

IBM DB2 Recovery Expert for Linux, UNIX, and
Windows



User's Guide

Version 4 Release 1

IBM DB2 Recovery Expert for Linux, UNIX, and
Windows



User's Guide

Version 4 Release 1

Note:

Before using this information and the product it supports, read the information in "Notices" on page 257.

This edition applies to Version 4 Release 1 Fix Pack 1 of IBM DB2 Recovery Expert for Linux, UNIX, and Windows for (product number 5724-H93) and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC19-3992-01.

© **Copyright IBM Corporation 2001, 2015.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

© Rocket Software, Inc. 2001, 2015

Contents

Figures v

Tables vii

About this information ix

Service updates and support information ix
Highlighting conventions ix
How to look up message explanations ix
How to send your comments x

Chapter 1. DB2 Recovery Expert overview 1

What's new in DB2 Recovery Expert for Linux, UNIX, and Windows 1
DB2 Recovery Expert features and benefits 3
 Log Analysis features and benefits 4
 Recovery features and benefits 4
 Object History features and benefits 5
 Full support for Database Partitioning Facility (DPF) environments 5
 Multiple object recovery support. 5
Terminology 5
DB2 Recovery Expert architecture and process flow 9
DB2 Recovery Expert interaction with DB2 and host 12
DB2 Recovery Expert documentation and updates 14
Accessibility features 14

Chapter 2. Preparing to install and configure DB2 Recovery Expert 17

Hardware and software prerequisites 17
Installation requirements and recommendations 22
Required privileges and authorities 26
Operational prerequisites 28
Trial and permanent licenses 39
DB2 Recovery Expert upgrade 39
Configuring Linux and UNIX systems to run the DB2 Recovery Expert Server. 40
 Changing resource limits 41

Chapter 3. Installing DB2 Recovery Expert 43

Installation roadmap 43
Installing or upgrading DB2 Recovery Expert 45
Installing or upgrading DB2 Recovery Expert using the non-interactive (silent) mode 48
Installing or upgrading the Command Line Processor (CLP) client by using the installation wizard or the console mode 50
Installing a permanent license with IBM DB2 Recovery Expert Activation Kit. 52
Backing up product application files 53
Configuring remote access to database server machines 53

Host name resolution and multiple IP addresses for a host name 53
Windows targets. 54
 Configuring specific versions of Windows 55
Uninstalling DB2 Recovery Expert. 57
Uninstalling the Command Line Processor (CLP) client by using the uninstallation wizard or the console mode. 58

Chapter 4. Configuring DB2 Recovery Expert 61

Administering the DB2 Recovery Expert Server 62
 Starting and stopping the DB2 Recovery Expert Server 62
 Changing the DB2 Recovery Expert Server settings and displaying server status 63
Configuring the DB2 Recovery Expert repository database 64
Configuring DB2 Recovery Expert security 65
Configuring DB2 Recovery Expert Services 66
Managing databases used with DB2 Recovery Expert 66
 Adding a managed database 66
 Editing managed database properties. 70
 Deleting a managed database from a datastore 71
Managing DB2 Recovery Expert datastores 71
 Creating a datastore by using the browser client 72
 Creating a datastore by using the command line tools. 72
 Editing datastore properties 73
 Changing database authorities of an existing datastore connection 74
 Deleting a datastore 75
 Datastore basic information and properties 75
Viewing and upgrading deployed database server components 75
 Installing or upgrading database server components from the command line 76
 Deployment details window. 77
Configuring DB2 Recovery Expert for Log Analysis and remote database processing 78
Modifying general settings 79
 Administration settings 79
Customizing view preferences 80

Chapter 5. DB2 Recovery Expert browser client overview 83

DB2 Recovery Expert workflow 83
 Database location 84
 Datastore repository database 84
 Schema Level Repository (SLR). 85
 Recovery plans 85
Main browser client view. 85
 Standard action buttons 87
 New Tab button 87

Sessions and specifications	88
Sessions	88
Specifications	89
Saving sessions or specifications	89
Searching for saved sessions or specifications	90
Tips for using the browser client	91

Chapter 6. Object History and the SLR 95

Connecting to a database to run the Object History tool	96
Reviewing SLR Information	96
Specifying additional locations to search for log files and backups	97
Performing SLR operations	98
Creating the SLR for a database	98
Rebuilding the SLR	100
Updating the SLR	101
Pruning SLR information	102
Dropping the SLR	102
Reviewing the SLR status	102
Updating the SLR from logs	103

Chapter 7. Log Analysis 105

Connecting to a database to run Log Analysis	106
Providing report details	106
Selecting objects for Log Analysis reporting	108
Providing filters for Log Analysis	111
Specifying additional options for Log Analysis processing	112
Reviewing Log Analysis status	113
Working with Log Analysis results	114
Creating column-level filters	118
Log Analysis Result step field definitions	119

Chapter 8. Recovery 123

Selecting a database for recovering objects	124
Selecting objects for recovery	124
Selecting a point in time for recovery	126
Specifying additional recovery options	126
Viewing dependent objects and specifying related tables for recovery	127
Selecting a recovery plan	129
Recovery plan definitions	131
Reviewing recovery status	137

Chapter 9. Command line processor (CLP) interface 139

Java Command Line Processor (CLP) overview	139
--	-----

Configuring the CLP	140
Command descriptions and usage details	141
Using the CLP to run Log Analysis against remote databases	160

Chapter 10. Considerations and restrictions 161

Operational limitations	161
Recovery, Log Analysis, and SQL limitations and considerations	163
SLR considerations	167
Limited or unsupported DB2 features, environments, and operations	168
Unsupported DDL statements	170
Considerations for DPF environments	171
HADR considerations	172
IDENTITY column considerations	173
Index values limitations	173
MDC table restrictions	173
Limitations for tables with LOB and LONG columns and extended row size	175
Materialized query table restrictions and considerations	175
Log chain restrictions and considerations	177
SET INTEGRITY statement considerations	178

Chapter 11. Troubleshooting 179

Information required to analyze errors	179
Information required to analyze performance issues	180
Gathering diagnostic information	182
Automatically generating trace files	182
Locating database server components information	182
Generating output from the db2support utility	183
Accessing server log files	183
DB2 Recovery Expert messages	184

Notices 257

Trademarks	259
----------------------	-----

Index 261

Figures

1. DB2 Recovery Expert architecture diagram 1 10
2. DB2 Recovery Expert architecture diagram 2 11

Tables

1. DB2 Recovery Expert interaction with DB2 and host	12	13. Commands for installing a permanent license	52
2. UNIX supported environments	19	14. NetBIOS Reserved Ports	55
3. Linux supported environments	20	15. Uninstalling DB2 Recovery Expert by using the uninstallation wizard or the console mode	57
4. Linux on Inspur K1 platform supported environments	20	16. Uninstalling the Command Line Processor (CLP) client by using the uninstallation wizard or the console mode	58
5. Windows supported environments	20	17. Commands for starting and stopping the server	63
6. DB2 supported versions	20	18. Keyboard shortcut keys	91
7. Language codes	26	19. Schema and object filtering include examples	109
8. Installation program start commands	46	20. Schema and object filtering exclude examples	110
9. Set the installation directory location	48	21. Alphabetical list of Log Analysis Result step fields and descriptions	119
10. Generate encrypted passwords during the installation procedure	49	22. Schema and object filtering include examples	126
11. Specify the path to the installation properties response file	49	23. Command summary	142
12. CLP installation program start commands	51		

About this information

IBM® DB2® Recovery Expert for Linux, UNIX, and Windows (also referred to as DB2 Recovery Expert for LUW or DB2 Recovery Expert) provides a simple, self-managing recovery solution that enables enhanced recovery operations with minimal disruption to DB2.

These topics are designed to help database administrators, system administrators and application administrators perform the following tasks:

- Plan for the installation of DB2 Recovery Expert for LUW
- Install and operate DB2 Recovery Expert for LUW
- Customize your DB2 Recovery Expert for LUW environment
- Diagnose and recover from DB2 Recovery Expert for LUW problems

To use these topics, you should have a working knowledge of:

- Linux, UNIX, or Windows operating systems
- DB2 Universal Database™

Specific changes since the previous edition of this book are indicated by a vertical bar (|) to the left of a change. Editorial changes that have no technical significance are not noted.

Always check the DB2 Tools Product page for the most current version of this publication:

<http://www.ibm.com/support/docview.wss?uid=swg27020910#db2ars-lib>

Service updates and support information

To find service updates and support information, including software fix packs, PTFs, Frequently Asked Question (FAQs), technical notes, troubleshooting information, and downloads, refer to the following web page:

http://www-947.ibm.com/support/entry/portal/Overview/Software/Information_Management/Integrated_Data_Management

Highlighting conventions

This information uses the following highlighting conventions:

- **Boldface** type indicates commands or user interface controls such as names of fields, folders, icons, or menu choices.
- Monospace type indicates examples of text that you enter exactly as shown.
- *Italic* type indicates variables that you should replace with a value, the titles of publications, and significant new terms.

How to look up message explanations

You can use any of the following methods to search for messages and codes:

Searching an information center

In the search box that is located in the top left toolbar of any Eclipse help system, such as the IBM DB2 Advanced Recovery Solutions Information Center <http://pic.dhe.ibm.com/infocenter/db2ars/v1r1>, enter the number of the message that you want to locate. For example, you can enter ARY03??E in the search field.

Use the following tips to help you improve your message searches:

- You can search for information on codes by entering the code; for example, enter -327.
- Enter the complete or partial message number. You can use the asterisk wildcard character (*) to represent multiple characters, and you can use the question mark wildcard character (?) to represent a single character.

The information center contains the latest message information for all of the information management products that are included in the information center.

Searching for messages on the web

You can use any of the popular search engines that are available on the web to search for message explanations. When you type the specific message number or code into the search engine, you will be presented with links to the message information in IBM information centers.

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book or any other DB2 Recovery Expert documentation, use either of the following options:

- Use the online reader comment form, which is located at:
www.ibm.com/software/data/rcf/
- Send your comments by email to comments@us.ibm.com. Be sure to include the name of the book, the part number of the book, the version of DB2 Recovery Expert, and, if applicable, the specific location of the text you are commenting on (for example, a page number or table number).

Chapter 1. DB2 Recovery Expert overview

IBM DB2 Recovery Expert for Linux, UNIX, and Windows is a data analysis and recovery tool that accesses DB2 systems through a browser client and a command-line processor interface. DB2 Recovery Expert performs detailed log analysis and recovery operations which helps organizations to protect their valuable data.

DB2 Recovery Expert helps customers protect data by:

- Reducing disaster recovery time and maximizing database availability.
- Eliminating data errors before they compound into costly business mistakes.
- Tracking and reporting data changes in response to auditing requests.

Key product features include:

- Ability to recover database objects to a user specified point-in-time without using resource intensive disaster recovery.
- Determining the most efficient technique for a recovery situation.
- Rolling back unwanted data changes throughout a database while the system remains online.
- Analyzing database logs, providing text reports and generating optional executable Redo or Undo SQL using user-specified filters and report range parameters.

Topics:

- “What’s new in DB2 Recovery Expert for Linux, UNIX, and Windows”
- “DB2 Recovery Expert features and benefits” on page 3
- “Terminology” on page 5
- “DB2 Recovery Expert architecture and process flow” on page 9
- “DB2 Recovery Expert interaction with DB2 and host” on page 12
- “DB2 Recovery Expert documentation and updates” on page 14
- “Accessibility features” on page 14

What’s new in DB2 Recovery Expert for Linux, UNIX, and Windows

This topic summarizes the technical changes for this edition.

New and changed information is indicated by a vertical bar (|) to the left of a change. Editorial changes that have no technical significance are not noted.

Security enhancement

The product now uses three-key triple Data Encryption Standard (3DES) to comply with security requirements from the National Institute of Standards and Technology (NIST).

When the server starts with the new components, it updates table `ARY_DATABASE` to replace passwords for all connections with passwords that use 3DES encryption. As a result, after you upgrade the database server components to use the new encryption, you must also upgrade the backend components for all connections on all instances.

Support for a remote repository for use with target databases

Previous versions of the product required that you create a datastore repository database to store information about the target databases (those with which you use DB2 Recovery Expert). The datastore repository had to be part of the same instance as the target database. This version of the product enables you to use a remote repository; that is, DB2 Recovery Expert can use a datastore repository database that was created in any instance on any host.

The installation process now automatically creates a DB2 database, and creates a default datastore repository and a *configuration repository* in that database. The configuration repository is a database that is used to store the list of datastores and managed databases, and related information that DB2 Recovery Expert requires to work with them. The product uses the default datastore for all managed databases unless the administrator creates a separate datastore. This enhancement enables you to get started with the product more quickly.

You can begin using the remote repository feature for all new managed databases. An automated migration process will integrate your existing datastores, and the existing datastores continue to hold existing content. You can use existing recovery information without performing any modifications.

You must run the upgrade script to upgrade the existing configuration repository in the application server to the new schema. During the upgrade, a migration process migrates existing installation and datastore information to the new schema. The migration process searches existing datastores for databases that have had the SLR built, and registers those databases as managed databases in the current repository.

