UniVerse

Using UniAdmin

Version 10.3
February, 2009

IBM Corporation
555 Bailey Avenue
San Jose, CA 95141

Licensed Materials – Property of IBM

© Copyright International Business Machines Corporation 2008, 2009. All rights reserved.

AIX, DB2, DB2 Universal Database, Distributed Relational Database Architecture, NUMA-Q, OS/2, OS/390, and OS/400, IBM Informix®, C-ISAM®, Foundation.2000 ™, IBM Informix® 4GL, IBM Informix® DataBlade® module, Client SDK™, Cloudscape™, Cloudsync™, IBM Informix® Connect, IBM Informix® Driver for JDBC, Dynamic Connect™, IBM Informix® Dynamic Scalable Architecture™ (DSA), IBM Informix® Dynamic Server™, IBM Informix® Enterprise Gateway Manager (Enterprise Gateway Manager), IBM Informix® Extended Parallel Server™, i.Financial Services™, J/Foundation™, MaxConnect™, Object Translator™, Red Brick® Decision Server™, IBM Informix® SE, IBM Informix® SQL, InformiXML™, RedBack®, SystemBuilder™, U2™, UniData®, UniVerse®, wIntegrate® are trademarks or registered trademarks of International Business Machines Corporation.

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

Windows, Windows NT, and Excel are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

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

Other company, product, and service names used in this publication may be trademarks or service marks of others.

Documentation Team: Claire Gustafson, Shelley Thompson, Anne Waite

US GOVERNMENT USERS RESTRICTED RIGHTS

Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Table of Contents

Preface

Organization of This Manual xii
UniVerse Documentation xvi
Related Documentation xviii
API Documentation xix

Chapter 1 Installing and Administering UniAdmin

Installation Process 1-3
Administering Servers 1-11
Registering Servers 1-11
Unregistering a Server 1-13
Editing Server Information 1-15
Managing Server Connections 1-16
Connecting to a Server 1-16
Disconnecting from a Server 1-17

Chapter 2 Administering Accounts

Accessing Accounts Admin 2-3
Adding a New Account 2-4
View UniVerse Account 2-5
Deleting a UniVerse Account 2-6
Retain Directory 2-6
Delete Directory 2-6

Chapter 3 Creating Backups from UniAdmin

Accessing the Backup Dialog Box 3-3
Choose Backup Location 3-4
Backing Up Data to Disk 3-4
Backing Up Data to Tape 3-5
Choose the Type of Backup 3-6

Choose Reporting Type 3-6
Label the Backup 3-6
Define Contents of Backup 3-6
Start Backup 3-7

Chapter 4 Configuring UniVerse Parameters

How UniVerse Uses Configurable Parameters 4-3
Access the Configuration Editor 4-4
Changing Parameter Values 4-5
Changing Configurable Parameters on UNIX 4-6
Changing Configurable Parameters on Windows Platforms 4-6
UniVerse Configurable Parameters 4-7

Chapter 5 Administering Data Sources

Administering Data Sources 5-3
Creating a New Data Source 5-4
Viewing Data Source Details 5-6
Deleting a Data Source 5-7

Chapter 6 Administering Deadlocks

Managing Dead Locks 6-3
Starting and Stopping the Dead Lock Manager 6-4
Configuring the Deadlock Manager 6-5
View Pending Deadlocks 6-7

Chapter 7 Administering Devices

Administering Tape Devices on UNIX 7-3
Defining a New Tape Device 7-3
Viewing or Modifying a Tape Device 7-8
Deleting a Tape Device 7-9
Administering Tape Devices on Windows Platforms 7-10
Defining a New Tape Device 7-10
Viewing or Modifying a Tape Device 7-14
Deleting a Tape Device 7-15
Using Test Tape 7-15
Administering Print Devices on UNIX 7-16
Defining a New Printer 7-17
Viewing or Modifying a Printer Definition 7-21
Deleting a Print Device Definition 7-22
Administering Print Devices on Windows Platforms 7-23

Configuring Other Devices 7-24
View Details about a Device 7-25
Delete a Device Definition 7-26

Chapter 8 Viewing Disk Space Usage

Chapter 9 Administering UniVerse Files

List Files in UniVerse Accounts 9-4
View File Properties 9-5
View File Statistics 9-11
Running File Diagnostics 9-13
Repairing Damaged Files 9-16
Select Display Options 9-20

Chapter 10 Importing Accounts

Importing Non-UniVerse Accounts to UniVerse 10-3
Manually Restoring Accounts from Tape 10-5
Restoring an Single-reel Account to a UNIX Server 10-5
Restoring a Multireel Account to a UNIX Server 10-5
Restoring Accounts to Windows Platforms 10-8

Chapter 11 Licensing UniVerse

Authorizing a New UniVerse License 11-3
Obtain Configuration Code 11-4
Obtain Authorization Code 11-4
Viewing License Information 11-5
Cleaning Up Current Licenses 11-6

Chapter 12 Administering Locks

File and Record Locks 12-4
Clearing Locks 12-5
Group Locks 12-7

Chapter 13 Administering Network Services

Administering the UniRPC on UNIX Systems 13-3
Specifying the RPC Port Number 13-4
Changing the RPC Port Number 13-5
Starting and Stopping the UniRPC Daemon on UNIX 13-5
Maintaining the *hosts* File 13-6
Managing a Windows Platform Telnet Session 13-9

Modifying the Telnet Session Parameters 13-11
Administering Users 13-14
Adding a New User 13-15
Starting Services 13-20

Chapter 14 UniVerse Replication

About Replication 14-4
About Hot Standby 14-4
Setting Up Data Replication 14-5
Configuring and Managing Data Replication 14-6
Managing a Publishing System 14-7
Configuring the Publishing System 14-9
The Publishing Config Dialog Box 14-9
Starting and Stopping the Subscribing System 14-11
Publishing Files 14-12
Publishing Information File 14-18
Managing the Subscribing System 14-20
Configuring the Subscribing System 14-20
Starting and Stopping the Subscribing System 14-22
Creating a List of Publishing Systems for the Subscriber 14-22
Subscribing Files 14-24
Managing Hot Standby Operations 14-32
Configuring a Hot Standby Subscriber 14-32
Turning On Fail-Over Mode 14-34
Reconciling the Hot Backup with the Publisher 14-34
Subscribing Information Files 14-36
Some Restrictions 14-38
What to Do When Disk Space Fills Up 14-39
Removing Obsolete Replication Log Files 14-40
What to Do When Replication Fails 14-41

Chapter 15 Restoring UniVerse Files

Restoring Files 15-3
Choosing the Restore Device 15-3
Checking the Backup Details 15-4
The UVRestore Window 15-5
Choosing What to Restore 15-6
Listing an Index of the Backup Image 15-8
Specifying How to Restore Files 15-8
. 15-9

Chapter 16 Managing Catalog Shared Memory

Setting Up Catalog Shared Memory 16-3
Defining Programs to Run in Shared Memory 16-5
Adding Programs to the SHM.TO.LOAD File 16-6
Removing a Program from the SHM.TO.LOAD File 16-6
Loading Programs into Catalog Shared Memory 16-7
Using Programs Stored in Catalog Shared Memory 16-9
Modifying Programs in Catalog Shared Memory 16-10
Adding a Program to Shared Memory 16-10
Updating a Program in Shared Memory 16-11
Removing a Program from Shared Memory 16-12
Removing the Catalog Shared Memory Segment 16-13

Chapter 17 Configuring SSL Through UniAdmin

Configuring SSL Through UniAdmin 17-2
Accessing UniVerse SSL Configuration Dialog Box 17-3
Creating a Certificate Request 17-4
Creating a Certificate 17-11
Creating a Security Context 17-19
Configuring SSL for UniObjects for Java or Telnet 17-35

Chapter 18 Transaction Logging

Configuring the Transaction Logging System 18-4
Setting Up or Changing Transaction Logging Configuration 18-4
Specifying Recoverable Files 18-7
Activating Files 18-7
Deactivating Files 18-8
Managing Transaction Logging 18-10
Managing Log Files 18-10
Managing Transaction Logging 18-14
Viewing and Deleting Information Files 18-16
Restoring Log Files 18-17
Recovering Files 18-20
Recovering Log Files from Disk 18-20
Recovering Files from a Full Log File on Tape 18-23
Deleting Restored Log Files 18-27

Chapter 19 Executing UniVerse Commands

Executing a UniVerse Command 19-3
UniVerse Administration Commands 19-5

Commands to Assign and Unassign Peripheral Devices 19-5
UniVerse BASIC Commands 19-5
File and Account Maintenance Commands 19-6
Managing Processes 19-8
Monitoring Resources 19-8
Task and Record Locking 19-9
Using the Command History 19-10
Reexecuting Commands 19-10
Editing a Command 19-11
Saving Commands 19-11
The UniVerse Command Output Window 19-12

Chapter 20 UniVerse User Administration

Interactive Users 20-4
Background Processes 20-4
Terminating a Process 20-5
Terminating a User Process 20-5
Terminating a Background Process 20-5
Sending Messages with UniVerse Admin 20-6
Sending Messages on UNIX Systems 20-6
Sending Message on Windows Platforms 20-7

Chapter 21 Automatic Data Encryption

Using UniAdmin for Encryption 21-2
Encryption Wallet Management 21-7

Preface

This manual describes the UniAdmin. This tool enables you to perform the following tasks:

- Administer Accounts
  - Create a new account
  - View details of an existing account
  - Delete an Account
- Create Backups
  - Backup UniVerse accounts
- Change UniVerse configuration parameters
- Administer UCI or UniVerse ODBC data sources
  - Create a data source
  - Delete an existing data source
  - View or modify information about a data source
- Administer UniVerse deadlocks
  - Start or stop the UniVerse Deadlock Manager
  - Configure the Deadlock Manager
  - Manually resolve file locking conflicts
- Administer Devices
  - Configure tapes
  - Configure printer
  - Configure other device
- View file system usage
  - Define and update log configuration table
  - Define and update archive configuration table
- File Tools
- View files in UniVerse accounts

- Import
  - Import non-UniVerse accounts from Prime Information or Pick systems
  - Manually import non-UniVerse accounts from tape
- License UniVerse or UV/Net
  - View information about licenses in use
  - Clean UniVerse license
  - Update UniVerse license
- Administer Locks
  - Monitor locks
  - Clear file locks
  - Clear record locks
  - Clear group locks
- Administer network services
  - Administer telnet services
  - Administer other services
- Monitor UniVerse processes
  - Monitor performance
- Administer Replication
- Restore files from backup
- Manage shared programs
  - Define programs to load into shared memory
  - Modify programs in shared memory
  - Load program into catalog shared memory
  - Remove catalog shared memory segment
- Configure SSL through UniAdmin
- Administer Transaction Logging
- Execute UniVerse commands
- Administer UniVerse users
  - View user and backgroung processes
  - Terminate a UniVerse user process
  - Send a message to a UniVerse user

Organization of This Manual

This manual is organized as follows:

Chapter 1, "Installing and Administering UniAdmin," describes how to install and administer UniAdmin.

Chapter 2, "Administering Accounts," describes how to create and administer UniVerse accounts.

Chapter 3, "Creating Backups from UniAdmin," describes how to back up UniVerse data from UniAdmin.

Chapter 4, "Configuring UniVerse Parameters," describes how to use UniAdmin to view and change UniVerse configuration parameters.

Chapter 5, "Administering Data Sources," describes how to administer data sources through UniAdmin. You can view and change the current set of defined UCI and UniVerse ODBC data sources.

Chapter 6, "Administering Deadlocks," describes how to manage dead locks through UniAdmin. You can start or stop the deadlock manager, configure the deadlock manager, or manually resolve file locking conflicts from the UniVerse Dead Lock Admin dialog box.

Chapter 7, "Administering Devices," describes how to configure and admin ister tape devices, printers, and line devices through UniAdmin.

Chapter 8, "Viewing Disk Space Usage," describes how to display file system disk usage through UniAdmin.

Chapter 9, "Administering UniVerse Files," describes how to configure and administer UniVerse files through UniAdmin.

Chapter 10, "Importing Accounts," describes how to import Prime INFORMATION or Pick accounts to UniVerse.

Chapter 11, "Licensing UniVerse," describes how to license UniVerse through UniAdmin.

Chapter 12, "Administering Locks," describes how to administer locks through UniAdmin.

Chapter 13, "Administering Network Services," describes how to administer network services through UniAdmin.

Chapter 14, “UniVerse Replication,” describes how to administer replication through UniAdmin.

Chapter 15, “Restoring UniVerse Files,” describes how to restore UniVerse files from a backup through UniAdmin.

Chapter 16, “Managing Catalog Shared Memory,” describes how to load cataloged BASIC programs into shared memory and run them from there. Shared memory reduces the amount of memory needed for multiple users to run the same program at the same time.

Chapter 17, “ Configuring SSL Through UniAdmin,” describes how to administer SSL through UniAdmin.

Chapter 18, “ Transaction Logging,” describes transaction logging, a way of securing a copy of your UniVerse file updates in a special set of log files. You can then use these log files to recover your system to a working state if you encounter machine or media failure.

Chapter 19, “ Executing UniVerse Commands,” describes how to use the **UniVerse Command** dialog box to execute UniVerse commands from UniAdmin.

Chapter 20, “ UniVerse User Administration,” describes how to use UniAdmin to view or terminate UniVerse processes and to send messages.

Documentation Conventions

This manual uses the following conventions:

Convention	Usage
Bold	In syntax, bold indicates commands, function names, and options. In text, bold indicates keys to press, function names, menu selections, and MS-DOS commands.
UPPERCASE	In syntax, uppercase indicates UniVerse commands, keywords, and options; BASIC statements and functions; and SQL statements and keywords. In text, uppercase also indicates UniVerse identifiers such as file names, account names, schema names, and Windows file names and paths.
Italic	In syntax, italic indicates information that you supply. In text, italic also indicates UNIX commands and options, file names, and paths.
`Courier`	Courier indicates examples of source code and system output.
`Courier Bold`	In examples, courier bold indicates characters that the user types or keys the user presses (for example, **`<Return>`**).
[ ]	Brackets enclose optional items. Do not type the brackets unless indicated.
{ }	Braces enclose nonoptional items from which you must select at least one. Do not type the braces.
itemA | itemB	A vertical bar separating items indicates that you can choose only one item. Do not type the vertical bar.
...	Three periods indicate that more of the same type of item can optionally follow.
‰	A right arrow between menu options indicates you should choose each option in sequence. For example, "Choose File ‰ Exit" means you should choose File from the menu bar, then choose Exit from the File pull-down menu.
`I`	Item mark. For example, the item mark (**`I`**) in the following string delimits elements 1 and 2, and elements 3 and 4: `1`**`I`**`2`**`F`**`3`**`I`**`4`**`V`**`5`

Documentation Conventions

Convention	Usage
F	Field mark. For example, the field mark (**F**) in the following string delimits elements FLD1 and VAL1: FLD1**F**VAL1**V**SUBV1**S**SUBV2
V	Value mark. For example, the value mark (**V**) in the following string delimits elements VAL1 and SUBV1: FLD1**F**VAL1**V**SUBV1**S**SUBV2
S	Subvalue mark. For example, the subvalue mark (**S**) in the following string delimits elements SUBV1 and SUBV2: FLD1**F**VAL1**V**SUBV1**S**SUBV2
T	Text mark. For example, the text mark (**T**) in the following string delimits elements 4 and 5: 1**F**2**S**3**V**4**T**5

Documentation Conventions

The following conventions are also used:

- Syntax definitions and examples are indented for ease in reading.
- All punctuation marks included in the syntax—for example, commas, parentheses, or quotation marks—are required unless otherwise indicated.
- Syntax lines that do not fit on one line in this manual are continued on subsequent lines. The continuation lines are indented. When entering syntax, type the entire syntax entry, including the continuation lines, on the same input line.

UniVerse Documentation

UniVerse documentation includes the following:

UniVerse Installation Guide: Contains instructions for installing UniVerse 10.3.

UniVerse New Features Version 10.3: Describes enhancements and changes made in the UniVerse 10.3 release for all UniVerse products.

UniVerse BASIC: Contains comprehensive information about the UniVerse BASIC language. It is for experienced programmers.

UniVerse BASIC Commands Reference: Provides syntax, descriptions, and examples of all UniVerse BASIC commands and functions.

UniVerse BASIC Extensions: Describes the following extensions to UniVerse BASIC: UniVerse BASIC Socket API, Using CallHTTP, and Using WebSphere MQ with UniVerse.

UniVerse BASIC SQL Client Interface Guide: Describes how to use the BASIC SQL Client Interface (BCI), an interface to UniVerse and non-UniVerse databases from UniVerse BASIC. The BASIC SQL Client Interface uses ODBC-like function calls to execute SQL statements on local or remote database servers such as UniVerse, DB2, SYBASE, or INFORMIX. This book is for experienced SQL programmers.

Administering UniVerse: Describes tasks performed by UniVerse administrators, such as starting up and shutting down the system, system configuration and maintenance, system security, maintaining and transferring UniVerse accounts, maintaining peripherals, backing up and restoring files, and managing file and record locks, and network services. This book includes descriptions of how to use the UniAdmin program on a Windows client and how to use shell commands on UNIX systems to administer UniVerse.

Using UniAdmin: Describes the UniAdmin tool, which enables you to configure UniVerse, configure and manage servers and databases, and monitor UniVerse performance and locks.

UniVerse Transaction Logging and Recovery: Describes the UniVerse transaction logging subsystem, including both transaction and warmstart logging and recovery. This book is for system administrators.

UniVerse Security Features: Describes security features in UniVerse, including configuring SSL through UniAdmin, using SSL with the CallHttp and Socket interfaces, using SSL with UniObjects for Java, and automatic data encryption.

UniVerse System Description: Provides detailed and advanced information about UniVerse features and capabilities for experienced users. This book describes how to use UniVerse commands, work in a UniVerse environment, create a UniVerse database, and maintain UniVerse files.

UniVerse User Reference: Contains reference pages for all UniVerse commands, keywords, and user records, allowing experienced users to refer to syntax details quickly.

Guide to RetrieVe: Describes RetrieVe, the UniVerse query language that lets users select, sort, process, and display data in UniVerse files. This book is for users who are familiar with UniVerse.

Guide to ProVerb: Describes ProVerb, a UniVerse processor used by application developers to execute prestored procedures called procs. This book describes tasks such as relational data testing, arithmetic processing, and transfers to subroutines. It also includes reference pages for all ProVerb commands.

Guide to the UniVerse Editor: Describes in detail how to use the Editor, allowing users to modify UniVerse files or programs. This book also includes reference pages for all UniVerse Editor commands.

UniVerse NLS Guide: Describes how to use and manage UniVerse's National Language Support (NLS). This book is for users, programmers, and administrators.

UniVerse SQL Administration for DBAs: Describes administrative tasks typically performed by DBAs, such as maintaining database integrity and security, and creating and modifying databases. This book is for database administrators (DBAs) who are familiar with UniVerse.

UniVerse SQL User Guide: Describes how to use SQL functionality in UniVerse applications. This book is for application developers who are familiar with UniVerse.

UniVerse SQL Reference: Contains reference pages for all SQL statements and keywords, allowing experienced SQL users to refer to syntax details quickly. It includes the complete UniVerse SQL grammar in Backus Naur Form (BNF).

Related Documentation

The following documentation is also available:

UniVerse GCI Guide: Describes how to use the General Calling Interface (GCI) to call subroutines written in C, C++, or FORTRAN from BASIC programs. This book is for experienced programmers who are familiar with UniVerse.

UniVerse ODBC Guide: Describes how to install and configure a UniVerse ODBC server on a UniVerse host system. It also describes how to use UniVerse ODBC Config and how to install, configure, and use UniVerse ODBC drivers on client systems. This book is for experienced UniVerse developers who are familiar with SQL and ODBC.

UV/Net II Guide: Describes UV/Net II, the UniVerse transparent database networking facility that lets users access UniVerse files on remote systems. This book is for experienced UniVerse administrators.

UniVerse Guide for Pick Users: Describes UniVerse for new UniVerse users familiar with Pick-based systems.

Moving to UniVerse from PI/open: Describes how to prepare the PI/open environment before converting PI/open applications to run under UniVerse. This book includes step-by-step procedures for converting INFO/BASIC programs, accounts, and files. This book is for experienced PI/open users and does not assume detailed knowledge of UniVerse.

API Documentation

The following books document application programming interfaces (APIs) used for developing client applications that connect to UniVerse and UniData servers.

Administrative Supplement for Client APIs: Introduces IBM's seven common APIs, and provides important information that developers using any of the common APIs will need. It includes information about the UniRPC, the UCI Config Editor, the *ud_database* file, and device licensing.

UCI Developer's Guide: Describes how to use UCI (Uni Call Interface), an interface to UniVerse and UniData databases from C-based client programs. UCI uses ODBC-like function calls to execute SQL statements on local or remote UniVerse and UniData servers. This book is for experienced SQL programmers.

IBM JDBC Driver for UniData and UniVerse: Describes UniJDBC, an interface to UniData and UniVerse databases from JDBC applications. This book is for experienced programmers and application developers who are familiar with UniData and UniVerse, Java, JDBC, and who want to write JDBC applications that access these databases.

InterCall Developer's Guide: Describes how to use the InterCall API to access data on UniVerse and UniData systems from external programs. This book is for experienced programmers who are familiar with UniVerse or UniData.

UniObjects Developer's Guide: Describes UniObjects, an interface to UniVerse and UniData systems from Visual Basic. This book is for experienced programmers and application developers who are familiar with UniVerse or UniData, and with Visual Basic, and who want to write Visual Basic programs that access these databases.

UniObjects for Java Developer's Guide: Describes UniObjects for Java, an interface to UniVerse and UniData systems from Java. This book is for experienced programmers and application developers who are familiar with UniVerse or UniData, and with Java, and who want to write Java programs that access these databases.

UniObjects for .NET Developer's Guide: Describes UniObjects, an interface to UniVerse and UniData systems from .NET. This book is for experienced programmers and application developers who are familiar with UniVerse or UniData, and with .NET, and who want to write .NET programs that access these databases.

Using UniOLEDB: Describes how to use UniOLEDB, an interface to UniVerse and UniData systems for OLE DB consumers. This book is for experienced programmers and application developers who are familiar with UniVerse or UniData, and with OLE DB, and who want to write OLE DB programs that access these databases.

Chapter

1

Installing and Administering UniAdmin

Installation Process . . . . . 1 -3
Administering Servers . . . . . 1 -11
Registering Servers . . . . . 1 -11
Unregistering a Server. . . . . 1 -13
Editing Server Information . . . . . 1 -15
Managing Server Connections . . . . . 1 -16
Connecting to a Server . . . . . 1 -16
Disconnecting from a Server. . . . . 1 -17

This chapter describes how to install and administer UniAdmin.

Installation Process

1. Log On As an Administrator

You must log on with Administrator privileges to install UniAdmin. Either log on to the Administrator account on the Windows system, or log on as a member of the local Administrators group.

2. Exit Other Applications

Before proceeding, exit any other Windows applications you may have open.

3. Load the UniVerse Client CD

Place the UniVerse Client CD in your CD-ROM drive. Make sure you have the UniVerse Product Configuration sheet that is included with your media. You will need this form when you license the product after installation.

4. Execute the Installation

From the **Setup** screen, double-click **UniAdmin**, as shown in the following example:

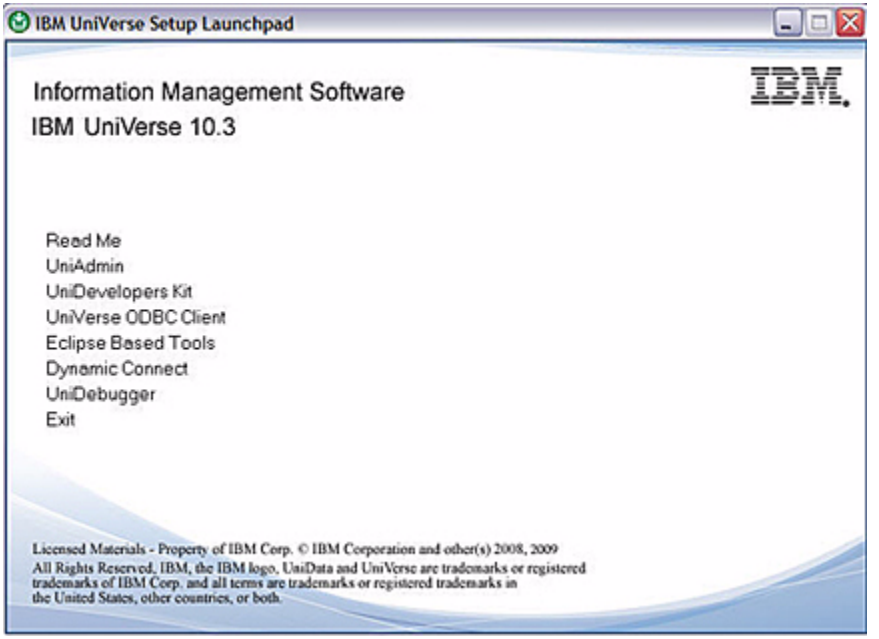

5. Proceed With Installation

From the **UniAdmin Setup** window, click **Next** to continue with the installation, or click **Cancel** to exit the installation process.

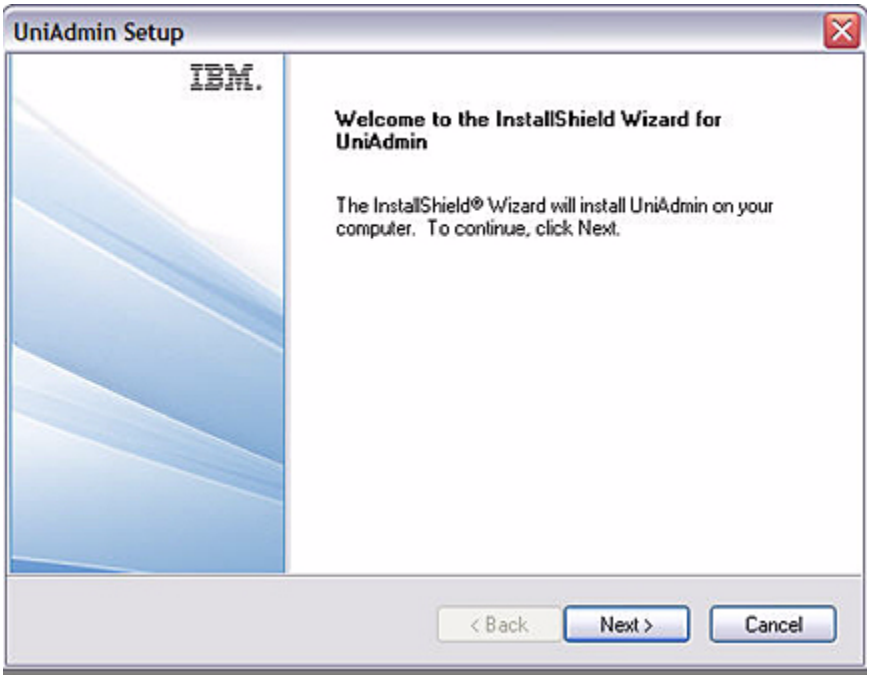

6. Review License Agreement

The **License Agreement** window appears, as shown in the following example:

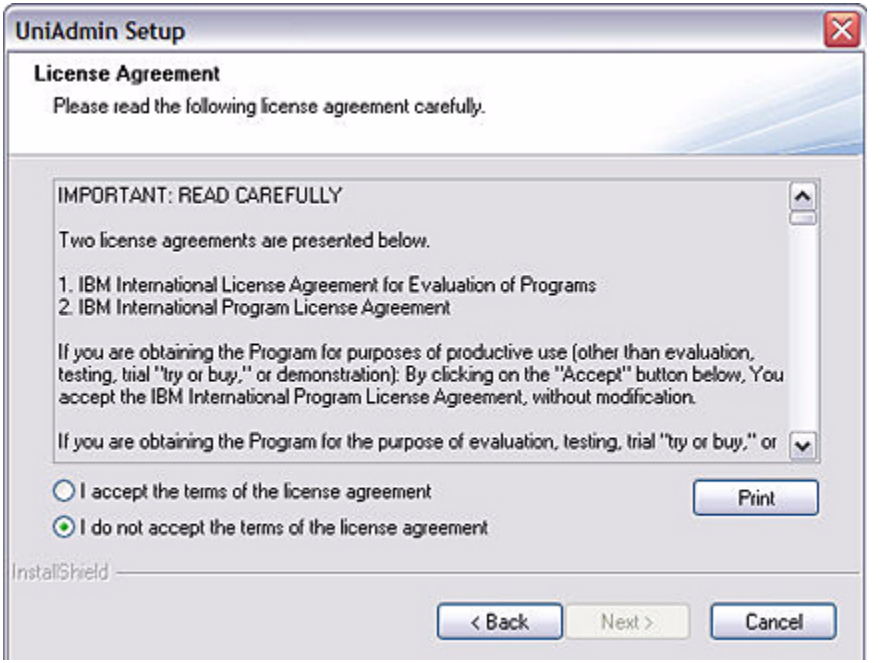

If you agree with the terms of the license agreement, choose **I accept the terms of the license agreement**. If you do not agree with the terms of the license agreement, select **I do not accept the terms of the license agreement**. If you choose this option, you will exit the installation.

7. Choose Destination Folder

Determine where you want to install UniAdmin. By default, it is installed in \IBM\UniAdmin. If you want to select a different directory, click **Browse.** The following example illustrates the **Choose Folder** dialog box:

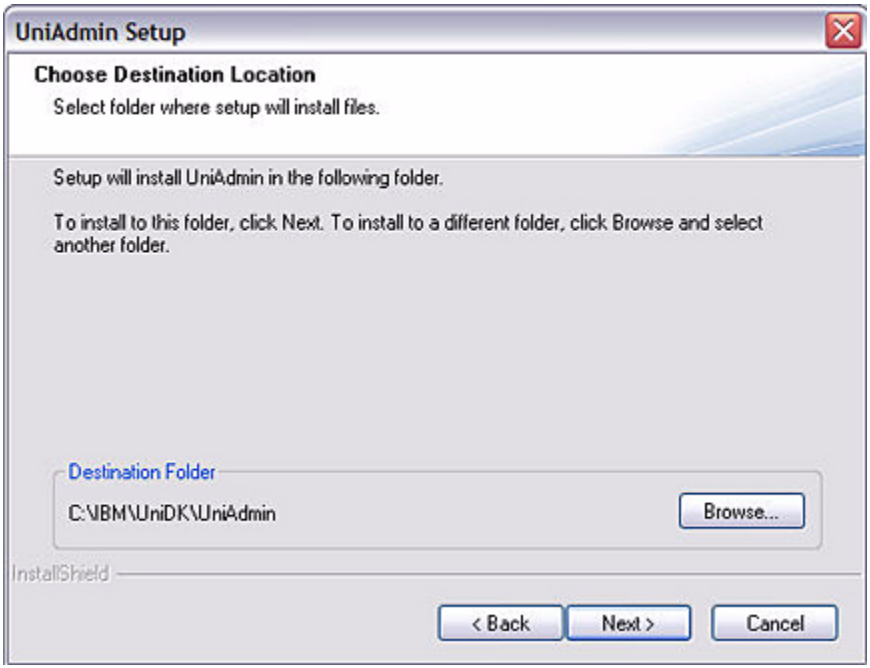

Click **Back** to return to the previous screen, **Next** to continue with the installation, or **Cancel** to exit the installation without saving changes.

7. Select Program Folder

Determine the name of the program folder where you want to install UniAdmin. The following example illustrates the **Select Program Folder** dialog box:

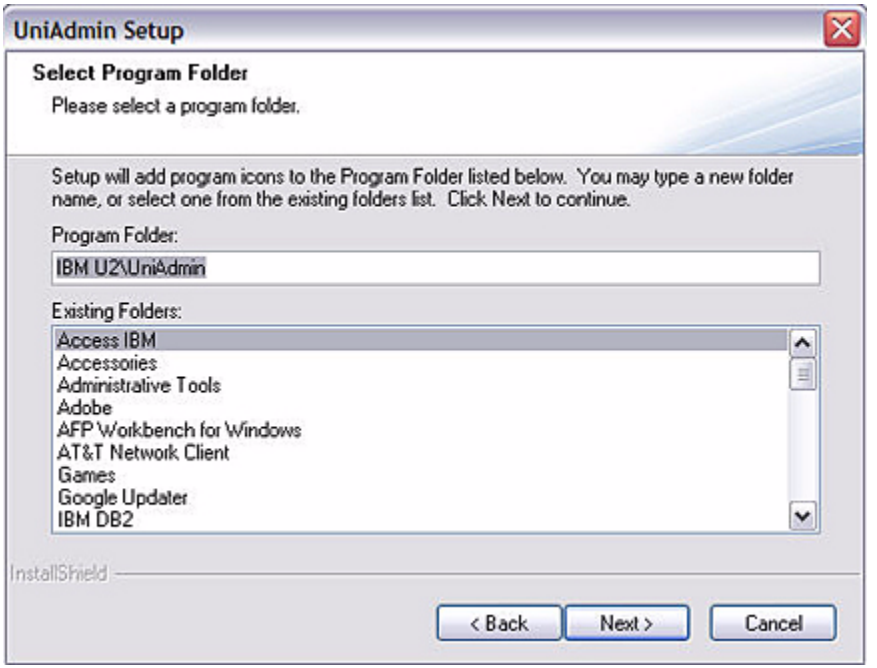

By default, the installation process installs UniAdmin in the IBM U2\UniAdmin folder. If you want to install it in a different folder, enter the name of the folder in the **Program Folders** box, or select the folder from the **Existing Folders** list.

Click **Back** to return to the previous screen, **Next** to continue with the installation, or **Cancel** to exit the installation without saving changes.

8. Copy Files

UniAdmin now has enough information to begin copying files. A dialog box similar to the following example appears:

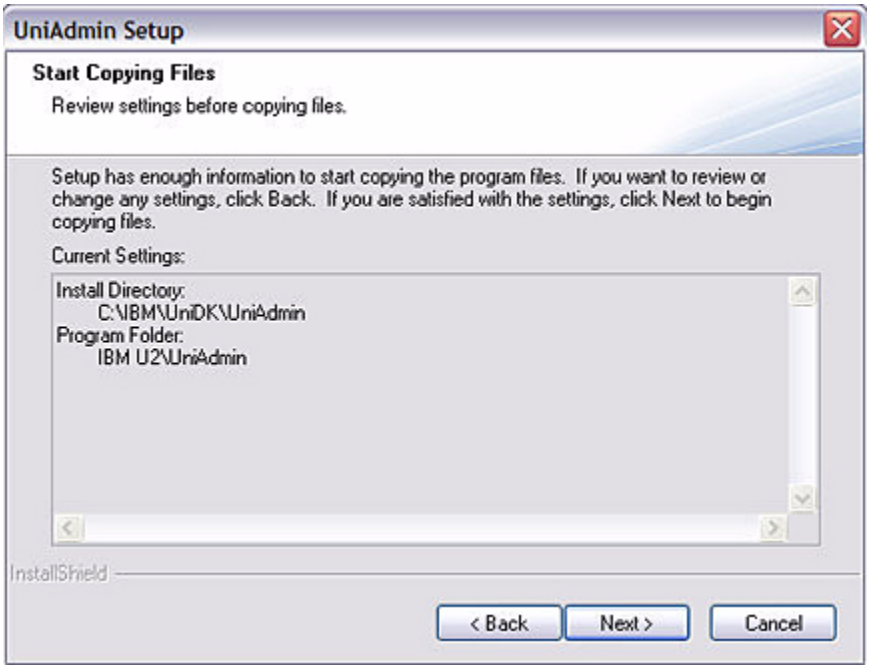

If you want to change any of the settings, click **Back**. Click **Next** to continue with the installation process, or click **Cancel** to exit the installation program.

9. Complete the Installation

After UniData copies the program files, the following dialog box appears:

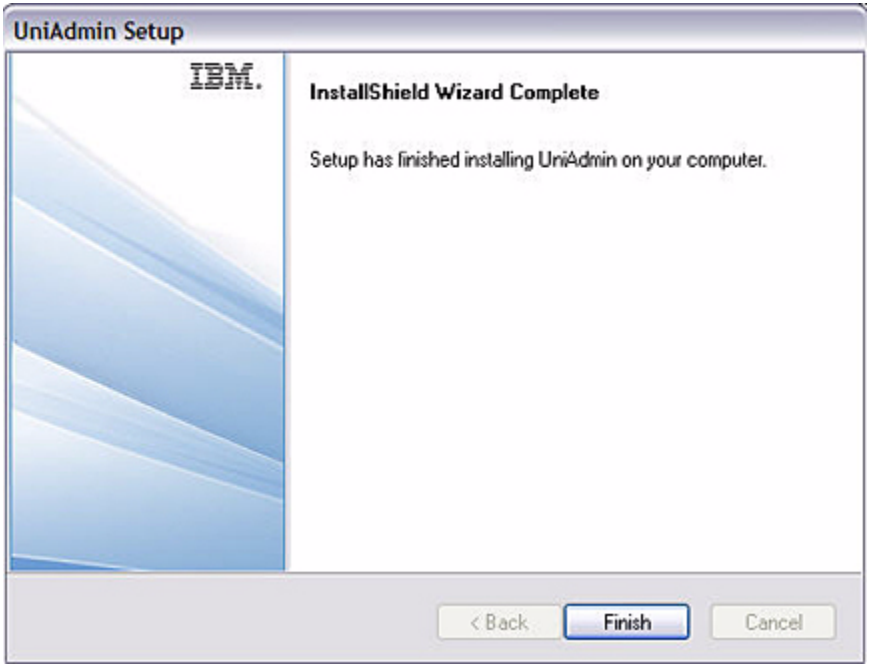

Click **Finish** to complete the installation.

Administering Servers

You must register each server that you want to access through UniAdmin. The server must be running UniVerse 10.0 or later.

Registering Servers

Complete the following steps to register servers for UniAdmin.

1. Start UniAdmin

From the **Start** menu, select **Programs**, then select the folder where UniAdmin resides, and then click **UniAdmin**. A window similar to the following example appears:

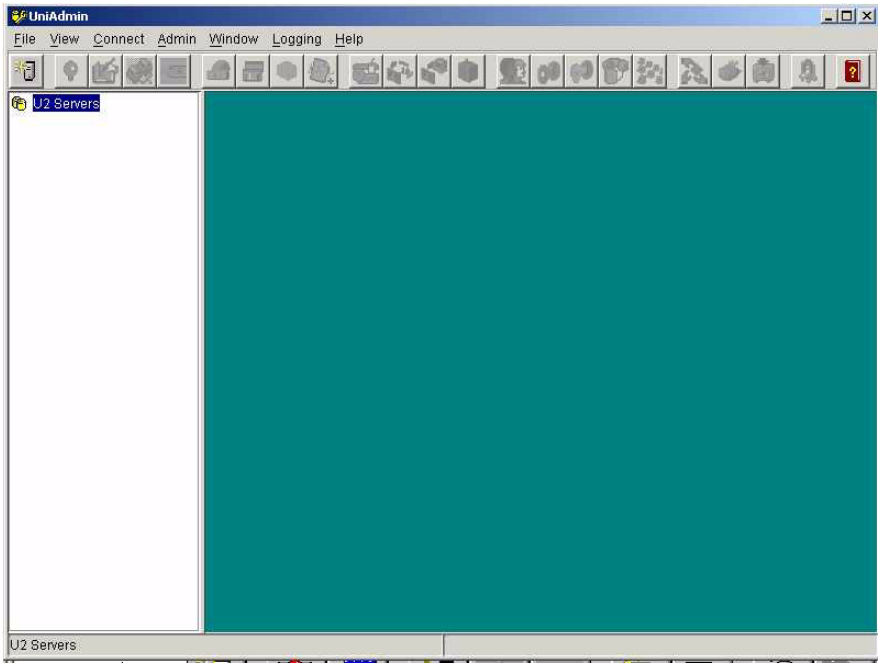

2. Define Server to Register

You can access the **Server Register** dialog box in any of the following ways:

- Using the right mouse button (right-click) **U2Servers**, and then click **Register Servers**.
- Click the server icon on the toolbar, as shown in the following example:

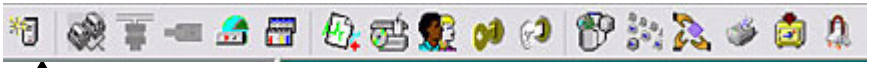

Server Icon

- From the **UniAdmin** menu, click **Connect**, and then click **New**.

A dialog box similar to the following appears:

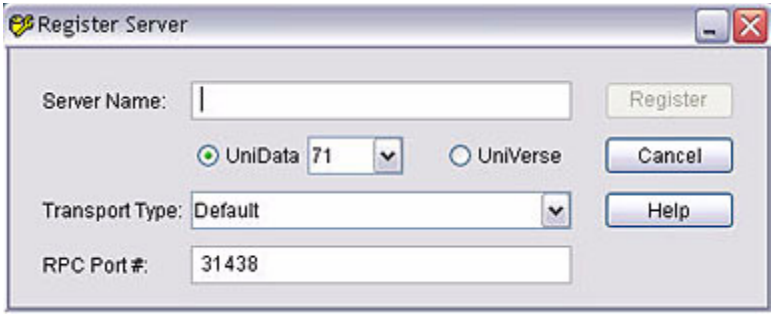

Enter the name of the server you want to register in the **Server Name** box.

3. Choose Database

Select the database against which you are using UniAdmin. The database is either UniData or UniVerse.

4. Choose the Transport Type

From the **Transport Type** box, choose the type of communication you are using to the server. You can choose Default, TCP/IP, or Lan Manager. The default is TCP/IP.

5. Register the Server

Click **Register** to register the server, or **Cancel** to exit the program without saving changes. The registered servers appear in the **UniAdmin** window, as shown in the following example:

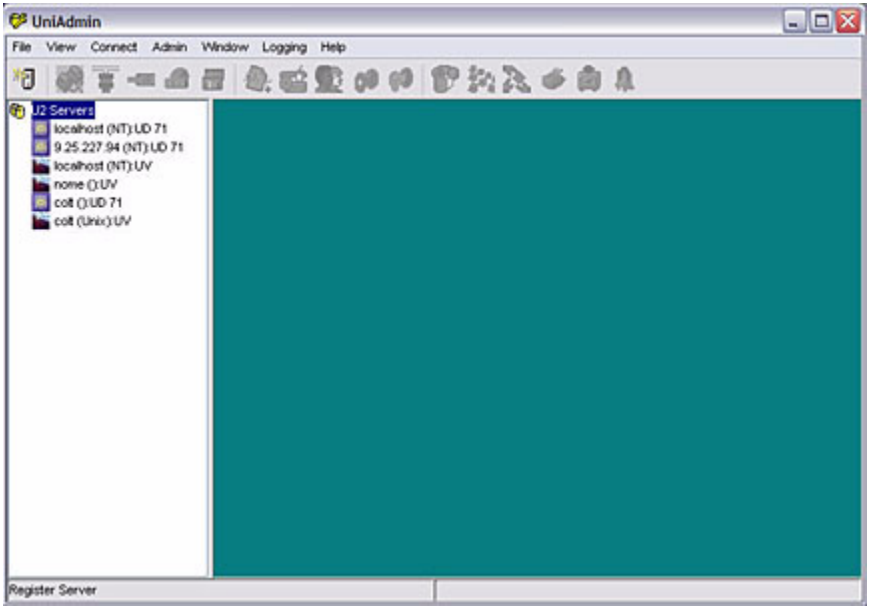

Unregistering a Server

Complete the steps described in this section to unregister a server.

From the **UniAdmin** window, right-click on the server you want to unregister, and then click **Un-Register**. A dialog box similar to the following appears:

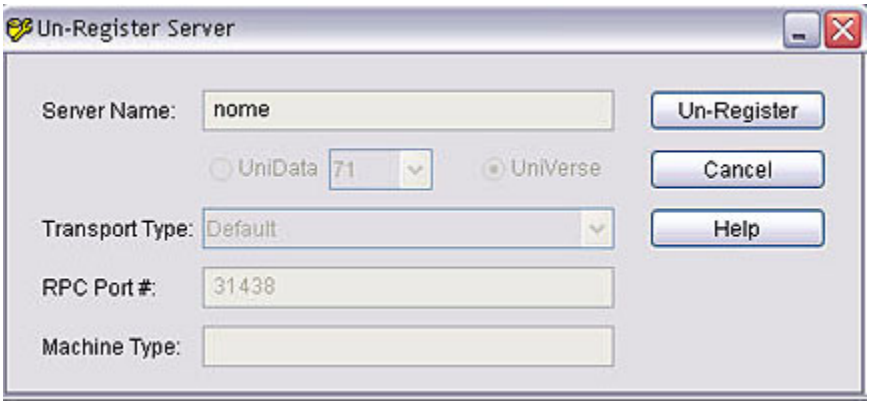

If the server information in the dialog box is accurate, click **Un-Register** to unregister the server, or click **Cancel** to exit the program without saving changes.

After you unregister a server, it no longer appears on the **UniAdmin** window.

Editing Server Information

If you want to change information you previously entered about a registered server, right-click on the server for which you want to edit information, and then click **Properties**. A dialog box similar to the following appears:

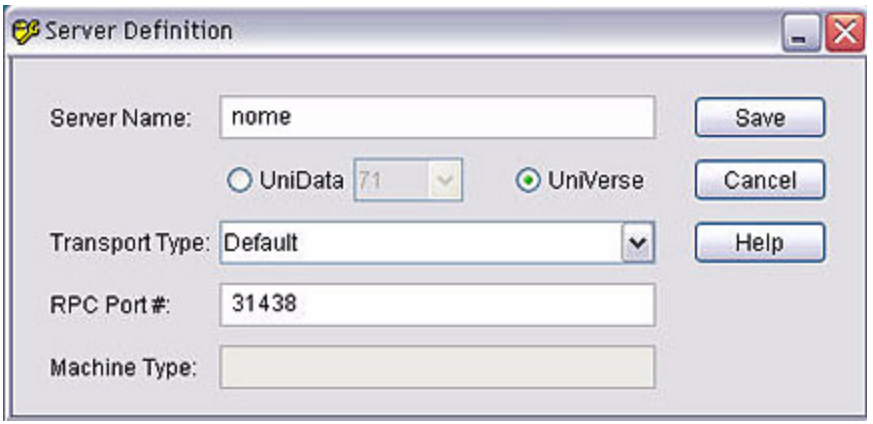

Edit the information you want to change. Click **Save** to save the changes, or click **Cancel** to exit the dialog box without saving changes.

Managing Server Connections

This section describes how to connect and disconnect from a UniVerse server.

Connecting to a Server

To connect to a server, double-click the server to which you want to connect, or right-click the server and then click **Connect,** or select Connect from the menu, and click the server to which you want to connect. If you selected the **Always prompt login panel** check box when you registered the server, a login dialog box similar to the following appears:

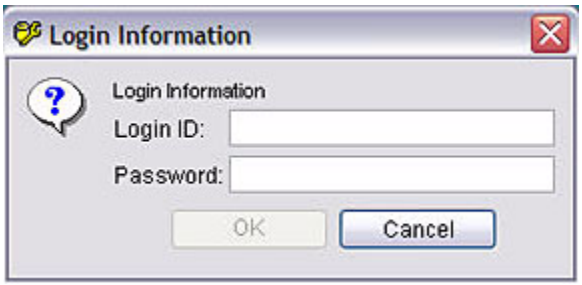

Enter the login name of the user you are connecting as in the **Login ID** box. Tab to the **Password** box and enter the password for the user.

***Note**: You must log on as root on UniVerse for UNIX, or Administrator on UniVerse for Windows platforms for many of the administration options available from UniAdmin. The necessary privileges associated with administrative processes are described in Administering UniVerse on UNIX and Administering UniVerse on Windows Platforms.*

Administration Options

After you connect to the server, a window similar to the following example appears:

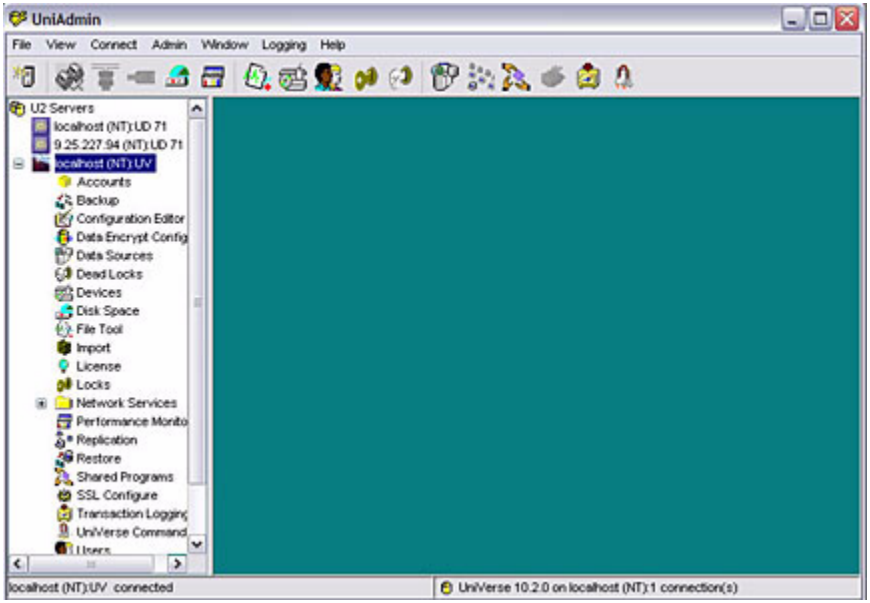

Double-click on the server folder to display a list of processes available, or review the options from the menu. Each option is described in detail in later chapters.

Disconnecting from a Server

To disconnect from a server, right-click on the server from which you want to disconnect, and then click **Disconnect**.

Chapter

2

Administering Accounts

Accessing Accounts Admin 2-3
Adding a New Account 2-4
View UniVerse Account 2-5
Deleting a UniVerse Account 2-6
Retain Directory 2-6
Delete Directory 2-6

This chapter describes how to create and administer UniVerse accounts.

Accessing Accounts Admin

You can access the *Accounts Admin* dialog box in any of the following ways:

- From the UniAdmin window, double-click *Accounts*.
- From the UniAdmin menu, select *Admin*, and then click *Accounts*

Regardless of the manner in which you administer accounts, a dialog box similar to the following appears:

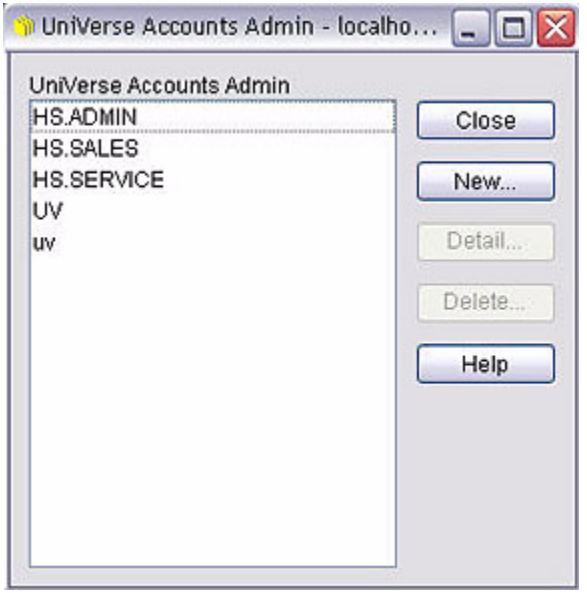

Adding a New Account

To create a new account, click the **New**. A dialog box similar to the following example appears:

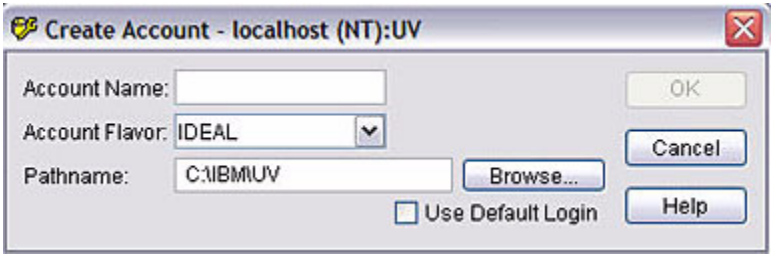

To create a new account, complete the following steps:

1. **Choose Account Name**

   In the *Account Name* box, enter the name the new account.

2. **Choose Account Flavor**

   In the Account Flavor box, select the flavor of the account. Valid flavors are:

   - IDEAL
   - INFORMATION
   - PI/OPEN
   - PICK
   - REALITY
   - IN2

   For more information about UniVerse account flavors, see *Administering UniVerse*.

3. **Define Path**

   In the ***Pathname*** box, enter the location of the new account, or click **Browse** to select the path.

4. **Use Default Login**

   Select the **Use Default Login** check box if you do not want to create a default login paragraph for the new account.

Note*: To prevent users from creating new accounts, create an entry for them in the UV.FLAVOR file, setting the access code to N.*

View UniVerse Account

To view information about a UniVerse account, select the account from the ***UniVerse Accounts*** window, and then click ***Detail***. A dialog box similar to the following example appears:

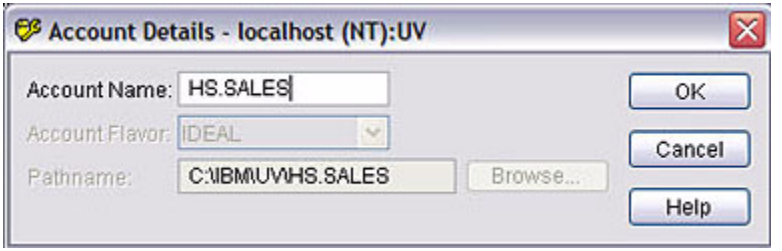

This dialog box displays the account name, the account flavor, and the path to the account.

Deleting a UniVerse Account

To delete a UniVerse account, select the account you want to delete in the ***UniVerse Accounts*** window, and then click ***Delete***. A dialog box similar to the following appears:

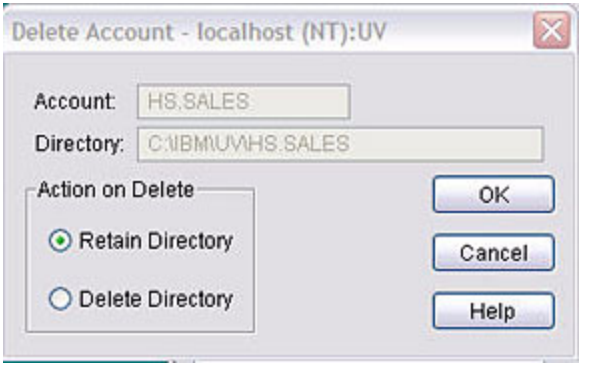

Retain Directory

If you select the **Retain Directory** option, the account is removed from the UV.ACCOUNTS file, but UniVerse retains the directory and its contents.

Delete Directory

If you select the **Delete Directory** option, UniVerse removes the account from the UV.ACCOUNTS file, in addition to deleting the account directory and its contents.

If other accounts defined in the UV.ACCOUNTS file use files in the account you want to delete, UniAdmin prompts if you want to delete these accounts, as well.

Click **OK** to delete the account, or click **Cancel** to exit the program without saving changes.

Chapter

3

Creating Backups from UniAdmin

Accessing the Backup Dialog Box 3-3
Choose Backup Location 3-4
- Backing Up Data to Disk 3-4
- Backing Up Data to Tape 3-5
- Choose the Type of Backup 3-6
- Choose Reporting Type 3-6
- Label the Backup 3-6
- Define Contents of Backup 3-6
- Start Backup 3-7

This chapter describes how to back up UniVerse data from UniAdmin.

Accessing the Backup Dialog Box

You can access the ***UniVerse Backup*** dialog box in any of the following ways:

- From the UniAdmin window, double-click **Backup**.
- From the UniAdmin menu, select **Admin**, and then click **Backup**.

Regardless of the manner you choose to access the **UniVerse Backup** dialog box, the following dialog box appears:

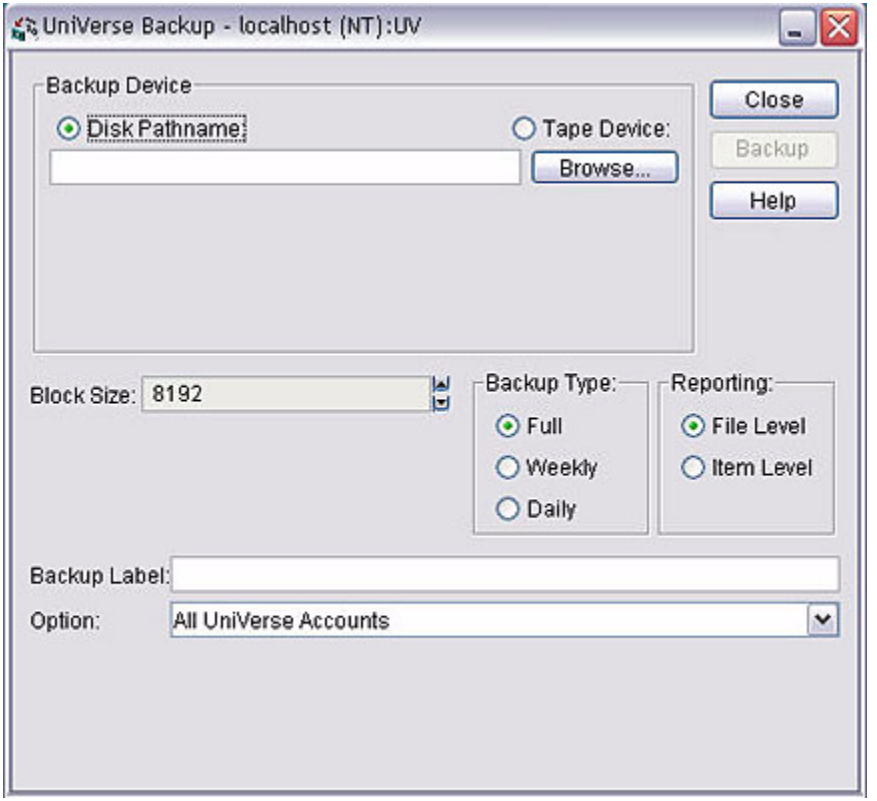

Choose Backup Location

You can choose to back up your data either to a disk path or tape.

Backing Up Data to Disk

If you want to back up your data to disk, click **Disk Pathname**.

In the **Disk Pathname** box, enter the full path of a disk file to which you want to back up the data, or click **Browse** to choose the location. The file to which you are backing up your data must already exist.

Backing Up Data to Tape

If you want to back up your data to tape, click **Tape Device**. A dialog box similar to the following example appears:

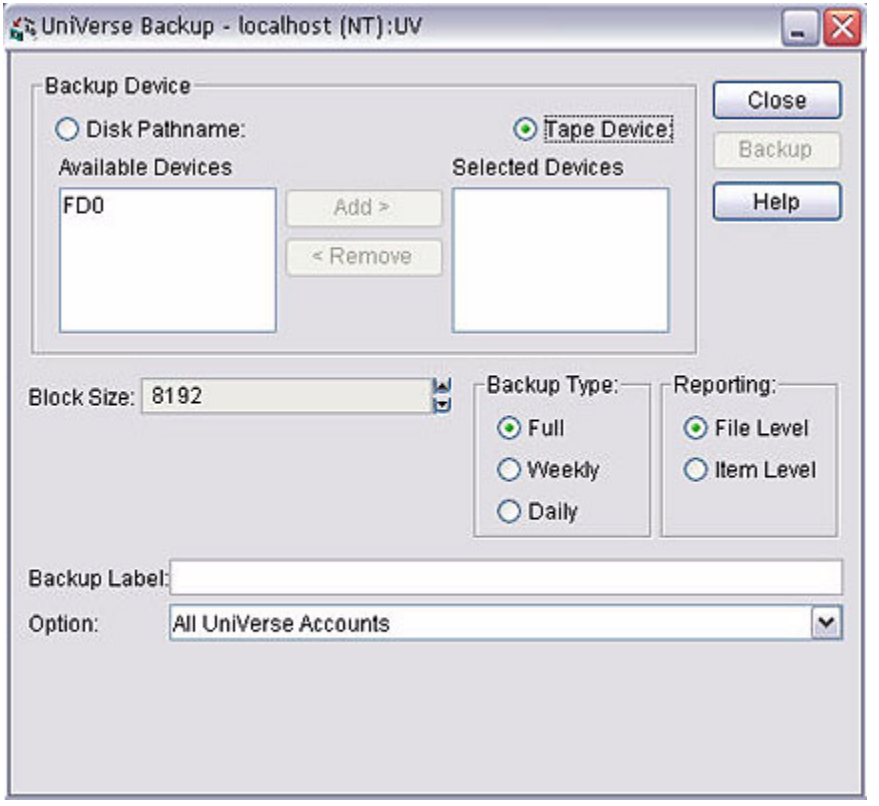

The **Available Devices** list displays all the tape devices defined in the &DEVICE& file. Click one or more devices you want to use, and then click **Add** to move it to the **Selected Devices** list.

The order of the devices in the **Selected Devices** list determines the order in which they are used by the backup process. You can reorder the list using a drag-and-drop operation.

To remove a tape device from the **Selected Devices** list, click the device you want to remove, and then click **Remove**.

Choose Block Size

In the **Block Size** box, enter the block size you want to use when backing up your data, or use the **Block Size** arrows to select a value. Values must be in increments of 512. The default is 8192.

Choose the Type of Backup

Click the type of backup you want to perform. Available backup types are:

- Full
- Weekly
- Daily

Choose Reporting Type

Click the type of report you want to display. Available reporting types are:

- None — No reports are generated during the backup process. UniVerse does report the start and end of the backup.
- File Level — The full paths of the files appear in the window during the backup process.
- Item Level — The full paths of the files appear in the window during the backup process, as well as record names for UniVerse hashed files.

Label the Backup

In the **Backup Label** box, enter a short description of the backup. This label helps to identify the backup image should you need to restore the backup.

***Note**: You cannot use single or double quotation marks in this field.*

Define Contents of Backup

In the **Option** box, use the **Option** arrow to select what you want to back up. The available options are:

- All UniVerse Accounts — UniVerse backs up all accounts defined in UV.ACCOUNT file.
- A UniVerse Account & Subdirectories — UniVerse backs up the UniVerse account you specify. If you select this option, an **Account** box appears. Use the **Account** arrow to select the account you want to back up.
- A File in a UniVerse Account — UniVerse backs up a file in the account you specify. If you select this option, an **Account** and **File** box appear. Uses the **Account** arrow to specify the account where the file resides, then specify the file using the **File** arrow. If you only want to back up the dictionary portion of the file (and not the data), select the **Dictionary** check box.
- A Directory and Subdirectories — UniVerse backs up a UNIX Windows directory you specify. Enter the full path of the directory you want to back up, or click **Browse** to select the directory.

Start Backup

To start the backup process, click **Backup**. The **UniVerse Backup** window appears, as shown in the following example.

