IBM

JD Edwards EnterpriseOne 8.9 High Availability and Storage

for IBM @server pSeries Clients

- Follow the step-by-step configuration and set up instructions
- Manage clusters of nodes to deliver high availability
- Review example windows and worksheets for easy installation

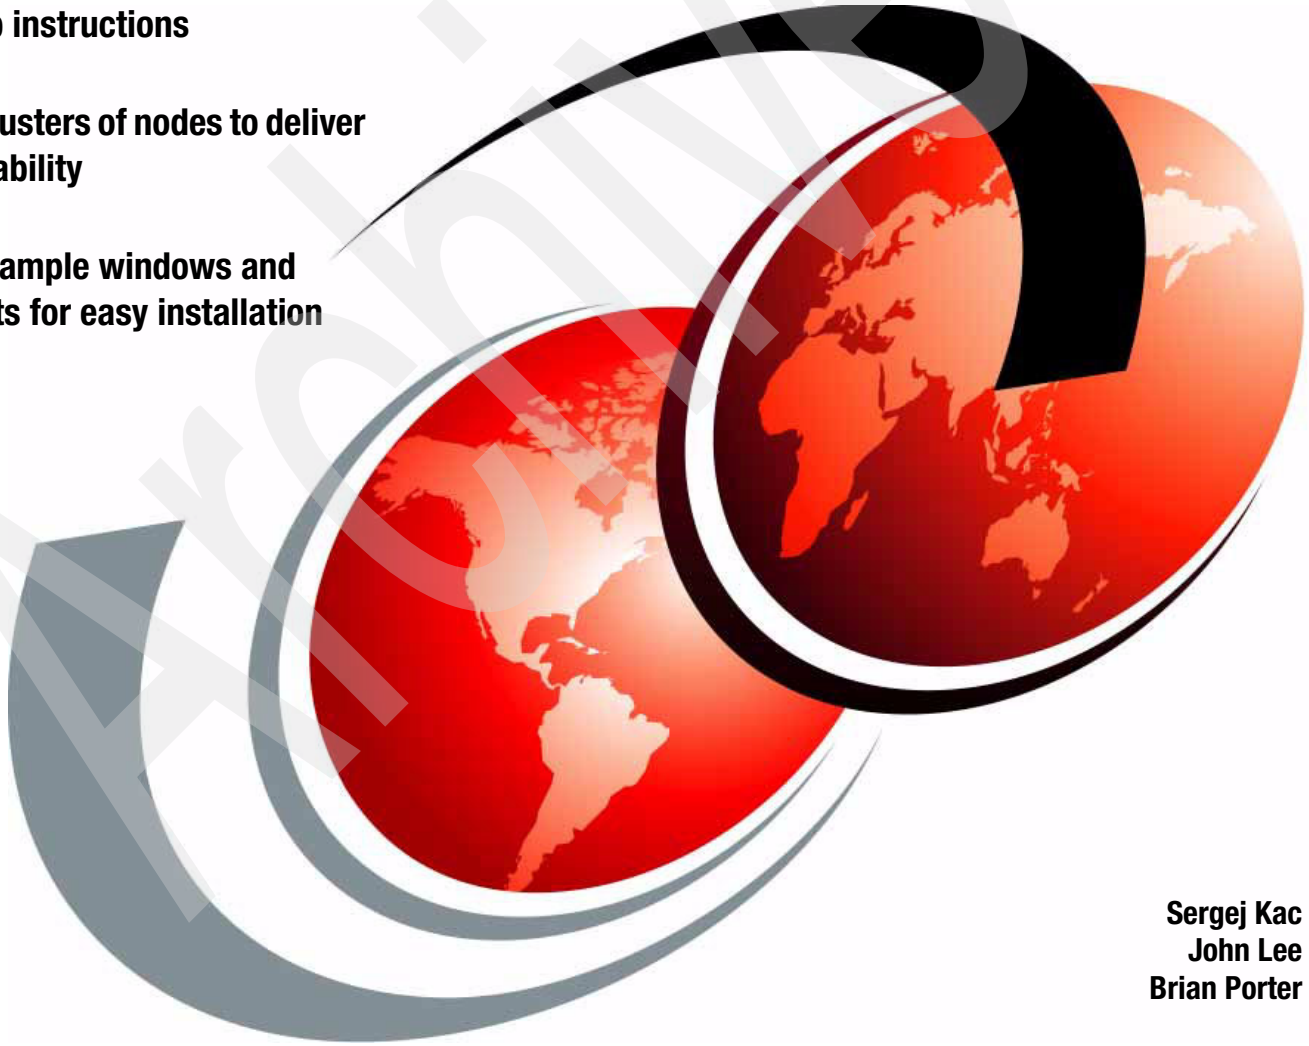

Sergej Kac
John Lee
Brian Porter

Redbooks

ibm.com/redbooks

International Technical Support Organization

JD Edwards EnterpriseOne 8.9 High Availability and Storage for IBM @server pSeries Clients

October 2005

SG24-6445-00

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (October 2005)

This edition applies to Version 4, Release 5, Modification 0.10 of HACMP/ES and Version 8.9, Service Pack 1, of JD Edwards EnterpriseOne.

Contents

Figures

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 described 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.

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 publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication 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.

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.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

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

AIX 5L™	eServer™	pSeries®
AIX®	FlashCopy®	Redbooks (logo) ™
AS/400®	HACMP™	Redbooks™
DB2 Universal Database™	HACMP/6000™	RS/6000®
DB2®	IBM®	Tivoli®
@server®	NetView®	TotalStorage®

The following terms are trademarks of other companies:

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Solaris, Sun, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Preface

As companies depend more and more on their enterprise software for daily operations, they require highly available systems. JD Edwards EnterpriseOne has long enabled their software with the IBM® High Availability Cluster Multi-Processing (HACMP) product for high availability.

This IBM Redbook describes the IBM HACMP/Enhanced Scalability (HACMP/ES) Version 4.5 solution to certify JD Edwards EnterpriseOne 8.9 in an IBM AIX® high availability environment. This book also describes HACMP/ES Version 5.1. It explains how to install and tune HACMP/ES, as well as reviews the setup steps to enable JD Edwards EnterpriseOne for HACMP/ES.

While writing this IBM Redbook, we conducted tests to simulate that the failures and observations made to ensure that the high availability environment performed properly. The results of these tests are also documented.

This redbook highlights and explains those characteristics that are specific to JD Edwards EnterpriseOne. It contains significant window examples and worksheets to help facilitate an easy step-by-step installation process. It does not replace the extensive installation and administration documentation that is already available for HACMP/ES.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Rochester Center.

Sergej Kac is a JD Edwards EnterpriseOne architect working for IBM Business Consulting Services (BCS). He has seven years of experience working on the JD Edwards OneWorld product. His technical experience includes database and system management for UNIX, AS/400®, and Microsoft® Windows® platforms. He is currently working on finishing his Master of Computer Science degree at Penn State University.

John Lee is a programmer for Oracle in Denver, Colorado. He works at Oracle as part of the UNIX Platform team, where he is involved with the verification and validation of the EnterpriseOne 8.9 product. John received a Bachelor of Science degree in computer science from the University of Wyoming in 1986. He worked for EDS in the State Operation Division from 1987 until 1999. While at EDS, he worked on diverse projects including OnStar and Title XIX automation in Wyoming and Rhode Island.

Brian Porter is an IT specialist for the IBM Advanced Technical Support Storage Solutions Group. He has 12 years experience in the UNIX systems, storage, and services industry. He joined IBM five years ago. He spent the first four years in IBM Global Services focused on client AIX systems and storage solutions on IBM @server pSeries® server. He currently supports the ATS Solutions-Storage Alliance independent software vendors (ISVs). He holds a Bachelor of Science degree in electrical and computer engineering with a minor in mathematics from Brigham Young University, Utah.

Thanks to the following people for their contributions to this project:

Leslie Clark
eServer™ Solutions Relationship Manager for JD Edwards EnterpriseOne and World Software, IBM Rochester, Minnesota

Susan Powers
ITSO Project Manager, IBM Rochester, Minnesota

Boyd Fenton
Oracle Internal Competency Center Manager, IBM Denver, Colorado

Mark Trbojevic
Advisory Software Engineer, IBM Beaverton, Oregon

Bob Stravinsky
Manager, UNIX Server Technology/Integration, Oracle, Denver, Colorado

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks™ to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. JLU Building 107-2
  3605 Highway 52N
  Rochester, Minnesota 55901-7829

1

Overview of IBM HACMP/ES

This chapter provides an overview of IBM High Availability Cluster Multi-Processing/Enhanced Scalability (HACMP/ES). It examines briefly the history of the product and describes its basic components. This chapter also introduces you to EnterpriseOne 8.9.

For a full explanation about how to plan, install, configure, and maintain an IBM HACMP/ES environment, see *HACMP for AIX 4.5: Enhanced Scalability Installation and Administration Guide*, SC23-4306.

1.1 High availability

HACMP was developed to enhance and extend AIX software and hardware, where reliability was of primary importance. In its simplest form, HACMP provides a mechanism to take over resources (primarily disks and IP addresses) and to restart applications. In doing so, you avoid losing time due to a possibly long outage required to fix hardware or the time to reboot a system in the event of a software failure. This helps you to quickly resume production.

The planning phase is important when setting up a high availability (HA) environment. Considerable work has been done to make the planning process easier and to eliminate sources of error. The HACMP Version 4.5 installation software provides a cluster worksheet planning program. You access the worksheets through a Web browser from a PC-based system. Using these worksheets, you can apply the resulting configuration directly to a cluster.

Management and maintenance of the cluster, after it is up and running, has been significantly developed in the HACMP software. It is now easier to set up and manage a complex HA environment. The publication *HACMP for AIX 4.5: Enhanced Scalability Installation and Administration Guide*, SC23-4306, which ships with the HACMP/ES product, documents the common administration tasks and tools that exist to help perform these tasks.

Among the tools that make up HACMP, the following tools contribute most to the manageability of a multinode environment:

- Version compatibility to support node-by-node migration for upgrading or migrating HA software
- The monitoring tools **clstat** for cluster status, event error notification, and the *HAView* add-on to use with the IBM Tivoli® NetView® product and log files
- Cluster Single Point of Control (C-SPOC) to allow an operation defined on one node to be executed on all nodes in the cluster (to maintain shared logical manager components, users, and groups)
- Dynamic reconfiguration of resources and cluster topology, known as Dynamic Automatic Reconfiguration (DARE)
- The verification tool, **clverify**, to run after any changes are made to the cluster to verify that the cluster topology and resource settings are valid
- The cluster snapshot utility, **snap**, to help save and restore cluster configurations

In a well-designed cluster, HACMP aims to eliminate any single point of failure, which may be due to disks, network adapters, power supply, processors, or other components.

For a comprehensive discussion about how high availability is achieved and what is possible in a cluster, see *HACMP for AIX 4.5: Enhanced Scalability Installation and Administration Guide*, SC23-4306. You can find additional background and useful information in *Configuring Highly Available Clusters Using HACMP 4.5*, SG24-6845.

1.1.1 A brief version history of HACMP

HACMP has been in the marketplace since 1992 when Version 1.1 of the HA product HACMP/6000™ was announced. Next Release 1.2, Version 2 (Release 2.1), and Version 3 (Release 3.1 and Modification Level 3.1.1) followed. These three releases were all based on AIX 3.2.5. Then four releases of HACMP Version 4 (4.1, 4.2, 4.3, and 4.5) ensued.

With the release of HACMP Version 4.5, the product line was expanded to include HACMP/ES. In July 2003, HACMP/ES Version 5 (Release 5.1) became available, and the earlier HACMP (Classic) product was dropped.

1.1.2 Scalability and HACMP/ES

Scalability, support of large clusters, and therefore, large configurations of nodes and disks lead to a requirement to manage *clusters* of nodes. To address management issues and take advantage of new disk attachment technologies, HACMP/ES was released. It was originally available for the SP hardware, where tools were already in place with PSSP hardware to manage larger clusters.

The technology that underlies HACMP/ES is the IBM RS/6000® Cluster Technology (RSCT) HA services. These services provide greater scalability, notify distributed subsystems of software failure, and coordinate recovery and synchronization among all subsystems in the software stack. Since these services are packed with HACMP/ES, you can run this software on all RS/6000s, not just on SP nodes.

The benefits of HACMP/ES arise from the technologies that underlie the *heart beating* mechanism supplied in RSCT. For example, HACMP/ES can support up to 32 nodes in a cluster. It can help to create a large HA cluster and take advantage of disk technologies, for example, such as storage area networks (SANs). SANs allow disks to be accessed by a large number of systems without a performance penalty.

The Event Management part of RSCT is a tool that is familiar to users of an SP system. Here are some useful sources for information about RSCT:

- *HACMP Enhanced Scalability: User-Defined Events*, SG24-5327
- *RS/6000 Cluster Technology: Event Management Programming Guide and Reference*, SA22-7354
- *RS/6000 Cluster Technology: Group Services Programming Guide and Reference*, SA22-7355

To learn about the family of disk products that is available for the pSeries server, see Chapter 8, “IBM TotalStorage FAStT in an HACMP for AIX environment” on page 47.

1.2 JD Edwards EnterpriseOne 8.9

JD Edwards EnterpriseOne is a complete suite of modular, pre-integrated, industry-specific business applications. It is designed for rapid deployment and ease of administration. It is ideally suited for organizations that manufacture, construct, distribute, service, or manage products or physical assets.

For a further discussion about EnterpriseOne 8.9, refer to:

`http://www.peoplesoft.com/corp/en/products/ent_one/index.jsp`

The remainder of this IBM Redbook discusses the configuration and installation of EnterpriseOne 8.9.

2

What's new in HACMP v5.1

This chapter presents some of the new features that you can find in High Availability Cluster Multi-Processing/Enhanced Scalability (HACMP/ES) 5.1. To learn about configuration and certification on HACMP/ES 5.1, see Chapter 3, "EnterpriseOne 8.9 on HACMP/ES v5.1 configuration and certification" on page 9.

2.1 Software prerequisites

You can install and run HACMP 5.1 in the following environments:

- AIX 5.1 ML3 or later with RSCT Version 2.2.1.30 or later
- AIX 5.2 ML1 or later with RSCT Version 2.3.1.0 or later (version 2.3.1.1 recommended)

2.2 Summary of changes

This section outlines the changes that are offered with HACMP v5.1 over v4.5.

2.2.1 Enhancements

Some of the major features that have improved in this release include:

- Cluster communication daemon (eliminates need to set up *.rhost*)
- Simplification of configuration and maintenance
- Online planning worksheet enhancements
- Forced “varyon” of volume groups
- Custom Resource Groups
- Heartbeat monitoring of service IP labels
- Heartbeat monitoring over aliases
- Heartbeat monitoring over disks
- Cluster Single Point of Control (C-SPOC) enhancements
- General Parallel File System (GPFS) integration
- Fast disk takeover
- Cluster verification and synchronization enhancements
- Resource group management

Additional enhancements include:

- All previous HACMP product lines, such as HAS/CRM, HACMP/ES/CRM, combined and replaced in HACMP 5.1
- New add-on feature HACMP XD (HA+PPRC, still to come HAGEO)
- JFS2 file system support on AIX 5.2B
- Two hundred and fifty-six network interfaces, devices, and labels (previously 120 via 40 of each type)
- Support for up to 64 resource groups
- Support for dynamic tracking of Fibre Channel devices

2.2.2 Features removed and restrictions

HACMP 5.1 still has some limitations, which include:

- Older classes and Customer Resource Management (CRM) variants have been removed.
- The following network types are not supported:
  - SOCC
  - Serial Line Internet Protocol (SLIP)
  - 802_ether (Ethernet protocol 802.3)
  - Generic IP
  - Fibre Channel switch (FCS)
  - Etherchannel
  - Virtual IP Addressing (VIPA)
  - IPv6
- AIX Connections Services is no longer supported.
- AIX Version 5.2 does not support Serial Storage Architecture (SSA) concurrent volume groups.
- Cluster Lock Manager works on 32-bit kernel only. There is no 64-bit support.
- AIX Version 5.2 defaults to Simple Network Management Protocol (SNMP) Version 3. Version 1 is required to run C-SPOC and **clinfo**.

3

EnterpriseOne 8.9 on HACMP/ES v5.1 configuration and certification

This chapter describes the environment that we used to certify EnterpriseOne 8.9. It outlines the software installed, the steps to set up the server cluster for High Availability Cluster Multi-Processing/Enhanced Scalability (HACMP/ES) v5.1, and the steps to install and set up the EnterpriseOne 8.9 application software.

The worksheets that we used to complete this work are documented in “HACMP/ES planning worksheets: Completed” on page 58.

3.1 Software versions

The JD Edwards EnterpriseOne HACMP/ES certification tests were performed using the following software packages and versions.

- AIX 5L™ Version 5.1 with the latest fix packs

  Note: The UNIX kernel must be booted in 64-bit mode for DB2® Universal Database™ (UDB) Version 8.1.

- HACMP/ES Version 5.1
- DB2 UDB 8.1 (64-bit) with Fix Pack 1A

  Note: Fix Pack 1A is a special fix pack that IBM provides for JD Edwards customers who use DB2 with EnterpriseOne 8.9. You can find this fix pack on the Web at:

  `ftp://ftp.software.ibm.com/ps/products/db2/fixes/english-us/jdedwards/v81fp1a/`

- JD Edwards EnterpriseOne 8.9 with Update 1

Chapter 7, "JD Edwards EnterpriseOne 8.9 HACMP/ES testing" on page 45, describes the testing procedure for the certification of the software.

3.2 Setting up the cluster for JD Edwards EnterpriseOne 8.9

Our test cluster was built using two pSeries Model 270 servers and an array of Serial Storage Architecture (SSA) disks. The ttclrs1 node hosted JD Edwards EnterpriseOne 8.9 applications. The ttclrs2 node hosted DB2 UDB databases used for EnterpriseOne 8.9. In case of failure, applications failed over to the other node.

Figure 3-1 illustrates the cluster configuration.