The following browser client changes are associated with this enhancement:

- The **Admin** perspective has been removed.
- The **Open** menu **Setup** section includes the following new perspectives:
 - Managed Databases
 - Datastores
 - Deployment Status
 - Admin Settings
- The Location step now shows the list of registered databases, rather than all databases for a given instance.
- By default, the product creates the DB2 database and also sets the security configuration to Repository database authentication. The topic *Repositories and security* is therefore obsolete and has been removed from the documentation.

Recovery enhancement

DB2 Recovery Expert can generate DDL to recover an object's privileges.

Log analysis reports enhancement

For Log Analysis processing, the **Log files range** option is now supported for the **Minimum Recovery Timestamp activity** report mode. For more information about specifying log files, see “Specifying log files ranges in MRT mode” on page 163.

Support for DB2 10.5 features

The following features are now supported:

- Extended row size: Log Analysis processing supports tables with extended rows. Recovery processing extracts data for tables with extended rows from a backup and reads data for such tables from the table space.
- Character-based columns: National language string and string unit values are included in the SLR, Log Analysis reports and DDL, and post filtering; Recovery includes table spaces that have column types that contain CODEUNITS16 and CODEUNITS32.

Support for KVM

DB2 Recovery Expert supports the Linux Kernel-based Virtual Machine (KVM) hypervisor.

DB2 Recovery Expert features and benefits

DB2 Recovery Expert offers several unique and significant features that you can use to improve your DB2 environment.

- Support for DB2 versions 9.5, 9.7, 10.1, and 10.5.
- Support for the following operating system platforms:
 - AIX® 6.1 64-bit
 - AIX 7.1 64-bit
 - HP-UX 11iv2 (11.23.0505) and 11iv3 (11.31)
 - Linux on xSeries (RHEL 5 and 6, SLES 10 and 11, Ubuntu 12.01) 32 and 64-bit
 - Solaris 10 on SPARC 64-bit
 - Windows 32 and 64-bit
- Browser client and command-line client are available.
- Distributed remote installation of database server components.
- Log Analysis feature with the ability to filter activity and generate Redo and Undo SQL and reports. Redo SQL can include optional DDL statements.
- Session support, including saving Recovery Scenarios and Log Analysis report options.

Users can save Recovery and Log Analysis object selections and settings. This allows the user to reuse scenarios to generate reports and recoveries.

- Backup and Log access on disk, TSM, and Vendor storage.
- Automatic, DMS, and SMS table space container support.
- Automatic archive log handling.

DB2 Recovery Expert handles recalling archived logs using the log archiving method specified in the Recovery History File (RHF).

- Support for Online and Offline database and table space backups.
- Support for table partitioning.
- DB2 Merge Backup for Linux UNIX and Windows support.

Merged backup images created by the DB2 Merge Backup for Linux UNIX and Windows product can be specified as the source backup image used for creating the SLR or for supplying data during Log Analysis or Recovery operations.

- Support for a single datastore repository database in any instance on any host.

Related concepts:

“Recovery, Log Analysis, and SQL limitations and considerations” on page 163
Review the following restrictions and considerations.

Log Analysis features and benefits

Log Analysis offers several features and benefits that aid in the recovery process.

The features and benefits provided by Log Analysis are as follows:

- For DB2 version 10.5 and later, extended row size tables
- For DB2 version 10.1 and later, you can run local Log Analysis processing in a pureScale® environment.
- Database Partitioning Facility (DPF) support
- Partitioned Table support
- Multi-dimensional clustered table support
- Optional inclusion of LOB/Long field reporting
- Row Compression
- Value Compression
- Summary, Detail, and Quiet Time reporting
- Filter by Application ID, Application Name, Table name and ID, Table space name and ID, Log Sequence Number (LSN), Log Record Identifier (LRI), Log Stream Offset (LSO), Transaction ID, Database Partition Group, Authorization ID, and Schema
- Include, exclude, and wildcard filtering
- Partial, Uncommitted and Rolled back transaction support
- DML and DDL generation, including ALTER TABLE
- Report export in CSV, XML, and text formats
- Currently Committed transaction support
- Timestamp with precision support
- Inline LOB support

Related concepts:

“Recovery, Log Analysis, and SQL limitations and considerations” on page 163
Review the following restrictions and considerations.

Recovery features and benefits

Recovery offers several features and benefits that aid in the recovery process.

The Recovery provides the following features and benefits:

- Database Partition Facility (DPF) support
- Partitioned table support
- Redirected restore
- Recover to point in time (PIT) or end of logs
- Multiple recovery path generation
- Recovery of multiple objects simultaneously
- Automatic recovery of related objects including referential integrity referenced tables, triggers, data types, functions, procedures, table spaces, indexes, views, database partition group, buffer pool, and schema
- Dropped table recovery without table space restore and roll forward
- Dropped table space recovery from offline backup using translated restore
- UNDO SQL recovery from the current time

- Recover renamed objects
- Recover objects with original privileges

Object History features and benefits

Object History offers several features and benefits that aid in the recovery process.

These features and benefits include:

- Generate DDL for supported objects at any point in time.
- Initialization from any point in time using a system catalog backup.
- Continuous Schema Level Repository (SLR) and object history using database logs.
- Create, rebuild, update, prune, and drop operations from the browser client and the command line.
- Allow Log Analysis reporting on changed and dropped tables.

Related concepts:

Chapter 6, “Object History and the SLR,” on page 95

Review Object History and Schema Level Repository (SLR) operations, such as how to view the current SLR range and how to create, rebuild, update, prune, or drop SLR information for user databases.

“Performing SLR operations” on page 98

On the **Operations** step, you can perform operations against a selected database.

Full support for Database Partitioning Facility (DPF) environments

Full support is provided for DPF environments.

The IBM DB2 Recovery Expert for Linux, UNIX, and Windows components perform coordinated recovery and Log Analysis operations against all of the database partitions without requiring user interaction for each partition. The user chooses to perform a Recovery or a Log Analysis operation for all database partitions, and DB2 Recovery Expert automatically initiates actions for each partition and coordinates the results after the various actions complete. This greatly simplifies the DPF object recovery process.

Multiple object recovery support

IBM DB2 Recovery Expert for Linux, UNIX, and Windows enables you to specify the recovery of multiple objects within the same database to the same point in time.

Object selection supports specifying individual objects or the use of pattern matching to select multiple objects at one time. Processing is optimized when recovering multiple objects to reduce recovery time.

Related concepts:

Chapter 8, “Recovery,” on page 123

On the **Recovery** perspective, you can recover data and dropped objects for a database.

Terminology

The following terms are used throughout the DB2 Recovery Expert product interfaces and documentation.

C

Configuration repository

A database that is used to store the list of datastores and managed databases, and related information that is required to work with them. Also known as the *metadata repository*.

D

Datastore repository

A product database that stores schema level tables, Log Analysis metadata and reports, recovery plans, and other data about a set of managed databases.

Database server components

Core executable components that are installed on each DB2 server/partition.

Deployment

An instance of database server components installed on a specific host for a specific instance.

Dropped object restore

You can restore database objects that have been dropped using DB2 Recovery Expert. Unlike the native DB2 method for restoring dropped tables, DB2 Recovery Expert does a selective restore and roll forward or Redo of transactions that relate specifically to the object in question using the DB2 logs.

J

Java Command Line Processor (CLP)

A stand-alone client that allows you to run major feature products from a command shell or in a batch mode environment. This does not require a connection to the DB2 Recovery Expert server. The CLP allows remote processing from any client supporting Java 5.

L

Log Analysis tool

The Log Analysis tool generates reports detailing the changes that have occurred in particular tables or table spaces within specific time periods. You can also generate Redo or Undo SQL statements that can be used to replicate or reverse changes that were previously made to database objects. Redo SQL can include optional DDL statements. DDL is not available with Undo SQL. A tool is provided to apply SQL output that was generated by Log Analysis.

LSN/LRI/LSO

The browser client and documentation use this combination of acronyms to represent the following terms:

- Log sequence number: 16-hexadecimal-character value, as defined in DB2 9.7 and earlier.
- Log record identifier: 34 hexadecimal-character value, as defined in DB2 10.1 and later.
- Log stream offset: 16-hexadecimal-character value, as defined in DB2 10.1 and later.

M

Managed database

A database that is configured for use with DB2 Recovery Expert. For a managed database, you can create a SLR, run Log Analysis processing, and recover objects.

Masked update reconstruction

The Log Analysis tool uses masked update reconstruction to reconstruct update statements. By default, when a row is updated, DB2 does not write a full before or after image of the row to the log. It only logs information about the columns that were changed. This means that the Log Analysis Tool cannot construct full before and after images of the row from only the log record for the update. To construct before and after images of the row, the tool performs one of the following actions:

- Forward update reconstruction, in which the tool starts with the image of the row from a backup image, previous insert log record, or previous currently committed log record and rolls forward through all of the updates to the row in the logs to get to the before image for the update in question. It can then reconstruct the after image by applying the change given in the log record.
- Backward update reconstruction, in which the tool starts with the image of the row that currently exists in the table and rolls back through all of the updates to the row in the logs to get to the after image for the update in question. It can then reconstruct the before image by reversing the change given in the log record.

Metadata repository

See *Configuration repository*.

Minimum Recovery Time (MRT)

A term used by DB2 to describe the timestamp of the last modification of any object stored in a particular table space. This timestamp is the earliest point in time to which the table space can be recovered and still match the current DB2 System Catalog.

DB2 Recovery Expert uses the MRT timestamp as a reference to determine when processing without the SLR is valid using only the System Catalog. To process further back in time beyond the MRT, it is necessary to create and update the SLR so that DB2 Recovery Expert has information about the database objects over the necessary time range. In cases where a set of table spaces or the whole database is processed, DB2 Recovery Expert uses the most recent MRT value from the set of table spaces that are being processed as the limit to guarantee that information is consistent for activity across the whole set of table spaces. If a report is requested for a time that exceeds the MRT for any processed table space, Log Analysis will report an error. In this case, a different report timeframe must be specified, or the SLR must be created or updated to cover the requested range of the report.

Multiple object recovery

DB2 Recovery Expert supports the recovery of more than one object at a time, if the objects are of the same type. If several related objects must be recovered to a specific point-in-time, DB2 Recovery Expert can recover them all at one time, saving the time and complexity of recovering each one separately.

O

Object History

The Object History perspective assists you through the steps of creating and maintaining history information in the Schema Level Repository (SLR) about objects in your database.

P

Point-in-time object recovery

DB2 Recovery Expert allows the recovery of an object (table, table space, or other types of object) to a specified point-in-time. DB2 Recovery Expert automatically determines the most efficient method of recovery, but permits the user to select alternate methods when more than one recovery method is available. DB2 Recovery Expert uses the Schema Level Repository (SLR) to track Object History during the life of the database and ensures that objects are created with the correct attributes for the chosen time frame.

Q

Quiet Time report

The Quiet Time report enables you to determine time ranges during which no active transactions were modifying a set of table spaces and tables.

R

Recovery plan

A collection of one or more recovery scenarios. When a user runs DB2 Recovery Expert to recover one or more objects, it processes the input information, reviews the available database resources (backups, logs and recovery history information), and generates a Recovery Plan. The user can select one of these Recovery Scenarios to perform the recovery. A generated Recovery Plan can be saved as a Specification so that it can be performed at a later time.

Recovery scenario

A collection of steps required to perform a particular recovery. When a particular scenario is selected, a user can review all of the recovery steps and modify some of the attributes of the selected step.

Remote repository

A datastore repository. The repository can reside in the same instance as the user or the target database, or in a different instance.

S

Schema Level Repository (SLR)

The Schema Level Repository (SLR) is a collection of tables that store information that is captured from the DB2 system catalog to record history about the objects in a database over a particular time period. This information allows DB2 Recovery Expert to re-create object DDL. The SLR is also used by Log Analysis to determine the structure for dropped or altered objects when analyzing the logs. The SLR is created from a user-specified backup image and is updated from that point forward by processing the DB2 logs to capture structural changes to the database objects. The SLR includes DB2 System Catalog information and DB2 Recovery Expert object recovery information. Only DB2 System Catalog

information that is required to re-create object DDL is maintained in the SLR. The SLR is a part of the datastore repository database for each DB2 instance.

Session

A Session is used for tracking currently running or finished tasks. For long running tasks, it is possible to disconnect and later reconnect to the Session for retrieving the results. Result information related to a Session is stored in the datastore repository database for future reference so, for example, after running Log Analysis, it is possible to display or export reports and generate Redo or Undo SQL.

Specification

A Specification is a set of user provided options and filters to run Log Analysis, Recovery or Object History tasks. The user can save their current state of processing as a Specification and reload it later to resume processing with the settings they had entered up to that point of work.

T

Translated table space recovery

The Translated table space recovery is one of the recovery scenarios that is generated by the product to recover one or more dropped table spaces to a specified point in time.

DB2 Recovery Expert architecture and process flow

Review the DB2 Recovery Expert architecture and process flow.

DB2 Recovery Expert architecture

The DB2 Recovery Expert architecture consists of the following components:

Client computers

One or more computers can run a web-browser client to access the DB2 Recovery Expert graphical application. Client computers can also run the Java™ Command Line Processor (CLP) client, if it is installed locally.