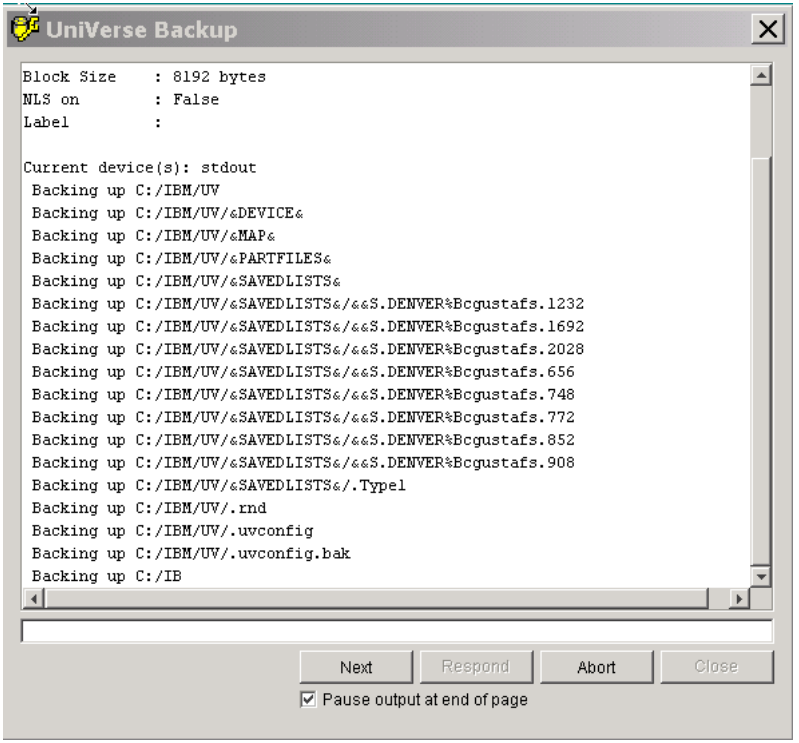

UniVerse displays messages generated during the backup process in this window.

When the backup is complete, click **Close**.

Chapter

4

Configuring UniVerse Parameters

How UniVerse Uses Configurable Parameters 4-3
Access the Configuration Editor 4-4
Changing Parameter Values 4-5
    Changing Configurable Parameters on UNIX 4-6
    Changing Configurable Parameters on Windows Platforms 4-6
UniVerse Configurable Parameters 4-7

This chapter describes how to use UniAdmin to view and change UniVerse configuration parameters.

How UniVerse Uses Configurable Parameters

When UniVerse starts, it uses the configurable parameters in the *uvconfig* file, located in the UV account directory, to specify certain UniVerse settings and limits.

When you modify *uvconfig* parameter settings, you must save the changes and restart UniVerse.

UniVerse creates the *.uvconfig* file from the *uvconfig* file using the *uvregen* command. UniVerse uses the *.uvconfig* file during the startup process. This file contains an encrypted version of the license details and current settings for configurable parameters.

On UNIX systems, the UniVerse installation process executes the *uv.rc* script, which invokes the *DBsetup* program. If you specify values for certain configurable parameters, *DBsetup* allocates the disk shared memory segment based on these values. *DBsetup* also checks for a valid serial number and expiration date, sets the limit for the number of users, and initializes the shared memory segment for UniVerse.

Access the Configuration Editor

Select one of the following methods to access the **Configuration Editor** dialog box:

- Right-click the server where you want to configure UniVerse, and then click **Configure**.
- From the **UniAdmin** window, double-click **Configuration Editor.**
- From the **UniAdmin** menu, click **Admin**, then click **Config Editor.**

Regardless of the method you chose to configure the server, a dialog box similar to the following appears:

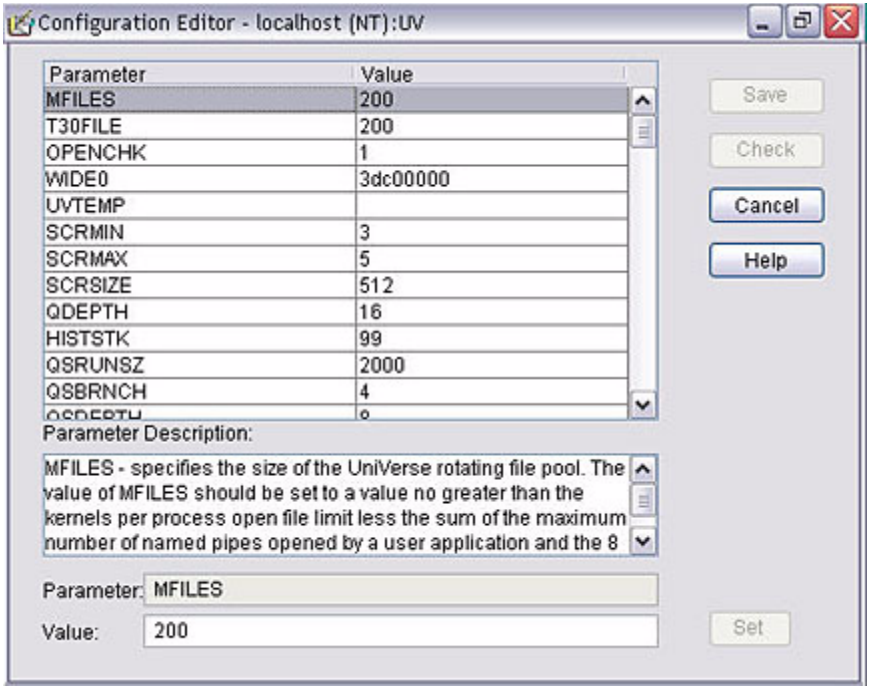

If you want to view the configuration parameters in ascending order, double-click **Parameter**. If you want to view the configuration parameters in descending order, double-click **Parameter** while holding down the SHIFT key.

Changing Parameter Values

Click the parameter for which you want information, or you want to change. A short description of the parameter appears in the **Parameter Description** box, as shown in the following example:

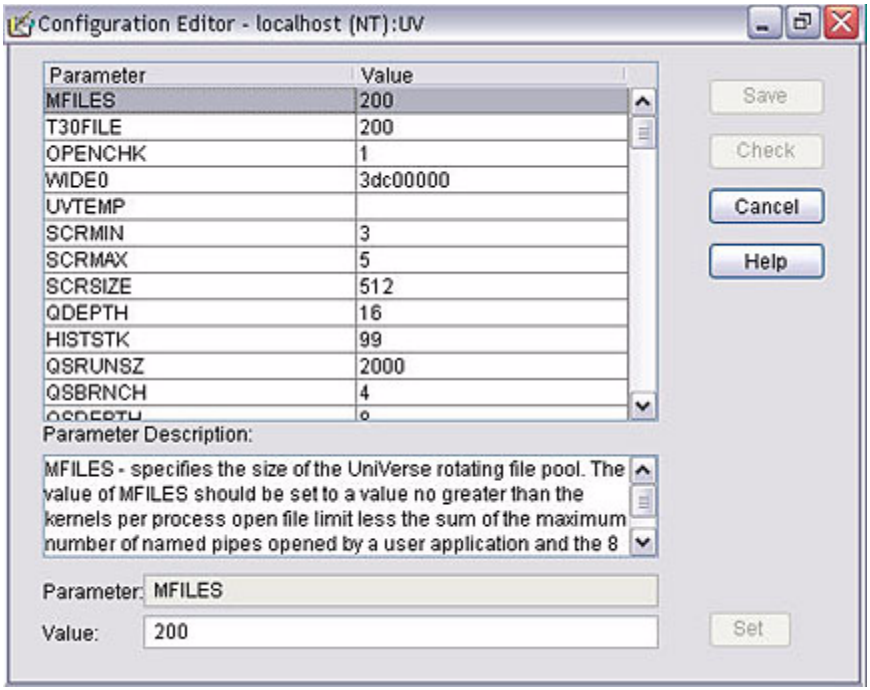

If you want to change the value for a parameter, click the parameter you want to change, enter the new value in the **Value** box, and then click **Set**.

To save the configuration, click **Save**. The file is written to the UV account file. Click **Cancel** to exit the program without saving changes.

***Note**: Changes you make to the configuration file do not take effect until you stop and restart UniVerse.*

Changing Configurable Parameters on UNIX

After you change parameters on a UNIX system, UniVerse prompts you to restart UniVerse. Choose Yes or No as a response to this prompt.

- Click **Yes** to execute the DBsetup command. This command reinitializes UniVerse with the new parameter settings. Make sure users are logged out of UniVerse before restarting.
- Click **No** to execute the DBsetup command at a later time. A second message appears to remind you to restart UniVerse manually at a more convenient time. Click **OK** to acknowledge this message.

Changing Configurable Parameters on Windows Platforms

Click **OK** to acknowledge the message UniVerse displays. Make sure users are logged out of UniVerse before restarting.

UniVerse Configurable Parameters

The following table describes the UniVerse configurable parameters. The default values may be different on your system.

Parameter	Description
64BIT_FILES	Specifies whether UniVerse uses 32-bit or 64-bit file systems. A value of 0 means UniVerse creates and resizes all UniVerse files as 32-bit files. A value of 1 means UniVerse creates and resizes all files as 64-bit files. The default value is 0.
ALLOWMARKS	The value of this parameter can be either 0 or 1. A value of 1 indicates to allow marks (SVM, VM, or AM) in record IDS for non-SQL tables. A value of 0 indicates not to allow marks in record IDs. **Warning**: Using multivalued record IDs can cause adverse effects on the database, and my corrupt index data. Reports may print improperly and the TRANS function will not work properly.
ALLOWNFS	Specifies whether UniVerse files stored on remote systems can be opened without UV/Net. A value of 0 means UV/Net is required to access remote files. Any other value allows remote files to be opened, but no locks are maintained on the remote system. Turn this parameter on when UniVerse is not running on the remote system. The default value is 0.
BGINPUTTIMEOUT	The amount of time, in seconds, a background process, including PHANTOMS, waits for input before it terminates.
BLKMAX	Sets the maximum block size for UVBACKUP and UVRESTORE. It must be greater than, and a multiple of, 512.
CENTURYPIVOT	Sets the Century Pivot year. This value determines how 1 and 2 digit years are interpreted by the ICONV command. A value of 1930 means that 30 through 90 are interpreted as the 1900s, and 00 through 29 are interpreted as 2000 - 2029. If this value is 2 digits, the Century Pivot year is based on the current year, as follows: A value of 30 means the Century Pivot year is 1930 in the year 2000, 1931 in 2001, 1932 in 2002, and so forth.

UniVerse Configurable Parameters

Parameter	Description
CONVERT_EURO	If this value is set to 0 (the default), UniVerse does not do any Euro conversion. If this value is set to 1, UniVerse performs Euro conversion.
CPLOGRLS	The amount of time, in seconds, UniVerse waits before releasing a log file. If this value is 0 (the default), the uvchkd process calls the fsync function for all data files that have updates in the current log file. If this value is greater than 0, uvchkd does not call the fsync function and waits to release the log file for the number of seconds you specify.
CSHDISPATCH	Defines the full path for the csh shell command. On UNIX systems, the default is */usr/bin/csh*. On platforms, the default is NOT_SUPPORTED.
DOSDISPATCH	Defines the full path for the DOS shell command. On UNIX systems, the default is NOT_SUPPORTED. On Windows platforms, the default is CMD.EXE.
EXACTNUMERIC	Specifies the number of digits of precision before rounding occurs. The maximum value is 57 digits. The minimum value (default) is 15 digits.
FLTABSZ	Sets the number of file lock entries in a file lock semaphore set. The default value is 11.
FSEMNUM	Sets the number of file lock semaphore sets used for concurrency control. The default value is 23.
GLTABSZ	Sets the number of group lock entries in a group lock semaphore set. The default value is 75.
GSEMNUM	Sets the number of group lock semaphore sets used for concurrency control. The default value is 97.
HISTSTK	Specifies the maximum number of sentences in a user's sentence stack. The default value is 99.
ISOMODE	Sets the SQL isoation level. See *UniVerse BASIC* for more information about possible values. The default value is 1.
JOINBUF	Specifies the size of the cache the optimizer uses for joins with explicit record IDs or indexes. The default value is 4095.

UniVerse Configurable Parameters (continued)

Parameter	Description
LAYERSEL	Determines whether a select list remains active when returning from a higher EXECUTE level. A value of 0 maintains an active select list. Any other value clears any select lists before returning to the previous layer, unless K mode is active. The default value is 0.
LOGBLNUM	Specifies the size of the log data buffer, in file system blocks. The default value is 8.
LOGBLSZ	Specifies the log buffer block size. It should be the same as the block size of the file system where the log directory is mounted. The default value is 512.
LOGSYCNT	Specifies the maximum number of commits allowed between log file syncs. If the specified value is less than 2, it is changed to 0. This parameter has no effect on a nontransactional environment. The default value is 0.
LOGSYINT	Specifies the maximum time interval allowed between log file syncs. The minimum value (other than 0) you can specify is 5. If the LOGSYCNT parameter is set to a value other than 0, LOGSYINT defaults to 120.
MALLOCTRACING	Turns on malloc tracing for UniVerse support analysts.
MAXERRLOGENT	Specifies the maximum number of entries that can be written to the *errlog* file. The default is 100.
MAXKEYSIZE	Sets the maximum number of characters that can be used within a primary key. The default value of 255 retains the historical maximum key length, as well as the historical encoding of record locks within a 64-byte entry. Any other value, which must be a multiple of 64 between 256 and 2048, extends the maximum key length to that value. In addition, UniVerse stores the full record ID within the record lock entry. This value, once raised, can only be reset by call IBM Support. DO NOT change this from the default value without understanding its effect on the record lock table entries.
MAXRLOCK	Sets the number of record locks held by an SQL transaction on a physical file (a device or an i-node) before a file lock is requested by the SQL interface.

UniVerse Configurable Parameters (continued)

Parameter	Description
MFILES	Specifies the size of the UniVerse rotating file pool. The value of MFILES must be at least 8 less than the kernel's open files per process limit. The default value is 12. Count UV/Net connections as files. Pipes are not part of the rotating file pool, but they do count as files.
MODFPTRS	Specifies whether file pointers in the VOC file can be modified by the COPY, DELETE, and EDIT commands. A value of 0 disallows modifications of VOC file pointers. The default value is 1.
NETTIME	Sets the timeout value in minutes for UV/Net. The default value is 5.
NLSDEFDEVMAP	Specifies the name of the default map to use for device input or output. This map is used for all devices, except printers, that do not have a map specified in the &DEVICE& file. The ASSIGN MAP command overrides this setting. The default value is ISO8859-1+MARKS.
NLSDEFDIRMAP	Specifies the name of the default map to use for type 1 and type 19 files without assigned maps. This occurs if a type 1 or type 19 file was not created on an NLS system, and has not had a map defined for it by the SET.FILE.MAP command. This map applies only to the data in records, not to record IDs. The default value is ISO8859-1+MARKS.
NLSDEFFILEMAP	Specifies the name of the default map to use for hashed files without assigned maps. This occurs if a hashed file was not created on an NLS system, and has not had a map defined for it by the SET.MAP.FILE command. The default value is ISO8859-1+MARKS.
NLSDEFGCIMAP	Specifies the name of the default map to use for string arguments passed to or from GCI subroutines. This map is used if the GCI subroutine did not explicitly define a map. The default value is ISO8859-1+MARKS.
NLSDEFPTRMAP	Specifies the name of the default map to use for printer output. This map is used if a printer does not have a map defined for it in the &DEVICE& file. The default is ISO8859-1+MARKS.

UniVerse Configurable Parameters (continued)

Parameter	Description
NLSDEFSEQMAP	Specifies the name of the default map to use for sequential input or output for files or devices without assigned maps. The SET.SEQ.MAP command overrides this setting. The default value is ISO8859-1+MARKS.
NLSDEFSOCKMAP	The name of the map to associate with sockets that are either explicitly created through UniVerse BASIC APIs, or implicitly created through other APIs, such as CallHTTP.
NLSDEFSRVLC	Specifies the name of the default locale to use for passing data to and from client programs. This locale is used if the client program does not specify a server locale. The default value is ISO8859-1+MARKS.
NLSDEFSRVMAP	Specifies the name of the default map to use for passing data to and from client programs. This map is used if the client program does not specify a server map. The default value is ISO8859-1+MARKS.
NLSDEFTERMMAP	Specifies the name of the default map to use for terminal input or output. This map is used if a terminal does not have a map defined for it in its *terminfo* definition. The SET.TERM.TYPE command overrides this setting. The default value is ISO8859-1+MARKS.
NLSDEFUSRLC	Specifies the default locale. The default value is OFF.
NLSLCMODE	Specifies whether locales are enabled. A value of 1 indicates that locales are enabled, a value of 0 indicates that locales are disabled. The default value is 0. This parameter has no effect unless NLSMODE is set to 1.
NLSMODE	Turns NLS mode on or off. A value of 1 indicates NLS is on, a value of 0 indicates NLS is off. If NLS mode is off, UniVerse does not check any other NLS parameters.
NLSNEWDIRMAP	Specifies the name of the map to use for new type 1 and type 19 files created when NLS mode is on. This map applies only to the data in records, not to record IDs. The default value is ISO8859-1+MARKS.

UniVerse Configurable Parameters (continued)

Parameter	Description
NLSNEWFILEMAP	Specifies the name of the map to use for new hashed files created when NLS mode is on. A value of NONE (the default value) indicates that the data is to be held in the internal UniVerse character set.
NLSOSMAP	Specifies the name of the map to use for file names or record IDs visible to the operating system. This chiefly affects CREATE.FILE and record ID written to type 1 or type 19 files. The default value is ISO8859-1.
NLSREADELSE	Specifies the action to take if characters cannot be mapped when a record is read by a READ statement. A value of 1 indicates that the READ statement takes the ELSE clause. A value of 0 indicates that unmappable characters are returned as the Unicode replacement character 0xFFFD. The default value is 1.
NLSWRITEELSE	Specifies the action to take if characters cannot be mapped when data is written to a record. A value of 1 indicates that the write aborts or takes the ON ERROR clause (if there is one). A value of 0 indicates that unmappable characters are converted to the file map's unknown character (for example, ?) before writing the record. When this happens, some data may be lost.
OCVDATE	Specifies whether UniVerse accepts partially bad internal dates. A value of 0 rejects any value that is not wholly numeric and supplied as an internal date to the ID conversion code; the date is not converted. Any value other than 0 accents a number followed by a nonnumeric data (for example, 9199-f); the number is treated as an internal date, and the STATUS function is set to 3. The default value is 0.
OPENCHK	Modifies the behavior of operations on files opened with the BASIC OPEN statement. When set to 0, no integrity constraints are observed. This parameter does not affect files opened using the OPENCHECK statement. The default value is 1.
OPTMEM	Specifies the amount of memory allocated for the query optimizer's workspace. This is specified by in 1K units. The default value is 64.
PAKTIME	Specifies the number of seconds the system waits at the Press Any Key to Continue message before releasing a pending group lock. The default value is 300.

UniVerse Configurable Parameters (continued)

Parameter	Description
PHANTOMSAMEUV	The value of this parameter can be either 0 or 1. If you specify 1, UniVerse enables the phantom using the same uv executable. A value of 0, the default, uses UVHOME/bin/uv.
PICK_MT	The value of this parameter can be either 0 or 1. 1 indicates the MT conversion works the same as the Pick flavor in INFORMATION. A value of 0 indicates UniVerse MT conversion handling.
PICKDATE	The value of this parameter can be either 0 or 1. A value of 1 indicates that UniVerse will accept Pick-style dates, and handle data in the same way as Pick. A value of 0 indicates the default UniVerse date format of yyyymmdd.
PICKNULL	Sets the masked decimal conversion for empty data. A value of 1 turns on Pick-style conversions, where empty data is converted to an empty string. A value of 0 indicates UniVerse-style conversions, where empty data is converted to 0.00.
PIOPENDEFAULT	Changes the default behavior of PIOpen-flavor accounts. When the value of this parameter is 0, UniVerse maintains the current functionality. When the value of this parameter is a nonzero value, the following defaults change: n $OPTIONS INFO.CONVERT is ON when UniVerse BASIC programs are recompiled. n $OPTIONS PIOPEN.EXECUTE is ON when UniVerse BASIC programs are recompiled. n RetrieVe queries in all flavors abort if the FROM list keyword references an inactive SELECT list.
PKRJUST	A value of 1 gives Pick-style right-justified behavior for LIST and SORT in all flavors. Pick-style right-justified behavior can overwrite data in previous columns if the data exceeds the column or the column header width. The default value is 0.
PROCACMD	Defines the action of the ProVerb A command. A value of 0 specifies that the A command quits when *m* characters are moved, or when a field mark or the end of the input buffer is reached. A positive nonzero value causes the A command to ignore the field mark. The default value is 0.

UniVerse Configurable Parameters (continued)

Parameter	Description
PROCRCMD	Determines the behavior of the ProVerb RI command. A value of 0 clears the input buffer and removes the preceding field mark. Any other value does not remove the field mark. The default value is 0.
PROCPRMT	Determines the effect of the UniVerse PROMPT keyword on the ProVerb prompt. A value of 0 lets the PROMPT keyword change the ProVerb prompt. Any other value retains the ProVerb prompt, which can be changed only by the IP ProVerb command. The default value is 0.
PSEMNUM	Sets the number of BASIC user process control locks. The default value is 64.
QBREAK	Selects the function of the keys **Q** and **Ctrl-X** at the Press Any Key to Continue message. A value of 0 means that the **Q** and **Ctrl-X** keys will be ignored while in BREAK OFF mode. A nonzero value allows the **Q** and **Ctrl-X** keys to quit at the Press Any Key to Continue message, even in BREAK OFF mode. There are security implications to selecting the latter mode of operation. The default value is 1.
QDEPTH	Specifies the maximum depth of nesting allowed in Q-pointer references. The default value is 16, and the minimum value is 0.
QSBRNCH	Specifies the number of runs which cause a sub-merge to be performed by the query processor's sorting algorithm. This is sometimes referred to as the branching factor of the sort. The efficiency of the sorting algorithm is very sensitive to the value of QSBRNCH. The default value is 4, and the minimum value is 2.
QSDEPTH	Specifies the maximum depth of the sort tree used by the query processor's sorting algorithm. The efficiency of the sorting algorithm is very sensitive to the value of QSDEPTH. The default is 8, and the minimum value is 2.
QSMXKEY	Specifies the maximum number of sort key components. The efficiency of the sorting algorithm is very sensitive to the value of QSMXKEY. The default is 32.

UniVerse Configurable Parameters (continued)

Parameter	Description
QSRUNSZ	Specifies the size of the initial sorting run used by the query processor's sorting algorithm. The efficiency of the sorting algorithm is very sensitive to the value of QSRUNSZ. The default value is 2000, and the minimum value is 2.
RLOWNER	Sets the number of lock owner entries maintained for shared record locks in a group semaphore set. The default value is 300.
RLTABSZ	Sets the number of update record lock entries in a group lock semaphore set. The default value is 75.
SCRMAX	Specifies the maximum size of the UniVerse scratch buffer pool. SCRMAX must be larger than SCRMIN, and must be specified after SCRMIN. The default value is 5.
SCRMIN	Specifies the minimum size of the UniVerse scratch buffer pool. SCRMIN must be at least 1 and must be specified before SCRMAX. The default value is 3.
SCRSIZE	Specifies the initial size of a scratch buffer. SCRSIZE must be from 512 through 2048. The default value is 512.
SELBUF	Specifies the size of the in-memory select list buffer. It is the amount of locally cached select data which can be stored before the select list starts using disk storage. It is specified in 1K units. The default value is 4.
SHDISPATCH	Defines the full path for the *sh* shell command. On UNIX systems, the default is */usr/bin/sh*. On Windows platforms, the default is NOT_SUPPORTED.
SMISDATA	Specifies whether a segment mark found in data should be treated as a data character (1) or a data terminator (0). This parameter affects data access in the file subsystem, as well as UniVerse BASIC functions such as EXTRACT and REPLACE when working with a dynamic array that contains one or more segment marks. If set to 0, UniVerse inserts data containing a segment mark, but truncated at the mark when retrieved.

UniVerse Configurable Parameters (continued)

Parameter	Description
SPINSLEEP	On an HP multiprocessor system, when spinlocks are in use, this value determines the time to sleep between successive attempts to lock a spinlock. The value, expressed in micro-seconds, should be between 50 and 50000. If spinlocks are in use, and the value of this parameter if 0, UniVerse defaults to a value of 5000. Misuse of this tunable can drastically affect system performance.
SPINTRIES	On an HP multiprocessor system, this value determines whether spin locking is used instead of regular semaphore locking. This value determines the number of attempts to obtain the spinlock before the process sleeps. If you want to use spinlocks, choose a value between 5 and 500. A value of 0 switches spinlocks OFF. Misue of this parameter can drastically affect system performance.
SQLNULL	You can assign this parameter any NULL value the system wants to use. The default value for SQLNULL is 128, which interferes with the EURO symbol value on Windows platforms. Change this value to another value if you are using EURO symbols.
SYNCALOC	A value of 1 causes creation of new UniVerse files to occur as soon as they are requested. The default value is 1.
SYSTEM_EURO	You can set this value to a system Euro codepoint value. The default is 128 on Windows platforms, and 164 on UNIX systems.
T30FILE	Specifies the number of dynamic files that can be opened. This is used to allocate shared memory concurrency control headers. The default value is 200.
TERM_EURO	This parameter can be set to a terminal Euro codepoint value. The default is 128 on Windows platforms, and 164 on UNIX systems.
THDR512	Specifies whether DR-type tapes are written with 512-byte labels. The default value is 0.
TSTIMEOUT	Sets the amount of time the UniVerse device licensing shell (uvdls) waits for a connection from a telnet client (wIntegrate or SBClient). Any value less than 1 defaults to 60.

UniVerse Configurable Parameters (continued)

Parameter	Description
TXMEM	Specifies the amount of memory allocated for the private transaction cache. This is specified in 1K units (1024 bytes). The default is 32.
TXMODE	Sets the transaction mode observed by the system. When set to 0, transactions are not logged by the log daemon. The default value is 0.
ULIMIT	Sets the maximum file size set by UniVerse. UniVerse uses either the value set by ULIMIT or the value set by the UNIX *ulimit*, whichever is larger. The default value of ULIMIT is 128000.
UDRBLKS	Specifies the size of the internal buffer used for caching replicated data before it is written to the log file. One block is equivalent to 4096 bytes. The minimum size is 10, the maximum size is system-dependent. A larger size can improve performance on larger systems. The default size is 10.
UDRMODE	Sets the data replication mode. When set to 0, replication is not activated. When set to 1, replication is activated. The default value is 0.
UVNET_CONNECT	The value of this parameter is either 0 or 1. A value of 1 implies to connect to the remote machine regardless of EURO data settings. A value of 0 implies communication takes place only if EURO settings match on both machines.
UVTSORT	The value of this parameter can be either 1 or 0. A value of 1 enables multithreaded sort. A value of 0 disables multithreaded sort.
UVSPOOL	Specifies the name of the directory to be used as the UniVerse spooler directory. This should be a fully qualified path of 112 characters or less.
UVSYNC	Determines if UniVerse uses the UNIX *sync()* call. A nonzero value allows a *sync()* to be performed if a leading process exits. Data loss can occur if *sync()* is not executed often enough.

UniVerse Configurable Parameters (continued)

Parameter	Description
UVTEMP	Specifies the name of the directory used to contain UniVerse select lists and other temporary files. This should be a fully qualified path of 112 characters or less.
VDIVDEF	Selects the default action of the vector divide operator when the divisor values are exhausted prematurely. A nonzero value returns the dividend. A zero value returns 0. The default value is 1.
WIDE0	Specifies the mask used internally to decide when the difference between two numeric values is to be considered 0. The default value is 0x3dc00000. For more information about the wide zero parameter, see *Administering UniVerse*.

UniVerse Configurable Parameters (continued)

Chapter

5

Administering Data Sources

Administering Data Sources 5-3
Creating a New Data Source 5-4
Viewing Data Source Details 5-6
Deleting a Data Source 5-7

This chapter describes how to administer data sources through UniAdmin. You can view and change the current set of defined UCI and UniVerse ODBC data sources.

Administering Data Sources

Select one of the following methods to access the UniVerse Data Sources dialog box:

- From the UniData Admin window, double-click **Data Sources**.
- From the UniAdmin menu, click **Admin**, then click **Data Sources**.
- From the UniAdmin toolbar, click the **Data Sources** icon, as shown in the following example.

The **UniVerse Data Sources** dialog box appears, as shown in the following example.

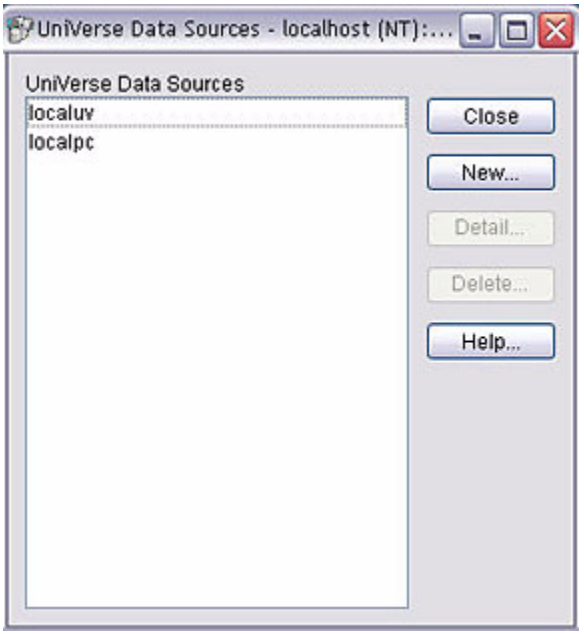

Creating a New Data Source

To create a new data source, click **New**. A dialog box similar to the following example appears.

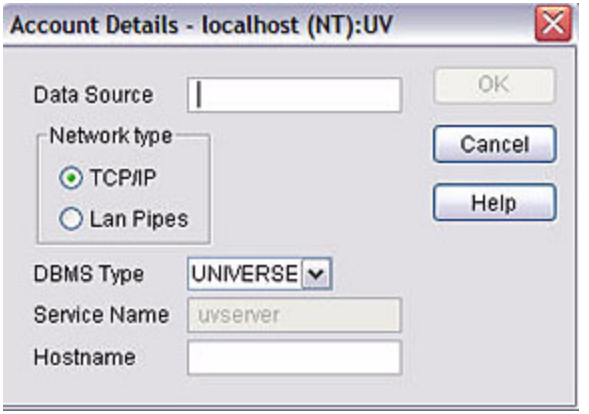

Enter Data Source Name

In the **Data Source** box, enter the logical name of the data source, stored in the uvodbc.config file.

Enter Network Type

The network type is the type of communication transport that the data source uses to communicate with the server. Click either TCP/IP or *LAN Pipes*.

***Note**: You can only use LAN pipes with Windows servers.*

Select the DBMS Type

Select the name of the database management system on the server. You can select UniVerse or ODBC.

Service Name

The Service Name is the name of the UniVerse server. You cannot change this field.

Enter the Host Name

Enter the name or IP address of the server on which the data source resides.

Click **OK** to save your entries, or **Cancel** to exit without saving changes.

Viewing Data Source Details

To view the details of an already established data source, select the data source for which you want to view details, then click **Details**. A dialog box similar to the following example appears:

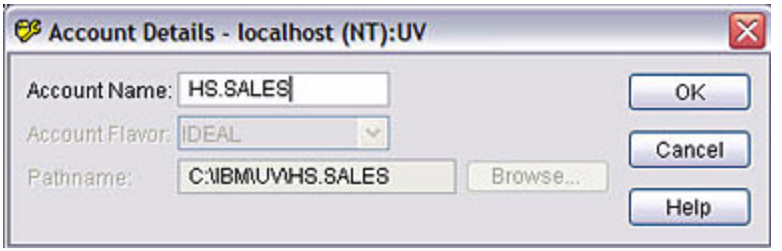

You can make changes to the data source through this dialog box. After making your changes, click **OK** to save them, or click **Cancel** to exit without saving them.

Deleting a Data Source

To delete a data source, select the data source you want to delete from the UniVerse Data Sources list, then click **Delete**. The following dialog box appears to confirm that you want to delete the data source.

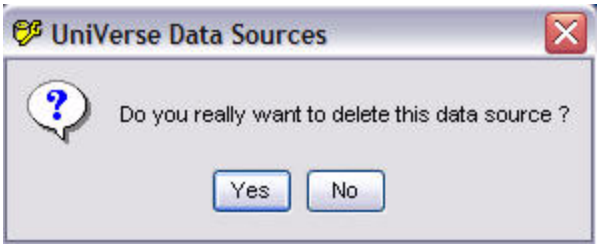

Click **Yes** to delete the data source, or click **No** to exit the program without saving changes.

Chapter

6

Administering Deadlocks

Managing Dead Locks 6-3
Starting and Stopping the Dead Lock Manager 6-4
Configuring the Deadlock Manager 6-5
View Pending Deadlocks 6-7

This chapter describes how to manage dead locks through UniAdmin. You can start or stop the deadlock manager, configure the deadlock manager, or manually resolve file locking conflicts from the **UniVerse Dead Lock Admin** dialog box.

When the deadlock system is running on the server, UniVerse automatically resolves dead locks. The deadlock manager writes a log file that records all dead locks that are automatically resolved.

Managing Dead Locks

Select one of the following methods to access the **UniVerse Dead Lock Admin** dialog box:

- From the UniData Admin window, double-click **Dead Locks**.
- From the UniAdmin menu, click **Admin**, then click **Dead Locks**.
- From the UniAdmin toolbar, click the **Manage UniVerse Dead Locks** icon, as shown in the following example.

UniVerse Dead Locks Manager

The **UniVerse Dead Lock Admin** dialog box appears, as shown in the following example.

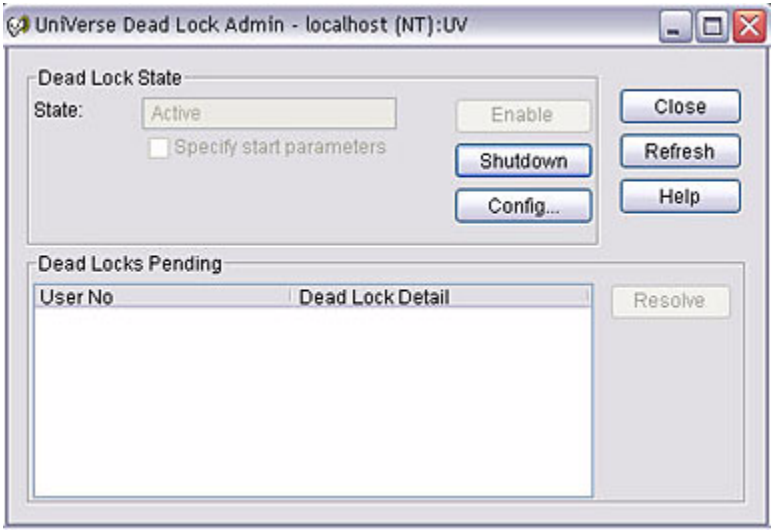

Starting and Stopping the Dead Lock Manager

To start or stop the Dead Lock Manager, click **Enable**.

To shutdown the Dead Lock Manager, click **Shutdown**.

Configuring the Deadlock Manager

To configure the Deadlock Manager, click **Config**. A dialog box similar to the following example appears:

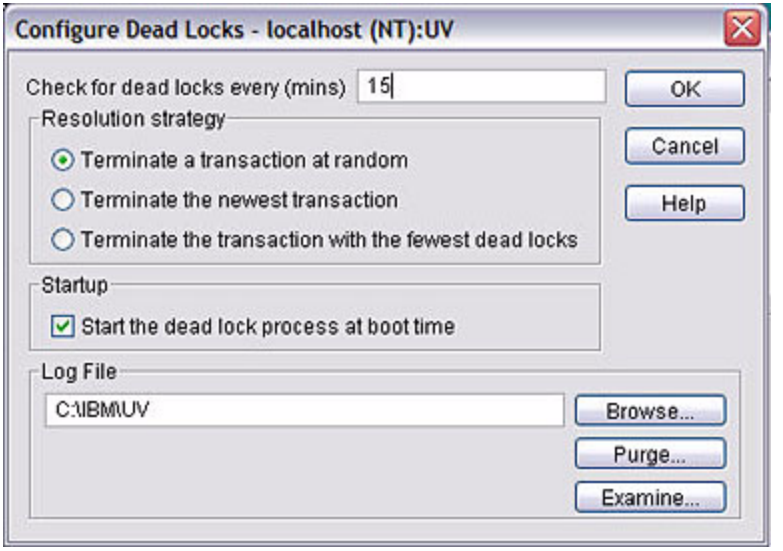

Determine Checking Interval for Dead Locks

In the **Check for dead locks every (mins)** box, enter the number of minutes the deadlock manager process waits before checking for deadlock conditions.

Determine Resolution Strategy

Determine how to automatically resolve a conflict by selecting a resolution strategy. The following options are available:

- Terminate a transaction at random
- Terminate the newest transaction
- Terminate the transaction with the fewest dead locks

Determine When to Start the Deadlock Manager

Select the **Start the dead lock process at boot time** check box to start the Deadlock Manager when the server boots. If you do not select this check box, you must enable the Deadlock Manager manually.

Define Location of Log File

In the **Log File** box, enter the directory of the log file to track dead lock events, or click **Browse** to select the directory for an existing log file.

Click **Purge** to clear an existing log file.

Click **Examine** to view the contents of the log file.

Click **OK** to implement your changes, or click **Cancel** to exit the dialog box without saving changes.

View Pending Deadlocks

The Dead Lock Administration window displays information about pending deadlocks.

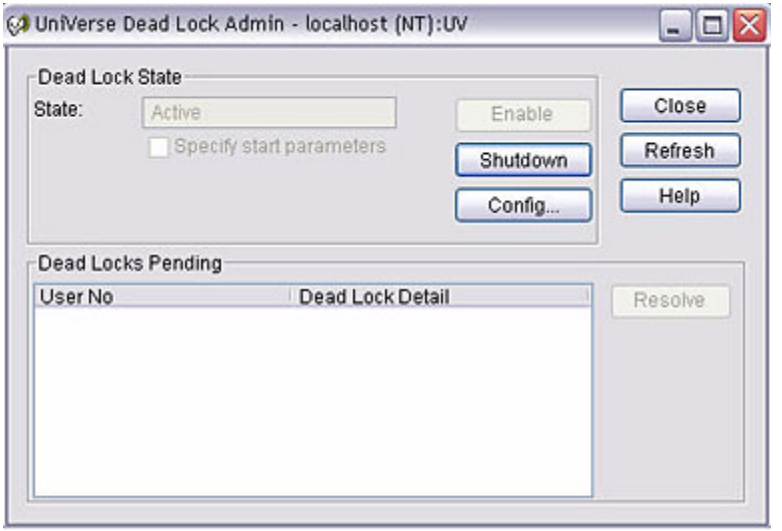

The **Dead Locks Pending** window displays the current dead lock processes.

To resolve a dead lock, select a process and click **Resolve**.

*Note: You can only resolve a dead lock manually when the Deadlock Manager is not running. If you want to resolve a dead lock manually, shut down the Deadlock Manager, then click **Refresh**.*

Chapter

7

Administering Devices

Administering Tape Devices on UNIX 7-3
Defining a New Tape Device 7-3
Viewing or Modifying a Tape Device 7-8
Deleting a Tape Device 7-9
Administering Tape Devices on Windows Platforms 7-10
Defining a New Tape Device 7-10
Viewing or Modifying a Tape Device 7-14
Deleting a Tape Device 7-15
Using Test Tape 7-15
Administering Print Devices on UNIX 7-16
Defining a New Printer 7-17
Viewing or Modifying a Printer Definition 7-21
Deleting a Print Device Definition 7-22
Administering Print Devices on Windows Platforms 7-23
Configuring Other Devices 7-24
View Details about a Device 7-25
Delete a Device Definition 7-26

This chapter describes how to configure and administer tape devices, printers, and line devices through UniAdmin. Select one of the following methods to access the **UniVerse Maintain Devices** dialog box:

- From the **UniData Admin** window, double-click **Devices**.
- From the **UniAdmin** menu, click Admin, then click **Devices**.
- From the **UniAdmin** toolbar, click the **Manage Devices** icon, as shown in the following example.

Manage Devices Icon

The **Maintain UniVerse Devices** dialog box appears, as shown in the following example:

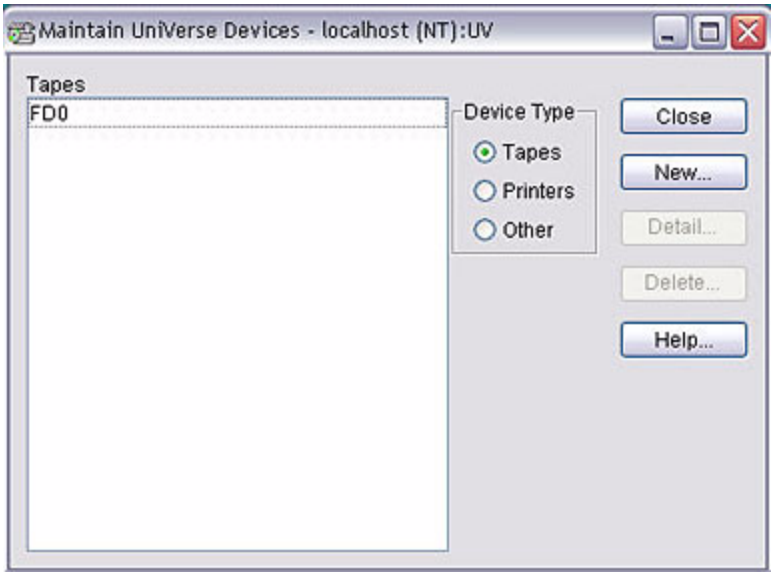

Administering Tape Devices on UNIX

From the **Maintain UniVerse Devices** dialog box, select **Tapes**. Any tape devices that have been defined in the &DEVICE& file appear in the **Tapes** window.

Defining a New Tape Device

To define a new tape device, click **New**. The **Tape Details** dialog box, similar to the following example, appears:

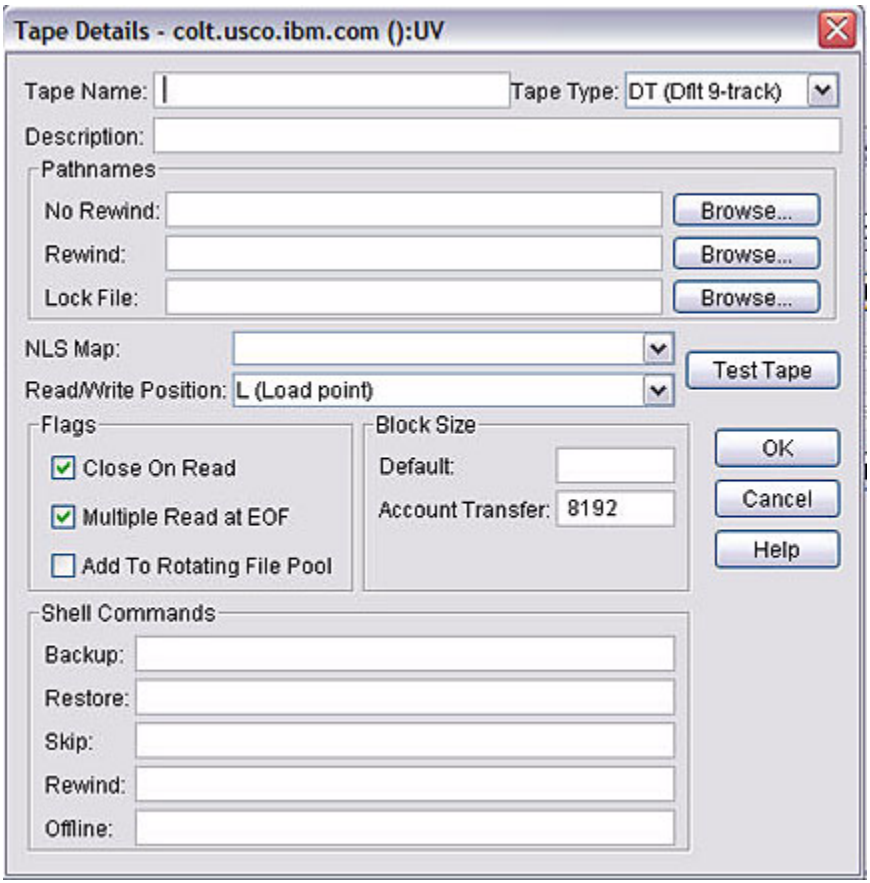

Enter Tape Name

Enter the name of the new tape drive in the **Tape Name** box. UniVerse uses this name for various UniVerse commands, such as ASSIGN, to refer to the device.

Select Tape Type

From the **Tape Type** list, select the type of tape drive. Valid types are:

- DT (Dflt 9-track)
- DC (Dflt cartridge)
- T (9-track)
- C (Cartridge)
- F (Floppy)

***Note**: Universe supports multireel tape handling for the T.DUMP and T.LOAD commands for device types DC, DT, and F only.*

Enter a Description

If desired, enter a description of the tape in the **Description** box.

Enter No Rewind Device Path

In the **Pathnames** area, enter the full path to the no rewind device in the **No Rewind** box, or click **Browse** to select the location.

Enter Rewind Device Path

In the **Pathnames** area, enter the full path to the rewind tape device in the **Rewind** box, or click **Browse** to select the location. You cannot use the same path used for the no rewind device.

***Note**: Make sure you assign the correct access permissions to the devices you specify. To set permissions for the device, use the UNIX chmod command.*

Enter Lock File

When UniVerse and UNIX system processes share a device, a special lock file is needed that coordinate access to the device when more than one process attempts to gain access to the device. The **Lock File** box contains the UNIX path used to implement the locking protocol used by the UniVerse spooler and UNIX facilities, such as the spooler and *uvcp*. For tape devices, this box is usually empty, but can be used to display ownership information. For more information about the form of the lock file name for a system, see the UNIX reference manual for the process that is sharing the device.

If desired, enter the name of a lock file to use in the **Lock File** box.

Specify NLS Map

If your system is NLS enabled, enter the name of a character set map for the device in the **NLS Map** box. For information about NLS maps, see the *UniVerse NLS Guide*.

Define Read/Write Position

Select the setting for the read/write position from the **Read/Write Position** list. This setting specifies the position on the tape where changing from read to write mode is allowed. Valid options are:

- L (Load point). This is the default position.
- E (Load point or EOF)
- A (Anywhere)

Note**: UniAdmin automatically sets the Read/Write position when you click **Test Tape.

Set Flags

You can select one of the following check boxes to meet your requirements:

- **Close on Read** – The setting of **Close on Read** determines the action taken when a tape was opened for reading and is then closed. If you select **Close on Read**, the default setting, the tape moves forward to the beginning of the next file. Most 1/4-inch tape devices use **Close On Read**. If you clear **Close on Read**, the tape does not move forward.

  Close on Read is reset to a suitable setting when you click **Test Tape**.

  ***Note**: Most Berkeley device drivers work with this flag turned off. Most System V device drivers work with this flag turned on.*

- **Multiple Read at EOF** – The setting of **Multiple Read at EOF** determines the behavior of the tape when it reaches an end-of-file mark. If you select **Multiple Read at EOF**, the default setting, the second read also returns an end-of-file. If you clear **Multiple Read at EOF**, the second read reads the next block/record after the end-of-file mark. Most 1/4-inch tape devices use **Multiple Read at EOF**.

  ***Note**: Most Berkeley device drivers work with this flag turned off. Most System V device drivers work with this flag turned on.*

  Multiple Read at EOF is automatically updated with a suitable setting when you click **Test Tape**.

- **Add to Rotating File Pool** – If you select the **Add to Rotating File pool** option, the device is included in the rotating file pool.

Define Block Size

In the **Block Size** area, enter the default block size in the **Default** box. This block size is used for normal tape operations, and needed only if the device is for cartridge tape (types DC and DT), or diskette (type F).

***Warning**: For diskettes, the default block size is 500. Do not change this setting. Any other block size may cause unpredictable results.*

If the tape device is a cartridge (type C), the block size value must be a multiple of 512.

For 9-track tapes (types T and DT), there is no default block size for IDEAL and INFORMATION flavor accounts. UniVerse reads or writes tape records with variable length. If the block size box is empty, UniVerse uses 8192 as the default block size for PICK and REALITY flavor accounts.

If you assign the device using the ASSIGN command, UniVerse uses the block size from the &DEVICE& file. If you assign the device using the T.ATT command, UniVerse uses the block size you specify in the **Default** box, unless it is empty. In this case, UniVerse uses the block size defined in the VOC entry for the T.ATT command.

Define Block Size for Account Transfer

To define the block size to use when importing accounts, click the **Account Transfer** box, and enter the appropriate block size.

Define UNIX Commands

If desired, enter appropriate UNIX shell commands in the Backup, Restore, Skip, Rewind, and Offline boxes. UniAdmin automatically updates the Backup and Restore commands if you use Test Tape.

Use Test Tape

You can run the tape device test program by clicking **Test Tape**. This program determines the following settings:

- What action to take when a tape file opened for reading is closed.
- Where to allow the mode to change from read to write mode.
- If a second read call at the end-of-file mark should return the end-of-file condition again.

UniAdmin automatically updates the appropriate fields with the results of the tape test program. If you do not run the test, you can enter your own values for those fields.

Click **OK** to save your changes and update the &DEVICE& file, or click **Cancel** to exit the dialog box without saving changes.

Viewing or Modifying a Tape Device

If you want to view details about a previously defined tape device, select the tape device in **Tapes** area of the **Maintain UniVerse Devices** dialog box, then click **Detail**. A dialog box similar to the following example appears:

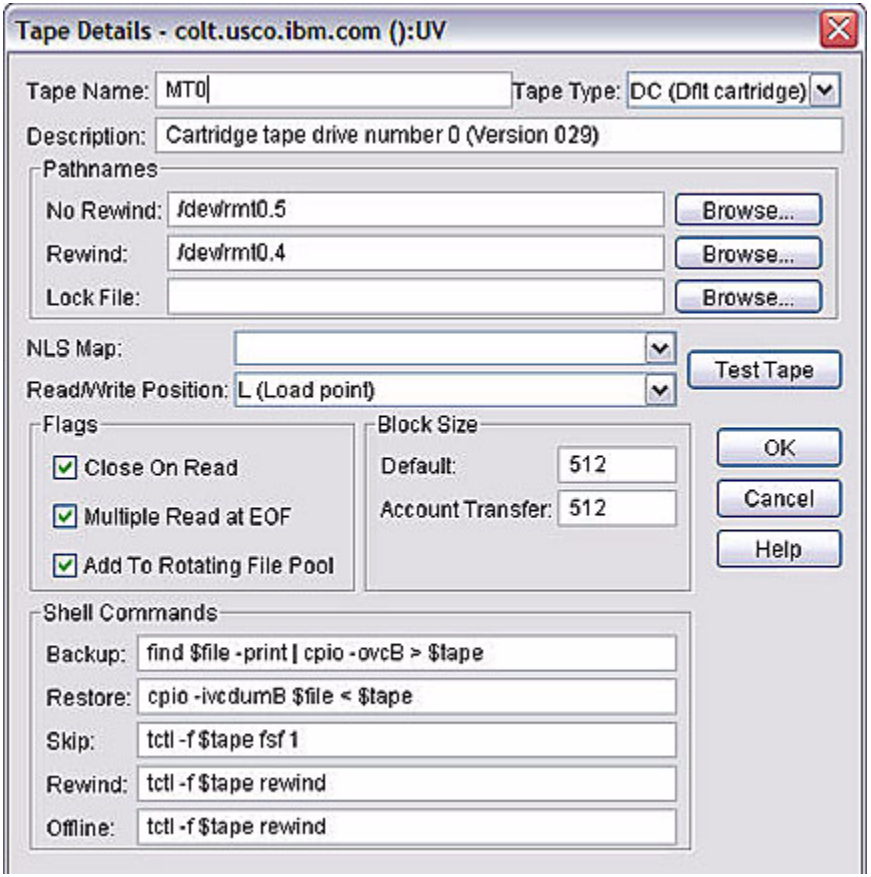

If you want to change a definition for the tape device, make the necessary changes, then click **OK** to update the tape device detail, or click **Cancel** to exit without saving changes.

Deleting a Tape Device

To delete a tape device, select the device in the **Tapes** area of the **Maintain UniVerse Devices** dialog box, then click **Delete**. A dialog box similar to the following example appears:

If you want to delete the tape device definition, click **Yes**. The tape device definition will be removed from the &DEVICE& file. If you want to exit without deleting the tape device definition, click **No**.

Administering Tape Devices on Windows Platforms

From the **Maintain UniVerse Devices** dialog box, select **Tapes**. Any tape devices that have been defined in the &DEVICE& file appear in the **Tapes** window.

Defining a New Tape Device

To define a new tape device, click **New**. A **Tape Details** dialog box similar to the following example appears:

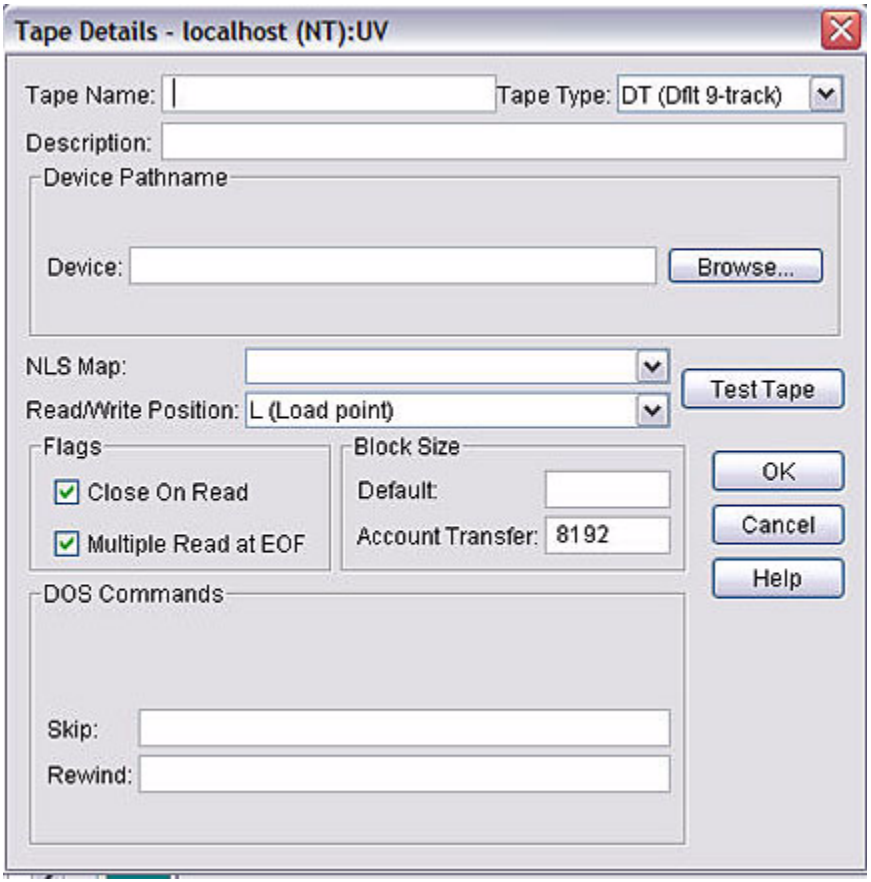

Enter Tape Name

Enter the name of the new tape drive in the **Tape Name** box. UniVerse uses this name for various UniVerse commands, such as ASSIGN, to refer to the device.

Select Tape Type

From the Tape Type list, select the type of tape drive. Valid types are:

- DT (Dflt 9-track)
- DC (Dflt cartridge)
- T (9-track)
- C (Cartridge)
- F (Floppy)

***Note**: Universe supports multireel tape handling for the T.DUMP and T.LOAD commands for device types DC, DT, and F only.*

Define Device Path

In the **Device** box, enter the path for the tape device, for example, \\.\tape0. Click **Browse** to select the location.

NLS Map

Select the NLS Map you want to use for this tape device from the NLS Map list.

Define Read/Write Position

Select the setting for the read/write position from the **Read/Write Position** list. This setting specifies the position on the tape where changing from read to write mode is allowed. Valid options are:

- L (Load point). This is the default position.
- E (Load point or EOF)
- A (Anywhere)

***Note**: UniAdmin automatically sets the Read/Write position when you click **Test Tape**.*

Set Flags

You can select one of the following check boxes to meet your requirements:

- **Close on Read** – The setting of **Close on Read** determines the action taken when a tape was open for reading and then is closed. If you select **Close on Read**, the default setting, the tape moves forward to the beginning of the next file. If you clear **Close on Read**, the tape does not move forward.

  Close on Read is reset to a suitable setting when you click **Test Tape**.

- **Multiple Read at EOF** – The setting of **Multiple Read at EOF** determines the behavior of the tape when it reaches an end-of-file mark. If you select **Multiple Read at EOF**, the default setting, the second read also returns an end-of-file. If you clear **Multiple Read at EOF**, the second read reads the next block/record after the end-of-file mark.

  Multiple Read at EOF is automatically updated with a suitable setting when you click **Test Tape**.

Define Block Size

To define the block size for the tape device, click the **Default** box in the **Block Size** area. This is the block size to use for normal tape operations. If the tape drive is a cartridge (C), this value must be a multiple of 512. A block size of 0 (variable) is normally used for DT and T type tape drives.

Define Block Size for Account Transfer

To define the block size to use when importing accounts, click the **Account Transfer** box, and enter the appropriate block size.

Specify Skip and Rewind Commands

Enter the DOS Skip and Rewind commands in the DOS commands, if necessary. The following example illustrates a skip command:

```
uvmt -d\\.tape0 -b512 fskip
```

Click **OK** to save your changes and update the &DEVICE& file, or click **Cancel** to exit the dialog box without saving changes.

Viewing or Modifying a Tape Device

If you want to view details about a previously defined tape device, select the tape device in **Tapes** area of the **Maintain UniVerse Devices** dialog box, then click **Detail**. A dialog box similar to the following example appears:

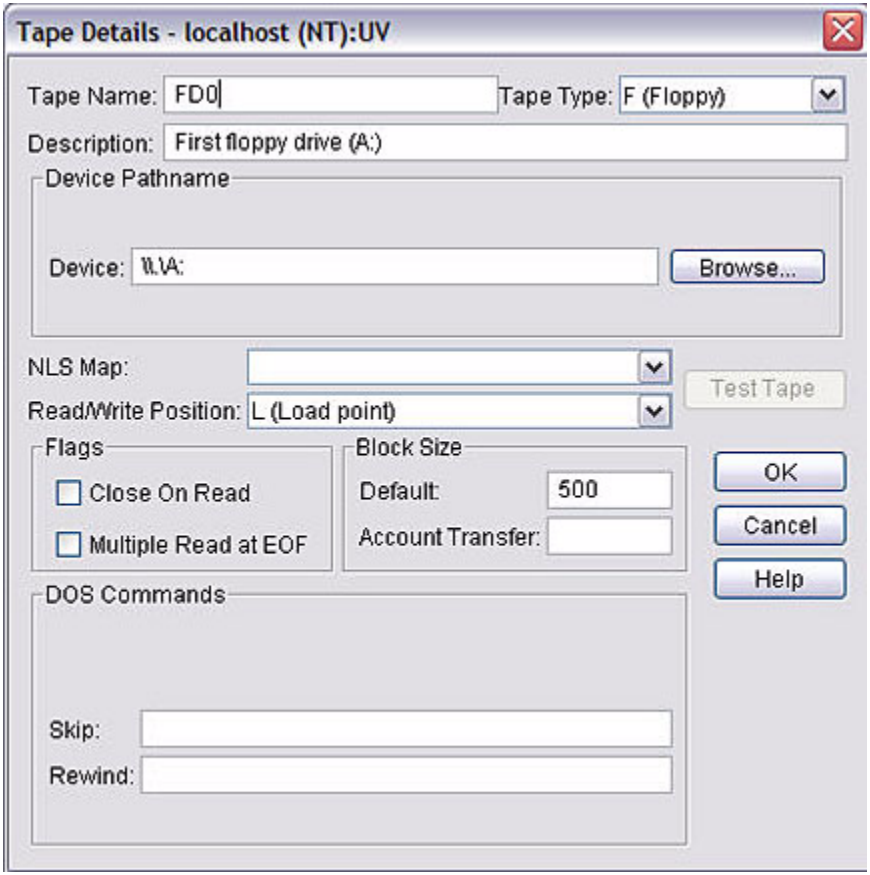

If you want to change a definition for the tape device, make the necessary changes, then click **OK** to update the tape device detail, or click **Cancel** to exit without saving changes.

Deleting a Tape Device

To delete a tape device, select the device in the **Tapes** area of the **Maintain UniVerse Devices** dialog box, then click **Delete**. The following dialog box appears:

If you want to delete the tape device definition, click **Yes**. The tape device definition will be removed from the &DEVICE& file. If you want to exit without deleting the tape device definition, click **No**.

Using Test Tape

You can use **Test Tape** with tape types DT and DC. The test tape process provides some of the tape drive definition details.

To start the tape tests, mount the tape and click **Test Tape**. A dialog box appears with suggested setting for the following tape device definitions:

- Read/Write Position
- Close on Read
- Multiple Read at EOF

The values for the above definitions are automatically updated when the tests complete, overwriting previous settings. Click **OK** to accept the new definitions, or click **Cancel** to exit without saving changes.

Administering Print Devices on UNIX

From the **Maintain UniVerse Devices** dialog box, select **Printers**. Any printers that have been defined in the &DEVICE& file appear in the **Printers** window.

Defining a New Printer

To define a new printer, from the **UniVerse Maintain Devices** dialog box, click **New**. The **Printers Detail** dialog box appears, as shown in the following example:

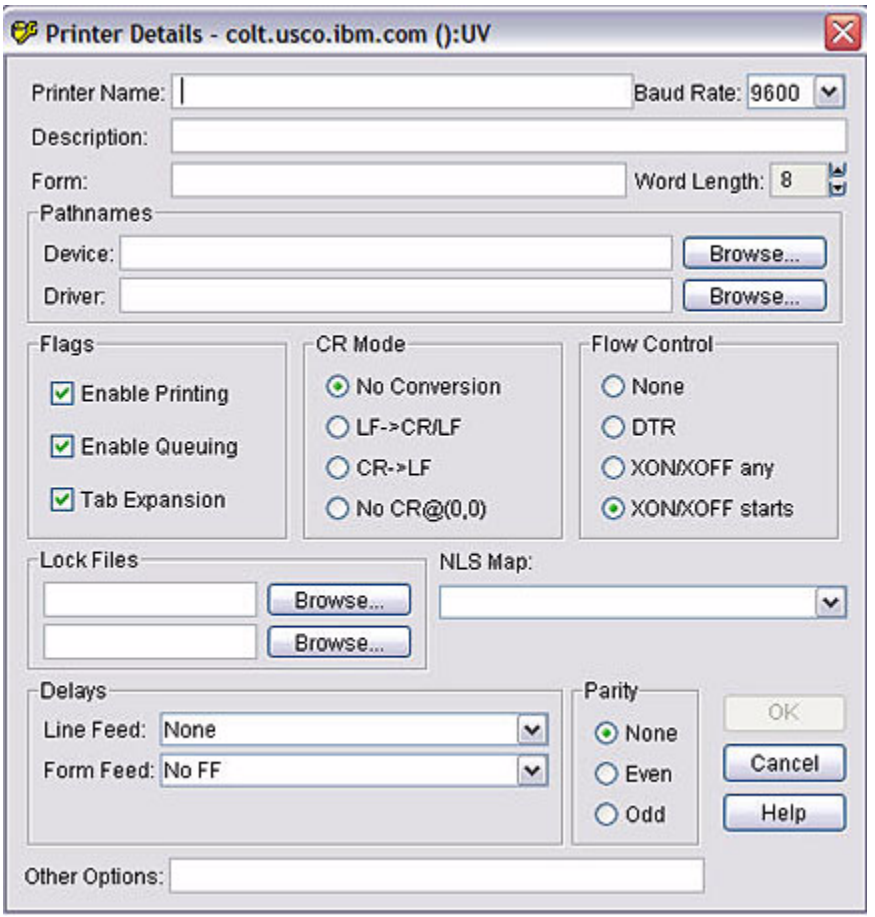

Enter Printer Name

Enter the name of the printer in the **Printer Name** box.

