize and authenticate access to the services. The servers where the Security and CDS services are located must be up and running for the DCE cell to operate and are critical to the DCE cell. HACMP provides DCE with the high-availability solution that DCE cells need.

DCE Security and CDS services provide replication for greater availability. This replication model is based on the idea that there is only one writable server and many read-only servers. An inoperative writable server is a more severe problem than an inoperative read-only server. When a writable server goes down, you must directly intervene to convert a read-only server to a writable server. Configuring a writable server in an HACMP environment provides the highest availability of writable servers for DCE applications running in a continuous 7-day-a-week, 24-hour-a-day operation.

CDS and the Security Service must be highly available since other services depend on them. If the master CDS server is down, DCE application servers cannot register their location in the CDS namespace. If the master security server is down, no changes can be made to the user / group registry, and DCE application servers cannot perform key management operations required by the DCE Security model.

## DFS in an HACMP Cluster Environment

The Distributed File System (DFS) is an application built on DCE services such as DCE security and DCE cell administration functions. DFS provides a rich set of tools and functionality that extends the view of a local file system to a distributed file system. With DFS, local AIX Journaled File Systems or local Enhanced DFS File Systems from individual machines can be exported to an enterprise-wide distributed file system.

DFS consist of and runs on DCE-based servers, where each server performs a specific machine role. When DFS is used in an HACMP environment, the DFS services become highly available in the event of hardware failure in nodes. The DFS machine roles that operate in an HACMP environment are System Control Machine, Fileset Database Machine, File Server Machine, and DFS Client Machine.

## HACMP Supported Configurations

This section list the services that are highly available for the following HACMP cluster configurations. Detailed information about HACMP cluster configurations can be found in the *HACMP for AIX Concepts and Facilities Guide*.

- One-for-One (Hot) Standby Configuration with Cascading Resources
  - DCE services: Cell Directory Service, Security Service, Distributed Time Service, DCE clients
  - DFS services: File Servers, System Control Machine, Fileset Database Machine, DFS clients
  - DCE and DFS applications
- One-Sided Takeover Configuration Using Cascading Resources
  - DCE services: Cell Directory Service, Security Service, Distributed Time Service, DCE clients
  - DFS services: File Servers, System Control Machine, Fileset Database Machine, DFS clients
  - DCE and DFS applications
- One-for-One (Hot) Standby Configuration with Rotating Resources
  - DCE services: Cell Directory Service, Security Service, Distributed Time Service, DCE clients
  - DFS services: File Servers, System Control Machine, Fileset Database Machine, DFS clients
  - DCE and DFS applications
- Mutual Takeover Configuration with Cascading Resources
  - DCE and DFS clients[2]
  - DCE and DFS applications

**Notes:**

1. DCE and DFS servers cannot reside on nodes configured as mutual takeover.
2. DCE and DFS clients cannot be configured as takeover in the HACMP mutual takeover configuration.

# Chapter 2. Integration Planning for DCE and DFS with HACMP

This section provides planning information for integrating DCE and DFS with HACMP. It includes information on setup and integration requirements.

Read and understand the HACMP for AIX Version 4.2.1 publications:

- *HACMP Concepts and Facilities*
- *HACMP Planning Guide*
- *HACMP Installation Guide*
- *HACMP Administration Guide*
- *HACMP Troubleshooting Guide*

Perform planning activities as identified in *HACMP for AIX Version 4.2.1 Planning Guide*, Chapters 1 through 9.

## Software Requirements

- AIX 4.2.1 or AIX 4.3
- HACMP 4.2.1 with PTFs
- DCE 3.1 (or higher) for AIX

## HACMP Cluster Event Processing

The HACMP environment is built on the concept of clustering. In a cluster, multiple server processors cooperate to provide a set of services or resources to other entities. HACMP defines relationships among cooperating processors where peer cluster nodes provide the services offered by a cluster node that becomes disabled.

The HACMP Cluster Manager runs on each cluster node, monitoring local hardware and software subsystems, tracking the state of the cluster peers, and triggering cluster events when the cluster status changes. A cluster event represents a change in a cluster’s operational state that the HACMP Cluster Manager recognizes and can respond to.