To run the web-based graphical application, you can use any supported browser client with a supported version of Adobe Flash Player. The Java CLP client software can be installed if a command-line interface is required. A Java Runtime Environment (JRE) is included in the installation for use by the CLP client.

DB2 Recovery Expert server

The DB2 Recovery Expert server computer runs the DB2 Recovery Expert software that hosts the main product application, and communicates to the browser client and to the remote database server systems.

An installation program is provided to install the DB2 Recovery Expert software, the DB2 Recovery Expert application, and supporting software files onto the DB2 Recovery Expert Server computer. Secondary installation files are placed on the DB2 Recovery Expert Server, which will be copied to the remote DB2 database servers when the administrator uses the DB2 Recovery Expert application to configure the product to access the required DB2 instances. A separate installation program is provided to install the Java CLP components onto client computers.

DB2 Recovery Expert database server components

One or more DB2 database server computers contain the data store repositories and user databases.

The figures in this topic illustrate the DB2 Recovery Expert architecture.

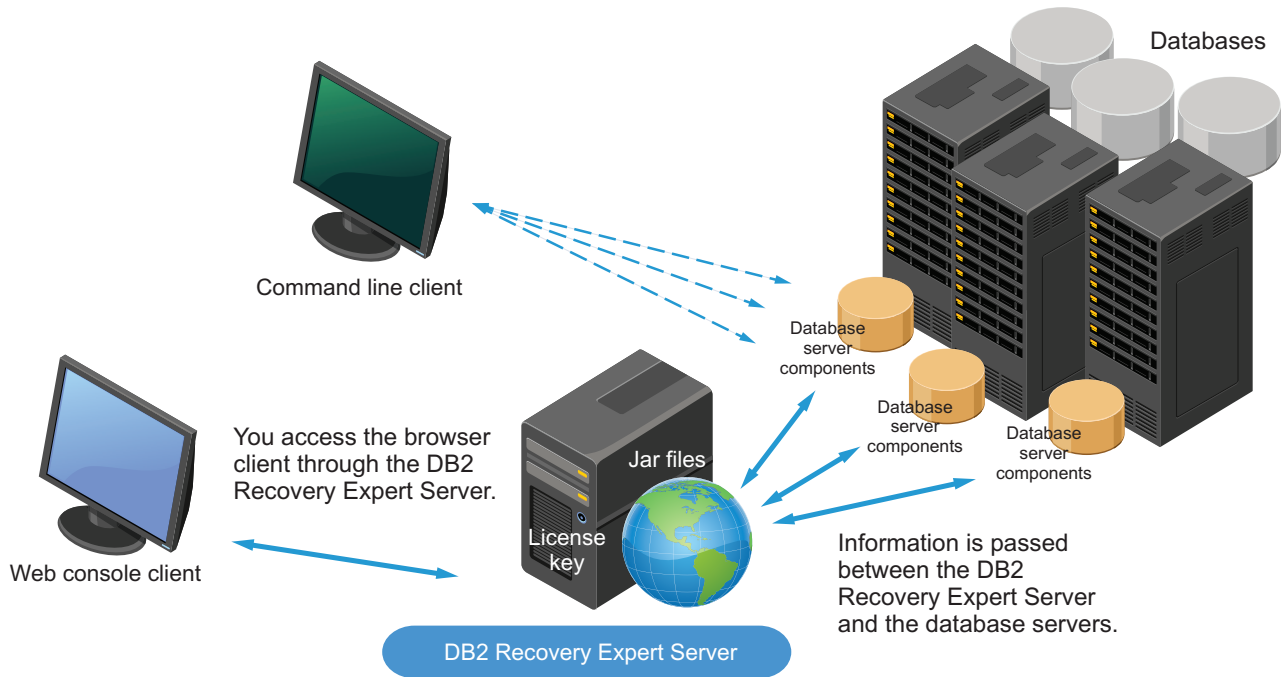


Figure 1. DB2 Recovery Expert architecture diagram 1

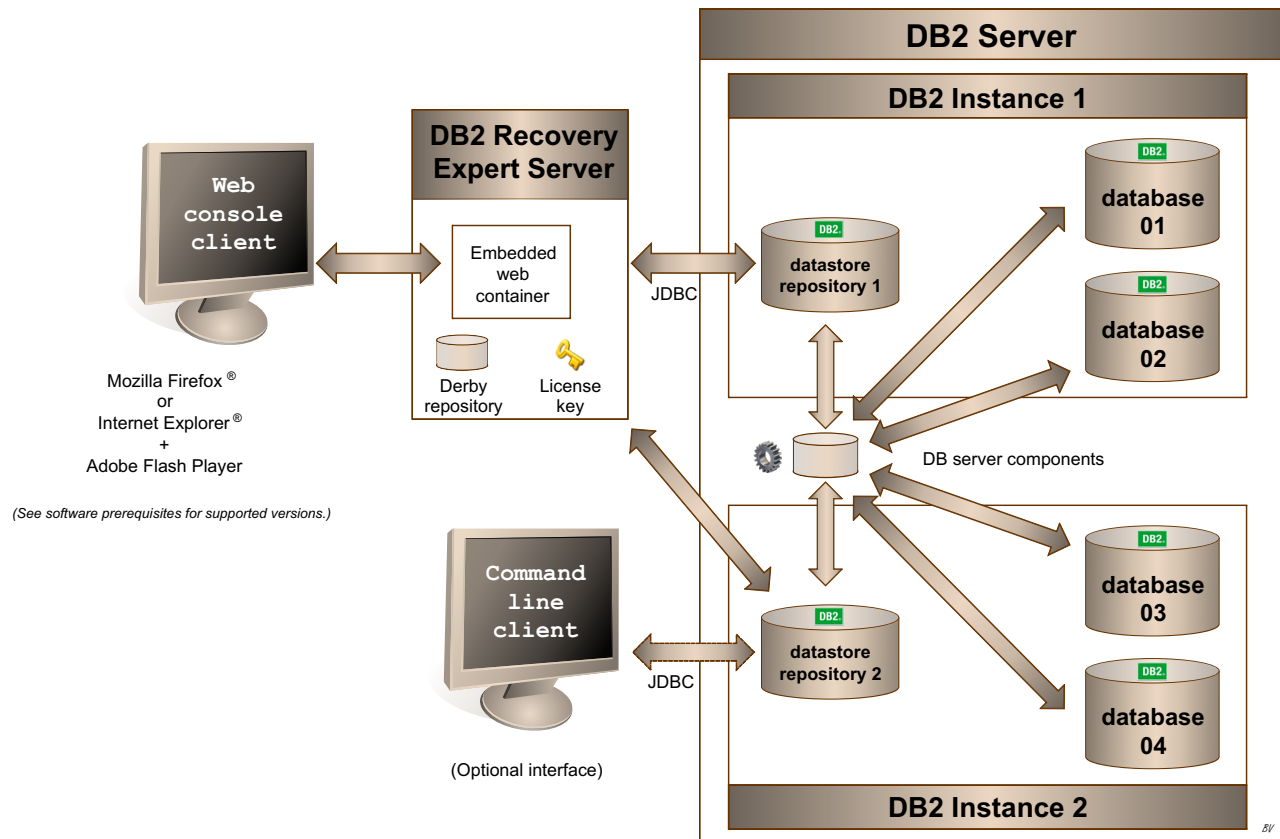


Figure 2. DB2 Recovery Expert architecture diagram 2

DB2 Recovery Expert process flow overview

The DB2 Recovery Expert administrator must perform the following high-level steps to install, configure, and run DB2 Recovery Expert:

1. Plan and prepare for the installation:
 - Review the prerequisites and dependencies.
 - Review the considerations and restrictions.
2. Install and configure components:
 - Install the software.
 - Configure DB2 Recovery Expert:
 - Configure application security.
 - Add and configure databases to work with DB2 Recovery Expert.
 - (optional) Create datastore repository databases for DB2 instances. Instead of creating a datastore, you can use the default datastore that is created during installation.
 - Deploy database server components for the instance of the managed database.
 - Define local data directories and modify other application parameters.

After the administrator has installed and configured the product, users can perform the following tasks from a client computer:

- Create an SLR in the Object History perspective. You cannot perform other DB2 Recovery Expert operations until you have created the SLR in the Object History perspective.

- Perform Log Analysis to generate summary or full reports, perform Redo or Undo SQL, apply filters and options for Log Analysis processing, and work with reports on the browser client.
- Perform a recovery of dropped data.
- Review saved Sessions and Specifications.
- Review command-line processing information.

Related concepts:

Chapter 2, “Preparing to install and configure DB2 Recovery Expert,” on page 17
 Before you install and configure DB2 Recovery Expert, review the information in this section to ensure that your environment and user IDs meet the product prerequisites.

Chapter 10, “Considerations and restrictions,” on page 161

This section lists the considerations and restrictions for installing, configuring, and using DB2 Recovery Expert.

Chapter 3, “Installing DB2 Recovery Expert,” on page 43

DB2 Recovery Expert uses a client/server design installation procedure to install or upgrade the product components.

Chapter 4, “Configuring DB2 Recovery Expert,” on page 61

Review the necessary steps to administer the DB2 Recovery Expert Server, install or upgrade, and configure the database server components, create datastore repositories for DB2 instances, define local data directories, and modify other application parameters.

Chapter 6, “Object History and the SLR,” on page 95

Review Object History and Schema Level Repository (SLR) operations, such as how to view the current SLR range and how to create, rebuild, update, prune, or drop SLR information for user databases.

Chapter 7, “Log Analysis,” on page 105

On the **Log Analysis** perspective of the browser client, you can generate reports, perform Redo or Undo SQL, apply filters and options for Log Analysis processing, and work with reports.

Chapter 8, “Recovery,” on page 123

On the **Recovery** perspective, you can recover data and dropped objects for a database.

“Sessions and specifications” on page 88

Review information about the **Sessions** and **Specifications** perspectives in the DB2 Recovery Expert browser client.

Chapter 9, “Command line processor (CLP) interface,” on page 139

In addition to the browser client, DB2 Recovery Expert provides a command line processor (CLP) interface that enables you to use the product on machines other than the DB2 server.

DB2 Recovery Expert interaction with DB2 and host

DB2 Recovery Expert features interact with and affect DB2 and the host as described in this topic.

Table 1. DB2 Recovery Expert interaction with DB2 and host

Recovery Plan	Recovery Plan tasks	Interaction with DB2	Interaction with host
Dropped Table Recovery	Data extract from backup	<ul style="list-style-type: none"> • Read of recovery history file (RHF) • Storage system contention 	CPU, network, and disk I/O

Table 1. DB2 Recovery Expert interaction with DB2 and host (continued)

Recovery Plan	Recovery Plan tasks	Interaction with DB2	Interaction with host
	Generate DDL	SLR access	Minimal
	Apply DDL	System catalog and table space containers for table and index definitions	Disk I/O
	LOAD table data	Table space containers modified	CPU, network, and disk I/O
	Generate Redo/Undo SQL	<ul style="list-style-type: none"> • Possible read of active logs using db2ReadLog API • Possible table space quiesce • Possible storage system contention 	CPU, network, and disk I/O
	Apply Redo/Undo SQL	Data manipulation language (DML) for recovered table	Disk I/O
	Object creation DDL	System catalog for definition of RI	Minimal
Translated Table Space Recovery	Generate DDL	SLR access	Minimal
	Apply DDL	System catalog and table space containers for table and index definitions	Disk I/O
	BACKUP	Utility contention, table space container access	Network and disk I/O
	RESTORE	Utility contention, table space container access	Network and disk I/O
	ROLLFORWARD	Utility contention, table space container access	Network and disk I/O
	Generate Redo/Undo SQL	<ul style="list-style-type: none"> • Possible read of active logs using db2ReadLog API • Possible table space quiesce • Possible storage system contention 	CPU, network, and disk I/O
	Apply Redo/Undo SQL	DML for recovered table	Disk I/O
DB2 RESTORE/ROLLFORWARD Recovery	RESTORE	Utility contention, table space container access	Network and disk I/O
	ROLLFORWARD	Utility contention, table space container access	Network and disk I/O

Table 1. DB2 Recovery Expert interaction with DB2 and host (continued)

Recovery Plan	Recovery Plan tasks	Interaction with DB2	Interaction with host
Undo from Current®	Generate Undo SQL	<ul style="list-style-type: none"> • Possible read of active logs using db2ReadLog API • Possible table space quiesce • Possible storage system contention 	CPU, network, and disk I/O
	Apply Undo SQL	DML for recovered table	Disk I/O
Redo from Logs	Generate DDL	SLR access	Disk I/O
	Apply DDL	DDL for recovered table	Disk I/O
	Generate Redo SQL	<ul style="list-style-type: none"> • Possible read of active logs using db2ReadLog API • Possible table space quiesce • Possible storage system contention 	CPU, network, and disk I/O
	Apply Redo SQL	DML for recovered table	Disk I/O

DB2 Recovery Expert documentation and updates

Review where you can locate DB2 Tools information on the Web, and find out how you can receive informational updates automatically.