Select Baud Rate

Select the appropriate baud rate from the **Baud Rate** list.

Enter Description

If you desire, enter a description of the printer in the **Description** box.

Enter Form Name

To mount a form on the printer, enter a form name in the **Form** box.

Define Word Length

Select the appropriate word length from the **Word Length** list. The word length defines the number of data bits that make up a word, not including the parity bit. Use the arrows to scroll between values. The minimum value is 5 bits, and the maximum value is 8 bits.

Specify Printer Path

In the **Pathnames** area, enter the full path of the print device in the **Device** box, or click **Browse** to select the location.

Specify Driver Path

In the **Pathnames** area, enter the full path to the printer driver, if necessary, or click **Browse** to select the location.

Set Flags

In the **Flags** area, select any of the following options:

- **Enable Printing** – Enable printing determines the initial behavior of the spooler daemon for the print device. By default, printing is enabled.

- **Enable Queuing** – Enable queuing determines the initial behavior of the spooler daemon for the print device. By default, queuing is enabled.
- **Tab Expansion** – Tab expansion determines whether tabs are expanded to spaces on output. By default, tabs are expanded.

Select Carriage Return/Line Feed Handling Mode

In the **CR Mode** area, select how you would like the printer to treat carriage returns and line feeds. The following options are available:

- No conversion
- **LF -> CR/LF** – Converts a line feed to carriage return, line feed
- **CR -> LF** – Converts a carriage return to a line feed
- **No CR @(0,0)** –

Select Flow Control Options

In the **Flow Control** area, select communication flow control for the printer. The following options are available:

- None – Cancels any DTR or XON/XOF flow control handshaking that may be set.
- DTR – Enables Data Terminal Ready handshaking. Loss of DTR is treated the same as the input of the stop character. If DTR is not enabled, loss of DTR is treated as a hangup, and foreground jobs are terminated.
- XON/XOFF Any – Enables XON/XOFF handshaking. Any character acts the same as the XON character.
- XON/XOFF Starts – Enables XON/XOFF handshaking. The XON character is required to *start* transmission. On Berkeley systems, the XON character is the start character; on System V, XON is always CTRL-Q.

Define Lock Files

Enter the full paths of lock files you want to use in the **Lock Files** boxes, or click **Browse** to select the location. You can specify up to two lock files that already exist on your system.

Specify NLS Map

Enter the path to the NLS map you want to use for the printer in the NLS Map box, or click **Browse** to select the location.

Specify Delays

You can specify line feed and form feed settings in the **Delays** area.

Select a line feed setting from the **Line Feed** list. The default setting is None.

Select a form feed setting from the **Form Feed** list. Valid options are:

- No form feed
- Form feed no delay
- Form feed 2 second delay

The default setting is No FF.

Select Parity

Select the appropriate parity in the Parity area. Valid options are:

- None
- Even
- Odd

Define Other Options

You can specify any of the UniVerse PTERM settings in the **Other Options** box. For a description of available PTERM settings, see the *UniVerse User Reference*.

Click **OK** to save the printer definitions and update the &DEVICE& file, or click **Cancel** to exit without saving changes.

Viewing or Modifying a Printer Definition

To view details or modify information for a printer previously defined in the &DEVICE& file, either double-click the printer, or select the printer in the **Printers** area of the **UniVerse Maintain Devices** dialog box, then click **Detail**. A dialog box similar to the following example appears:

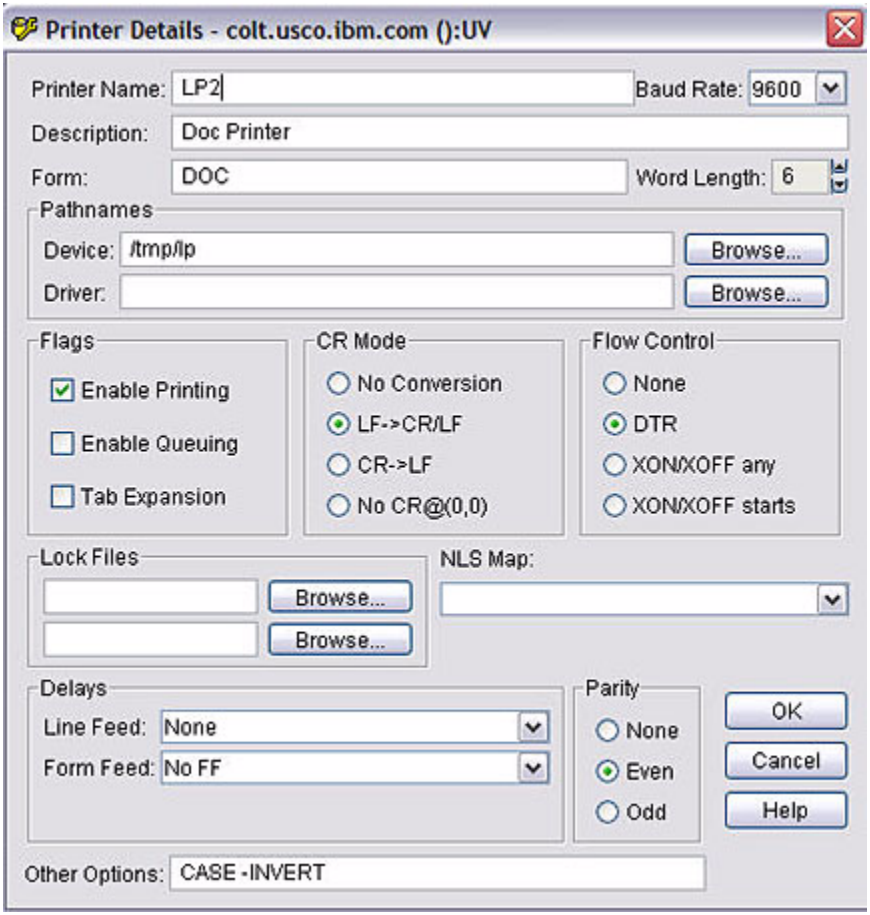

If you want to change a definition for the print device, make the necessary changes, then click **OK** to update the tape device detail, or click **Cancel** to exit without saving changes.

Deleting a Print Device Definition

To delete a print device definition, select the device in the Printers area of the **Maintain UniVerse Devices** dialog box, then click **Delete**. The following dialog box appears:

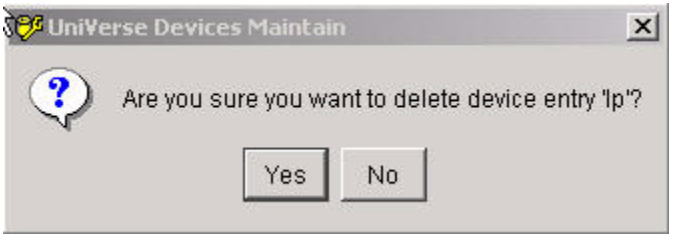

Administering Print Devices on Windows Platforms

On Windows platforms, you must configure printers using the Windows Print Manager. You cannot use the device option in UniAdmin to configure printers if you are connected to a Windows system.

Configuring Other Devices

To configure a device other than a tape device or a print device, from the **UniVerse Maintain Devices** dialog box, select **Other**, then click **New**. A dialog box similar to the following example appears:

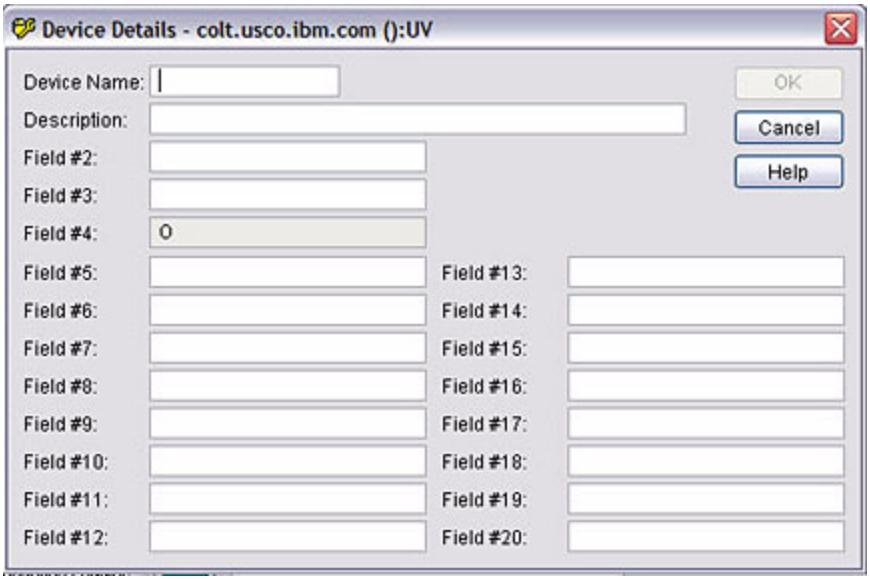

Enter Device Name

In the **Device Name** box, enter the name of the new device.

Enter Description

In the **Description** box, enter a description of the device, if desired.

Enter Path

In the **Field #2** box, enter an operating system path for the device.

Complete Device Definition

Enter other necessary information in the **Field #*nn*** boxes. Field #4 is automatically set to 0 and cannot be modified. The other fields correspond to the same field numbers in the &DEVICE& file, and most apply only to tape device and print device definitions. For information about the contents of the other fields, see the *UniVerse User Reference*.

Click **OK** to save the device definition and update the &DEVICE& file, or click **Cancel** to exit the program without saving changes.

View Details about a Device

To view or modify details about a device other than a printer or tape device, from the **UniVerse Maintain Devices** dialog box, select **Other**, then select the device for which you want to view or modify details, then click **Details**. A dialog box similar to the following example appears:

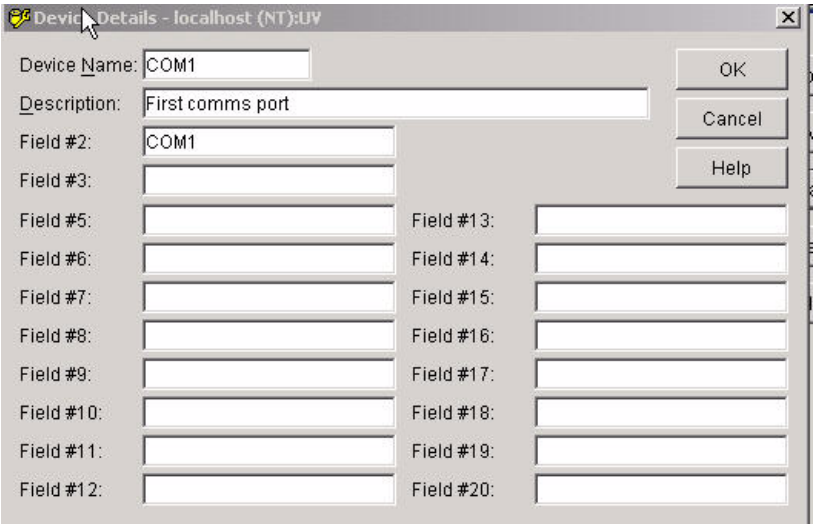

Complete any necessary modifications, then click **OK** to save your changes and update the &DEVICE& file, or click **Cancel** to exit the program without saving changes.

Delete a Device Definition

To delete a device definition other than a tape or print device definition, select **Other** from the **UniVerse Maintain Device** dialog box, then select he device you want to delete, and then click **Delete**. A dialog box similar to the following example appears:

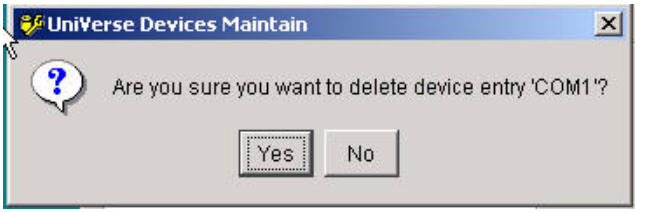

Click **Yes** if you want to delete the device, or click **No** to exit without making changes.

Chapter

8

Viewing Disk Space Usage

This chapter describes how to display file system disk usage through UniAdmin.

Select one of the following methods to access the **File System Usage** dialog box:

- From the **UniData Admin** window, double-click **Disk Space**.
- From the **UniAdmin** menu, click Admin, then click **Disk Space**.
- From the **UniAdmin** toolbar, click the **Show Disk Space Information** icon, as shown in the following example.

The **File System Usage** dialog box appears, as shown in the following example:

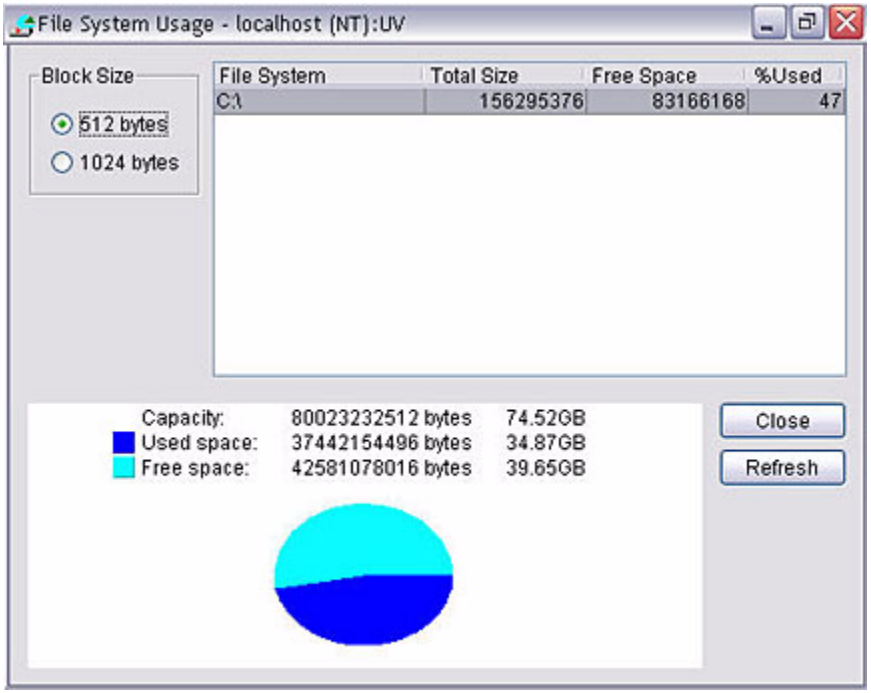

Chapter

9

Administering UniVerse Files

List Files in UniVerse Accounts 9-4
View File Properties 9-5
View File Statistics 9-11
Running File Diagnostics 9-13
Repairing Damaged Files 9-16
Select Display Options 9-20

This chapter describes how to configure and administer UniVerse files through UniAdmin. You can perform the following tasks from the **UniVerse File Tool** dialog box:

- List all files in all UniVerse accounts
- List file properties
- List file statistics
- Run file diagnostics
- Repair damaged files

Select one of the following methods to access the **UniVerse File Tool** dialog box:

- From the **UniData Admin** window, double-click **File Tool**.
- From the **UniAdmin** menu, click **Admin**, then click **File Tool**.
- From the **UniAdmin** toolbar, click the **Manage Devices** icon, as shown in the following example.

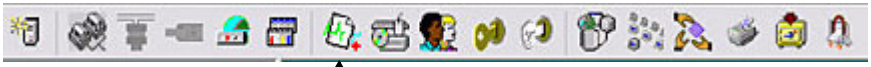

File Tools Icon

The **UniVerse File Tool** dialog box appears, as shown in the following example:

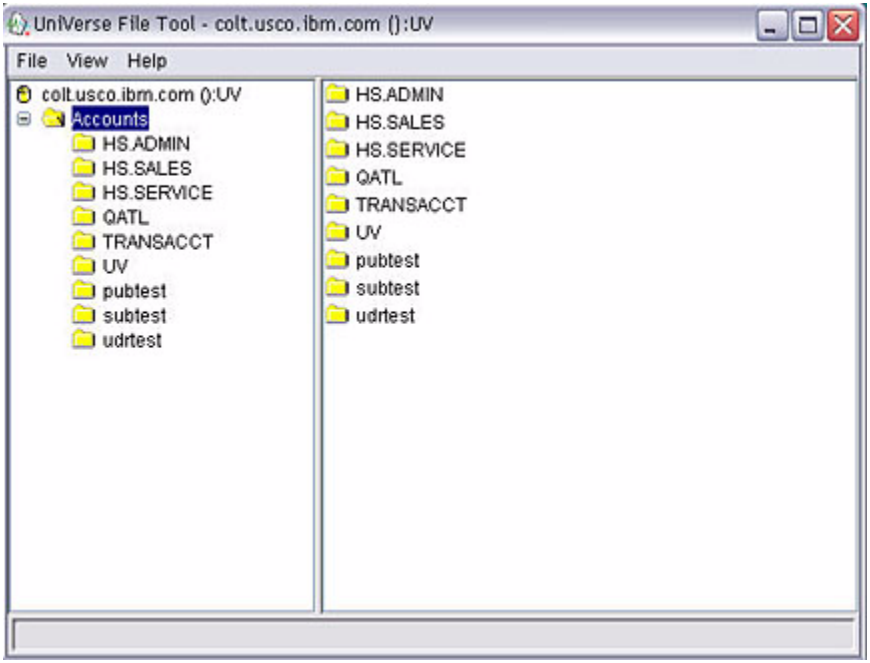

List Files in UniVerse Accounts

From the **UniVerse File Tool** window, select the account for which you want to view files. All files for that account appear in the right pane of the window, as shown in the following example:

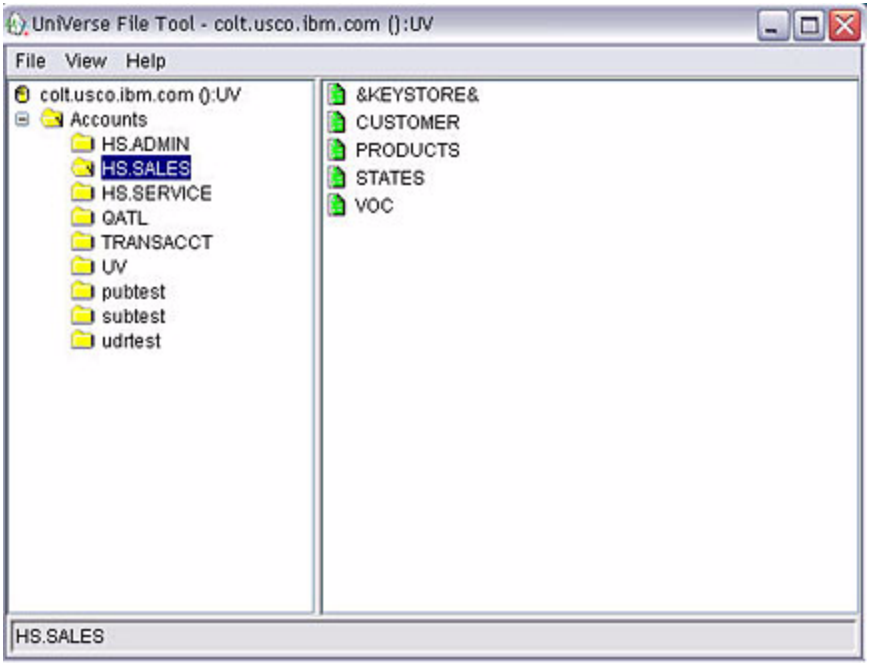

View File Properties

To view the properties of a file, select the file for which you want to view properties in the right pane of the **UniVerse File Tool** dialog box, then from the **File** menu, click **Properties**. The **Properties** dialog box appears, as shown in the following example:

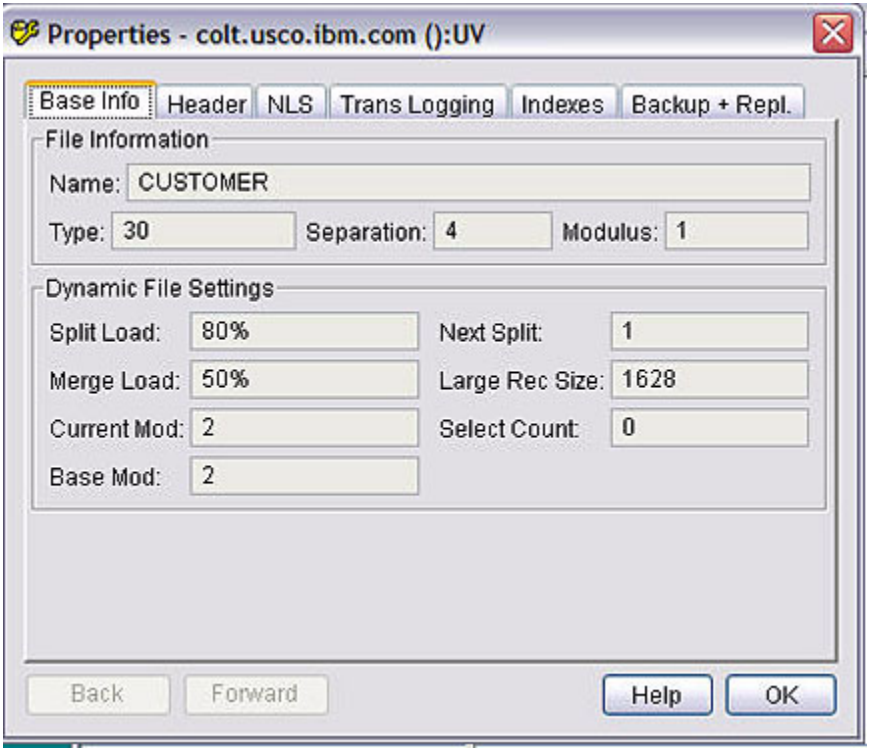

Base Information

From the **Properties** dialog box, select the **Base Info** tab. UniAdmin displays the following information about the file:

- File name
- File type
- Separation and modulus of static hashed file
- Dynamic file parameters

For a detailed description of UniVerse files, see the *UniVerse System Description*.

Header Information

From the **Properties** dialog box, select the **Header Info** tab. A dialog box similar to the following example appears:

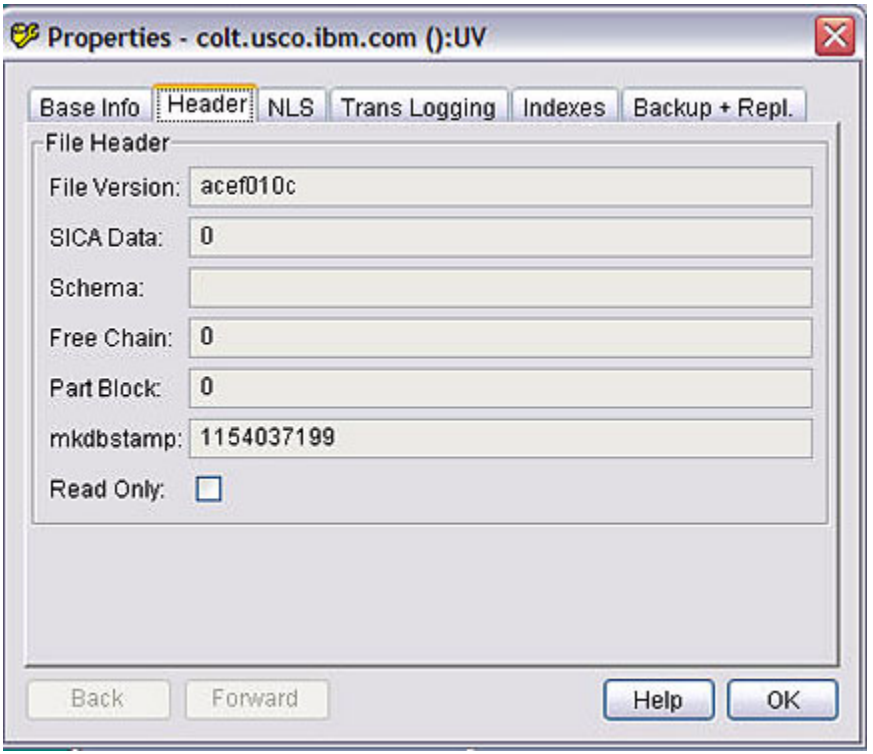

UniAdmin displays the following header information for the file:

- File version
- SICA and schema name, if the file is a table
- Free chain
- Part block
- mkdbstamp

If the **Read Only** check box is selected, the file is read-only.

National Language Support (NLS) Information

From the **Properties** dialog box, click the **NLS** tab. A dialog box similar to the following example appears:

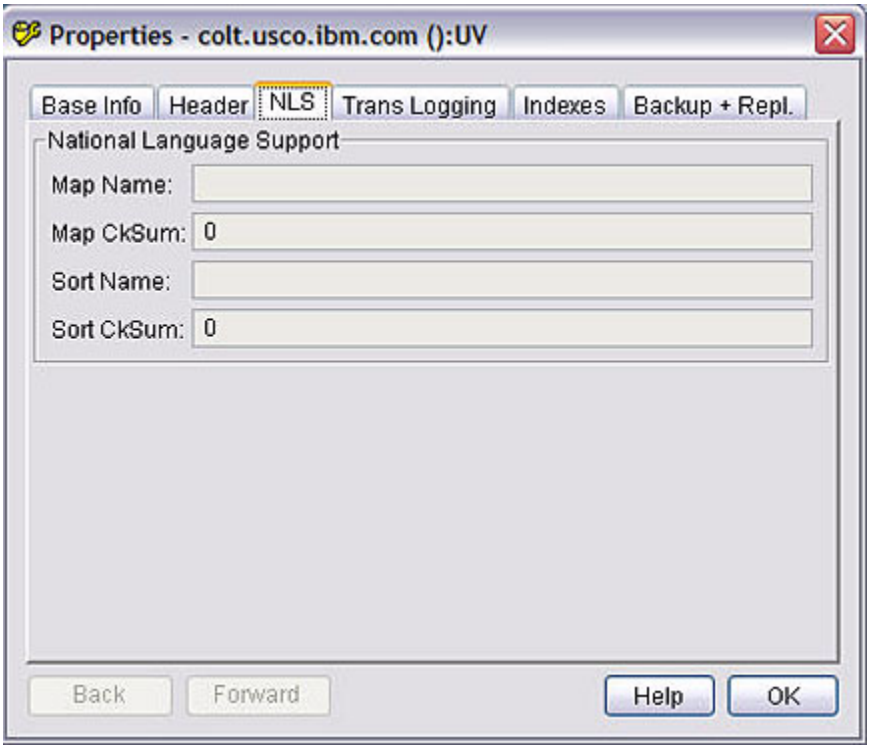

UniAdmin displays the following information about NLS:

- In the **Map Name** box, the name of the character set map associated with the file
- Map checksum
- In the **Sort Name** box, the Collate convention that determines how to sort file data
- Sort checksum

For detailed information about NLS, see the *NLS Guide*.

Transaction Logging Information

From the **Properties** dialog box, click the **Trans Logging** tab. A dialog box similar to the following example appears:

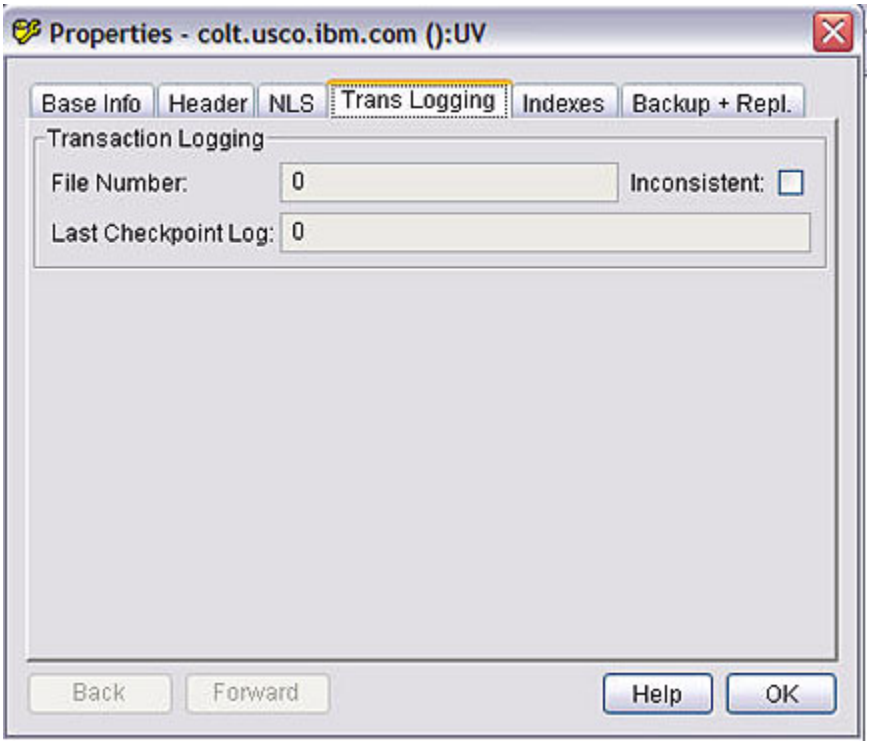

UniAdmin displays the following information about transaction logging:

- File number
- Number of the last checkpoint log

If the **Inconsistent** check box is selected, and file is inconsistent.

Indexes Information

From the **Properties** dialog box, click the **Indexes** tab. A dialog box similar to the following example appears:

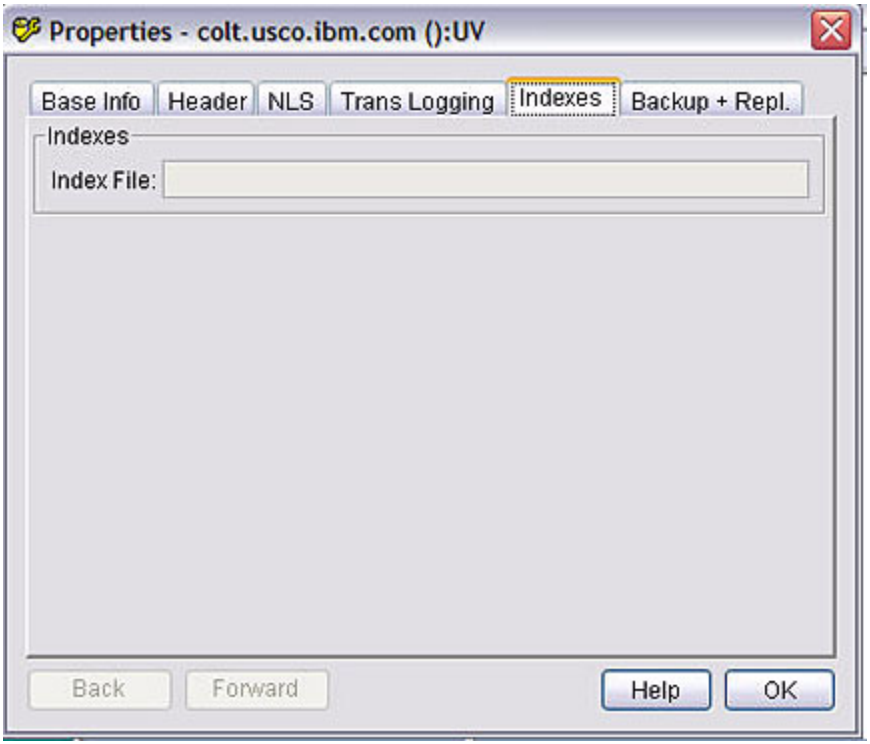

If the file has a secondary index, UniAdmin displays the name of the index in the **Index File** box.

Backup and Replication Information

From the **Properties** dialog box, click the **Backup + Repl.** tab. A dialog box similar to the following example appears:

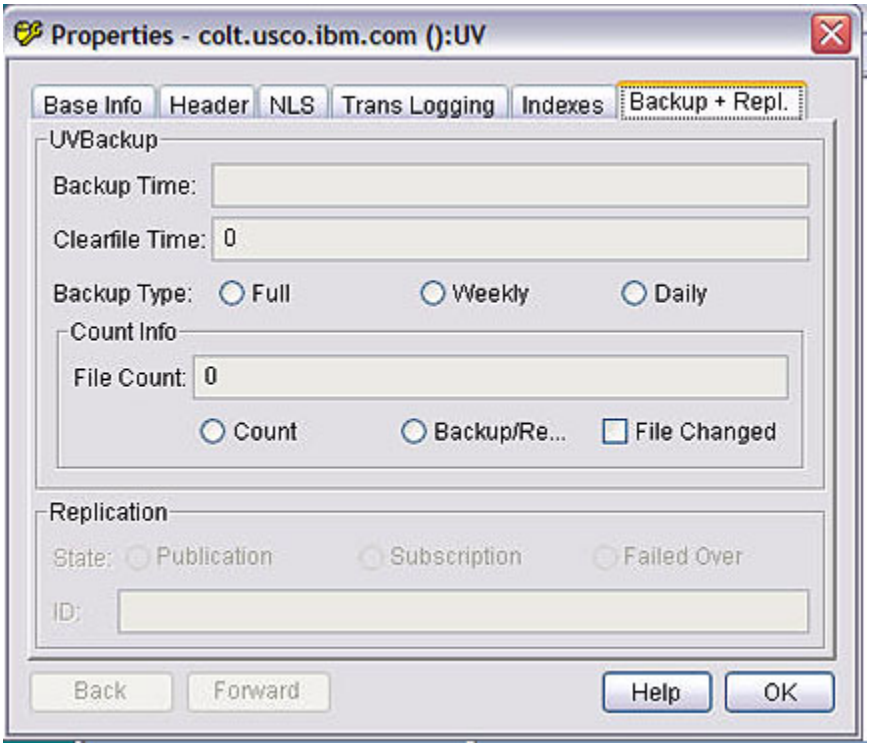

UniAdmin displays the following information about backup and replication:

- In the **Backup Time** box, the date and time of the last backup.
- In the **Clearfile Time** box, the date and time the last CLEARFILE command was executed against the file.
- The type of backup, either full, weekly, or daily.
- In the **File Count** box, the number of records in the file, counted by either the last COUNT command executed against the file, the last full backup, or the last restore.

  If the **File Changed** check box is selected, the file count may be out of date because the file has been changed since the last file count.

- In the **Replication** area, **Stat** indicates whether the file is a published file, a subscription file, or a failed-over file. This area also lists the replication ID in the **ID** box.

View File Statistics

To view statistics about a file, from the **UniVerse File Tool** window, select the file for which you want to view statistics, then click **File**, then click **Statistics**. A dialog box similar to the following example appears:

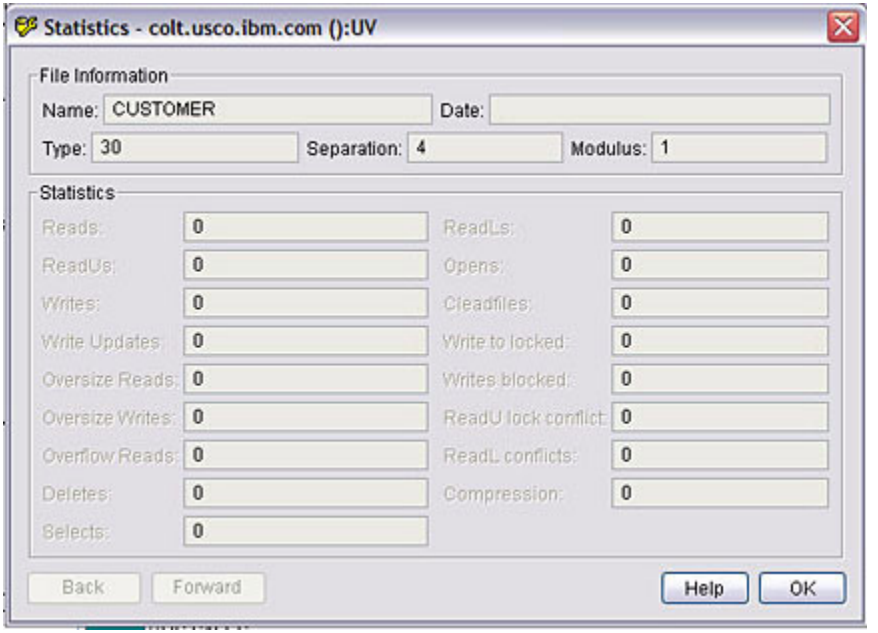

File Information

In the **File Information** area of the **Statistics** dialog box, UniAdmin displays the following information:

- File Name
- Date of the last update
- File separation if a static hashed file

- Modulus if a static hashed file

File Statistics

UniAdmin displays the following statistics about the file you selected:

Field	Description
Reads	Total number of READ operations on the file.
ReadUs	Total number of READU operations on the file.
Writes	Total number of WRITE operations on the file.
Write Updates	Total number of WRITEU operations on the file.
Oversize Reads	Total number of READ operations executed against large records.
Oversize Writes	Total number of WRITE operations executed against large records.
Overflow Reads	Total number of READ operations that accessed overflow buffers.
Deletes	Total number of DELETE operations on the file.
Selects	Total number of SELECT operations on the file.
ReadLs	Total number of READL operations on the file.
Opens	Total number of OPEN operations on the file.
Clearfiles	Total number of CLEARFILE operations on the file.
Write to Locked	Total number of WRITE operations executed against a locked record.
Writes Blocked	Total number of WRITE operations blocked by a record lock.

UniAdmin File Statistics

Field	Description
ReadU lock conflict	Total number of READU operations that failed because of an existing record lock.
ReadL conflicts	Total number READL operations that failed because of an existing record lock.
Compressions	Total number of free operations that compacted a group after a record was deleted.

UniAdmin File Statistics (continued)

Running File Diagnostics

To run file diagnostics against a file, from the **UniVerse File Tool** window, select the file for which you want to run diagnostics, then click **File**, then click **Diagnostics**. A dialog box similar to the following example appears:

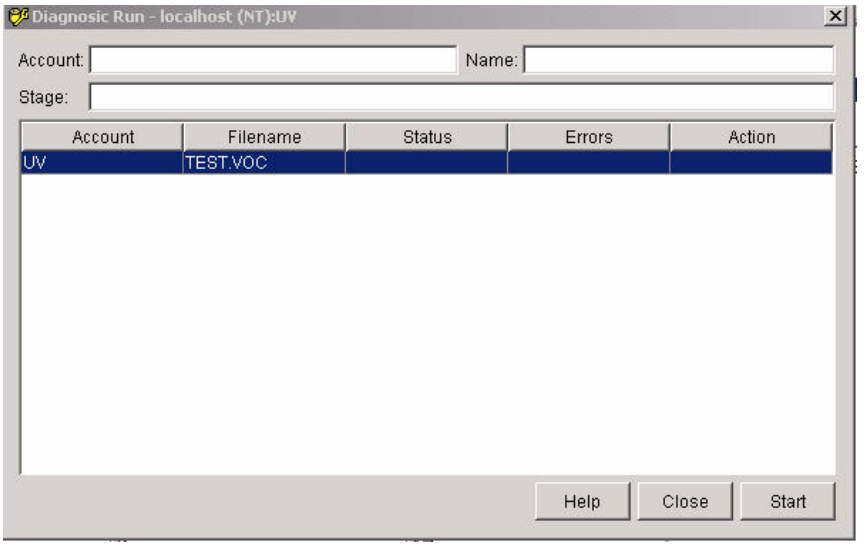

To begin running diagnostics, click **Start**. The **Diagnostics Setup** dialog box appears, as shown in the following example:

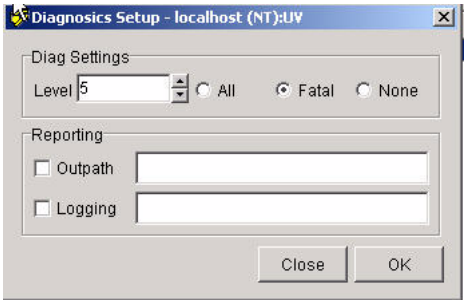

Use the **Diagnostics Setup** dialog box to specify how much diagnostic testing to perform on the file, how much diagnostic detail to list in the error report window, and where to store output from the diagnostic test.

Determine Diagnostic Level

In the **Diag Settings** area, select the level of diagnostic detail you want to produce in the **Level** box. The lowest diagnostic level is 1, while the highest diagnostic level is 10. The default value is 5. The higher this level setting, the longer the diagnostics test takes to complete. Use the **Level** arrows to select the diagnostic level.

Specify Types of Errors Report

In the **Diag Settings** area, select the types of errors you want to appear in the **Error Report** window. Select one of the following values:

- **All** – List all diagnostic details in the **Error Report** window.
- **Fatal** – List only fatal errors in the **Error Report** window.
- **None** – Do not list any errors in the **Error Report** window.

Specify Output Location

If you want to save a report of irreparable groups and record blocks detected by the diagnostic test, select **Output**, then enter the full path to a directory where you want to store the output.

If you want to store a copy of the error report, select **Logging**, then enter the full path to a directory where you want to store the report.

If you do not specify **Output** or **Logging**, the output and error report are stored in the directory where the file currently resides.

Click **OK** to save your changes, or click **Close** to exit without saving changes. The diagnostics program starts.

Diagnostics Test

After you run the diagnostics program, the **Diagnostic Run** dialog box displays the account name, the file name, and the progress of the program. It also displays the number of errors encountered and specifies what action to take if an error is detected.

The following example illustrates the **Diagnostic Run** dialog box:

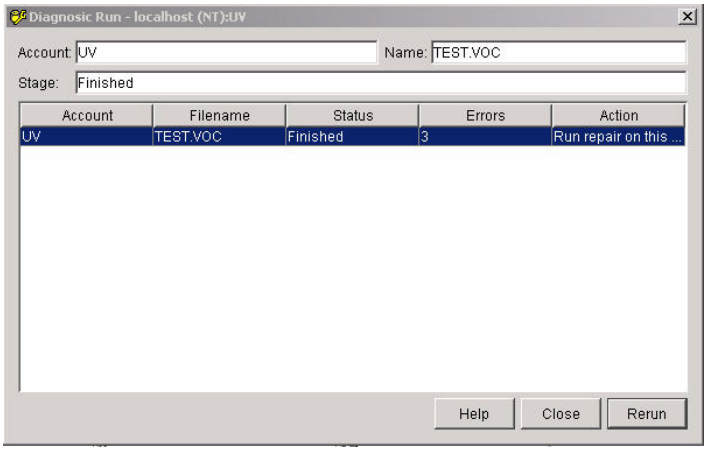

Viewing Errors

If the diagnostics test program detects an error, click the account name to display the Error Report window.

The **Error Report** windows displays the following information:

- The error type

- A description of the error
- The group number where the error occurred
- The record block number where the error occurred

After repairing the damaged file, click **Rerun** to rerun the diagnostic program.

Repairing Damaged Files

To run file diagnostics against a file, from the **UniVerse File Tool** window, select the file you want to repair, then click **File**, then click **Repair**. A dialog box similar to the following example appears:

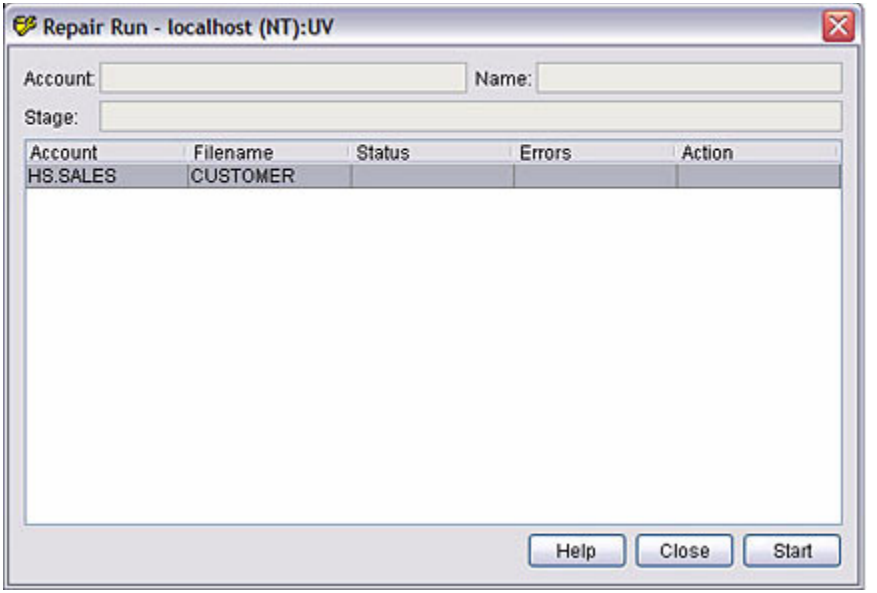

Click **Start**. The **Repair Setup** dialog box appears, as shown in the following example:

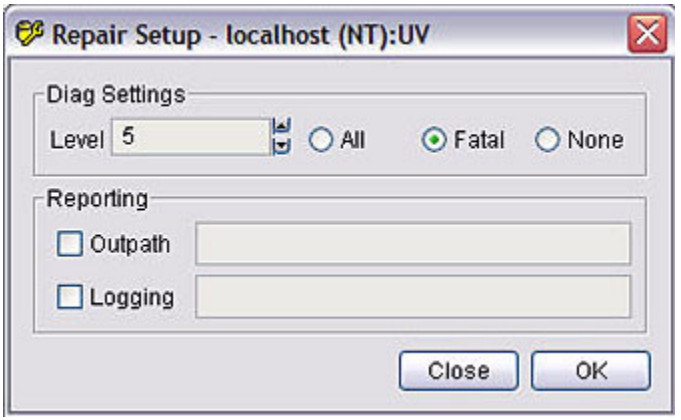

Use the **Repair Setup** dialog box to specify how much diagnostic testing to perform on the file, how much diagnostic detail to list in the error report window, and where to store output from the diagnostic test.

Determine Diagnostic Level

In the **Diag Settings** area, select the level of diagnostic detail you want to produce in the **Level** box. The lowest diagnostic level is 1, while the highest diagnostic level is 10. The default value is 5. The higher this level setting, the longer the diagnostics test takes to complete. Use the **Level** arrows to select the diagnostic level.

Specify Types of Errors Report

In the **Diag Settings** area, select the types of errors you want to appear in the **Error Report** window. Select one of the following values:

- **All** – List all diagnostic details in the **Error Report** window.
- **Fatal** – List only fatal errors in the **Error Report** window.
- **None** – Do not list any errors in the **Error Report** window.

Specify Output Location

If you want to save a report of irreparable groups and record blocks detected by the diagnostic test, select **Output**, then enter the full path to a directory where you want to store the output.

If you want to store a copy of the error report, select **Logging**, then enter the full path to a directory where you want to store the report.

If you do not specify **Output** or **Logging**, the output and error report are stored in the directory where the file currently resides.

Click **OK** to save your changes and run the Repair program, or click **Close** to run the Repair program with default settings. The repair program starts.

When the Repair program completes, a dialog box similar to the following example appears:

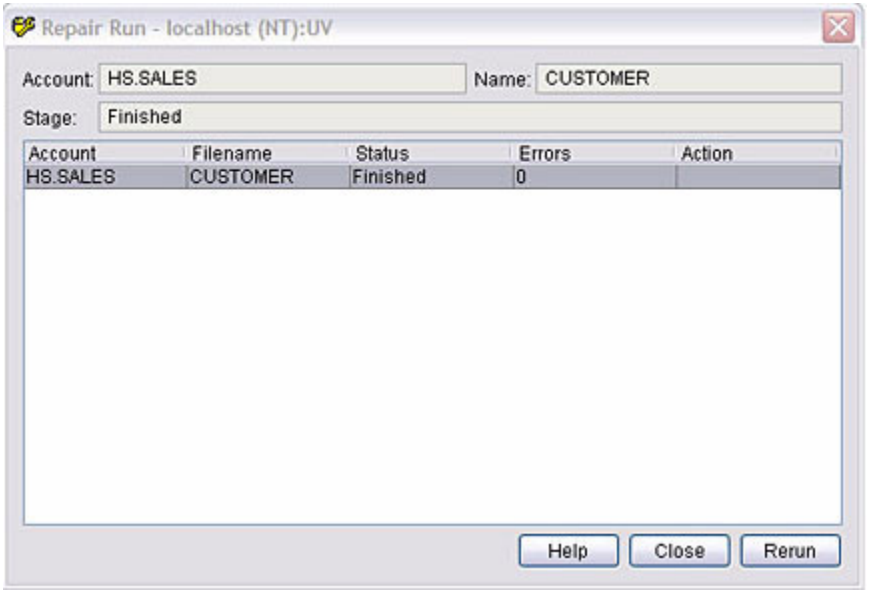

The **Repair** dialog box displays the following information:

- The account name where the file resides
- The file name
- The current stage of the program

- The number of errors encountered
- What action to take if an error is found

If an error is encountered, click the account name to display the **Error Report** window.

Rerun Repair Program

After you execute the Repair program, click **Rerun** to reexecute the program to ensure no errors remain in the file.

Exit the Program

Click **Close** the exit the dialog box and return to the **File Tool** window.

Select Display Options

To select how information is displayed in the **UniVerse File Tool** window, click **View**. The following options are available:

- **Large Icons** – Displays account and file information using large icons.
- **Small Icons –** Displays account and file information using small icons.
- **List** – Displays account and file information in a list format.
- **Details** – Lists account and file information in a list format, and displays file type and modulus for files.

Chapter

10

Importing Accounts

Importing Non-UniVerse Accounts to UniVerse 10-3
Manually Restoring Accounts from Tape 10-5
Restoring an Single-reel Account to a UNIX Server 10-5
Restoring a Multireel Account to a UNIX Server 10-5
Restoring Accounts to Windows Platforms 10-8

This chapter describes how to import Prime INFORMATION or Pick accounts to UniVerse.

Select one of the following methods to access the **UniVerse Import Account** dialog box:

- From the **UniAdmin** window, double-click **Import**.
- From the **UniAdmin** menu, click Admin, then click **Import**.

The **UniVerse Import Account** dialog box appears, as shown in the following example:

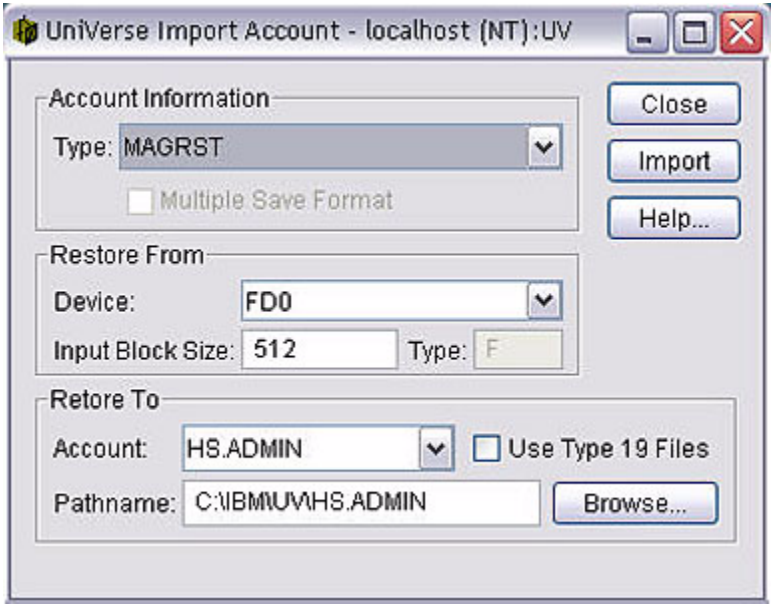

Importing Non-UniVerse Accounts to UniVerse

Complete the following steps to import a non-UniVerse account to UniVerse.

Select Tape Type

In the **Type** box, select the tape type. Valid tape types are:

- MAGRST
- PICK
- REALITY (Microdata)
- IN8000
- IN5000

The tape type you choose determines which import program UniVerse uses, and the required tape positioning.

If you are using a REALITY tape type, you can select the **Multiple Save Format** check box, if required.

Select Device

In the **Restore From** area, select the tape device to use in the **Device** box. UniAdmin displays devices previously defined in the &DEVICE& file, and populates the **Input Block Size** and **Type** fields with the settings stored in the &DEVICE& file. You may change these values, if necessary.

Choose Account Location

In the **Restore To** area, select the account to which you want the account restored in the **Account** box. UniAdmin populates the **Pathname** box with the full path to the account.

Alternatively, you can enter the full path to the account in the **Pathname** box, or click **Browse** to search the system for a suitable directory.

Select the **Use Type 19 Files** check box if you want to restore the account using type 19 files instead of type 1 files (the default).

Import Account

Click **Import** to begin importing the account, or click **Close** to exit the program. The results of the transfer appear in the **UniVerse Command Output** window.

Manually Restoring Accounts from Tape

You can manually restore non-UniVerse accounts directly from tape. This method may be necessary if additional import options, not available in the **UniVerse Import Account** dialog box, are required.

Restoring an Single-reel Account to a UNIX Server

You can manually restore a Pick or Prime INFORMATION account from tape to a UNIX server.

Complete the following steps if you are restoring a single-reel ACCOUNT-SAVE or MAGSAV tape:

1. Log on as root.
2. Create a new directory to be the parent directory for the account.
3. Change directories to the directory created in step 2.
4. If you are restoring a Pick account, use the *acct.restore* command. If you are restoring a PRIME information account, use the *magrst* command. Both of these commands create a UniVerse account in the current directory, as well as subdirectories that contain the accounts loaded from tape.

Restoring a Multireel Account to a UNIX Server

Use the *tapein* program to restore multireel ACCOUNT-SAVE and MAGSAV tapes to a UNIX system. You can execute this program directly from a UNIX shell.

tapein uses the standard UNIX filter to take data from an input device or file and pipe the data to *acct.restore* or *magrst*.

The *dd* filter correctly handles end-of-media on most devices, eliminating many of the problems of multivolume tape handling. The *tapein* filter provides the necessary handshaking and signaling between *dd* and the *acct.restore* and *magrst* procedures. *tapein* prompts for the next volume of input and performs the necessary cleanup when the restoration procedures complete.

***Note**: It is not necessary to use the tapein filter for diskettes.*

The syntax for *tpein* follows:

tapein -p*restore.fileter* **-f***filename* [*options*]

Do not separate variables from the option codes by a space.

The following table describes each parameter of the syntax.

Parameter	Description
-p*restore.filter*	The full UNIX path for either *acct.restore* or *magrst.*
-f*filename*	The input file name for the *if* option of the *dd* command. The input file name is the path of the input device or file.

tapein Parameters

The following table describes the *tapein* options.

Option	Description
-a*pathname*	Reads input from the specified UNIX path. *pathname* is the full UNIX path.
-b*blocksize*	The input block size for the *ib* option of the *dd* command. For 1/2-inch tape input, *blocksize* should be larger than the largest physical block on tape. For other devices, the most likely value is 512 bytes. The default is 8192, which should be correct for most 1/2-inch tapes.
-c	Specifies that *dd* should use the *conv=swab* option to swap data bytes on input.
-D *device*	The path of a terminal (tty) device used when multiple instances of *acct.restore* are executed.
-g*pathname*	The path of the rewind tape device.
-h*frhdr*	The number of header files to skip over at the beginning of the first reel only. The default is 0. The value of this parameter should be 0 for *magrst* and REALITY ACCOUNT-SAVE tapes, and 2 for most Pick ACCOUNT-SAVE tapes.
-ice	Generates the input prompt character for use with the **ic_execute** function.

***tapein* Options**

Option	Description
-l*labels*	The number of label records to skip over at the beginning of each reel. The default value is 0. *labels* is used as the value of the *skip* option of the *dd* command. For MAGSAV tapes, the value of this parameter should be 0. For ACCOUNT-SAVE tapes, it depends on the type of media and the type of ACCOUNT-SAVE tape. For 1/2-inch ACCOUNT-SAVE tapes, the value of this parameter should be 1. For most cartridge tapes, if the input block size is 512 and the ACCOUNT-SAVE tape has a tape label blocked at 512 characters, this parameter should be 1. If the ACCOUNT-SAVE tape is on a cartridge and the tape label is blocked at greater than 512 characters, the input block size times the number of label records should equal the total number of characters in the tape label block. For example, if the cartridge tape has a label blocked at 8192, an input block size of 512 skipping 16 label records, will work correctly. This same formula works with UNIX files containing a tape label.
-m	Restores data from a Microdata REALITY tape.
-mas	Restores data from a Microdata 7.0 M-A-S tape.
-n	Specifies that the separation is a factor of 2048.
-r*srhdr*	The number of header files to skip over at the beginning of each subsequent reel for multireel input tapes. The default value is 0. The value of this parameter should be 1 for *magrst* tapes and should probably be 0 for Pick ACCOUNT-SAVE tapes.
-t	Specifies that *acct.restore* and *magrst* should create type 19 files instead of type 1 files.
-u	Restore data using Ultimate cartridge format.

tapein* Options *(continued)

The source for the *tapein* filter is in the *sample/tapein.c* file in the UV account directory. It can be modified to your specific needs. To compile a new version of *tapein*, save the original file, make your changes, and compile using the following command:

```
# cc tapein.c -o tapein
```

Because *tapein* is a user-modifiable program, it is not linked with the UniVerse library, and does not trap the **BREAK** key in the same was as UniVerse. If you press the **BREAK** key, *tapein* tries to clean up all its subprocesses before exiting.

Restoring Accounts to Windows Platforms

You can manually restore Pick or Prime INFORMATION accounts from tape to a Windows server.

Complete the following steps if you are restoring a single-reel ACCOUNT-SAVE or MAGSAVE tape.

1. Log on as Administrator.
2. Create a new directory to be the parent directory for the account.
3. Move to the directory created in step 2.
4. If you are loading a Pick account, use the *acct.restore* executable. If you are loading a Prime INFORMATION account, use the *magrst* executable. These executables create a UniVerse account in the current directory as well as subdirectories that contain the account loaded from the tape.

Using the acct.restore or magrst Executables

If you are restoring a multireel ACCOUNT-SAVE or MAGSAV tape, use *tapein.exe*.

The *acct.restore* executable transfers a Pick System ACCOUNT-SAVE tape to a UniVerse PICK flavor account on a Windows server.

The *magrst* executable transfers a Prime INFORMATION MAGSAV tape to a UniVerse INFORMATION flavor account on a Windows server.

Both of these executables are located in the *bin* directory of the UV account directory. You must run these executables from an MS-DOS window.

Syntax:

acct.restore.exe[*options*]

magrst.exe[*options*]

The following table describes the options of the syntax.

Option	Description
-	Reads input from standard input.
-a *pathname*	Reads input from the disk file specified in the path.
-m	This option is available for *acct.restore.exe* only. Restores REALITY Microdata tapes.
-n	This option is available for *acct.restore.ext* only. Creates files with a default separation of 4 (each group buffer holds 2048 bytes). Use this option when you are transferring accounts from systems such as ADDS Mentor or PICK 370, where a separation of 1 means 2K bytes instead of 512 bytes.
-s	This option is available for *acct.restore.exe* only. Used with the *tapein* executable. Send the appropriate signals back to the *tapein* filter so *tapein* can prompt for the next volume of input and exit cleanly after *acct.restore.exe* completes.
-t *uv.device*	Reads input from a tape device defined in the &DEVICE& file. *uv.device* is the ID of the entry in the &DEVICE& file.
-19	Creates type 19 files instead of type 1 files.

***acct.restore.ext* and *magrst.exe* Options**

The following example illustrates the *magrst.exe* command:

```
d:\>magrst.exe -t MT0
```

If you use *acct.restore.exe* or *magrst.exe* without any arguments, the executables use MT0 as the default tape device.

If *acct.restore.exe* is run by *tapein.exe*, it does not prompt the user when an error condition exists, such as unrecognized data or a file cannot be opened. It ignores the data or error and continues with the restoration.

If the MAGSAV tape is constructed on a PRIMOS environment of Release 20 or beyond, you must specify the -NO -ACL option of the MAGSAV command. This makes tapes that are compatible with Release 18. At Release 21, you must also specify the -REV19 option to create tape formats compatible with Release 19. UniVerse does not support tape formats compatible with Release 20 or later formats.

Using tapein.exe

The *tapein* executable is in the *bin* directory of the UV account directory. Use *tapein.exe* to restore multireel ACCOUNT-SAVE and MAGSAV tapes to a Windows system.

tapein.exe takes data from a tape device and pipes the data to *acct.restore.exe* or *magrst.exe*. This program handles end-of-media on most devices, eliminating many of the problems of multivolume tape handling. *tapein.exe* prompts for the next volume of input and performs the necessary cleanup when the restoration procedures complete.

***Note**: It is not necessary to use tapein.exe for diskettes.*

Syntax:

> **tapein.exp -p***restore.filter* **-f***filename*[-**b***blocksize*] [**-t**] [**-c**] [**-l**] [**-m**] [**-n**] [**-u**] [**-i**] [**-a**] [**-r**] [**-norew**]

Do not separate variables from the option codes by a space.

The following table describes each parameter of the syntax.

Parameter	Description
-p*restore.filter*	The full Windows path for either *acct.restore.exe* or *magrst.exe*.
-f*filename*	The input file name. This is the path of the tape device, for example, \\.\tape0.
-b*blocksize*	The input block size. For 1/2-inch tape input, *blocksize* should be larger than the largest physical block on tape. For other devices, the most likely value is 512 bytes. The default is 8192, which should be correct for most 1/2-inch tapes.
-t	Specifies that *acct.restore.ext* and *magrst.exe* should create type 19 files instead of type 1 files.
-c	Specifies the use of cartridge tape format.
-l	Activates logging to the *tapein.log* file.
-m	Specifies Microdata REALITY format.
-n	Sets the separation as a factor of 2048.

***tapein.exe* Parameters**

Parameter	Description
-u	Specifies Ultimate cartridge format.
-i	Specifies an IN2 type tape.
-a	Specifies the input file is an ASCII file, not a tape device.
-r	Verifies the reel number
-norew	Specifies not to rewind the tape after the restoration is complete.

tapein.exe* Parameters *(continued)

The following example illustrates the *tapein.exe* command:

```
d:\> tapein.ext -pd:\uv\uv\bin\magrst.exe -f\\.\tape0 -1
```

Using uvmt.exe

The *uvmt* executable is in the *bin* directory of the UV account directory. Use *uvmt.exe* to restore accounts from a tape created on an IN2 system to a Windows system.

Syntax:

uvmt.exe -d*devicename*[**-b***blocksize*] [**-c**] [**-t**] {*command*}

Do not separate variables from the option codes by a space.

The following table describes each parameter of the syntax.

Parameter	Description
-d*devicename*	Reads input from the tape device, for example, \\.\tape0.
-b*blocksize*	The input block size. For 1/2-inch tape input, *blocksize* should be larger than the largest physical block on tape. For other devices, the most likely value is 512 bytes. The default is 512, which should be correct for most 1/2-inch tapes.
-c	Specifies that the device is a cartridge tape device. This is the default setting.
-t	Specifies that the device is a magnetic tape device.

uvmt.ext Parameters

The following table describes valid *commands*:

Command	Description
rew	Rewinds the tape device.
bskip	Skips a tape block.
fskip	Skips a tape file (EOF skip).
bread	Reads a block from tape and writes it to standard out.
fread	Reads a file from tape and writes it to standard out.

uvmt.exe Commands

Restoring IN8000 Tapes

Complete the following steps to restore tapes created on an IN8000 machine (in SMA format).