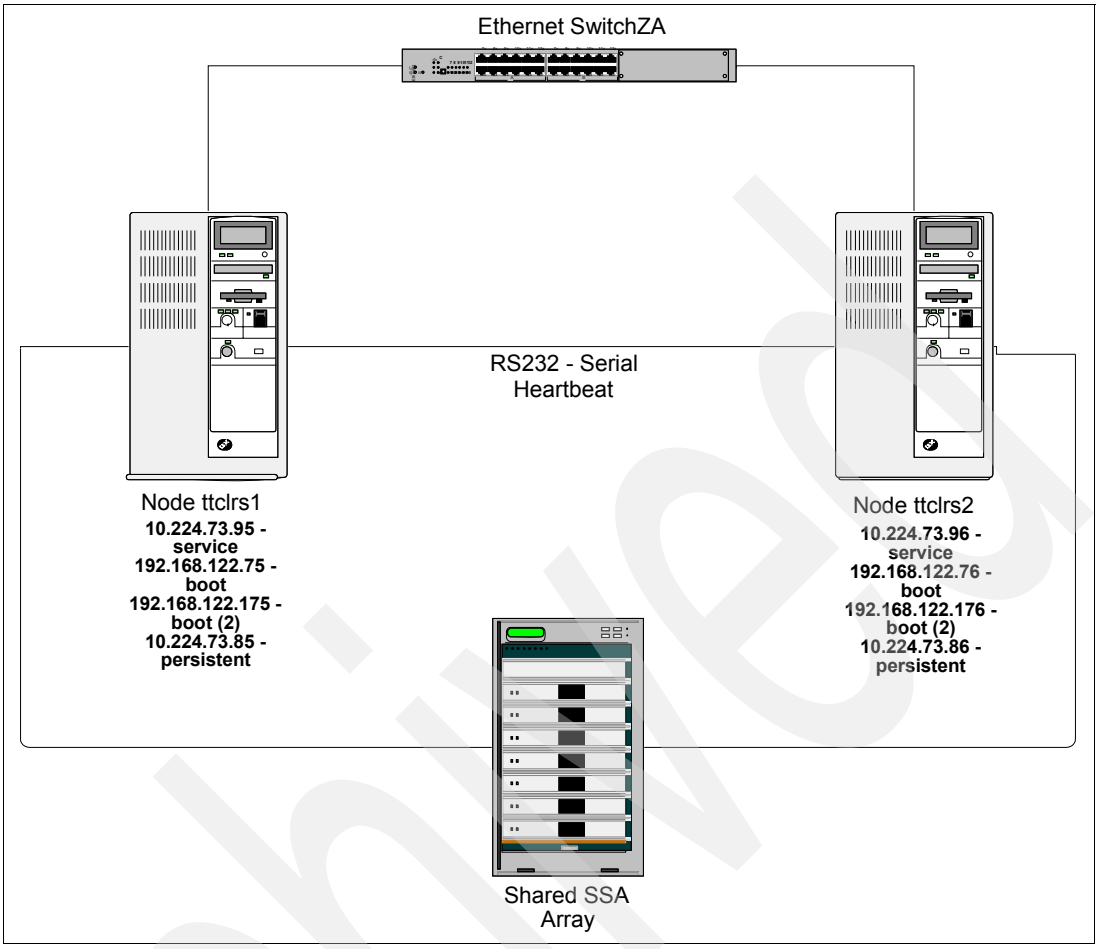

Figure 3-1 HACMP/ES v5.1 cluster topology

3.2.1 Network configuration

In our cluster configuration, we defined the following network addresses:

- Two boot addresses for each node
- Service IP address for each application server
- Persistent IP address for each node
- RS232 serial link between nodes used for heartbeat

Figure 3-2 illustrates these parameters.

False failures can occur during heavy TCP/IP traffic when heartbeat information is delayed from reaching the other node in a predetermined time frame. To avoid false failures, define another type of secondary network, either the RS232 network or a network over the SSA connection.

3.2.2 Setting up the cluster

We set up the clustered servers based on the following sequence of actions:

1. Prepare AIX 5L Version 5.1 for clustering. Install the necessary AIX fix packs.
2. Set up HACMP/ES v5.1.
   a. Install and set up HACMP/ES base software on both nodes.
   b. Install and set up the latest clustering fix pack for 5.1 on both nodes prior to the initial cluster setup.
   c. Complete the HACMP/ES planning worksheets.

      Refer to “HACMP/ES planning worksheets: Completed” on page 58 to see the worksheets used for this test configuration.
   d. Set up the cluster name, cluster nodes, and network interfaces to be used for the EnterpriseOne cluster. Changes were replicated to the other node using topology and cluster synchronization.
3. Install the JD Edwards EnterpriseOne database.
   a. Install the database from the distribution media.
   b. Apply the latest database fix packs that are supported by Oracle.
4. Install JD Edwards EnterpriseOne 8.9.
   a. Install the EnterpriseOne 8.9 software on the deployment and enterprise servers.
   b. Configure EnterpriseOne software to work with the DB2 UDB databases.
   c. Create shared databases using supplied JD Edwards database creation scripts.
5. Set up HACMP/ES v5.1 for EnterpriseOne 8.9 software,
   a. Application servers and resource groups for DB2 UDB and EnterpriseOne 8.9 software (see Figure 3-2) were included.

      Resource Group rgJDE controls the application server asJDE, service IP addresses and shared SSA vgJE and vgJDE2 volume groups dedicated to the ttclrs1 node.

      Resource Group rgDB2 controls the application server asDB2, service IP addresses and shared SSA vgLData, vgJRData, and vgIndex volume groups dedicated to the ttclrs2 node.
   b. Application scripts to start up and shut down application servers in the event of failure or resource group migration were added.

      The resource groups were defined as cascading without fallback. See Chapter 6, “Application script examples for EnterpriseOne 8.9 with HACMP/ES” on page 41.

For HACMP cluster setup, follow the guidelines in *HACMP for AIX 4.5: Enhanced Scalability Installation and Administration Guide*, SC23-4306.

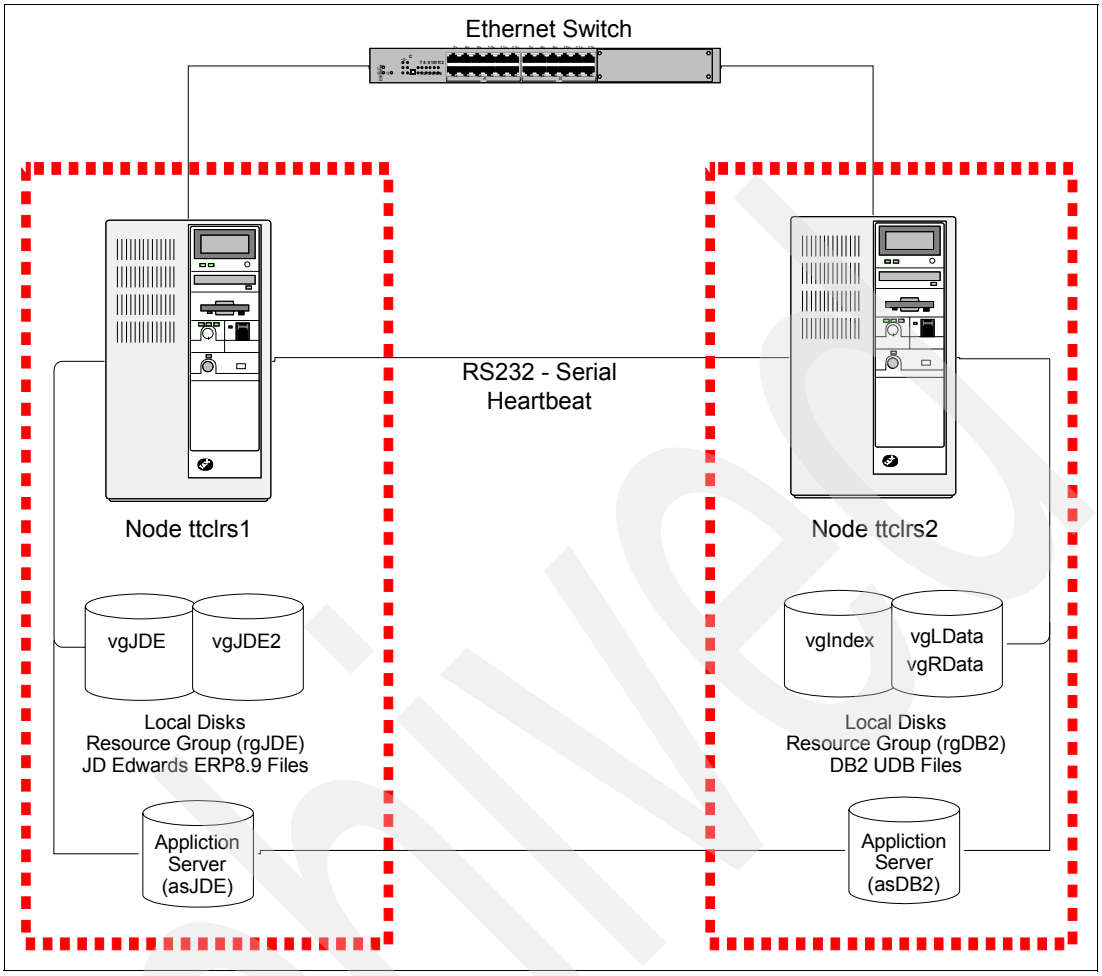

Figure 3-2 Logical layout on HACMP/ES cluster

Table 3-1 identifies the volume groups defined for the ttclrsHA cluster.

Table 3-1 Volume groups defined for cluster ttclrsHA

vgRData	**lvRData**	**/RData**	**rgDB2**
vgJDE	lvJDE	/u07	rgJDE
vgJDE2	lvJDE2	/u08	rgJDE
vgIndex	lvIndex	/Index	rgDB2
vgRData	lvRData	/RData	rgDB2
vgLData	lvLData	/LData	rgDB2

When configuring volume groups for use in the clustering environment, be sure the volume group does not automatically activate at system restart. HACMP-used volume groups are activated when the resource group that controls that volume group is activated.

Example 3-1 shows our cluster configuration using the **cltopinfo** command.

Example 3-1 Printout of the ttclrsHA cluster configuration

```
Cluster Description of Cluster: ttclrsHA
Cluster Security Level: Standard
There are 2 node(s) and 2 network(s) defined

NODE ttclrs1:
        Network IPnet
                ttclrs2_svc     10.224.73.96
                ttclrs1_svc     10.224.73.95
                ttclrs1_boot    192.168.122.75
                ttclrs1_boot2   192.168.122.175
        Network rsNet
                ttclrs1tty1     /dev/tty1

NODE ttclrs2:
        Network IPnet
                ttclrs2_svc     10.224.73.96
                ttclrs1_svc     10.224.73.95
                ttclrs2_boot2   192.168.122.186
                ttclrs2_boot    192.168.122.86
        Network rsNet
                ttclrs2tty1     /dev/tty1

Resource Group rgJDE
        Behavior                cascading
        Participating Nodes     ttclrs1 ttclrs2
        Service IP Label                ttclrs1_svc

Resource Group rgDB2
        Behavior                cascading
        Participating Nodes     ttclrs2 ttclrs1
        Service IP Label                ttclrs2_svc
```

Notice that for the rgJDE resource group, the participating nodes are ttclrs1 and ttclrs2. For the resource group rgDB2, the participating nodes are ttclrs2 and ttclrs1. The node listed first is the primary node.

The resource groups have the additional settings listed in Table 3-2.

Table 3-2 Additional settings for resource groups

Resource group	rgJDE	rgDB2
Application server	asJDE	asDB2
Application server start	local_start_ow	local_start_db2
Application server stop	local_stop_ow	local_stop_db2
Volume groups	vgJDE vgJDE2	vgIndex vgRData vgLData
File systems	/u07 /u08	/Index /RData /LData

The nodes themselves have the following additional information configured:

```
Node ttclrs1 persistent IP address ttclrs1_prs 10.224.73.85
Node ttclrs2 persistent IP address ttclrs2_prs 10.224.73.86
```

Note: Persistent IP addresses are not required for cluster functionality, but are useful in administering the systems in the cluster.

3.3 Installing and setting up the EnterpriseOne 8.9 deployment and enterprise servers

To set up the deployment and enterprise servers, follow these steps:

1. Install the necessary prerequisites for EnterpriseOne 8.9. See the *Minimum Technical Requirements* (MTR) topic on the Oracle customer resource Web site for EnterpriseOne 8.9. You can find this information on the Oracle PeopleSoft Customer Connection site, which you can access on the Web at:

   `http://www.peoplesoft.com/corp/en/public_index.jsp`

   When you reach this site, complete these tasks:

   a. Click the **LOG IN** button for customers and partners in the upper right corner.

   b. Sign in with your user name and password.

   c. Select Implement, **Optimize + Upgrade** →**Optimization Guide** →**Supported Platforms** →**JD Edwards EnterpriseOne**.

2. Plan the installation of the EnterpriseOne 8.9 software on the AIX cluster.

   In our lab, we decided to create two application servers using two different resource groups. Our first application server runs EnterpriseOne services, and the other hosts EnterpriseOne databases.

3. Install EnterpriseOne 8.9 software on the deployment server.

4. Install the latest Planner Electronic Software Update (ESU) for EnterpriseOne 8.9.

5. Log in to EnterpriseOne 8.9.

6. In the Fast Path, type `P98402` and press Enter.

7. The Database Creation Scripts Director form (Figure 3-3) opens. For Select the System Type, select **UNIX**. Then for Select the Database Type, specify your database type. For our testing, we used the DB2 UDB database. Since only UDB/DB2 is available on the form, we select **UDB/DB2**. Click **Next**.

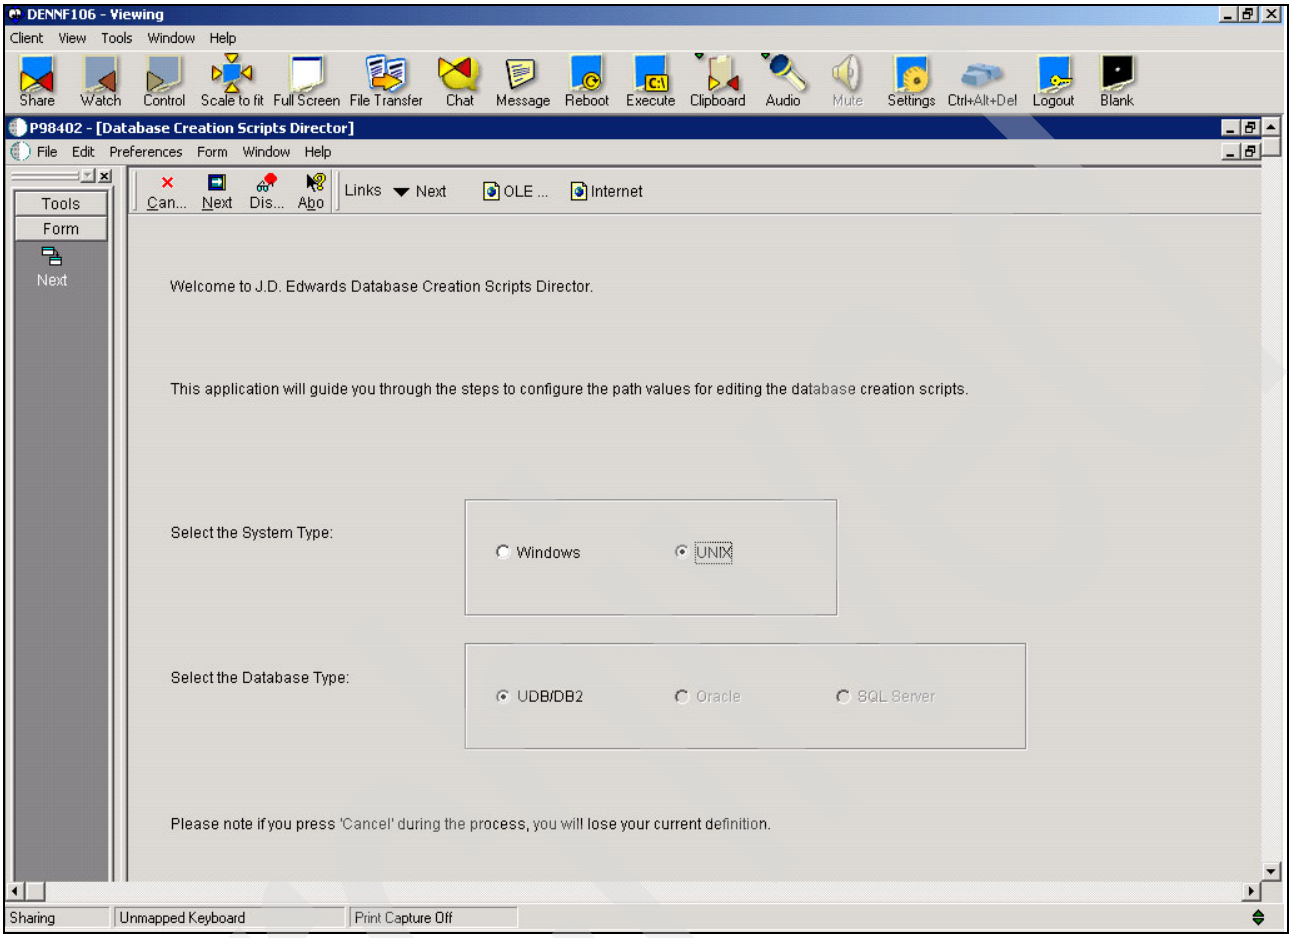

Figure 3-3 Database Creation Scripts Director (Part 1 of 2)

8. The form shown in Figure 3-4 opens. Make sure that the INSTANCE name, SERVICE port, and NODENAME are correct. Verify your NODENAME and SERVICE by entering the following command:

   ```
   db2 list node directory
   ```

   See Figure 3-2 on page 13 for information about the database file locations.

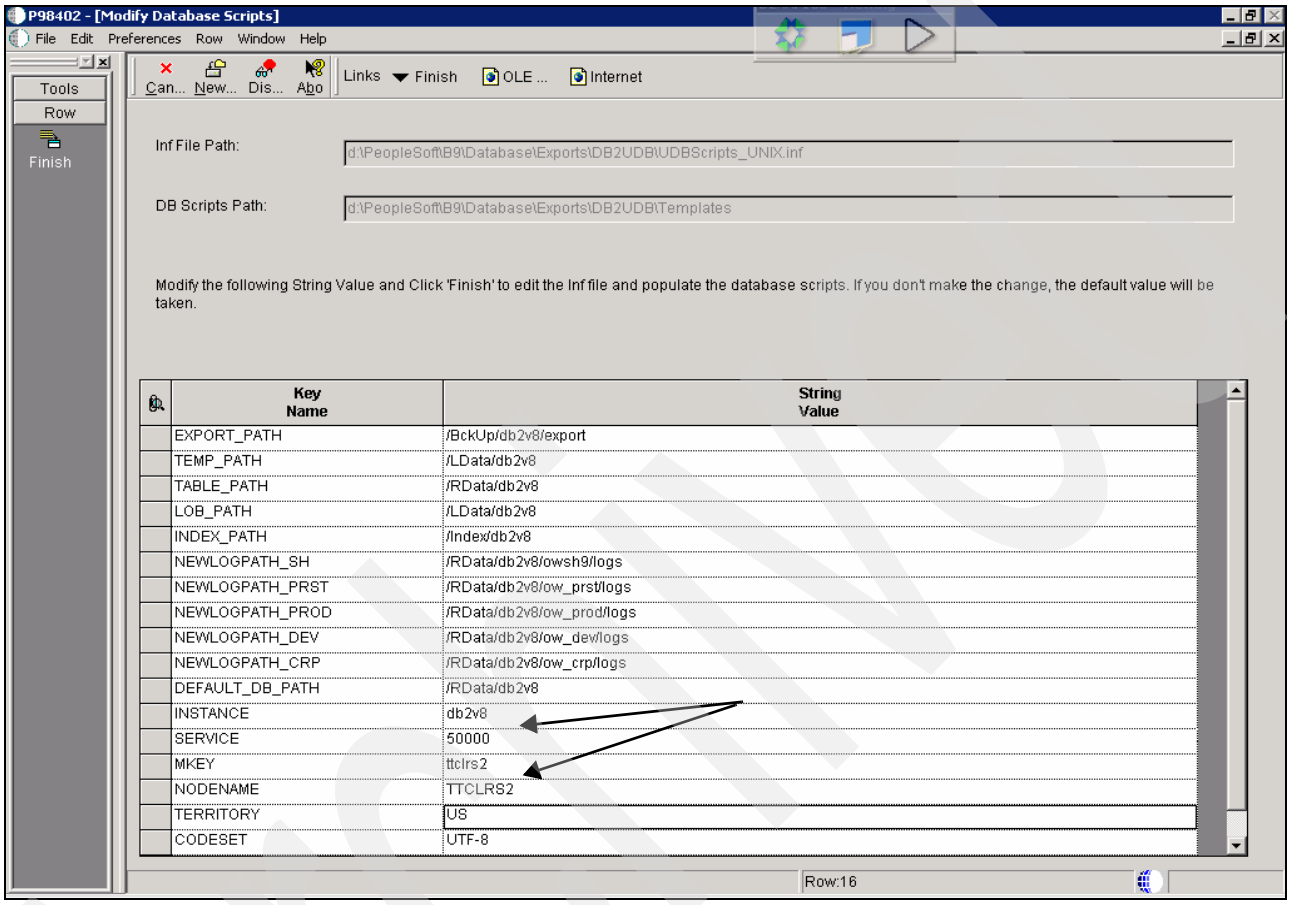

Figure 3-4 Database Creation Scripts Director (Part 2 of 2)

9. Changes in the form update the necessary database creation scripts for your environment. Click **Finish** when the form is complete.
10. Create the necessary database structures using the database scripts modified in step 9.