DB2 Recovery Expert information on the Web

The DB2 Tools Library web page provides current product documentation that you can view, print, and download. To locate publications with the most up-to-date information, refer to the following web page:

<http://www-01.ibm.com/software/data/db2/linux-unix-windows/tools/>

Documentation for many DB2 Tools that run on Linux, UNIX, and Windows systems can be found in the Integrated Data Management Information Center:

<http://publib.boulder.ibm.com/infocenter/idm/docv3/index.jsp>

IBM Redbooks® publications that cover DB2 Tools are available from the following web page:

<http://www.redbooks.ibm.com/>

Accessibility features

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use a software product successfully.

The major accessibility features in IBM DB2 Recovery Expert for Linux, UNIX, and Windows enables users to:

- Use assistive technologies, such as screen readers and screen magnifier software.
- Customize display attributes such as color, contrast, and font size.
- Operate specific or equivalent features by using only the keyboard.

Related reference:

“Tips for using the browser client” on page 91

This topic provides basic how-to information for new Recovery Expert users and tips for experienced users.

Chapter 2. Preparing to install and configure DB2 Recovery Expert

Before you install and configure DB2 Recovery Expert, review the information in this section to ensure that your environment and user IDs meet the product prerequisites.

Topics:

- “Hardware and software prerequisites”
- “Installation requirements and recommendations” on page 22
- “Required privileges and authorities” on page 26
- “Operational prerequisites” on page 28
- “Trial and permanent licenses” on page 39
- “DB2 Recovery Expert upgrade” on page 39
- “Configuring Linux and UNIX systems to run the DB2 Recovery Expert Server” on page 40

Hardware and software prerequisites

Before you install and configure DB2 Recovery Expert, ensure that your environment meets the minimum hardware and software requirements.

Always check the system requirements for DB2 Recovery Expert for Linux, UNIX, and Windows web page for the most current version of this information: <http://www-01.ibm.com/support/docview.wss?uid=swg27038720>.

Hardware prerequisites

The hardware requirements for running DB2 Recovery Expert are as follows:

DB2 Recovery Expert server

- Operating system platform: The DB2 Recovery Expert server software can be installed on any supported Linux, UNIX, or Windows operating system platform. Optionally, it can be installed on the same system as an existing DB2 database server.
- Storage and memory requirements:
 - A minimum of 1 GB of disk space should be available on the file system where the DB2 Recovery Expert server components will be installed.
 - Additional space may be required for storing activity logs and intermediate result data that is made available to the DB2 Recovery Expert browser clients. The amount will depend on the type of product activity.
 - At least 1 GB of temporary space should be available during installation on the system temporary file system (the /tmp file system on Linux/UNIX systems, or the location specified by the TEMP variable on Windows systems).
 - A minimum of 2 GB of available memory is recommended.

- Processor requirements: Sufficient processing power should be available to support the demands of any concurrent DB2 Recovery Expert browser client sessions.

DB2 Recovery Expert browser client

- Operating system platform: No product software needs to be installed to run the DB2 Recovery Expert browser client application. The browser client application runs in a supported web browser on any of the supported operating system platforms. It accesses the application software on the system where the DB2 Recovery Expert server is installed.
- Storage and memory requirements:
 - A reasonable amount of local disk space should be available for the user to save Log Analysis reports, SQL output, or other result data.
 - No local disk storage is required for product intermediate or result data because all processing occurs remotely on the DB2 Recovery Expert server or on the DB2 database server systems.
 - A minimum of 2 GB of available memory is recommended.

DB2 Recovery Expert Database Server components

- Operating system platform: The DB2 Recovery Expert Database Server Components must be installed on any supported Linux, UNIX, or Windows operating system platform where the user wants to use DB2 Recovery Expert to access DB2 databases.
- Storage and memory requirements:
 - A minimum of 1 GB of disk space should be available on the file system where the DB2 Recovery Expert Database Server components will be installed.
 - Additional space will be required for storing intermediate working data while the product is running, and for storing result data into the product datastore repository. For more information about disk space requirements, see “Operational prerequisites” on page 28.
 - A minimum of 2 GB of memory should be available for running DB2 Recovery Expert processes on the DB2 database server systems.

DB2 Recovery Expert Database Command Line Processor client

- Operating system platform: The DB2 Recovery Expert Command Line Processor Client can be installed on any supported Linux, UNIX, or Windows operating system platform where the user wants to use the command line feature.
- Storage and memory requirements:
 - A minimum of 500 MB of disk space should be available on the file system where the DB2 Recovery Expert Command Line Processor Client components will be installed.
 - At least 500 MB of temporary space should be available during the installation on the system temporary file system (the /tmp file system on Linux/UNIX systems, or the location specified by the TEMP variable on Windows systems).
 - A minimum of 2 GB of available memory is recommended.

DB2 Recovery Expert License Activation Kit

- Operating system platform: The DB2 Recovery Expert License Activation Kit can be installed on any supported Linux, UNIX, or Windows operating system platform where the DB2 Recovery Expert server has been installed.
- Storage requirements:
 - Only a small amount of disk space is used on the file system where the DB2 Recovery Expert Database Server components were installed.
 - At least 500 MB of temporary space should be available during the installation on the system temporary file system (the /tmp file system on Linux/UNIX systems, or the location specified by the TEMP variable on Windows systems).

Display recommendations

The minimum recommended screen resolution for using the DB2 Recovery browser client interface is 1024 x 768 (width x height). Adjust the following settings to increase the viewable area:

- Disable search bars and Toolbars in the browser.
- Decrease the font size in the browser.
- Decrease the zoom setting in the browser.
- Increase the resolution of the display.

Software prerequisites

The software prerequisites for running the DB2 Recovery Expert server and the DB2 Recovery Expert database server components vary depending on the type of operating system that you are using and other environmental considerations.

Supported operating system platforms

UNIX

Table 2. UNIX supported environments

Operating system	64-bit	Considerations
AIX 6.1	Yes	None
AIX 7.1	Yes	None
Solaris 10	Yes	SPARC 64-bit
HP-UX 11iv2 (11.23.0505)	Yes	Required maintenance: <ul style="list-style-type: none"> • May 2005 Base Quality (QPKBASE) bundle • May 2005 Applications Quality (QPKAPPS) bundle • PHCO_38637 - libc cumulative patch 32-bit is not supported. DB2 Recovery Expert supports only the Itanium-based HP Integrity Series Systems hardware for HP-UX.
HP-UX 11iv3 (11.31)	Yes	Required maintenance: PHCO_38658 - libc cumulative patch. 32-bit is not supported. DB2 Recovery Expert supports only the Itanium-based HP Integrity Series Systems hardware for HP-UX.

Linux on xSeries and zSeries

Note: Because Java version 7.1.1 does not support Red Hat Enterprise Linux (RHEL) 5.0 or SUSE Linux Enterprise Server (SLES) 10, you cannot install the DB2 Recovery Expert server on these versions of Linux. DB2 Recovery Expert still supports working with DB2 servers on any of the supported versions of Linux.

Table 3. Linux supported environments

Operating system	x86-32	x86-64	s390x (zSeries)
Red Hat Enterprise Linux (RHEL) 5.0	Yes	Yes	Yes
Red Hat Enterprise Linux 6.0	Yes	Yes	Yes
SUSE Linux Enterprise Server (SLES) 10	Yes	Yes	Yes
SUSE Linux Enterprise Server (SLES) 11	Yes	Yes	Yes
Ubuntu Linux 12.04.2 LTS	Yes	Yes	No

Table 4. Linux on Inspur K1 platform supported environments

Operating system	64-bit
K-UX Linux	Yes

Microsoft Windows

Table 5. Windows supported environments

Operating system	x86-32	x86-64
Windows XP Professional Service Pack 2 [SP2]	Not applicable	Not applicable
Windows 7 Professional, Enterprise, Ultimate	Yes	Yes
Windows 8 Standard, Professional	Yes	Yes
Windows Server 2003 Standard, Enterprise, DataCenter and DataCenter R2	Yes	Yes
Windows Server 2008 Standard, Enterprise, DataCenter and DataCenter R2	No	Yes
Windows Server 2012 Standard, DataCenter	No	Yes
Windows Server 2012 R2	No	Yes

Supported DB2 versions

DB2 Recovery Expert supports the following versions of DB2. Note that DB2 version 9.8 is *not* supported.

Table 6. DB2 supported versions

DB2 version	DB2 Enterprise Server Edition	DB2 Workgroup Server Edition
Version 9.5	Yes (Fix Pack 8 and later)	Yes (Fix Pack 8 and later)
Version 9.7	Yes (Fix Pack 3a and later) DB2 Recovery Expert uses the db2readlog API to extract log records. In DB2 Version 9.7 Fix pack 6, this API might trigger a private memory leak. For more information about the problem and the resolution, see the IBM Flash http://www-01.ibm.com/support/docview.wss?uid=swg21600160 .	Yes (Fix Pack 3a and later) DB2 Recovery Expert uses the db2readlog API to extract log records. In DB2 Version 9.7 Fix pack 6, this API might trigger a private memory leak. For more information about the problem and the resolution, see the IBM Flash http://www-01.ibm.com/support/docview.wss?uid=swg21600160 .

Table 6. DB2 supported versions (continued)

DB2 version	DB2 Enterprise Server Edition	DB2 Workgroup Server Edition
Version 10.1 (including pureScale support)	Yes	Yes
Version 10.5 (including pureScale support)	Yes	Yes

Supported storage management software

- DB2 with IBM Tivoli® Storage Manager (TSM) client software communicating to compatible versions of a TSM Server.
- DB2 Backup and Restore Vendor API library supported by the LOAD option of BACKUP or the log archival method VENDOR option.

Library requirements for Linux

The Command Line Processor (CLP) client and the DB2 Recovery Expert Activation Kit installation executables are 32-bit and require the 32-bit libraries to be installed. On Linux distributions, the 32-bit executable environment is not installed by default. To enable the 32-bit executable environment, install the following packages:

For RHEL:

- glibc-i686
- libstdc++-i686

For Ubuntu:

- glibc-2.13-1:i386
- libstdc+6:i386
- libstdc+5:i386

The DB2 Recovery Expert database server components for Linux depend on library version `libstdc++.so.5`. Ensure that all available `compat-libstdc++` RPM packages are installed on Linux DB2 server systems where DB2 Recovery Expert will be used.

For more information, see the DB2 installation documentation.

Library requirements for Windows

The latest Microsoft Visual C++ 2010 libraries must be installed on target Microsoft Windows systems where DB2 Recovery Expert database server components will be installed.

The redistributable libraries are available from the Microsoft Download Center on the Microsoft support website. Links for the packages are as follows, with separate packages for 32-bit and 64-bit systems. Before installing, verify that these packages are the latest versions.

32-bit systems

Microsoft Visual C++ 2010 SP1 Redistributable Package (x86)

<http://www.microsoft.com/en-us/download/details.aspx?id=8328>

64-bit systems

Microsoft Visual C++ 2010 SP1 Redistributable Package (x64)

<http://www.microsoft.com/en-us/download/details.aspx?id=13523>

After applying the libraries, it might be necessary to restart the system to complete the installation.

Runtime environment requirement

The IBM XL C/C++ for AIX, V10.1 compiler or runtime environment must be installed on systems on which you will run DB2 Recovery Expert. For more information and to download the Runtime Environment package, see <http://www-01.ibm.com/support/docview.wss?uid=swg24019829>.

Supported web browsers

- Mozilla Firefox 3.6 and later
- Google Chrome 25 and later
- Microsoft Windows Internet Explorer 8 and later.

Considerations for Internet Explorer:

- Compatibility View for Internet Explorer 7.0 is not supported. To disable Compatibility View, complete the following steps:
 1. From an Internet Explorer window, select **Tools > Compatibility View Settings**.
 2. In the Compatibility View Settings window, take the following actions:
 - a. From **Websites you've added to Compatibility View**, remove the host on which the DB2 Recovery Expert server is running.
 - b. Clear the following check boxes:
 - **Display intranet sites in Compatibility View**
 - **Display all websites in Compatibility View**
 3. Select **Tools > Internet Options**.
 4. On the **Advanced** tab, clear **Automatically recover from page layout errors with Compatibility View**.
 5. From an Internet Explorer window, press **F12** and click **Document Mode**. Choose **Internet Explorer 8 standards** or a later version of Internet Explorer.

For more information about Compatibility View, see http://answers.microsoft.com/en-us/ie/forum/ie8-windows_7/turn-off-compatibility-view/33bb7aaf-ab73-47e6-8b5d-d466162ee1cc.

- Internet Explorer 10: Only Internet Explorer 10 desktop mode is supported; the Windows 8 Style UI is not supported.
- Internet Explorer 11: The user-agent string must be set to Default.

Other software requirements

- JavaScript. If your browser has JavaScript disabled, you must re-enable it before using DB2 Recovery Expert.
- Adobe Flash Player Version 10 or later.

Related concepts:

“Operational prerequisites” on page 28

Review the operational prerequisites that are necessary to run DB2 Recovery Expert.

Installation requirements and recommendations

Review the following installation requirements and recommendations before proceeding with the DB2 Recovery Expert installation procedure.