1. Rewind the tape using the following command:

   uvmt.exe -d\\.\tape0 -b512 -c rew

2. Skip the first two files on tape using the following commands:

   uvmt.exe -d\\.\tape0 -b512 -c fskip
   uvmt.exe -d\\.\tape0 -b512 -c fskip

3. Run the restoration process using tapein.exe:

   tapein.exs -p$UVHOME\bin\acct.restore.exe -f\\.\tape0 -b512 -i -c -r

Note: *$UVHOME represents the UV account directory. The device name shown is an example.*

If the tape contains multiple accounts, the -norew option should be included in the *tapein.exe* command line. This stops the tape from rewinding at the end of the first account restoration.

Complete the following steps to restore a tape containing two accounts:

1. Restore the first account using the following commands:

   uvmt.exe -d\\.\tape0 -b512 -c rew
   uvmt.exe -d\\.\tape0 -b512 -c fskip
   uvmt.exe -d \\.\tape0 -b512 -c fskip
   tapein.exe -p$UVHOME\bin\acct.restore.exe -f\\.\tape0 -i -c-r -norew

2. Restore the second account using the following commands:

   uvmt.exe -d\\.\tape0 -b512 -fskip
   tapein.exe -p $UVHOME\bin\acct.resstore.exe -f\\.\tape0 -b512 -i -c -r

Restoring IN5000 Tapes

Complete the following steps to restore tapes created on an IN5000 machine.

1. Restore the tape using the following command:

   uvmt.exe -d\\.\tape0 -b512 -c rew

2. Skip the header files using the following commands:

   uvmt.exe -d\\.\tape0 -b512 -c fskip
   uvmt.exe -d\\.\tape0 -b512 -c fskip

3. Run the restoration process using *INfilter.exe*, located in the *bin* directory of the UV account. This executable filters the data before it is passed to *acct.restore.exe*. The pipe symbol (|) pipes the data between each program.

   -uvmt -d\\.\tape0 -b512 -c fread |INfilter.exe | acct.restore.exe -i

If the tape contains two account, restore the first account using the following commands:

uvmt -d\\.\tape0 -b512 -c rew
uvmt -d\\.\tape0 -b512 -c fskip
uvmt -d\\.\tape0 -b512 -c fskip
uvmt -d\\.\tape0 -b512 -c fread |INfilter.exe | acct.restore.exe -i

Restore the second account using the following commands:

uvmt -d\\.\tape0 -b512 -c fskip
uvmt -d\\.\tape0 -b512 fread |INfilter.exe | acct.restore.exe -i

Chapter

11

Licensing UniVerse

Authorizing a New UniVerse License 11-3
Obtain Configuration Code 11-4
Obtain Authorization Code 11-4
Viewing License Information 11-5
Cleaning Up Current Licenses 11-6

This chapter describes how to license UniVerse through UniAdmin.

Select one of the following methods to access the **UniVerse License** dialog box:

- From the **UniAdmin** window, double-click **License**.
- From the **UniAdmin** menu, click Admin, then click **License**.

The **UniVerse License** dialog box appears, as shown in the following example:

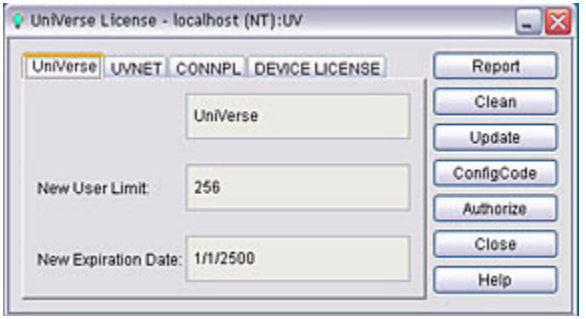

Authorizing a New UniVerse License

Verify that the configuration displayed in the **UniVerse Licensing** dialog box matches the configuration on the Product Configuration sheet shipped with UniVerse. If you need to update any information, click **Update**. The **UniVerse License Update** dialog box appears, as shown in the following example:

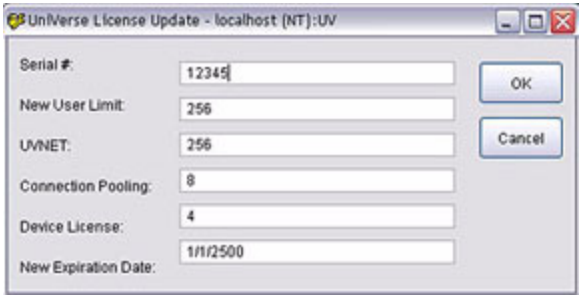

1. Enter your UniVerse serial number in the **Serial #** box.
2. Enter the number of users for which you are licensed in the **New User Limit** box.
3. Enter the number of UVNet users for which you are licensed in the **UVNET** box.
4. Enter the number of Connection Pooling licenses in the **Connection Pooling** box.
5. Enter the number of device licenses for which you are authorized in the **Device License** box.
6. If the expiration date of you license is incorrect, enter the correct date in the **New Expiration Date** box.

When the license information is correct, click **OK**.

Note: *If you are using UV/NET, you must authorize both the UniVerse database and UV/NET.*

Obtain Configuration Code

Click **ConfigCode** to obtain the configuration code you will need to authorize UniVerse. A window similar to the following example appears with the configuration code:

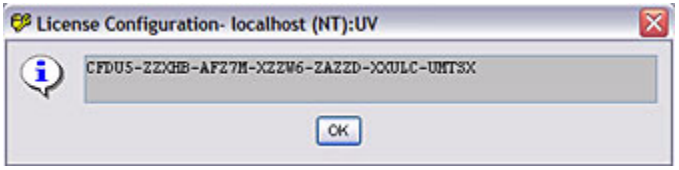

Obtain Authorization Code

To obtain your authorization code, go to:

http://www-306.ibm.com/software/data/u2/universe/

Click **Authorize Products**. Follow the instructions on the website to obtain you authorization code.

Once you have your authorization code, click **Authorize** from the **UniVerse License** dialog box. The **License Authorization** dialog box appears, as shown in the following example:

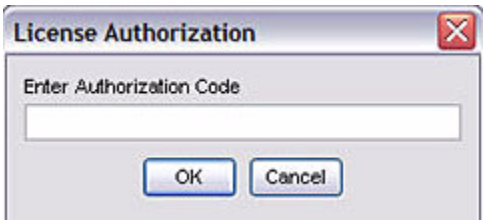

Enter your authorization code, then click **OK**.

Viewing License Information

From the **UniVerse License** dialog box, click **Report** to view license information. A window similar to the following example appears:

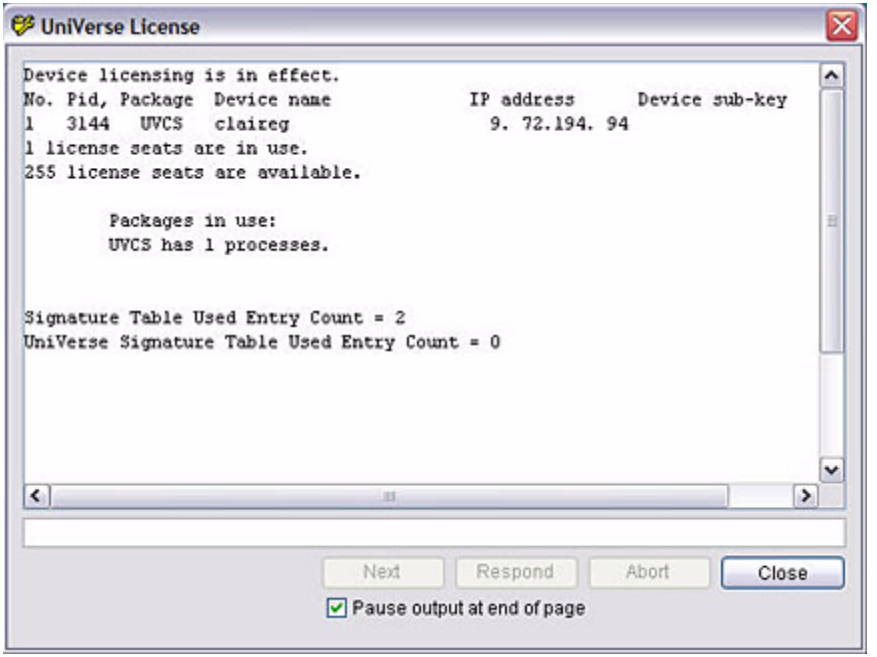

The **UniVerse License** window displays the following information:

- Licenses in use at both the UniVerse and the package level.
- Identifies the process that owns the license, and lists package licenses it holds.
- Identifies the remote device holding the license.

Cleaning Up Current Licenses

To clean up a current license, from the **UniVerse License** dialog box, click **Clean**.

On UNIX systems, UniAdmin cleans up the current licenses based on shared memory segments associated with dead processes.

On Windows platforms, UniData cleans up the current licenses based on dead entries in the process table.

After cleaning up the licenses, UniVerse recomputes license counts at the UniVerse, package, and seat levels.

When the UniVerse finishes cleaning up licenses, the following message appears:

Chapter

12

Administering Locks

File and Record Locks 12-4
Clearing Locks 12-5
Group Locks 12-7

This chapter describes how to administer locks through UniAdmin.

Locks are set on UniVerse files by certain BASIC statements and UniVerse commands. The type of lock determines what a process can access while other processes hold locks on records or files.

Select one of the following methods to access the **Lock Administration** dialog box:

- From the **UniAdmin** window, double-click **Locks**.
- From the **UniAdmin** menu, click **Admin**, then click **Locks**.
- From the **UniAdmin** toolbar, click the **Manage Locks** icon, as shown in the following example.

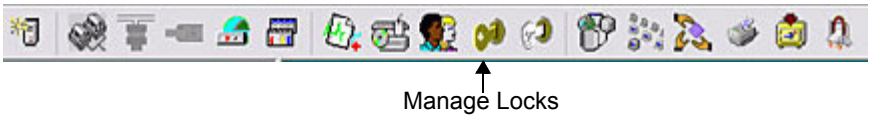

The **Lock Administration** dialog box appears, as shown in the following example:

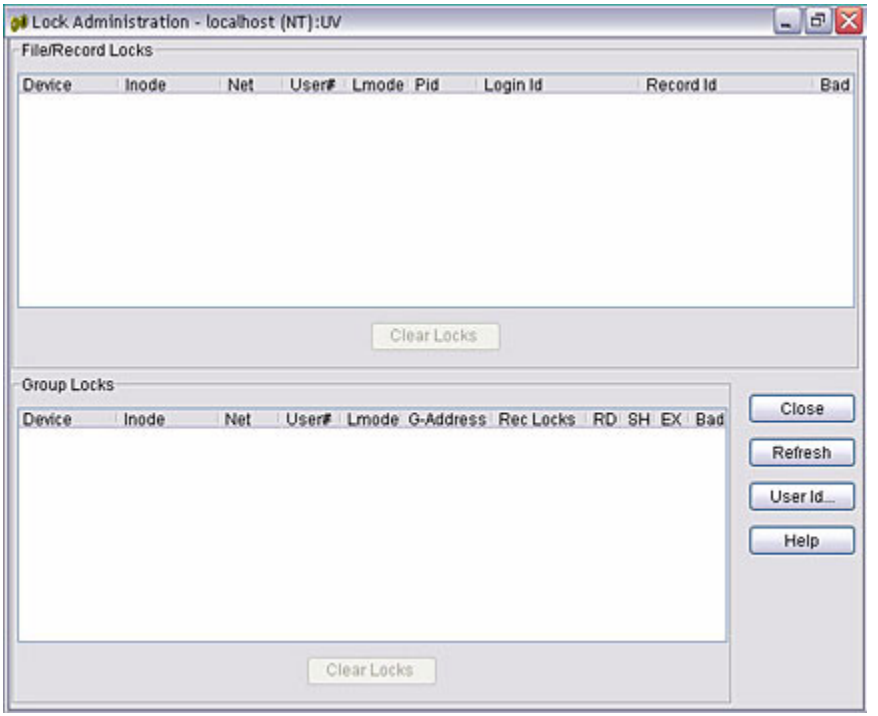

The **Lock Administration** dialog box is divided into to main areas, **File and Record Locks** and **Group Locks**.

Note: *The information displayed in the* ***Lock Administration*** *window is a snapshot of the file, record, and group locks when you activated the* ***Locks*** *option. To view the current state of locks, click* ***Refresh****.*

File and Record Locks

UniAdmin displays the following information in the File/Records Locks portion of the **Lock Administration** dialog box.

Field	Description
Device	A number that identifies the logical partition of the disk where the file system is located.
inode	A number that identifies the file that is being accessed.
Net	A number that identifies the host from which the lock originated. Zero (0) indicates a lock on the local machine.
User #	The user ID.
Lmode	The lock semaphore number and the type of lock. For record locks, there are two settings: n **RU** for an update (READU) lock n **RL** for a shared (READL) lock For file locks, there are six settings: n FS for a shared (READL) lock. n IX for a shared (READL) lock with intent to acquire an exclusive file lock. n FX for an exclusive (READU) file lock. n XU for an exclusive (READU) lock set by CLEAR.FILE. n CR for a shared (READL) file lock set by RESIZE. n XR for an exclusive (READU) file lock set by RESIZE or UVFIXFILE.
Pid	The process ID number.
Login ID	The login ID.
Record ID	The name of the record that is locked.
Bad	The bad lock indicator. If this column is empty, the lock is good. If this column contains an asterisk (*), the lock is bad.

File/Record Lock Display

Clearing Locks

You can clear a single file, record, or group lock, or all the locks for a particular user using the **Lock Administration** dialog box.

Complete the following steps to clear a file or record lock:

1. Select the lock from the **File/Record Locks** list.
2. Click **Clear Lock**. The **Lock Administration** dialog box is updated.

Complete the following steps to clear a group lock:

1. Select the lock from the **Group Locks** list.
2. Click **Clear Group Lock**. The **Lock Administration** dialog box is updated.

Complete the following steps to clear all the lock for a particular user:

1. Click **User Id...** The **Clear User Locks** dialog box appears, as shown in the following example:

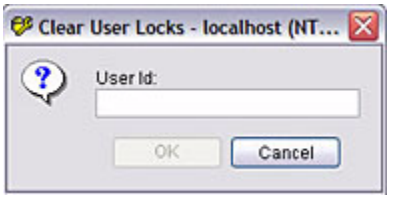

2. Enter the user ID for which you want to clear locks in the **User Id** box.
3. Click **OK** to return to the **Lock Administration** dialog box, or click **Cancel** to exit the dialog box without saving changes.

4. From the **Lock Administration** dialog box, click **Clear Locks**. The following window appears:

Click **Yes** to clear the locks, or click **No** to exit. The **Lock Administration** dialog box is updated.

Group Locks

UniAdmin displays the following information in the **File/Records Group Locks** portion of the **Lock Administration** dialog box.

Field	Description
Device	A number that identifies the logical partition of the disk where the file system is located.
Inode	A number that identifies the file that is being accessed.
Net	A number that identifies the host from which the lock originated. Zero (0) indicates a lock on the local machine.
User #	The user ID.
Lmode	The lock semaphore number and the type of lock. There are five settings: n EX for an exclusive update (READU) lock. n SH for a shared (READL) lock. n RD for a read lock. n WR for a write lock. n IN for an information lock.
G-Address	The logical disk address of the group. The value is 1 for a type 1 or type 19 file. Any other value is represented in hexadecimal format.
Rec Locks	The number of locked records in the group.
RD	The number of readers in the group.
SH	The number of shared (READL) group locks.
EX	The number of exclusive update (READU) locks.
Bad	The bad lock indicator. If this column is empty, the lock is good. If this column contains an asterisk (*), the lock is bad.

Group Locks Information

Chapter

13

Administering Network Services

Administering the UniRPC on UNIX Systems 13-3
Specifying the RPC Port Number 13-4
Changing the RPC Port Number 13-5
Starting and Stopping the UniRPC Daemon on UNIX 13-5
Maintaining the *hosts* File 13-6
Managing a Windows Platform Telnet Session 13-9
Modifying the Telnet Session Parameters 13-11
Administering Users 13-14
Adding a New User 13-15
Starting Services 13-20

This chapter describes how to administer network services through UniAdmin.

The tasks you can perform through Network Services Administration depend on the server to which you are connected. On UNIX servers, you can change the UniRPC daemon port number, and maintain the *hosts* file. On Windows platforms, you can change the port number and configuration parameters for the telnet session.

Administering the UniRPC on UNIX Systems

Use the **UniVerse Network Services** dialog box to define the RPC port number and maintain the *hosts* file for a UNIX server.

Select one of the following methods to access the **Network Services** dialog box:

- From the **UniAdmin** window, double-click **Network Services**, then click **UniRPC**.
- From the **UniAdmin** menu, select ***Network Services,*** then click **UniRPC**.
- From the **UniAdmin** toolbar, click the **Network Services** icon, as shown in the following example.

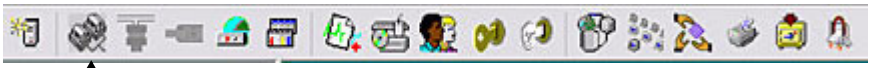

Network Services

The **UniVerse Network Services** dialog box appears, as shown in the following example:

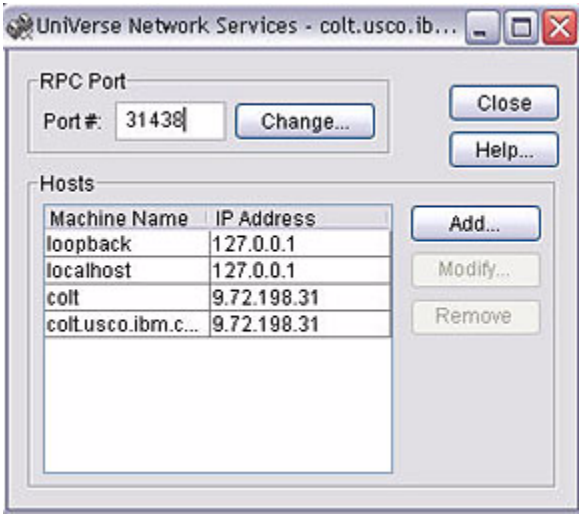

Specifying the RPC Port Number

The UniRPC (remote procedure call) utility enables UniVerse to communicate with server processes on remote systems. UniAdmin uses this utility to communicate with the UniVerse servers you previously defined.

UniRPC comprises the following UniVerse components:

- UniRPC daemon, *unirpcd*, which receives requests from remote machines for services, and starts those services.
- UniVerse BASIC programs for administering the UniRPC.

Before you can use the UniRPC daemon, you must specify the number of the port for it to use.

In the **Port #** box, enter a valid port number for use by the UniRPC daemon.

Changing the RPC Port Number

If you want to change the port number, click **Change**. A dialog box similar to the following example appears:

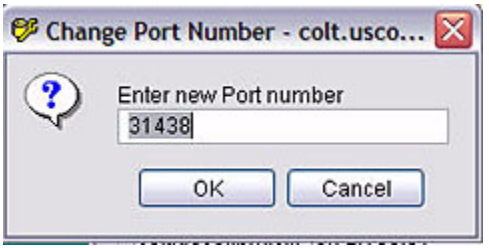

Enter the new port number, then click **OK**, or click **Cancel** to exit without saving changes.

If you change the UniRPC port number, you must also change it on all other systems that communicate via UniRPC.

To use the new port number, you must restart the UniRPC daemon.

Starting and Stopping the UniRPC Daemon on UNIX

You cannot start or stop the UniRPC daemon from within UniAdmin, because it is used to connect to UniVerse servers. You must start and stop the UniRPC daemon from the UniVerse server itself. On UNIX systems you can use the System Administration menus to start and stop the UniRPC daemon.

Starting the UniRPC Daemon

Complete the following steps to start the UniRPC daemon:

1. Choose **Rpc Administration** from the **Package** menu, then select **Start the rpc daemon**.

2. At the prompt, execute one of the following tasks to handle any error messages:
   - Enter the name of the file to which you want to send all error and system messages.
   - Enter a space to display messages on your terminal screen.
   - Press ENTER if you do not want to display or save messages.
3. At the next prompt, click **Yes** to start the UniRPC daemon, or click **No** to return to the **Rpc Administration Menu**.

Once you start the UniRPC daemon, it automatically restarts whenever you start UniVerse.

Stopping the UniRPC Daemon

Complete the following steps to stop the UniRPC daemon:

1. Choose **Rpc administration** from the **Package** menu, then choose **Halt the rpc daemon**.
2. At the prompt, check **Yes** to stop the UniRPC daemon, or **No** to return to the **Rpc administration** menu.

Note: *Stopping the UniRPC daemon does not interrupt active UniRPC processes.*

Maintaining the *hosts* File

You can add, modify, or remove nodes in the *hosts* file from the **UniVerse Network Services** dialog box.

Adding a Node

To add a node to the *hosts* file, in the **Hosts** area of the **UniVerse Network Services** dialog box, click **Add**. The following dialog box appears:

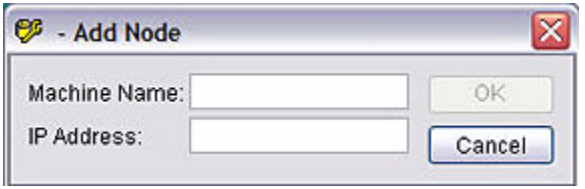

In the **Machine Name** box, enter the name of the node.

Enter the address of the node in the **IP Address** box.

Click **OK**. UniVerse checks the machine name and IP Address of the new node against existing entries in the *hosts* file. If the new node matches an existing entry, a message box appears. You must acknowledge the message before you can enter alternative values. If the new node details are unique, the new node definition is added to the *hosts* file, and the **Network Services** dialog box is updated.

Modifying a Node

To modify information about a node already defined in the *hosts* file, select the machine name you want to modify, then click **Modify**. A dialog box similar to the following example appears:

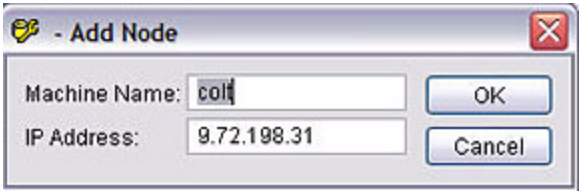

Make any necessary changes, then click **OK** to save the changes, or click **Cancel** to exit the program without saving changes. If the new node matches an existing entry, a message box appears. You must acknowledge the message before you can enter alternative values. If the new node details are unique, the new node definition is added to the *hosts* file, and the **Network Services** dialog box is updated.

Removing a Node

To remove a node from the *hosts* file, select the node you want to remove, then click **Remove**. A dialog box similar to the following example appears.

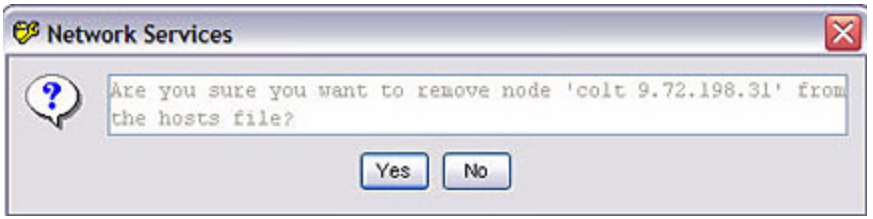

If you want to remove the node, click **Yes**. If you do not want to remove the node, click **No**.

Managing a Windows Platform Telnet Session

Select one of the following methods to access the **UniVerse Network Services** dialog box:

- From the **UniAdmin** window, double-click **Network Services**, then click **Telnet**.
- From the **UniAdmin** menu, select **Network Services,** then click **Telnet**.
- From the **UniAdmin** toolbar, click the **Network Services** icon, as shown in the following example.

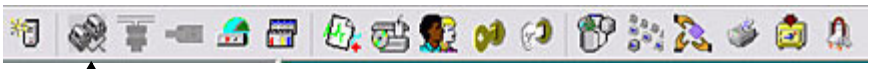

The **UniVerse Network Services** dialog box appears, as shown in the following example:

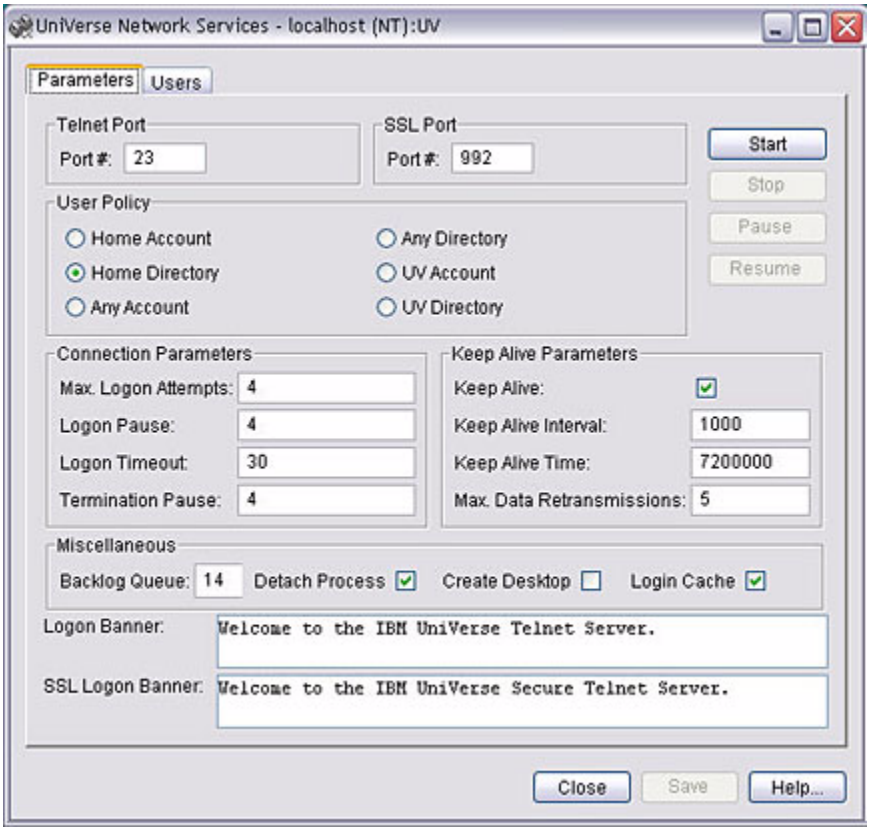

The **Network Services** dialog box contains the following fields and options:

- **Telnet Port #** – This field displays the TCP port that the telnet session uses. This is taken from the *services* file. If a *uvtelnet* entry exists in the *services* file, this is the number UniAdmin displays. If these entries do not exist in the *services* file, UniAdmin displays the default port number, 23.
- **User Policy** – The **User Policy** setting determines how the telnet session is used when a user makes a telnet connection.
- **Connection Parameters** – **Connection Parameters** are the current connection values for the telnet service. UniVerse stores these parameters in the Windows Registry on the Server.

- **Keep Alive Parameters** – **The Keep Alive** parameters determine intervals when UniVerse checks the viability of a network connection between the client and server.

Modifying the Telnet Session Parameters

You can modify any of the telnet session parameters from the **Network Services** dialog box.

Note: *To use the new settings, you must stop and restart the uvtelnet service.*

Changing the Telnet Session Port Number

To change the port number for the telnet session, enter the new port number in the **Port #** box. UniVerse stores the new port number as a *uvtelnet* entry in the *services* file when you click **Save**.

Defining the User Policy

As a UniVerse administrator, you can specify how all users use the telnet session. Valid user policies are:

- **Home Account** – On connection, users attach to their home directory. The home directory must be a valid UniVerse account.
- **Home Directory –** This is the default setting. Users connect to their home directory, but if the home directory is not a UniVerse account, UniVerse prompts users to set up the account.
- **Any Account –** Users can connect to any valid UniVerse account.
- **Any Directory –** Users can connect to any directory, but if the directory is not a UniVerse account, UniVerse prompts to set up the account.
- **UV Account** – Specifies that the user connects to an existing UniVerse account defined in the UV.LOGINS file.
- **UV Directory** – Specifies that the user connects to a directory defined in the UV.LOGINS file, and can create a UniVerse account in that directory if the directory is not already configured for UniVerse.

Note: *Administrators are prompted for the account to which they want to connect regardless of the User Policy setting.*

Setting the Telnet Connection Parameters

The four valid telnet connection parameters are:

- **Max. Logon Attempts** – Defines the number of failed log in attempts a user is allowed before the telnet connection is dropped. The default setting is 4.
- **Logon Pause** – If a logon attempt fails, the pause between logon attempts (in seconds). The default setting is 4 seconds.
- **Logon Timeout** – The time (in seconds) the system waits for a response to a logon prompt. As soon as this time limit is reached, the telnet connection is dropped. The default value is 30 seconds.
- **Termination Pause** – The amount of time UniVerse pauses after the final failed logon attempt before dropping the telnet connection. The default value is 4 seconds.

Setting Keep Alive Parameters

The Keep Alive feature determines when inactive connections can be disconnected. When a connection becomes inactive, keep-alive packets are periodically exchanged. When a number of consecutive packets remain unanswered, by default 20, the connection is broken.

- **Keep Alive Interval** – The interval, in milliseconds, separating keep alive retransmissions until a response is received. Once a response is received, the delay until the next keep alive transmission is controlled by the value of **Keep Alive Time**. After the number of retransmissions specified by **Max. Data Retransmissions** are unanswered, the connection aborts. The default value is 1000 (one second).
- **Keep Alive Time** – This parameter specifies how often TCP attempts to verify that an idle connection is still valid by sending a keep alive packet. If the connection is still valid, the remote system will acknowledge the keep alive transmission. The default value is 7,200,000 milliseconds (two hours).
- **Max. Data Retransmissions** – This parameter specifies the number of times TCP retransmits an individual data segment before aborting the connection. The retransmission timeout is doubled with each successive retransmission on a connection. It is reset when responses resume.

Miscellaneous Parameters

Set Backlog Queue Value

The maximum length of the queue of pending telnet connections. The default value is 14.

Select Detach Process

If you select this option, the UniVerse Telnet Service creates the UniVerse process (uv_tl_server.exe) as a detached process.

Select Create Desktop

If you select this option, the UniVerse Telnet Service creates its own WinStation/Desktop and assigns it to the UniData process.

Specify Logon Banner

You can specify the banner that users will see when they telnet to a host in the **Logon Banner** box.

Administering Users

The UV.LOGINS file resides in the UV account. It contains a list of users and the directories or UniVerse accounts they log on to when they first invoke UniVerse from a telnet session.

To maintain entries in the UV.LOGINS file, click the **Users** tab. The **UniVerse Users** dialog box appears, as shown in the following example.

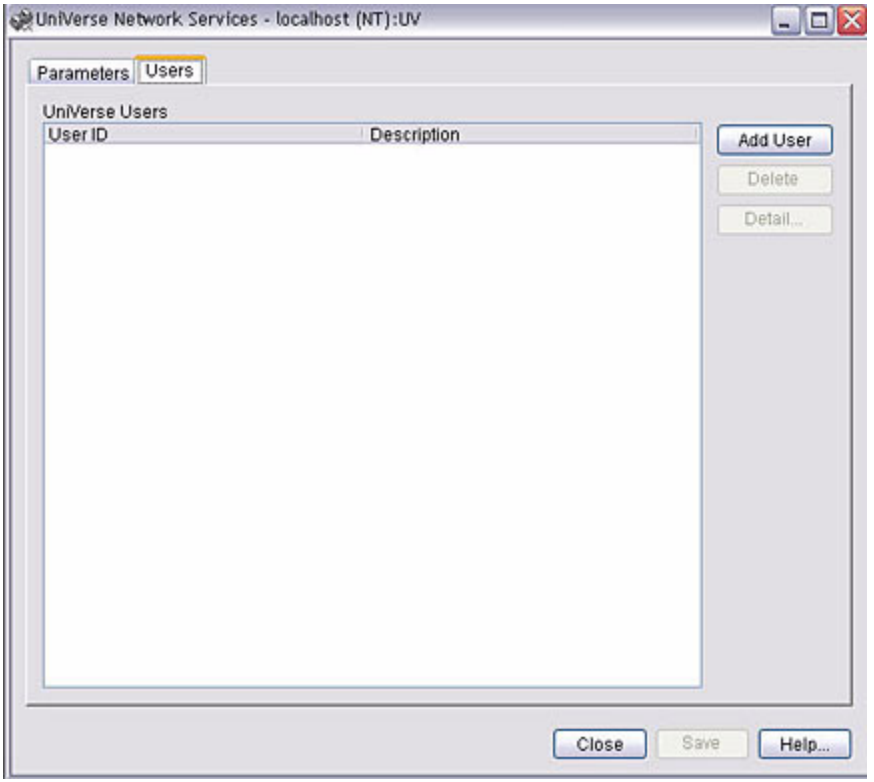

You can enter users logging on to the system both from the local machine and from domains. You can also maintain entries for users who have accounts on multiple domains with access to this system. You can specify the user's account either as a case-sensitive entry in the UV.ACCOUNTS file, or as a fully qualified path.

If the user logs on to the system using a local machine login ID, UniVerse uses the Local Machine entry. If the user logs on to the system through a domain, UniVerse uses the entry for the domain. If the user enters a login ID without a machine or domain name, UniVerse first uses a local machine login ID if it exists, and then checks domain login IDs.

Adding a New User

To add a new user, click **Add User** from the **UniVerse Users** dialog box. A dialog box similar to the following example appears:

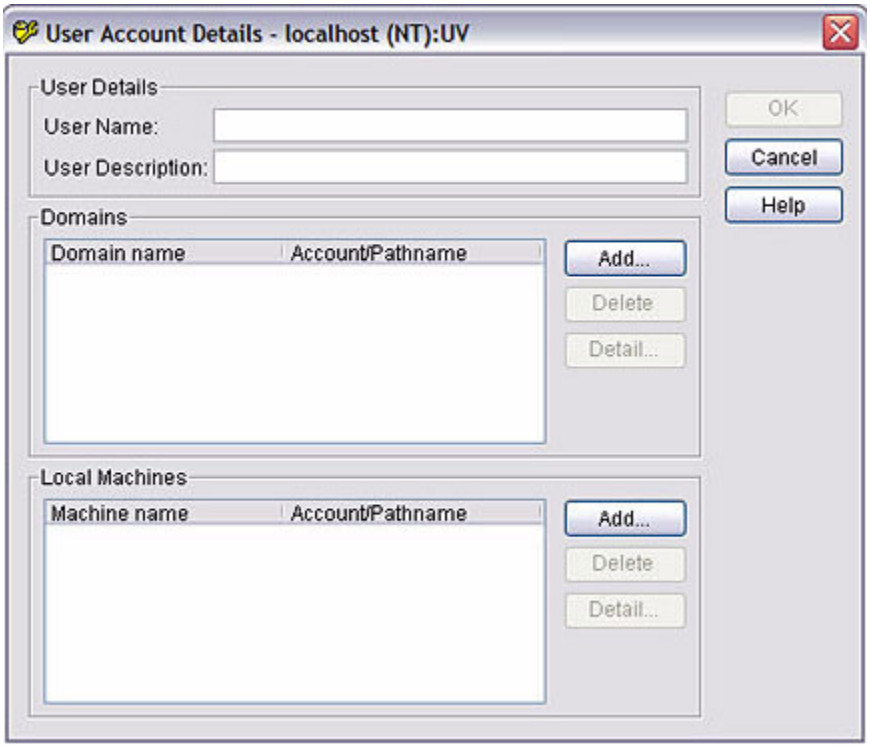

Add a Domain User

To add a domain user, in the **Domain** area, click **Add**. The **Domain Account Details** dialog box appears, as shown in the following example:

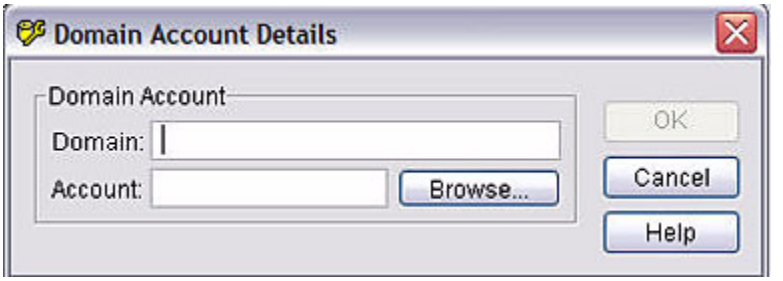

Enter the name of the domain to which you want the user to connect in the **Domain** box.

Enter the full path to the account to which the user is to connect, or click **Browse** to search for the account.

Click **OK** to save the information, or click **Cancel** to exit without saving changes. The user appears in the **Domain** area of the **User Account Details** dialog box, as shown in the following example:

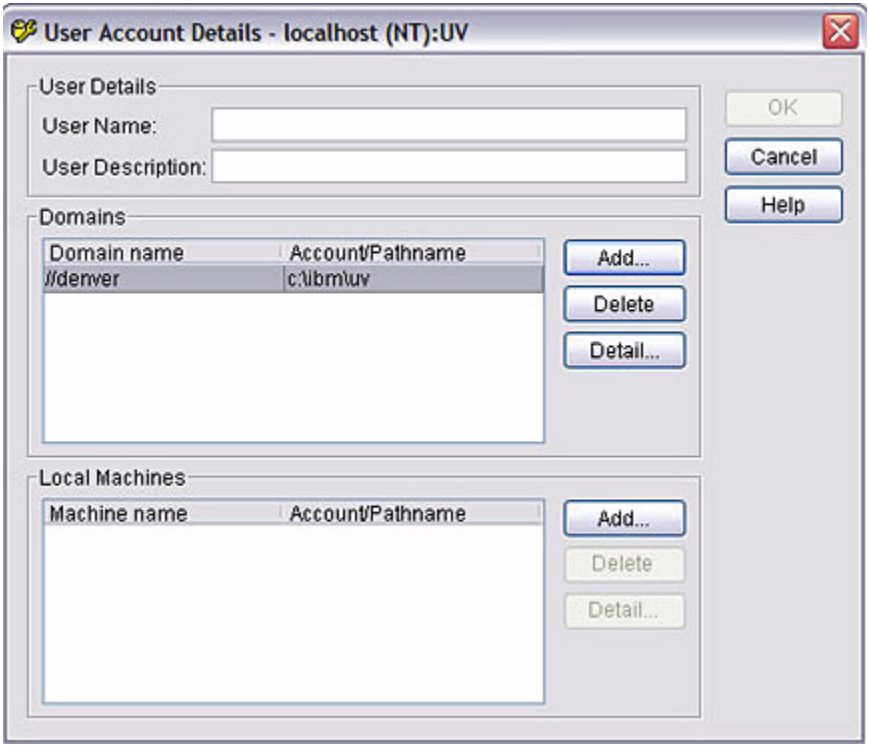

Adding a Local Machine User

To add a user to a local machine, in the **Local Machines** area of the **User Account Details** dialog box, click **Add**. The **Local Machine Account Details** dialog box appears, as shown in the following example:

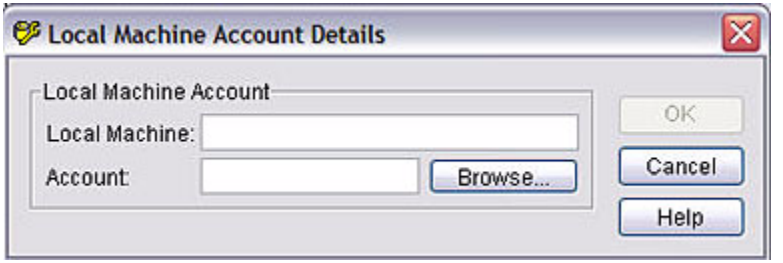

Enter the name of the local machine in the **Local Machine** box.

Enter the full path to the account to which the user is to connect, or click **Browse** to search for the account.

Click **OK** to save the information, or click **Cancel** to exit without saving changes. The user appears in the **Local Machines** area of the **User Account Details** dialog box, as shown in the following example:

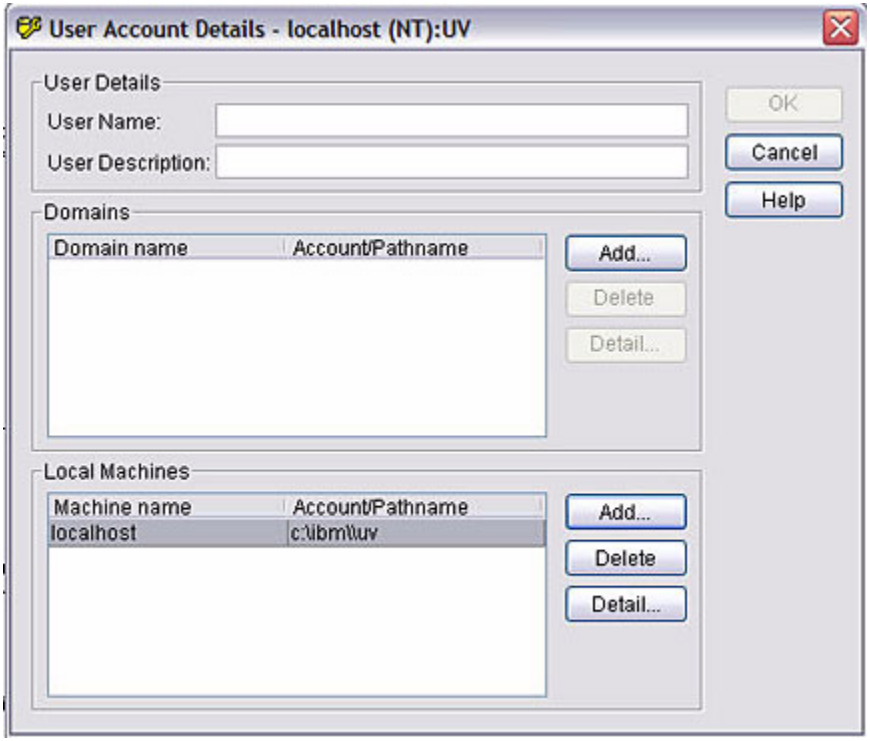

Starting Services

On Windows platforms, you can start and stop the following services from the **UniVerse Services** dialog box:

- universe
- uvtelnet
- unirpc
- hsrexec

Select one of the following methods to access the **UniVerse Services** dialog box:

- From the **UniAdmin** window, double-click **Network Services**, then click **NT Services**.
- From the **UniAdmin** menu, select **Network Services,** then click **NT Services**.

The **UniVerse Services** dialog box appears, as shown in the following example:

You can perform the following tasks from the **UniVerse Services** dialog box:

- To start a service, click the service you want to start, then click **Start.**
- To stop a service, click the service you want to stop, then click **Stop**.
- If you want to start all services, click **Start All**.

- If you want to stop all services, click **Stop All**.
- If you want to refresh the information displayed in the dialog box, click **Refresh**.

Click **Close** to exit the **UniVerse Services** dialog box.

Chapter

14

UniVerse Replication

About Replication 14-4
About Hot Standby 14-4
Setting Up Data Replication 14-5
Configuring and Managing Data Replication 14-6
Managing a Publishing System 14-7
Configuring the Publishing System 14-9
The Publishing Config Dialog Box 14-9
Starting and Stopping the Subscribing System 14-11
Publishing Files 14-12
Publishing Information File 14-18
Managing the Subscribing System 14-20
Configuring the Subscribing System 14-20
Starting and Stopping the Subscribing System 14-22
Creating a List of Publishing Systems for the Subscriber 14-22
Subscribing Files 14-24
Managing Hot Standby Operations 14-32
Configuring a Hot Standby Subscriber 14-32
Turning On Fail-Over Mode 14-34
Reconciling the Hot Backup with the Publisher 14-34
Subscribing Information Files 14-36
Some Restrictions 14-38
What to Do When Disk Space Fills Up 14-39
Removing Obsolete Replication Log Files 14-40
What to Do When Replication Fails 14-41

This chapter describes how to administer replication through UniAdmin.

The UniVerse data replication service provides two operations:

- Replication – Maintains one or more read-only copies of UniVerse files for the purpose of data backup or distribution. The copies can be on one or more computer systems.
- Hot Standby – A special case of replication in which a system that maintains copies of active UniVerse files can provide read/write versions of the replicated files to users, should the original files be unavailable.

Use the **UniVerse Replication Admin** dialog box to administer replication.

Select one of the following methods to access the **UniVerse Replication Admin** dialog box:

- From the **UniAdmin** window, double-click **Replication**.
- From the **UniAdmin** menu, select ***Admin,*** then click **Replication**.
- From the **UniAdmin** toolbar, click the **Replication Admin** icon, as shown in the following example.

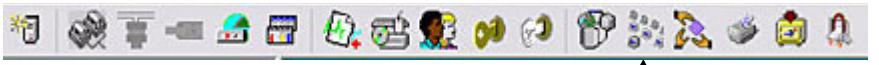

Note: *Replication works only with files created or resized on Release 9.4 or later of UniVerse. To replicate files created on older systems, resize the files on a valid release first.*

The **UniVerse Replication Admin** dialog box appears, as shown in the following example:

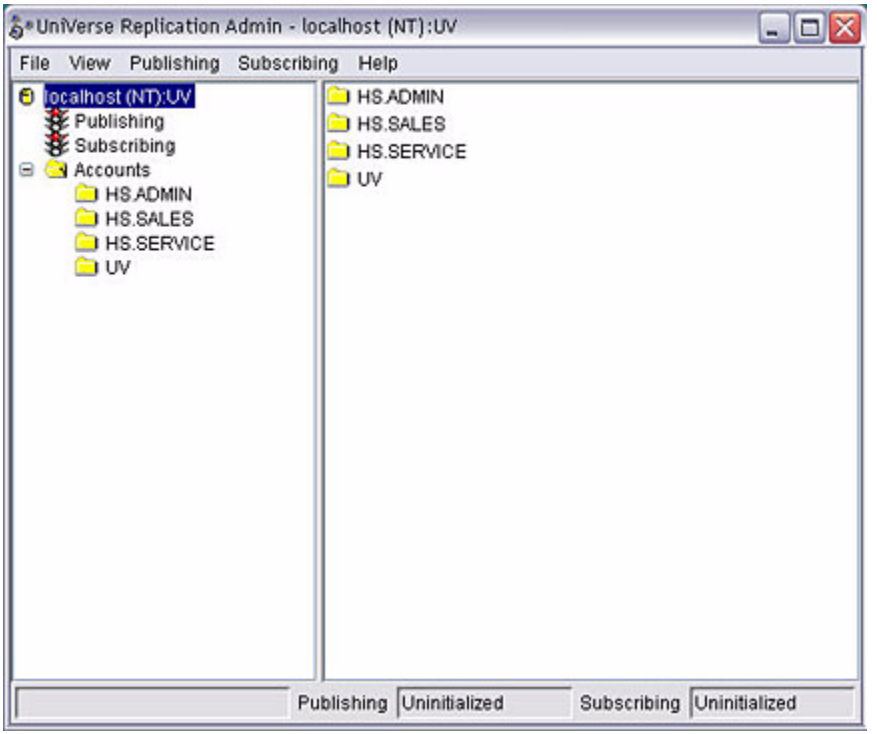

About Replication

UniVerse data replication provides an automatic and reliable way to deliver read-only copies of UniVerse files to other UniVerse systems.

The system where the source data resides is called the *publisher*. A system requesting copies of file updates from the publisher is called a *subscriber*.

Subscribing files on a subscriber are read-only. Users logged on to the subscriber cannot modify the replicated data.

Note: *A publisher can also be a subscriber of another publisher's files. However, a publisher cannot publish replicated (subscribing) files.*

About Hot Standby

You can configure data replication to create hot standby (also referred to as hot backup) functionality, with the publisher acting as the primary UniVerse server and the subscriber acting as the development system or dormant backup system. The hot standby subscriber must be a different computer system from the publisher.

If the publisher's disk integrity is compromised, or the hardware is unusable, you use UniAdmin to switch from using the publisher to using the hot standby system.

If a crash occurs on the publisher, you use UniAdmin on the hot standby system to switch to fail-over mode, thus making the subscriber read/write-enabled. Users can then log on to the hot standby system and run their applications as usual. Later, the administrator can reconcile the originally published files on the publisher with all updates and changes made to the files on the hot standby system.

Setting Up Data Replication

Before you can configure a subscriber, you must set up the replicated (subscribing) database on the subscriber. Next you configure the publisher and all subscribers. Finally, you start publishing on the publisher and subscribing on all subscribers.

Note: *You must configure the publisher and all subscribers before you start publishing and subscribing, otherwise the replication system will start up in the crashed state.*

Complete the following steps to set up a data replication system:

1. Before setting up and configuring your publishing and subscribing systems, make sure that no users can access the files you intend to publish.
2. On the publisher, resize all files you want to replicate that were created on UniVerse Release 9.3 or earlier.
3. On the subscriber, create the UniVerse accounts to contain the replicated (subscribing) files.
4. On the subscriber, create the subscribing UniVerse files that will store the replicated data. Permissible file types are types 2 through 18 (static hashed files), type 25 (B-tree files), type 30 (dynamic hashed files), and distributed files. These files need not have the same file names as the corresponding files on the publisher, nor need they be of the same file type. The structures of distributed files on the publisher and subscriber should match.

   Note: *You cannot replicate type 1 or type 19 files, or secondary indexes.*
5. Before configuring and starting the publishing system, ensure that:
   - Data in the subscriber's files is identical to data in the publisher's files.
   - No users can access the replicated files.
6. Set the value of the UDRMODE configuration parameter to 1, then stop and restart UniVerse.
7. Configure the publisher and all subscribers.
8. Start publishing on the publisher, and start subscribing on all subscribers.

Configuring and Managing Data Replication

There are two parts to data replication:

- A publishing system keeps track of updates to a published database and makes information about those updates available to subscribing systems across the network. You replicate read-only copies of published UniVerse files on subscribing systems.
- A subscribing system receives information about updates on published source files across the network.

Note: *A special case of subscribing called hot standby lets you track updates to published UniVerse files on a subscribing system and make those subscribing copies into fully read/write files in the event the primary publishing system becomes temporarily unavailable.*

Managing a Publishing System

You can publish the following file types:

- Types 2 through 18 (static hashed files)
- Type 25 files (B-tree files)
- Type 30 files (dynamic hashed files)
- Multiple data files
- Distributed files

You cannot publish the following:

- Type 1 or type 19 files
- Secondary indexes

Before configuring and starting the publishing system, make sure that:

- Data in the subscriber's files is identical to data in the publisher's files.
- No users can access the replicated files.

Complete the following steps to configure the publishing system:

1. Set the value of the UDRMODE configurable parameter to 1. For this procedure, see *The UniVerse Configuration Editor*.
2. To unload shared memory, shut down UniVerse, and execute one of the following steps:
   - From the **Start** menu, choose **Programs**, then choose **IBM U2**, then **UniVerse**, then click **UniVerse Control**. Click **Stop All Services** to stop all UniVerse services, then click **OK** to exit the UniVerse Control panel.
   - From the **Start** menu, choose **Settings**, then choose **Control Panel**, then choose **Services**. Double-click **UniVerse Resource Service**, then click **Stop**. Click **OK** in response to the message prompting you to confirm that you want to shut down other UniVerse services.

   Wait for all processes related to UniVerse to shut down. You can use the Windows Task Manager to monitor these processes.

3. Start UniVerse again by executing one of the following tasks:
   - From the **Start** menu, choose **Programs**, then choose **IBM U2**, then **UniVerse**, then click **UniVerse Control**. Click **Start All Services** to start all UniVerse services, then click **OK** to exit the UniVerse Control panel.
   - From the **Start** menu, choose **Settings**, then choose **Control Panel**, then choose **Services**. Double-click **UniVerse Resource Service**, then click **Start**. In the **Startup Type** list, choose **Automatic**. Repeat these steps for the UniRPC and Telnet services.
4. Next, configure the publisher and all subscribers.
5. Start publishing on the publisher, and start subscribing on all subscribers.

Note: *After you configure publishing and subscribing, no published or subscribing files can be modified until you start publishing and subscribing on all systems.*

Configuring the Publishing System

You must configure the publishing system before you can publish its files. You can configure the following options for publishing systems:

- The location of the directory containing replication log files.
- The size of the replication log files.
- The number of the current replication log file.
- Whether or not publishing services start up when UniVerse starts.

The Publishing Config Dialog Box

From the **UniVerse Replication Admin** menu, click **Publishing**, then click **Configure Publishing**. The **Publishing Config** dialog box appears, as shown in the following example:

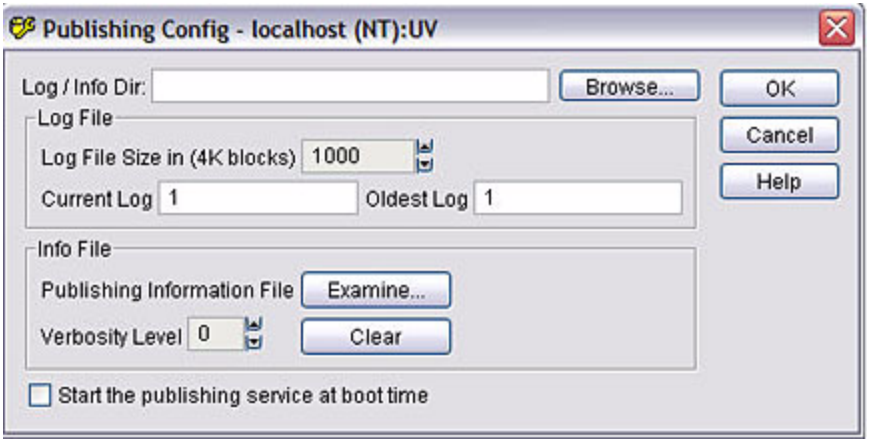

The following table describes the fields in the **Publishing Config** dialog box.

Field	Description	Action
Log/Info Dir.	Indicates the location of the replication log files. This directory also includes a publishing information file, which stores messages pertaining to the publishing system.	Click **Browse** to locate and specify a directory for replication log files and the publishing information file.
Log File Size	Specifies the current size of the log file in multiples of 4K blocks.	Set the size of the log file.
Current Log	Indicates the number of the current log file.	Set the number of the current log file.
Oldest Log	Indicates the number of the oldest log file.	Set the number of the oldest log file.
Info File	The publishing information file is useful for debugging. Normally you need not log information to this file. We recommend you log information to this file only on the advice of an IBM support specialist.	Click **Examine** to view the contents of the publishing information file. Click **Clear** to clear the contents of the publishing information file.
Verbosity Level	Determines how much information to log to the publishing information file.	Set the verbosity level to a number between 1 and 9 to activate logging to the publishing information file. Set the level to 0 to deactivate information logging.
Start the publishing service at boot time	Toggles whether or not to start the publishing service at boot time.	Select the check box to start the publishing service at boot time. If you do not want to do this, make sure the check box is clear.

Publishing Config Dialog Box

Complete the following steps to configure a publishing system:

1. In the **Log/Info Dir.** box, enter the full path of an existing directory to use as the log file directory, or click **Browse** to search for the directory.

   The log directory contains the log files that store updates to published files before they are sent to subscribers for replication. The log directory also contains an information log file.

   Note: *UniVerse creates log files automatically.*

2. In the **Log File Size in (4K blocks)** box, select the size of the replications log files from the list. The default is 1000.
3. In the **Current Log** box, you can optionally specify the number of the current log file. If UniVerse finds an existing log file, it increments its number.
4. In the **Oldest Log** box, you can optionally specify the number of the oldest log file. To determine which log file is the oldest, complete the following steps:
   a. In an MS-DOS window, change directories to the location of your log/info directory.
   b. List all files named uvdrlg*n*, where *n* is a unique number. The oldest log file is the uvdrlg*n* file, where *n* is the lowest number in the set.
5. In the **Verbosity Level** box, you can optionally select the verbosity level for the publishing information file. Select a value from 0 through 9. The default is 0.

   Note: *Use the publishing information file only for debugging.*

6. If you want to start the publishing service at boot time, select the **Start the Publishing Service at boot time** check box.
7. Click **OK** to save your changes and exit the **Publishing Config** dialog box. Click **Cancel** to exit the dialog box without saving changes.

Starting and Stopping the Subscribing System

To start the subscribing system, from the **Subscribing** menu, click **Start Subscribing**. The status bar displays **Enabled**.

To stop the subscribing system, from the **Subscribing** menu, click **Stop Subscribing**. The status bar displays **Uninitialized**.

Publishing Files

You can publish:

- One file at a time
- Selected files
- All files in an account

Publishing One UniVerse File

Complete the following steps to publish one UniVerse file in an account:

1. From the left pane of the **UniVerse Replication Admin** window, click **Accounts**, then select one of the accounts listed. A list of UniVerse files you can replicate appears in the right pane.

2. From the right pane, select a file, then execute one of the following steps:
   - Choose **Publish** from the **Publishing** menu.
   - Click **Publish** on the toolbar.
   - Right-click the file, then choose **Publish** from the menu.

   The **Publish** dialog box appears with the Filename, Dictname, Account, and File Type fields populated, as shown in the following example:

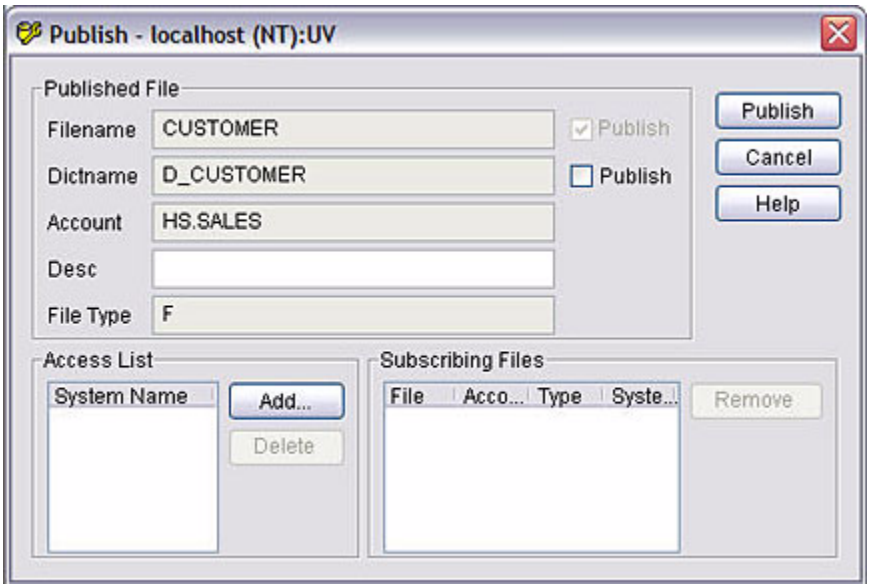

3. (Optional) Select the **Publish** check box next to the **Dictname** box if you want to replicate the file dictionary as well as the data file.
4. (Optional) Enter a description of the file in the **Desc** box.

5. By default, all subscribing systems can access any files you publish. To restrict access to a finite list of published files, add a system name to the **Access List**. Complete the following steps to add a system name:

   - Click **Add**. The **Add System** dialog box appears, as shown in the following example.

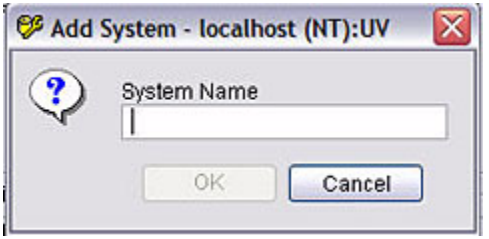

   - Enter the name of the system in the **System Name** box, then click **OK**, or click **Cancel** to exit without saving changes.
   - Repeat these steps until you specify all systems you want to have access to the publication.
   - Click **OK** to return to the **Publish** dialog box.

6. Click **Publish** to publish the file and exit the **Publish** dialog box. The letter "P" appears on the file icon in the **Replication** window.

   Click **Cancel** to exit the **Publish** dialog box without publishing the file.

Publishing Multiple UniVerse Files

Complete the following steps to publish selected UniVerse files:

1. To publish selected UniVerse files:

   - From the left pane, double-click **Accounts**.
   - Select on the accounts listed. A list of UniVerse files you can publish appears in the right pane.
   - To select a range of files, select the first file in the range, then hold down the **Shift** key and select the last file in the range.
   - To select random files, hold down the **Ctrl** key and select the files you want to publish.

2. To publish all files in an account:
   - From the left pane, select **Accounts**. A list of accounts appears in the right pane.
   - From the right pane, select the account whose files you want to publish.
3. Access the **Account Publishing** dialog box in one of the following ways:
   - Choose **Publish** from the **Publishing** menu.
   - Click **Publish** on the toolbar.
   - From the right pane, right-click the selected files or the selected account, then choose **Publish** from the menu.

   The **Account Publishing** dialog box appears, as shown in the following example:

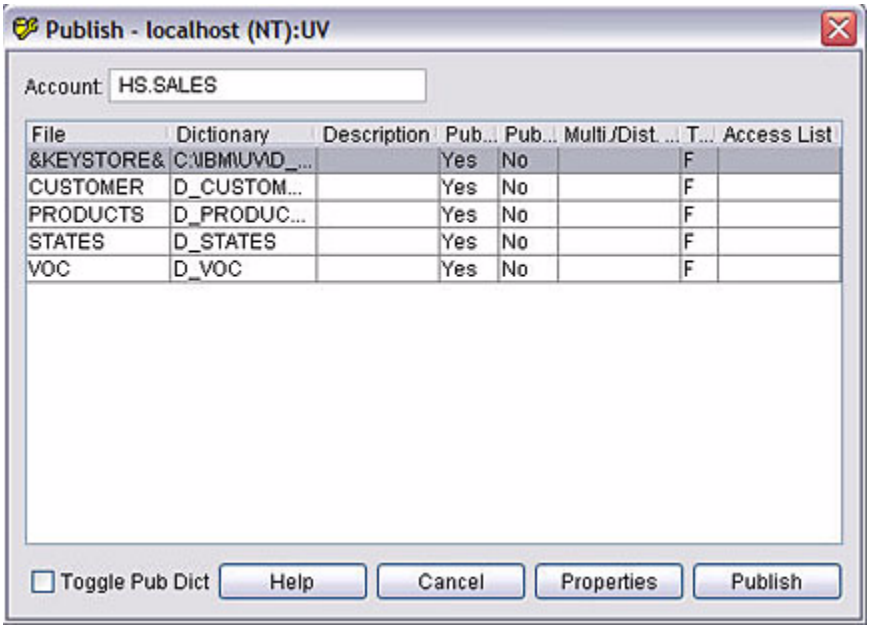

4. You can now execute any of the following tasks:
   - Click **Publish** to publish all selected files.
   - Select the **Toggle Pub Dict** check box to publish all selected data files and their file dictionaries.
   - Double-click a file in the list, or select a file and click **Properties** to display the **Publish** dialog box. This enables you to customize publication of the selected file. You can modify the following:
     - The file description.
     - Whether or not to publish the file dictionary.
     - The list of systems that can subscribe to this file.
5. Click **Cancel** to exit the **Account Publishing** dialog box without publishing any files.

Viewing and Modifying Published Files

Complete the following steps to view a list of currently published files:

1. From the left pane, double-click **Publishing**, then select one of the accounts listed. A list of published UniVerse files in the account appears in the right pane. The **Published Dictionary** column lists any published file dictionaries.

   File type is one of the following:

Type	Description
File	Hashed file (types 2 through 18), B-tree file (type 25), dynamic file (type 30).
Q Pointer	File referenced by a Q-pointer in the VOC file.