11. Create an installation plan with the cluster service names for the enterprise server and database servers. See Figure 3-5, Figure 3-6, and Figure 3-7 on page 20.

    a. During the enterprise server definition part of the installation plan, be sure to select **Form →Advanced**.

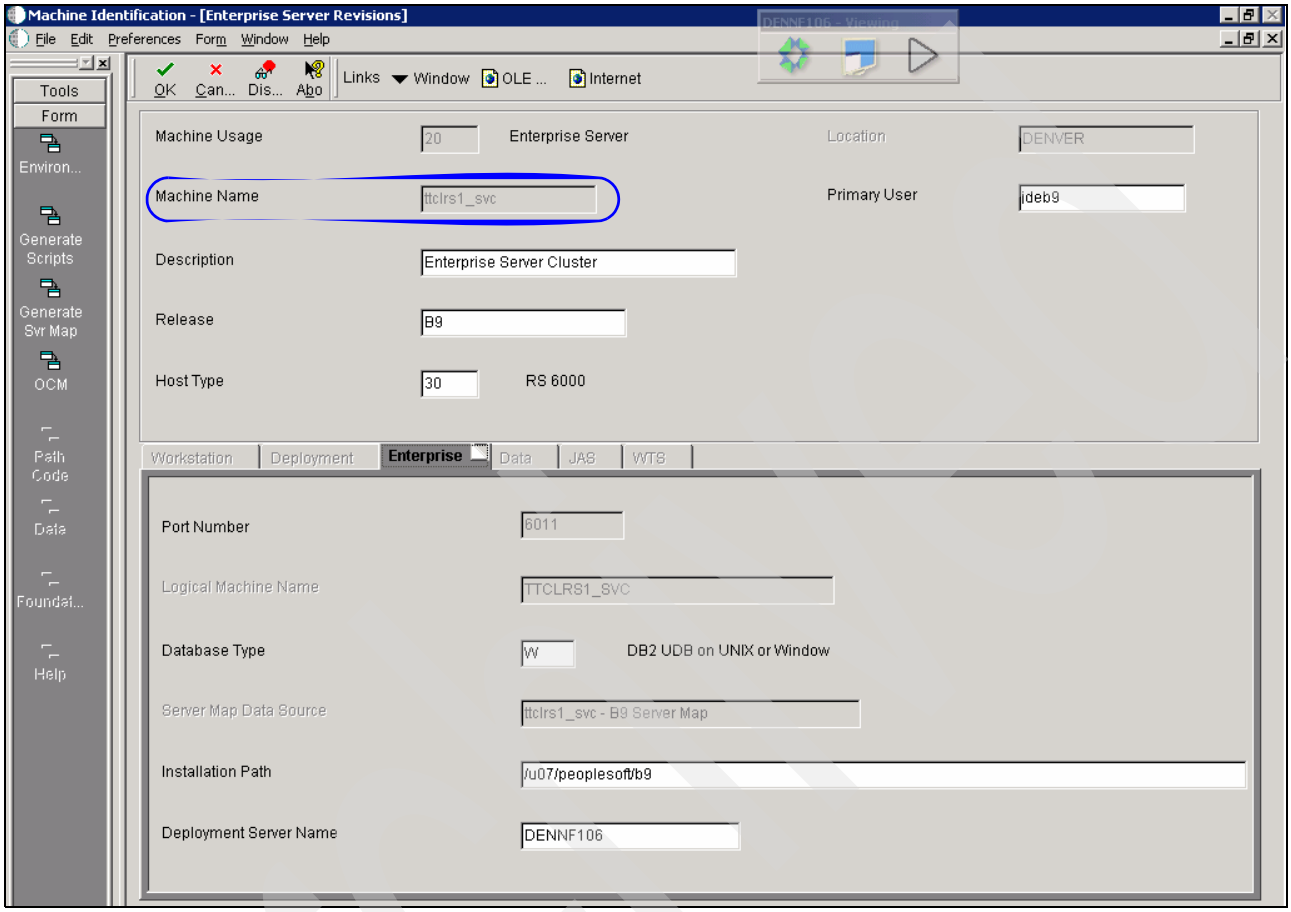

Figure 3-5 Clustered enterprise server

b. For the logical server SVR data source, make sure that **Clustered** is selected (Figure 3-6). This option disables failover to a LOCAL logic data source and continues retrying until clustered resources become available.

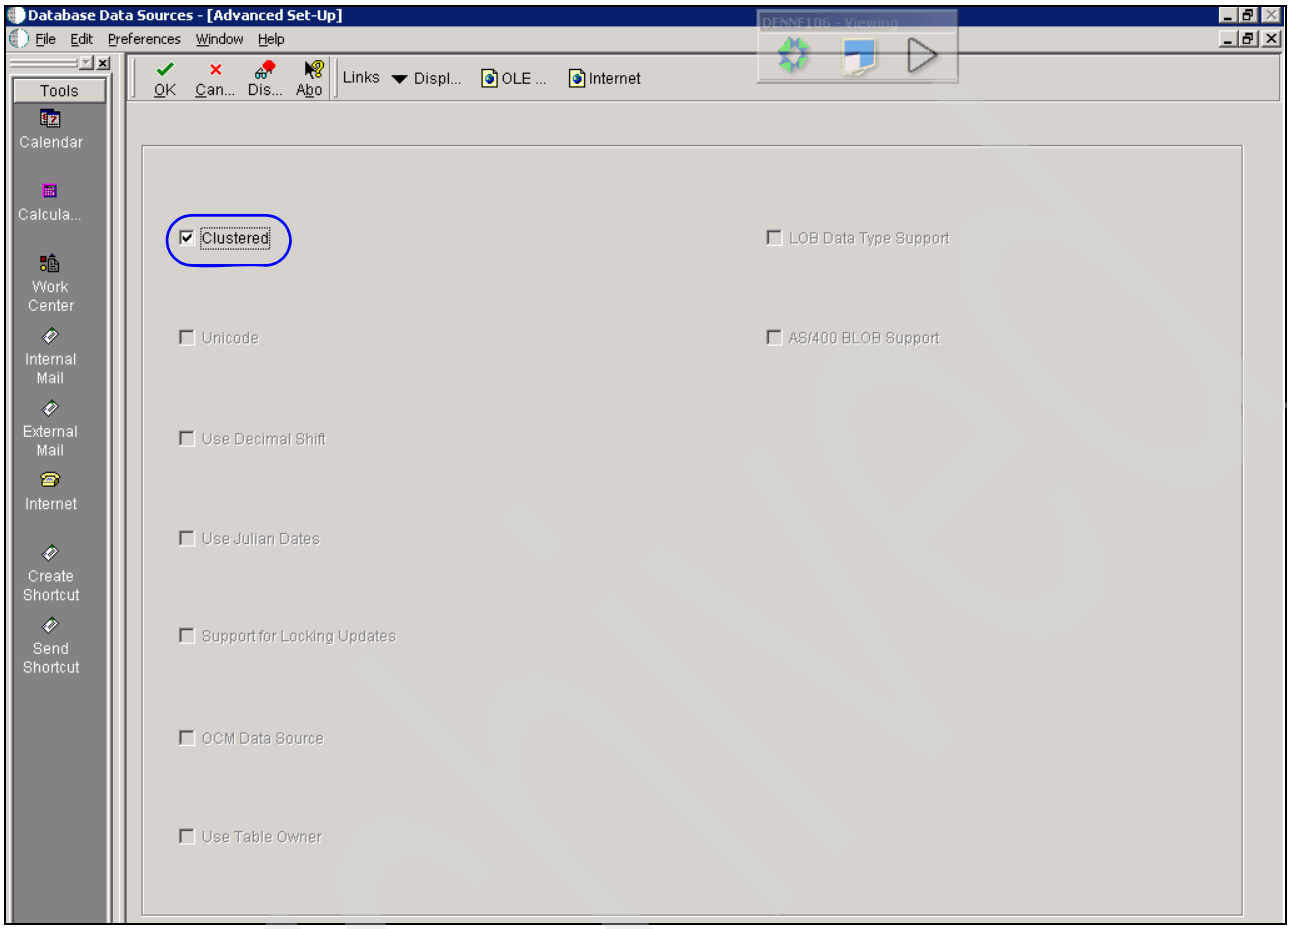

Figure 3-6 Clustered enterprise server: Advanced setup

c. Add an additional database server using the clustered service name of application server that hosts databases. See Figure 3-7.

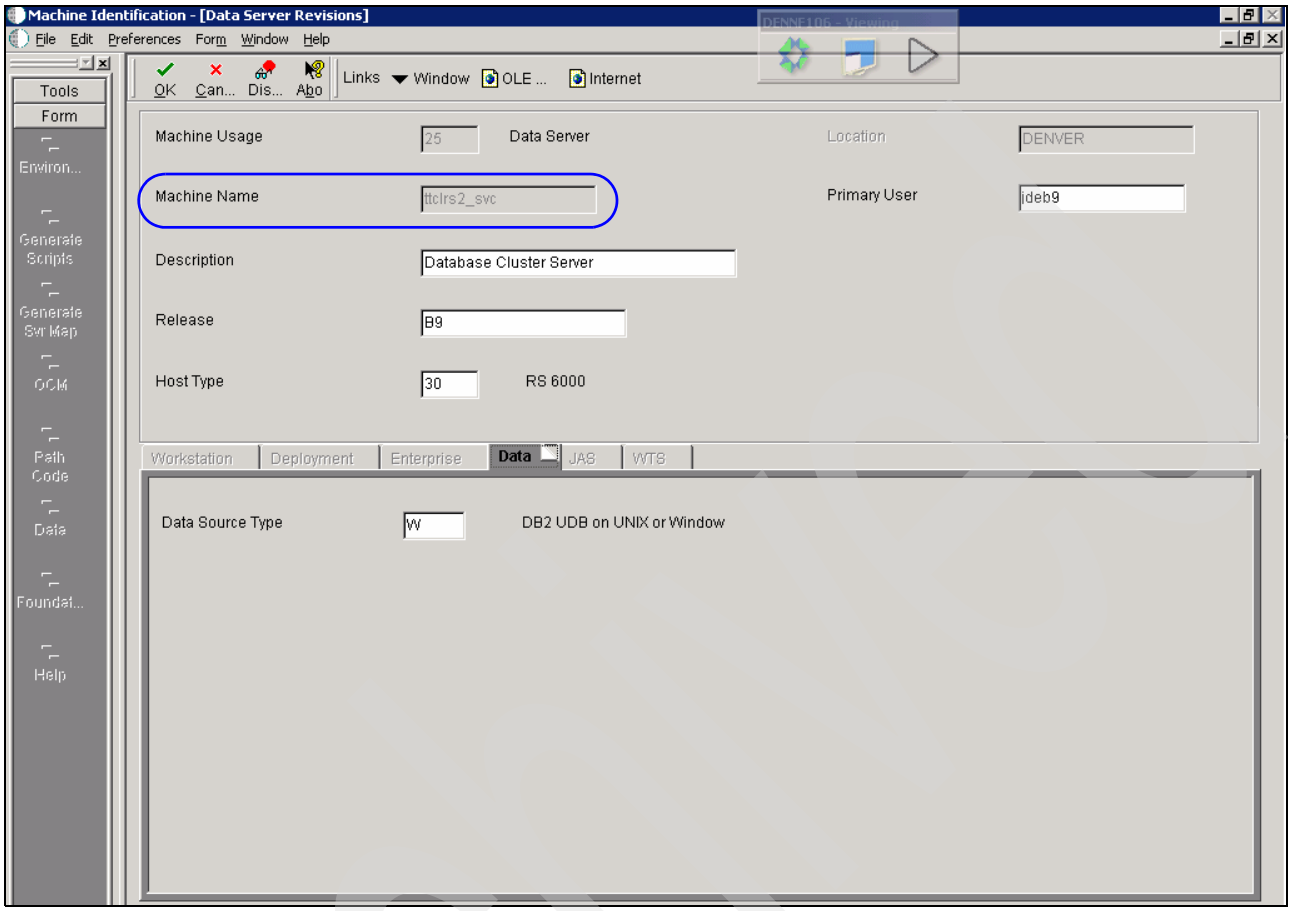

Figure 3-7 Clustered database server

d. During the creation of database data sources (for example, System, Data Dictionary, and Object Librarian), ensure that the Database Server Name is the same one that you entered in step c. See Figure 3-8.

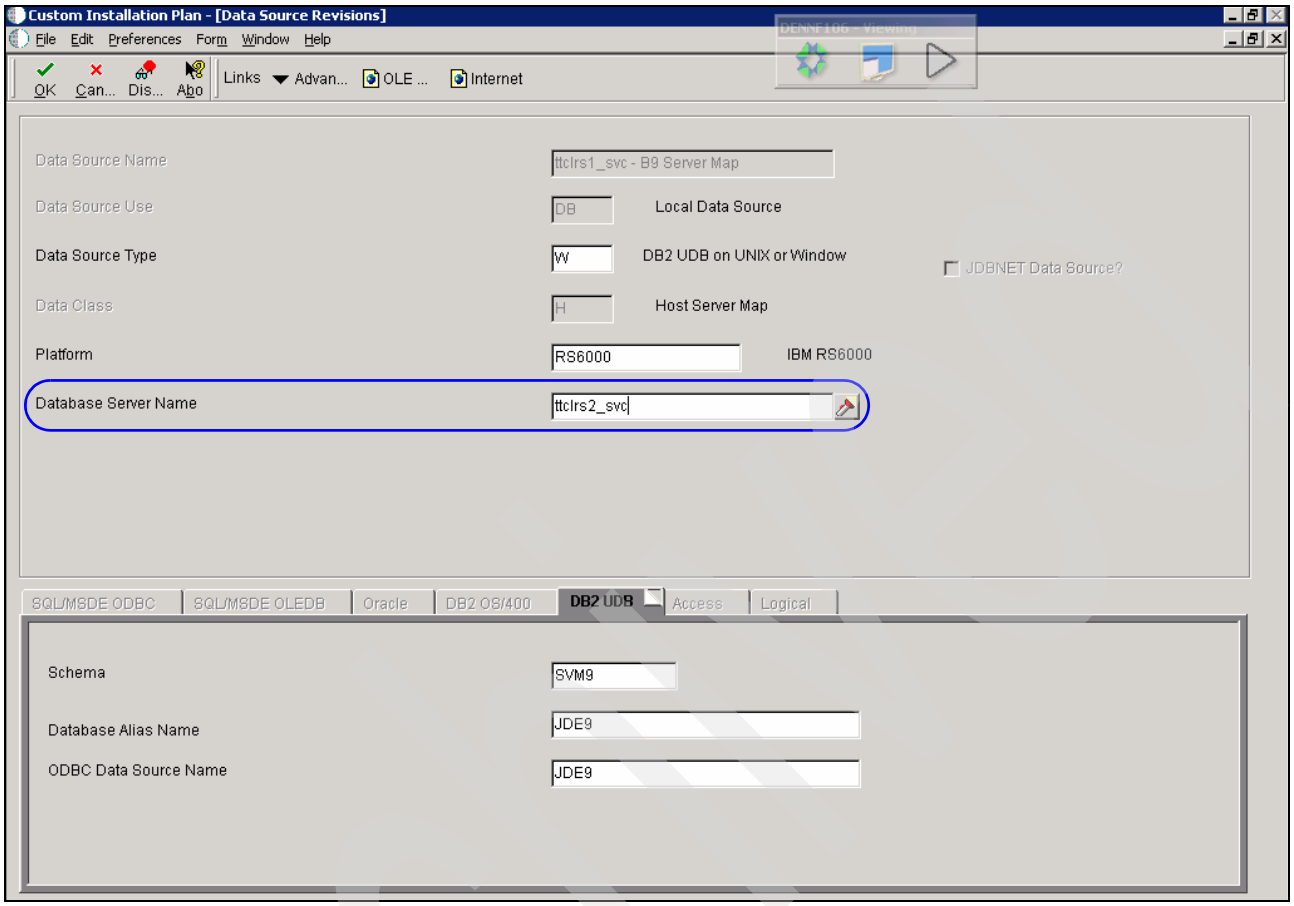

Figure 3-8 Server map data source example

12. Continue with the normal EnterpriseOne 8.9 installation steps to install software on the enterprise server.

13. Modify the JDE.INI file on the enterprise server with the correct cluster service name or names.

    In our example, we documented the DB settings as shown in Example 3-2. The clustered database service name is ttclrs2_svc.

Example 3-2 DB system settings

```
[DB SYSTEM SETTINGS]
Default Role=*ALL
Default Env=DV9
Default PathCode=DV9
Base Datasource=ttclrs1_svc - B9 Server Map
Object Owner=SVM9
Server=ttclrs2_svc
….
```

Example 3-3 shows the security settings for our example. The clustered enterprise service name is ttclrs1_svc.

Example 3-3 Security settings

```
[SECURITY]
Default Role=*ALL
DefaultEnvironment=DV9
DataSource=System - B9
SecurityServer=ttclrs1_svc
….
```

Example 3-4 shows the cluster settings for our example.

Example 3-4 Cluster settings

```
[CLUSTER]
PrimaryNode=ttclrs1_svc
```

14. After the cluster services are running on both nodes and the databases are up and running, run PORTTEST. If you run PORTTEST before cluster services are up, it fails to find the correct host name and IP address.