Disk space requirements

DB2 Recovery Expert requires temporary disk storage on the DB2 server systems where it runs to store intermediate results. The amount of required disk space for a single Log Analysis, Object History, or Recovery process execution is proportional to the set of table spaces that are processed by the request. Concurrent processing and saved session data will require additional working storage.

Additional space requirements are necessary for the datastore repository, which holds Schema Level Repository data and result information from the Log Analysis report and SQL generation for any processed databases in the instance. For more information, see the topic about temporary disk space configuration details in the Operational prerequisites section.

Path name requirements

Path names are entered during the installation process to specify the location of the product components. Path names must be entered as absolute (not relative) path names. Path names can contain spaces. Do not include quotes when entering a path name; quoting is handled by the installation application. Path names cannot include environment variables or embedded command shell syntax.

Use only the letters defined in the basic ASCII character set for directory and file names. While your computer's operating system might support different code pages, non-ASCII characters might not work reliably. Using non-ASCII characters can be a particular problem in distributed environments, where different computers might be using different code pages.

UNIX systems

The path name must start with a forward slash ("/") character. The path name cannot include the dollar ("\$") character.

Windows systems

The path name must start with a drive letter (uppercase or lowercase [A-Z]) followed by a colon (":") character and the rest of the path name. Universal Naming Convention (UNC) style path names that include a system hostname are not supported. Path component separators can be specified using either a backslash ("\") or a forward slash ("/") character. Do not mix path separator types in the same path name.

GNU tar required to unpack install archives

The GNU tar command must be used on AIX or Solaris systems to unpack the install packages. The GNU tar command must also be accessible using the command search path so that the installer can use it to unpack embedded packages. If you do not want to make GNU tar the permanent default, you can set the path temporarily when running the installer command.

For example, if GNU tar is available in /usr/local/bin, you can run the installer command using this format:

```
PATH=/usr/local/bin:$PATH <installer-command-line>
```

GNU tar and other GNU tools for AIX are freely available on the Web from the AIX Toolbox for Linux applications. Visit the following website for more information:

<http://www-03.ibm.com/systems/power/software/aix/linux/>

GNU tar is available on Solaris 10 as part of the SUNWgtar package that can be installed on the system. The command is available using the path /usr/sfw/bin/gtar. Because the command name must be "tar" for the installer to locate it, you must create a symbolic link to the command named "tar" and ensure that the link is found first through the search path when running the installer command.

Microsoft Windows system temporary directory requirement

Temporary files used while installing DB2 Recovery Expert components onto the database server system are created in and executed from the default system temporary directory. The user performing the installation must have access to write and execute files created under the system temporary directory. This might be a problem on Windows Vista or Windows 2008 Server systems when installing as a user that is not a domain user account unless the User Account Control feature is disabled. If necessary, the permission settings for the temporary directory can be changed to enable access, or the installation user can specify a different temporary directory path during the installation dialog.

Temporary disk space requirements for the product installers

The installation programs for the DB2 Recovery Expert product, the Command Line Processor, and the license Activation Kit have certain requirements for temporary disk space. The first step of each installer is to self-extract the contents to a temporary directory. The following list describes the temporary directory details listed by the operating system platform:

Windows

The installer extracts to the Windows temporary directory specified by the TEMP environment variable. The self extractor checks for free disk space three times the size of the installer on the volume where the temporary directory is located. If not enough free space is available, the installer will prompt the user for an alternate location for extraction. After the installation has been successfully completed or canceled by the user, the temporary space will be released.

Linux/UNIX

On Linux/UNIX platforms, the installer extracts to the /tmp directory. The self extractor checks for adequate free disk space on the file system where the /tmp directory is located. On Linux/UNIX platforms, the installer does not prompt to pick another location if there is not enough space available. However, the user may set the IATEMPDIR environment variable to a location with sufficient free space and repeat the installation command.

For example:

```
IATEMPDIR=/data/tmp
export IATEMPDIR
./DB2RecoveryExpertLUW.bin
```

After the installation has been successfully completed or canceled by the user, the temporary resources used by the installer will be deleted.

Note: Even if a different temporary directory is specified by using the IATEMPDIR variable, a very small amount of disk space is still required in the /tmp file system. The installation might report errors if the /tmp file system is completely full.

Configuration (metadata) repository database requirements

Determine where you want to store the DB2 Recovery Expert configuration repository and create a DB2 database. DB2 Recovery Expert will use this database to create the configuration repository that it uses to store the list of datastores and managed databases and related information that it requires to work with them. DB2 Recovery Expert will also create a default datastore in the same repository. (Optionally, you can create one or more datastores elsewhere after installation.)

During installation, you will need the following database connection information:

- host name
- database name
- port number
- user ID (must be authorized to connect to the database and create tables)
- password

(DB2 10.5 only)

DB2 10.5 supports the database configuration parameter **dft_table_org**, which controls the default data organization for newly created tables. For the following databases, DB2 Recovery Expert requires the value ROW for this parameter:

- configuration repository database
- any databases in which a datastore repository will be created after installation

Recommendations for installing on an IBM Smart Analytics System

If you are installing DB2 Recovery Expert in an IBM Smart Analytics System environment, review the following recommendations:

- Install the DB2 Recovery Expert Server on the management node to provide access to the internal application network. This step allows the server to support installation of the database server components directly in the browser client. If the DB2 Recovery Expert Server is installed outside of the IBM Smart Analytics System, the database server components must be manually installed by using the command line.
- DB2 Recovery Expert supports shared and non-shared installations of the database server components. For performance reasons, install the components locally on each database partition. In a high availability (HA) configuration, install the DB2 Recovery Expert database server components locally on the administration, data, and standby nodes.
- In an HA configuration, the DB2 Recovery Expert working storage must be configured so that it is accessible on the standby administration node if a fail-over occurs. The DB2 Recovery Expert working storage contains cached information required for processing. When a fail-over occurs, this information must be accessible for DB2 Recovery Expert to continue processing. The DB2 Recovery Expert working storage location is configured by setting the DATA_DIR property.

Changing the language used for installation

By default, the language used for installation is the language that is specified by the locale setting in the environment variables. If you want to use a different language for the installation, run the installation program with an explicit language code using the following command syntax:

```
DB2RecoveryExpertLUW.exe -l <language-code>
```

For example: DB2RecoveryExpertLUW.exe -l ru

The following table lists valid language codes.

Table 7. Language codes

Language	Code
Chinese (Simplified)	zh_CN
Chinese (Traditional)	zh_TW
Czech	cs
English	en
French	fr
German	de
Hungarian	hu
Italian	it
Japanese	ja
Korean	ko
Polish	pl
Portuguese (Brazil)	pt_BR
Russian	ru
Spanish	es

Related concepts:

“Operational prerequisites” on page 28

Review the operational prerequisites that are necessary to run DB2 Recovery Expert.

Required privileges and authorities

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges, higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects.

Requirements for installation

Linux and UNIX systems: The product installation procedure does not need to be run as root. You can choose to install either as root or as another user. If you run the installation procedure as a non-root user, ensure that you have permission to write to the installation directory. The default installation directory for non-root users is set to a path located under the user's home directory.

You can use the browser client to install the DB2 Recovery Expert database server components either as the root user, or as another user by selecting the option to use "su" or "sudo" to enable root authority during the installation.

When performing the following tasks, you must use the same user identity that was used for the original installation (root or non-root user):

- Starting, stopping, and administering the DB2 Recovery Expert Server.
- Installing a permanent license with the IBM DB2 Recovery Expert for LUW Activation Kit.

Microsoft Windows systems: To install or start the DB2 Recovery Expert server, you must belong to the DB2 administrators group, and also have system administrator or equivalent authority.

You can run the command line processor (CLP) client installation procedure as any user.

Requirements for database access

Users can access only those objects for which they have the appropriate authorization; that is, the required privilege or authority.

To use DB2 Recovery Expert with DB2 version 9.5, you must have one of the following authorities to use DB2 Recovery Expert with a DB2 instance:

- SYSADM
- SYSCTRL
- SYSMANT

These authorities include the DBADM authority, which permits access to the database information needed to use DB2 Recovery Expert with any databases in the DB2 instance.

Starting in DB2 version 9.7, certain authorities have been split or extended to give the administrator greater security control. SYSADM authority for the instance no longer implies DBADM authority for each database, and a new DATAACCESS authority was defined to isolate certain abilities from the DBADM authority. In addition, the SECADM authority was enhanced and separated from the other authorities. By default, only the user who created a database has SECADM authority (unless that user grants authority to other users), and only a user with SECADM authority can grant DBADM or DATAACCESS.

To use all DB2 Recovery Expert features with a DB2 instance of version 9.7 or later, users must be granted explicit DBADM, DATAACCESS for the product datastore repository database. User must have necessary permissions to execute DDL, DML, DCL, and TCL statements in the database. DB2 Recovery Expert provides methods to grant these authorities as a set to users.

After a datastore repository is created, authority settings can be updated by using the **Datastores** perspective. Equivalent functionality is available when using the command line to create or update the datastore. Authorities can only be granted or revoked by the user who originally created the datastore unless the administrator has granted SECADM authority to other users.

If the datastore is created in an existing user database, only users who already have the required authorities can use DB2 Recovery Expert. Any users with SECADM authority can propagate the authority to use DB2 Recovery Expert to others by using the methods described previously in this section. DB2 Recovery Expert does not propagate the SECADM authority to other users; it only sets the authorities necessary to permit using DB2 Recovery Expert features.

In addition to setting authorities on the datastore repository database, users of DB2 version 9.7 and later must also be granted DATAACCESS authority on any databases they need to access with DB2 Recovery Expert. Users do not require administrative authorities on the user databases unless they need to perform recovery operations that involve dropping and creating of database objects (such as

tables and table spaces). For additional information about the requirements for using the various DB2 authority settings, see the DB2 documentation.

Related concepts:

“Administering the DB2 Recovery Expert Server” on page 62

To administer the server, you must be the user who originally installed DB2 Recovery Expert, or have equivalent authority. You can start, stop, and display the status of the DB2 Recovery Expert Server at any time with simple commands.

“Trial and permanent licenses” on page 39

During the IBM DB2 Recovery Expert for Linux, UNIX, and Windows installation procedure, you are required to install either the permanent license or the trial license. If you install the trial license, the product is fully functional for the 30 day trial license period.

Related tasks:

“Changing database authorities of an existing datastore connection” on page 74

You can add or revoke authorities on a datastore repository database to allow or disallow user access to DB2 Recovery Expert features.

“Starting and stopping the DB2 Recovery Expert Server” on page 62

The DB2 Recovery Expert Server must be running to use the DB2 Recovery Expert browser client application. You can start and stop the DB2 Recovery Expert Server at any time with simple commands.

“Installing or upgrading DB2 Recovery Expert” on page 45

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the product on a computer that does not have a graphical user interface, you can perform the installation by using the console mode.

“Installing or upgrading the Command Line Processor (CLP) client by using the installation wizard or the console mode” on page 50

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the CLP on a computer that does not have a graphical user interface, you can perform the task by using the console mode.

“Installing a permanent license with IBM DB2 Recovery Expert Activation Kit” on page 52

Upgrade your trial license to a permanent license with the IBM DB2 Recovery Expert for LUW Activation Kit.

Operational prerequisites

Review the operational prerequisites that are necessary to run DB2 Recovery Expert.

Archive log files must be accessible using DB2 configuration

DB2 Recovery Expert locates database archive logs by using information that DB2 records in the Recovery History File (RHF), and by the current database configuration settings for archiving logs. If the user database configuration was recently changed to enable log archiving, DB2 might not have moved all of the archive logs to the archive location. In this case, DB2 Recovery Expert cannot find the files. To avoid problems during this timeframe, force DB2 to archive the logs, or copy the archive logs to a separate location and make the location known to DB2 Recovery Expert. No special action is required after the logs have been copied by DB2 to the archive location.

Archive logging

All databases and database partitions must have archived logging enabled. Circular logging is not supported.

AUTO_DEL_REC_OBJ configuration setting recommendation

When the DB2 Recovery History File (RHF) is pruned and the AUTO_DEL_REC_OBJ configuration setting is in effect, DB2 recovery log files, backup images, and load copy files can be automatically deleted. This action might cause problems for DB2 Recovery Expert if the Schema Level Repository (SLR) references resources that no longer exist. In this case, the SLR should also be pruned to remain synchronized with the DB2 recovery history.

Capturing AUTHID, application name, application ID information, and transaction start time

To capture the authorization ID, application name, and application ID information, use one of the following methods to instruct DB2 to write an additional informational log record at the start of each transaction. To filter on AUTHID, DB2 must be capturing AUTHID information in the log.

(Using the DATA CAPTURE CHANGES [DCC] schema or table option also causes authorization ID information to be recorded at the start of each transaction in the logs. However, the following methods are preferred because they cover all of the tables in the database, and include additional information.)

- *DB2 version 9.7 and earlier:* Set the **DB2_LOGGING_DETAIL** variable by issuing the following command: `db2set DB2_LOGGING_DETAIL=APPLINFO`. The instance must be restarted after setting this registry variable.
- *DB2 10.1 and later:* Set the **log_app1_info** database configuration parameter to YES. The database must be deactivated before this change will take effect.