Cluster nodes exchange keep-alive messages with peer nodes so that the Cluster Manager can track the availability of the nodes in the cluster. If a node stops sending keep-alives, the peer nodes drive the recovery process. The peer nodes take the necessary actions to get critical applications up and running and to ensure that data has not been corrupted. This relationship between nodes is the basis for a *failover of services*. A failover of services occurs when an HACMP cluster environment experiences a change that requires stopping services on one node and resuming those services on the standby or peer node.

When a node that had previously failed rejoins the cluster, the Cluster Manager running on the existing nodes acknowledges that the new node is up. In HACMP configurations using a standby node, a special sequence of actions called *reintegration* occurs. A reintegration of services occurs when the node originating a failover of services becomes operational. After reintegration completes, the standby node is ready for subsequent service failover. For DFS, you must shutdown or reboot the standby node to complete reintegration before you can restart services on the owner node that originally provided the service.

You can customize the response to service or system failures that result in a failover or reintegration. Each event can have a registered event script that is started when that particular event occurs. You can customize the sample scripts provided by HACMP or can write completely new scripts based on the set of applications on the node.

## Application Server Scripts

An application server is an application that runs on a cluster node and is a cluster resource that can be made highly available by the HACMP software. The application server is configured as part of an HACMP resource group. The application is then started or stopped when the resource group is acquired or released by the joining or leaving cluster node. For more information, see Chapter 10 “Managing Application Servers” in the *HACMP for AIX Administration Guide.*

Applications can alternately be managed by pre and postcluster event processing. You can specify commands that execute before and after a specific event is generated by the cluster manager. See Chapter 9 “Maintaining Cluster Events Processing” in the *HACMP for AIX Administration Guide.*

DCE and DFS have specific requirements to achieve failover and reintegration in HACMP-supported configurations. Sample start and stop scripts for application servers are provided to assist automating the failover and reintegration of DCE and DFS services in HACMP configurations.

## Planning Considerations for DCE

When planning to integrate DCE with HACMP, keep in mind the following:

- DCE must be installed in a supported HACMP configuration, as listed in “HACMP Supported Configurations” on page 2.

For the HACMP mutual takeover configuration, keep in mind the following:

- DCE servers cannot operate on the mutual takeover peer nodes. For applications using DCE services to be highly available, DCE clients must be operating on both of the mutual takeover peer nodes.
- DCE clients and DCE applications can operate on the mutual takeover peer nodes.
- IP addresses for the DCE services are configured as service IP labels in the cascading resource group.
- Initially, the **RPC_UNSUPPORTED_NETADDRS** environment variable must be set to exclude all addresses on the system which are not the service address. The only supported network address should be the service address. Before a mutual takeover node can take over a DCE application, set **RPC_UNSUPPORTED_NETADDRS** to exclude the original service address on the node which will be taking over the DCE application. You must export this environment variable in the session you will use to configure DCE.

  This will allow the DCE clients and DCE applications which were started on the takeover node to remain running on the original service address and will allow the DCE applications which are being taken over to continue running on the service address for the node which failed. This allows applications which are taken over to remain associated with the same IP address regardless of which node is running the applications.

Put this variable into the **/etc/environment** file so that DCE application servers do not use these interfaces. The environment variable is initialized with up to 10 addresses to exclude, each separated by a colon. For example, if the boot adapter has an address of 9.3.240.1, the service adapter has an address of 9.3.240.2, and the standby adapter has an address of 9.3.240.66, the environment variable would be the following:

```
RPC_UNSUPPORTED_NETADDRS=9.3.240.1:9.3.240.66
```

For the HACMP one-for-one (hot) standby configuration with cascading resources, one-sided takeover configuration using cascading resources, and one-for-one (hot) standby configuration with rotating resources, keep in mind the following:

- The filesystems **/krb5**, **/var/dce**, and **/etc/dce** must reside on shared DASD and must be configured as file systems in the cascading or rotating resource groups.
- IP addresses for the DCE services are configured as file systems in the cascading or rotating resource groups.
- DCE services must not be configured on the standby/takeover node.
- Reintegration does not occur in the HACMP one-for-one (hot) standby configuration with rotating resources. When the original failing node is restarted in this configuration, it assumes the standby role and does not try to take over the resources as it would in the one-for-one (hot) standby configuration with cascading resources.
- When a DCE server application initializes the RPC runtime to be able to register its services, the RPC runtime queries the operating system to find out the network adapters that are configured to use TCP/IP. If there are multiple network adapters, the RPC runtime will attempt to use all interfaces available. The **RPC_UNSUPPORTED_NETADDRS** environment variable is used to exclude unwanted addresses from being enabled by the DCE RPC runtime. All addresses that are not defined as service IP labels in the cascading or rotating resource groups should be excluded by means of the environment variable on all nodes of the HACMP cluster.