For further information about the installation and setup of EnterpriseOne software, see the *EnterpriseOne Release 8.9 UNIX Installation Guide* provided by Oracle.

4

Application script examples for EnterpriseOne 8.9 with HACMP/ES v5.1

High Availability Cluster Multi-Processing (HACMP) requires that the administrator define a series or script for starting, stopping, and monitoring applications. In our example, we used the scripts provided by DB2 to start and stop the database. This chapter explains the scripts used to start, stop, and monitor JD Edwards EnterpriseOne 8.9.

4.1 Application startup scripts

Oracle provides the RunOneWorld.sh and EndOneWorld.sh scripts for the EnterpriseOne administrator to start and stop the EnterpriseOne 8.9 server. HACMP executes the configured scripts in the root directory. Consequently, an intermediate script is needed to switch user execution to the EnterpriseOne administrator user in order to execute the RunOneWorld.sh and EndOneWorld.sh scripts.

In our testing, we found that, in some failovers, HACMP tried to start EnterpriseOne before the disks where ready. A timing loop was added to the script to prevent this.

Unlike HACMP application monitor scripts, parameters can be freely passed into application startup and shutdown scripts. Three parameters are passed in:

- Maximum Number of Tries
- Sleep Time
- JDE User

The local_start_ow script starts EnterpriseOne services. This script is saved in the /usr/local/bin directory on both nodes. Example 4-1 shows the local_start_ow script.

Example 4-1 The local_start_ow script

```
#! /usr/bin/ksh
(( TRUE = 0 ))
(( FALSE = 1 ))
(( COUNT = 0 ))
OW_STARTED=${FALSE}

if [[ $# -ne 3 ]]
then
  echo "Usage: $0 <Maximum Number of Tries> <Sleep time> <JDE User> "
  exit 1
fi
MAX_TRIES=$1
SLEEP_4=$2
JDE_USER=$3
echo "$0 is Running at `date`"
while [ ${OW_STARTED} -ne ${TRUE} ]
do
  if [[ ${COUNT} -ge ${MAX_TRIES} ]]
  then
    echo "\nTimed out waiting for /home/${JDE_USER} to be mounted"
    break
  else
    if [[ -f /home/${JDE_USER}/.profile ]]
    then
      . /home/${JDE_USER}/.profile
      if [[ -d ${EVRHOME}/log ]]
      then
        LOGDATE=`date +"%m%d%y%H%M%S"`
        mkdir -m 777 ${EVRHOME}/log/${LOGDATE}
        chown ${JDE_USER} ${EVRHOME}/log/${LOGDATE}
        mv ${EVRHOME}/log/jde*log ${EVRHOME}/log/${LOGDATE}
        echo "Starting OneWorld as ${JDE_USER}"
        /usr/bin/su - ${JDE_USER} "-c ${SYSTEM}/bin32/RunOneWorld.sh"
        OW_STARTED=${TRUE}
        echo "\nScript RunOneWorld.sh ran.  Check logs for results"
      else
        echo "EVRHOME not found at ${EVRHOME}"
```

```
        (( COUNT = ${COUNT} + 1 ))
        sleep ${SLEEP_4}
fi
    else
       (( COUNT = ${COUNT} + 1 ))
       sleep ${SLEEP_4}
       echo "$0 Could not start OneWorld. Tried $COUNT times"
    fi
  fi
done
exit ${OW_STARTED}
```

4.2 Application shutdown scripts

The EnterpriseOne 8.9 shutdown script is similar in form to the startup script. As with the startup script, we pass the following three parameters from the command line.

- Maximum Number of Tries
- Sleep Time
- JDE User

The EnterpriseOne 8.9 shutdown script needs only the EnterpriseOne 8.9 administrator user ID to be passed. The local_stop_ow script is saved in the /usr/local/bin directory on both nodes. Example 4-2 shows the local_stop_ow script.

Example 4-2 The local_stop_ow script

```
#! /usr/bin/ksh
(( TRUE = 0 ))
(( FALSE = 1 ))
(( COUNT = 0 ))
MAX_TRIES=$1
SLEEP_4=$2
OW_STOP=${FALSE}
JDE_USER=$3

if [[ $# -ne 3 ]]
then
  echo "Usage: $0 <Maximum Number of Tries> <Sleep time> <User>"
  exit 1
fi
echo "$0 is Running at `date`"
while [ ${OW_STOP} -ne ${TRUE} ]
do
  if [[ ${COUNT} -ge ${MAX_TRIES} ]]
  then
    echo "\nTimed out waiting for /home/${JDE_USER} to be mounted"
    break
  else
    if [[ -f /home/${JDE_USER}/.profile ]]
    then
      . /home/${JDE_USER}/.profile
      echo "Ending OneWorld as ${JDE_USER}"
      /usr/bin/su - ${JDE_USER} "-c ${SYSTEM}/bin32/EndOneWorld.sh"
      OW_STOP=${TRUE}
      echo "\nScript EndOneWorld.sh ran.  Check logs for results"
else
       (( COUNT = ${COUNT} + 1 ))
```

```
      sleep ${SLEEP_4}
      echo "$0 Could not stop OneWorld. Tried $COUNT times"
    fi
  fi
done

exit ${OW_STOP}
```

4.3 Application monitor script for application server asJDE

Example 4-3 shows the ow_app_monitor script.

Example 4-3 The ow_app_monitor script

```
#! /usr/bin/ksh

count=`ps -ef | grep jdenet_n | grep -v grep |wc -l`
if [[ "$count" -gt 0 ]]
then
  return_code=0
else
  return_code=1
fi
exit ${return_code}
```

4.4 Application monitor settings for application server asJDE

HACMP provides a simple single process monitor setting. This option does not work with EnterpriseOne 8.9. The main process for EnterpriseOne 8.9, jdenet_n, forks additional processes to handle additional load. HACMP detects these additional processes and triggers the defined response. Therefore, a simple script is created to monitor whether one or more jdenet_n processes are present. If no jdenet_n processes are present, HACMP considers this a failure and triggers the defined responses.

Note: After this monitoring is configured, you should only start and stop EnterpriseOne 8.9 through HACMP.

The local_mon_ow script is saved in the /usr/local/bin directory on both nodes. Example 4-4 shows the local_mon_ow script.

Example 4-4 The local_mon_ow script

```
Change/Show Custom Application Monitor

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                        [Entry Fields]
  Application Server Name                asJDE
* Monitor Method                        [/usr/local/bin/local_mon_ow>
  Monitor Interval                      [10]                        #
  Hung Monitor Signal                   [9]                         #
* Stabilization Interval                [60]                        #
  Restart Count                         [3]                         #
  Restart Interval                      [231]                       #
* Action on Application Failur          [fallover]                 +
  Notify Method                         []
  Cleanup Method              [/usr/local/bin/local_stop_ow 5 5 jdeb9>
  Restart Method              [/usr/local/bin/local_start_ow 5 5 jdeb9>
```

5

EnterpriseOne 8.9 on HACMP/ES v4.5 configuration and certification

This chapter describes the environment used to certify EnterpriseOne 8.9. It discusses the software that was installed, the steps to set up the server cluster for High Availability Cluster Multi-Processing/Enhanced Scalability (HACMP/ES) v4.5, and the steps to install and set up the EnterpriseOne 8.9 application software.

The worksheets used to complete this work are documented in “HACMP/ES planning worksheets: Completed” on page 58.

5.1 Software versions

The JD Edwards EnterpriseOne HACMP/ES certification tests were performed using the following software packages and versions.

- AIX 5L Version 5.1 with Maintenance Level 3

  Note: The UNIX kernel must be booted in 64-bit mode for DB2 Universal Database (UDB) Version 8.1.

- HACMP/ES Version 4.5.0.10
- DB2 UDB 8.1 (64-bit) with Fix Pack 1A

  Note: Fix Pack 1A is a special fix pack that IBM provides for JD Edwards customers who use DB2 with EnterpriseOne 8.9. You can download the fix pack from the Web at:

  `ftp://ftp.software.ibm.com/ps/products/db2/fixes/english-us/jdedwards/v81fp1a/`

- JD Edwards EnterpriseOne 8.9 with Update 1

Chapter 7, "JD Edwards EnterpriseOne 8.9 HACMP/ES testing" on page 45, explains the testing procedure for the certification of the software.

5.2 Setting up the cluster for JD Edwards EnterpriseOne 8.9 certification

Our test cluster was built using two pSeries Model 270 servers and array of Serial Storage Architecture (SSA) disks. Both servers participated in the ttclrsHA cluster, with each server hosting one resource group (EnterpriseOne 8.9 or DB2). Each node was equipped with four network cards (two used for testing) and a serial connection between them.

Figure 5-1 illustrates the cluster configuration.

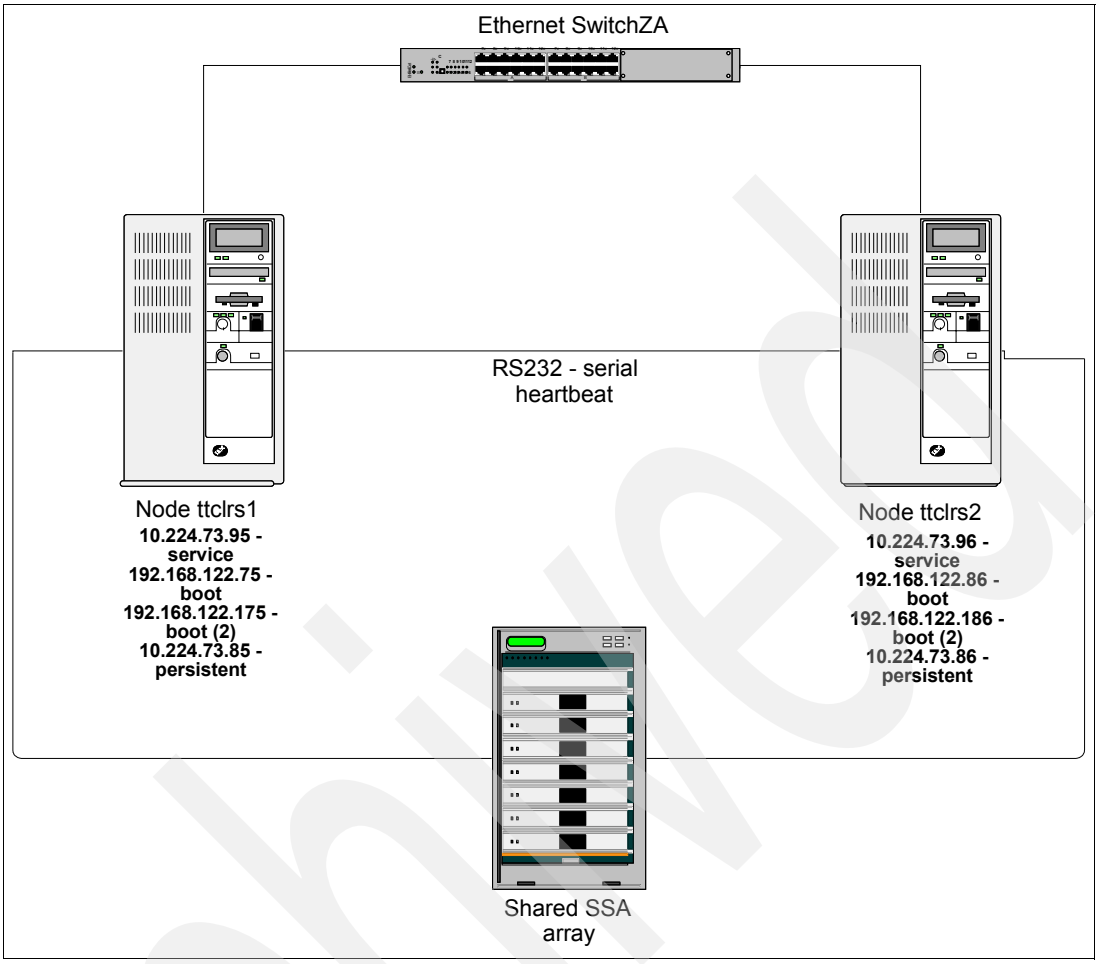

Figure 5-1 HACMP/ES v4.5 cluster topology

5.2.1 Network configuration

One of the new features of HACMP/ES v4.5 is that application resource IP addresses can be moved around through IP address aliasing. This feature is supported by AIX 5L Version 5.1 and later.

Traditional HACMP IP configurations

Prior to HACMP/ES v4.5, IP address takeover (IPAT) was managed by physically moving a service IP address from its service adapter to a standby adapter. In addition, to support the move, either the Media Access Control (MAC) address of the original service adapter had to move to the standby adapter, or the Address Resolution Protocol (ARP) cache had to be updated. This process is known as *Hardware Address Takeover* (HWAT).

Enhanced HACMP IP configurations

With the introduction of HACMP/ES v4.5 on AIX 5L Version 5.1, there is an option to support IPAT through *IP aliasing*. IP aliasing is a mechanism whereby a single network adapter can support multiple IP addresses simultaneously.

There are restrictions that limit the use of IPAT through IP aliasing. The biggest one is that IP aliasing cannot support HWAT. If an application resource has a dependency on adapter MAC

addresses and requires HWAT, the traditional IPAT mechanism described earlier is still available. The two methods of IPAT cannot be mixed in a single cluster.

In our cluster configuration, we defined the following network addresses:

- Two boot addresses for each node
- Service IP address for each application server
- Persistent IP address for each node
- RS232 serial link between nodes used for heartbeat

Figure 5-2 illustrates these parameters.

5.2.2 Setting up the cluster

We set up the clustered servers by using the following steps:

1. Prepare AIX 5L Version 5.1 for clustering. Install the necessary AIX fix packs.
2. Set up HACMP/ES.
   a. Install the HACMP/ES base software and set it up on both nodes.
   b. Download the latest clustering Fix Pack 4.5.0.10 and apply it prior to the initial cluster setup.
   c. Complete the HACMP/ES planning worksheets.

      Refer to Appendix A, “HACMP/ES planning worksheets” on page 55, to see the worksheets used for this test configuration.
   d. Set up the cluster name, cluster nodes, and network interfaces to use for the EnterpriseOne cluster. Replicate changes to the other node by using topology and cluster synchronization.
3. Install the JD Edwards EnterpriseOne database.
   a. Install the database from the distribution media.
   b. Apply the latest database fix packs that are supported by Oracle.
4. Install JD Edwards EnterpriseOne 8.9.
   a. Install the EnterpriseOne 8.9 software on the deployment and enterprise servers.
   b. Configure the EnterpriseOne software to work with DB2 UDB databases.
   c. Create shared databases using supplied JD Edwards database creation scripts.
5. Set up HACMP/ES for the EnterpriseOne 8.9 software.
   a. Include the application servers and resource groups for DB2 UDB and EnterpriseOne 8.9 software. See Figure 5-2.
   b. Add the application scripts to start and shut down application servers in the event of failure or resource group migration. We defined the resource groups as cascading without fallback. See Chapter 4, “Application script examples for EnterpriseOne 8.9 with HACMP/ES v5.1” on page 23.

To set up an HACMP cluster, follow the guidelines that are explained in *HACMP for AIX 4.5: Enhanced Scalability Installation and Administration Guide*, SC23-4306.

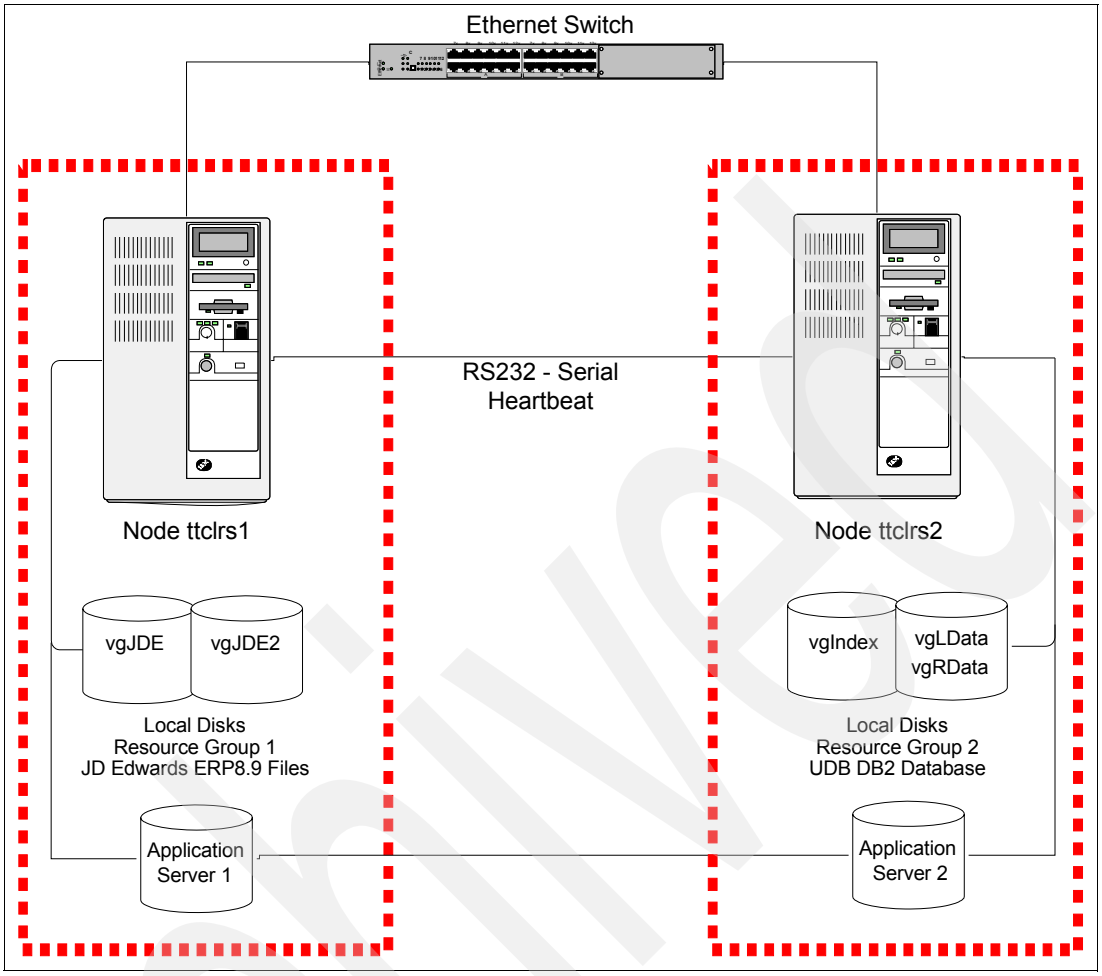

Figure 5-2 Logical layout on the HACMP/ES cluster

5.3 Installing and setting up the EnterpriseOne 8.9 deployment and enterprise servers

To set up the deployment and enterprise servers, follow these steps:

1. Install the necessary prerequisites for EnterpriseOne 8.9. See the Minimum Technical Requirements (MTR) topic on the Oracle customer resource Web site for EnterpriseOne 8.9. You can find this information on the Oracle PeopleSoft Customer Connection site, which you can access on the Web at:

   `http://www.peoplesoft.com/corp/en/public_index.jsp`

   When you reach this site, complete these tasks:

   a. Click the **LOG IN** button for customers and partners in the upper right corner.

   b. Sign in with your user name and password.

   c. Select Implement, **Optimize + Upgrade →Optimization Guide →Supported Platforms →JD Edwards EnterpriseOne**.