Published UniVerse File Types

2. From the right pane, right-click a file and choose **Properties** from the menu to view information about the file. The **Publish** dialog box appears. This dialog box displays information about the published file. You can modify the following fields:
   - Publish file dictionary check box
   - File description
   - Access list
   - Subscriptions list
3. Select **Publish** next to **Dictname** to replicate the file dictionary as well as the data file.
4. Enter a description of the file in the **Desc** field.
5. By default, all subscribing systems can access any files you publish. To restrict access to a finite list of published files, add a system name to the **Access List**.
   - Click **Add.** The **Add System** dialog box appears, and shown in the following example.

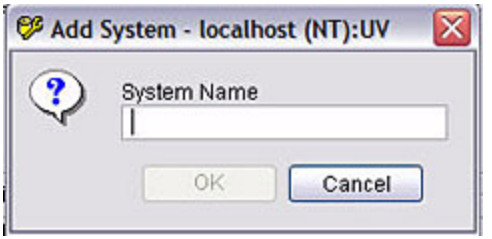

   - Enter the name of the system in the **System Name** box, then click **OK.**
   - Repeat these steps until you have added all systems you want to have access to the publication.
6. All files that subscribe to this are listed in the **Subscribing Files** list. To unsubscribe a file, select it and click **Remove**. If there are no other subscribing files, logging of file updates ceases immediately.
7. Click **OK** to save your changes and exit the **Publish** dialog box. Click **Cancel** to exit the **Publish** dialog box without saving changes.

Unpublishing Files

When you unpublish a file, all subscriptions to that file are cancelled. Complete the following steps to unpublish files:

1. Execute one of the following steps:

   To select individual files:

   - From the left pane, double-click **Publishing** or **Accounts**.
   - Select one of the accounts listed. A list of published UniVerse files in that account appears in the right pane.
   - To select a range of files, select the first file in the range, hold down the **Shift** key, then select the last file in the range.

     To select random files, hold down the **Ctrl** key and select the files you want to unpublish.

   To select all files in an account:

   - From the left pane, select **Publishing** or **Accounts**. A list of accounts appears in the right pane.
   - From the right pane, select the account whose files you want to unpublish.

2. Next, do one of the following:

   - Choose **Unpublish** from the **Publishing** menu.
   - Click **Unpublish/Unsubscribe** on the toolbar.
   - From the right pane, right-click the selected files or the selected account, then choose **Unpublish** from the menu.

3. When prompted if you are sure you want to unpublish the files, chick **Yes** or **No**. If you click **Yes**, all selected files are unpublished, the file names are removed from the publishing list, and all subscriptions to the files are cancelled.

Publishing Information File

The publishing information file is useful for debugging. Normally you do not need to log information to this file. We recommend you log information to this file only on the advice of an IBM support specialist.

To activate logging using the publishing information file, set the verbosity level to a number between 1 and 9. To deactivate logging, set the verbosity level to 0. 0 is the default setting.

Setting the Verbosity Level

Complete the following steps to set the verbosity level:

1. Select **Configure** from the **Publishing** menu. The **Publishing Config** dialog box appears.
2. Under **Info File**, set the **Verbosity Level** to a number between 0 and 9. The higher the verbosity level, the more information is logged.
3. Click **OK** to exit the **Publishing Config** dialog box.

Viewing Information Files

Complete the following steps to view the contents of the publishing information files:

1. Choose **Configure** from the **Publishing** menu. The **Publishing Config** dialog box appears.
2. Under **Info File**, click **Examine**. The **Publishing Info File** window appears, listing the contents of the Publishing Information file.

Clearing Information Files

Complete the following steps to clear the contents of the Publishing Information file:

1. Choose **Configure** from the **Publishing** menu. The **Publishing Config** dialog box appears.
2. Under **Info File**, click **Clear**. The contents of the Publishing Information file are cleared.

Managing the Subscribing System

The subscribing system handles the configuring and enabling of subscribing operations and the distribution of replicated records.

Configuring the Subscribing System

You must configure the subscribing system before you can subscribe to published files. You can configure the following options for the subscribing system:

- The location of the subscribing information files.
- Whether or not the subscribing service starts up when UniVerse starts.

From the **UniVerse Replication Admin** window, from the **Subscribing** menu, click **Configure Subscribing**. The **Subscription Config** dialog box appears, as shown in the following example:

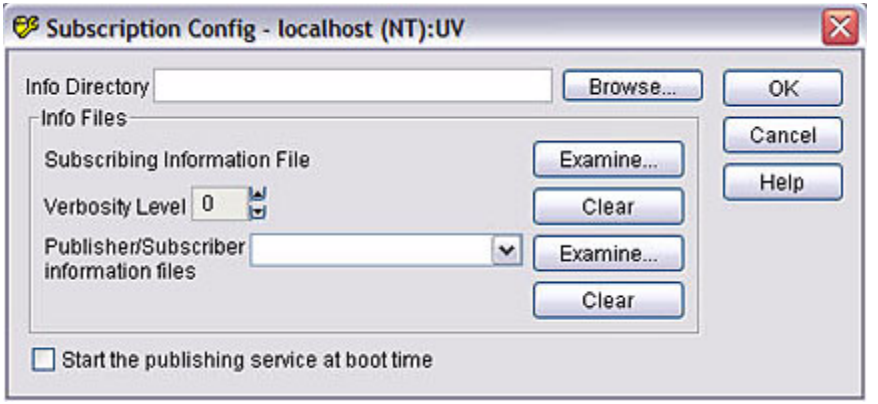

The following table describes the settings you can configure in the **Subscription Config** dialog box.

Field	Description	Action
Info Directory	Specifies the location of the subscribing information files, which store messages pertaining to the subscribing system.	Click **Browse** to locate and specify a directory for the subscribing information files.
Subscribing Information File	Logs information about the subscriber. This file is useful for debugging. Normally, you need not log information to this file. We recommend you log information to this file only on the advice of an IBM support specialist.	Click **Examine** to view the contents of the subscribing information file. Click **Clear** to clear the contents of the subscribing information file.
Verbosity Level	Determines how much information to log to the subscribing information files.	Set the level to a number between 1 and 9 to activate logging to the subscribing information files. Set the level to 0 to deactivate information logging.
Publisher/Subscriber Information Files	Logs information about connections between subscribing and publishing systems. This file is useful for debugging. Normally, we recommend you log information to this file only on the advice of an IBM support specialist.	Click **Examine** to view the contents of the publisher/subscriber information file. Click **Clear** to clear the contents of the publisher/subscriber information file.
Start the subscribing service at boot time	Toggles whether or not to start the subscribing service at boot time.	Select the check box to start the subscribing service at boot time.

Subscription Config Dialog Box Settings

Complete the following steps to configure a subscribing system:

1. In the **Info Directory** box, enter the path of an existing directory to store the information file directory, or click **Browse** to search for the directory. The Info directory contains the subscription information files.
2. You can optionally specify the verbosity level for the subscribing information files. Valid values are 0 through 9. The default is 0.

Note: *Use the subscribing information files only for debugging purposes. For details about these files, see Subscribing Information Files.*

3. If you want to start the subscribing service at boot time, select the **Start the publishing service at boot time** check box.
4. Click **OK** to save your changes and exit the **Subscription Config** dialog box, or click **Cancel** to exit without saving changes.

Starting and Stopping the Subscribing System

To start the subscribing system, choose **Start Subscribing** from the **Subscribing** menu. The status bar displays Enabled.

To stop the subscribing system, choose **Stop Subscribing** from the **Subscribing** menu. The status bar displays Uninitialized.

Creating a List of Publishing Systems for the Subscriber

Before a subscribing system can subscribe to published files, you must first create a list of publishing systems available to the subscriber. Next, you subscribe to the files published by these systems.

Complete the following steps to add a publishing system to the subscriber's list of publishing systems:

1. Choose **Systems** from the **Subscribing** menu. The **Publishing Systems** dialog box appears, as shown in the following example:

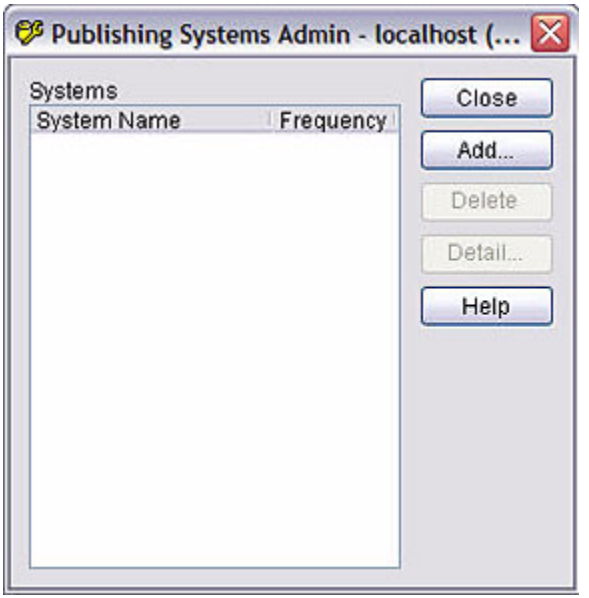

2. Click **Add.** The **System Details** dialog box appears, as shown in the following example:

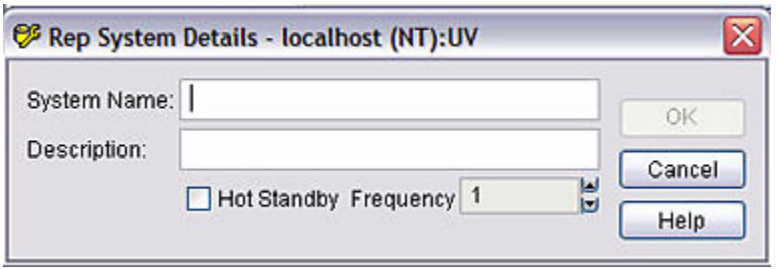

3. Enter the name of the publishing system in the **System Name** box.
4. (Optional) Enter a description of the system in the **Description** box.
5. Select a file update frequency from the **Frequency** list. This is the number of minutes to elapse between logging updates. Valid values are between 1 and 1,440 minutes (24 hours). The default is 1.

6. (Optional) If you want the subscriber to be the hot standby system for this publishing system, select the **Hot Standby** check box. For information about hot standby operations, see "Configuring a Hot Standby Subscriber" on page 32.
7. Click **OK**. The publishing system is added to the **Systems** list, as shown in the following example:

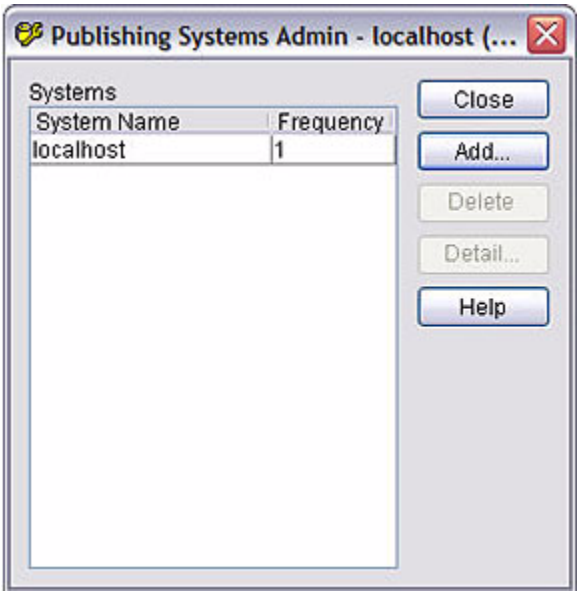

8. Repeat steps 2 through 7 to add the names of all publishing systems to which you want to subscribe.
9. When you finish, click **Close** to exit the **Publishing Systems** dialog box.

Subscribing Files

You can subscribe:

- One file at a time
- Selected files
- All files in an account

Subscribing One UniVerse File

Complete the following steps to subscribe one UniVerse file in an account:

1. From the left pane, double-click Accounts, then select one of the accounts listed. A list of UniVerse files you can subscribe appears in the right pane.
2. From the right pane, select a file, then do one of the following:
   - Choose Subscribe from the Subscribing menu.
   - Click Subscribe on the toolbar.
   - Right-click the file and choose **Subscribe** from the menu.

   The Subscribe dialog box appears with the Filename, Dictname, Account, and File Type fields filled in for you, as shown in the following example:

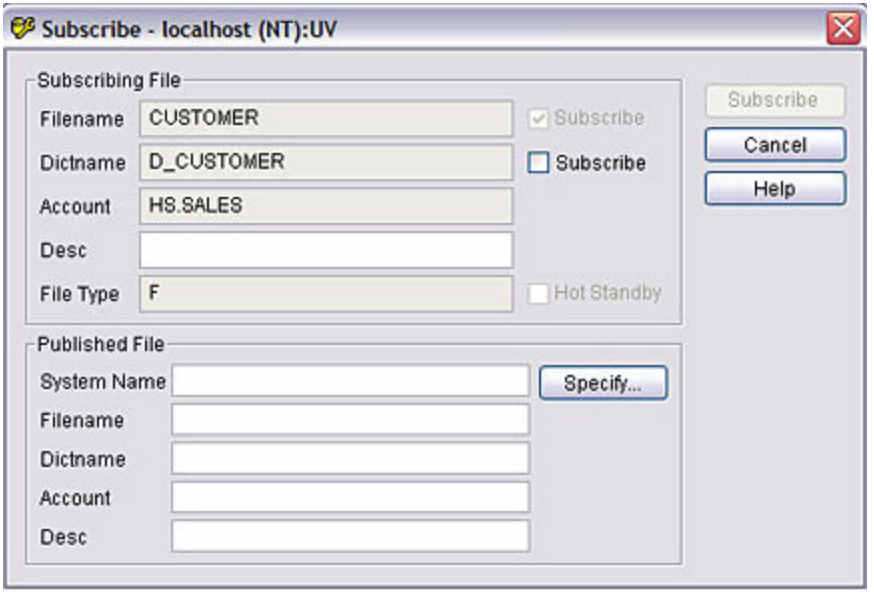

3. (Optional) Select the **Subscribe** check box, located next to the **Dictname** box, if you want to subscribe the file dictionary as well as the data file.
4. (Optional) Enter a description of the file in the **Desc** box.

5. Click **Specify** to display the **Available Publications** dialog box.
   - Double-click the name of a publisher to list all accounts with published files.
   - Double-click the name of an account to list all its published files.
   - Click the name of the file to which you want to subscribe. Information about the published file appears at the bottom of the **Subscribe** dialog box.
6. Click **OK** to subscribe the file and exit the **Subscribe** dialog box. The letter "S" appears on the file icon in the **Replication** window.

   Click **Cancel** to exit the **Subscribe** dialog box without subscribing the file.

Subscribing Multiple UniVerse Files

To subscribe selected UniVerse files or all files in an account:

1. Execute one of the following steps:

- To subscribe selected files:
  - From the left pane, double-click Accounts.
  - Select one of the accounts listed. A list of UniVerse files you can subscribe appears in the right pane.
  - To select a range of files, select the first file in the range, then hold down the Shift key and select the last file in the range.
  - To select random files, hold down the Ctrl key and select the files you want to subscribe.
- To subscribe all files in an account:
  - From the left pane, select Accounts. A list of accounts appears in the right pane.
  - From the right pane, select the account whose files you want to subscribe.

2. Execute one of the following steps:

- Choose Subscribe from the Subscribing menu.
- Click Subscribe on the toolbar.

- From the right pane, right-click the selected files or the selected account, then choose Subscribe from the menu.

  The Auto subscribe dialog box appears, as shown in the following example:

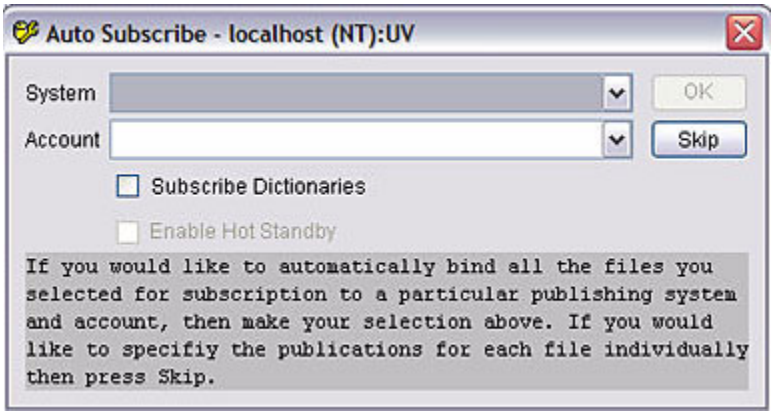

***Note**: Use this dialog box only if the files on the publisher are identical to the files on the subscriber. If they are not, press **Skip** and proceed to step 7.*

3. Select a publishing system.
4. Select the UniVerse account containing the files to which you want to subscribe.
5. (Optional) Select Subscribe Dictionaries to subscribe all selected data files and their file dictionaries.
6. (Optional) Select Enable Hot Standby if your subscriber is to be a hot standby system.

7. Click OK. The **Account Subscribing** dialog box appears, as shown in the following example:

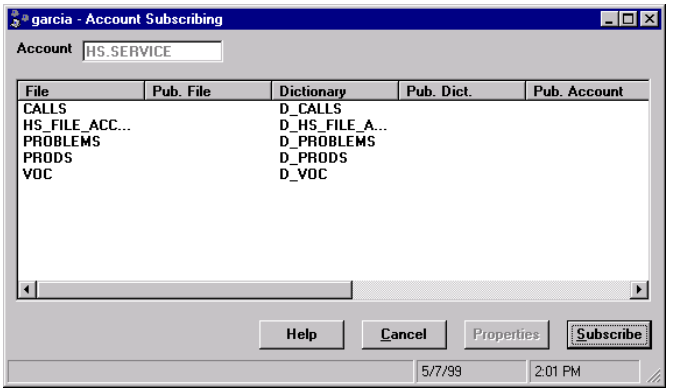

8. You can now do any of the following:

- Click Subscribe to subscribe all selected files.
- Double-click a file in the list, or select a file and click Properties to display the **Subscribe** dialog box. This lets you customize subscription of the selected file. You can modify the following:
- The file description
- Whether or not to subscribe the file dictionary
- The published file this file subscribes to
- Follow the procedure described in "Subscribing One UniVerse File" on page 25.
- Click Cancel to exit the **Account Subscribing** dialog box without subscribing any files.

Viewing and Modifying the Subscribing Files

Complete the following steps to view a list of currently subscribing files:

1. From the left pane, double-click Subscribing or Accounts, then select one of the accounts listed. A list of subscribing UniVerse files in the account appears in the right pane. The **Subscribing Dict** column lists any subscribing file dictionaries. The **Pub. System** column lists the name of the system where the published file resides.

   File type is one of the following:

Type	Description
File	Hashed file (types 2 through 18), B-tree file (type 25), dynamic file (type 30).
Q Pointer	File referenced by a Q-pointer in the VOC file.

Subscribing UniVerse File Types

2. From the right pane, double-click a file to view information about the file. The **Subscribe** dialog box appears. This dialog box displays information about the subscribing file.
3. (Optional) Select the Publish check box, located next to the Dictname box, if you want to subscribe the file dictionary as well as the data file.
4. (Optional) Enter or modify the description of the file in the Desc field.
5. Click OK to save your changes and exit the **Subscribing** dialog box. Click Cancel to exit the **Subscribing** dialog box without saving changes.

Unsubscribing Files

Complete the following steps to unsubscribe files.

1. Execute one of the following steps:

- To select individual files:
  - From the left pane, double-click Subscribing or Accounts.
  - Select one of the accounts listed. A list of subscribing UniVerse files in that account appears in the right pane.

    To select a range of files, select the first file in the range, then hold down the Shift key and select the last file in the range.

    To select random files, hold down the Ctrl key and select the files you want to unsubscribe.

- To select all files in an account:
  - From the left pane, select Subscribing or Accounts. A list of accounts appears in the right pane.
  - From the right pane, select the account whose files you want to unsubscribe.

2. Execute one of the following steps:
   - Choose Unsubscribe from the **Subscribing** menu.
   - Click Unpublish/Unsubscribe on the toolbar.
   - From the right pane, right-click the selected files or the selected account, then choose Unsubscribe from the menu.
3. When asked if you are sure you want to unsubscribe the selected files, click Yes or No. If you click Yes, the files are unsubscribed and the letter "S" is removed from the file icon.

Subscribing Information Files

Information files are useful for debugging. The subscribing system has two information files:

- Subscribing information file, which logs information about the subscribing system.
- Publisher/Subscriber information file, which logs information about connections between subscribing and publishing systems.

Normally you do not need to log information to these information files. We recommend you log information to these files only on the advice of an IBM support specialist.

To activate logging to the subscribing information files, set the verbosity level to a number between 1 and 9. To deactivate logging, set the verbosity level to 0. 0 is the default setting.

Setting the Verbosity Level

Complete the following steps to set the verbosity level:

1. Choose Configure from the **Subscribing** menu. The **Subscribing Config** dialog box appears.

2. Under Info File, set the Verbosity Level to a number between 0 and 9. The higher the verbosity level, the more information is logged.
3. Click OK to exit the **Subscribing Config** dialog box.

Viewing Information Files

Complete the following steps to view the contents of the subscribing information file:

1. Choose Configure from the **Subscribing** menu. The **Subscribing Config** dialog box appears.
2. Under Info File, click Examine. The **Subscribing Info File** window appears, listing the contents of the subscribing information file.

Clearing Information Files

Complete the following steps clear the contents of the subscribing information file:

1. Choose Configure from the **Subscribing** menu. The **Subscribing Config** dialog box appears.
2. Under Info File, click Clear. The contents of the subscribing information file are cleared.

Managing Hot Standby Operations

You can configure only one subscriber to be the hot standby system for a publisher. The hot standby system must be a different computer system from the publishing system it backs up.

A hot standby subscriber can take over for a publisher that suffers a fault or failure rendering it inoperable. When the publisher fails, you can turn on fail-over mode on the hot standby subscriber. Fail-over mode converts all replicated files from being read-only files to read/write-enabled files.

***Note**: If you want users to log in to the hot standby subscriber and run their applications as they run them on the publisher, you must use the same file and account names on the hot standby subscriber as the ones the publisher uses.*

Configuring a Hot Standby Subscriber

Complete the following steps to configure a hot standby subscriber:

1. Add the name of the publishing system you want to hot-backup to the subscriber's Publishing Systems list. Select the Hot Standby check box on the **Publishing Systems** dialog box.

- Choose Systems from the **Subscribing** menu. The **Publishing Systems** dialog box appears, as shown in the following example.

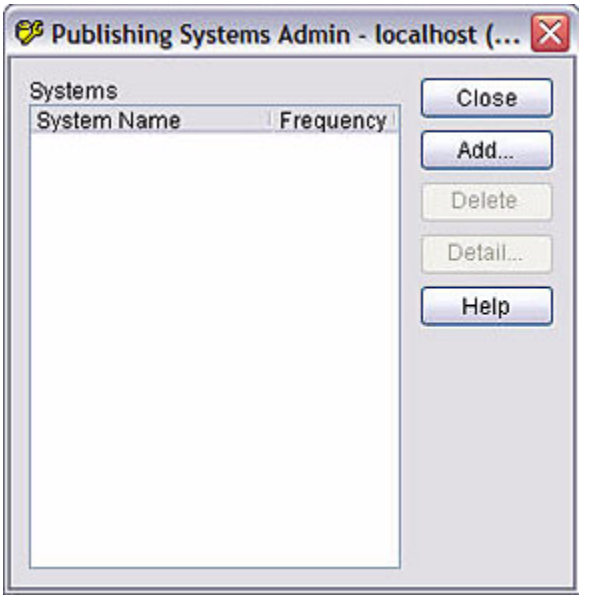

- Click Add. The **Systems Details** dialog box appears, as shown in the following example.

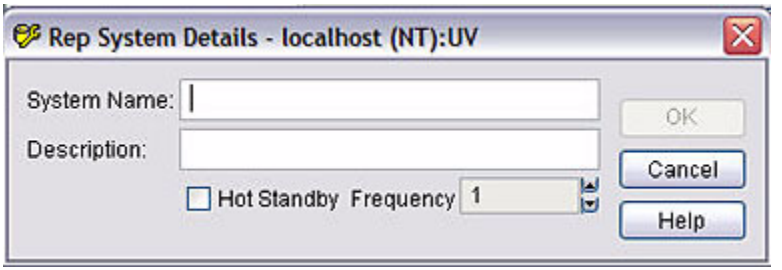

- Enter the name of the publishing system in the System Name field.
- (Optional) Enter a description of the system in the Description field.
- Select Hot Backup.
- Click OK. The publishing system is added to the **Systems** list.

2. You subscribe to one or more published files, selecting the Hot Backup check box on the **Subscribe** dialog box.

Turning On Fail-Over Mode

If your primary publisher fails, you manually enable fail-over mode on the hot backup subscriber. To do this, complete the following steps:

1. Choose Fail Over from the **Subscribing** menu. The **Fail Over** dialog box appears.
2. Choose the name of the publishing system for which to activate fail-over mode.
3. Click OK. The files on the subscriber are now read/write-enabled.

Reconciling the Hot Backup with the Publisher

When the primary publisher is available again, you must reconcile the contents of the subscribing files on the hot backup subscriber with the published files on the original publishing system.

***Note**: When reconciling subscribing files with published files, you need at least the same amount of space as the largest single file you are reconciling. You may need even more space than that, perhaps as much as twice the size of the largest file.*

Here is an outline of the reconciliation process:

1. Stop publishing on the original publisher.
2. Prepare the hot backup subscriber to be reconciled with the publisher.
3. Run the **Sync** command on the original publisher.
4. Reset configurations on the original publisher.

***Note**: You can reconcile only those published files for which a hot backup subscriber system is configured. You cannot reconcile any files published from the failed publishing system that were only replicated but not hot-backed-up.*

Stopping Publishing on the Original Publisher

Complete the following steps to stop publishing on the original publisher.

1. Choose Stop Publishing from the **Publishing** menu.
2. Set the value of the UDRMODE configurable parameter to 0.
3. To unload shared memory, shut down UniVerse.

4. Start UniVerse again.

Preparing the Hot Backup Subscriber to Be Reconciled

Complete the following steps to prepare the hot backup subscriber for reconciliation:

1. On the hot backup subscriber, make sure there are no users logged on to the system.
2. Wait for all file updates to complete.
3. Choose Stop Subscribing from the **Subscribing** menu.

Running the Sync Command on the Original Publisher

To run the Sync command on the original publisher, complete the following steps.

1. On the original publisher, choose Sync from the **Publishing** menu. The **Sync** dialog box appears, listing the hot backup system.
2. Click Sync Up to start the reconciliation process. The original publisher connects to the hot backup subscriber and compares each record of each hot-backed-up file with the originally published version, and implements the updates in the original published files.

When the reconciliation process is complete, the subscribing files on the hot backup system are made read-only, and all subscriptions to the original publisher are restored.

Resetting Configurations on the Original Publisher

To return the original publisher to full publishing status:

1. Set the value of the UDRMODE configurable parameter to 1.
2. To unload shared memory, shut down UniVerse.
3. Start UniVerse again.

Subscribing Information Files

Information files are useful for debugging. The subscribing system has two information files:

- Subscribing information file – This file logs information about the subscribing system.
- Publisher/Subscriber information file – This file logs information about connections between subscribing and publishing systems.

Normally you do not need to log information to these information files. We recommend you log information only on the advice of an IBM support specialist.

To activate logging to the subscribing information files, set the verbosity level to a number between 1 and 9. To deactivate logging, set the verbosity level to 0. 0 is the default setting.

Some Restrictions

You cannot use the following UniVerse commands or statements on published or subscribing files:

- ALTER.TABLE
- CLEAR.FILE
- CNAME
- DELETE.FILE
- DROP.TABLE
- RESIZE

You cannot use triggers with published or subscribing files.

What to Do When Disk Space Fills Up

If replication log files on a publishing system are not regularly cleared, they can fill up all available disk space. When this happens, the state of replication on the publisher changes from Enabled to Full, and all published files become inaccessible.

One reason your log files might fill up is that one or more subscribing systems are not receiving replicated data for some reason. The publisher's log files are cleared only when all subscribing systems have received all replicated data. To resume replication activity in such cases, you need to fix the problem on the subscribing system so that it is receiving replicated data again. This clears the publisher's log files, thus freeing up disk space. When there is enough space to continue, you can resume replication.

To resume replication activity, complete the following steps:

1. Do one or both of the following:
   - Remove unneeded files from the publishing system's disk to free up disk space.
   - Fix the problem on the subscribing systems so they can continue to receive replicated data.
2. On the publisher, choose Resume Publishing from the Publishing menu.
3. On the subscriber, choose Start Subscribing from the Subscribing menu.

Removing Obsolete Replication Log Files

When you start up UniVerse on a machine where replication is enabled, normally all existing replication log files are removed. Replication log files are automatically removed only under the following conditions:

- The replication log file is not being used
- All replication log file entries have been sent to all subscribers
- The replication log file is not marked as “crashed”

In some cases, some replication log files may remain on the system. For example, a subscriber may not receive all its updates because it is off-line.

You can remove obsolete replication log files manually.

What to Do When Replication Fails

When either a publishing or a subscribing system fails for any reason, you need to ensure that the replicated data on all subscribing systems is identical with the published data on the publishing system before you restart replication. In the event of any system crash, do the following:

1. Fix the problem on the systems that crashed.
2. Restart UniVerse.
3. Ensure that published data is identical with replicated data on all subscribers.
4. (Optional) Do one or both of the following:

- On a publishing system, restart the publishing system if it is not configured to start automatically. Choose Start Publishing from the **Publishing** menu.
- On a subscribing system, restart the subscribing system if it is not configured to start automatically. Choose Start Subscribing from the Subscribing menu.

***Note**: If you are fixing a problem on a subscriber only, you need not stop the publishing system on the publisher. Once the subscriber resumes operation, all of the publisher's replicated data is sent to it.*

Chapter

15

Restoring UniVerse Files

Restoring Files 15-3
Choosing the Restore Device 15-3
Checking the Backup Details 15-4
The UVRestore Window 15-5
Choosing What to Restore 15-6
Listing an Index of the Backup Image 15-8
Specifying How to Restore Files 15-8
. 15-9

This chapter describes how to restore UniVerse files from a backup through UniAdmin.

Use the **UniVerse Restore Device** dialog box to restore UniVerse files.

Select one of the following methods to access the **UniVerse Restore Device** dialog box:

- From the **UniAdmin** window, double-click **Restore**.
- From the **UniAdmin** menu, select ***Admin,*** then click **Restore**.

The following example illustrates the **UniVerse Restore Device** dialog box:

Restoring Files

Use the **UniVerse Restore Device** dialog box to choose the restore device and to check the backup details.

When you choose the restore device and confirm that the backup details are correct, the **UVRestore** dialog box appears. From this dialog box, you can execute the following tasks:

- Choose the level of reporting
- Choose what to restore
- Edit the restore selection
- Generate an index of the backup
- Specify how to restore the files
- Start the restoration

Choosing the Restore Device

You can restore files from disk or from tape using the **UniVerse Restore Device** dialog box.

Complete the following steps to choose the restore device:

1. Click the appropriate device type:
   - **Restoring from Disk.** If you are restoring from disk, select **Disk Pathname**. Enter the name of the file that contains the backup in the text entry box, or click **Browse** to search the system for the file.
   - **Restoring from Tape**. Select one or more devices from the **Available Devices** list (this contains all the tape devices defined in the &DEVICE& file), then click **Add**. The devices you choose are listed in the **Selected Devices** list. The order of devices in this list determines the order in which they are used during the restoration if you backed up to multiple tapes. You can reorder devices in this list by dragging and dropping them. To remove a device, select is from the **Selected Devices** list, then click **Remove**.

Note: *Options available in the dialog box change dynamically according to your choice to back up to disk or tape.*

2. If you are restoring from a tape device, check that the backup tapes have been mounted on the selected devices.

   Click **Continue**. The **Check Backup Label** window appears, as shown in the following example:

Checking the Backup Details

If you are restoring from tape, the **Check Backup Label** window contains the backup details for the first device in the **Selected Devices** list.

This window displays the following backup details:

- **Backup Date**. The date and time the backup was performed.
- **Reel Number.** The tape number. If you backed up to a single tape or to a disk path, the value displayed is 1. If you backed up to multiple tapes, this field displays the number of the tape.
- **Compression**. The level of compression during the backup. This features in not supported at this release, so this field always displays **false**.
- **Image Type**. The backup type.
- **Block Size**. The block size (in bytes) used for the backup.
- **NLS on**. The state of NLS when the backup took place.
- **Label**. The backup label. If you backed up the data using the **Backup** options of UniAdmin, this is the text entered in the **Backup Label** field in the UVBackup window.

If the backup details are correct, click **Yes**. The UVRestore window appears. If the backup details are incorrect, click **No**. The **Restore Device** window reappears and you can choose an alternative restore device.

The UVRestore Window

The **UVRestore** dialog box appears when you confirm the backup details are correct, as shown in the following example:

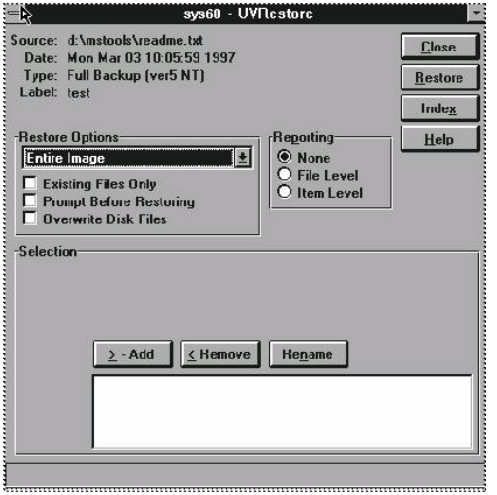

The **UVRestore** window has four main buttons:

- **Close.** Exits the **Restore** option.
- **Restore**. Starts the restoration.
- **Index**. Displays the index of the backup image.
- **Help**. Invokes the Help system.

Choosing What to Restore

Select what to restore from the **Restore Options** list on the **UVRestore** dialog box. When you select an option, a list of accounts or files to restore appears. To add the name of an account or file, select it, then click Add. To remove the name of an account or file from the restore selection, select it, then click Remove.

- Entire Image. This option restores the entire backup image.
- All Accounts in the UV.ACCOUNT File. This option restores all the accounts listed in the UV.ACCOUNT file. The accounts (except the UV account) are automatically added to the restore selection.
- Selected UniVerse Accounts. This option restores accounts selected from those listed in the UV.ACCOUNT file.
- Files in a UniVerse Account. This option restores selected files from a UniVerse account. If you want to restore the file dictionary (not the data file), select the **Dictionary** check box before you click Add.
- Records in a UniVerse File. This option restores selected records in a UniVerse file. Enter the record name in the Record field, then click Add. If you want to add a record from a file dictionary, select the Dictionary check box.
- Selected Directory. This option restores selected UNIX or Windows NT directories. To add a directory to the restore selection, enter a directory path in the Directory field, or click Browse to search the system for a suitable directory. Click Add.

***Note**: The items in the Selection list use paths on the server file system. These paths must match the paths on the tape.*

Editing the Selection List

After you choose what to restore, the record, file, or account details are listed under **Selection** at the bottom of the window (except the Entire Image option). You can add or remove items from this list using Add and Remove. You can edit the list using the Rename button.

To edit an entry, select it from the list and click Rename. The **Edit Restore Selection** dialog box appears, as shown in the following example.

Use this dialog box to define a different item to restore from the backup, to specify an alternative name or destination for the restored item, or both. This dialog box has two fields:

- Search For. Contains the record, file, or account details of the item chosen from the **Selection** list.
- Restore As. This field is empty when the dialog box appears.

To define a different item to restore, edit the Search Fo**r** field. To specify an alternative name or destination for the item to be restored, edit the Restore As field.

Click OK to update the entry in the Selection list.

If you edited the item in the Search Fo**r** field, the new item replaces the original one in the Selection list.

If you chose an alternative name or destination, the Selection list entry is updated to include an equal sign between the original and new details.

***Note**: The Selection list items use paths on the server file system. You must be certain that these paths match the paths on the tape.*

Listing an Index of the Backup Image

To display an index of the backup image, click Index on the **UVRestore** dialog box. The backup index appears in the **UniVerse Command Output** window.

The backup information displayed in the index depends on the level of reporting you choose:

- None. Displays the backup details.
- File Level. Displays the paths of the files in the backup.
- Item Level. Displays the paths of the files, and for UniVerse hashed files, the record IDs.

Specifying How to Restore Files

You can control how the files are restored by selecting any of these check boxes on the **UVRestore** dialog box:

- Existing Files Only. This option restores only files that already exist in the target account. If you don't select this option, all files on the backup are restored, and any files that do not exist in the target account are created.

- Prompt Before Restoring. If you choose Item Level, this option prompts you to restore each file in turn. To restore a file, enter **Y**. The file is restored, overwriting the existing file on disk. To skip a file, enter **N**. You are then prompted for the next file in the **Selection** list. If you don't select this option, files are restored to the chosen account without prompting.
- Overwrite Disk Files. This option determines whether files from the backup overwrite existing files in the chosen account, regardless of the last modification date. If you select this option, the files on the backup overwrite the files on the disk (if the names match). If you clear this option (the default), existing files are overwritten only if the files on the backup have the same (or a more recent) last modification date.

Under Reporting you can choose how much detail to display on the screen during the restoration:

- None. No reporting is done, but you will be notified at the start and end of the restoration.
- File Level. The paths of the files are displayed on the screen.
- Item Level. The paths of the files are displayed, and for UniVerse hashed files, record IDs are also displayed.

All the reports appear in the **UniVerse Command Output** window when restoration starts.

Chapter

16

Managing Catalog Shared Memory

Setting Up Catalog Shared Memory 16-3
Defining Programs to Run in Shared Memory 16-6
Adding Programs to the SHM.TO.LOAD File 16-7
Removing a Program from the SHM.TO.LOAD File 16-7
Loading Programs into Catalog Shared Memory 16-8
Using Programs Stored in Catalog Shared Memory 16-10
Modifying Programs in Catalog Shared Memory 16-11
Adding a Program to Shared Memory 16-11
Updating a Program in Shared Memory 16-12
Removing a Program from Shared Memory 16-13
Removing the Catalog Shared Memory Segment 16-14

This chapter describes how to load cataloged UniVerse BASIC programs into shared memory and run them from there. Shared memory reduces the amount of memory needed for multiple users to run the same program at the same time. The program also starts a little faster because it is already in memory.

For example, if 21 users are running the same UniVerse BASIC program at the same time without catalog shared memory, and the program code requires 50 kilobytes of memory, the total amount of memory used by everyone running that program is 21x50, or 1050, kilobytes. On the other hand, if the program is loaded into catalog shared memory, all 21 users can run one copy of the program, which uses only 50 kilobytes of memory. In this example, catalog shared memory saves 1000 kilobytes (1 megabyte) of memory.

Setting Up Catalog Shared Memory

Before users can use programs in catalog shared memory, the administrator must designate which programs are available and load them into shared memory. UniAdmin simplifies this process. You can specify any BASIC programs to be run in shared memory, including those listed in the system catalog space.

Use the **UniVerse Catalog Shared Memory** dialog box to manage catalog shared memory.

Select one of the following methods to access the **UniVerse Catalog Shared Memory** dialog box:

- From the UniAdmin window, double-click **Shared Programs**.
- From the UniAdmin menu, select **Admin,** then click **Shared Programs**.
- From the UniAdmin toolbar, click the **Manage Shared Programs** icon, as shown in the following example.

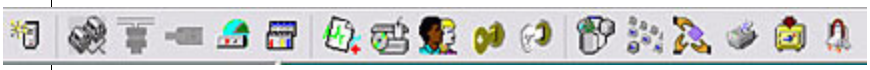

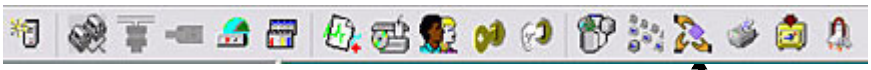

Manage Shared Programs

The following example illustrates the **UniVerse Catalog Shared Memory** dialog box:

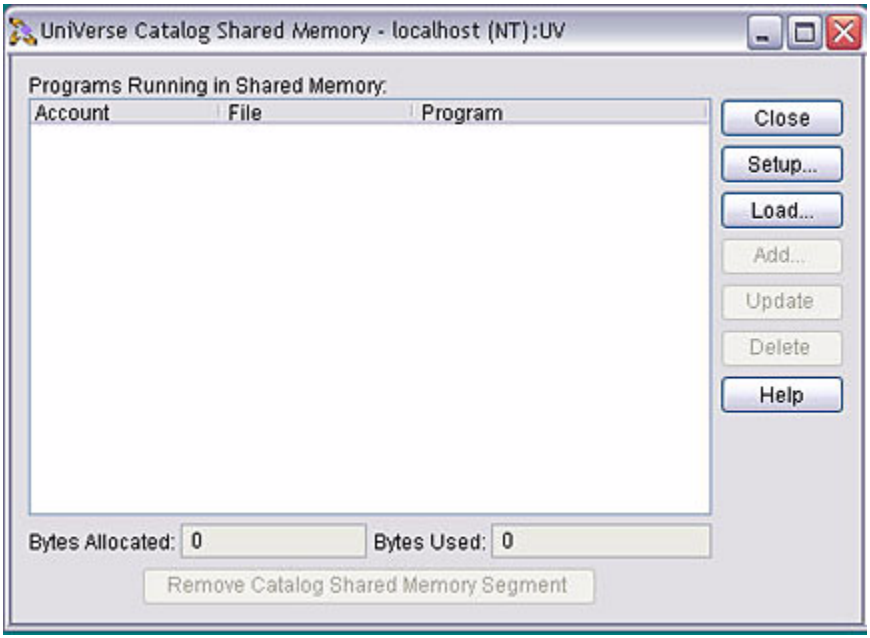

The **UniVerse Catalog Shared Memory** dialog box contains a list of the programs currently running in shared memory. It also displays the amount of shared memory allocated, in bytes, and how much of this memory is currently being used.

Note: *The Catalog Shared Memory window is empty if shared memory is not loaded.*

The tasks you can perform from this window include:

- Defining programs to run in shared memory
- Loading catalog shared memory
- Modifying programs in shared memory
- Removing the catalog shared memory segment

Defining Programs to Run in Shared Memory

The amount of memory available for catalog shared memory is limited. This limit varies from machine to machine. The administrator determines how best to use this space by designating the programs to run from catalog shared memory.

The SHM.TO.LOAD file in the UV account directory contains details, such as account, file, and program names, of the programs selected to run in shared memory. When shared memory is loaded, this file is read to determine which programs are to be loaded. An appropriate amount of shared memory is then loaded.

***Note**: Changes made to the SHM.TO.LOAD file take effect only the next time you start UniVerse or load shared memory.*

To define the programs to be run in shared memory, click **Setup** from the **UniVerse Catalog Shared Memory** dialog box. A dialog box similar to the following example appears:

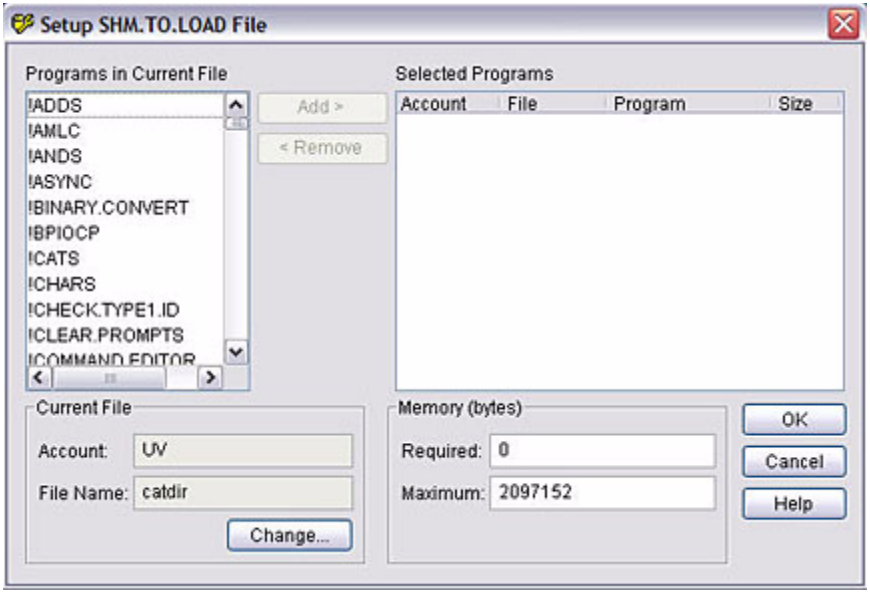

This dialog box enables you to choose programs to add to the SHM.TO.LOAD file. If the SHM.TO.LOAD file does not exist, a warning message window appears. You must acknowledge this message before you can choose any programs. UniVerse creates the SHM.TO.LOAD file when you click **OK**.

Adding Programs to the SHM.TO.LOAD File

Complete the following steps to add a program to the SHM.TO.LOAD file:

1. Choose one or more programs from the **Programs in Current File** list. This list displays the programs in the catalog space.
2. Click **Add**. The updated **Selected programs** list includes the additional programs.
3. If you want to choose a program in a file other than *catdir* (the catalog space), click **Change**. The **Change Current File** dialog box appears, as shown in the following example:

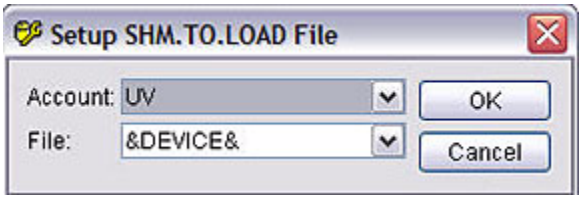

4. Select a new account and file from the **Account** and **File** lists.
5. Click **OK**. The **Setup SHM.TO.LOAD File** dialog box reappears with the updated program list.
6. Select one or more programs from the **Programs in Current File** list.
7. Click **Add**. The updated **Selected Programs** list includes the additional programs.
8. Click **OK** to save the changes and close the **Setup SHM.TO.LOAD File** dialog box.

Removing a Program from the SHM.TO.LOAD File

Complete the following steps to remove a program from the SHM.TO.LOAD file:

1. Click **Setup** from the **UniVerse Catalog Shared Memory** dialog box. The **Setup SHM.TO.LOAD File** dialog box appears.
2. Select one or more programs to remove from the **Selected Programs** list.
3. Click **Remove**. The **Selected Programs** list and the memory details are updated to reflect the deletions.

4. Click **OK** to save the changes and close the **Setup SHM.TO.LOAD File** dialog box.

Loading Programs into Catalog Shared Memory

After designating the programs to be loaded into catalog shared memory, you must create the shared memory segment and load the programs into memory. This is also routinely done each time UniVerse starts.

When you load shared memory, you can do one of the following:

- Load just enough memory to run the programs in the SHM.TO.LOAD file.
- Load additional memory, which is then available if you want to modify the programs you are running.

Complete the following steps to load shared memory:

1. From the **UniVerse Catalog Shared Memory** dialog box, click **Load**. The **Load Catalog Shared Memory** dialog box appears, as shown in the following example:

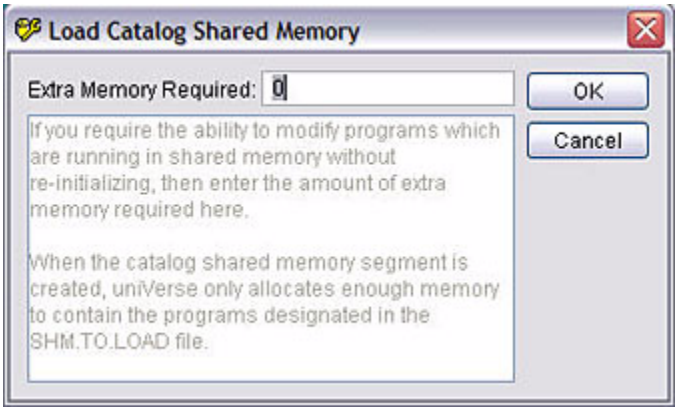

2. If you want to load additional (spare) memory, enter a new value in the **Extra Memory Required** box. This defines the additional amount of shared memory, in bytes, to load.
3. Click **OK**. The existing shared memory is destroyed and the shared memory is reloaded.

The Waiting for Catalog Shared Memory Message

When UniVerse starts on UNIX systems, it tests to see if catalog shared memory is being loaded. UniVerse waits for loading to complete before allowing users to log in to the system. While waiting, UniVerse displays the message “Waiting for Catalog Shared Memory...”

If this message appears and catalog shared memory is **not** being loaded, you should execute the **uv** command with the *-admin -c* options. To execute this command, log in as a UniVerse administrator, then entering the following command from the shell prompt:

```
# /usr/ibm/uv/bin/uv -admin -c
```

This command tells UniVerse that catalog shared memory is not in use, and lets users log in to the UniVerse system.

Using Programs Stored in Catalog Shared Memory

UniVerse invokes programs in the same way, whether they are stored in catalog shared memory or in a disk file. UniVerse attaches to the catalog shared memory only when it starts. Therefore, anyone who starts UniVerse before catalog shared memory is loaded always uses programs from disk. To take advantage of catalog shared memory, those users must leave UniVerse (by entering **Q** at the prompt), and start it again. Even the system administrator must leave UniVerse and restart it to take advantage of the programs in shared memory.

Changing a UniVerse BASIC program and running the BASIC command changes the copy of the object code stored in the system catalog, but does not change the copy stored in catalog shared memory.

Modifying Programs in Catalog Shared Memory

Once programs are installed in catalog shared memory, they are not affected by any changes made to the UniVerse BASIC program on disk. If you modify a program, you must install it again.

You can add, remove, and update programs running in shared memory without having to reload it. However, to do this you must have additional (spare) shared memory loaded. When you modify programs in shared memory, it acts only as a temporary measure, and no changes are made to the SHM.TO.LOAD files.

You can perform the following modification tasks from the **Catalog Shared Memory** dialog box:

- Add a program to shared memory
- Update a program in shared memory
- Remove a program from shared memory

Adding a Program to Shared Memory

You can add any UniVerse BASIC program to shared memory without reloading the program. However, a program is added and run in shared memory only if there is enough spare memory available. The program remains in shared memory until the shared memory is reloaded.

Complete the following steps to add a program to shared memory:

1. From the **Catalog Shared Memory** dialog box, click **Add**. The **Add Program** dialog box appears, as shown in the following example:

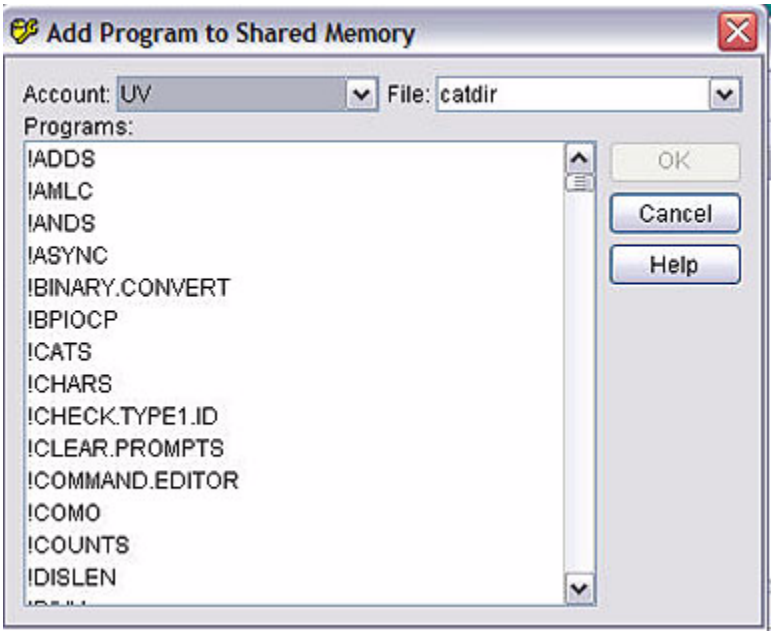

2. Select the appropriate account and file from the **Account** and **File** lists. The default settings for these fields point to *catdir* in the UV account directory (the catalog space).
3. Select the program you want to add from the **Programs** list.
4. Click **OK** to load this program to shared memory. UniVerse checks to see if enough memory is available to load this additional program. If not, a message box appears. You must acknowledge the message. If there is enough space, the program is loaded into shared memory and the **Catalog Shared Memory** dialog box is updated.

Updating a Program in Shared Memory

If a program has been edited and recompiled, you can load the latest version of the program into shared memory. The memory used for the “old” version of the program is not made available, so the updated program will be loaded only if enough spare memory is available.

Complete the following steps to update a program in shared memory:

1. Select the program to update from the **Catalog Shared Memory** dialog box.
2. Click **Update**. A message window appears.
3. Click **Yes**. UniVerse checks that there is enough memory available to load the updated program. If there is not, a message window appears. You must acknowledge the message. If there is enough space, UniVerse loads the program into shared memory and updates the **Catalog Shared Memory** dialog box.

Removing a Program from Shared Memory

You can remove programs from shared memory without reloading. However, the memory used by this program is not made available for reuse.

Complete the following steps to remove a program from shared memory:

1. Select the program to remove from the **Catalog Shared Memory** dialog box.
2. Click **Delete**. A message window appears.
3. Click **Yes**. UniVerse updates the **Catalog Shared Memory** dialog box.

Removing the Catalog Shared Memory Segment

You can remove the catalog shared memory segment without reloading.

Note: *You must be sure you want to remove the catalog shared memory segment, because it can have an impact on other UniVerse users.*

If any users are running programs in shared memory when you delete the catalog shared memory segment, they can continue to run the program from the standard UniVerse catalog.

Complete the following steps to delete the catalog shared memory segment:

1. Click **Remove Catalog Shared Memory Segment** from the **Catalog Shared Memory** dialog box. The following dialog box appears:

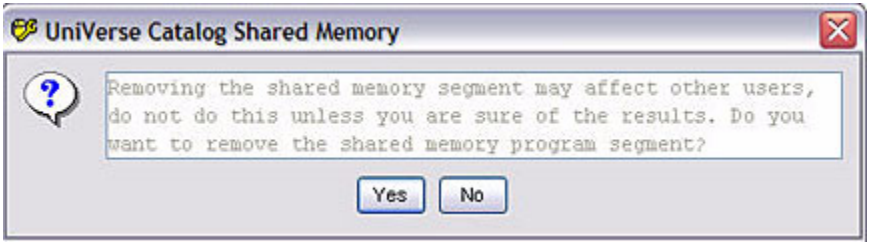

2. Click **Yes** to delete the shared memory segment, or click **No** to exit.

Chapter

17

Configuring SSL Through UniAdmin

Configuring SSL Through UniAdmin 17-2
Accessing UniVerse SSL Configuration Dialog Box. 17-3
Creating a Certificate Request 17-4
Creating a Certificate 17-11
Creating a Security Context 17-19
Configuring SSL for UniObjects for Java or Telnet 17-35

Configuring SSL Through UniAdmin

Secure Sockets Layer (SSL) is a transport layer protocol that provides a secure channel between two communicating programs over which you can send arbitrary application data securely. It is by far the most widely deployed security protocol used on the World Wide Web.

SSL provides server authentication, encryption, and message integrity. It can also support client authentication.

UniData currently supports CallHTTP and the Sockets API. SSL support is important for both of these protocols in order to deploy commercial applications and securely process sensitive data, such as credit card transactions.

This chapter assumes that users who want to use SSL have a basic knowledge of public key cryptography.

Accessing UniVerse SSL Configuration Dialog Box

Use the **UniVerse SSL Configuration** dialog box to administer SSL.

Select one of the following methods to access the **UniVerse SSL Configuration** dialog box:

- From the **UniAdmin** window, double-click **SSL Configure**.
- From the **UniAdmin** menu, select **Admin**, then click **SSL Configure**.

A dialog box similar to the following example appears:

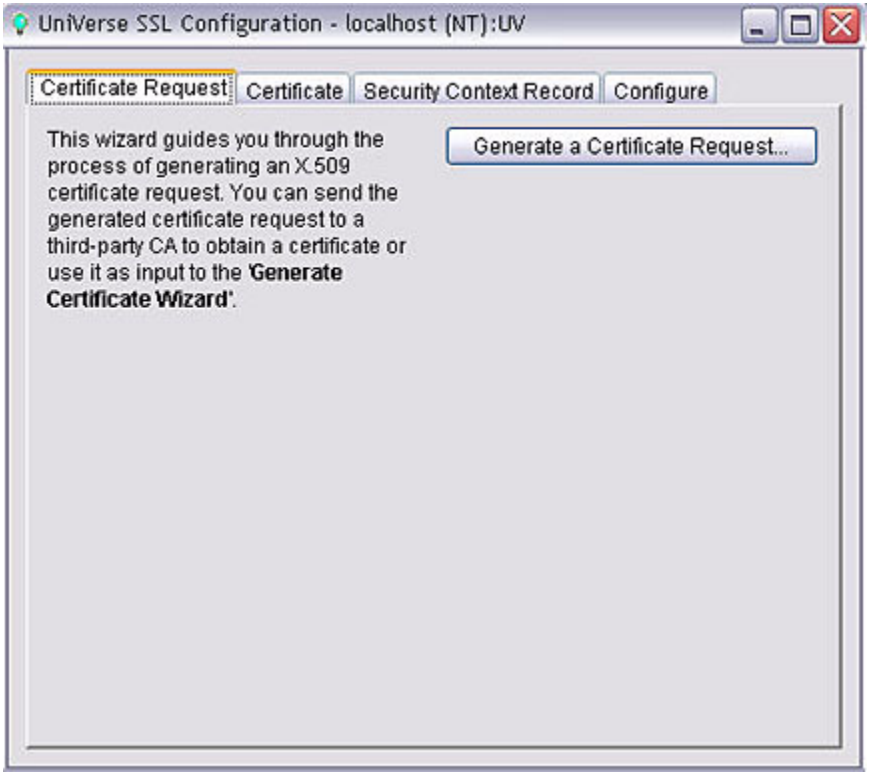

Creating a Certificate Request

Complete the following steps to generate an X.509 certificate request, which you can send to a third-party CA to obtain a certificate, or use as input to the **Generate Certificate** wizard.

1. Click **Generate a Certificate Request**. A dialog box similar to the following example appears:

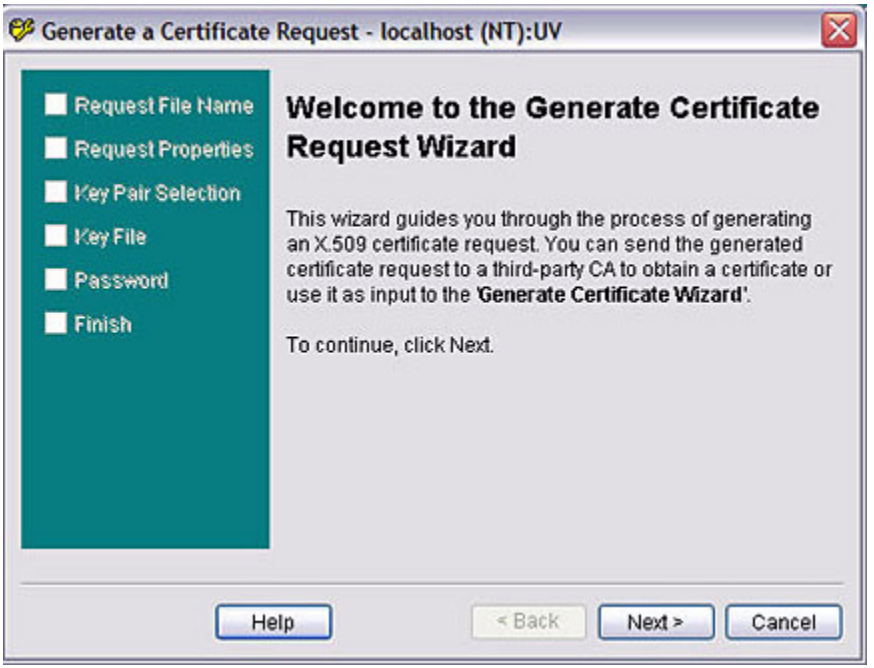

   Click **Next**.

2. The **Certficate Request file** dialog box appears, as shown in the following example:

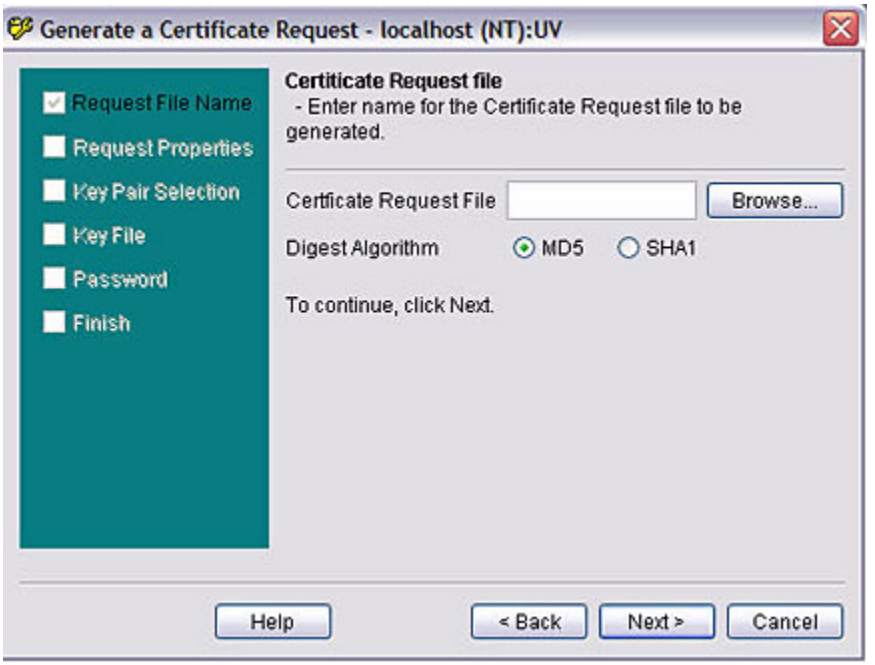

In the **Certficate Request File** box, enter the full path of the operating system-level file to hold the certificate request, or click **Browse** to search for the location.

Click the type of Digest Algorithm for the certificate request. The following types of algorithms are available:

- MD5 – MD5 hash function
- SHA1 – SHA1 hash function

Click **Next**.

3. Next, choose the properties for the certificate from the **Request Properties** dialog box, as shown in the following example:

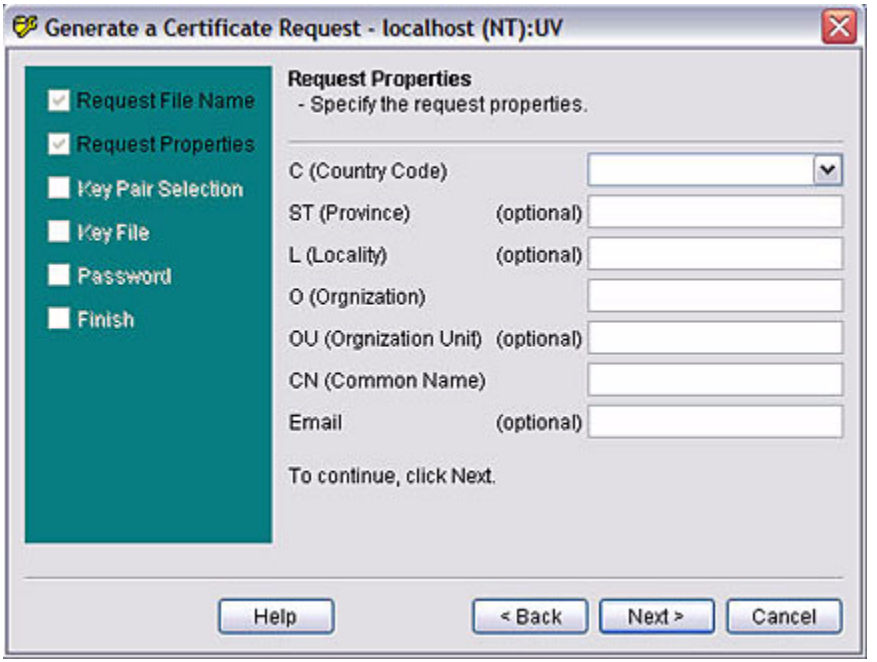

The available properties are:

- C – Country Code
- ST – State or Province
- L – Locality (city)
- O – Organization
- OU – Organization Unit
- CN – Common Name
- Email – Email address

You must define the Country Code, Organization, and Common Name.