  You can put this variable into the **/etc/environment** file so that server programs do not use these interfaces. The only interfaces that should be used by DCE are the ones defined to HACMP as service IP labels. The environment variable is initialized with up to 10 addresses to exclude, each separated by a colon. For example, if the boot adapter has an address of 9.3.240.1, the service adapter has an address of 9.3.240.2, and the standby adapter has an address of 9.3.240.66, the environment variable would be the following:

  ```
  RPC_UNSUPPORTED_NETADDRS=9.3.240.1:9.3.240.66
  ```

**Note:** The DCE daemon (**dced**) and any daemon managed by **dced** will recognize the **RPC_UNSUPPORTED_NETIFS** variable and ignore the **RPC_UNSUPPORTED_NETADDRS** variable. If you want any daemon managed by **dced** such as **dtsd** to exclude certain interfaces, use the **RPC_UNSUPPORTED_NETIFS** variable, or start the daemon on the command line.

## DCE Failover Actions

**Note:** DCE failover actions listed below do not apply to the HACMP mutual takeover configuration.

You must perform the following actions sequentially for failover of DCE services. If you perform these actions manually, you must complete each action before you perform the next action. However, direct intervention defeats the purpose of using HACMP to achieve automated failover of services.

- HACMP on the standby/takeover node must take over IP address and shared DASD.
- For the HACMP one-for-one (hot) standby configuration with rotating resources, save the hostname environment value by running the **hostname** command. Then set hostname to the name associated with the IP address of the rotating resource.
- From the standby/takeover node, delete the following endpoint map and security credentials files from shared DASD:

  **/opt/dcelocal/var/dced/Ep.db**

  **/opt/dcelocal/var/security/creds/***

  **Note:** All endpoints for any running DCE application servers are removed when **/opt/dcelocal/var/dced/Ep.db** is deleted. When the **/opt/dcelocal/var/security/creds/*** files are removed, any application client or application server or DCE authenticated user will no longer have valid DCE credentials. The application server or client must reauthenticate to obtain new credentials. Users must log in again to DCE to obtain new credentials.
- Start the DCE daemons on the standby/takeover node by running **/usr/bin/start.dce**.
- For the HACMP one-for-one (hot) standby configuration with rotating resources, reset the hostname environment value by running the **hostname** command.

You can automate these actions through the use of HACMP cluster event scripts or an HACMP application server start script on the standby/takeover node. When the standby/takeover node detects a condition requiring a failover of services, the standby/takeover node takes over the shared resources and then starts services that have registered their application server start scripts.

A sample application server start script for DCE failover actions on the standby/takeover node is provided in "Application Server Start Script Example for DCE" on page 15. You can define this script as an HACMP application server start script, or as a **node_down** postevent on the standby/takeover node.

You can stop the owner node with a shutdown mode of graceful with takeover to force a failover of services to the standby/takeover node. If you have defined application server stop scripts for this event, these stop scripts will be run on the owner node before actions on the standby/takeover node begin. A sample application server stop script for DCE failover actions on the owner node is provided in "Application Server Stop Script Example for DCE" on page 16. You can define this script as an HACMP application server stop script, or as a **node_down** preevent on the owner node. In the case of a nongraceful takeover, this stop script would not be started.

**Notes:**

1. You must tailor the sample application server start and stop scripts for additional user applications.
2. You can use several methods to automate failover. However, these actions must be sequential. Application server scripts can be used alone or in conjunction with event scripts. The actions relating to stopping on one node must be sequential, and the actions relating to starting on the other node must be

sequential. You can do this by putting all actions relating to stopping DCE and DCE-related services in one application server stop script and all actions relating to starting DCE and DCE-related services in one application server start script.