2. Plan the installation of the EnterpriseOne 8.9 software on the AIX cluster.

   In our lab, we decided to create two application servers using two different resource groups. Our first application server runs EnterpriseOne services, and the other hosts EnterpriseOne databases.

3. Install the EnterpriseOne 8.9 software on the deployment server.
4. Install the latest Planner Electronic Software Update (ESU) for EnterpriseOne 8.9.
5. Log in to EnterpriseOne 8.9.
6. In the Fast Path, type P98402 and press Enter.
7. The Database Creation Scripts Director form (Figure 5-3) opens. For Select the System Type, select **UNIX**. Then for Select the Database Type, specify your database type. For our testing, we used the DB2 UDB database. Since only UDB/DB2 is available on the form, we select **UDB/DB2**. Click **Next**.

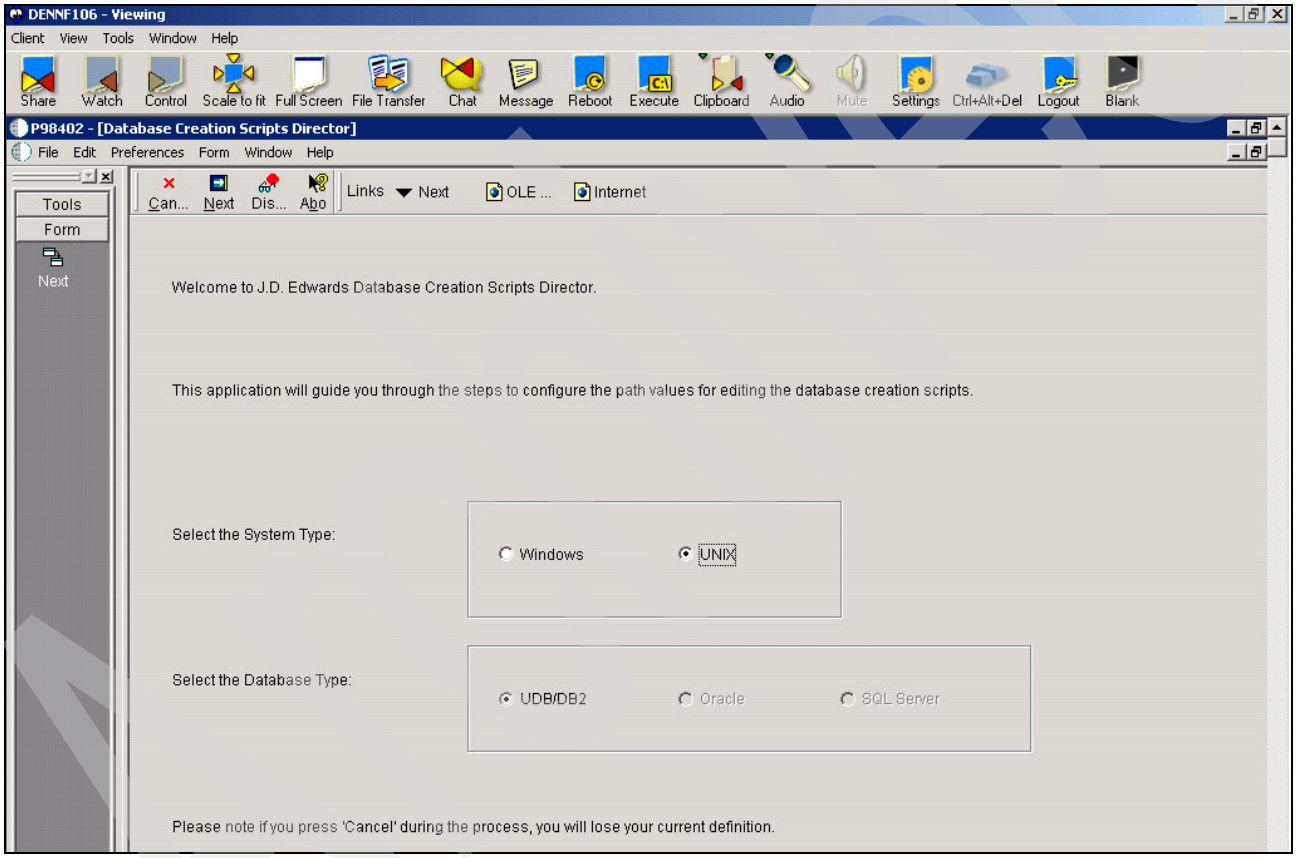

Figure 5-3 Database Creation Scripts Director (Part 1 of 2)

8. The form shown in Figure 5-4 opens. Make sure that the INSTANCE name, SERVICE port, and NODENAME are correct. Verify your NODENAME and SERVICE by entering the following command:

   ```
   db2 list node directory
   ```

   See Figure 5-2 on page 33 for information about the database file locations.

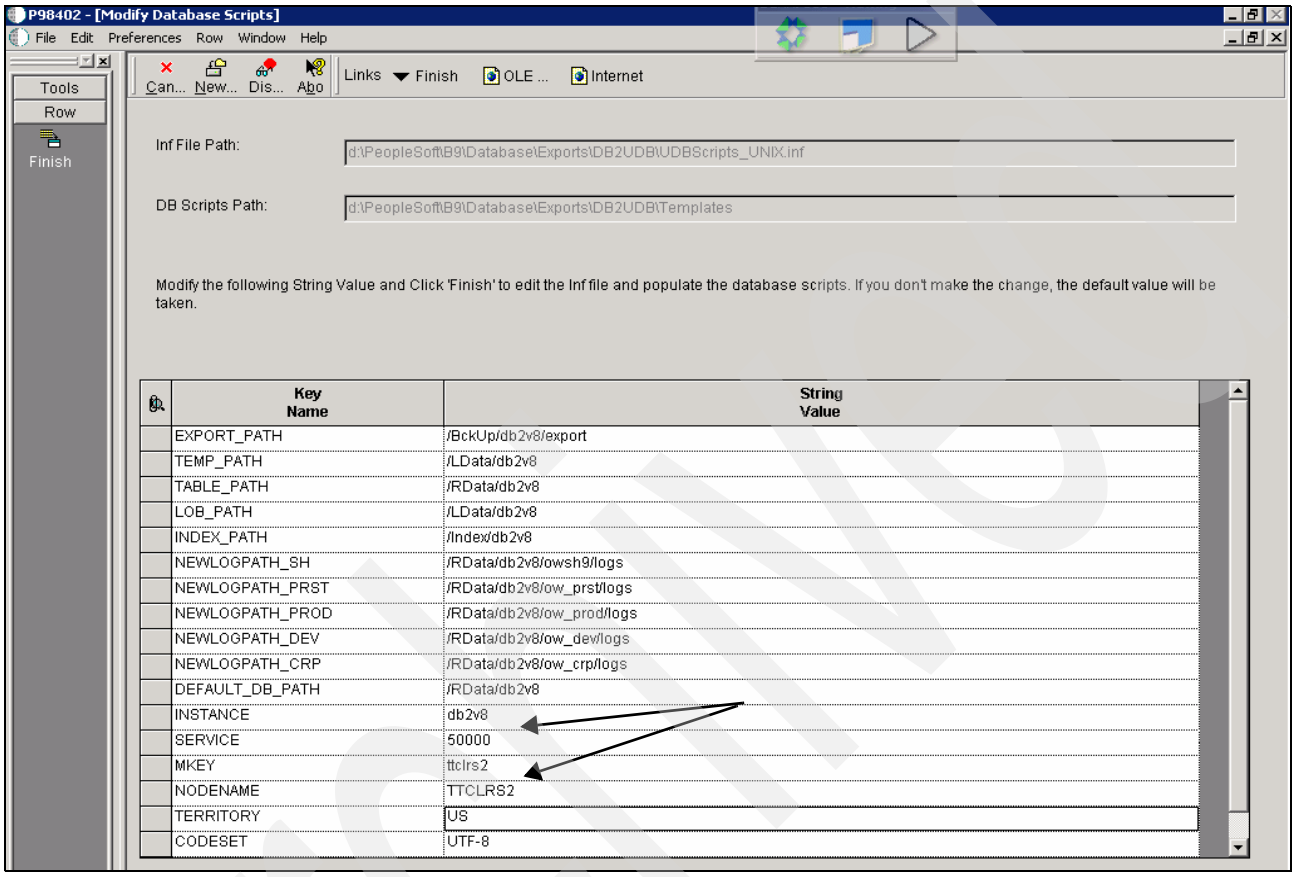

Figure 5-4 Database Creation Scripts Director (Part 2 of 2)

9. Changes in the form updates necessary database creation scripts for your environment. Click **Finish** when the form is complete.
10. Create the necessary database structures using the database scripts modified in step 9.

11. Create an installation plan with cluster service names for the enterprise server and database servers. See Figure 5-5, Figure 5-6, and Figure 5-7 on page 38.

    a. During the enterprise server definition part of the installation plan, be sure to select **Form →Advanced**.

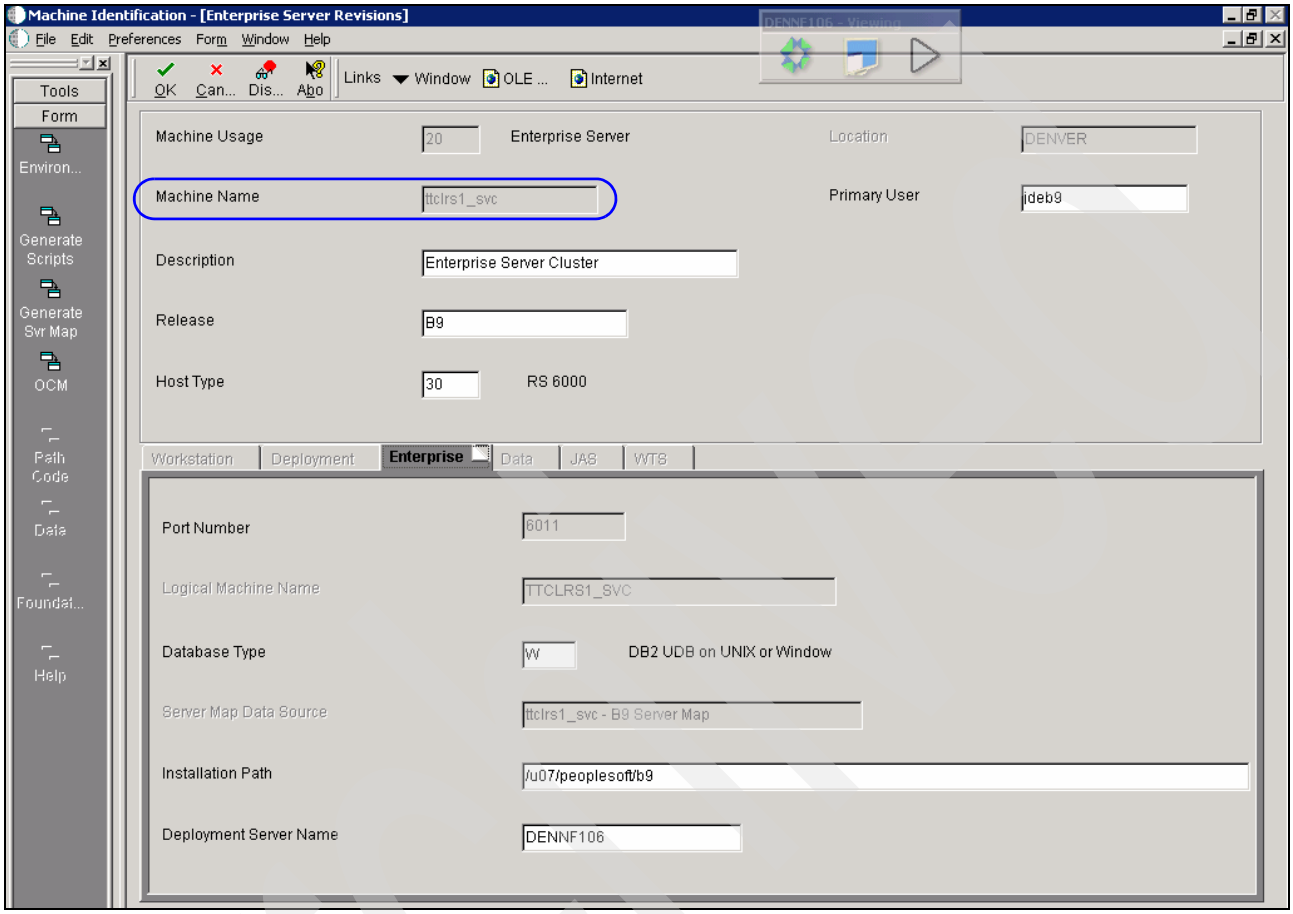

Figure 5-5 Clustered enterprise server

b. For the logical server SVR data source, make sure that **Clustered** is selected (Figure 5-6). This option disables failover to a LOCAL logic data source and continues retrying until clustered resources become available.

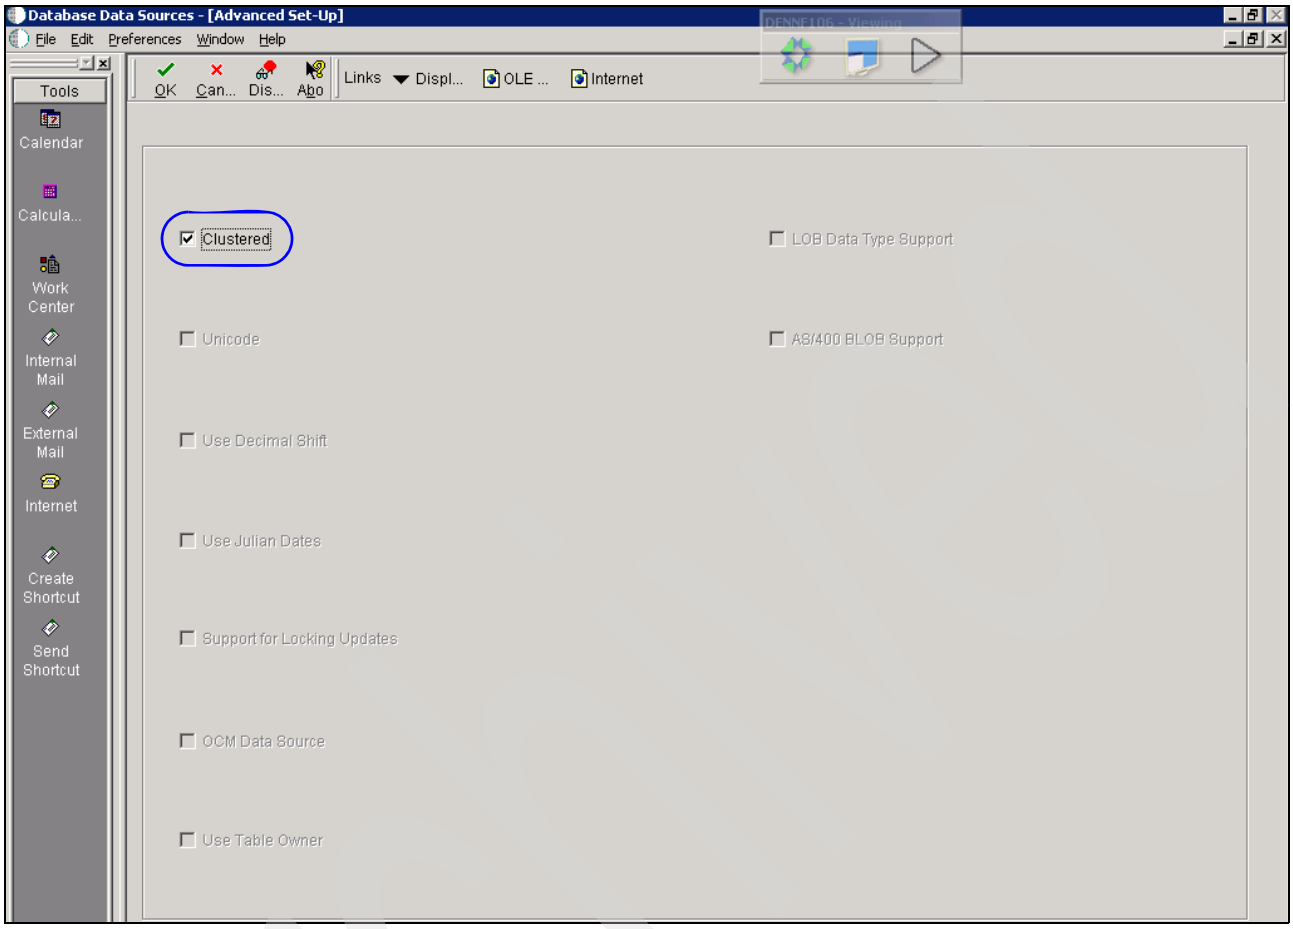

Figure 5-6 Clustered enterprise server: Advanced setup

c. Add an additional database server using the clustered service name of application server that hosts databases. See Figure 5-7.