Click **Next**.

4. From the **Key Pair Selection** dialog box, select the type of Key Pair.

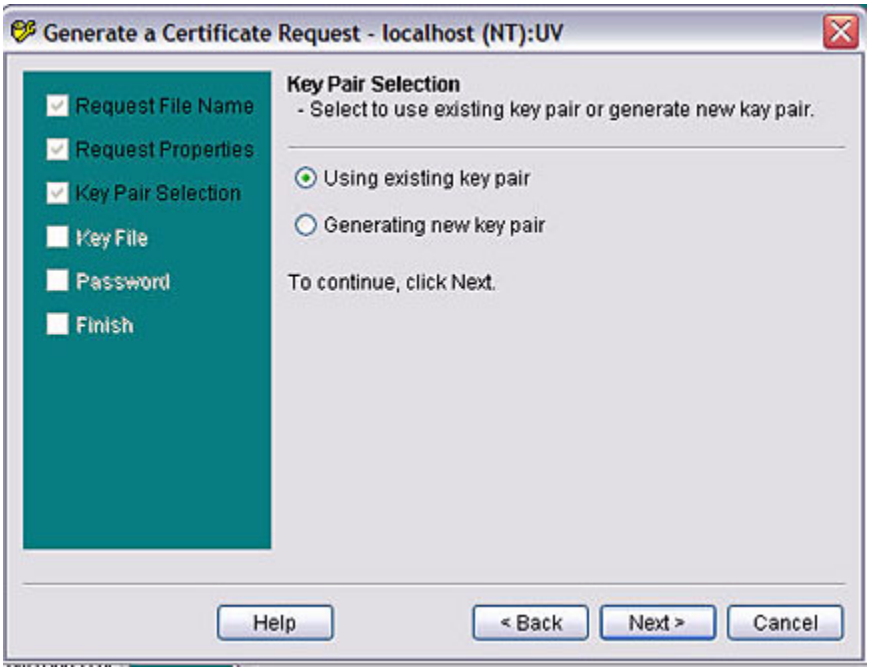

If you are using a previously generated key pair for the certificate request, select **Using existing key pair**. If you are creating a new key pair, select **Generating new key pair**.

Click **Next** to continue.

5. The **Key Pair Info** dialog box appears, as shown in the following example:

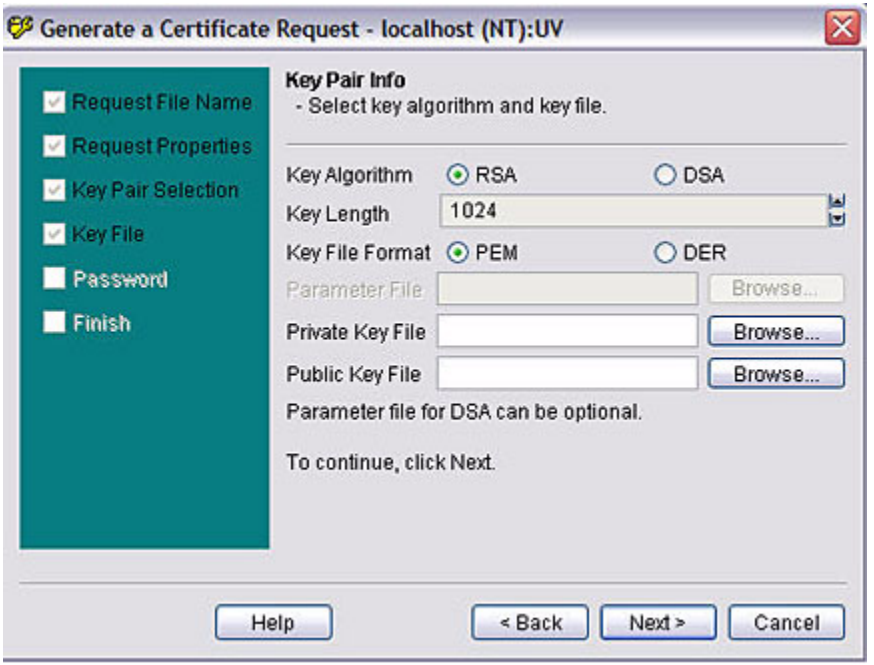

Specify the **Key Algorithm**. Select **RSA** if you want to use an RSA key algorithm, or **DSA** if you want to use a DSA key algorithm.

Select the key length from the **Key Length** list. Valid values range from 512 to 2048.

Select the Key File Format. Select **PEM** for a Privacy Enhanced Mail format, or **DER** for a Distinguished Encoding Rules format.

If you selected a Key File Format of DER, enter the patch to the parameter file in the **Parameter File** box. UniVerse uses this file to generate a new key pair. If you leave this box empty, UniVerse uses the default.

In the **Private Key File** box, enter the name of the file in which you want to store the generated private key, or click **Browse** to search for the existing key if you selected **Use Existing Key Pair**.

In the **Public Key File** box, enter the name of the file in which you want to store the generated public key, or click **Browse** to search for the existing key if you selected **Use Existing Key Pair**.

Click **Next** to continue.

6. Next, define a password for the private key from the **Password** dialog box, as shown in the following example:

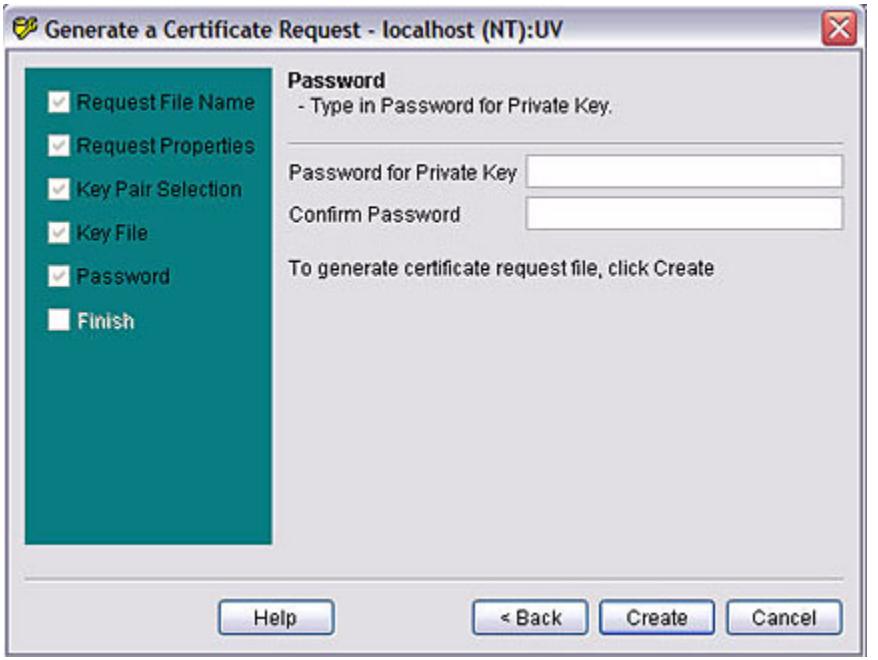

In the **Password for Private Key** box, enter a password for the private key. Confirm the password by reentering it in the **Confirm Password** box.

7. Click **Create** to generate the certificate request file. The following dialog box appears after the certificate request is successfully generated:

Click **Finish** to return to the **UniVerse SSL Configuration** wizard.

Creating a Certificate

You can create two types of certificates from UniAdmin:

- Self-signed certificates as a root CA that can be used later to sign other certificates.
- CA signed certificates.

To create a certificate, from the **UniVerse SSL Configuration** dialog box,click **Certificate**. A dialog box similar to the following example appears:

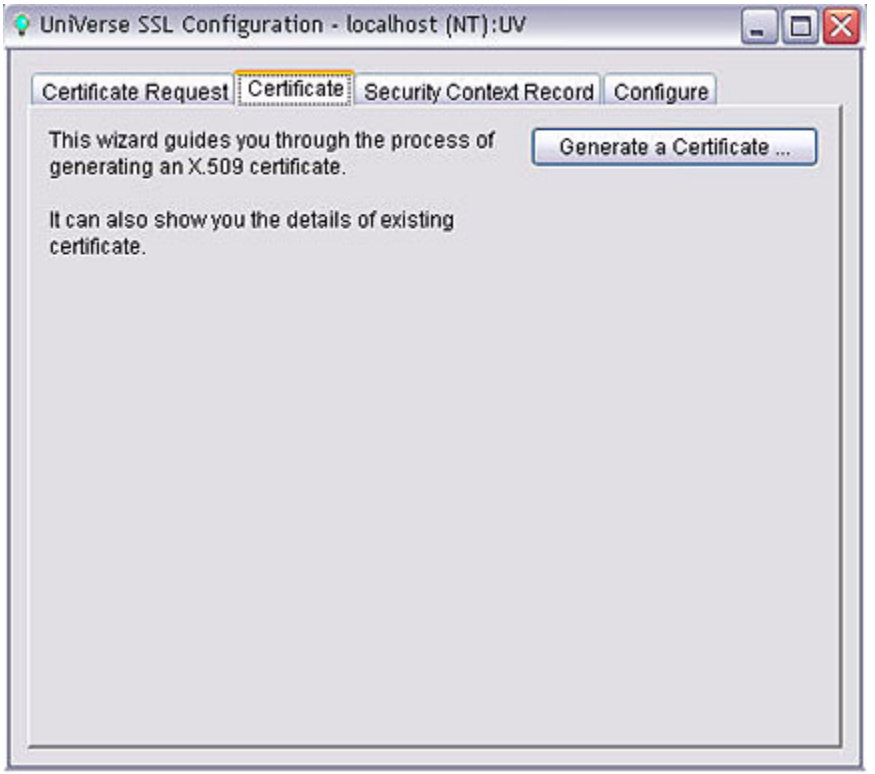

To begin generating an X.509 certificate, or to view details of an existing certificate, click **Generate a Certificate**. A dialog box similar to the following example appears:

Click **Next**, then complete the following steps to create a certificate:

1. The **Certificate File** dialog box appears, as shown in the following example:

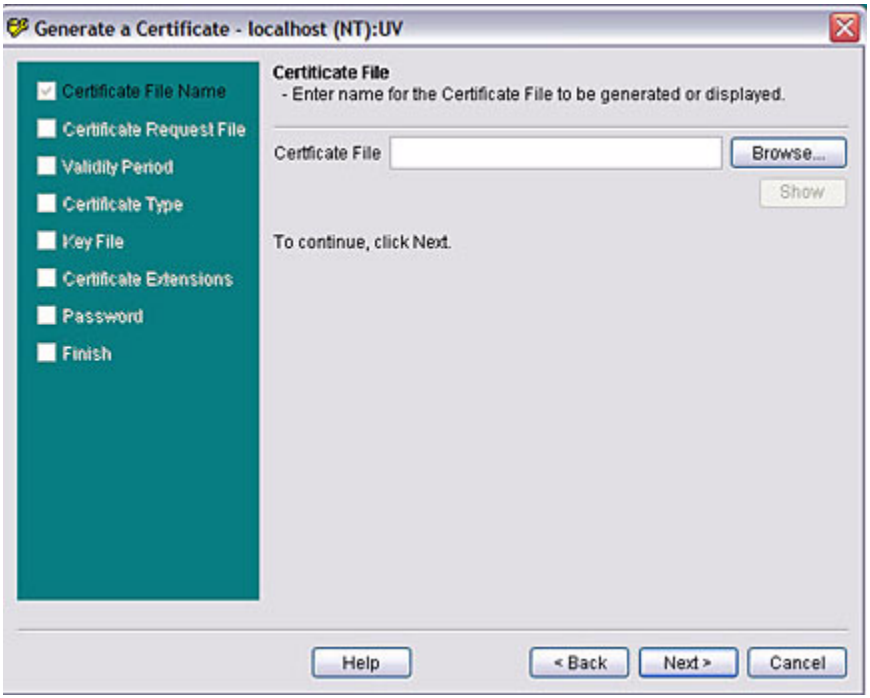

In the **Certificate File** box, enter the certificate file name, or click **Browse** search for the file. Click **Next**.

2. The **Certificate Request File** dialog box appears, as shown in the following example:

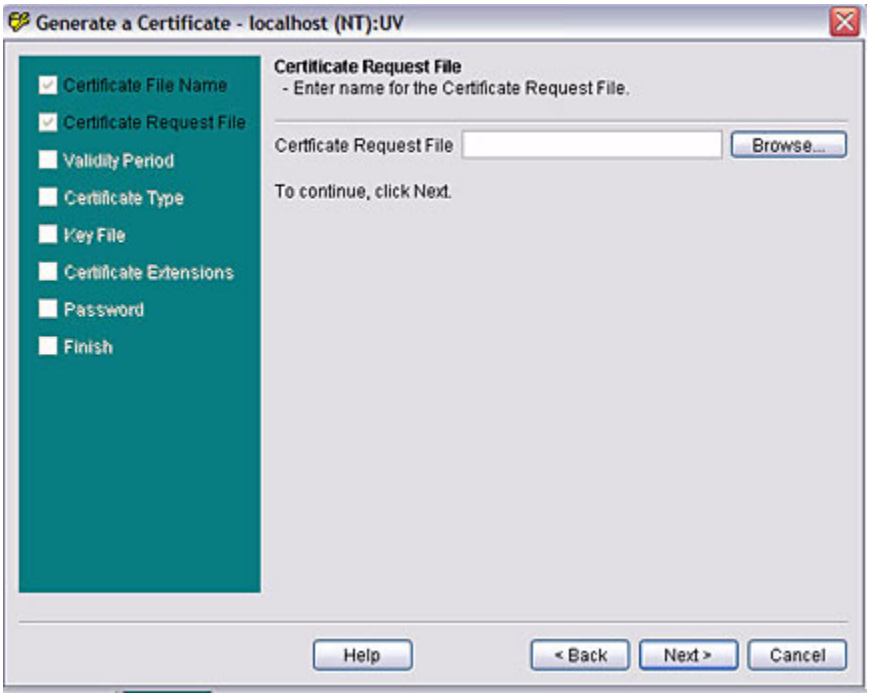

In the **Certficate Request File** box, enter the name of the file to write the generated certificate, or click **Browse** to search for the file. Click **Next**.

3. The **Validity Period** dialog box appears, as shown in the following example:

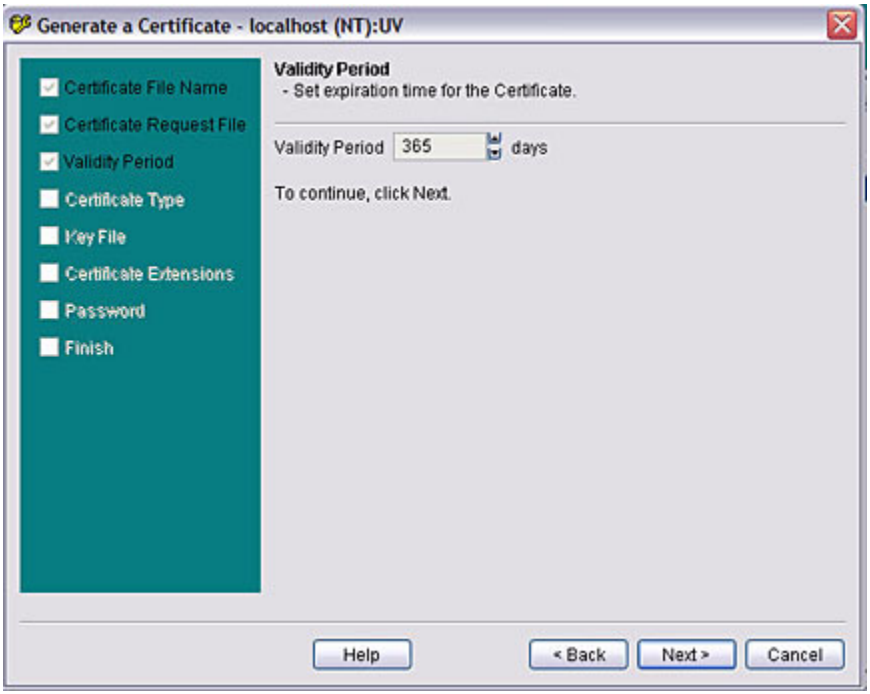

Select the number of days for which certificate is valid from the **Validity Period** list. The certificate is valid starting from the current date until the number of days you specify expires. The default value is 365 days.

Click **Next.**

4. The **Certificate Type** dialog box appears, as shown in the following example:

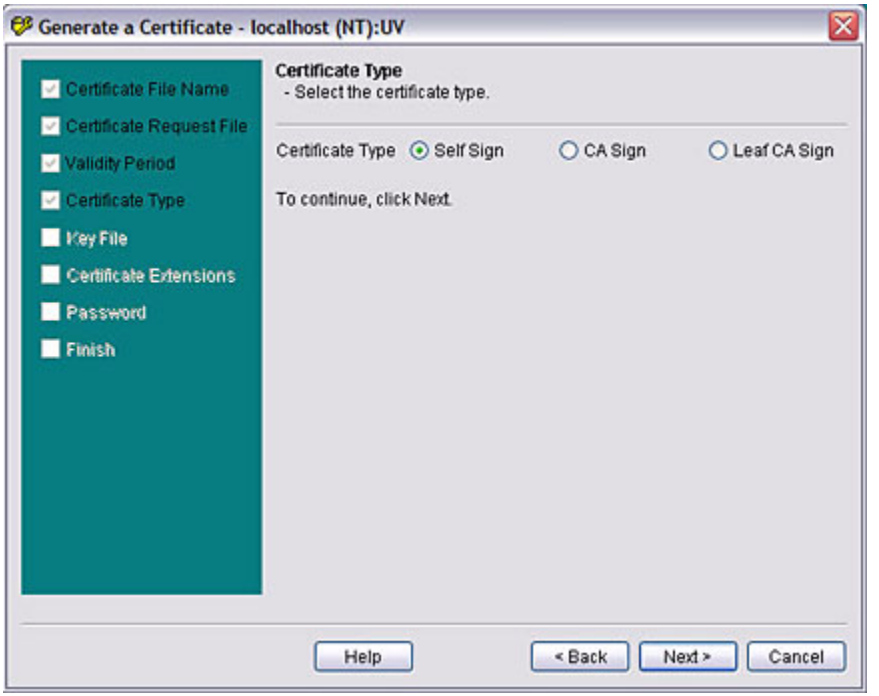

Select the type of certificate. The following types of certificates are available:

- Self Sign
- CA Sign
- Leaf CA Sign

Click **Next** to continue.

5. If the type of certficate is CA or Leaf CA, you must specify a private key file associated with the signing CA certificate. Define the private key file from **Private Key File** dialog box, as shown in the following example:

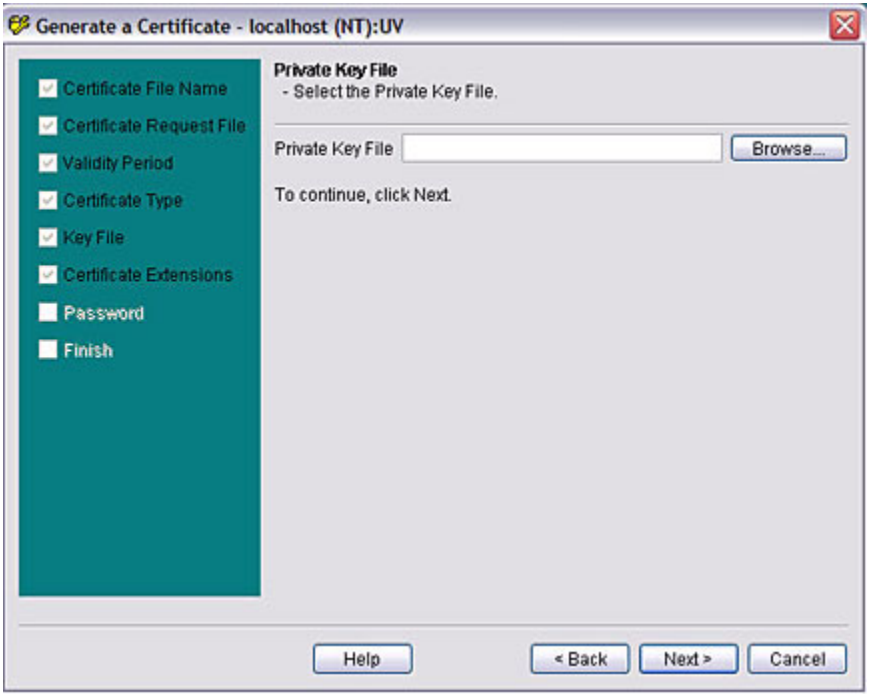

Enter the name of the private key file in the **Private Key File** box, or click **Browse** to search for this file, then click **Next**.

6. Enter the password for the private key file in the **Password** dialog box, as shown in the following example:

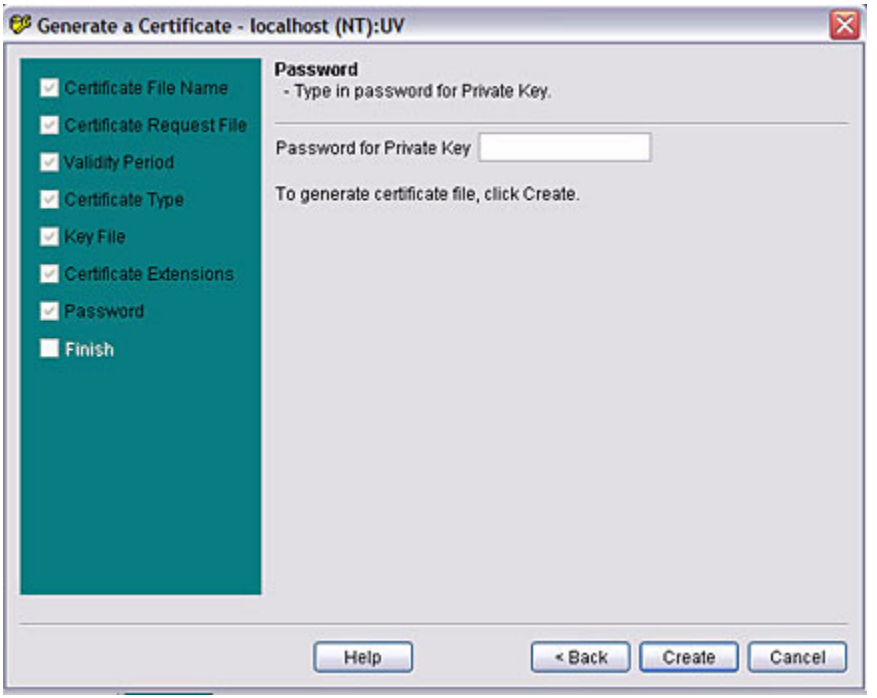

7. Click **Create** to create the certificate.

Creating a Security Context

A security context is a data structure that holds all aspects of security characteristics that the application intends to associate with a secured connection.

To create a security context, from the **UniVerse SSL Configuration** dialog box, click the **Security Context Record** tab. A dialog box similar to the following example appears:

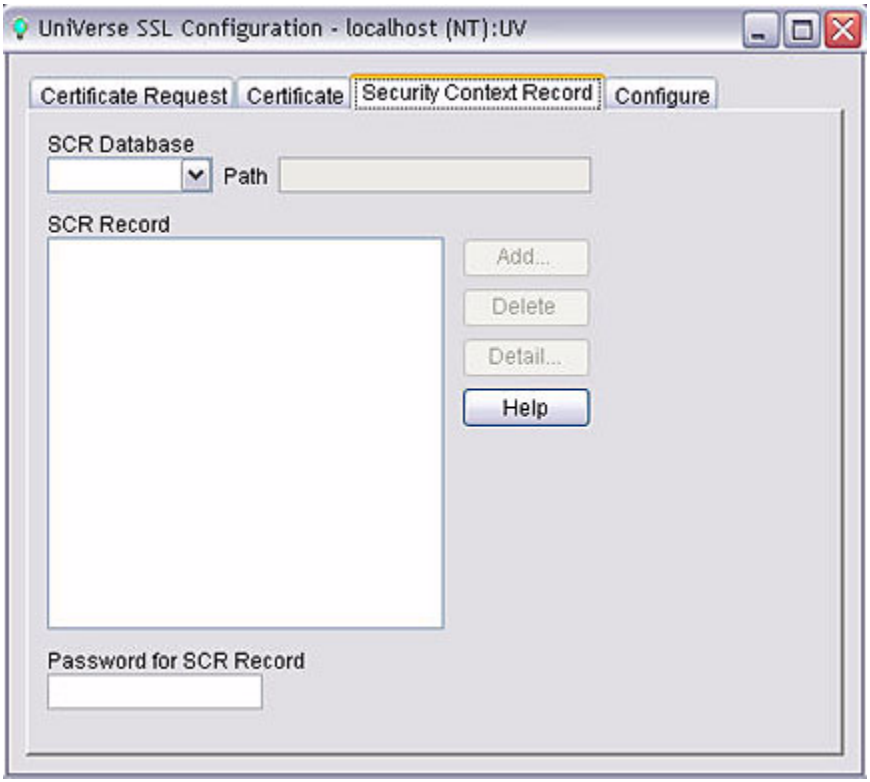

Complete the following steps to create a Security Context Record:

1. Select the account where you want to create or view the security context record (SCR) from the **SCR Database** list. UniAdmin populates the **Path** box with the full path of the database.

2. To add an SCR record, click **Add**. A dialog box similar to the following example appears:

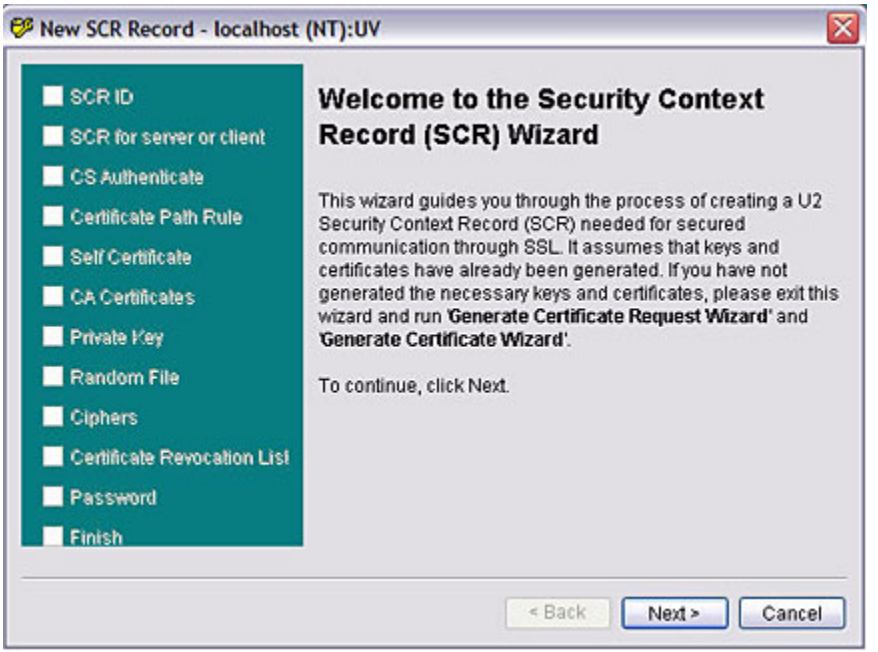

Make sure you have generated the necessary keys and certificates needed before proceeding.

Click **Next** to continue.

3. The **Security Context Record ID** dialog box appears, as shown in the following example:

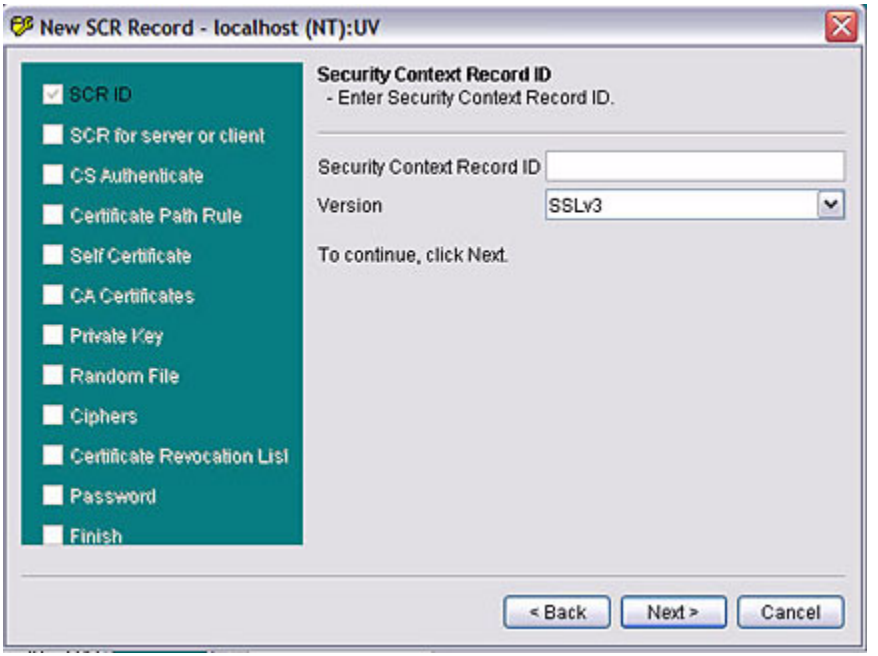

Enter an ID for the SCR in the **Security Context Record ID** box.

Select the appropriate version for the SCR record in the **Version** box. Valid versions are:

- SSLv2
- SSLv3
- TLSv1

Note: *IBM recommends that you only use SSLv3 or TLSv1.*

Click **Next** to continue.

4. Next, define if the SCR for the server or client from the **SCR for Server or Client** dialog box, as shown in the following example:

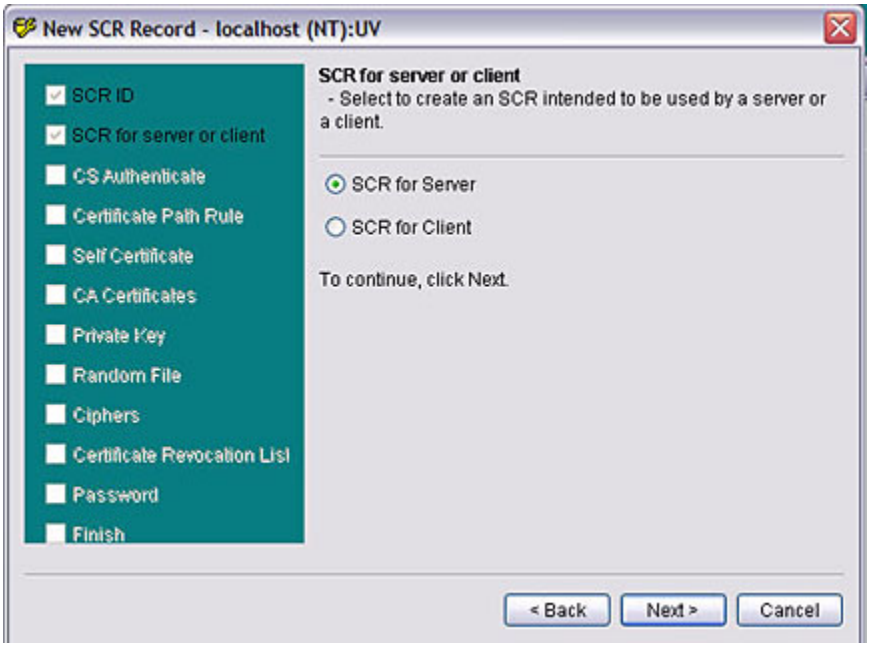

If the SCR is for use by a server, select **SCR for Server**. If the SCR is for use by a client, select **SCR for Client**.

Click **Next** to continue.

5. When you select Client Authentication, during the initial SSL handshake, the server sends the client authentication request to the client. It also receives the client certificate and performs authentication according to the issuer's certificate (or certificate chain) set in the security context.

   Set authentication parameters from the **Server Authenticate** dialog box, as shown in the following example:

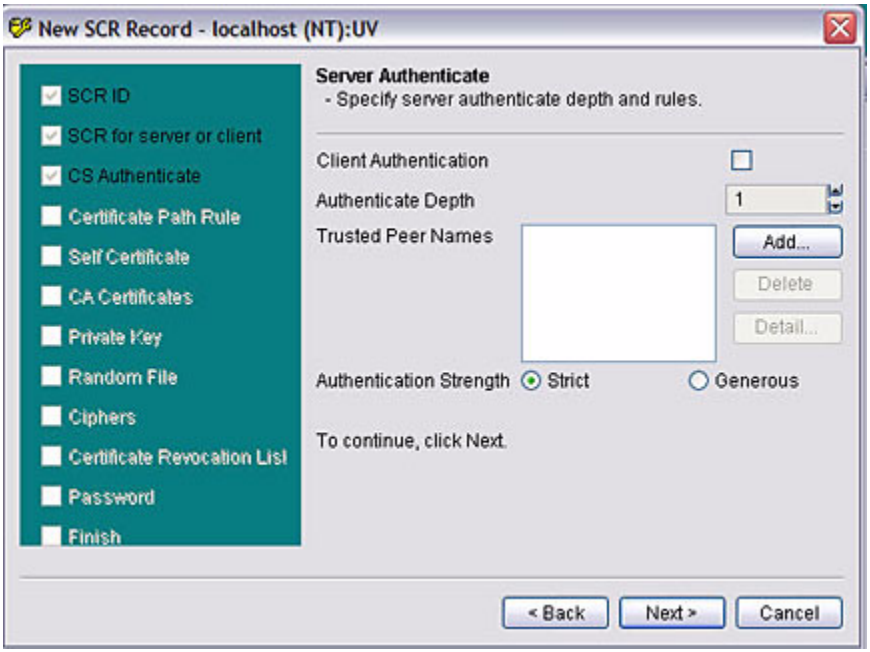

   If you want to set authentication, select the **Client Authentication** check box.

   The **Authentication Depth** value determines how deeply UniVerse verifies before determining that a certificate is not valid.

   Depth is the maximum number of intermediate issue certificates, or CA certificates, UniVerse must examine while verifying an incoming certificate. A depth of 0 indicates that the certificate must be self-signed. A depth of 1 means that the incoming certificate can be either self-signed, or signed by a CA known to the security context.

   You can set Authentication Depth on the server and the client. The default value for both is 1.

   You can add an authentication rule to a security context. UniVerse uses the

rules during SSL negotiation to determine whether or not to trust the peer. UniVerse supports the following rules:

- **Verification Strength** rule – This rule governes the SSL negotiation and determines whether or not an authentication process is considered successful. There are two levels of security, ***generous*** and ***strict***. If you specify ***generous***, the certificate need only contain the subject name (common name) that matches one specified by "PeerName" to be considered valid. If you specify ***strict***, the incoming certificate must pass a number of checks, including signature check, expiry check, purpose check, and issuer check.

Note: *IBM recommends setting the rule to generous only for development or testing purposes.*

- **PeerName** rule – By specifying the **PeerName** rule and attribute mark separated common names in ruleString, trust server/client names will be stored in the context.

  During the SSL handshake negotiations, the server sends its certificate to the client. By specifying trust server names, the client can control with which server or servers it should communicate. During the handshake, once the server certificate has been authenticated by way of the issuer (CA) certificate(s). UniVerse compares the subject name contained in the certificate against the trusted server names set in the context. If the server subject name matches one of the trusted names, communication continues, otherwise UniVerse does not establish the connection.

If no trusted peer name is set, any peer is considered legitimate.

To add an authentication rule, click **Add**. A dialog box similar to the following example appears:

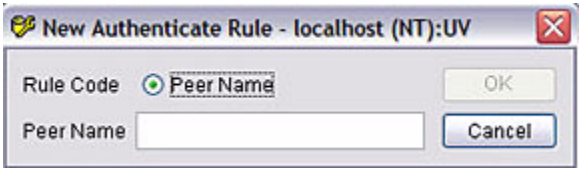

Enter the information for the new rule, then click **OK**.

Select the Authentication Strength, then click **Next**.

6. The **Certificate Path Rule** enables you to specify locations to search for certificates. Define the Certificate Path Rule from the **Certificate Path Rule** dialog box, as shown in the following example:

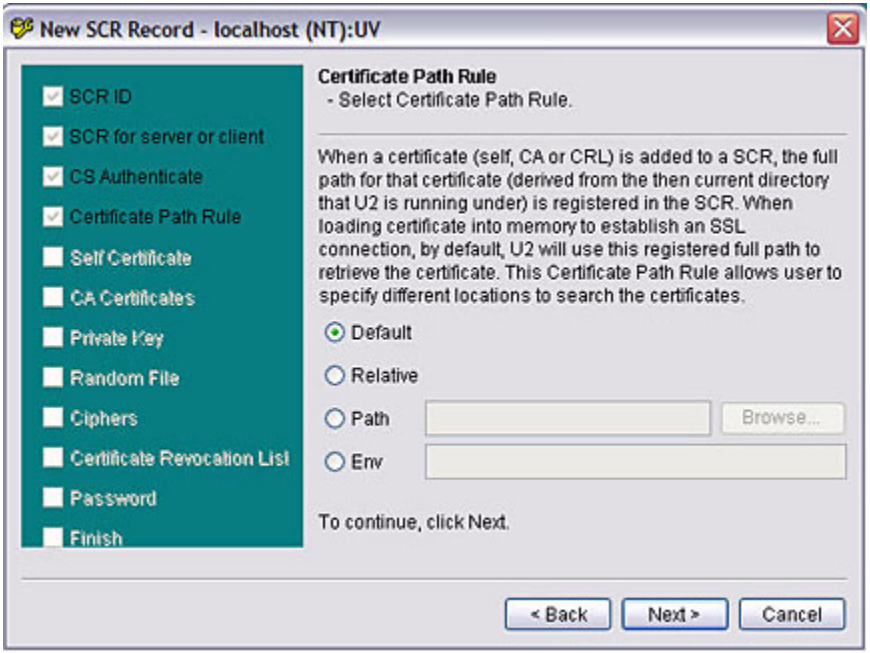

When you add a certificate to an SCR, the full path for that certificate is registered in the SCR. This path is derived from the current directory where UniVerse is running. When the certificate is loaded into memory to establish the SSL connection, UniVerse uses this registered full path to retrieve the certificate. You can change the path in one of the following ways:

- **Relative** – If you select **Relative**, UniVerse looks in the current directory where UniVerse is running for the certificate.

  Note: *Some of the UniVerse processes, such as the telnet server processes, run from the system directory.*

- **Path** – If you select **Path**, UniVerse uses the path you specify for loading the certificate added to this SCR. You can specify either an absolute path or a relative path.

- **Env** – If you select **Env**, enter an environment variable name in the **Env** box. If you specify this option, the UniVerse process first obtains the value of the environment variable you specify and uses that value as the

path to load the certificates.

***Note**: UniVerse only evaluates the environment variable when the first SSL connection is made. The value is cached for later reference.*

7. You can load a certificate, or multiple certificates, into a security context for use as a UniVerse server certificate or client certificate. Alternatively, you can specify a directory that contains the certificates to use as a Certificate Authority (CA) certificate to authenticate incoming certificates or act as a revocation list, checking for expired or revoked certificates.

   The purpose of a certificate is to bind the name of an entity with its public key. It is basically a means of distributing public keys. A certificate always contains the following three pieces of information:

   - Name
   - Public Key
   - Digital signature signed by a trusted third party called a Certificate Authority (CA) with its private key.

   If you have the public key of the CA, you can verify that the certificate is authentic.

   SSL protocol specifies that when two parties start a handshake, the server must always send its certificate to the client for authentication. It may also require the client to send its certificate to the server for authentication. UniVerse servers that act as HTTP clients are not required to maintain a client certificate. UniVerse applications that act as SSL socket servers must install a server certificate. UniObjects for Java servers and telnet servers also require server certificates.

   There can be only one server/client certificate per specific security context. Adding a new certificate automatically replaces an existing certificate. However, for issuer certificates, UniVerse chains a new one with existing certificates so UniVerse applications can perform chained authentication.

   If the issuer certificate is in PEM format, it can contain multiple certificates generated by concatenating certificates together.

   ***Note**: All certificates that form an issuer chain must be of the same type.*

Select the self certificate file from the **Self Certificate** dialog box, as shown in the following example:

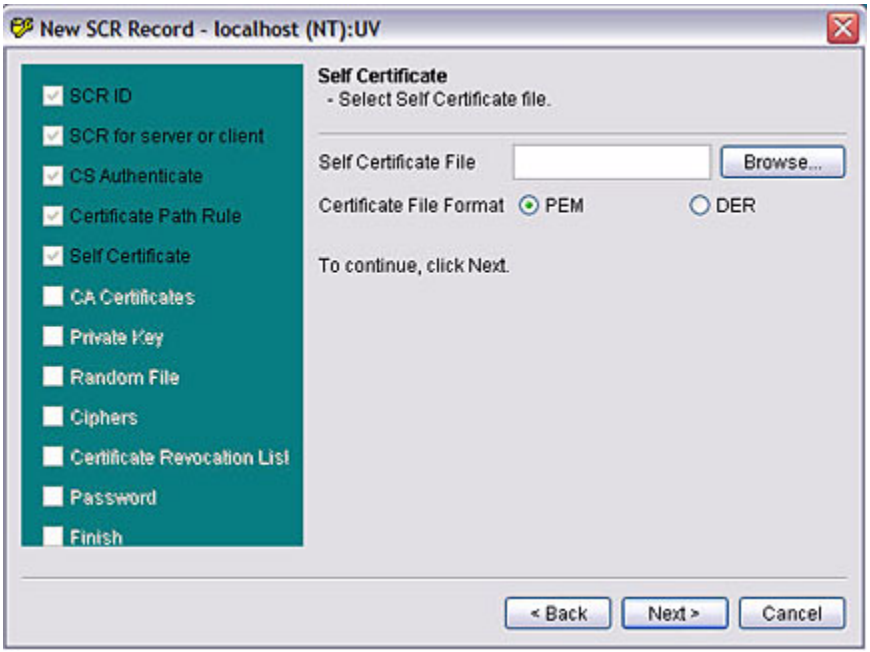

In the **Self Certificate File** box, enter the path to the file containing the self certificate, or click **Browse** to search for the file.

Next, select the format for the certificate. Select **PEM** for Base64 encoded format, or **DER** for ASN.1 binary format.

Click **Next** to continue.

8. If you are defining an SCR record intended to be used by a server, you must install a private key. If you are defining an SRC record intended to be used by a client and you do not install a self-certificate, you do not need to install a private key.

   You can load a private key into a security context so it can be used by SSL functions. Setting a private key replaces an existing private key.

   UniVerse uses a private key to digitally sign a message or encrypt a symmetric secret key to use for data encryption.

   Select the private key associated with self certificate from the **Private Key** dialog box, as shown in the following example:

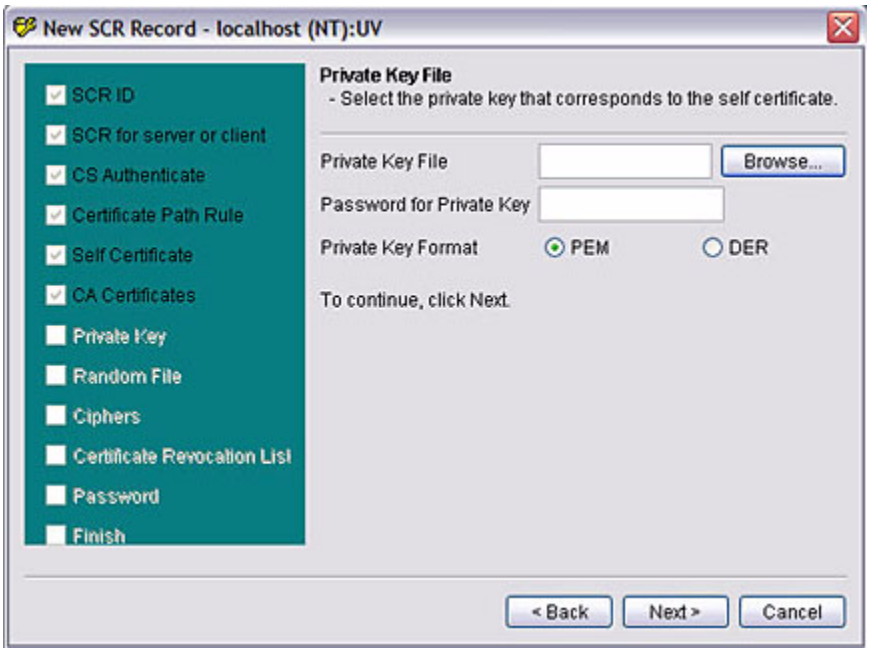

   In the **Private Key File** box, enter the file that contains the private key, or click **Browse** to search for the file.

   In the **Password for Private Key** box, enter the password for the private key.

   Next, select the format for the private key. Click **PEM** for Base64 encoded format, or **DER** for ASN.1 binary format.

   Click **Next** to continue.

9. To define a CA certificate, in the **CA Certificate** dialog box, click **Add**. A dialog box similar to the following example appears:

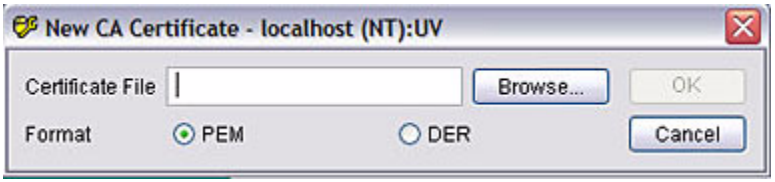

In the **Certificate File** box, enter the path to the file containing the certificate, or click **Browse** to search for the file.

Next, select the format for the certificate. Click **PEM** for Base64 encoded format, or **DER** for ASN.1 binary format.

Click **OK** to add the certificate, or **Cancel** to exit.

Click **Next** to continue.

10. You can generate a random file from a series of source files and set that file as the default random file for the SCR record.

    The strength of cryptographic functions depends on the true randomness of the keys. Be default, UniVerse uses the .rnd file in the current account. You can override the default by adding a random seed file.

    Note: *IBM recommends you use the default .rnd file.*

    To select a random seed file other than the default, click **Browse** to search for the file, as shown in the following example:

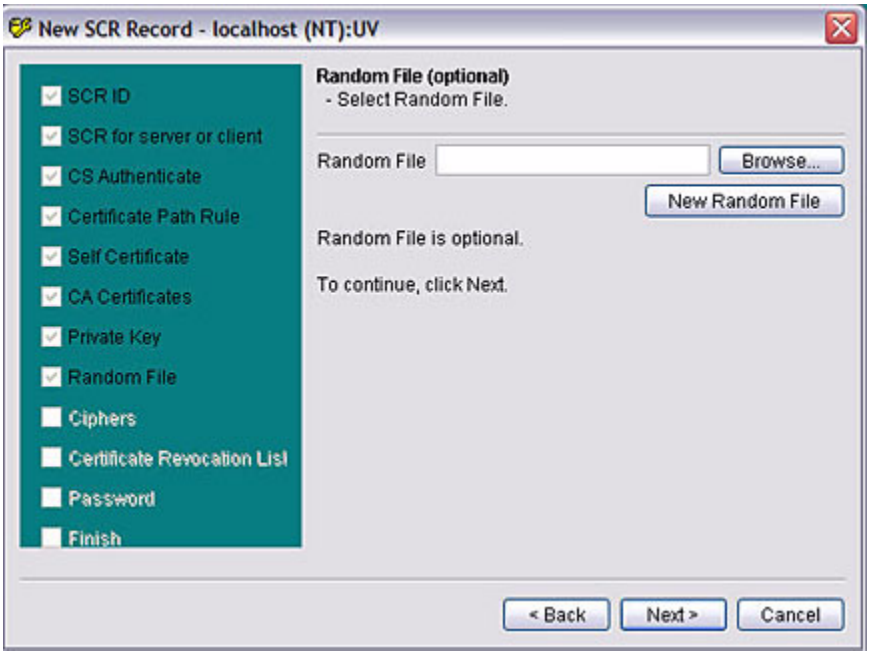

To create a new random seed file, click **New Random File**. A dialog box similar to the following example appears:

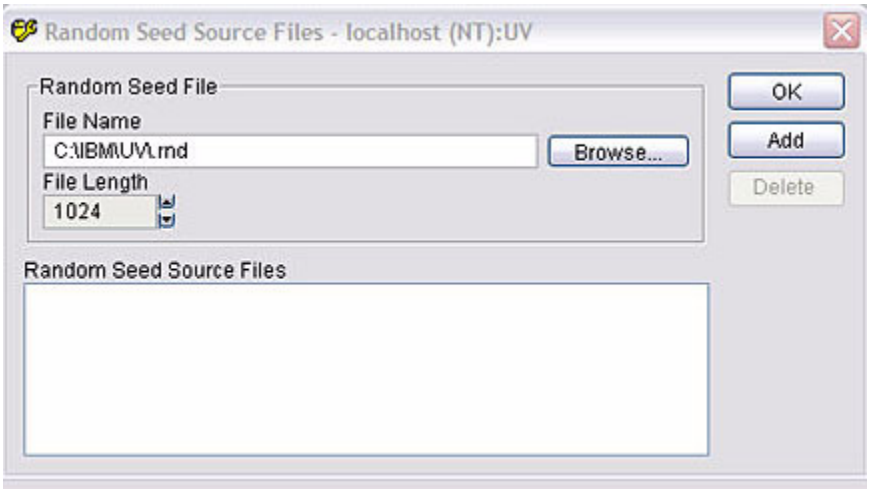

Enter the path to random file, or click **Browse** to select the random source file.

In the **File Length** list, choose the file length.

In the **Random Seed Source Files** box, enter the random seed source files.

Click **OK**.

11. You can define ciphers from the **Ciphers** dialog box, as shown in the following example:

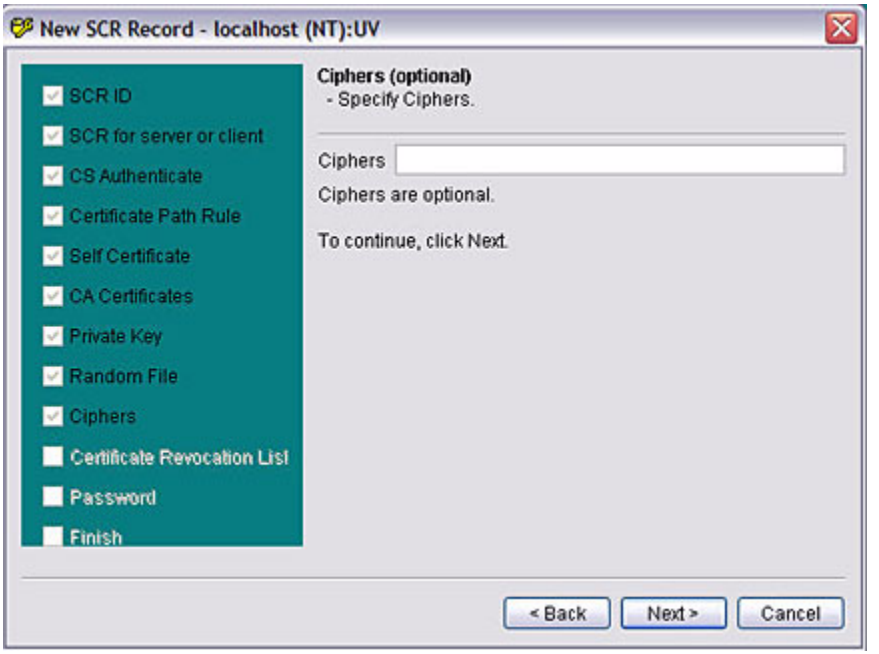

Ciphers enable you to identify which cipher suites should be supported for the specified context. It affects the cipher suites and public key algorithms supported during the SSL/TLS handshake and subsequent data exchanges.

When a context is created, its cipher suites will be set to SSLv3 suites supported by the SSL version you selected.

In the **Ciphers** box, enter the Cipher Suite for this SCR record.

The *CipherSpecs* parameter is a string containing *cipher-spec* separated by colons. An SSL cipher specification in *cipher-spec* is composed of 4 major attributes as well as several, less significant attributes. For detailed information about Cipher Suites, see "*UniVerse BASIC Extensions.*"

12. From the **Certificate Revocation List** dialog box, click **Add** to add a CRL file as part of a revocation list, as shown in the following example:

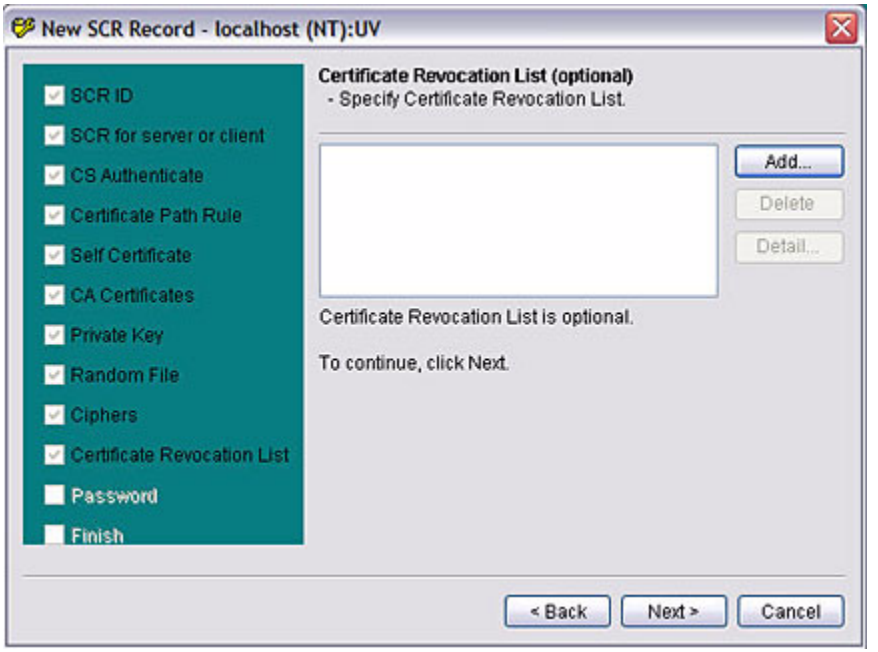

A dialog box similar to the following example appears:

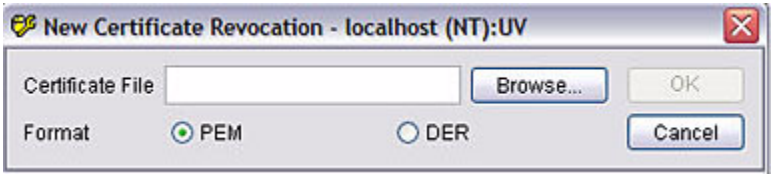

In the **Certificate File** box, enter the path to the file containing the CRL, or click **Browse** to search for the file.

Next, select the format for the CRL file. Click **PEM** for Base64 encoded format, or **DER** for ASN.1 binary format.,

Click **OK** to create the SCR record, or click **Cancel** to exit.

Click **Next** to continue.

13. A dialog box similar to the following example appears:

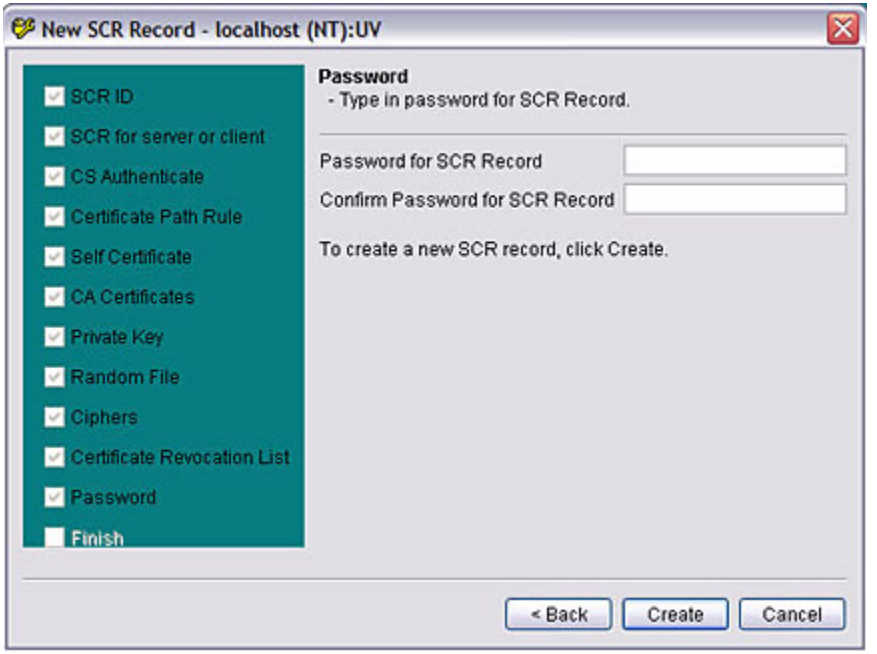

In the **Password for SCR Record** box, enter a password to access the record. Reenter the password in the **Confirm Password for SCR Record** box.

Click **Create** to create the SCR record, or click **Cancel** to exit.

Configuring SSL for UniObjects for Java or Telnet

After you create an SCR record, you need to configure SSL for UniObjects for Java or SSL for Telnet.

An SCR record contains all SSL-related properties necessary for the server to establish a secured connection with an SSL client. The properties include the server's private key certificate, client authentication flag and strength, and trusted entities. For more information, see "*UniVerse Security Features*."

From the **UniVerse SSL Configuration** dialog box, select the **Configure** tab. A dialog box similar to the following example appears:

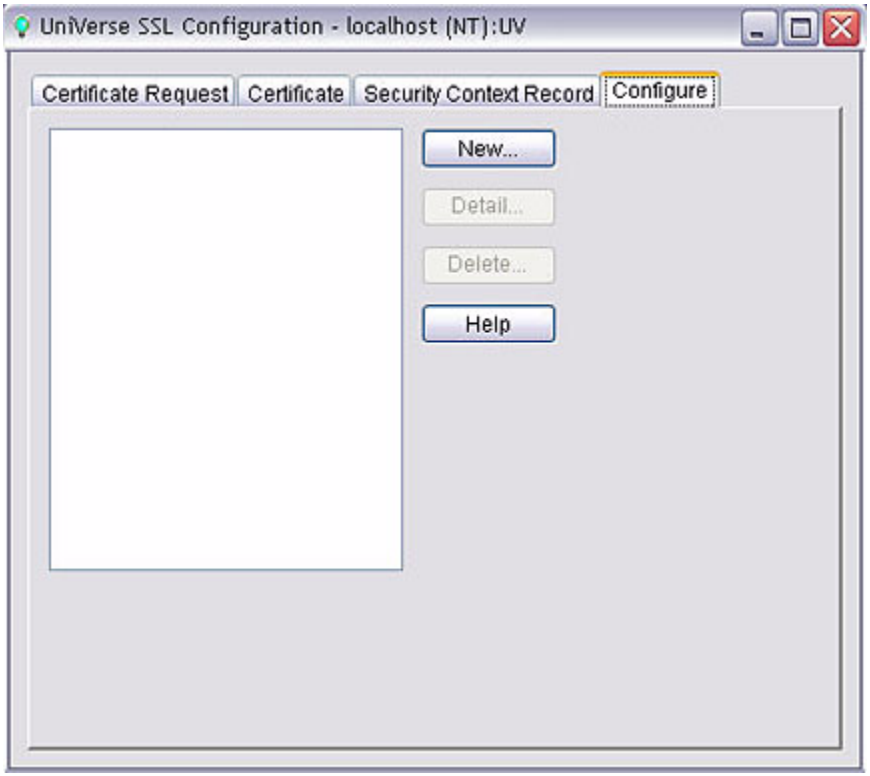

Complete the following steps to configure SSL:

1. Click **New**. A dialog box similar to the following example appears:

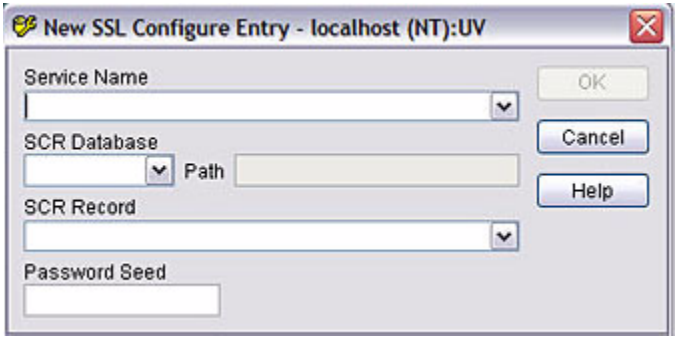

From the **Service Name** list, select a service name.

From the **SCR Database** list, select the database for this configuration entry. UniAdmin automatically populates the **Path** box.

Select the SCR record from the **SCR Record** list, then enter the password for this record in the **Password Seed** box.

Chapter

18

Transaction Logging

Configuring the Transaction Logging System 18-4
Setting Up or Changing Transaction Logging Configuration 18-4
Specifying Recoverable Files 18-7
Activating Files 18-7
Deactivating Files 18-8
Managing Transaction Logging 18-10
Managing Log Files 18-10
Managing Transaction Logging 18-14
Viewing and Deleting Information Files 18-16
Restoring Log Files 18-17
Recovering Files 18-20
Recovering Log Files from Disk 18-20
Recovering Files from a Full Log File on Tape 18-23
Deleting Restored Log Files 18-27

Transaction logging is a way of securing a copy of your UniVerse file updates in a special set of log files. You can then use these log files to recover your system to a working state if you encounter machine or media failure. Transaction logging occurs only if you have activated and enabled transaction logging on your system.

Use the **UniVerse Transaction Logging** dialog box to set up and manage transaction logging.

Select one of the following methods to access the **UniVerse Transaction Logging** dialog box:

- From the **UniAdmin** window, double-click **Transaction Logging**.
- From the **UniAdmin** menu, select ***Admin,*** then click **Transaction Logging**.
- From the **UniAdmin** toolbar, click the **Manage Transaction Logging** icon, as shown in the following example.