## DCE Reintegration Actions

**Note:** DCE reintegration actions listed below do not apply to the HACMP one-for-one (hot) standby configuration with rotating resources and the HACMP mutual takeover configuration.

You must perform the following actions sequentially to reintegrate DCE services. If you perform these actions manually, you must complete each action before you perform the next action.

- Stop the DCE services on the standby/takeover node by running **/usr/bin/stop.dce**.
- Start HACMP on the owner node such that IP address and shared DASD takeover occurs on owner node.
- From the owner node, delete the following endpoint map and security credentials files from shared DASD:

  **/opt/dcelocal/var/dced/Ep.db**

  **/opt/dcelocal/var/security/creds/***

  **Note:** All endpoints for any running DCE application servers are removed when **/opt/dcelocal/var/dced/Ep.db** is deleted. When the **/opt/dcelocal/var/security/creds/*** files are removed, any application client or application server or DCE authenticated user will no longer have valid DCE credentials. The application server or client must reauthenticate to obtain new credentials. Users must log in again to DCE to obtain new credentials.
- Start the DCE services on the owner node by running **/usr/bin/start.dce**.

You can automate these actions through the use of HACMP cluster event scripts or HACMP application server start and stop scripts. When the owner node is restarted, the standby/takeover node is notified so that the application server scripts can then stop DCE services on the standby/takeover node. Then the owner node can start DCE services on the owner node.

A sample application server stop script for DCE reintegration actions on the standby/takeover node is provided in “Application Server Stop Script Example for DCE” on page 16. You can define this script as an HACMP application server stop script, or as a **node_up** preevent on the standby/takeover node.

A sample application server start script for DCE reintegration actions on the owner node is provided in “Application Server Start Script Example for DCE” on page 15. This is the same application server start script for DCE failover actions. You can customize it to the user’s environment. In this case, you can define the sample script as an HACMP application server start script, or as a **node_up** postevent on the owner node.

**Notes:**

1. You must tailor the sample application server start and stop scripts for additional user applications.

2. You can use several methods to automate reintegration. However, these actions must be sequential. Application server scripts can be used alone or in conjunction with event scripts. The actions relating to stopping on one node must be sequential, and the actions relating to starting on the other node must be sequential. You can do this by putting all actions relating to stopping DCE and DCE-related services in one application server stop script and all actions relating to starting DCE and DCE-related services in one application server start script.

## Planning Considerations for DFS

When planning to integrate DFS with HACMP, keep in mind the following:

- Planning considerations for DCE. Refer to “Planning Considerations for DCE” on page 4 for additional information.
- DCE and DFS must be installed in a supported HACMP configuration, as listed in “HACMP Supported Configurations” on page 2.
- DFS requires DCE to be up and running prior to startup. Startup scripts used to start DFS must previously start DCE.
- Set up the DFS cache in a separate local AIX Journaled File System not on shared DASD. By default, the DFS cache uses the **/var/dce/adm/dfs/cache** directory. If you use a different directory, specify that directory when you configure the DFS client. Refer to *IBM DCE Version 3.1 for AIX: Quick Beginnings* and *IBM DCE Version 3.1 for AIX: Administration Guide—Core Components* for additional information.

For the HACMP mutual takeover configuration, keep in mind the following:

- DFS servers cannot operate on the mutual takeover peer nodes. For applications using DFS services to be highly available, DFS clients must be operating on both of the mutual takeover peer nodes.
- DFS clients and DFS applications can operate on the mutual takeover peer nodes.

For the HACMP one-for-one (hot) standby configuration with cascading resources, one-sided takeover configuration using cascading resources, and one-for-one (hot) standby configuration with rotating resources, keep in mind the following:

- All DFS aggregates, filesets, and other supported filesystems must reside on shared DASD and must be configured as volume groups or file systems in the cascading or rotating resource groups.
- The DFS node that had failed must be shutdown or rebooted to release shared DASD before the other node can be started. The shutdown can occur from the node failure or by your actions.
- DFS services cannot be configured on the standby or takeover node.
- Reintegration does not occur in the HACMP one-for-one (hot) standby configuration with rotating resources. When the original failing node is restarted in this configuration, it assumes the standby role and does not try to takeover the resources as it would in the one-for-one (hot) standby configuration with cascading resources.