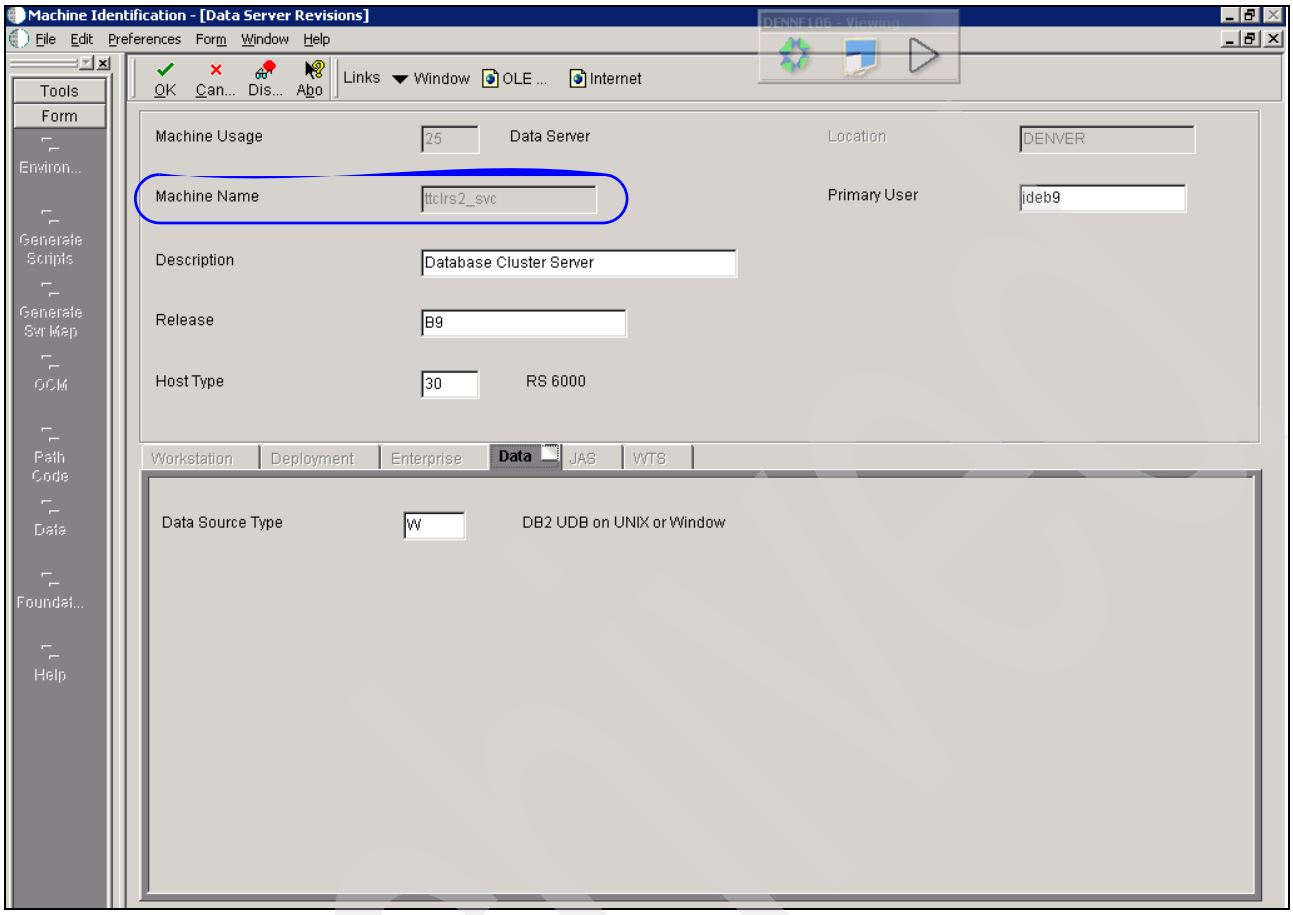

Figure 5-7 Clustered database server

d. During the creation of database data sources (for example, System, Data Dictionary, and Object Librarian), ensure that the Database Server Name is the same one that you entered in step c. See Figure 5-8.

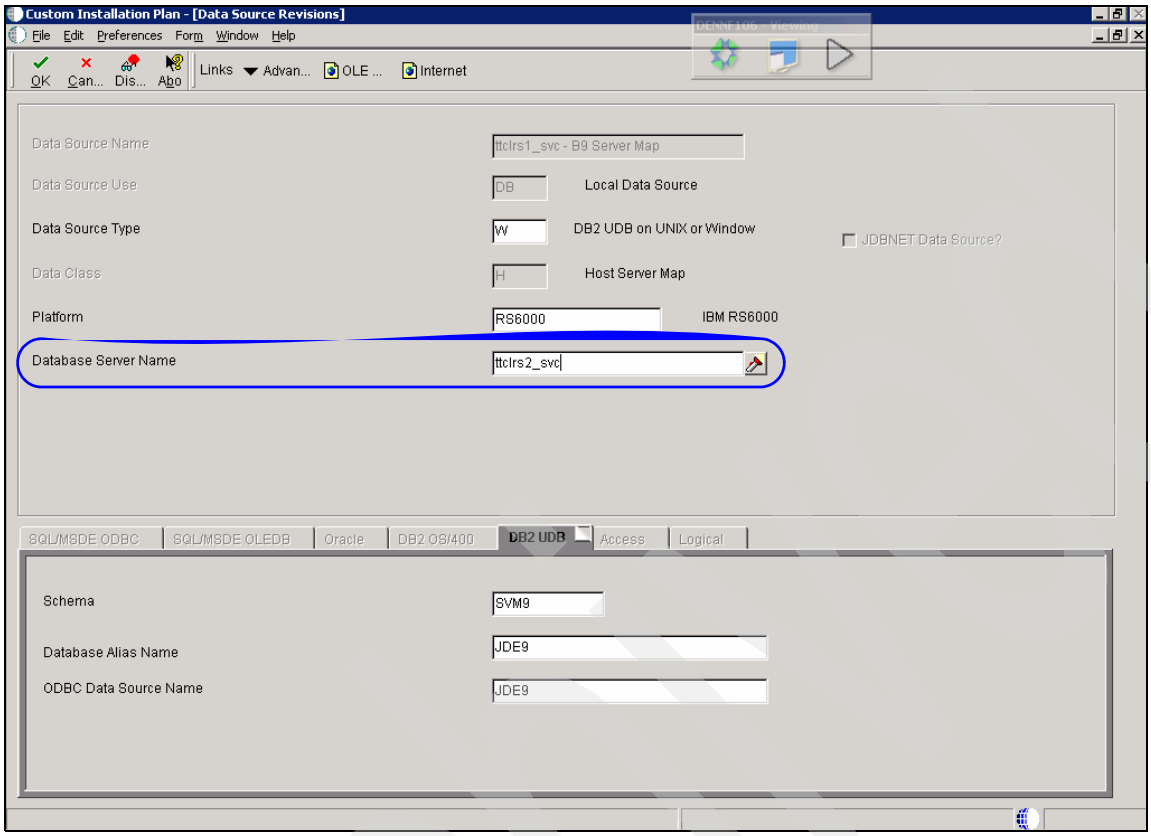

Figure 5-8 Server map data source example

12. Continue with the normal EnterpriseOne 8.9 installation steps to install the software on the enterprise server.
13. Modify the JDE.INI file on the enterprise server with the correct cluster service name or names.

    In our example, Example 5-1 shows the DB settings. The clustered database service name is ttclrs2_svc.

Example 5-1 DB system settings

```
[DB SYSTEM SETTINGS]
Default Role=*ALL
Default Env=DV9
Default PathCode=DV9
Base Datasource=ttclrs1_svc - B9 Server Map
Object Owner=SVM9
Server=ttclrs2_svc
….
```

Example 5-2 shows the security settings for our example. The clustered enterprise service name is ttclrs1_svc.

Example 5-2 Security settings

```
[SECURITY]
Default Role=*ALL
DefaultEnvironment=DV9
DataSource=System - B9
SecurityServer=ttclrs1_svc (Clustered Enterprise service name)
```

Example 5-3 shows the cluster settings for our example.

Example 5-3 Cluster settings

```
[CLUSTER]
PrimaryNode=ttclrs1_svc
```

14. After the cluster services are running on both nodes and the databases are up and running, run PORTTEST. If you run PORTTEST before the cluster services are up, it fails to find the correct host name and IP address.

You can find further information to help you install and set up EnterpriseOne software in the *EnterpriseOne Release 8.9 UNIX Installation Guide* provided by Oracle.

6

Application script examples for EnterpriseOne 8.9 with HACMP/ES

High Availability Cluster Multi-Processing (HACMP) v4.5 and v5.1 require that the administrator defines a series or script for starting, stopping, and monitoring applications. In our example, we used the scripts provided by DB2 to start and stop the database. This chapter explains the scripts used to start, stop, and monitor JD Edwards EnterpriseOne 8.9.

6.1 Application startup scripts

Oracle provides the RunOneWorld.sh and EndOneWorld.sh scripts for the EnterpriseOne administrator to start and stop the EnterpriseOne 8.9 server. HACMP executes the configured scripts in the root directory. Consequently, an intermediate script is needed to switch user execution to the EnterpriseOne administrator user in order to execute the RunOneWorld.sh and EndOneWorld.sh scripts.

In our testing, we found that, in some failovers, HACMP tried to start EnterpriseOne before the disks where ready. A timing loop was added to the script to prevent this.

Unlike HACMP application monitor scripts, parameters can be passed freely into application startup and shutdown scripts. Three parameters are passed in:

- Maximum Number of Tries
- Sleep Time
- JDE User

The local_start_ow script is responsible for starting EnterpriseOne services. This script is saved in the /usr/local/bin directory on both nodes. Example 6-1 shows the local_start_ow script.

Example 6-1 The local_start_ow script

```
#! /usr/bin/ksh
(( TRUE = 0 ))
(( FALSE = 1 ))
(( COUNT = 0 ))
OW_STARTED=${FALSE}
if [[ $# -ne 3 ]]
then
echo "Usage: $0 <Maximum Number of Tries> <Sleep time> <JDE User> "
  exit 1
fi
MAX_TRIES=$1
SLEEP_4=$2
JDE_USER=$3
echo "$0 is Running at `date`"
while [ ${OW_STARTED} -ne ${TRUE} ]
do
if [[ ${COUNT} -ge ${MAX_TRIES} ]]
  then
    echo "\nTimed out waiting for /home/${JDE_USER} to be mounted"
    break
else
    if [[ -f /home/${JDE_USER}/.profile ]]
    then
. /home/${JDE_USER}/.profile
      if [[ -d ${EVRHOME}/log ]]
      then
        LOGDATE=`date +"%m%d%y%H%M%S"`
        mkdir -m 777 ${EVRHOME}/log/${LOGDATE}
        chown ${JDE_USER} ${EVRHOME}/log/${LOGDATE}
        mv ${EVRHOME}/log/jde*log ${EVRHOME}/log/${LOGDATE}
        echo "Starting EnterpriseOne as ${JDE_USER}"
        /usr/bin/su - ${JDE_USER} "-c ${SYSTEM}/bin32/RunEnterpriseOne.sh"
        OW_STARTED=${TRUE}
echo "\nScript RunEnterpriseOne.sh ran.  Check logs for results"
      else
        echo "EVRHOME not found at ${EVRHOME}"
```

```
        (( COUNT = ${COUNT} + 1 ))
        sleep ${SLEEP_4}
      fi
    else
       (( COUNT = ${COUNT} + 1 ))
       sleep ${SLEEP_4}
       echo "$0 Could not start EnterpriseOne. Tried $COUNT times"
    fi
  fi
done
exit ${OW_STARTED}
```

6.2 Application shutdown scripts

The EnterpriseOne 8.9 shutdown script needs only the EnterpriseOne 8.9 administrator User ID to be passed. The local_stop_ow script is saved in the /usr/local/bin directory on both nodes. Example 6-2 shows the local_stop_ow script.

Example 6-2 The local_stop_ow script

```
#! /usr/bin/ksh
(( TRUE = 0 ))
(( FALSE = 1 ))
(( COUNT = 0 ))
OW_STOP=${FALSE}
JDE_USER=$1


if [[ $# -ne 1 ]]
then
  echo "Usage: $0 <User>"
  exit 1
fi

echo "$0 is Running at `date`"

while [ ${OW_STOP} -ne ${TRUE} ]
do

    if [[ -f /home/${JDE_USER}/.profile ]]
    then
      . /home/${JDE_USER}/.profile
      echo "Ending OneWorld as ${JDE_USER}"
      /usr/bin/su - ${JDE_USER} "-c ${SYSTEM}/bin32/EndOneWorld.sh"
      OW_STOP=${TRUE}
      echo "\nScript EndOneWorld.sh ran.  Check logs for results"
    else
       echo "$0 Could not stop OneWorld."
    fi
done

exit
${OW_STOP
```

6.3 Application monitor script for application server asJDE

Example 6-3 shows the ow_app_monitor script.

Example 6-3 The ow_app_monitor script

```
#! /usr/bin/ksh
count=`ps -ef | grep jdenet_n | grep -v grep |wc -l`
if [[ "$count" -gt 0 ]]
then
  return_code=0

else
  return_code=1
fi
exit
${return_code}
```

6.4 Application monitor settings for application server asJDE

HACMP provides a simple single process monitor setting. This option does not work with EnterpriseOne 8.9. The main process for EnterpriseOne 8.9, jdenet_n, forks additional processes to handle additional load. HACMP detects these additional processes and triggers the defined response. Therefore, a simple script is created to monitor whether one or more jdenet_n processes are present. If no jdenet_n processes are present, HACMP considers this a failure and triggers the defined responses.

Note: After this monitoring is configured, you should start and stop only EnterpriseOne 8.9 through HACMP.

The local_mon_ow script is saved in the /usr/local/bin directory on both nodes. Example 6-4 shows the local_mon_ow script.

Example 6-4 The local_mon_ow script

```
                Change/Show Custom Application Monitor
Type or select values in entry fields.
Press Enter after making all desired changes.
[Entry Fields]
  Application Server Name                  asJDE
* Monitor Method                          [/usr/local/bin/local_mon_ow>
  Monitor Interval                        [10]                        #
  Hung Monitor Signal                     [9]                         #
* Stabilization Interval                  [60]                        #
  Restart Count                           [3]                         #
  Restart Interval                        [231]                       #
* Action on Application Failur            [fallover]                 +
  Notify Method                           []
  Cleanup Method              [/usr/local/bin/local_stop_ow 5 5 jdeb9>
  Restart Method              [/usr/local/bin/local_start_ow 5 5 jdeb9>
```

7

JD Edwards EnterpriseOne 8.9 HACMP/ES testing

Table 7-1 outlines the test methodologies used during this certification project. Our cluster failure testing is not exhaustive for all EnterpriseOne 8.9 applications.

Table 7-1 Test methodologies

Test	Steps	Outcome
Fail rgJDE (resource group)	1. Enter a new sales order. 2. While in the sales order header, initiate an rgJDE failure.	No data is written to the tables. Rerun the transaction.
Fail rgJDE	1. Enter a new sales order. 2. While in the sales order detail, initiate an rgJDE failure.	You must exit back to the menu. Sales Order header and details are lost. The next number is advanced past the failed sales order number. Redo the transaction.
Fail rgJDE	1. Submit a Universal Batch Engine (UBE). 2. Fail rgJDE while UBE is running.	Processing UBE ends in an *Error* status. Resubmit failed UBE.
Fail rgJDE	1. Submit several UBEs. 2. Fail rgJDE while some UBEs are running with others in *waiting* status.	Running UBEs end in *Error* status in most cases. This can take up to five minutes to occur while the UBE clean process runs. Some UBEs may remain in "P" status even though they are not processing. You must delete those records manually. The waiting UBEs process after the rgJDE becomes available.
Fail rgDB	1. Enter a new sales order. 2. While in the sales order header, initiate an rgDB failure.	There is no data loss. The transaction completed after rgDB became available.

Test	Steps	Outcome
Fail rgDB	1. Enter a new sales order. 2. While in the sales order detail, initiate an rgDB failure.	There is no data loss. Transaction completed normally, including writing data to header and detail tables.
Shut down the cluster node that owns rgJDE	Initiate the hardware failure by turning off one of the nodes.	The failed resource group migrates to the other node. This can take up to five minutes to occur while the UBE clean process runs. Application monitoring is required for this to occur automatically. Running UBEs ends in *Error* status. New users see the "`Waiting for clustered server`" message upon logon. After rgJDE becomes available, users can log in into EnterpriseOne.
Shut down the cluster node that owns rgDB	Initiate a hardware failure by turning off one of the nodes.	Failed resource group migrates to the other node. Application monitoring is required for this process to occur automatically. Interactive users will experience a short delay before rgDB becomes available.
Fail one of the network cards	Unplug the network cable from one of the network interfaces that is used in the cluster.	Other boot adapters automatically take over missing service IP address. There is no loss of data.
Fail an RS232 connection between nodes	Unplug the serial connection between two nodes (interrupting heartbeat).	Heartbeat is routed through other available network interfaces. There is no data loss.
Fail one network card and an RS232 connection	1. Unplug network cable from the network interface. 2. Unplug the serial connection.	The node with a surviving network connection takes over the isolated node's resource group. If the ttclrs1 node is isolated, the rgJDE resource group should start on the ttclrs2 node.

Note: The Universal Batch Engine is the system on EnterpriseOne to submit batch jobs, usually reports.

8

IBM TotalStorage FAStT in an HACMP for AIX environment

This chapter gives a brief overview of the IBM TotalStorage® Fibre Array Storage Technology (FAStT) family of products. It discusses some of the strengths of using FAStT storage in an IBM High Availability Cluster Multi-Processing (HACMP) for AIX in a pSeries environment.

The information in this chapter provides reference links to additional FAStT documentation. Therefore, it is not meant to replace or reproduce the existing FAStT installation and configuration documentation. Nor does it attempt to explain how to install or configure HACMP for AIX.

For additional information about HACMP, refer to:

- *HACMP for AIX 4.5: Concepts and Facilities*, SC23-4276
- *HACMP for AIX 4.5: Planning Guide*, SC23-4277
- *HACMP for AIX 4.5: Installation Guide*, SC23-4278
- *HACMP for AIX 4.5: Administration Guide*, SC23-4279

The hardware and software requirements are documented in the *Release Notes* in the /usr/lpp/cluster/doc/release_notes directory.

8.1 IBM TotalStorage FAStT product overview

As of the time of this writing, the IBM TotalStorage FAStT family of products includes the FAStT200, DS4300, DS4400, and DS4500 storage arrays, as illustrated in Figure 8-1.

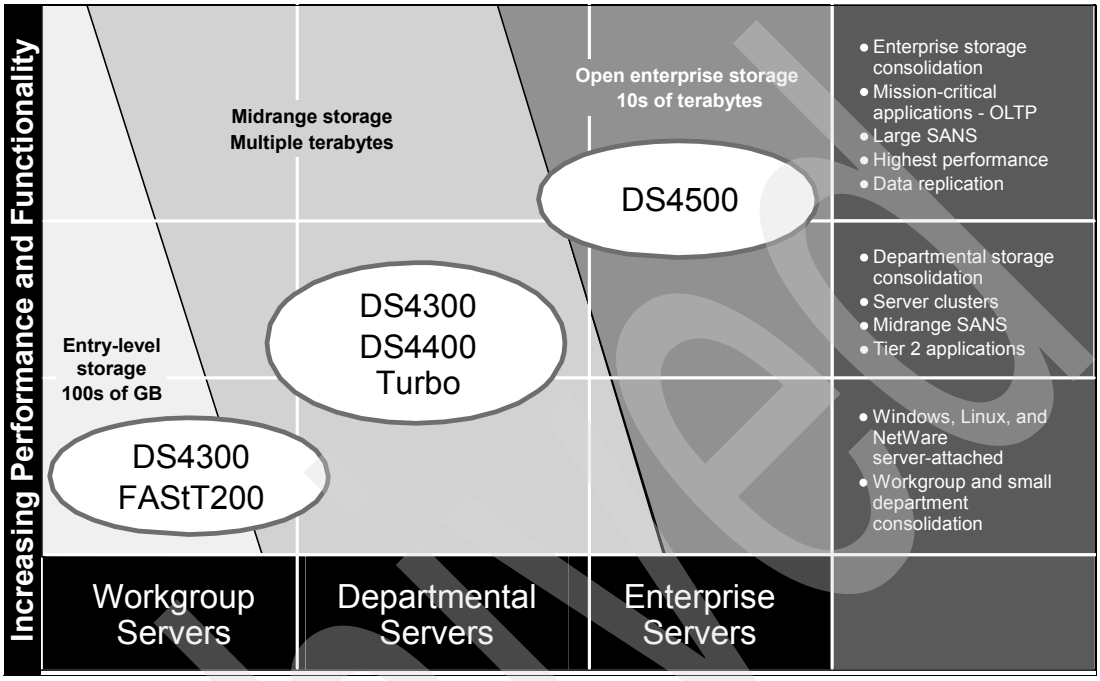

Figure 8-1 FAStT family

The FAStT products have various capabilities and capacities ranging from entry level to high end. However, for this discussion, the FAStT200, DS4300, DS4400, and DS4500 products are referred to as *FAStT*. The features that are discussed can vary by FAStT subsystem, operating system, and FAStT Storage Manager software version.