Tip: Activate application information logging to ensure the most accurate results for the Quiet Time report.

Related dependencies changes

Changes to the related dependencies are not saved in the Sessions and Specifications. Therefore, they must be re-selected for each additional run.

Character columns with invalid character codes

DB2 Recovery Expert processes character column data and converts it to Unicode for display and processing. When invalid character codes are encountered, the conversion to Unicode will replace the invalid codes with the Unicode replacement character code. The Unicode replacement character code might cause the character value to increase in length, causing processing problems. A common cause for this issue is storing binary data in character columns that are not defined as FOR BIT DATA.

Constraints might not be enforced during recovery

Table constraints are defined at the end of recovery processing after all table data has been recovered and SQL operations are processed. In certain situations, this can result in rows that violate the constraints to exist in the table. This might cause the constraint definitions to fail. If the constraints were added and dropped during the recovery timeframe, the table might contain extra rows that were removed during SET INTEGRITY processing. The table data should be checked to verify that it is correct.

Data must be logged for Log Analysis or recovery SQL generation

DB2 Recovery Expert Log Analysis and Recovery processing can only

generate report details and Redo or Undo SQL for tables when complete transaction details are logged. If a table is created or altered using the NOT LOGGED INITIALLY clause, changes made to the table by an Insert, Delete, Update, Create Index, Drop Index, or Alter Table operation in the same unit of work in which the table is created or altered are not written to the DB2 transaction log. This prevents DB2 Recovery Expert from offering certain recovery plans that depend on processing the missing transaction data. It also prevents the Log Analysis tool from generating report details or Redo or Undo SQL statements that depend on the missing data. The same issues apply when the NOT LOGGED clause is used with LOB columns.

Database backups are located using the RHF

DB2 Recovery Expert locates backup images by using information that DB2 records in the Recovery History File (RHF). If a backup does not currently exist in the original location that was recorded in the RHF, complete one of the following actions:

- Update the RHF entry to specify the new location.
- Specify the current location of the file.

Datastore repository database STMTHEAP requirement

The datastore repository database must have a minimum statement heap size (STMTHEAP parameter) of 4096 4-KB pages. If the datastore repository database is created in a database with a smaller statement heap size, creation of the Schema Level Repository (SLR) will fail. The default STMTHEAP value for 32-bit instances is 2048, and must be increased.

DB2 memory management for DB2 Recovery Expert datastore repository

For single partition environments in DB2 Version 9.5 and later, the datastore repository memory resources can be managed automatically by the DB2 Self Tuning Memory Manager (STMM). When DB2 Recovery Expert creates a separate datastore repository, STMM is enabled automatically in single partition environments. STMM is not enabled by default for DPF environments. If a user chooses to create the datastore in an existing database, DB2 Recovery Expert does not change the existing STMM settings. The database administrator can adjust the database resource configuration parameters to improve performance as needed. The DB2 documentation for STMM provides guidelines for activating the feature in DPF environments.

DB2 Version 9.7 consideration

The DB2 Recovery Expert table space ARYSPACE, located in the datastore repository, might be taken offline unexpectedly by DB2 when using DB2 Recovery Expert with DB2 9.7 Versions earlier than Fix Pack 3a. DB2 will log message ECF_FAILED for buffer pool services in the db2diag.log when this issue is encountered. To avoid this problem, apply DB2 Version 9.7 Fix Pack 3a or later. If you encounter this problem, connect to the datastore repository and issue command ALTER TABLESPACE ARYSPACE SWITCH ONLINE. Then, apply DB2 Version 9.7 Fix Pack 3a or later.

DDL activity during online backup

When an online backup is required to obtain initial data for processing data during Log Analysis or Recovery, the DDL operations ALTER TABLE ... ADD COLUMN or ALTER TABLE ... DROP COLUMN cannot occur during the timeframe of the online backup.

If such activity takes place during an online backup, the Log Analysis report or Recovery processing might fail and report an error condition.

Direct access to tape storage

Accessing tape storage directly for archived logs and backup images is not supported. The data must be restored to a local disk or a storage manager might be used.

Full backup is required after a table LOAD or REORG

Log Analysis requires a backup of the table data after a LOAD or REORG to obtain row images for processing of certain logged operations. When DATA CAPTURE CHANGES (DCC) or currently committed semantics of the Cursor Stability isolation level are used, a backup is not required.

The Schema Level Repository update process is based on Log Analysis. After a System Catalog table REORG for processing of logged System Catalog operations, a backup is required.

Graphic strings

A graphic string is a sequence of bytes that represents double-byte character data. The length of the string is the number of double-byte characters in the sequence. If the length is zero, the value is called the empty string. This value should not be confused with the null value.

Graphic strings are not checked to ensure that their values contain only double-byte character code points. (The exception to this rule is an application precompiled with the WCHARTYPE CONVERT option. In this case, validation does occur.) Rather, the database manager assumes that double-byte character data is contained in graphic data fields. The database manager does check that a graphic string value is an even number of bytes long.

Host name requirements for installation

Host names or IP addresses for DB2 servers that are specified during the installation process must be accessible to the DB2 Recovery Expert server. If a host name is not resolvable by Domain Name Services (DNS), it must be added to the local hosts file on the DB2 Recovery Expert server.

Host name requirements for partitioned DB2 instances

Host name entries for partitioned instances in the `db2nodes.cfg` configuration file must be resolvable from the DB2 Recovery Expert server. If host names in the `db2nodes.cfg` file are not resolvable by DNS, they must be added to the local hosts file on the DB2 Recovery Expert server.

URL requirement for installation

The product issues the URL for the browser client based on the information that you provide during installation and your target system configuration. The browser client cannot connect to the server if both of the following conditions exist:

- The host name that is specified in the URL is not a fully qualified domain name.
- The computer that runs the browser client is located in a different domain than the DB2 Recovery Expert server.

To use the browser client under these conditions, manually change the host name in the URL to a fully qualified domain name.

Incremental and delta backups are not used

Log Analysis does not use incremental or delta backup images when locating row images for use in processing. Merge Backup might be used to create a full backup that can be used by Log Analysis.

IPv6 environment requirement

When using the DB2 Recovery Expert browser client to install the database server components onto DB2 server systems in an Internet Protocol version 6 (IPv6) environment, you must identify the target system by using a host name. The installation will fail if you specify an IPv6 address.

Java heap size requirements

DB2 Recovery Expert requires Java heap memory as follows:

- DB2 instance settings: When processing large application messages, if sufficient memory is not available, an XML error might be displayed during Schema Level Repository (SLR) creation for a database with a large number of table spaces, or on the SLR **Operations** step when trying to list a large number of backups.

To avoid this problem, set the DB2 instance Java heap size (JAVA_HEAP_SZ database manager configuration parameter) to 16384 or greater.

To change the JAVA_HEAP_SZ parameter value, issue the following DB2 CLP command, where *value* is 16384 or greater:

```
UPDATE DBM CFG USING JAVA_HEAP_SZ value
```

Note: DB2 instance settings apply to a specific instance. You must separately update the instance configuration parameter for the Java heap size for each instance with which you will use DB2 Recovery Expert.

- DB2 Recovery Expert server settings: When downloading LOB files from the DB2 server to the DB2 Recovery Expert server, if sufficient memory is not available, the download might fail. To avoid this problem, increase the Java heap size to 1 GB (1024 MB) or greater.

To change the Java options that are used when starting the DB2 Recovery Expert server:

1. Add a statement to set the initial Java heap size to 1 GB. For example:
 - Linux and UNIX: In file `<server-install-dir>/serverenv.sh` add the statement

```
JAVA_OPTS="-Xmx1024m $JAVA_OPTS"
```
 - Windows: In file `<server-install-dir>\serverenv.bat` add the statement

```
set JAVA_OPTS=-Xmx1024m %JAVA_OPT%
```
2. Stop and then restart the DB2 Recovery Expert server.

Note: The settings apply to all users that access the DB2 Recovery Expert server, so you only need to set them once.

LOAD copy image processing

LOAD copy images are not used in Log Analysis masked update processing due to the inability to locate LOB data in the LOAD copy image based on the LOB descriptor stored in the data row.

Naming rules

DB2 Recovery Expert requires that you follow DB2 naming rules.

Rules exist for the naming of all database objects, user names, passwords, groups, files, and paths. Some of these rules are specific to the platform that you are working on.

Unless otherwise specified, all names can include the following characters:

- The letters A through Z, and a through z, as defined in the basic (7-bit) ASCII character set. When used in identifiers for objects created with SQL statements, lowercase characters “a” through “z” are converted to uppercase unless they are delimited with quotes (“”).
- 0 through 9.
- ! % () { } . - ^ ~ _ (underscore) @, #, \$, and space.
- \ (backslash).

Restrictions:

Use only the letters defined in the basic ASCII character set for directory and file names. While your computer's operating system might support different code pages, non-ASCII characters might not work reliably. Using non-ASCII characters can be a particular problem in distributed environments, where different computers might be using different code pages.

Partial partition group recovery to end of logs

Recovery processing does not support the recovery of a single partition of a multi-partition database or table space using DB2 utilities or enhanced recovery. Recovery is supported at the complete table space level.

Password prompting for storage vendor access to backup images and archived logs is not supported

When using the DB2 Backup and Restore API for accessing backup images and archived log files, password prompting is not supported. When using TSM, the password generate configuration must be used.

Perform a backup of recovered objects after recovery

It is recommended to take a backup of the recovered objects after the recovery has completed. The backup can be a full backup or a table space backup. If it is a table space backup, it should include the catalog table space (SYSCATSPACE) and all of the user table spaces affected by the recovery. This backup will create a new recovery point for use by future recoveries. If a backup is not taken, it may not be possible to recover the same objects in the future by using certain recovery plans.

PIT selection in a DPF environment

When choosing a Recovery Point In Time (PIT) for table space or table recovery, ensure that backups exist prior to the PIT on all partitions for the table spaces that are participating in the recovery. For Single System View (SSV) backups on DB2 Version 9.5 and later, this not an issue because the backups on all partitions share the same timestamp. For independent partition backups, ensure that the Recovery PIT is after the backup end timestamp of the most recent backup.

Recovering the System Catalog or Recovery Expert datastore is not supported

Only standard DB2 utilities are allowed for the recovery of the System Catalog or DB2 Recovery Expert datastore. Log Analysis reporting is supported. The DB2 Recovery Expert datastore can be rebuilt and the SLR re-created without loss of functionality.

Recovery of a table with LOB data in a separate table space requires a backup that contains both the table data and LOB data

Dropped object recovery of DMS table spaces where the index and LOB data is stored in separate table spaces is only supported when all three table spaces are contained in the same backup image.

Recovery and Log Analysis considerations for tables with large object (LOB) data or extended row size

The following considerations and requirements apply to tables with LOB or LONG columns or with the extended row size feature that was introduced in DB2 10.5:

- A source backup image that contains the table and was taken before the selected recovery point in time (PIT) must exist.
- The definition of the table in the backup must contain the LOB or LONG columns or use the extended row size feature.
- The backup image must fall within the current SLR range.
- For DPF instances, the source backup must be newer than the backup that was used to create the SLR. For non-DPF DB2 instances, the source backup can be the same backup that was used to create the SLR.

When performing a recovery, if UPDATE or DELETE log operations occurred for this table between the time of the source backup and the selected recovery PIT, the data must be reconstructed because DB2 does not log the complete "before image" information. Therefore, all log files for this time frame must be available for processing. The requirements are similar to those for *masked update reconstruction* (see "Terminology" on page 5.)

When performing Log Analysis processing for these tables, Redo or Undo SQL can be generated. However, reconstruction of UPDATE or DELETE operations requires accessing a source backup for initial data, unless the original CREATE TABLE event exists in the logs that are to be processed. Reconstruction cannot be performed in reverse from the current database state for these cases. Reconstruction also cannot cross a REORG or LOAD event that affects the table.

As a best practice, take backup images often to ensure that it will be possible to recover and reconstruct UPDATE and DELETE operations for these tables.

Schema Level Repository (SLR) built for any user database

If the SLR was previously built for any user database, and that database is later restored from a backup, then the SLR for the database is no longer valid. The SLR must be rebuilt from a new backup. If the database was restored to the same name, you can create the SLR from the same backup that was used for the restore. Otherwise, you must create the SLR from a new backup that was taken after the restore was performed.

Selecting schema or mixed object types only recovers DDL definitions

When you use the DB2 Recovery Expert browser client, you can select different types of database objects to be recovered on the **Objects** step on the **Recovery** perspective. By choosing to recover tables or table spaces, you can recover object definitions and the associated table data. Selecting any of the other types of object only allows recovery of the object definitions.

When selecting objects, you must select objects of the same type. For example, select one or more tables, or one or more SQL procedures. Do not select both tables and procedures.

Note: You can select a combination of tables and table spaces; however, only the DDL definitions will be recovered. If you want to recover table or table space data, select a set of the same type of objects (tables or table spaces, but not both).

Selecting the schema object type enables you to recover the object definitions (DDL) for the various objects that are defined with the selected schema names. You can also use a wildcard pattern to match multiple schema names. When recovering using the schema object, only the DDL definitions are recovered for the selected schema contents. Table data and index data are not recovered. In this case, the generated recovery plan is named Re-create object definitions from SLR (DDL only; no data will be restored). The recovery will drop existing objects, including the current table data, before re-creating the objects that are associated with the selected schema, because they existed at the specified recovery point in time (PIT). Tables and associated indexes will be empty after the recovery.