### DFS Failover Actions

**Note:** DFS failover actions listed below do not apply to the HACMP mutual takeover configuration.

You must perform the following actions sequentially for failover of DFS services. If you perform these actions manually, you must complete each action before you perform the next action. However, direct intervention defeats the purpose of using HACMP to achieve automated failover of services.

- Ensure that the owner node has been shutdown. The shutdown can occur from the node failure or by your actions.
- HACMP on the standby/takeover node must take over IP address and shared DASD.
- For the HACMP one-for-one (hot) standby configuration with rotating resources, save the hostname environment value by running the **hostname** command. Then set hostname to the name associated with the IP address of the rotating resource.
- From the standby/takeover node, delete the following endpoint map and security credentials files from shared DASD:

  **/opt/dcelocal/var/dced/Ep.db**

  **/opt/dcelocal/var/security/creds/***

  **Note:** All endpoints for any running DCE application servers are removed when **/opt/dcelocal/var/dced/Ep.db** is deleted. When the **/opt/dcelocal/var/security/creds/*** files are removed, any application client or application server or DCE authenticated user will no longer have valid DCE credentials. The application server or client must reauthenticate to obtain new credentials. Users must log in again to DCE to obtain new credentials.
- Start the DCE daemons on the standby/takeover node by running **/usr/bin/start.dce**.
- Start the DFS services on the standby/takeover node by executing **/usr/bin/start.dfs**.
- For the HACMP one-for-one (hot) standby configuration with rotating resources, reset the hostname environment value by running the **hostname** command.

You can automate these actions through the use of HACMP cluster event scripts or an HACMP application server start script on the standby/takeover node. When the standby/takeover node detects a condition requiring a failover of services, the standby/takeover node takes over the shared resources and then starts services that have registered their application server start scripts.

A sample application server start script for DFS failover actions on the standby/takeover node is provided in “Application Server Start Script Example for DFS” on page 16. You can define this script as an HACMP application server start script, or as a **node_down** postevent on the standby/takeover node.

You can stop the owner node with a shutdown mode of graceful with takeover to force a failover of services to the standby/takeover node. If you have defined application server stop scripts for this event, these stop scripts will be run on the owner node before actions on the standby/takeover node begin. A sample application server stop script for DFS failover actions on the owner node is provided in “Application Server Stop Script Example for DFS” on page 17. You can define this script as an HACMP application server stop script, or as a **node_down** preevent on the owner node. In the case of a node failure (nongraceful takeover), this stop script would not be run.

**Notes:**

1. You must tailor the sample application server start and stop scripts for additional user applications.

2. You can use several methods to automate failover. However, these actions must be sequential. Application server scripts can be used alone or in conjunction with event scripts. The actions relating to stopping on one node must be sequential, and the actions relating to starting on the other node must be sequential. You can do this by putting all actions relating to stopping DCE, DFS, DCE-related, and DFS-related services in one application server stop script and all actions relating to starting DCE, DFS, DCE-related, and DFS-related services in one application server start script.

## DFS Reintegration Actions

**Note:** DFS reintegration actions listed below do not apply to the HACMP one-for-one (hot) standby configuration with rotating resources and the HACMP mutual takeover configuration.

You must perform the following actions sequentially to reintegrate DFS services. If you perform these actions manually, you must complete each action before you perform the next action.

- Stop the DFS services on the standby/takeover node by running **/usr/bin/stop.dfs**.
- Stop the DCE services on the standby/takeover node by running **/usr/bin/stop.dce**.
- Shutdown or reboot the standby/takeover node. This is necessary to stop the DFS daemons and release the shared DASD and must be performed to complete reintegration.
- Restart the cluster by starting HACMP on the owner node such that IP address and shared DASD takeover occurs.
- From the owner node, delete the following endpoint map and security credentials files from shared DASD:

  **/opt/dcelocal/var/dced/Ep.db**

  **/opt/dcelocal/var/security/creds/***

  **Note:** All endpoints for any running DCE application servers are removed when **/opt/dcelocal/var/dced/Ep.db** is deleted. When the **/opt/dcelocal/var/security/creds/*** files are removed, any application client or application server or DCE authenticated user will no longer have valid DCE credentials. The application server or client must reauthenticate to obtain new credentials. Users must log in again to DCE to obtain new credentials.
- Start the DCE services on the owner node by running **/usr/bin/start.dce**.
- Start the DFS services on the owner node by running **/usr/bin/start.dfs**.