The following example illustrates the **UniVerse Transaction Logging** dialog box:

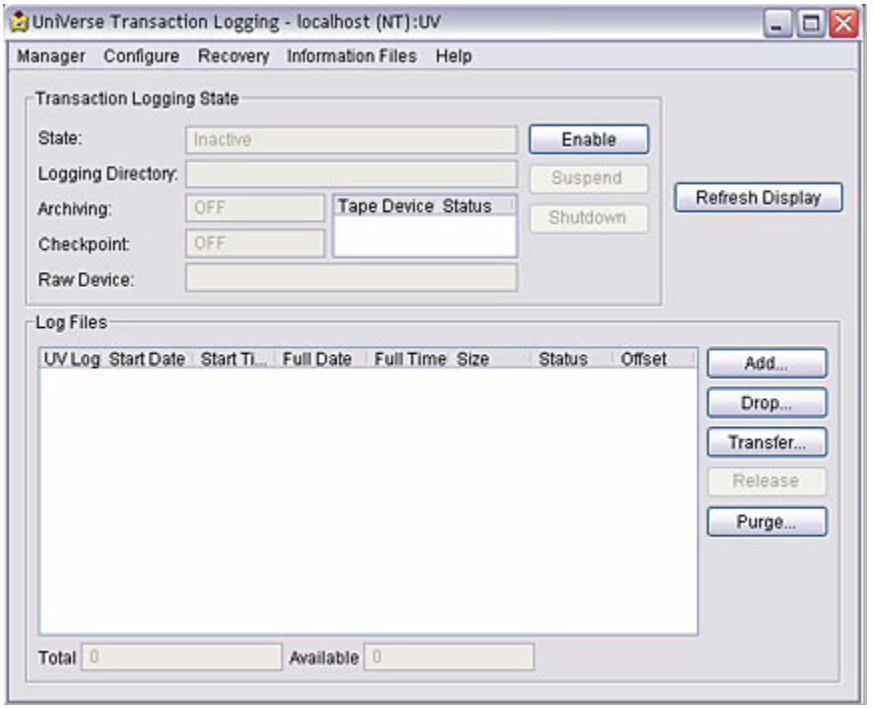

You can perform the following tasks from the **UniVerse Transaction Logging** dialog box:

- Configure transaction logging
- Specify recoverable files
- Manage transaction logging
- Recover files
- View and delete information files

***Note**: The **UniVerse Transaction Logging** dialog box displays the state of transaction logging when the task was activated. To display the current logging state, click **Refresh Display**.*

Configuring the Transaction Logging System

Before you enable and use the transaction logging system, you must configure the system. Steps to configure the system may include:

- Creating a log directory where log files are to be stored.
- Specifying a raw device where log files are to be stored.
- Setting the transaction logging mode.

***Note**: You also need to specify which UniVerse files are recoverable and create log files before you can enable logging.*

Setting Up or Changing Transaction Logging Configuration

Complete the following steps to set up or change the configuration for transaction logging:

1. From the **UniVerse Transaction Logging** dialog box, from the **Configure** menu, choose **Logging**. The **Configure Logging** dialog box appears, as shown in the following example:

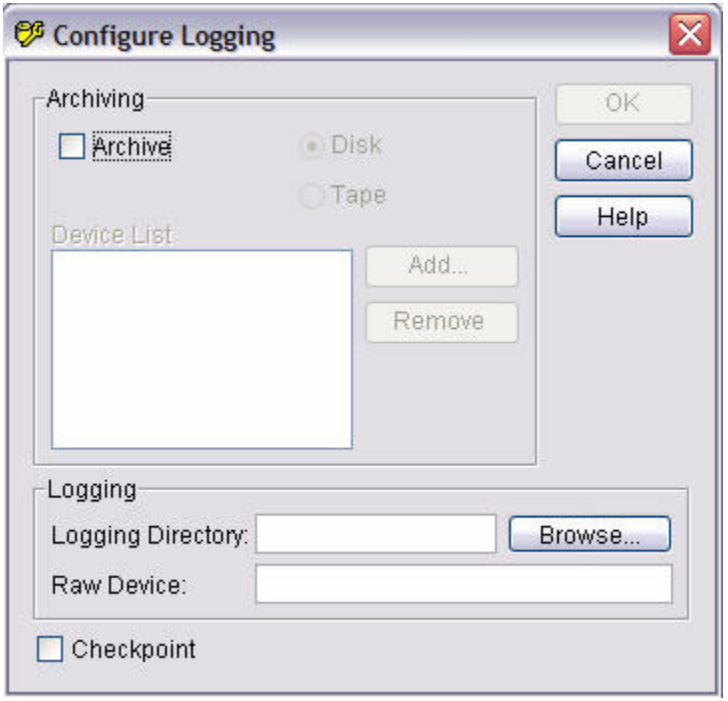

2. If you want to perform archiving, select the **Archive** check box. This is the only type of transaction logging available if you are logging to tape.
3. Choose where to create the log files by selecting either **Disk** or **Tape**.

- **Logging to Tape**. If you are logging to **Tape**, select one or more tape devices from the **Device List**. Initially, this list is empty. To select devices to use, click **Add**. The **Add Device** dialog box appears, which lists all the DC or DT type devices defines in the &DEVICES& file. Select one or more devices from the **Tape Device** list, then click **OK**. UniAdmin adds the devices you choose to the **Device List**. Click **Remove** to remove a tape device from the list.

- **Logging to Disk**. If you want to log to disk, select the **Checkpoint** check box. You can choose this type of transaction in addition to or instead of the **Archive** option.

  Enter the path of the logging directory in the **Logging Directory** box, or click **Browse** to search the system for a suitable directory. If the directory you enter exists, a message box appears to confirm that this is the directory you want to use. Click **Yes** to use this directory.

  Click **OK** to save the settings and close the dialog box.

Note: *To reconfigure transaction logging on a system where it is already enables, you must shut down transaction logging before you click* ***Logging***.

Specifying Recoverable Files

As part of the setup for transaction logging, you must specify which UniVerse files you want to make recoverable.

Activating Files

Complete the following steps to activate files for transaction processing:

1. From the **UniVerse Transaction Logging** dialog box, click **Configure**, then click **Recoverable Files**. The **Configure Recoverable Files** dialog box appears, as shown in the following example:

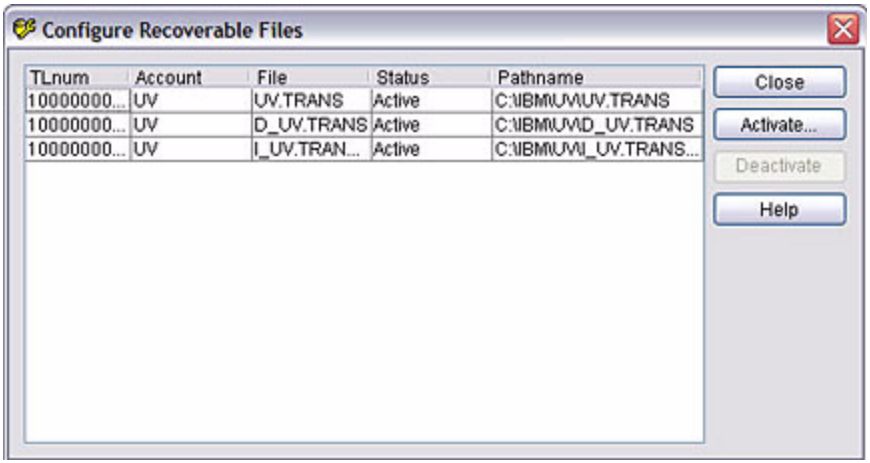

2. To add a file to the **Active Files** list, click **Activate**. The **Activate Files** dialog box appears, as shown in the following example:

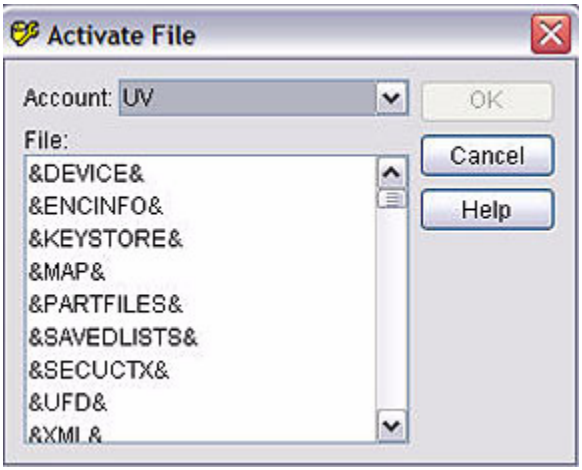

3. Select a UniVerse account from the **Account** list.
4. Select a file to activate from the **Files** list. You cannot perform transaction logging against a type 1 or type 19 file. If you select a file with either of these types, UniAdmin displays an error message. You must acknowledge this message and select another file.
5. Click **OK**. The **Activate File** dialog box closes, and UniAdmin updates the **Configure Recoverable Files** dialog box with a list of activated UniVerse files.

Note: *You can also activate a deactivated file. To do this, choose the file from the Active Files list, then click* ***Activate****. This details of the file automatically appear in the* ***Configure Recoverable Files*** *dialog box. Click* ***OK****. A message box informs you that this file has been activated on a previous occasion. Click* ***Yes****. The file entry in the* ***Configure Recoverable Files*** *dialog box is updated with an Active status.When you activate a deactivated file, any transactions that occurred on that file when it was deactivated are lost.*

Deactivating Files

Complete the following steps to deactivate files:

1. From the **Configure Recoverable Files** dialog box, click **Configure** then click **Recoverable Files**. The **Configure Recoverable Files** dialog box appears, as shown in the following example:

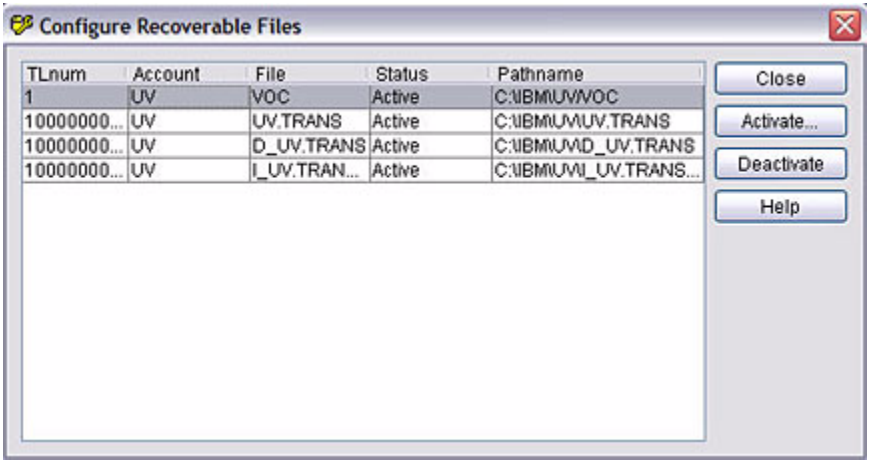

2. Select the file you want to deactivate from the list.
3. Click **Deactivate**. A message box confirming that you want to deactivate the file appears.
4. Click **Yes.** UniAdmin updates the **Configure Recoverable Files** dialog box to show the file is now deactivated.
5. Click **Close** to exit this dialog box, saving the changes.

Note: *Deactivated files remain in the* ***Active File*** *list. You cannot remove them from the list.*

Managing Transaction Logging

Transaction logging management falls into two main areas, managing log files and managing transaction logging.

Managing Log Files

This section describes how to manage your log files, including:

- Adding Log Files
- Deleting Log Files
- Transferring Log Files
- Purging Log Files
- Releasing Log Files
- Releasing a Tape

Adding Log Files

Before you can enable transaction logging, your must configure the system, including specifying the path of the logging directory and creating the log files. You can create additional log files for a system that has logging enabled.

Complete the following steps to create log files:

1. From the **UniVerse Transaction Logging** dialog box, click **Add**, or from the **Manager** menu, click **Add Logs**. A dialog box similar to the following example appears:

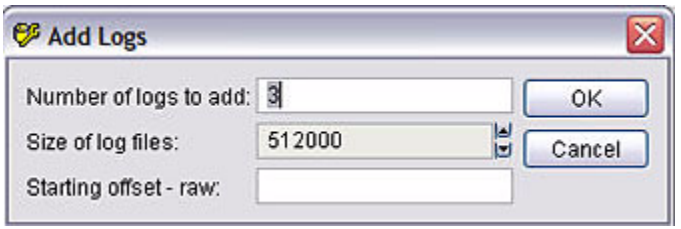

2. Enter the number of log files to add or create in the **Number of logs to add** box. The default number is 3.

3. Enter the size of the log files, in bytes, in the **Size of log files** box. This number is always rounded to the nearest increment of 512 bytes. You can also select the size of the log files from the list. The default size is 512000 bytes.
4. If you are using raw disk for your log files, enter the starting offset for the log files in the **Starting offset - raw** box.
5. Click **OK**. UniVerse creates the log files, with a status of Available, and adds them to the **Log Files** list in the **UniVerse Transaction Logging** dialog box.

***Note**: This option is available only if you are logging to disk. If you are logging to tape, UniVerse creates a single log file automatically.*

Deleting Log Files

You can delete unneeded log files if the status of the log file is inactive, and the log files have not been used. Once you have deleted log files, the log file sequence numbers become available and are used by the next log files you create.

Complete the following steps to delete log files:

1. Do one of the following to access the **Drop Logs** dialog box:
   - From the **UniVerse Transaction Logging** dialog box, click **Drop.**
   - From the **Manager** menu, click **Drop Log Files**.

   The **Drop Logs** dialog box appears, as shown in the following example:

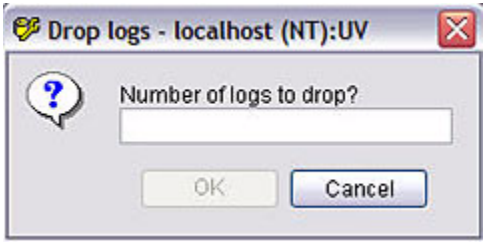

2. Enter the number of logs files you want to delete in the **Number of logs to drop** box.
3. Click **OK.** UniVerse deletes the logs files with the highest log file specification numbers. If you specify a number of log files greater than the number available, UniVerse only deletes those available.

Note: *This option is available only if you are logging to disk.*

Transferring Log Files

You can back up full log files to tape before releasing them. The log files are backed up to a device you choose and are automatically released on disk. You can then use the backed up copies of the log files to recover data at a later date.

Complete the following steps to transfer a log file:

1. Do one of the following to access the **Transfer Logfiles** dialog box:
   - From the **UniVerse Transaction Logging** dialog box, click **Transfer.**
   - From the **Manager** menu, click **Transfer Log files**.
2. Select a tape device from the **Tape Device** list. This list contains all the DC and DT tape devices defined in the &DEVICES& file. After you choose a device, make sure that a tape is mounted.
3. Specify the log files to transfer in the **First Log to Backup** and **Last Log to Backup** boxes. Note that these boxes are automatically updated with the numbers of the full log files currently on the system. If you try to enter numbers for log files that do not have a status of Full, an error message appears. You must acknowledge this message and reenter suitable values.
4. Click **OK**. The full log files are backed up to the chosen device. Once a check has been made to make sure the backup was successful, the log files are released to disk.

Purging Log Files

If you previously transferred or released a full log file, remove its entry from the **UV.LOGS** file using the **Purge** option. Regularly purging released files reduces the number of files that are searched during data recovery.

Purging is based on a date you specify. UniVerse purges any log files with a Full date before the date you specify.

Complete the following steps to purge a log file:

1. Execute one of the following to access the **Purge Log Entries** dialog box:
   - From the **UniVerse Transaction Logging** dialog box, click **Purge**.
   - From the **Manager** menu, click **Purge Log Files**.

2. Enter a date in the **Date (MM/DD/YY)** box.
3. Click **OK**. UniVerse removes the entries before the date you specified from the UV_LOGS file, and updates the **UniVerse Transaction Logging** dialog box.

Note: *This option is only available if you are logging to disk.*

Releasing Log Files

When a log file is full and you do not need the contents of the log file, you can release it. UniVerse deletes the contents of the log file, and updates the status in the **Log Files** box. The space used by the released file is reallocated to a new log file to maintain the number of available log files.

To keep a copy of the contents of the log file, in case you want to recover data from the file later, use the **Transfer** option instead.

Complete the following steps to release a full log file:

1. Select the full log file from the **Log Files** list in the **UniVerse Transaction Logging** dialog box.
2. Execute one of the following to release the log file:
   - From the **UniVerse Transaction Logging** dialog box, click **Release**.
   - From the **Manager** menu, click **Release Log File**.
3. A message box appears. Click **Yes** to release the log files.

Once you release the log file, you can remove its entry from the UV.LOGS file using the **Purge** option.

Releasing a Tape

When you log to tape, you can choose to log to one or more devices.

If you log to a single device, UniVerse suspends transaction logging when the tape is full. If you log to several devices and a tape is full, logging automatically continues to the next available device. UniVerse suspends transaction logging when all the tapes are full.

In both of these cases, you must mount a new tape and reenable logging. UniVerse creates a new log file and logging continues.

You can prevent the suspension of transaction logging when you log to multiple tapes by replacing the tape when it is full. When the new tape is mounted, you can release the tape, making it available again when the next device is full.

Complete the following steps to release a tape:

1. Select the device that is full from the **Tape Device** list on the **Transaction Logging** dialog box.
2. From the **Manager** menu, select **Release Tape**. A message box appears.
3. Click **Yes** to release the tape. UniVerse updates the state of the tape device.

Managing Transaction Logging

As a UniVerse administrator, you can enable, suspend, or disable the transaction logging process.

Enabling Transaction Logging

To enable transaction logging system-wide, do one of the following:

- From the **UniVerse Transaction Logging** dialog box, click **Enable**.
- From the **Manager** menu, click **Enable Logging**.

It may take some time for the dialog box to be updated with the Enabled status. Click **Refresh Display** to display the new transaction logging state.

***Note**: Transaction logging is successfully enabled only if you set the TXMODE configurable parameter to 1 and restart UniVerse.*

In the enabled state, updates to recoverable files are written to your UniVerse files and the current log file.

Suspending Transaction Logging

When you suspend transaction logging, UniVerse disallows updates to recoverable files for the duration of the suspension. To suspend transaction logging system-wide, do one of the following:

- From the **UniVerse Transaction Logging** dialog box, click **Suspend**.
- From the **Manager** menu, click **Suspend Logging**.

Note: *Transaction logging is automatically suspended when all log files are full.*

It may take some time for the dialog box to be updated with the suspended status. Click **Refresh Display** to display the new transaction logging state.

You must use the **Enable Logging** option to reinstate transaction logging after it has been suspended.

Disabling Transaction Logging

If you disable transaction logging, updates to recoverable files are written to your UniVerse files, but not to the log files. This means that any changes made to the UniVerse files while transaction logging is disabled cannot be recovered.

To disable transaction logging, do one of the following:

- From the **UniVerse Transaction Logging** dialog box, click **Shutdown**.
- From the **Manager** menu, click **Shutdown Logging**.

It may take some time for the dialog box to be updated with the disabled status. Click **Refresh Display** to display the new transaction logging state.

You must use the **Enable Logging** option to reinstate transaction logging after shutdown.

Viewing and Deleting Information Files

Transaction logging contains the following three information files:

- *uvlogd.info* – The logging information file.
- *uvchkd.info* – The checkpoint information file.
- *uvrolf.info* – The roll-forward information file.

You can view or delete the contents of an information file using the **Information Files** menu options.

Viewing an Information File

To view an information file, from the **Information Files** menu in the **UniVerse Transaction Logging** dialog box, select one of the following choices:

- **View Logging Information**
- **View Checkpoint Information**
- **View Rollforward Information**

The contents of the file you chose appears in an output window. Click **Close** to close this window.

Note: If a file is too large to display in the output window, UniVerse displays a message box and displays only the last 16K bytes of the information file.

Deleting the Contents of an Information File

To delete the contents of an information file, from the **Information Files** menu in the **UniVerse Transaction Logging** dialog box, select one of the following choices:

- **Delete Logging Information**
- **Delete Checkpoint Information**
- ***Delete Rollforward Information***

When the message box appears indicating the contents have been deleted, click **Yes.**

Restoring Log Files

To recover files from transferred log files, you must first restore the log files from tape to disk. When the log files have been restored, you can recover the files you need.

Complete the following steps to restore log files from tape:

1. From the **Recovery** menu in **UniVerse Transaction Logging** dialog box, select **Rollforward from Disk.** The **Recover Files** dialog box appears, as shown in the following example:

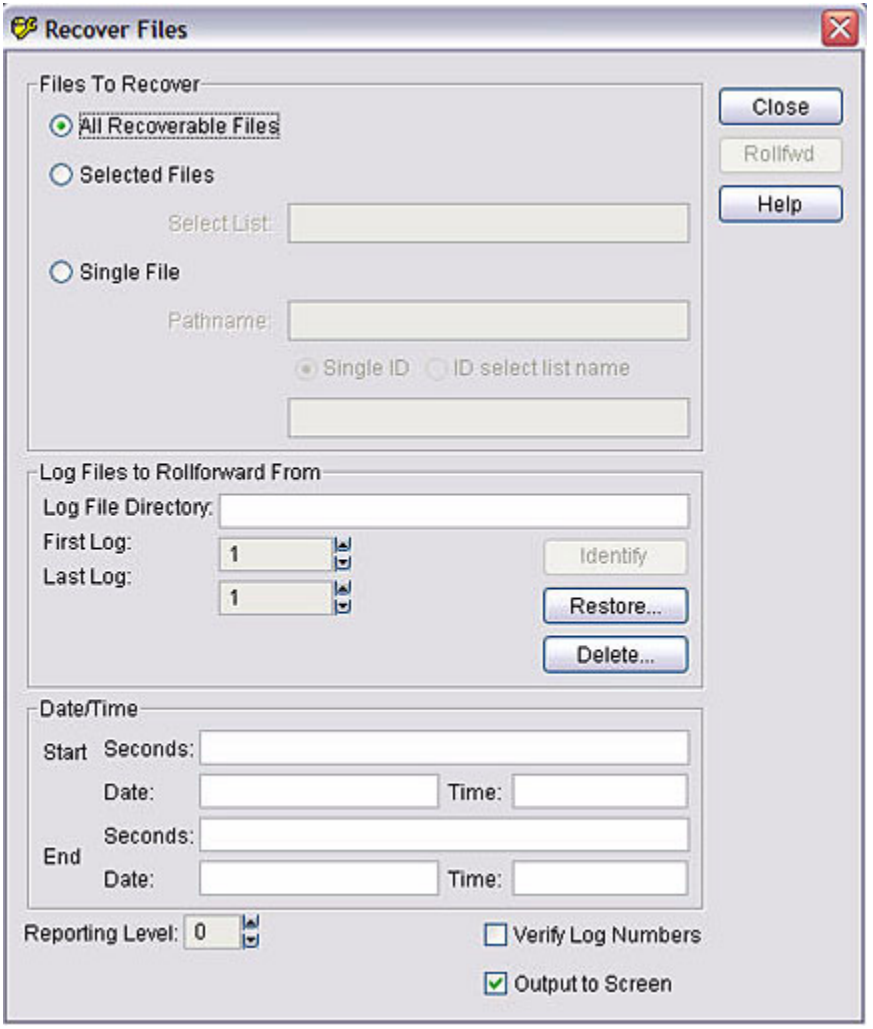

2. In the **Log Files to Rollforward From** area of the **Recover Files** dialog box, click **Restore**. The **Restore Logs** dialog box appears, as shown in the following example:

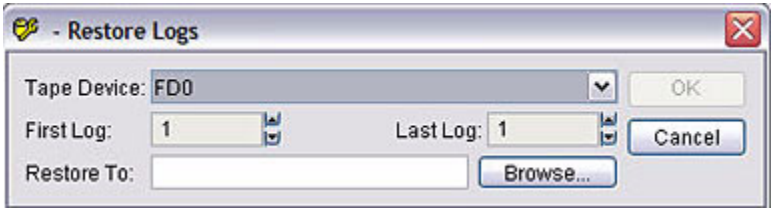

3. In the **Tape Device** list, select a tape device.
4. Choose the log files you want to restore by entering a range of numbers in the **First Log** and **Last Log** boxes. Use the arrows to increase or decrease these values. The default is 1.
5. Choose the directory to which you want to restore the log files in the **Restore to** box, or click **Browse** to search the system for the directory.
6. Click **OK**. UniVerse restores the chosen log files from tape and copies them to the chosen directory.

Recovering Files

This section provides step-by-step instructions for recovering log files from disk and log files from tape.

Recovering Log Files from Disk

Complete the following steps to recover files from disk:

1. Chose the file or files you want to recover by clicking the appropriate option in **Files to Recover** area of the **Recover Files** dialog box:
   - **All Recoverable Files** – UniVerse recovers all the activated files.
   - **Selected Files** – UniVerse uses the select list name you specify to recover files. You must create the select list in the UV account directory, and the select list must contain the paths of the files you want to recover. Enter the name of the select list in the **Select List** box.
   - **Single File** – UniVerse recovers a single file. Enter the full path of the file in the **Pathname** box.
   - **Specific IDs** – UniVerse recovers a specific record ID or list of record IDs. Click **Single File**. In the **Pathname** box, enter the full path to the account where the record ID or list of record IDs exists. Click **Single ID** if you want to recover one record ID, then enter the ID in the box beneath **Single ID**.

     Click **ID select list name** if you want to recover a list of record IDs, then enter the name of the saved list containing the list of IDs from the &SAVEDLIST& file in the box beneath **ID select list name**.
2. Specify the location of the log file or files by entering the path in the **Log File Directory** box.
3. Enter the numbers of the first and last log files in the **First Log** and **Last Log** boxes. You can use the arrows to increase or decrease these values. The default is 1.
4. If you enabled checkpointing, you can use the **Identify** option. UniVerse searches the log file directory and identifies the log files containing the required files. UniVerse automatically updates the **First Log** and **Last Log** boxes with recommended values. This option is available only if you are restoring all recoverable files or selected files.

5. If you want to specify the starting date and time and ending date and time to recover log files, complete the following steps:
   - In the **Date/Time** area of the **Recover Files** dialog log, enter the number of seconds from January 1, 1970 at 00:00:00 GMT at which you want to start recovery in the **Start Seconds** box. Do not specify this field if you are defining the Date and Time.
   - Enter the date on which you want to start recovery in the **Start Date** box. The date must be entered in the yyyy-MM-dd format. Do not enter a date if you specified seconds.
   - Enter the time at which you want to start recovery in the **Start Time** box. The time must be entered in the HH:mm:ss format. This field is necessary if you are defining the date to start recovery. The default value for the time field is 00:00:00. Do not enter a date if you specified seconds.
   - In the **Date/Time** area of the **Recover Files** dialog log, enter the number of seconds from January 1, 1970 at 00:00:00 GMT at which you want to end recovery in the **End Seconds** box. Do not specify this field if you are defining the Date and Time.
   - Enter the date on which you want to end recovery in the **End Date** box. The date must be entered in the yyyy-MM-dd format. Do not enter a date if you specified seconds.
   - Enter the time at which you want to end recovery in the **End Time** box. The time must be entered in the HH:mm:ss format. This field is necessary if you are defining the date to end recovery. The default value for the time field is 23:59:59. Do not enter a date if you specified seconds.
6. Choose how much detail to report during the recovery by entering an appropriate number in the **Reporting Level** box. You can use the arrows to increase or decrease this value. The minimum setting is 0 (no reporting), and the maximum setting is 3 (highest level of reporting).
7. Select the **Verify Log Numbers** check box if you want the rollforward program to verify the log numbers during recovery.
8. Choose where to report recovery information. If you select the **Output to Screen** check box, all reports appear in the **UniVerse Command Output** window. This is the default setting. If you clear the **Output to Screen** check box, UniVerse writes the recovery information to the *uvrolf.info* file (the roll-forward information file), which you can review and delete at a later date.

9. Click **Rollfwd** to start the recovery. The **UniVerse Command Output** window appears.
10. When the roll-forward is complete, click **Close** to close the **UniVerse Command Output** window. You can now delete any restored log files.

Completing the Roll-Forward

If you need to restore more log files from tape, UniVerse prompts you to do so now. Using another window, complete steps 2 through 4 of the Recovering Files steps until all require log files are processed.

1. From the **Recovery** menu in the **UniVerse Transaction Logging** dialog box, select **Rollforward from Disk**. In the **Log Files to Rollforward From** area of the **Recover Files** dialog box, click **Delete** to delete all of the log files just restored and rolled forward to free up disk space.
2. From the **Recovery** menu in the **UniVerse Transaction Logging** dialog box, select **Rollforward from Disk**. In the **Log Files to Rollforward From** area of the **Recover Files** dialog box, click **Restore** to restore as many of the still required log files as possible.
3. From the **Recovery** menu in the **UniVerse Transaction Logging** dialog box, select **Rollforward from Disk**, the click **Rollfwd** to roll forward the log files.
4. If possible, make a full backup of all files.
5. Enable transaction logging by clicking **Enable** in the **UniVerse Transaction Logging** dialog box, or, from the **Manager** menu, select **Enable Logging.**
6. Resume normal operations.

Recovering Files from a Full Log File on Tape

Complete the following steps to recover files from a full log file on tape:

1. From the **Recovery** menu in the **UniVerse Transaction Logging** dialog box, select **Rollforward from Tape**. The **Recover Files From Tape** dialog box appears, as shown in the following example:

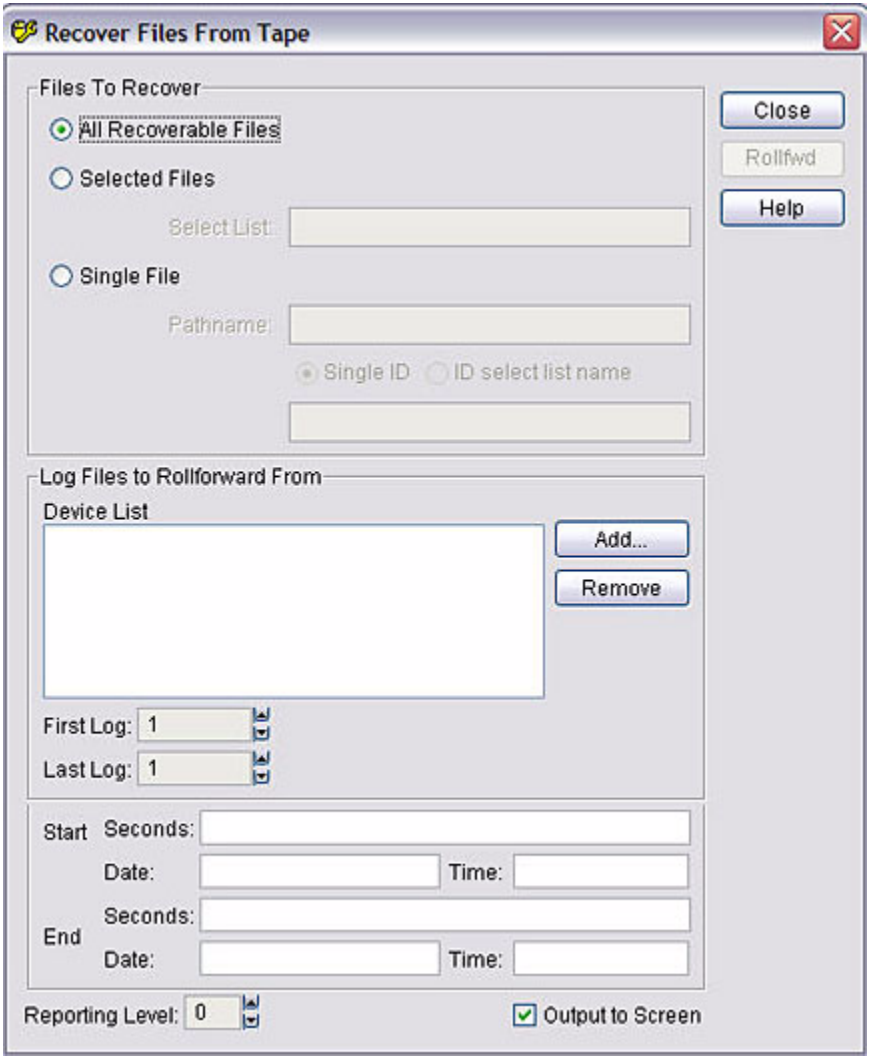

2. Chose the file or files you want to recover by clicking the appropriate option in **Files to Recover** area of the **Recover Files** dialog box:

   - **All Recoverable Files** – UniVerse recovers all the activated files.
   - **Selected Files** – UniVerse uses the select list name you specify to recover files. You must create the select list in the UV account directory, and the select list must contain the paths of the files you want to recover. Enter the name of the select list in the **Select List** box.
   - **Single File** – UniVerse recovers a single file. Enter the full path of the file in the **Pathname** box.
   - **Specific IDs** – UniVerse recovers a specific record ID or list of record IDs. Click **Single File**. In the **Pathname** box, enter the full path to the account where the record ID or list of record IDs exists. Click **Single ID** if you want to recover one record ID, then enter the ID in the box beneath **Single ID**.

     Click **ID select list name** if you want to recover a list of record IDs, then enter the name of the saved list containing the list of IDs from the &SAVEDLIST& file in the box beneath **ID select list name**.

3. Select the device or devices to use from the **Device List**. Initially this list is empty. To select devices to use, click **Add**. The **Add Device** dialog box appears, as shown in the following example:

The **Add Device** dialog box lists all the DC or DT devices defined in the &DEVICE& file.

4. Select a device or devices from the tape devices list, then click **OK**. UniVerse adds the devices to the **Device List** on the **Recover Files From Tape** dialog box. Click **Remove** to remove a file from this list.

   Note: *You must choose the devices in the order in which they will be used. If you choose a single device and the log file spans more than one tape, UniVerse prompts to change the tape at the appropriate time. If you use more than one device and the log file spans more than one tape, you must load reel 1 on the first device listed in the Device List.*

5. Enter the numbers of the first and last log files in the **First Log** and **Last Log** boxes. You can use the arrows to increase or decrease these values. The default is 1.

6. If you want to specify the starting date and time and ending date and time to recover log files, complete the following steps:

   - In the **Date/Time** area of the **Recover Files** dialog log, enter the number of seconds from January 1, 1970 at 00:00:00 GMT at which you want to start recovery in the **Start Seconds** box. Do not specify this field if you are defining the Date and Time.
   - Enter the date on which you want to start recovery in the **Start Date** box. The date must be entered in the yyyy-MM-dd format. Do not enter a date if you specified seconds.
   - Enter the time at which you want to start recovery in the **Start Time** box. The time must be entered in the HH:mm:ss format. This field is necessary if you are defining the date to start recovery. The default value for the time field is 00:00:00. Do not enter a date if you specified seconds.
   - In the **Date/Time** area of the **Recover Files** dialog log, enter the number of seconds from January 1, 1970 at 00:00:00 GMT at which you want to end recovery in the **End Seconds** box. Do not specify this field if you are defining the Date and Time.
   - Enter the date on which you want to end recovery in the **End Date** box. The date must be entered in the yyyy-MM-dd format. Do not enter a date if you specified seconds.
   - Enter the time at which you want to end recovery in the **End Time** box. The time must be entered in the HH:mm:ss format. This field is necessary if you are defining the date to end recovery. The default value for the time field is 23:59:59. Do not enter a date if you specified seconds.

7. Choose where to report recovery information. If you choose **Output to Screen**, all reports appear in the **UniVerse Command Output** window. This is the default setting. If you clear this option, UniVerse writes the recovery information to the *uvrolf.info* file (the roll-forward information file), which you view and delete at a later date.
8. Choose how much detail to report during the recovery process by entering a suitable value in the **Reporting Level** box. You can use the arrows to increase or decrease this value. The minimum setting is 0 (no reporting), and the maximum setting is 3 (the highest level of reporting).
9. Click **Rollfwd** to start the recovery. The **UniVerse Command Output** window appears.
10. When the roll-forward is complete, click **Close** to close the **UniVerse Command Output** window.
11. Click **Close** to close the **Recover Files** dialog box.

Deleting Restored Log Files

If you recover files from a transferred log file, you can delete the restored log files as soon as the recovery is complete. Complete the following steps to delete restored log files:

1. From the **Recovery** menu in the **UniVerse Transaction Logging** dialog box, select **Rollforward from Disk**. In the **Log Files to Rollforward From** area of the **Recover Files** dialog box, click **Delete**. The **Delete Restored Logs** dialog box appears, as shown in the following example:

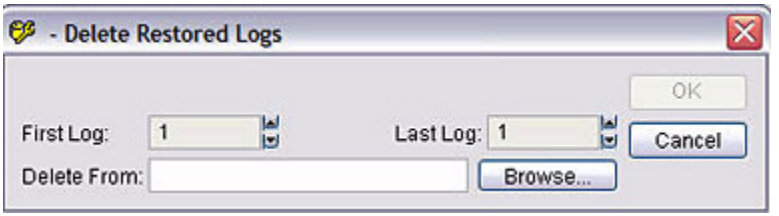

2. Enter the range of log files you want to delete in the **First Log** and **Last Log** boxes. You can use the arrow to increase or decrease these values.
3. Enter the path of the directory containing the restored log files in the **Delete From** box, or click **Browse** to search for the directory.

4. Click **OK**. The **Delete Restored Logs** dialog box closes, and a message box appears.
5. Click **Yes** to delete the chosen log files.
6. Click **Close** to close the **Recover Files** dialog box.

Chapter

19

Executing UniVerse Commands

Executing a UniVerse Command 19-3
UniVerse Administration Commands 19-5
Commands to Assign and Unassign Peripheral Devices 19-5
UniVerse BASIC Commands 19-5
File and Account Maintenance Commands 19-6
Managing Processes 19-8
Monitoring Resources 19-8
Task and Record Locking 19-9
Using the Command History 19-10
Reexecuting Commands 19-10
Editing a Command 19-11
Saving Commands 19-11
The UniVerse Command Output Window 19-12

Use the **UniVerse Command** dialog box to execute UniVerse commands from UniAdmin. As you issue commands, UniVerse creates a command history, which you can use to reexecute a command or save commands in the VOC file.

Select one of the following methods to access the **UniVerse Command** dialog box:

- From the **UniAdmin** window, double-click **UniVerse Command**.
- From the **UniAdmin** menu, select ***Admin,*** then click **UniVerse Command**.
- From the **UniAdmin** toolbar, click the **Manage Transaction Logging** icon, as shown in the following example.

The following example illustrates the **UniVerse Command** dialog box:

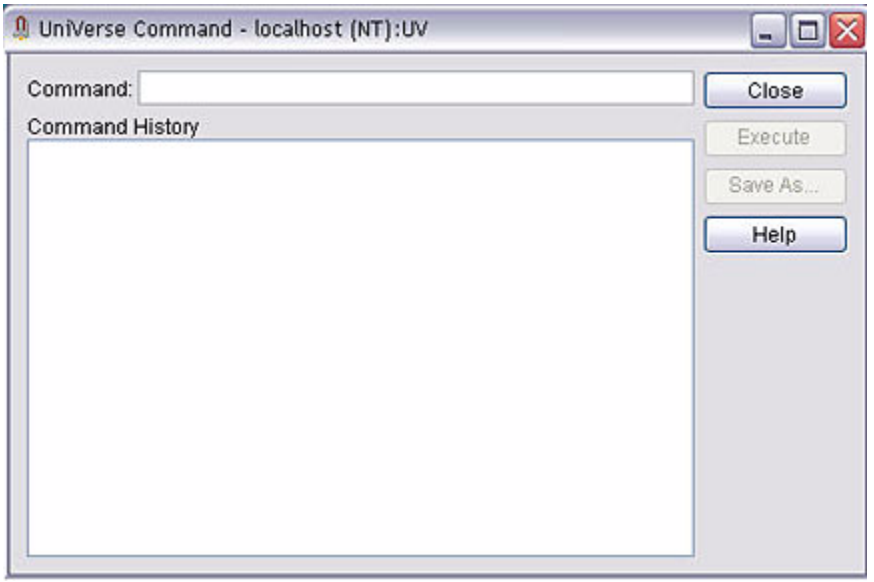

Executing a UniVerse Command

Complete the following steps to execute a UniVerse command:

1. Enter the UniVerse command you want to execute in the **Command** box.
2. Click **Execute**. The result of the command appears in the **UniVerse Command Output** window.

   The following example illustrates the output from the LIST VOC command:

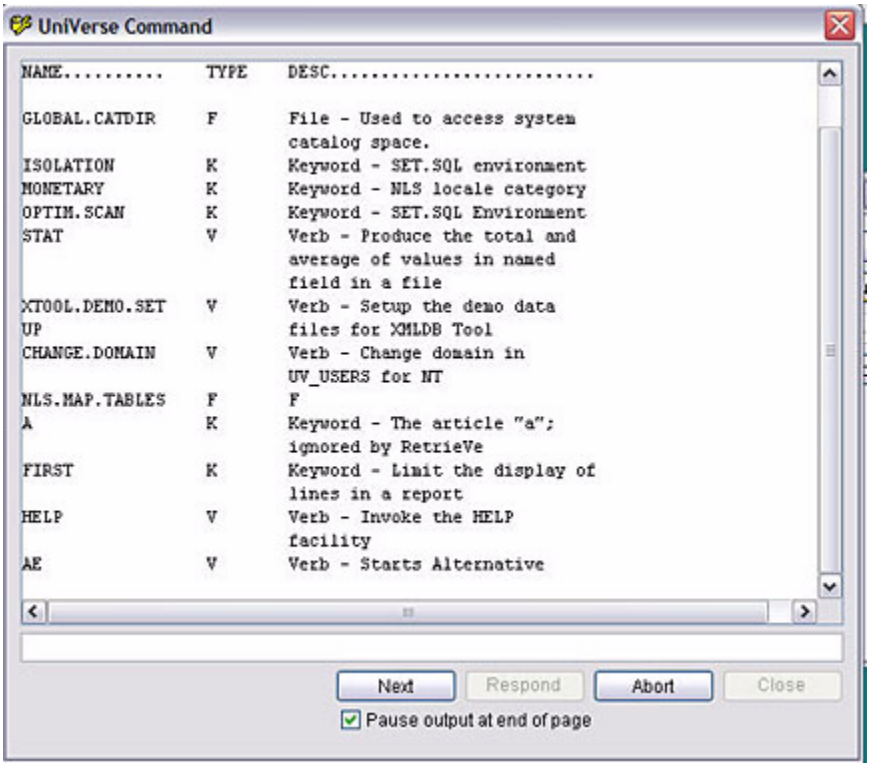

3. If the command output requires more than one page, click **Next** to continue the display. Clear the **Pause output at end of page** check box if you do not want the output to pause at the end of every page.
4. If the command requires user input, enter an appropriate reply in the text entry box, then click ***Respond***. Repeat entering input until no further input is required.

5. To quit from command output, click **Abort**.
6. Click **Close** to close the **UniVerse Command Output** window. You can enter more commands, use the command history, or exit the dialog box.

UniVerse Administration Commands

You can perform UniVerse administration tasks by issuing specific UniVerse administration commands. Many of these commands can be used only by a user logged in to the UV account as a UniVerse Administrator. UniAdmin is an interface to some of the commands listed in this section. For ease of reference, the commands are organized in functional groups and are listed alphabetically.

Commands to Assign and Unassign Peripheral Devices

The following table describes commands to assign and unassign peripheral devices.

Command	Description
ASSIGN	Assigns a device for your exclusive use. Usually used for handling tape and disk devices.
UNASSIGN	Relinquishes control of a physical device that has been assigned to you.

Commands to Assign and Unassign Peripheral Devices

UniVerse BASIC Commands

The following table describes commands used with UniVerse BASIC.

Command	Description
BASIC	Compiles a UniVerse BASIC program.
CATALOG	Copies the compiled object code to the system catalog space.
DELETE.CATALOG	Deletes programs from the catalog space.
ED	Edits the contents of data files, file dictionaries, select lists, and basic programs.
INITIALIZE.CATALOG	Initializes the catalog space.
LIMIT	Sets the maximum size of memory storage for a user's active UniVerse BASIC routines.

BASIC Commands

Command	Description
MAKE.MAP.FILE	Creates the &MAP& file from the catalog contents.
MAP	Displays information about the contents of the catalog space.
RAID	Debugs a UniVerse BASIC program.
VCATALOG	Compares the object code of a program in the catalog to object code in the original file.
VLIST	Lists UniVerse BASIC object code.

BASIC Commands (continued)

File and Account Maintenance Commands

The following table describes commands to use for file and account maintenance.

Command	Description
ACCOUNT.FILE.STATS	Gathers file statistics on the current state of selected files.
ANALYZE.FILE	Displays statistics about a dynamic file.
CLEAN.ACCOUNT	Performs routine maintenance and verifies the condition of files in an account.
CONFIGURE.FILE	Changes the parameters of a dynamic file.
FILE.STAT	Displays statistical information about the file size and record partitioning in a file.
FILE.USAGE	Displays statistics on the use patterns of a file.
FILE.USAGE.CLEAR	Resets statistics displayed by the FILE.USAGE command.
FORMAT.CONV	Changes the storage format of UniVerse files or BASIC object code.
GROUP.STAT	Displays information about the record distribution in a file.
GROUP.STAT.DETAIL	Displays a detailed record distribution summary for a file.
HASH.AID	Displays statistical information about the hypothetical file size and record partitioning in the file.

File and Account Maintenance

Command	Description
HASH.HELP	Displays a recommendation for a file type, modulo, and separation based on the current record IDs and file size.
HASH.HELP.DETAIL	Displays the same information as HASH.HELP and also includes details on the record ID sizes and record size in bytes.
HASH.TEST	Displays how the record distribution would work with a hypothetical file type, modulo, and separation.
HASH.TEST.DETAIL	Displays the same hypothetical information as HASH.TEST, and also includes the hypothetical number of bytes per record, and the number of bytes per group.
LIST.FILE.STATS	Displays the file statistics gathered by ACCOUNT.FILE.STATS.
RECORD	Determines the group that a record should reside in, or if the record exists, verifies that the record is there.
RESIZE	Changes the structure of a file with a new file type, modulo, or separation.
SUSPEND.FILES	Suspends UniVerse processes that make changes to files, without terminating user processes.
VVOC	Compares the contents of the VOC file in the current account to those of NEWACC and reports the differences.

File and Account Maintenance (continued)

Managing Processes

The following tables describes commands used to manage processes.

Command	Description
AUTOLOGOUT	Logs the user out of UniVerse after a period of inactivity.
ENVIRONMENT	Sets and displays environment variables.
JOBS	Lists active phantom processes.

Commands for Managing Processes

Command	Description
MESSAGE	Sends a message from your terminal to another user.
PHANTOM	Starts a phantom process.
SLEEP	Suspends a process.

Commands for Managing Processes (continued)

Monitoring Resources

The following tables describes commands used for monitoring resources.

Command	Description
CONFIG	Displays information about current authorization parameters and configurable parameter values.
LISTU	Displays information about the users currently logged in to the system.

Commands for Monitoring Resources

Task and Record Locking

The following table describes commands used for task and record locking.

Command	Description
CLEAR.LOCKS	Clears a specific lock number or all the locks set by an account.
LIST.LOCKS	Lists the 64 task synchronization locks, informing you which are still available.
LIST.READU	Displays a list of locked files and records.
LOCK	Reserves one of the 64 task synchronization locks to guarantee that you can process a file record without interference from others.
MASTER	Releases task synchronization locks set with the LOCK command.

Commands Used for Task and Record Locking

Command	Description
RELEASE	Releases record locks that were set by the BASIC commands READU, MATREADU, and READVU.
SEMAPHORE.STATUS	Displays information about the status of system semaphores.
UNLOCK	Clears file, group, and update locks.

Commands Used for Task and Record Locking (continued)

Using the Command History

You can reexecute, edit, or create sentences or paragraphs from commands listed in the **Command History** area of the **UniVerse Command** dialog box. The following example illustrates command history:

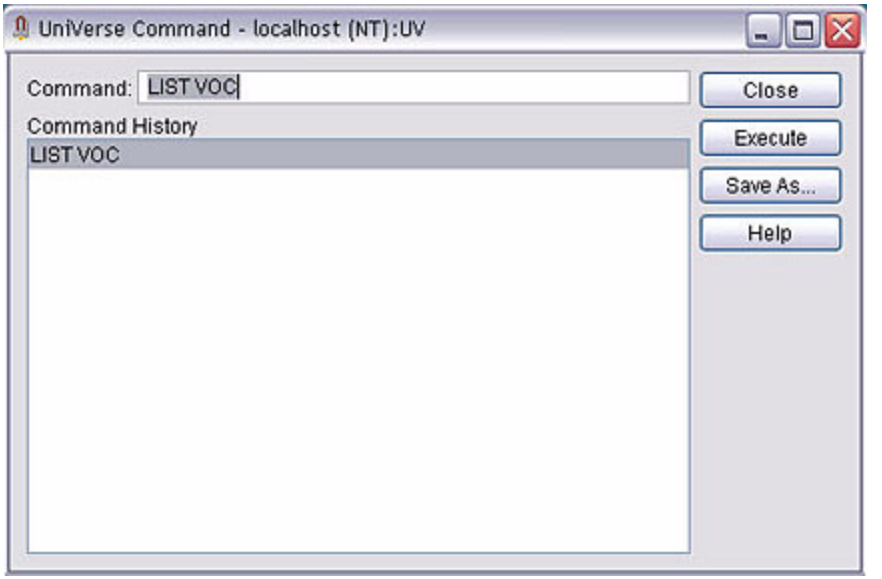

Reexecuting Commands

Complete the following steps to reexecute a command listed in **Command History**:

1. Double-click the command in the **Command History** portion of the **UniVerse Command** dialog box. The command appears in the **Command** box.
2. Click **Execute**. The result of the command appears in the **UniVerse Command Output** window.

Editing a Command

Complete the following steps to edit a command that appears in the **Command History** area of the **UniVerse Command** window.

1. Double-click the command in the **Command History** area. The command appears in the **Command** box.
2. Click the **Command** field and make your changes to the command.
3. Click **Execute** to run the new command. UniVerse adds the new command to the **Command History** area, and the result appears in the **UniVerse Command Output** window.

Saving Commands

Complete the following steps to save commands to the VOC file:

1. Select the commands you want to save from the **Command History** area of the **UniVerse Command** window.
2. Click **Save As**. The **Save Command(s)** dialog box appears, as shown in the following example:

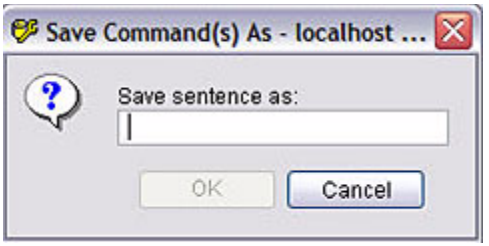

3. If you are saving one command, enter a unique name for the sentence. If you are saving multiple commands, enter a unique name for the paragraph.
4. Click **OK** to save the new sentence or paragraph in the VOC file.

The UniVerse Command Output Window

The **UniVerse Command Output** window displays the results of the execution of a UniVerse command. The output displays can be the result of using the **UniVerse Command, Restore,** or **Transaction Logging** options of UniAdmin.

The following example illustrates the **UniVerse Command Output** window:

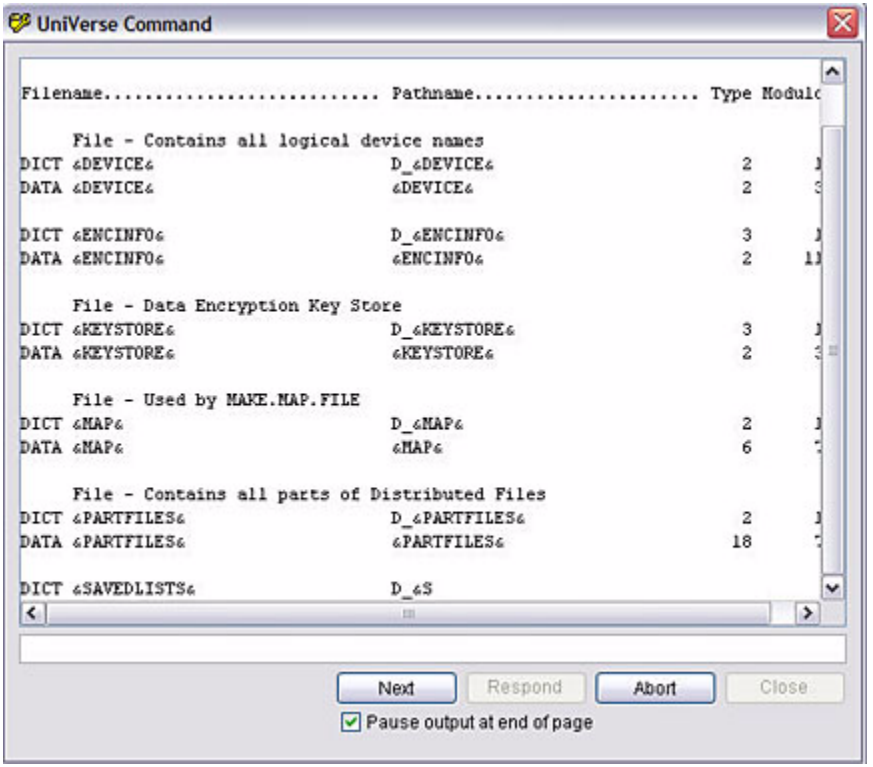

The output display area of the **UniVerse Command Output** window displays the results of the UniVerse command. The number of lines UniAdmin displays in this area depends on the size of the window. When the number of lines exceeds the display area, UniAdmin pauses the output. Click **Next** to continue the display.

Use the text entry box to enter further user input if required by the command. Enter the appropriate information, then click **Reply**.

The four action button on the **UniVerse Command Output** window are:

- **Next** – Continues paused output.
- **Respond** – Enters the response in the text entry box.
- **Abort** – Terminates a command. The **Abort** button is active when output is paused or when user input is required.
- **Close** – Closes the **UniVerse Command Output** window.

If you select the **Pause output at end of page** check box, UniAdmin pauses the output at the end of each page, and you must click **Next** to continue displaying the output. If you clear this check box, the output scrolls continuously.

Chapter

20

UniVerse User Administration

Interactive Users 20-4
Background Processes 20-4
Terminating a Process 20-5
Terminating a User Process 20-5
Terminating a Background Process 20-5
Sending Messages with UniVerse Admin 20-6
Sending Messages on UNIX Systems 20-6
Sending Message on Windows Platforms 20-7

To view or terminate UniVerse processes and to send messages, use the **User Administration** dialog box. The **User Administration** dialog box contains a snapshot of the user and background processes at the time when you invoked the dialog box.

Select one of the following methods to access the **User Administration** dialog box:

- From the **UniAdmin** window, double-click **Users**.
- From the **UniAdmin** menu, select ***Admin,*** then click **Users**.
- From the **UniAdmin** toolbar, click the **Users** icon, as shown in the following example.

The following example illustrates the **User Administration** dialog box:

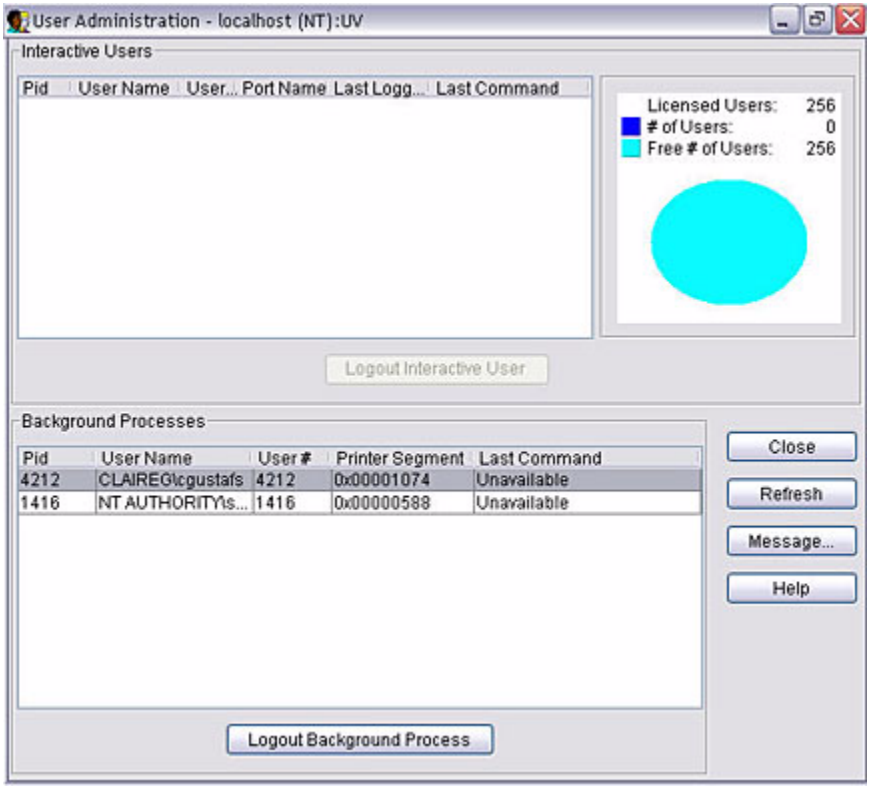

The tasks you can perform from the **User Administration** dialog box include:

- Terminating a user process
- Sending messages

The **User Administration** dialog box is divided into two main areas:

- Interactive Users
- Background Processes

Interactive Users

The following table describes the information that appears in the **Interactive Users** area of the **User Administration** dialog box for each user process.

Parameter	Description
Pid	The process ID. On a Windows platform, this is the UniVerse user number. It is the same as the value shown in the User # column.
User Name	The user's login name.
User #	The user's UniVerse user number.
Port Name	On a UNIX system, this is the device path of the session. On a Windows platform, this is the user type and UniVerse user number, for example, console: 124.
Last Logged In	The date and time the user logged on.
Last Command	The last command the user issued (if known).

Interactive User Display

Background Processes

The following table describes the information that appears in the **Background Processes** area of the **User Administration** dialog box.

Parameter	Description
Pid	The process ID. On a Windows platform, this is the UniVerse user number. It is the same as the value shown in the User # column.
User Name	The user's login name.
User #	The user's UniVerse user number. This column is displayed for Windows platforms only.
Printer Segment	The address of the printer shared memory segment.
Last Command	The last command issued (if known).

Background Processes Display

Terminating a Process

You can terminate a user or a background process from the **User Administration** dialog box.

Terminating a User Process

Complete the following steps to terminate a user process:

1. Choose the user from the **Interactive Users** list.
2. Click **Logout Interactive User**. A message box appears confirming that you want to log out the user.
3. Click **Yes**. UniVerse attempts to log the user off of the server, and updates the **User Administration** dialog box.

Terminating a Background Process

Complete the following steps to terminate a background process:

1. From the **User Administration** dialog box, select the process you want to terminate in the **Background Processes** list.
2. Click **Logout Background Process**. A message box appears confirming that you want to log out the background process.
3. Click **Yes**. The process you chose is immediately terminated, and the **User Administration** dialog box is updated.

Sending Messages with UniVerse Admin

You can send broadcast messages to one or more users from the **User Administration** dialog box. The options available from this dialog box vary depending on whether you are sending a message to users using a UNIX or Windows platform server.

Sending Messages on UNIX Systems

If you are connected to a UNIX server, you can send messages to a single user, all users, or to the system console. Complete the following steps to send a message.

1. From the **User Administration** dialog box, click **Message**. The **Send Messages** dialog box appears, as shown in the following example:

2. Choose the users that you want to receive the message. Valid options are:
   - **All Users** – The message will be sent to all users logged on to the system, not just those logged on to UniVerse.
   - **System Console**
   - **User name or Tty** – Using this option, you can enter the name of a user or tty, or select a user from this list. This list contains all of the users in the **Interactive Users** list.

   Note: *If you select a user before you click* ***Message,*** *the* ***User name or Tty*** *option is automatically set, and the user's name is selected.*
3. Enter your message in the **Message Text** box.
4. Click **OK**. The message appears in the user's current session if the user has messaging enabled.

Sending Message on Windows Platforms

If you are connected to a Windows server, you can send messages to a single user, all UniVerse users, or the UniVerse users in a domain you specify. Complete the following steps to send a message:

1. Click **Message**. The **Send Message** dialog box appears, as shown in the following example:

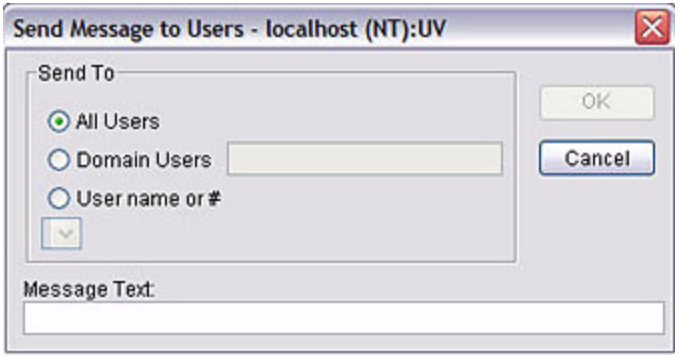

2. Choose the users to which you want to send the message. Valid options are:
   - **All Users** – All the UniVerse users.
   - **Domain Users** – When you choose this option, you must enter the name of a domain.
   - **User Name or #** – When you choose this option, you can enter the name or number of a user, or you can choose a user from the list. This list contains all users in the **Interactive Users** list.

   Note: *If you select a user before you click* ***Message,*** *the* ***User name or #*** *option is automatically set, and the user's name is selected.*
3. Enter your message in the **Message Text** box.
4. Click **OK**. The message appears in the user's current session if the user has messaging enabled.

Chapter

21

Automatic Data Encryption

Using UniAdmin for Encryption 21-2
Encryption Wallet Management 21-7

Using UniAdmin for Encryption

You can use UniAdmin to manage data encryption on your system.

From the UniAdmin main window, select **Data Encrypt Configure**. The **UniVerse Data Encrypt Configure** dialog box appears, as shown in the following example:

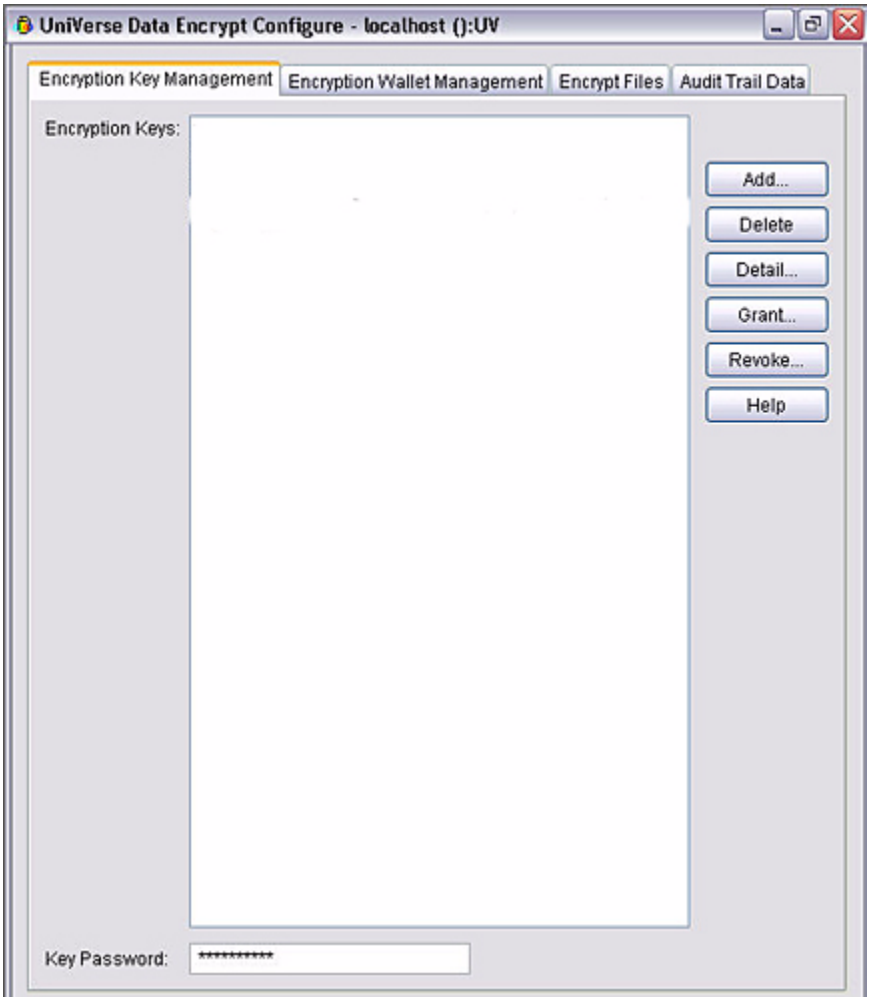

Adding an Encryption Key

To create and encryption key, click **Add**. The **New Encryption Key** dialog box appears, as shown in the following example:

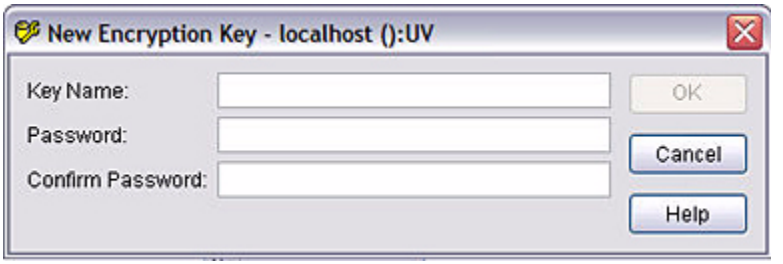

Enter the name of the encryption key in the **Key Name** box. Although not required, you can enter a password for the new key in the **Password** box. Reenter the password in the **Confirm Password** box.