In open systems general purpose computing, performance and faster storage create more responsive applications. This improves employee productivity and, therefore, client satisfaction. FAStT storage arrays are designed to achieve these goals. In addition, FAStT can support large configurations and high workloads. FAStT is ready to and capable of expanding to meet tomorrow's storage and performance demands.

FAStT storage arrays are optimized for disk-based performance, which is required by open systems applications. In general, there are two constraining types of applications. The first is a transaction-based or input/output per second (IOPS) application. IOPS applications have distributed random-access data and rarely experience "cache hit." The second type is a throughput application. Throughput applications need to sustain MB/s to and from disk. FAStT storage arrays are dynamically tunable for changing conditions to meet the needs of the application better, regardless of the type.

FAStT storage arrays are designed to maximize storage usage, which lowers the total cost of ownership (TCO) and increases the return on investment (ROI). Such configurations are highly flexible to support various server and application requirements. They can be optimized for capacity, performance, or both. In addition, FAStT capacity and performance scaling can be done incrementally if needed.

FAStT storage allows the consolidation of storage to match application requirements with virtual partitioning, improves capacity usage and storage flexibility, and provides high-end

storage functionality for smaller servers. It offers FlashCopy® data replication, which allows capacity efficient, point-in-time, “snapshot” copies.

Figure 8-2 shows a possible FAStT client and server environment. This configuration can be managed by the FAStT Storage Manager software. This software lowers storage management costs and provides centralized administration of all FAStT storage servers.

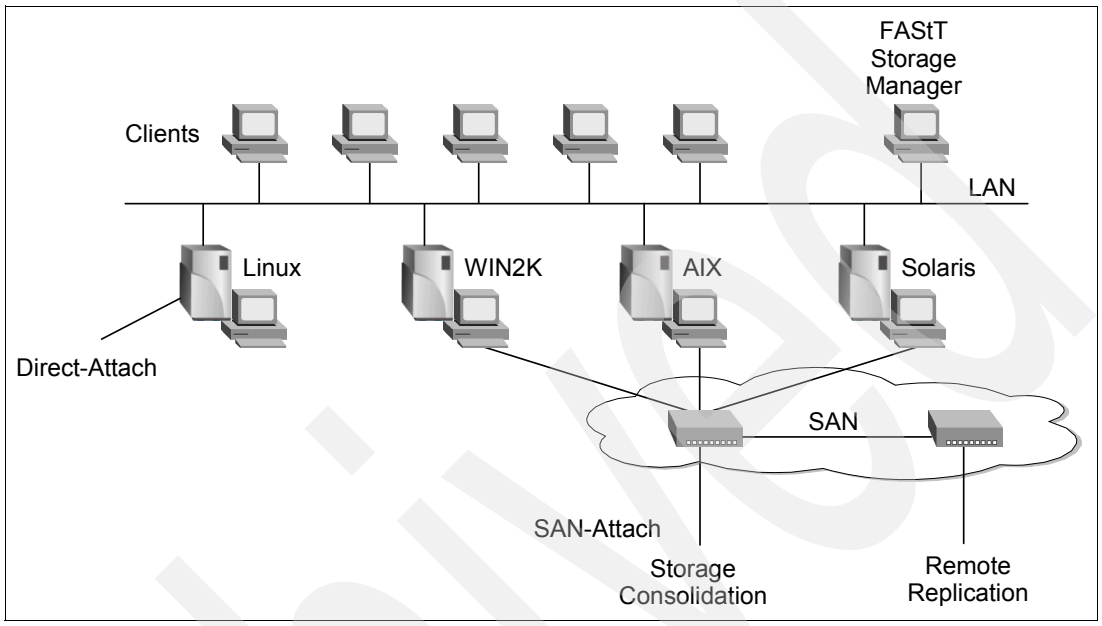

Figure 8-2 FAStT Storage Manager

FAStT Storage Manager provides a consistent interface across all platforms, which lowers training and support costs. In addition, it provides configuration flexibility with an online “anytime administration” capability to maximize performance, capacity, and availability.

As illustrated in Figure 8-3, FAStT storage arrays offer modular “Pay-As-You-Grow” scalability. This lowers acquisition costs while avoiding over-configuration without sacrificing future scalability. The cost of FAStT expansion is lowered, which enables optimal just-in-time purchasing.

FAStT has built-in expansion or EXP hot-add technology, which supports capacity expansion while system I/O continues. This capacity enables you to cable new expansion units, configure new drives, and access new logical drives without the expected downtime or necessity to reboot.

FAStT provides high data availability, which ensures valuable information is accessible. Data I/O continues through capacity expansion, configurations, reconfigurations, tuning, and maintenance. It decreases the risk of downtime-related revenue loss.

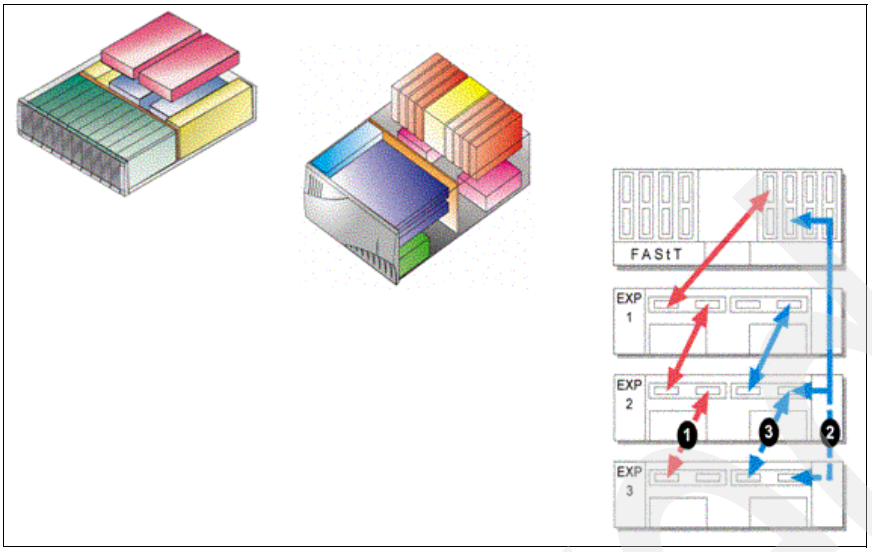

Figure 8-3 FAStT modular scalability

As illustrated in Figure 8-4, FAStT allows local and remote storage-based replication. Locally, FlashCopy and VolumeCopy provide fast restorations and application usage (for example, backups, application testing, data mining, and analysis). Remote mirroring permits offsite disaster recovery.

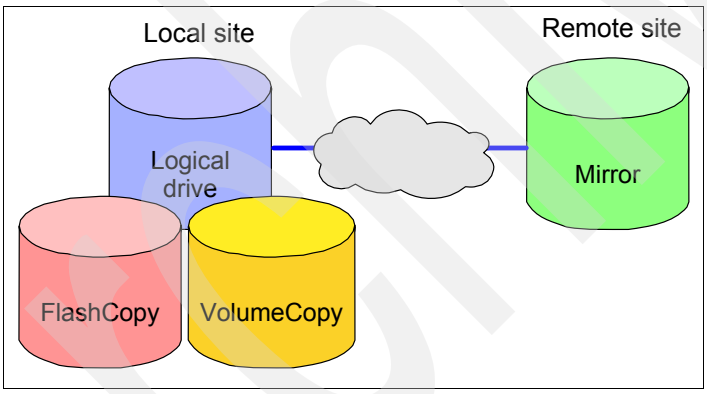

Figure 8-4 FAStT remote data replication

Choosing FAStT protects the storage investment of hardware, software, and data. The EXP700 expansion unit is supported in all IBM TotalStorage FAStT systems. This enables common spares and module migration within the FAStT family of products. FAStT offers “data intact” upgrades and migration of data. Common firmware ensures that data and configuration information are recognized by other FAStT controllers. Plus, the unique DACstore technology enables array-level migrations from one system to another.

FAStT storage arrays offer superior supportability. FAStT Service Alert provides monitoring and notification of hardware issues. The Recovery Guru analyzes problems and provides step-by-step resolution instructions. FAStT offers performance monitoring, automatic code synchronization, and background media scans, which detect media errors before they cause a read/write failure. It also provides a major event log.

In summary, the FAStT storage servers are a fully-featured family of storage systems that:

- Provide the flexibility to easily manage growth from entry level to the enterprise
- Offer unmatched performance value
- Incorporate powerful remote and centralized management tools to create efficiencies in storage systems
- Maximize the return on storage investment
- Ensure that the storage strategy supports the business strategy
- Offer improved efficiencies
- Lower the storage TCO
- Provide investment protection

The HACMP cluster with attachment to IBM TotalStorage FAStT shown in Figure 8-5 is taken from the IBM Redpaper *IBM TotalStorage: FAStT Best Practices Guide*, REDP-3690. It is included to show a basic two-node AIX HACMP cluster with only two FAStT host port connections and to illustrate the required Fibre Channel (FC) switch zoning to the FAStT controllers.

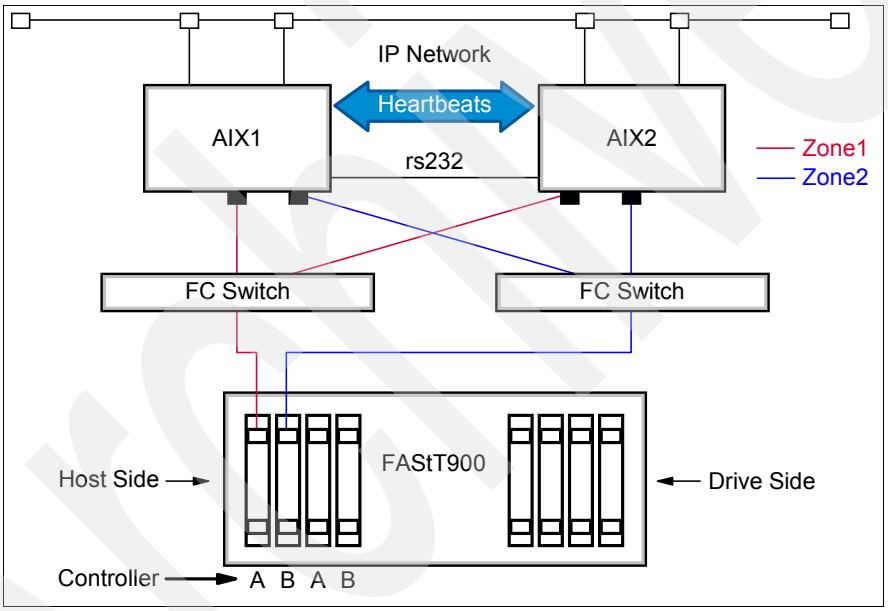

Figure 8-5 HACMP cluster with attachment to FAStT

In an HACMP environment, the cluster is designed to eliminate any single point of failure. In the case of our example, each AIX host has two FC adapters. Each AIX FC adapter connects to FAStT through a different FC switch. Each FC switch provides a path to one of the two FAStT storage array controllers. If any one of these components fails to function, there is still another path to the FAStT storage array data.

Additional FAStT host mini-hubs are not necessary. However, you can order them for added redundancy or specific security requirements as long as they are zoned as shown. The AIX-specific requirement is that each pSeries adapter should be zoned so that it can see one, and only one, FAStT controller. The HACMP cluster is not limited to only two AIX hosts, but is shown as such for this simple example.

Important: Before you install FAStT in an HACMP environment, always read the AIX readme file, the FAStT readme for the specific Storage Manager version and model, and the HACMP configuration information. You can download the latest FAStT documents from the following sites based on model number:

- FAStT200: http://www.ibm.com/servers/storage/support/disk/fastt200/
- FAStT500: http://www.ibm.com/servers/storage/support/disk/fastt500/
- DS4300: http://www.ibm.com/servers/storage/support/disk/ds4300/
- DS4400: http://www.ibm.com/servers/storage/support/disk/ds4400/
- DS4500: http://www.ibm.com/servers/storage/support/disk/ds4500/

8.2 HACMP configuration information for the IBM TotalStorage FAStT server

The supported hardware, software, and operating systems generally change as the products progress in versions and models. At the time of this writing, the combinations of HACMP, FAStT, AIX, and FAStT Storage Manager software listed in Table 8-1 are supported together.

Note: FAStT Storage Manager 8.4 is not currently supported with HACMP.

Table 8-1 HACMP software configuration

HACMP software version	AIX version	FAStT 200	FAStT 500	DS4300	DS4400	DS4500	Storage Manager version
HACMP v4.4.1 ES and ESCRM	v4.3.3, v5.1.0	X	X	X	X	X	v8.21, v8.3
HACMP v4.5 ES and ESCRM	v5.1.0, v5.2.0	X	X	X	X	X	v8.21, v8.3
HACMP v5.1	v5.1.0, v5.2.0	X	X	X	X	X	v8.3

You can download the complete and most recently updated *HACMP configuration information for FAStT on pSeries/AIX Servers* document from:

ftp://ftp.software.ibm.com/storage/fastt/fastt500/HACMP_config_info.pdf

This includes the latest supported configurations of AIX, HACMP, and FAStT; recommended versions of HACMP software; FAStT software and firmware requirements; AIX and HACMP configuration limitations; and other FAStT information that is specific to an HACMP environment.

For the most current information about AIX and device driver versions for FAStT for Storage Manager v8.21, see the *aixreadme* file on the Web at:

http://www.ibm.com/pc/support/site.wss/MIGR-43839.html

For the most current information about AIX and device driver versions for FAStT for Storage Manager v8.30, see the *aixreadme* file on the Web at:

http://www.ibm.com/pc/support/site.wss/MIGR-50177.html

8.2.1 Configuration limitations

Current restrictions and guidelines are presented in the *HACMP configuration information for FAStT on pSeries/AIX Servers* document referenced in 8.2, “HACMP configuration information for the IBM TotalStorage FAStT server” on page 52. You should know and understand these limitations and guidelines when you install the FAStT storage array in an HACMP environment. The guidelines are provided to help you avoid unpredictable results when you implement FAStT with HACMP. There are pSeries and AIX limitations, as well as HACMP-specific limitations.

The general limitations and restrictions for HACMP include:

- Only switched fabric connections allowed between host node and FAStT. No direct-attach connections are allowed between the host node and FAStT.
- HACMP Cluster Single Point of Control utility (C-SPOC) cannot be used to add a FAStT disk to AIX through the “Add a Disk to the Cluster” facility.
- Single Node Quorum is not supported in a two-node General Parallel File System (GPFS) cluster with FAStT disks in the configuration.
- HACMP v5.1 is not supported in heterogeneous server environments at this time.

Applicable pSeries and AIX limitations (not HACMP) include:

- Each AIX host attaches to FAStT storage servers using *pairs* of FC adapters host bus adapters (HBA). For each adapter pair, configure one HBA to connect to controller “A,” and the other to connect to controller “B.” Configure each HBA pair to connect to a single partition in a FAStT storage server or multiple FAStT storage servers (fan-out).
- To attach an AIX host to single or multiple FAStT storage servers with two partitions, use two HBA pairs.
- Single HBA configurations are allowed, but each single HBA configuration requires that you connect both controllers in the FAStT to a switch within the same storage area network (SAN) zone as the HBA.

  Note: Although single HBA configurations are supported, we recommend that you do not use them for HACMP environments because it introduces a single point-of-failure in the storage I/O path.

- The maximum number of FAStT partitions (host groups) per AIX host per FAStT storage subsystem is two.
- In regard to SAN switch zoning rules, configure each HBA within a host in a separate zone from other HBAs within that same host, when connected to the same FAStT controller port. You can configure only one HBA within a host in the same zone with a given FAStT controller port.
- Hosts within a cluster can share zones with each other.
- For highest availability, distributing the HBA and FAStT connections across separate FC switches minimizes the effects of a SAN fabric failure.

Here are other usage notes specific to HACMP environments:

- Concurrent and non-concurrent modes are supported with HACMP and FAStT running Storage Manager v8.21 or v8.3, including hot standby and mutual take-over.
- HACMP is supported on pSeries 615 Entry, 630 Entry (earlier), 650 Midrange, 670 Midrange, 670 Midrange, and 655 Technical LPAR clustered configurations.

- HACMP clusters can support two to 32 servers per FAStT partition. In this environment, be sure to read and understand the AIX device drivers queue depth settings as documented in *IBM TotalStorage FastT Storage Manager V8.30 Installation and Support Guide*, GC26-7519.
- You can connect non-clustered AIX hosts to the same FAStT that is attached to an HACMP cluster, but must be configured on separate FAStT host partitions.
- HACMP v4.4.1 and v4.5.0 are supported in heterogeneous server environment with the following operating systems:
  - Novell NetWare 6.0
  - Sun™ Solaris™ 2.6, 7, and 8
  - Sun Solaris 8 with Veritas Cluster Server 3.5
  - Sun Solaris 2.6, 7, and 8 with Veritas Volume Manager 3.5
  - HP-UX 11.0 and 11i
  - HP-UX 11.0 and 11i with HP-UX MC/Serviceguard Clustering
  - Linux Red Hat Advanced Server 2.1
  - SUSE LINUX Enterprise Server 8
  - Microsoft Windows 2000 and NT 4.0
  - Microsoft Windows 2000 and NT 4.0 with Microsoft Cluster Services

For more information about a particular operating system environment, refer to the specific *Installation and Support Guide* associated with the software.

For information about various platform, operating system, HBA, and applications support with FAStT, see the FAStT Storage interoperability matrix at:

http://www.storage.ibm.com/disk/fastt/supserver.htm

A

HACMP/ES planning worksheets

This appendix identifies the name of worksheets recommended to plan for a High Availability Cluster Multi-Processing/Enhanced Scalability (HACMP/ES) v4.5 or HACMP/ES v4.5 implementation. It includes the worksheets that we completed for the certification and testing performed with the JD Edwards EnterpriseOne 8.9 application.

Note: The worksheets apply to both v5.1 and v4.5. However, v5.1 has more options.

HACMP/ES planning worksheets

Table A-1 contains a complete list of the HACMP/ES planning worksheets that are available to plan for an HACMP/ES v4.5 implementation.

You can find instructions for completing the Web-based worksheets in *HACMP for AIX 4.5: Planning Guide*, SC23-4277, which is provided with the HACMP software. You can find blank worksheets in Appendix A and the online version of the worksheets in Appendix B of the same *Planning Guide*.