You can automate most of these actions through the use of HACMP cluster event scripts or HACMP application server start and stop scripts. The standby/takeover node must be shutdown or rebooted manually to release shared DASD before the owner node can be started. After the standby/takeover node is completely shutdown, HACMP can be started on the owner node, and HACMP can start the application server start scripts.

A sample application server stop script for DFS reintegration actions on the standby/takeover node is provided in “Application Server Stop Script Example for DFS” on page 17. You can define this script as an HACMP application server stop script, or as a **node_up** preevent on the standby/takeover node.

A sample application server start script for DFS reintegration actions on the owner node is provided in “Application Server Start Script Example for DFS” on page 16. This script performs the same actions as the application server start script for DFS Failover actions. In this case, you can define the sample script as an HACMP application server start script, or as a **node_up** postevent on the owner node. This script will be run when HACMP is started on the owner node.

**Notes:**

1. You must tailor the sample application server start and stop scripts for additional user applications.
2. You can use several methods to automate most of the DFS reintegration. However, these actions must be sequential. Application server scripts can be used alone or in conjunction with event scripts. The actions relating to stopping on one node must be sequential, and the actions relating to starting on the other node must be sequential. You can do this by putting all actions relating to stopping DCE, DFS, DCE-related, and DFS-related services in one application server stop script and all actions relating to starting DCE, DFS, DCE-related, and DFS-related services in one application server start script. A manual shutdown or reboot of the standby/takeover node must be done before starting HACMP on the owner node.

# Chapter 3. Integration of DCE and DFS with HACMP

This section provides detailed information on setting up DCE and DFS to operate in an HACMP cluster environment.

## Setting Up DCE and DFS with HACMP

Prepare for integrating DCE for supported HACMP environments as listed in “HACMP Supported Configurations” on page 2:

1. Read and understand the HACMP for AIX Version 4.2.1 publications:
   a. *HACMP for Aix Concepts and Facilities Guide*
   b. *HACMP for AIX Planning Guide*
   c. *HACMP for AIX Installation Guide*
   d. *HACMP for AIX Administration Guide*
   e. *HACMP for AIX Troubleshooting Guide*
2. Perform planning activities identified in the *HACMP for AIX Planning Guide*, Chapters 1 through 9.

### DCE and DFS Integration on Highest Priority Node

**Note:** It is assumed that DCE and DFS are not already installed on nodes being installed and configured.

Perform the following actions to integrate DCE and DFS for the HACMP environments on the highest priority node:

1. Install and configure HACMP based on the planning activities.

   **Note:** The following action (Step 2) does not apply to the HACMP mutual takeover configuration.
2. DCE with HACMP requires that logical volumes and associated filesystems **/krb5, /var/dce**, and **/etc/dce** be created and mounted on shared DASD as file systems in the cascading or rotating resource groups. To set up shared DASD, see Chapter 6 of *HACMP for AIX Installation Guide.*
   a. Create logical volumes with non-default names for **/var/dce**, **/etc/dce**, and **/krb5**.
   b. Create filesystems using previously defined logical volumes.
   c. Mount the filesystems.
   d. Rename the log file to a non-default name.
3. Based on the network interface planning, assign a string and export the environment variable **RPC_UNSUPPORTED_NETADDRS**. For further information, see “Planning Considerations for DCE” on page 4.
4. If you are going to install and configure DFS, create an AIX Journaled File System on a nonshared DASD for the DFS client cache and mount it at **/var/dce/adm/dfs/cache**. If you use a different mount point, you must specify that mount point when configuring the DFS client. Refer to *IBM DCE Version 3.1 for AIX: Quick Beginnings* and *IBM DCE Version 3.1 for AIX: Administration Guide—Core Components* for additional information.