After you create the encryption key, it appears in the **Encryption Keys** area of the **Data Encrypt Configure** dialog box.

Deleting an Encryption Key

To delete an encryption key, from the **Data Encrypt Configure** dialog box, click the encryption key you want to delete, then click **Delete**. The following dialog box appears:

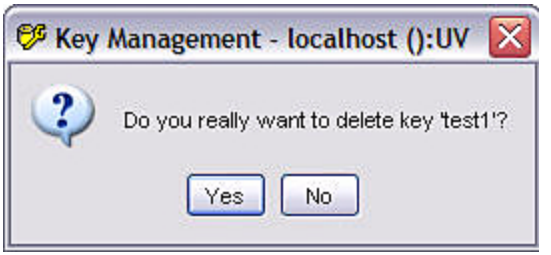

If you want to delete the encryption key, click **Yes**. If not, click **No**.

Viewing Encryption Key Details

To view details about an encryption key, click the encryption key for which you want to view details from the **Data Encrypt Configure** dialog box, then click **Detail**. The **Encryption Key Details** dialog box appears, as shown in the following example:

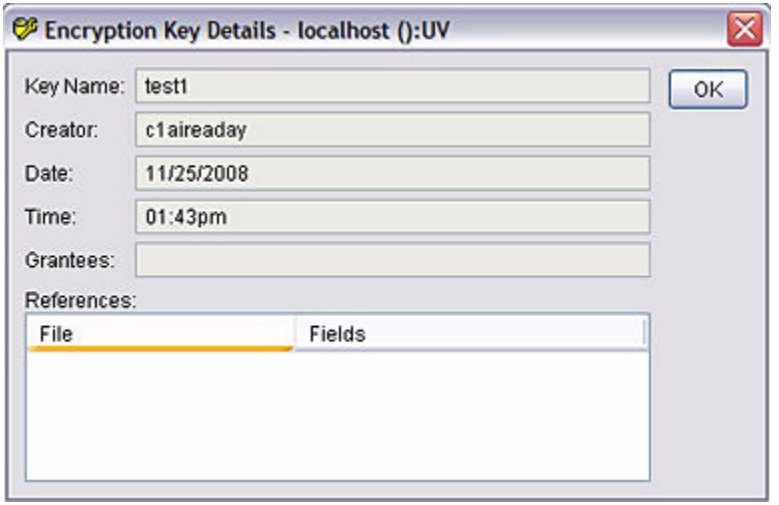

The **Encryption Key Details** dialog box displays the following information about an encryption key:

- **Key Name** –The name of the encryption key.
- **Creator** – The user ID of the user who created the key.
- **Date** – The date the encryption key was created.
- **Time** – The time the encryption key was created.
- **Grantees** – The users or groups who have access to the encryption key.
- **References** – The files and fields for which the encryption key is being used.

If the key is password protected, you must enter the password in the **Key Password** box at the bottom of **Encryption Key Management** dialog box.

Granting Privileges

To grant privileges to the encryption key to a user or group, click **Grant**. The **Grant Encryption Key** dialog box appears, as shown in the following example:

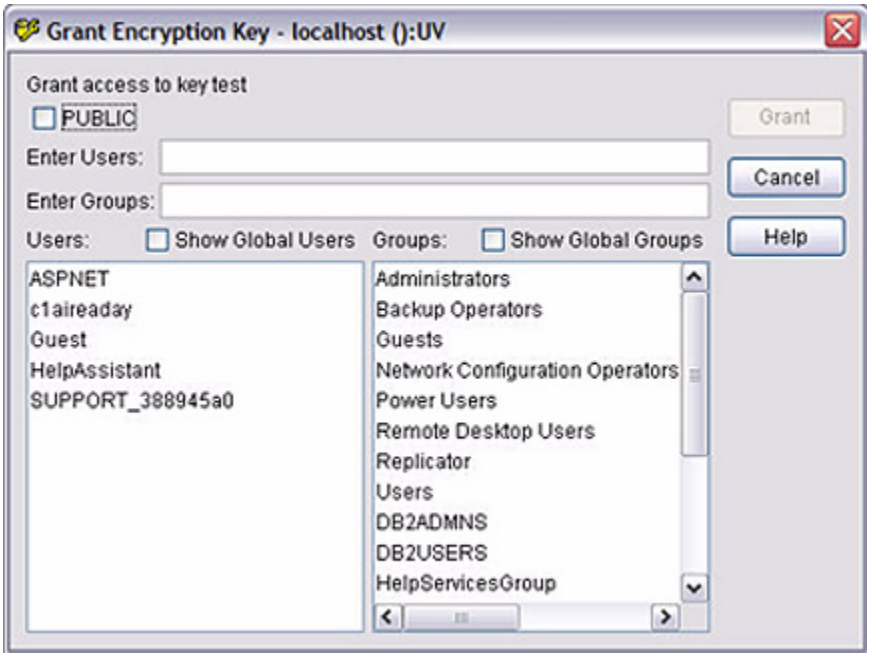

To grant public privileges, click the **PUBLIC** check box.

To grant privileges to individual users, enter the user ID in the **Enter Users** box. Separate user IDs with a comma. You can also select a user from the **Users** box.

To grant privileges to groups, enter the group ID in the **Enter Groups** check box. Separate each group ID with a comma. You can also select a group from the **Group** box.

If you are granting access to keys on a Windows system, the Users/Groups show local users and groups.

To show global users, select the **Show Global Users** check box. Click the user ID of each user for which you want to grant privileges. Alternatively, select the **Show Global Groups** check box and click the group ID for which you want to grant privileges, then click **Grant**. To select multiple users or groups, hold the CTRL key down while selecting the users or groups.

***Note**: You can only grant privileges to Public on Windows platforms.*

Revoking Privileges

To revoke privileges from an encryption key from a user or group, click **Revoke**. The **Revoke Encryption Privilege** dialog box appears, as shown in the following example:

To revoke privileges from PUBLIC users, click the **PUBLIC** check box.

To revoke privileges from individual users or groups, click the user ID of each user for which you want to revoke privileges, or click the group ID for which you want to revoke privileges, then click **Revoke**. To select multiple users or groups, hold the CTRL key down while selecting the users or groups.

Encryption Wallet Management

To manage encryption wallets, click the **Encryption Wallet Management** tab. A window similar to the following example appears:

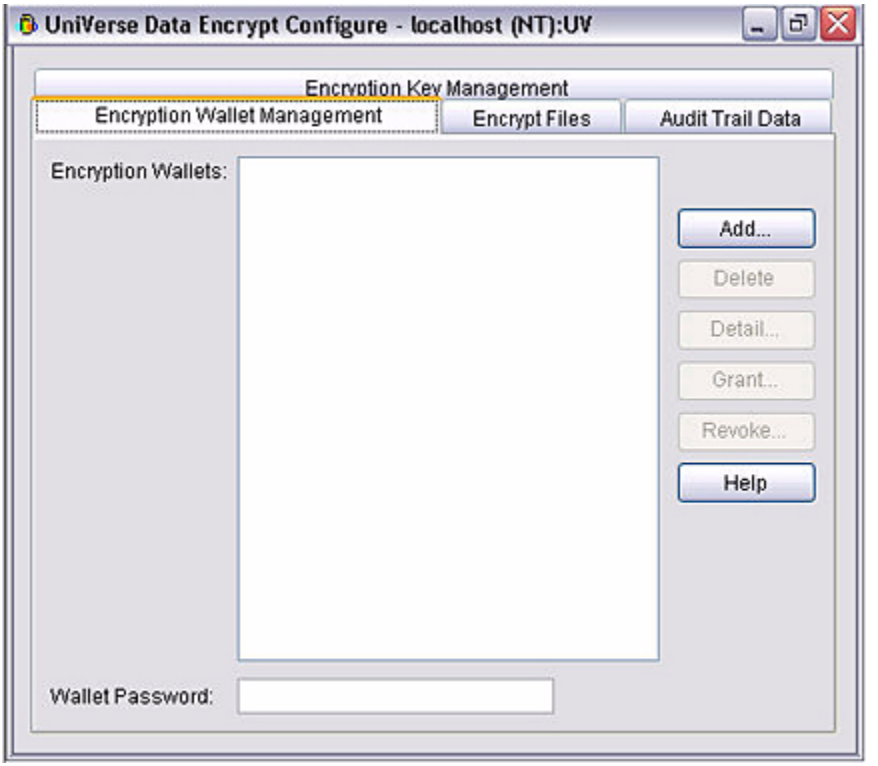

Adding an Encryption Wallet

To add an encryption wallet, click **Add**. The **New Encryption Wallet** dialog box appears, as shown in the following example:

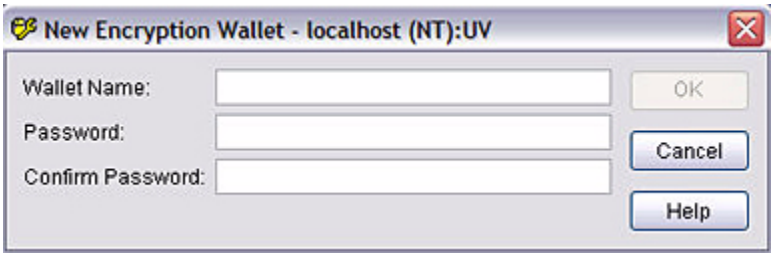

Enter the name for the wallet in the **Wallet Name** box. Enter a password for the wallet in the **Password box**. Reenter the password in the **Confirm Password** box, then click **OK**.

Deleting an Encryption Wallet

To delete an encryption wallet, from the **Encryption Wallet Management** dialog box, click the encryption wallet you want to delete, then click **Delete**. The following dialog box appears:

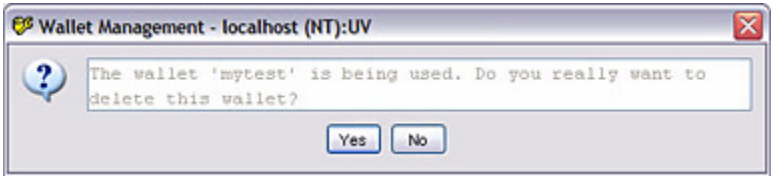

Viewing Encryption Wallet Details

To view details about and encryption wallet, click the encryption wallet for which you want to view details from the **Data Encrypt Configure** dialog box, then click **Detail**. The **Encryption Wallet Details** dialog box appears, as shown in the following example:

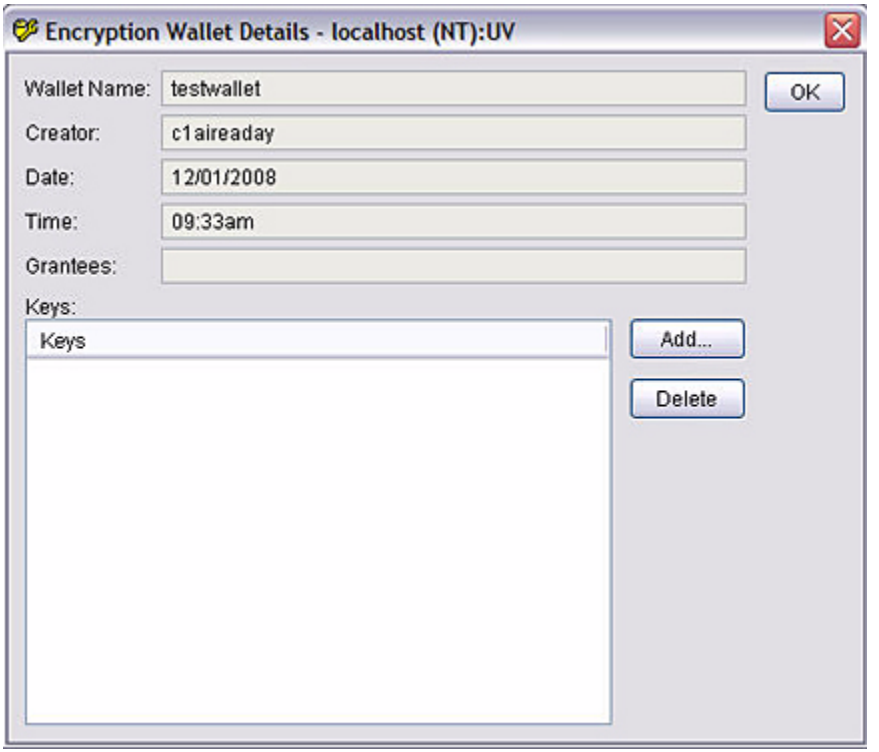

The **Encryption Wallet Details** dialog box displays the following information about an encryption key:

- **Wallet Name** – The name of the encryption key.
- **Creator –** The user ID of the user who created the key.
- **Date** – The date the encryption wallet was created.
- **Time** – The time the encryption wallet was created.
- **Grantees** – The users or groups who have access to the encryption wallet.
- **Keys** – The keys included in the encryption wallet.

If the key is password protected, you must enter the password in the **Key Password** box at the bottom of the **Encryption Key Management** dialog box.

Adding a Key to the Encryption Wallet

To add a key to the encryption wallet, click **Add**. The **Add Encryption Key Into Wallet** dialog box appears, as shown in the following example:

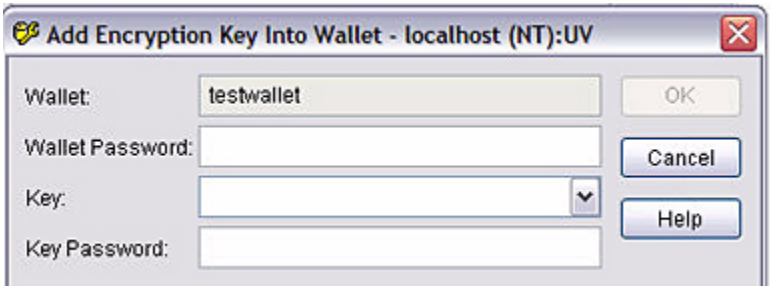

Enter the password for the encryption wallet in the **Wallet Password box**.

Select the key to include in the wallet in the **Key** box.

Enter the password for the key in the **Key Password** box.

Click **OK** to save the encryption key in the wallet, or click **Cancel** to exit without saving changes.

Deleting a Key From an Encryption Wallet

To delete a key from an encryption wallet, from the **Encryption Wallet Management** dialog box, click the encryption wallet where the key you want to delete resides, then click **Detail**. Select the key you want to delete from the wallet, then click **Delete**. The following dialog box appears:

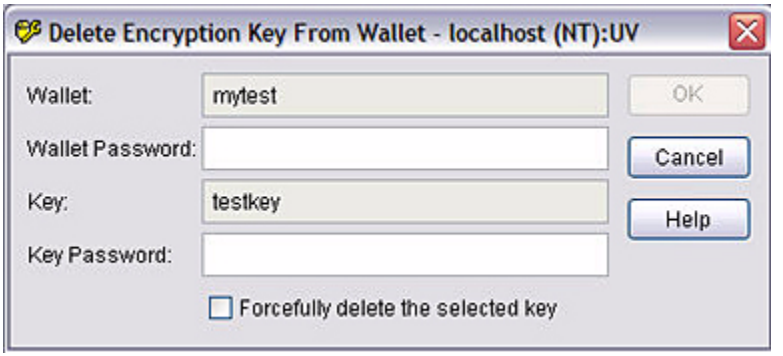

If you select the **Forcefully delete the selected key** check box, a root or uvadmin user can delete an encryption key from an encryption wallet without knowing the key password.

Granting Privileges to Encryption Wallet

To grant privileges to the encryption wallet to a user or group, click **Grant**. The **Grant Encryption Wallet** dialog box appears, as sown in the following example.

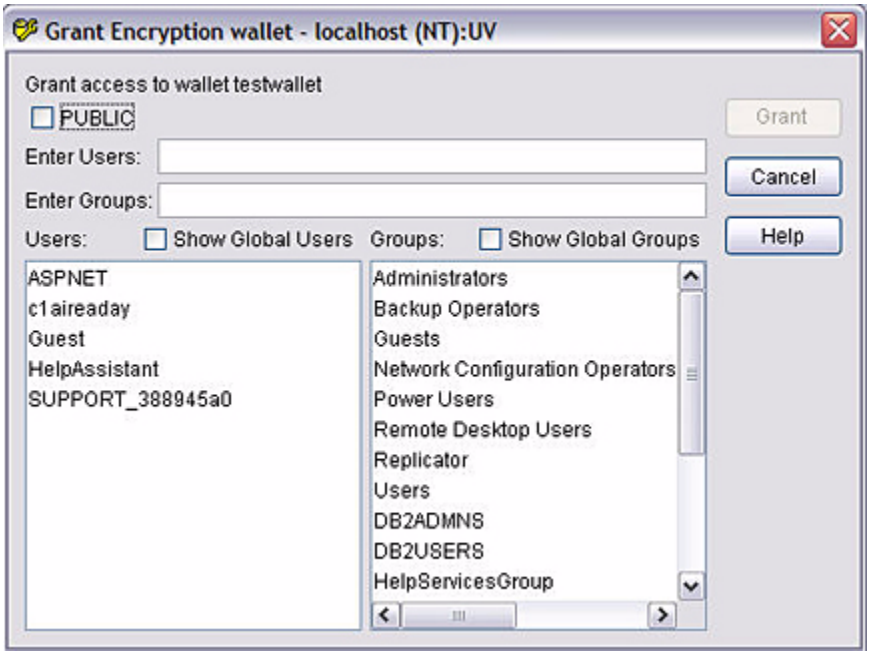

To grant Public privileges, click the **PUBLIC** check box.

To grant privileges to individual users, enter the user ID in the **Enter Users** box. Separate user IDs with a comma. You can also select a user from the **User** box.

To grant privileges to groups, enter the group ID in the **Enter Groups** check box. Separate each group ID with a comma. You can also select a group from the **Group** box.

If you are granting access to keys on a Windows system, the Users/Groups area shows local user and groups.

To show system users, select the **Show Global Users** check box. Click the user ID of each user for which you want to grant privileges. Alternatively, select the **Show Global Groups** check box and click the group ID for which you want to grant privileges, then click **Grant**. To select multiple users or groups, hold the CTRL key down while selecting the users or groups.

Revoking Encryption Wallet Privileges

To revoke privileges from an encryption wallet from a user or group, click **Revoke**. The **Revoke Encryption Wallet** dialog box appears, as shown in the following example:

Waiting for issue 38622 to get fixed for screen shot

To revoke privileges from Public users, click the **Public** check box.

To revoke privileges from individual users, click the user ID of each user for which you want to revoke privileges, or click the group ID for which you want to revoke privileges, then click **Revoke**. To select multiple users or groups, hold the CTRL key down while selecting the users or groups.

Encrypting a File

To encrypt a file or fields in a file, check the **Encrypt File** tab. A window similar to the following example appears:

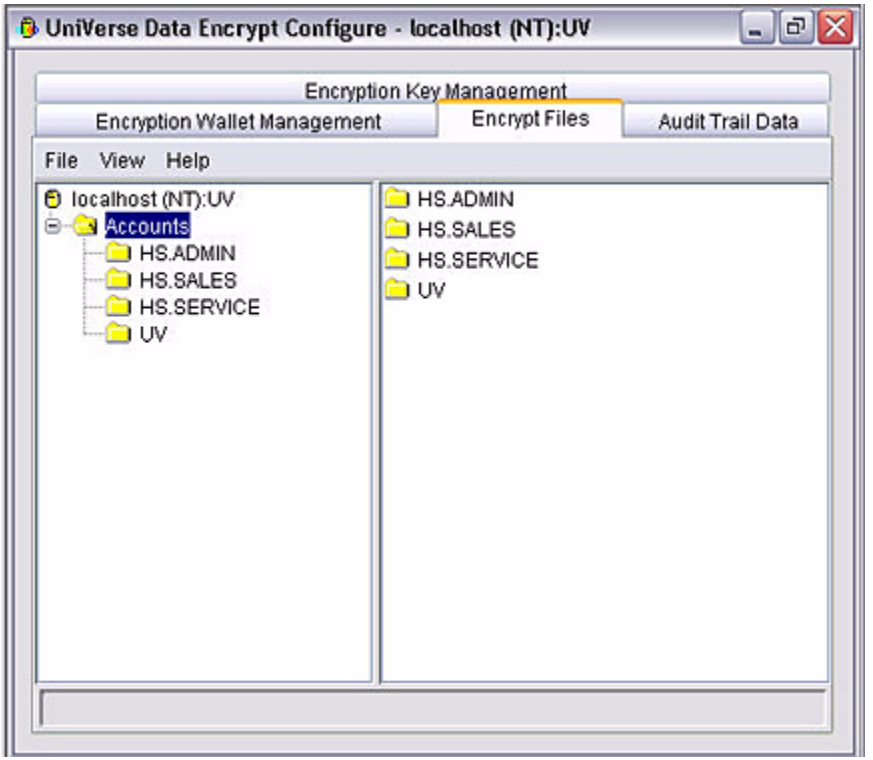

In the **Accounts** area of the screen, click the account where you want to encrypt files. With the right mouse button, click the file in which you want to encrypt files.

The **Encrypt File** dialog box appears, as shown in the following example:

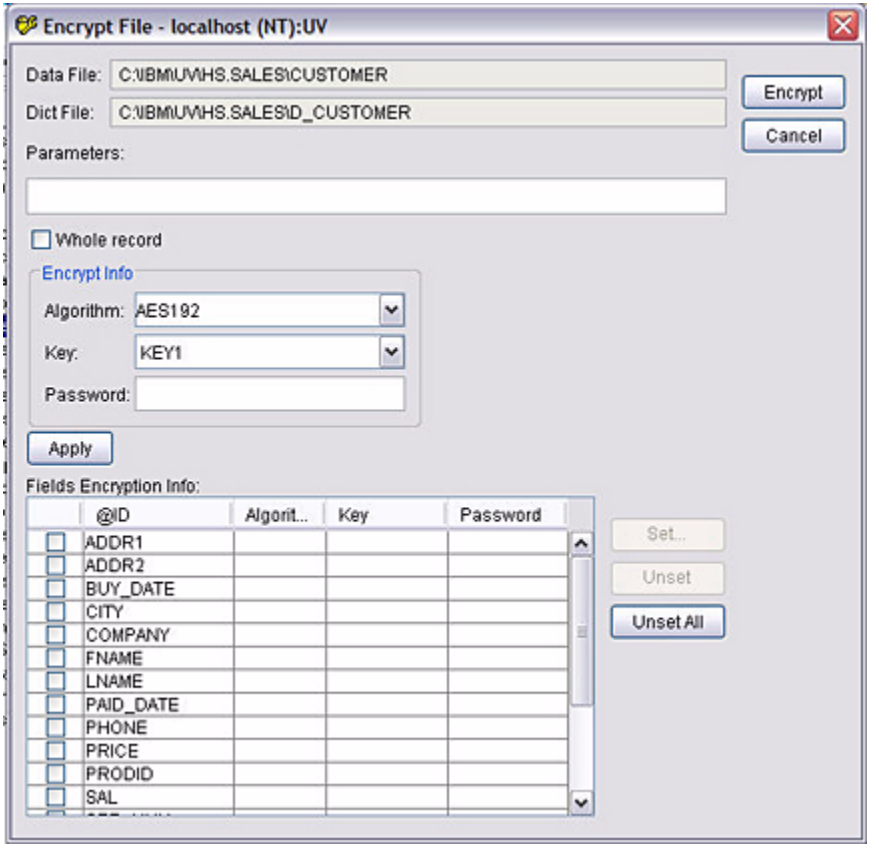

Encrypting an Entire Record

Define the following information to encrypt an entire record:

- **Data File** – The full path to the data file where you want to encrypt data.
- **Dict File** – The full path to the dictionary file where you want to encrypt data.
- **Parameters** – The RESIZE parameters to use when encrypting the file. For a list of valid parameters, see RESIZE in the *UniVerse User Reference*. Leave this box empty if you do not want to resize the file.

- **Whole record** – If you want to encrypt each field in the record, click the **Whole record** check box.
- **Encrypt Info** – Define the following information in the **Encrypt Info** area if you are encrypting and entire record:
  - **Algorithm** – Enter the algorithm to use for encrypting the record. For a list of valid algorithms, see "UniVerse Encryption Algorithms" on page 14.
  - **Key** – Select the encryption key you want to use when encrypting the data from the **Key** list.
  - **Password** – Enter the password corresponding to the encryption key, if one exists.

Click **Apply**. UniVerse encrypts every field for every record in the file.

Encrypting Specific Fields In a Record

Define the following information to encrypt specific fields in a record:

- **Data File** – The full path to the data file where you want to encrypt data.
- **Dict File** – The full path to the dictionary file where you want to encrypt data.
- **Parameters** – The RESIZE parameters to use when encrypting the file. For a list of valid parameters, see RESIZE in the *UniVerse User Reference*. Leave this box empty if you do not want to resize the file.

- **Fields Encryption Info** – Click the name of the field you want to encrypt, then click **Set**. The **Field Encrypt Info** dialog box appears, as shown in the following example:

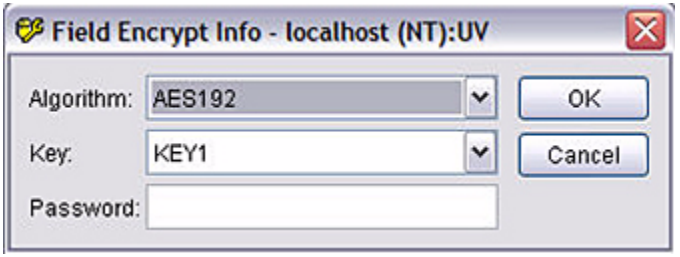

  - **Algorithm** – Select the algorithm to use when encrypting the field. For a list of valid algorithms, see **"UniVerse Encryption Algorithms"** on page 14**.**
  - **Key** – Select the key to use when encrypting the file.
  - **Password** – Enter the password corresponding to the key.

When you have defined all the fields you want to encrypt, click **Encrypt**.

Decrypting a File

To decrypt a file or fields in a file, check the **Encrypt Files** tab. A window similar to the following example appears:

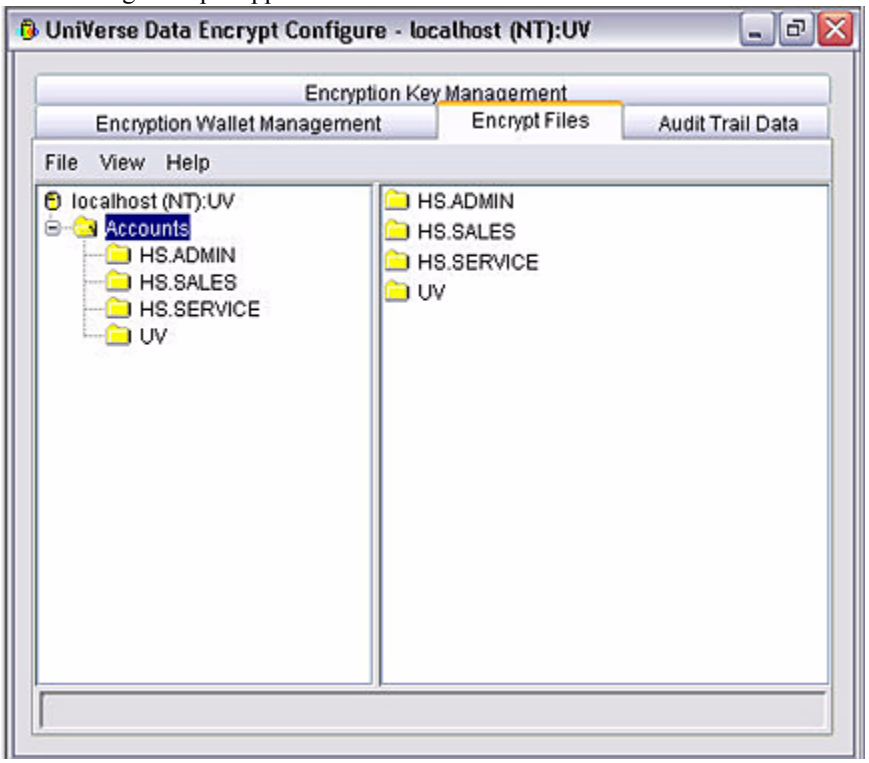

In the **Accounts** area of the screen, click the account where you want to decrypt files. With the right mouse button, click the file in which you want to decrypt fields.

The **Decrypt File** dialog box appears, as shown in the following example:

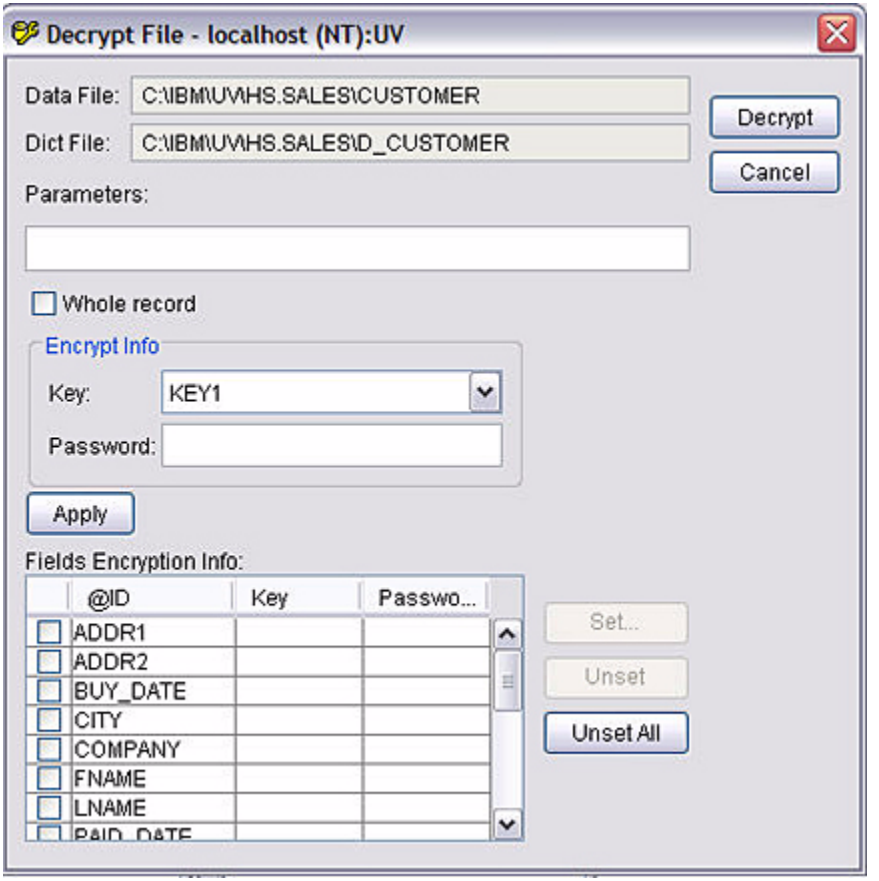

Decrypting an Entire Record

Define the following information to decrypt an entire record:

- **Data File** – The full path to the data file where you want to decrypt data.
- **Dict File** – The full path to the dictionary file where you want to decrypt data.
- **Parameters** – The RESIZE parameters to use when decrypting the file. For a list of valid parameters, see the *UniVerse User Reference*. Leave this box empty if you do not want to resize the file.

- **Whole Record** – If you want to decrypt each field in the record, click the **Whole record** check box.
- **Encrypt Info** – Define the following information in the **Encrypt Info** area if you are encrypting an entire record:
  - **Key** – Select the encryption key you want to use when decrypting the data from the **Key** list.
  - **Password** – Enter the password corresponding to the encryption key, if one exists.

Click **Apply**. UniVerse decrypts every field for every record in the file.

Decrypting Specific Fields in a Record

Define the following information to decrypt specific fields in a record:

- **Data File** – The full path to the data file where you want to decrypt data.
- **Dict File** – The full path to the dictionary file where you want to decrypt data.
- **Parameters** – The RESIZE parameters to use when decrypting the file. For a list of valid parameters, see the *UniVerse User Reference*. Leave this box empty if you do not want to resize the file.
- **Fields Encryption Info** – Click the name of the field you want to decrypt, then click **Set**. The **Field Encrypt Info** dialog box appears, as shown in the following example:

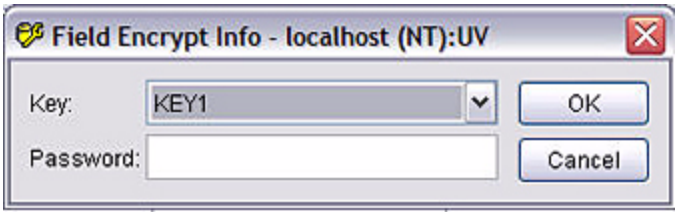

  - **Key** – Select the key to use when decrypting the file.
  - **Password** – Enter the password corresponding to the key.

When you have defined all the fields you want to decrypt, click **Decrypt**.

Listing Encryption Information

To list encryption information for a file, click the **Encrypt Files** tab. A window similar to the following example appears:

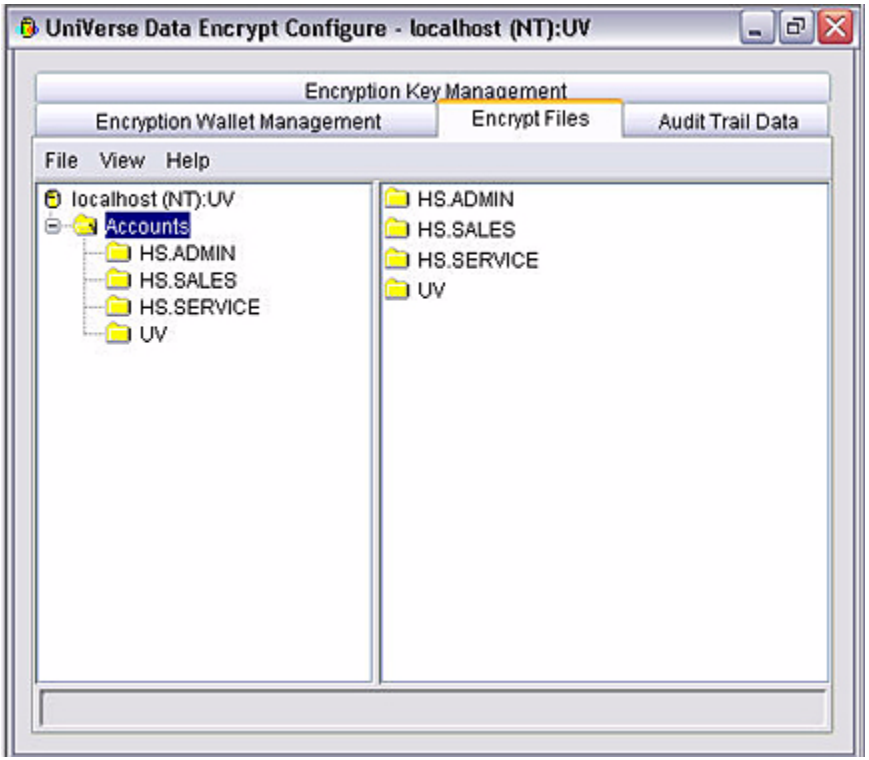

In the **Accounts** area of the screen, click the account where you want to view encryption information. With the right mouse button, click the file for which you want to view encryption information, then click **List Encrypt Info**.

Each field that has been encrypted is listed in the **Encrypted Fields** area of the dialog box. If the entire record is encrypted, Whole record appears under the **Field** column.

Viewing Audit Information

To view audit information, from the **UniVerse Data Encrypt Configure** dialog box, click **Audit Trail Data**. A dialog box similar to the following example appears:

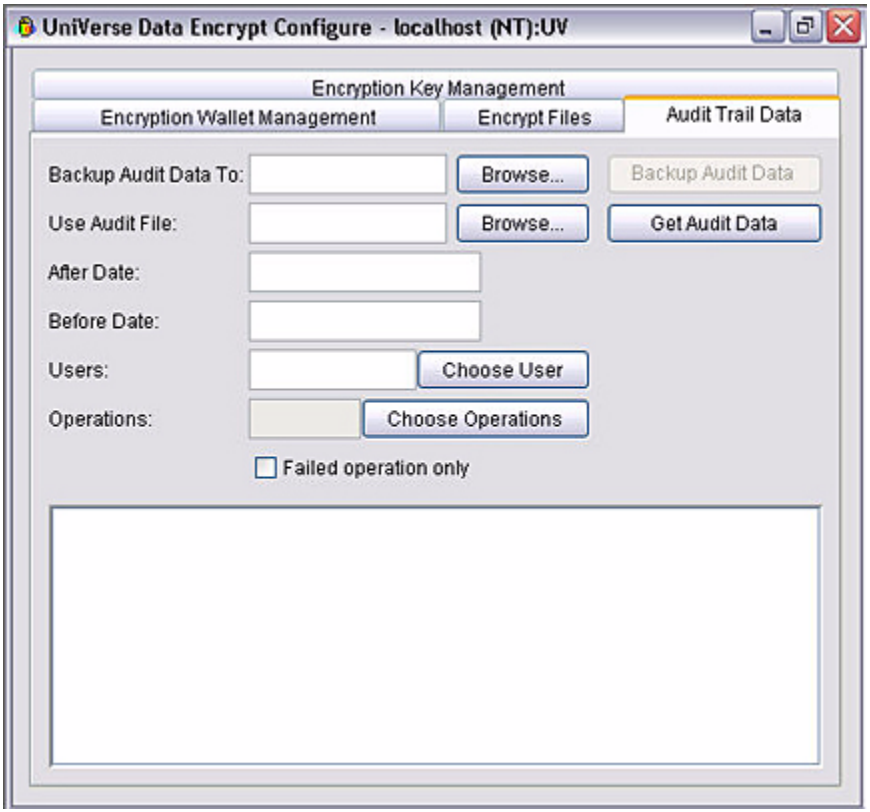

The **Audit Trail Data** dialog box offers the following options:

- **Backup Audit Data To** – If you want to back up the current audit file, enter the patch to the file where you want to back up the file, or click **Browse** to select the file. After backing up the file, the current audit file is cleared.
- **Use Audit File** – If you want to display audit data located in a file different from the current audit file, enter the full patch to the file you want to display, or click **Browse** to select the file.
- **After Date** – UniAdmin will display the audit information after the date you specify. Enter the date in the mm/dd/yyyy format.

- **Before Date** – UniAdmin will display the audit information before the date you specify. Enter the date in the mm/dd/yyyy format.
- **Users** – UniAdmin will display audit trail data for the users you specify. Click **Choose User**. A dialog box similar to the following example appears:

To select multiple users, hold the CTRL key down while selecting the desired users.

- **Operations** – UniAdmin will display audit trail information or the operation you specify. Click **Choose Operations**. The following dialog box appears:

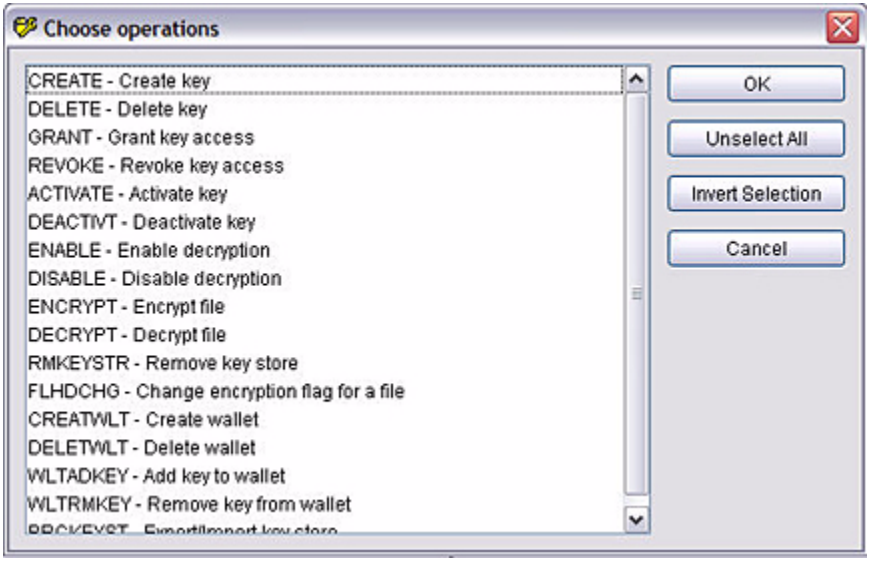

To select multiple operations, hold the CTRL key down while selecting the desired operation. Valid operations are:

- CREATE – Create encryption key
- DELETE – Delete encryption key
- GRANT – Grant key access
- REVOKE – Revoke key access
- ACTIVATE – Activate encryption key
- DEACTIVT – Deactivate encryption key
- ENABLE – Enable encryption key
- DISABLE – Disable encryption key
- ENCRYPT – Encrypt a file
- DECRYPT – Decrypt a file
- RMKEYSTR – Delete a key store
- FLHDCHG – Change encryption flag in file header
- CREATEWLT – Create a wallet

- DELETWLT – Delete a wallet
- WLTADKEY – Add a key to a wallet
- WLTRMKEY – Remove a key from a wallet
- PRCKEYST – Key store import or export

If you only want to display audit trail data for failed operations, select the **Failed operations only** check box.

Click **Get Audit Data**. UniAdmin displays the audit trail data for the criteria you specified, as shown in the following example:

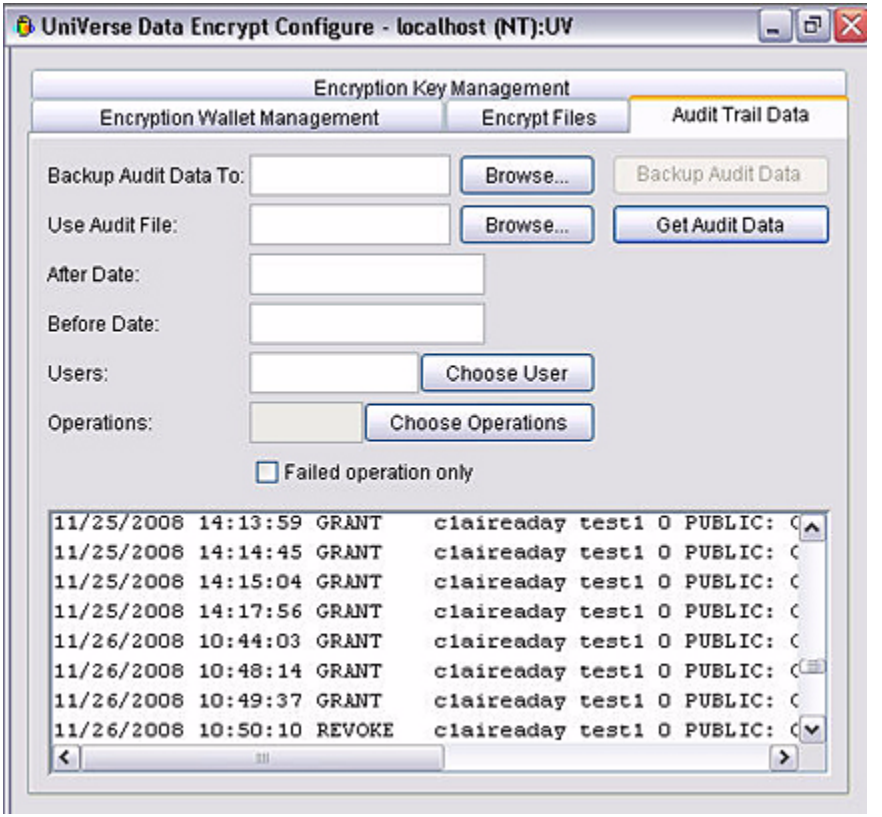

Index

Numerics

64BIT_FILES 4-8

A

A command
  defining action of 4-15
account
  creating 2-3
  deleting 2-6
  viewing 2-5
accounts
  restoring from tape 10-5
Accounts Admin dialog box 2-3
acct.restore executable 10-8
activating
  NLS 4-12
adding
  local machine user 13-15
  new user 13-14
administering
  background processes 21-4
  data sources 5-3
  network services 13-2
  replication 14-2
  servers 1-10
  user processes 21-2
  users 13-13
ALLOWNFS configuration
    parameter 4-8
ALTER.TABLE command 14-38

B

background processes
  administering 21-4
  terminating 21-5
backup
  checking details 15-4
  choosing type of 3-5
  data to disk 3-4
  data to tape 3-4
  defining contents of 3-6
  defining reporting type 3-5
  information about 9-9
  listing index of 15-8
BASIC
  commands 19-5
BASIC program
  defining which to run in shared
      memory 16-5
  loading into shared memory 16-2
  removing from shared memory 16-
    12
  updating in shared memory 16-11
BASIC programs
  loading into shared memory 16-7
BLKMAX configuration
    parameter 4-8
block size
  defining for backup 3-5
  defining for tape device 7-6
  setting log buffer 4-10
  setting maximum 4-8
buffer
  setting initial size of scratch
      buffer 4-16
  setting maximum size of scratch
      pool 4-16
  setting minimum size of scratch
      pool 4-16
  setting size of select list 4-16

C

cache
  setting memory for 4-18
  setting size of 4-9
catalog shared memory
  loading programs into 16-7
  modifying programs in 16-10
  setting up 16-3
  using programs in 16-9
catalog shared memory segment
  removing 16-13
changing
  configuration parameter value 4-6
CLEAR.FILE command 14-38
clearing
  locks 12-5
  publishing information file 14-19
CNAME command 14-38
command history
  using 19-10
command output window 19-12
commands
  BASIC 19-5
  editing 19-10
  executing 19-2
  file & account maintenance 19-6
  managing processes 19-8
  monitoring resources 19-8
  peripheral devices 19-5
  reexecuting 19-10
  saving 19-11
  task and record locking 19-9
commits
  setting maximum number of 4-10
configuration editor
  accessing 4-4
configuration parameter
  changing value of (UNIX) 4-6
  changing value of (Windows platforms) 4-7
configuration parameters
  changing value of 4-6
  defined 4-8
  how UniVerse uses 4-3
  sorting 4-5
configuring
  hot standby subscriber 14-32
  publishing system 14-9
  subscribing system 14-20
configuring transaction logging 18-4
connecting
  to server 1-14
control locks
  number of BASIC user process 4-15
creating
  new data source 5-4
  publishing system for subscriber 14-22
  UniVerse account 2-3
csh shell command
  defining full path for 4-9
CSHDISPATCH configuration parameter 4-9

D

damaged files
  repairing 9-15
data
  backing up 3-3
data replication
  setting mode 4-18
data source
  administering 5-3
  creating new 5-4
  deleting 5-7
  viewing details 5-6
date
  accepting partially bad 4-13
DBsetup program 4-3
deactivating
  NLS 4-12
dead lock manager
  starting 6-4
dead locks
  managing 6-2
  setting checking interval 6-5
  setting resolultion strategy 6-5
  viewing pending 6-7
deadlock manager
  configuring 6-5
debugging
  replication 14-18
defining
  block size for back up 3-5
  contents of backup 3-6
  number of semaphore sets 4-9
deleting
  contents of information file 18-15
  data source 5-7
  restored log files 18-22
  UniVerse account 2-6
devices
  configuring other 7-23
  setting default map for 4-11
diagnostic level
  setting 9-14
disabling
  transaction logging 18-14
disconnecting
  from server 1-15
disk space and replication 14-39
disk usage
  displaying 8-2
display options
  for UniVerse file tool 9-19
domain user
  adding 13-14
DOS command
  defining path for 4-9
DOSDISPATCH configuration parameter 4-9
DROP.TABLE command 14-38
dynamic files
  setting number opened 4-17

E

editing
  commands 19-10
  server information 1-13
enabling
  transaction logging 18-13
errlog
  setting maximum entries 4-10
executing
  UniVerse commands 19-2

F

fail-over mode 14-32
  turning on 14-33
failover-mode
  activating 14-33
file

locks 12-4
SHM.TO.LOAD 16-5
file and account maintenance commands 19-6
file diagnostics
running 9-12
viewing errors 9-15
file lock
defining number of in semaphore set 4-9
file pointers
modifying in VOC file 4-11
file size
setting maximum 4-18
file statistics
viewing information 9-10
files
activating for transaction logging 18-6
administering 9-2
base information 9-5
deactivating for transaction logging 18-7
header information 9-5
listing 9-4
modifying published 14-16
NLS information 9-6
opening remote without UV/NET 4-8
properties 9-4
recovering 18-18
repairing damaged 9-15
restoring 15-3
subscribing 14-24, 14-26
subscribing multiple 14-26
unpublishing 14-17
unsubscribing 14-29
viewing information about 9-11
viewing published 14-16
FLTABSZ configuration parameter 4-9
FSENUM configuration parameter 4-9
function keys
setting 4-15

G

GLTABSZ configuration parameter 4-9
group lock
setting release time 4-13
group lock entries
defining number in semaphore set 4-9
group lock semaphore sets
setting 4-9
group locks
administering 12-7
setting number in semaphore set 4-16
GSEMNUM configuration parameter 4-9

H

hashed files
defining NLS map for 4-12
setting default map for 4-11
HISTSTK configuration parameter 4-9
hosts file
maintaining 13-6
hot backup
reconciling with publisher 14-34
hot backup subsriber
preparing to be reconciled 14-34
hot standby
configuring subscriber 14-32
defined 14-4
fail-over mode 14-32
reconciling with publisher 14-34
turning on fail-over mode 14-33
hot standby operations
managing 14-32
hsrexec
starting 13-17

I

importing
Pick accounts 10-2
Prime INFORMATION accounts 10-2
IN5000 tapes
restoring 10-13
IN8000 tapes
restoring 10-12
indexes
information about 9-8
information file
deleting contents of 18-15
viewing 18-15
installing
UniAdmin 1-3
interactive users 21-3
ISOMODE configuration parameter 4-9

J

JOINBUF configuration parameter 4-9
justification
setting Pick-style 4-14

K

keep alive parameters 13-10

L

LAYERSEL configuration parameter 4-9
licenses
cleaning up 11-5
licensing
authorization code 11-3
viewing license information 11-4
licensing UniVerse 11-2
locale
enabling 4-12
setting default 4-12
locks
administering 12-2
clearing 12-5
file and record 12-4
group 12-7
log buffer
setting block size 4-10
setting size of 4-10
log file sync
setting time interval 4-10
log files

adding 18-9
deleting 18-10
deleting restored 18-22
purging 18-11
releasing 18-12
restoring 18-16
transferring 18-11
LOGBLNUM configuration parameter 4-10
LOGBLSZ configuration parameter 4-10
logging level
setting for telnet session 13-12
logon banner
specifying 13-12
LOGSYCNT configuration parameter 4-10
LOGSYINT configuration parameter 4-10

M

magrst executable 10-8
managing
dead locks 6-2
hot standby operations 14-32
publishing system 14-7
server connections 1-14
transaction logging 18-9
masked decimal
setting for empty data 4-14
MAXERRLOGENT configuration parameter 4-10
MAXRLOCK configuration parameter 4-10
memory
allocation for query 4-13
messages
sending 21-2
sending on UNIX 21-7
sending on Windows platforms 21-8
waiting for cataloged shared memory 16-8
MFILES configuration parameter 4-10
MODFPTRS configuration parameter 4-11
modifying
published files 14-16
modifying programs
in catalog shared memory 16-10

N

nesting
setting depth of 4-15
NETTIME configuration parameter 4-11
network services
administering 13-2
new account
creating 2-3
NLS
activating 4-12
and READ statement 4-13
default map for GCI 4-11
defining map for type 1 & type 19 files 4-12
enabling locale 4-12
map for tape device(UNIX) 7-5
map for type 1 & 19 files 4-11
setting default locale 4-12
setting default map for client programs 4-12
setting default map for devices 4-11
setting default map for hashed files 4-11
setting default map for printer 4-11
setting default map for sequential input 4-11
setting default map for terminal processing 4-12
setting OS map 4-13
specifying map for hashed files 4-12
NLS map
defining for tape device(Windows) 7-11
NLSDEFDEVMAP configuration parameter 4-11
NLSDEFDIRMAP configuration parameter 4-11
NLSDEFFILEMAP configuration parameter 4-11
NLSDEFGCIMAP configuration parameter 4-11
NLSDEFPTRMAP configuration parameter 4-11
NLSDEFSEQMAP configuration parameter 4-11
NLSDEFSRVLC configuration parameter 4-12
NLSDEFSRVMAP configuration parameter 4-12
NLSDEFTERMMAP configuration parameter 4-12
NLSDEFUSRLC configuration parameter 4-12
NLSLCMODE configuration parameter 4-12
NLSMODE configuration parameter 4-12
NLSNEWDIRMAP configuration parameter 4-12
NLSNEWFILEMAP configuration parameter 4-12
NLSOSMAP configuration parameter 4-13
NLSREADELSE configuration parameter 4-13
NLSWRITEELSE configuration parameter 4-13
node
adding 13-6
modifying 13-7
removing 13-7

O

OCVDATE configuration parameter 4-13
OPEN statement
modifying behavior 4-13
OPENCHK configuration parameter 4-13
OPTMEM configuration parameter 4-13

P

PAKTIME configuration parameter 4-13
path
defining for csh shell command 4-9
defining for DOS command 4-9

for sh shell command 4-16
peripheral devices
UniVerse commands for 19-5
Pick
setting Pick-style justification 4-14
Pick account
importing to UniVerse 10-2
PICKNULL configuration parameter 4-14
PKRJUST configuration parameter 4-14
Prime INFORMATION account
importing to UniVerse 10-2
print device(UNIX)
delays 7-18
deleting definition 7-20
flags 7-17
flow control 7-18
NLS map for 7-18
options 7-19
parity 7-19
setting baud rate 7-16
setting form name 7-17
specifying path 7-17
viewing definition 7-19
print device(Windows)
administering 7-22
print devices(UNIX)
defining new 7-15
printer
setting default NLS map 4-11
private transaction cache
setting memory for 4-18
procedure
installation 1-3
processes
managing 19-8
PROCPRMT configuration parameter 4-15
PROCRCMD configuration parameter 4-15
PROMPT keyword
defining effect of 4-15
ProVerb A command
defining action of 4-15
Proverb prompt
PROMPT keyword 4-15
Proverb RI command
defining behavior of 4-15
PSEMNUM configuration parameter 4-15
publisher
defined 14-4, 14-6
reconciling hot standby with 14-34
publishing
multiple files 14-14
one UniVerse file 14-12
stopping on original publisher 14-34
publishing information file 14-18
publishing system
configuring 14-9
creating for subscriber 14-22
managing 14-7
purging
log files 18-11

Q

QBREAK configuration parameter 4-15
QDEPTH configuration parameter 4-15
Q-pointer
setting nesting level 4-15
QSBRNCH configuration parameter 4-15
QSDEPTH configuration parameter 4-15
QSMXKEY configuration parameter 4-16
QSRUNSZ configuration parameter 4-16
query statement
allocating memory for 4-13

R

READ statement
and NLS 4-13
reconciling
hot backup subscriber 14-34
hot standby with publisher 14-34
record
locks 12-4
record locks
setting number in group semaphore set 4-16
setting number in SQL transaction 4-10
recoverable files
specifying 18-6
reexecuting commands 19-10
registering
servers 1-10
releasing
log files 18-12
remote files
opening without UV/Net 4-8
removing
catalog shared memory segment 16-13
obsolete replication log files 14-40
replicated data
setting buffer size 4-18
replication
activating fail-over mode 14-33
administering 14-2
and ALTER.TABLE 14-38
and CLEAR.FILE 14-38
and CNAME 14-38
configuring 14-6
configuring hot standby subscriber 14-32
configuring publishing system 14-9
debugging 14-18
defined 14-4
and disk space 14-39
and DROP.TABLE 14-38
failure 14-41
managing subscribing system 14-20
modifying subscribing files 14-28
publisher defined 14-6
publishing information file 14-18
publishing multiple files 14-14
publishing one file 14-12
reconciling hot backup with publisher 14-34
removing obsolete log files 14-40
and RESIZE 14-38
restrictions 14-38
setting mode 4-18
setting up 14-5

stopping and starting subscribing 14-22
subscriber defined 14-6
subscribing files 14-24, 14-26
subscribing information file 14-30
and triggers 14-38
unpublishing files 14-17
unsubscribing files 14-29
viewing subscribing files 14-28
when disk fills up 14-39
reporting type
for backup 3-5
RESIZE command 14-38
resources
monitoring commands 19-8
restore device
choosing 15-3
restoring
account to UNIX server 10-5
account to Windows platforms 10-8
files
choosing what to restore 15-6
reporting options 15-8
viewing the backup index 15-8
IN5000 tapes 10-13
IN8000 tapes 10-12
log files 18-16
restoring data
editing selection list 15-7
options 15-6
restoring files
specifying method 15-8
restoring files from backup 15-3
RI command
defining behavior of 4-15
RLOWNER configuration parameter 4-16
RLTABSZ configuration parameter 4-16
rotating file pool
setting size of 4-10
rpc port number
changing 13-5
specifying 13-4

S

saving
commands 19-11
scratch buffer pool
setting initial size of 4-16
setting maximum size 4-16
setting minimum size of 4-16
SCRMAX configuration parameter 4-16
SCRMIN configuration parameter 4-16
SCRSIZE configuration parameter 4-16
SELBUF configuration parameter 4-16
select list
in EXECUTE levels 4-9
setting buffer size 4-16
semaphore set
defining lock entries in 4-9
semaphore sets
defining number of 4-9
setting number of 4-9
sentences
maximum in sentence stack 4-9
sequential input
setting default map for 4-11
server
connecting to 1-14
disconnecting from 1-15
editing information 1-13
servers
administering 1-10
registering 1-10
unregistering 1-12
services
starting 13-17
setting
maximum file size 4-18
temp directory 4-19
transaction mode 4-18
sh shell command
defining path of 4-16
shared memory
adding program to 16-10
loading BASIC programs into 16-2
removing program from 16-12
updating program in 16-11
SHDISPATCH configuration parameter 4-16
SHM.TO.LOAD file
adding programs to 16-6
defined 16-5
removing programs from 16-6
sorting
configuration parameters 4-5
setting initial sorting run 4-16
setting maximum depth of tree 4-15
setting maximum key components 4-16
setting sub-merge level 4-15
spooler directory
specifying 4-19
SQL isolation level
setting 4-9
SQL transaction
setting number of record locks 4-10
subscriber
configuring hot standby 14-32
defined 14-4, 14-6
subscribing
multiple files 14-26
one file 14-24
subscribing files 14-24, 14-26
modifying 14-28
viewing 14-28
subscribing information file 14-30, 14-36
clearing 14-31
viewing 14-31
subscribing system
managing 14-20
starting and stopping 14-11
stopping and starting 14-22
suspending
transaction logging 18-14
sync command 14-36
SYNCALOC configuration parameter 4-17

T

T30FILE configuration parameter 4-17
tape
releasing 18-12
tape device(UNIX)
add to rotating file pool flag 7-6

close on read flag 7-6
defining block size 7-6
defining new 7-3
defining UNIX commands 7-7
deleting 7-8
lock file 7-5
multiple read at EOF flag 7-6
types 7-4
using test tape 7-7
viewing definition 7-7
tape device(Windows)
close on read flag 7-12
defining block size 7-12
defining multiple read at EOF flag 7-12
defining new 7-10
deleting definition 7-13
specifying skip and rewind commands 7-12
using test tape 7-14
viewing definition 7-13
tape devices
administering 7-2
tape devices(Windows)
selecting type 7-11
tapein executable 10-10
task and record locking commands 19-9
telnet connection
setting parameters 13-11
telnet port number
displaying 13-10
telnet session
changing port number 13-10
connection parameters 13-10
defining user policies 13-10
keep alive parameters 13-10
managing 13-9
modifying parameters 13-10
user policy 13-10
temp directory
setting 4-19
terminal processing
setting default map for 4-12
terminating
background process 21-5
user process 21-5
THDR512 configuration parameter 4-18
timeout value
setting for UV/NET 4-11
transaction logging
adding log files 18-9
completing roll-forward 18-19
configuring 18-4
deactivating files 18-7
defined 18-2
deleting restored log files 18-22
disabling 18-14
enabling 18-13
information 9-7
managing 18-9
purging log files 18-11
recovering files 18-18
releasing log file 18-12
releasing tape 18-12
restoring log files 18-16
specifying recoverable files 18-6
suspending 18-14
viewing information file 18-15
transaction mode
setting 4-18
transaction processing
activating files 18-6
deleting log files 18-10
transferring log files 18-11
triggers 14-38
TXMEM configuration parameter 4-18
TXMODE configuration parameter 4-18

U

UDRBLKS configuration parameter 4-18
UDRMODE configuration parameter 4-18
UDRMODE parameter 14-34, 14-36
ULIMIT configuration parameter 4-18
UniAdmin
destination folder 1-5
installing 1-3
unirpc
administering on UNIX 13-3
starting 13-17
unirpc daemon
starting on UNIX 13-5
stopping on UNIX 13-6
UniVerse
importing accounts to 10-3
licensing 11-2
starting 13-17
UniVerse account
deleting 2-6
viewing 2-5
UniVerse Command Output window 15-8
UniVerse commands
BASIC 19-5
editing 19-10
file and account maintenance 19-6
for peripheral devices 19-5
managing processes 19-8
monitoring resources 19-8
reexecuting 19-10
saving 19-11
task and record locking 19-9
UNIX commands
defining for tape device 7-7
unpublishing files 14-17
unregistering
server 1-12
unsubscribing files 14-29
user
adding 13-14
adding domain 13-14
adding local machine 13-15
user policies
defining for telnet session 13-10
user process
terminating 21-5
users
administering 13-13
interactive 21-3
uv.rc script 4-3
UV/NET
setting timeout value 4-11
uvbackup
setting maximum block size 4-8
uvconfig file 4-3
uvmt executable 10-11
uvregen command 4-3
uvrestore
setting maximum block size 4-8

uvrestore window 15-5
UVSPOOL configuration
parameter 4-19
UVSYNC configuration parameter 4-19
uvtelnet
starting 13-17
UVTEMP configuration parameter 4-19

V

VDIVDEF configuration
parameter 4-19
viewing
publishing information file 14-19
UniVerse account 2-5
VOC file
modifying file pointers 4-11

W

WIDE0 configuration parameter 4-19
Windows
managing telnet session 13-9