Table A-1 HACMP/ES planning worksheets

Worksheet	Purpose
TCP/IP Networks	Use this worksheet to record the TCP/IP network topology for a cluster. Complete one worksheet per cluster.
TCP/IP Network Adapters	Use this worksheet to record the TCP/IP network adapters connected to each node. You need a separate worksheet for each node defined in the cluster, so begin by photocopying a worksheet for each node and filling in a node name on each worksheet.
Serial Networks	Use this worksheet to record the serial network topology for a cluster. Complete one worksheet per cluster.
Serial Network Adapters	Use this worksheet to record the serial network adapters connected to each node. You need a separate worksheet for each node defined in the cluster. To begin, photocopy a worksheet for each node and fill in the node name on each worksheet.
Shared SCSI-2 Differential or Differential Fast/ Wide Disks	Use this worksheet to record the shared SCSI-2 Differential or Differential Fast/Wide disk configuration for the cluster. Complete a separate worksheet for each shared bus.
Shared IBM SCSI Disk Arrays	Use this worksheet to record the shared IBM SCSI disk array configurations for the cluster. Complete a separate worksheet for each shared SCSI bus.
Shared IBM SCSI Tape Drives	Use this worksheet to record the shared IBM SCSI tape drive configurations for the cluster. Complete a separate worksheet for each shared tape drive.
Shared IBM Fibre Tape Drives	Use this worksheet to record the shared IBM fibre tape drive configurations for the cluster. Complete a separate worksheet for each shared tape drive.
Shared IBM Serial Storage Architecture (SSA) Disk Subsystems	Use this worksheet to record the IBM 7131-405 or 7133 SSA shared disk configuration for the cluster.
Non-Shared Volume Groups (Non-Concurrent Access)	Use this worksheet to record the volume groups and file systems that reside on a node's internal disks in a non-concurrent access configuration. You need a separate worksheet for each volume group. To begin, photocopy a worksheet for each volume group and fill in a node name on each worksheet.

Worksheet	Purpose
Shared Volume Groups/Filesystems (Non-Concurrent Access)	Use this worksheet to record the shared volume groups and file systems in a non-concurrent access configuration. You need a separate worksheet for each shared volume group. To begin, photocopy a worksheet for each volume group and fill in the names of the nodes sharing the volume group on each worksheet.
NFS-Exported Filesystems/Directories	Use this worksheet to record the file systems and directories NFS-exported by a node in a non-concurrent access configuration. You need a separate worksheet for each node defined in the cluster. To begin, photocopy a worksheet for each node and fill in a node name on each worksheet.
Non-Shared Volume Groups (Concurrent Access)	Use this worksheet to record the volume groups and file systems that reside on a node's internal disks in a concurrent access configuration. You need a separate worksheet for each volume group. To begin, photocopy a worksheet for each volume group and fill in a node name on each worksheet.
Shared Volume Groups (Concurrent Access)	Use this worksheet to record the shared volume groups and filesystems in a concurrent access configuration. You need a separate worksheet for each shared volume group. To begin, photocopy a worksheet for each volume group and fill in the names of the nodes sharing the volume group on each worksheet.
Applications	Use these worksheets to record information about applications in the cluster.
Fast Connect	Use this worksheet to record Fast Connect resources.
AIX Connections	Use this worksheet to record AIX connections realm/service pairs.
Communication Links (SNA-Over-LAN)	Use this worksheet to record information about SNA-over-LAN communications links in the cluster.
Communication Links (X.25)	Use this worksheet to record information about X.25 communication links in the cluster.
Communication Links (SNA-Over-X.25)	Use this worksheet to record information about SNA-over-X.25 communication links in the cluster.
Application Servers	Use this worksheet to record information about application servers in the cluster.
Application Monitors (Process)	Use this worksheet to record information for configuring a process monitor for an application.
Application Monitors (Custom)	Use this worksheet to record information for configuring a custom (user-defined) monitor method for an application.
Resource Groups	Use this worksheet to record the resource groups for a cluster.
Cluster Events	Use this worksheet to record the planned customization for an HACMP/ES cluster event.

HACMP/ES planning worksheets: Completed

The worksheets in this appendix reflect the implementation of the HACMP/ES v4.5 cluster, on which the EnterpriseOne 8.9 code was installed, configured, and certified. The worksheets were completed according to the guidelines in *Configuring Highly Available Clusters Using HACMP 4.5*, SG24-6845.

TCP/IP network

This TCP/IP Networks worksheet records the TCP/IP network topology for the cluster named ttclrsHA.

Cluster ID	1
Cluster name	ttclrsHA

Network name	Network type	Network attribute	Netmask	Node names	Use IP aliasing
IPnet	Ethernet	Public	255.255.255.0	ttclrs1, ttclrs2	Yes
rsNet	rs232	Serial		ttclrs1, ttclrs2	No

TCP/IP network adapters

These TCP/IP Network Adapters worksheets record the TCP/IP network adapters connected to the ttclrs1 and ttclrs2 nodes.

Note: The TCP/IP Network Adapters worksheet defines a boot or standby interface for a network on the ttclrs1 and ttclrs2 nodes. If you configure a network adapter for a network that uses IP address takeover (IPAT) through IP aliasing, a node must not have standby interfaces defined.

Node name ttclrs1

Interface IP label	Network type	Network name	Network attribute	Interface function	IP address	Netmask
ttclrs1_svc	Ethernet	IPnet	Public	Service	10.224.73.95	255.255.255.128
ttclrs1_boot	Ethernet	IPnet	Public	Boot	192.168.122.75	255.255.255.128
ttclrs1_boot2	Ethernet	IPnet	Public	Boot	192.168.122.175	255.255.255.128
ttclrs1_prs	Ethernet	IPnet	Public	Persistent	10.224.73.85	255.255.255.128

Node name ttclrs2

Interface IP label	Network type	Network name	Network attribute	Interface function	IP address	Netmask
ttclrs2_svc	Ethernet	IPnet	Public	Service	10.224.73.96	255.255.255.128
ttclrs2_boot	Ethernet	IPnet	Public	Boot	192.168.122.86	255.255.255.128
ttclrs2_boot2	Ethernet	IPnet	Public	Boot	192.168.122.186	255.255.192.128
ttclrs2_prs	Ethernet	IPnet	Public	Persistent	10.224.73.86	255.255.255.128

Serial network

This Serial Network worksheet records the serial network topology for the ttclrsHA cluster.

Cluster ID 1

Cluster name ttclrsHA

Network name	Network type	Network attribute	Node names
rsNet	rs232	serial	ttclrs1,ttclrs2

Note: RS232 serial lines, target mode SCSI-2 buses, and target mode SSA serial links do not use the TCP/IP protocol and do not require a netmask or an IP address.

Miscellaneous data

Any extra information about devices used to extend serial links, for example modem number or extender information, is recorded in the Miscellaneous Data portion of the Serial Networks worksheet.

Serial network adapters

This Serial Network Adapters worksheet records the serial network adapters connected to the ttclrs1and ttclrs2 nodes.

Node name ttclrs1

Slot number	Interface name	Adapter label	Network name	Network attribute	Adapter function
1	/dev/tty1	ttclrs1tty1	rsNet	Serial	Service

Node name ttclrs2

Slot number	Interface name	Adapter label	Network name	Network attribute	Adapter function
1	/dev/tty1	ttclrs2tty1	rsNet	Serial	Service

Note: Serial networks do not carry TCP/IP traffic. Therefore, no boot addresses, adapter identifiers (IP addresses), or adapter hardware addresses are required to maintain keep-alives and control messages between nodes.

Shared IBM Serial Storage Architecture disk subsystems

This Shared IBM Serial Storage Architecture (SSA) Disk Subsystems worksheet records the IBM 7131-405 or 7133 SSA shared disk configuration for the ttclrsHA cluster.

Host and adapter information

	Node 1	Node 2	Node 3	Node 4	Node 5
Node name	ttclrs1	ttclrs2			
SSA adapter label	ssa0				
Slot number					
Dual-port number					

SSA logical disk drive

Logical device name _____

Node 1	Node 2	Node 3	Node 4	Node 5
hdisk9	hdisk2			
hdisk3	hdisk4			
	hdisk9			
	hdisk17			

Shared volume groups/file systems (non-concurrent access)

This Shared Volume Groups and Filesystems worksheet records the shared volume groups and file systems in the non-concurrent access configuration. The ttclrs1and ttclrs2 nodes share five volume groups: lvJDElog, lvJDElog2, vgLData, vgRData, and vgIndex.

	Node 1	Node 2	Node 3	Node 4	Node 5
Node name	ttclrs1	ttclrs2			
Shared volume group name	vgJDE vgJDE2	vgLData vgRData vgIndex			
Major number	50 52	68 69 70			
Log logical volume name	lvJDElog lvJDElog2	lvIndexlog lvLDatalog lvRDatalog			
Physical volumes	hdisk9	hdisk2			
	hdisk3	hdisk4			
		hdisk17			
		hdisk9			

Logical volume name	lvJDE
Number of copies of logical partition	1
On separate physical volumes?	No
Filesystem mount point	/u07
Size (in 512-byte blocks)	256pp

Logical volume name	lvJDE2
Number of copies of logical partition	1
On separate physical volumes?	No
Filesystem mount point	/u08
Size (in 512-byte blocks)	66pp

Logical volume name	lvIndex
Number of copies of logical partition	1
On separate physical volumes?	No
Filesystem mount point	/Index
Size (in 512-byte blocks)	66pp

Logical volume name	lvRData
Number of copies of logical partition	1
On separate physical volumes?	No
Filesystem mount point	/RData
Size (in 512-byte blocks)	66pp

Logical volume name	lvLData
Number of copies of logical partition	1
On separate physical volumes?	No
Filesystem mount point	/LData
Size (in 512-byte blocks)	132pp

Application: JDE

This Application worksheet records information about the JDE application that is used in the ttclrsHA cluster.

Application name	JDE
Cluster name	ttclrsHA
Node relationship (cascading/concurrent/rotating) type of node priority policy used	Cascading

Key application files

	Folder/path	**File system**	**Location**	**Sharing**
Executables	/u07/PeopleSoft/b9	/u07	External	No
Configuration files	$JDE_BASE/JDE.INI			
Data files/devices				
Log files/devices				

Application failover strategy (P=Primary; T=Takeover)

	Node 1	**Node 2**	**Node 3**	**Node 4**	**Node 5**
Node name	ttclrs1	ttclrs2			
Strategy	Primary	Takeover			

Start application: Commands/procedures

/u07/PeopleSoft/b9/system/bin32/RunEnterpriseOne.sh

Start application: Verification commands/procedures

netwm

porttest JDE JDE DV7334

Start application: Node reintegration/takeover caveats

Node	Node name	Caveat
1	ttclrs1	N/A
2	ttclrs2	N/A
3		
4		
5		
6		

Stop application: Commands/procedures

/u07/PeopleSoft/b9/system/bin32/EndEnterpriseOne.sh

Stop application: Verification commands/procedures

netwm

Stop application: Node reintegration/takeover caveats

Node	Node Name	Caveat
1	ttclrs1	No failback
2	ttclrs2	N/A
3		
4		
5		
6		

Application: DB2

This Application worksheet records information about the DB2 application that is used in the ttclrsHA cluster.

Application name	DB2
Cluster name	ttclrsHA
Node relationship (cascading/concurrent/rotating) type of node priority policy used	Cascading

Key application files

	Folder/path	File system	Location	Sharing
Executables	/home/db2v8	/home	Internal	No
Configuration files				
Data files/devices	/Index/db2v8 /RData/db2v8 /LData/db2v8	/Index /RData /LData	External	Yes
Log files/devices				

Application failover strategy (P=Primary; T=Takeover)

	Node 1	Node 2	Node 3	Node 4	Node 5
Node name	ttclrs2	ttclrs1			
Strategy	Primary	Takeover			

Start application: Commands/procedures

local_start_db2

Start application: Verification commands/procedures

ps –ef|grep db2v8 > 1

Start application: Node reintegration/takeover caveats

Node	Node name	Caveat
1	ttclrs2	N/A
2	ttclrs1	N/A
3		
4		
5		
6		

Stop application: Commands/procedures

local_stop_db2

Stop application: Verification commands/procedures

No

Stop application: Node reintegration/takeover caveats

Node	Node name	Caveat
1	ttclrs2	No failback
2	ttclrs1	
3		
4		
5		
6		

AIX connections

This AIX Connections worksheet records the AIX connections realm and service pairs in the ttclrsHA cluster.

Cluster ID	1
Cluster name	ttclrsHA

Resource group	Nodes	Realm (NB,NW,AT)	Service name	Service type (file, print, term, nvt, atls)
rgJDE	ttclrs1		ttclrs1_svc	
	ttclrs2			

Resource group	Nodes	Realm (NB,NW,AT)	Service name	Service type (file, print, term, nvt, atls)
rgDB2	ttclrs2		ttclrs2_svc	
	ttclrs1			

Resource group	Nodes	Realm (NB,NW,AT)	Service name	Service type (file, print, term, nvt, atls)

Resource group	Nodes	Realm (NB,NW,AT)	Service name	Service type (file, print, term, nvt, atls)

Application servers

This Application Servers worksheet records information about the asJDE and asDB2 application servers in the ttclrsHA cluster.

Cluster ID	1
Cluster name	ttclrsHA

Note: Use full path names for all user-defined scripts.

Server name	asJDE
Start script	local_start_ow 10 30 JDE
Stop script	local_stop_ow 10 30 JDE

Server name	asDB2
Start script	local_start_db2
Stop script	local_stop_db2

Server name	
Start script	
Stop script	

Resource group: rgJDE

This Resource Groups worksheet records the rgJDE resource group in the ttclrsHA cluster.

Cluster ID	1
Cluster name	ttclrsHA

Resource group name	rgJDE
Node relationship	Cascading
Site relationship	
Participating node names	ttclrs1 ttclrs2
Dynamic node priority policy	
Processing order	
Service IP label	ttclrs1_svc
File systems	/u07 /u08
Filesystems consistency check	fsck
Filesystems recovery method	seq
Filesystems/folders to export	
Filesystems/folders to NFS mount	
Network for NFS mount	
Volume groups	vgJDE vgJDE2
Concurrent volume groups	
Raw disk PVIDs	
Connections services	
Fast connect services	
Tape resources	
Application servers	asJDE
Highly available comm links	
Primary WLM class	
Secondary WLM class	
Miscellaneous data	
Auto import volume groups	
Inactive takeover activated	False
Cascading without fallback enabled	True
Disk fencing activated	False
Filesystems mounted before IP configured	True

Resource group: rgUDB

This Resource Groups worksheet records the rgUDB resource group in the ttclrsHA cluster.

Cluster ID	1
Cluster name	ttclrsHA

Resource group name	rgUDB
Node relationship	Cascading
Site relationship	
Participating node names	ttclrs2 ttclrs1
Dynamic node priority policy	
Processing order	
Service IP label	ttclrs2_svc
File systems	/Index /RData /LData
Filesystems consistency check	fsck
Filesystems recovery method	seq
Filesystems/folders to export	
Filesystems/folders to NFS mount	
Network for NFS mount	
Volume groups	vgIndex vgRData vgLData
Concurrent volume groups	
Raw disk PVIDs	
Connections services	
Fast connect services	
Tape resources	
Application servers	asUDB
Highly available comm links	
Primary WLM class	
Secondary WLM class	
Miscellaneous data	
Auto import volume groups	
Inactive takeover activated	False
Cascading without fallback enabled	True
Disk fencing activated	False
Filesystems mounted before IP configured	True

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 74. Note that some of the documents referenced here may be available in softcopy only.

- *HACMP Enhanced Scalability: User-Defined Events*, SG24-5327
- *Configuring Highly Available Clusters Using HACMP 4.5*, SG24-6845
- *IBM TotalStorage: FAStT Best Practices Guide*, REDP-3690

Other publications

These publications are also relevant as further information sources:

- *IBM TotalStorage FastT Storage Manager V8.30 Installation and Support Guide*, GC26-7519
- *RS/6000 Cluster Technology: Event Management Programming Guide and Reference*, SA22-7354
- *RS/6000 Cluster Technology: Group Services Programming Guide and Reference*, SA22-7355
- *HACMP for AIX 4.5: Concepts and Facilities*, SC23-4276
- *HACMP for AIX 4.5: Planning Guide*, SC23-4277
- *HACMP for AIX 4.5: Installation Guide*, SC23-4278
- *HACMP for AIX 4.5: Administration Guide*, SC23-4279
- *HACMP for AIX 4.5: Enhanced Scalability Installation and Administration Guide*, SC23-4306
- *Release Notes* describing hardware and software requirements in the /usr/lpp/cluster/doc/release_notes directory

Online resources

These Web sites are also relevant as further information sources:

- Discussion of EnterpriseOne 8.9

  http://www.oracle.com/peoplesoft/integration.html

- Latest FAStT documents based on model number
  - FAStT200: http://www.ibm.com/servers/storage/support/disk/fastt200/
  - FAStT500: http://www.ibm.com/servers/storage/support/disk/fastt500/
  - DS4300: http://www.ibm.com/servers/storage/support/disk/ds4300/

  - DS4400: http://www.ibm.com/servers/storage/support/disk/ds4400/
  - DS4500: http://www.ibm.com/servers/storage/support/disk/ds4500/

- The aixreadme file that describes the most current information about AIX and device driver versions for FAStT for Storage Manager v8.21

  http://www.ibm.com/pc/support/site.wss/MIGR-43839.html

- The *aixreadme* file that describes the most current information about AIX and device driver versions for FAStT for Storage Manager v8.30

  http://www.ibm.com/pc/support/site.wss/MIGR-50177.html

- IBM TotalStorage FAStT Storage Interoperability matrix for information about various platforms, operating systems, host bus adapter, and applications support with IBM TotalStorage FAStT

  http://www.storage.ibm.com/disk/fastt/supserver.htm

- *HACMP configuration information for FAStT on pSeries/AIX Servers*

  ftp://ftp.software.ibm.com/storage/fastt/fastt500/HACMP_config_info.pdf

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

S

T

U

V

W

JD Edwards EnterpriseOne 8.9 High Availability and Storage

for IBM @server pSeries Clients

Follow the step-by-step configuration and set up instructions

Manage clusters of nodes to deliver high availability

Review example windows and worksheets for easy installation

As companies depend more and more on their enterprise software for daily operations, they require highly available systems. JD Edwards EnterpriseOne has long enabled their software with the IBM® High Availability Cluster Multi-Processing (HACMP) product for high availability.

This IBM Redbook describes the IBM HACMP/Enhanced Scalability (HACMP/ES) Version 4.5 solution to certify JD Edwards EnterpriseOne 8.9 in an IBM AIX® high availability environment. This book also describes HACMP/ES Version 5.1. It explains how to install and tune HACMP/ES, as well as reviews the setup steps to enable JD Edwards EnterpriseOne for HACMP/ES.

While writing this IBM Redbook, we conducted tests to simulate that the failures and observations made to ensure that the high availability environment performed properly. The results of these tests are also documented.

This redbook highlights and explains those characteristics that are specific to JD Edwards EnterpriseOne. It contains significant window examples and worksheets to help facilitate an easy step-by-step installation process. It does not replace the extensive installation and administration documentation that is already available for HACMP/ES.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6445-00 ISBN 0738491357