5. Start HACMP. Initial service IP address swap must occur before configuration of DCE can begin. For HACMP configurations other than mutual takeover, you must mount the file systems (shared DASD must be accessible before installing DCE).
6. Install DCE (and DFS if desired). Refer to *IBM DCE Version 3.1 for AIX: Quick Beginnings* for additional information.
7. Configure DCE (and DFS if desired).
8. The following action does not apply to the HACMP mutual takeover configuration. If DFS has been installed and configured, create all DFS aggregates and filesets on shared DASD as volume groups or file systems in the cascading or rotating resource groups.
9. If event handling and application server scripts are used, put them in the location as specified in the HACMP event and application server configurations.
10. After successful configuration, stop DFS (if installed and configured) by running **/usr/bin/stop.dfs**, and then stop the DCE daemons by running **/usr/bin/stop.dce**.
11. The following action does not apply to the HACMP mutual takeover configuration. If DFS was installed and configured, shutdown the owner node to stop DFS daemons and release shared DASD. If DFS was not installed and configured, unmount the filesystems **/krb5, /var/dce**, and **/etc/dce**, and vary off the volume groups. This must be done for configuration of the standby/takeover node to have access to the shared DASD.

## DCE and DFS Integration on Standby/Takeover Node

**Note:** It is assumed that DCE and DFS are not already installed on nodes being installed and configured.

Perform the following actions to integrate DCE and DFS for the HACMP environments on the standby/takeover node:

1. Install and configure HACMP based on the planning activities.
2. The following action does not apply to the HACMP mutual takeover configuration. Set up shared DASD created on the highest priority node created in Step 1. To set up shared DASD, see Chapter 6 of *HACMP for AIX Installation Guide*.
3. Based on the network interface planning, assign a string and export the environment variable **RPC_UNSUPPORTED_NETADDRS**. For further information, see "Planning Considerations for DCE" on page 4.
4. If you are going to install and configure DFS, create an AIX Journaled File System on a nonshared DASD for the DFS client cache and mount it at **/var/dce/adm/dfs/cache**. If you use a different mount point, you must specify that mount point when configuring the DFS client. Refer to *IBM DCE Version 3.1 for AIX: Quick Beginnings* and *IBM DCE Version 3.1 for AIX: Administration Guide—Core Components* for additional information.
5. Start HACMP for possible IP address takeover. For HACMP configurations other than mutual takeover, shared DASD must be accessible before installing DCE. For HACMP configurations other than mutual takeover, this will verify that HACMP can vary on volume groups and mount the filesystems **/krb5, /var/dce**, and **/etc/dce** and any filesystems created for DFS.
6. Install DCE (and DFS if desired). Refer to *IBM DCE Version 3.1 for AIX: Quick Beginnings* for additional information.

7. For the HACMP configuration mutual takeover, configure DCE and DFS clients if desired. For other HACMP configurations, do not configure DCE (or DFS).
8. If DFS was installed, you must perform the following actions to create a mount point for DFS in event of failover:
   a. Create the directory ″**/...**″.
   b. Create the soft links as follows:
      - **ln -s /.../***cell_name***/fs /:**
      - **ln -s /.../***cell_name* **/.:**

        where cell_name is the name of the DCE cell configured during owner node integration.
9. If the port for the Kerberos key distribution center for security is not defined in the **/etc/services** file, then the following Bourne Shell command sequence makes the required changes to the **/etc/services** file.

   ```
   grep "kerberos5[b|t.]" /etc/services >/dev/null
   if [ $? –ne 0 ]; then
      echo "kerberos5\t88/udp\t\tkdc" >>/etc/services
   fi
   ```

   **Note:** In the above example, the ″t″ in the bracketed pair represents the ″tab″ character. The ″b″ in the bracketed pair represents a single space.
10. If event handling and application server scripts are used, put them in the location as specified in the HACMP event and application server configurations.
11. The following action does not apply to the HACMP mutual takeover configuration. Stop HACMP to unmount the filesystems **/krb5, /var/dce**, and **/etc/dce**, and vary off the volume groups. You must confirm that the filesystems have been unmounted and that the volume groups have been varied off. This must be done to release shared DASD in preparation for starting DCE on the highest priority node.