To recover a set of tables (including both definitions and table data) that are defined within the same schema, select the table object type on the **Objects** step. You can select one or more table names from the expanded list of tables, or you can use a wild card pattern to match multiple tables.

To match a set of tables that exist within the same schema, specify a pattern in the format `schema.expression`, where the expression uses the DB2 pattern that matches characters "%" or "?". For example, to match all tables in the schema ORANGE, specify the pattern `ORANGE.%`. You can add more than one wild card pattern to the **Selected objects** list to simplify the pattern expressions or to allow recovering tables from different schemas during the same recovery.

SQL generation considerations

Consider the following information when generating Redo or Undo SQL from Log Analysis and during recoveries. DB2 Recovery Expert and Log Analysis processing cannot mimic application program logic. For generating Redo or Undo SQL, it is strongly recommended that there is a primary key or unique constraint on the table for which you are generating SQL. Without a primary key, you might not get the expected results because it is not possible to uniquely identify rows that were changed by the original SQL. For example, if an application program exists that does UPDATE statements that are WHERE CURRENT OF CURSOR, the DB2 log merely records the fact that the update occurred to a particular row. Any SQL generated by this product cannot be performed against such a particular row; the WHERE clause only determines the rows against which the UPDATE will be performed. This can cause problems against non-uniquely keyed tables. Such an UPDATE statement, when redone, might update many rows. The same holds true for DELETE operations.

Temporary disk space configuration details

DB2 Recovery Expert requires temporary disk storage to store intermediate results on the DB2 server systems where it runs. Predicting storage requirements can be difficult because requirements depend on the user database structure and contents, the nature of the database log activity, and the requested Log Analysis report or recovery scenario that you perform.

The amount of disk space that is required for a single Log Analysis, Object History (SLR), or Recovery process execution is usually related to the set of table spaces that are processed by the request. Data that is extracted from a backup image to perform a recovery or to reconstruct masked update operations from the database log files is temporarily cached on disk, and

later removed after the request is complete. Most requests will only require space that is a small proportion of the total database size.

Creating the Schema Level Repository (SLR) reads the system catalog table space from a backup image and stores the SLR data into the datastore repository.

The initial permanent storage requirements for the SLR are approximately the same size as the user database catalog, plus about 25 percent additional data that remains on disk in the working storage area to optimize subsequent processing.

The size of the SLR data that is saved in the datastore grows over time as user database objects are created, modified, or dropped, and the SLR is updated.

If a Recovery or Log Analysis request needs to process a large number of tables or table spaces, the temporary disk space requirements can approach the size of the entire database. Additional space will be needed by Log Analysis and Recovery for creating SQL and generating temporary load files. If SQL is generated for tables that contain LOB columns, temporary LOB files are created for reference by the generated SQL. These temporary files are released after processing is complete.

Concurrent processing of requests for other databases and saved session data will require additional working storage to account for data that will be kept until the sessions complete or until any saved sessions are marked for cleanup.

The datastore repository contains SLR data for user databases and various status information and Log Analysis result data and reports for current and saved DB2 Recovery Expert sessions. This data exists on the file systems where the user databases are defined. The size requirements will vary depending on the size and number of the results; it will grow and shrink over time as sessions are run, saved, or released.

The location of temporary working storage is controlled by the `DATA_DIR` property in the DB2 Recovery Expert configuration file named `resex.properties`. You set it up when you add the first managed database for a given instance. All managed databases for an instance (on all applicable hosts) use the same `DATA_DIR`.

For performance reasons, it is recommended to use local or fast networked storage for the working data, and to keep the data separate from the file systems where database containers and log files reside to avoid disk or controller contention. In partitioned (DPF) database environments, best performance will occur when each partition accesses independent storage areas, as opposed to writing all product data to a common NFS-mounted file system.

The configuration file is located in the `conf` subdirectory under the database server components installation directory. In physically partitioned database environments where the DB2 Recovery Expert database components are installed onto a common shared file system (a "shared install"), there is one copy of the configuration file. If the components are installed separately onto a local file system for each machine (a "distributed install"), then each machine has its own copy of the configuration file.

The `DATA_DIR` setting indicates the directory path on the DB2 server for storing intermediate data. This directory should point to a file system that

has sufficient free space to run DB2 Recovery Expert. The default setting for DATA_DIR is folder data in the installation directory path.

Changing the location of temporary working storage:

You specify paths to the DATA_DIR, deployment, and temporary directories when adding a managed database if database components must be deployed. The directories are created if they do not exist. (The new directory might not appear immediately, but is available when you perform a subsequent operation such as Log Analysis.) If necessary, and if you have the proper permissions, you can modify the configuration file named `recex.properties` to set the DATA_DIR path to a different location. You cannot modify the path by using the browser client.

Considerations for modifying the configuration file

- Only the administrator who performed the initial installation (root authority) or the instance owner for a non-root installation can modify the configuration file, unless the administrator changed the file permissions to allow write access by other users. To modify the configuration file, you must specify the correct user credentials.
- In physically partitioned database environments using a distributed installation, update the configuration file as needed for each machine where the DB2 Recovery Expert components are installed. Normally, all partitions use the same settings, however, that is not required.
- Choose the DB2 Recovery Expert data directory location with care. If insufficient free space is available and the file system becomes full while running DB2 Recovery Expert, the behavior of DB2 Recovery Expert, and in some cases, of the entire system, is unpredictable and can affect other running processes. This applies in particular to UNIX or Linux systems where the default DATA_DIR setting is in effect and the /data file system is small or shares storage with other file systems such as the root or /tmp file systems. It can also occur if the user-specified DATA_DIR value references a file system that includes other system file systems, such as root, /var, or /tmp. For best results, set DATA_DIR to an independent file system with a large amount of free space that cannot affect the rest of the system if it becomes full.

Timestamps display in current wall clock time

The following timezone options are available for timestamps:

Server time

The timestamps are displayed in the timezone of the DB2 server.

Client time

The timestamps are displayed in the client browser's timezone.

UTC time

The timestamps are displayed in Universal Time Coordinated, which is the same as GMT.

Client time or server timestamp values are displayed in the user interface using the current client or server timezone accordingly. If the timezone observes daylight savings time (DST), the timestamp is displayed using the current daylight savings value.

The time values entered are based on the time zone setting in the client. The user is responsible for making adjustments for daylight savings time (DST).

Note the following points:

- If the time representation in the UI has changed from UTC/Client time to Server time or from Server time to UTC/Client time, then all of the values with date/time data will be dynamically changed in accordance with the current DB2 time offset of the active datastore repository for all open **Log Analysis**, **Recovery** and **Object history** tabs.
- Some error messages might be reported in UTC time and in this case will show the timezone as UTC.

Undo From Current recovery requires tables to be inactive

Performing an Undo From Current recovery requires that the tables participating in the recovery are not actively being modified. DB2 Recovery Expert does not quiesce the table spaces containing the tables or lock them in any way to prevent the tables from being modified. It is up to the user to ensure that the tables are not actively being modified.

Unicode application considerations when using DB2 Recovery Expert

The DB2 Recovery Expert application and datastore repository process characters using the Unicode code page. When the database that is being recovered uses a code page other than Unicode, recovery processing may be affected. Code page conversion of identifiers may cause the number of bytes required to encode the identifier to increase. The increase in size might cause the maximum identifier length to be exceeded. DB2 Recovery Expert will not be able to recover the object.

As a work around, rename the object to an identifier that when encoded in UTF-8, does not exceed the byte length limit of the object type. Once the objects are renamed, perform a backup of the System Catalog, and other affected table spaces on all database partitions. Recovery will then be possible from the point in time of the backup image.

Using a backup image of a table in the SET INTEGRITY PENDING state

Recovery of table data from a backup image where the table is in the SET INTEGRITY PENDING state can result in invalid data. The SET INTEGRITY PENDING state of the source table is not processed. A warning message is issued in this case. Any invalid data in the source table will not be removed in the extracted data or load operation.

Variables cannot be used in the directory paths

Variables cannot be used in the directory paths from the graphical user interface. The actual value must be used.

Related concepts:

“Installation requirements and recommendations” on page 22

Review the following installation requirements and recommendations before proceeding with the DB2 Recovery Expert installation procedure.

Related tasks:

“Specifying additional locations to search for log files and backups” on page 97

On the **Options** step, you can specify additional locations to search for log files and disk backup files that are not located in the default location recorded in the DB2 database history. By default, DB2 Recovery Expert uses backups and logs that DB2 records in the Recovery History File (RHF).

“Providing report details” on page 106

On the **Details** step, you can specify several settings that affect the Log Analysis output.

“Starting and stopping the DB2 Recovery Expert Server” on page 62

The DB2 Recovery Expert Server must be running to use the DB2 Recovery Expert browser client application. You can start and stop the DB2 Recovery Expert Server

at any time with simple commands.

Related reference:

“Hardware and software prerequisites” on page 17

Before you install and configure DB2 Recovery Expert, ensure that your environment meets the minimum hardware and software requirements.

Trial and permanent licenses

During the IBM DB2 Recovery Expert for Linux, UNIX, and Windows installation procedure, you are required to install either the permanent license or the trial license. If you install the trial license, the product is fully functional for the 30 day trial license period.

When the trial license period expires, the product is no longer functional and the permanent license must be installed with the IBM DB2 Recovery Expert for LUW Activation Kit. For more information about upgrading from the trial license to the permanent license with the IBM DB2 Recovery Expert for LUW Activation Kit, see “Installing a permanent license with IBM DB2 Recovery Expert Activation Kit” on page 52.

Related concepts:

“Required privileges and authorities” on page 26

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges, higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects.

Related tasks:

“Installing a permanent license with IBM DB2 Recovery Expert Activation Kit” on page 52

Upgrade your trial license to a permanent license with the IBM DB2 Recovery Expert for LUW Activation Kit.

DB2 Recovery Expert upgrade

Implemented in DB2 Recovery Expert V3.1.0.1 (Fix Pack 1), the product provides an upgrade path from the previous release.

When you launch the product installer, you can choose to upgrade an existing DB2 Recovery Expert installation. The upgrade replaces DB2 Recovery Expert files with new versions, but retains existing configuration information. You can choose to retain the older database server component files, but you can only install the new database server components with the new version of the product.

The database server components that were previously installed on any database servers must be upgraded to match the new version of the product. Usually, the upgrade can be performed from the browser client. However, if the existing components were installed from the command line, the administrator must upgrade the database server components by using the command line.

The datastore repository databases must also be upgraded to match the new product version. This upgrade can be performed from the browser client.

Considerations for upgrading from version 3.1.x.x version 4.1

The licenses that are used by the version 3.1 are not compatible with version 4.1. After upgrading from version 3.1.x.x to version 4.1, you must apply a permanent license by using version 4.1 of the DB2 Recovery Expert for Linux, UNIX, and Windows Activation Kit before using the product.

Upgrading a "Try and Buy" version of the product from 3.1 to 4.1 and continuing to use the product in "Try and Buy" mode is not supported. To use version 4.1 in "Try and Buy" mode, you must uninstall the earlier version, and then perform a new installation of version 4.1.

Related concepts:

Chapter 3, "Installing DB2 Recovery Expert," on page 43

DB2 Recovery Expert uses a client/server design installation procedure to install or upgrade the product components.

Related tasks:

"Installing a permanent license with IBM DB2 Recovery Expert Activation Kit" on page 52

Upgrade your trial license to a permanent license with the IBM DB2 Recovery Expert for LUW Activation Kit.

"Uninstalling DB2 Recovery Expert" on page 57

The uninstallation program removes product files that were installed by the installation program.

Configuring Linux and UNIX systems to run the DB2 Recovery Expert Server

Review the recommended resource limit values for AIX, HP-UX, Linux, and Solaris for optimum performance.

Recommended resource limits for the DB2 Recovery Expert Server

The following table lists the minimum recommended value for each resource:

Resource	Recommended minimum	Considerations
Process number of open files	8000 files	See the next section for default values.
Process file size	10 GB	unlimited is preferred.
Process data segment size	2 GB	The data segment size limit is the data or data seg size parameter in <code>ulimit -a</code> output. unlimited is preferred.
Process virtual memory	2 GB	None

Default number of open files

The recommended minimum number of open files is 8000. The default number of open files are as follows:

- AIX: 2000

If the value of this parameter is too low, errors might occur when opening files or establishing connections. Because this value limits the number of file descriptors that a server process might open, a value that is too low prevents optimum performance.

- HP-UX 11iv2 and HP-UX 11iv3: 2048

If the value is too low, a file open error, memory allocation failure, or connection establishment error might be displayed.

- SUSE Linux Enterprise Server 9 (SLES 9): 1024

If the value is too low, a file open error, memory allocation failure, or connection establishment error might be displayed.

- Solaris: 1024

If the value is too low, the following error is displayed in the DB2 Recovery Expert Server stderr.log file: Too many files open.