## Starting DCE and DFS with HACMP

Start the HACMP cluster manager on the highest priority node. After HACMP has started on the highest priority node, start the HACMP cluster manager on the standby/takeover node. If application server start scripts have been defined, DCE and DFS will automatically be started. Otherwise, start DCE by running **/usr/bin/start.dce** and then start DFS, if installed and configured, by running **/usr/bin/start.dfs**.

# Application Server Script Samples

The following application server script samples can be used to automate required failover and reintegration actions. You may need to customize these scripts to achieve automated failover and reintegration of other user applications.

## Application Server Start Script Example for DCE

Figure 1 on page 16 is an example of an application server start script for DCE. It deletes the endpoint map and security credentials files and then runs **/usr/bin/start.dce** to start the configured DCE processes.

```
#!/bin/ksh
#
#For the HACMP one-for-one (hot) standby configuration with rotating
#resources, save the hostname, associate it with the IP address
#which is the rotating resource and reset it at the end.  Note the
#command substitution string for hostname and the grave accents
#around it.
#
save_hostname = `hostname`
hostname xxx      # replace xxx with your rotating resource hostname
#
#

echo "Removing DCE Files\n"
rm /opt/dcelocal/var/dced/Ep.db
rm /opt/dcelocal/var/security/creds/*

echo "Starting DCE\n"
/usr/bin/start.dce
echo "Finished Starting DCE\n"
#
#For the HACMP one-for-one (hot) standby configuration with rotating
#resources, reset the hostname.
hostname $save_hostname
#
```

*Figure 1. Application Server Start Script Example for DCE*

## Application Server Stop Script Example for DCE

Figure 2 is an example of an application server stop script for DCE. It runs **/usr/bin/stop.dce** to stop the DCE services.

```
#!/bin/ksh
#
echo "Stopping DCE\n"
/usr/bin/stop.dce
echo "Finished Stopping DCE\n"
```

*Figure 2. Application Server Stop Script Example for DCE*

## Application Server Start Script Example for DFS

Figure 3 on page 17 is an example of an application server start script for DFS. It deletes the endpoint map and security credentials files, runs **/usr/bin/start.dce** to start the configured DCE processes, and then runs **/usr/bin/start.dfs** to start the configured DFS processes.

```
#!/bin/ksh -x
#
#For the HACMP one-for-one (hot) standby configuration with rotating
#resources, save the hostname, associate it with the IP address
#which is the rotating resource and reset it at the end.  Note the
#command substitution string for hostname and the grave accents
#around it.
#
save_hostname = `hostname`
hostname xxx      # replace xxx with your rotating resource hostname
#
#
echo "Removing DCE Files\n"
rm /opt/dcelocal/var/dced/Ep.db
rm /opt/dcelocal/var/security/creds/*

echo "Starting DCE\n"
/usr/bin/start.dce
echo "Finished Starting DCE\n"

echo "Starting DFS\n"
/usr/bin/start.dfs
echo "Finished Starting DFS\n"
#
#For the HACMP one-for-one (hot) standby configuration with rotating
#resources, reset the hostname.
hostname $save_hostname
#
```

*Figure 3. Application Server Start Script Example for DFS*

## Application Server Stop Script Example for DFS

Figure 4 is an example of an application server stop script for DFS. It runs **/usr/bin/stop.dfs** to stop the DFS services, and then runs **/usr/bin/stop.dce** to stop the DCE services. You must also shutdown or reboot the standby/takeover node to stop DFS daemons and release shared DASD before starting HACMP on the owner node.

```
#!/bin/ksh
echo "Stopping DFS\n"
/usr/bin/stop.dfs
echo "Finished Stopping DFS\n"
#
echo "Stopping DCE\n"
/usr/bin/stop.dce
echo "Finished Stopping DCE\n"
```

*Figure 4. Application Server Stop Script Example for DFS*

# Appendix. Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the information. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this information at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Department LZKS
11400 Burnet Road
Austin, TX 78758
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM’s future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM’s suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

## Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

- AIX
- IBM
- RISC System/6000

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

DFS is a trademark of the Transarc Corporation.

Other company, product, and service names may be trademarks or service marks of others.

# Index

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.