Related tasks:

“Changing resource limits”

If necessary, you can increase the data segment size and the number of open files from the operating system defaults for AIX, HP-UX, Linux, and Solaris.

Changing resource limits

If necessary, you can increase the data segment size and the number of open files from the operating system defaults for AIX, HP-UX, Linux, and Solaris.

Procedure

1. To change the data segment size limit for the KornShell shell (ksh), issue the following command: `ulimit -d <limit>`.
2. To change the number of open files from the operating system defaults:
 - To set or view AIX file descriptors (ulimit), perform the following steps:
 - a. Open the command window.
 - b. Edit the `/etc/security/limits` file by adding the following lines to the user account on which the DB2 Recovery Expert Server process runs:

```
nofiles = limit
nofiles_hard = limit
```
 - c. Save the changes.
 - d. Restart the AIX system.
 - e. To verify the result, type one of the following commands on the command line:
To display all ulimit limits: `# ulimit -a`.
To display only the number of open files that are permitted: `# ulimit -n`.
 - To set or view HP-UX, Linux, or Solaris file descriptors (ulimit):
 - a. Check the reference pages on the `ulimit` command for the syntax of different shells and on the file descriptor limits for parameters and commands that are used.
 - b. To set the `ulimit` value for the KornShell shell (ksh), issue the following command:

```
ulimit -n limit
```
 - c. To display all current `ulimit` settings, issue the following command:

```
ulimit -a
```

Related reference:

“Configuring Linux and UNIX systems to run the DB2 Recovery Expert Server” on page 40
Review the recommended resource limit values for AIX, HP-UX, Linux, and Solaris for optimum performance.

Chapter 3. Installing DB2 Recovery Expert



DB2 Recovery Expert uses a client/server design installation procedure to install or upgrade the product components.

DB2 Recovery Expert supports a distributed architecture. An automated installation process is provided for copying the DB2 Recovery Expert installation file to the DB2 Recovery Expert Server.

Instructions for a manual installation using the command line are also provided if the browser client installation process is not compatible with the user's operating environment.

Topics:

- “Installation roadmap”
- “Installing or upgrading DB2 Recovery Expert” on page 45
- “Installing or upgrading DB2 Recovery Expert using the non-interactive (silent) mode” on page 48
- “Installing or upgrading the Command Line Processor (CLP) client by using the installation wizard or the console mode” on page 50
- “Installing a permanent license with IBM DB2 Recovery Expert Activation Kit” on page 52
- “Backing up product application files” on page 53
- “Configuring remote access to database server machines” on page 53
- “Host name resolution and multiple IP addresses for a host name” on page 53
- “Windows targets” on page 54
- “Configuring specific versions of Windows” on page 55
- “Uninstalling DB2 Recovery Expert” on page 57
- “Uninstalling the Command Line Processor (CLP) client by using the uninstallation wizard or the console mode” on page 58

Related concepts:

“DB2 Recovery Expert upgrade” on page 39

Implemented in DB2 Recovery Expert V3.1.0.1 (Fix Pack 1), the product provides an upgrade path from the previous release.

Installation roadmap

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

This roadmap describes typical installation and configuration processes that the administrator performs.

Pre-installation steps and tasks

1. Verify that your environment meets the minimum requirements and that your user ID has the appropriate authorizations for installing the product, and complete any pre-installation tasks.
2. Review installation and configuration considerations and restrictions.

3. If necessary, complete pre-installation tasks to prepare your environment.
4. Copy the installation image to a directory in your environment.

Installation and configuration steps and tasks

1. Install or upgrade the DB2 Recovery Expert Server software, the DB2 Recovery Expert application, and supporting software files onto the DB2 Recovery Expert Server computer by using one of the following methods:
 - Installation wizard.
 - Console mode (Use console mode to install the product on a computer without a graphical user interface.
 - Non-interactive (silent) mode.
2. Complete configuration tasks:
 - (Required) If you did not start the DB2 Recovery Expert Server as part of the installation process, you must start it manually before you can use the browser client application.
 - Configure application security.
 - (Required) Install and configure the database server components.
 - (Required) Create datastore repository databases for DB2 instances.
 - Define local data directories.
 - Install the Command Line Processor (CLP).
 - Configure remote access to database server machines.

Post-installation tasks

- Log into the browser client.
- Review installation details in the log file, `DB2_Recovery_Expert_LUW_InstallLog.txt`, which is located in the product installation directory.
- Perform post-installation tasks, if necessary:
 - Upgrade a trial license to a permanent license.
 - Uninstall the product.
 - Uninstall the CLP.

Related concepts:

Chapter 2, “Preparing to install and configure DB2 Recovery Expert,” on page 17
Before you install and configure DB2 Recovery Expert, review the information in this section to ensure that your environment and user IDs meet the product prerequisites.

Chapter 10, “Considerations and restrictions,” on page 161

This section lists the considerations and restrictions for installing, configuring, and using DB2 Recovery Expert.

“Administering the DB2 Recovery Expert Server” on page 62

To administer the server, you must be the user who originally installed DB2 Recovery Expert, or have equivalent authority. You can start, stop, and display the status of the DB2 Recovery Expert Server at any time with simple commands.

“Managing DB2 Recovery Expert datastores” on page 71

A *datastore* is a database that DB2 Recovery Expert uses for storing information about databases that are configured to work with it.

“Configuring remote access to database server machines” on page 53

Review how to configure remote access to database server machines.

Chapter 4, “Configuring DB2 Recovery Expert,” on page 61

Review the necessary steps to administer the DB2 Recovery Expert Server, install or upgrade, and configure the database server components, create datastore repositories for DB2 instances, define local data directories, and modify other application parameters.

Chapter 5, “DB2 Recovery Expert browser client overview,” on page 83

The DB2 Recovery Expert browser client is a Rich Internet Application (RIA) that runs in any supported web browser.

Related tasks:

“Installing or upgrading DB2 Recovery Expert”

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the product on a computer that does not have a graphical user interface, you can perform the installation by using the console mode.

“Installing or upgrading DB2 Recovery Expert using the non-interactive (silent) mode” on page 48

DB2 Recovery Expert can be installed or upgraded by using the non-interactive (silent) mode to support remote or unattended installations.

“Configuring DB2 Recovery Expert security” on page 65

You can configure the DB2 Recovery Expert security settings so that the product authenticates against the DB2 database that you specify.

“Installing or upgrading the Command Line Processor (CLP) client by using the installation wizard or the console mode” on page 50

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the CLP on a computer that does not have a graphical user interface, you can perform the task by using the console mode.

“Installing a permanent license with IBM DB2 Recovery Expert Activation Kit” on page 52

Upgrade your trial license to a permanent license with the IBM DB2 Recovery Expert for LUW Activation Kit.

“Uninstalling DB2 Recovery Expert” on page 57

The uninstallation program removes product files that were installed by the installation program.

“Uninstalling the Command Line Processor (CLP) client by using the uninstallation wizard or the console mode” on page 58

DB2 Recovery Expert provides an uninstallation wizard to uninstall the Command Line Processor (CLP) client. If you want to uninstall the Command Line Processor (CLP) client from a computer that does not have a graphical user interface, you can perform this uninstallation by using the console mode.

Installing or upgrading DB2 Recovery Expert

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the product on a computer that does not have a graphical user interface, you can perform the installation by using the console mode.

Before you begin

Ensure that your environment meets the product prerequisites, and review product considerations, restrictions, and requirements that are described in the product documentation.

Determine the database where you want to store the DB2 Recovery Expert configuration (metadata) repository. Have the following database connection information available during installation:

- host name
- database name
- port number
- user ID (must be authorized to connect to the database and create tables)
- password

Have available the port numbers for http and https connections for the browser client. Default values for port numbers are 13080 for HTTP and 13081 for HTTPS. If you have more than one copy of the DB2 Recovery Expert server on your computer, to avoid conflicts, you must choose unique port numbers during installation. Make sure that the selected port numbers do not conflict with any other applications on the system.

About this task

When installing on Linux and UNIX systems, for installations that are shared by multiple instances, use the root user. For a single instance installation, specify the instance owner. If you run the installation procedure as a non-root user, make sure that you have permission to write to the installation directory. The default installation directory for non-root users is set to a path located under the user's home directory.

When installing on Windows systems:

- If you are not the administrator, you cannot start the DB2 Recovery Expert server during installation. After installation, the administrator must install service IBMDB2RELUWServer by using the installservice.bat file that is located in the bin directory.
- The installation might require that you restart the computer if certain files are locked and cannot be updated by the installation system.

Do not make changes to your computer environment while the installation program is running. Examples of changes include changes to users or user groups or to DB2. If you make any of these changes while the installation program is running, you must end the installation program and restart it so that the installation program will recognize the changes.

Procedure

1. To start the installation program, from the directory of the installation image for DB2 Recovery Expert, run one of the following commands, depending on your operating system environment:

Table 8. Installation program start commands

Installation method	System	Command
Installation wizard	Linux/UNIX	./DB2RecoveryExpertLUW.bin
Installation wizard	Windows	DB2RecoveryExpertLUW.exe
Console mode	Linux/UNIX	./DB2RecoveryExpertLUW.bin -i console
Console mode	Windows	DB2RecoveryExpertLUW.exe -i console

2. Complete the pages of the wizard or the questions of the console. The wizard or console indicates whether the installation was successful.

3. If the installation failed, look for details in the installation log file, correct the problem, and then try again. The installation log file, `DB2_Recovery_Expert_LUW_InstallLog.txt`, is located in the product installation directory.
4. If prompted, the administrator must replace the service `IBMDB2RELUWServer`.
If you are installing the product again after it was uninstalled and you receive a message stating that the server cannot be updated because the service `IBMDB2RELUWServer` exists from a previous installation, a user with administrator privileges must replace the service:
 - a. Run `removeservice.bat`.
 - b. Run `installservice.bat`.

What to do next

After successfully completing the installation, perform the following tasks:

- *(Required)* You must start the DB2 Recovery Expert Server before you can use the browser client application. If you do not start the server as part of the installation process, you can start it manually.
- *(Required)* add and configure the managed databases that are to be used with the product.
- Log in to the browser client with one of the following URLs:

The hostname that is used in the URL must be accessible from the system that is running the browser client. If the server and the browser client are in different domains, you must enter the fully qualified hostname in the URL.

- `http://hostname_or_IP_address:port_number/recoveryexpert`
- `https://hostname_or_IP_address:port_number/recoveryexpert`

Default values for port numbers are 13080 for HTTP and 13081 for HTTPS. If you have more than one copy of the DB2 Recovery Expert Server on your computer, to avoid conflicts, you must choose unique port numbers during installation. Make sure that the selected port numbers do not conflict with any other applications on the system.

Related concepts:

“Administering the DB2 Recovery Expert Server” on page 62

To administer the server, you must be the user who originally installed DB2 Recovery Expert, or have equivalent authority. You can start, stop, and display the status of the DB2 Recovery Expert Server at any time with simple commands.

“Managing databases used with DB2 Recovery Expert” on page 66

Before you can begin working with DB2 Recovery Expert, you must add and configure the *managed databases* that are to be used with the product. For a managed database, you can create a Schema Level Repository (SLR), run Log Analysis, perform object recovery. A datastore is created by default, and the product uses this datastore by default for all managed databases unless the administrator explicitly creates a separate datastore and sets up the managed database to use the alternate datastore.

Chapter 5, “DB2 Recovery Expert browser client overview,” on page 83

The DB2 Recovery Expert browser client is a Rich Internet Application (RIA) that runs in any supported web browser.

“Required privileges and authorities” on page 26

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges, higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects.

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

Related tasks:

“Uninstalling DB2 Recovery Expert” on page 57

The uninstallation program removes product files that were installed by the installation program.

Installing or upgrading DB2 Recovery Expert using the non-interactive (silent) mode

DB2 Recovery Expert can be installed or upgraded by using the non-interactive (silent) mode to support remote or unattended installations.

Before you begin

Ensure that your environment meets the product prerequisites, and review product considerations, restrictions, and requirements that are described in the product documentation.

About this task

Note: Do not make changes to your computer environment while the installation program is running. Examples of changes include changes to users or user groups or to DB2. If you make any of these changes while the installation program is running, you must end the installation program and restart it so that the installation program will recognize the changes.

Procedure

1. Locate an installation properties response file to copy or edit, and then use as input for a non-interactive installation. A sample installation properties file named `re.properties` is provided in the directory of the installation image. Alternatively, you can use an installation properties file from a previously installed copy of the product.
2. In the installation properties file, set the installation directory location by entering one of the commands provided in the following table.

When you enter the commands, the directory will be created if it does not exist. If the directory already exists, it must be empty when the installation is performed. On Linux/UNIX systems, use a forward slash character as the path separator. On Windows systems, use two backslash characters as the path separator.

Table 9. Set the installation directory location

System	Command
Linux/UNIX	<code>USER_INSTALL_DIR=/opt/ibm/RE</code>
Windows	<code>USER_INSTALL_DIR=C:\\ibm\\RE</code>

3. Choose the components to be installed. To include the component, set the value to 1. To exclude the component, set the value to 0. If you do not intend to use DB2 Recovery Expert with certain platforms, you can skip installing those components to save space.

For example:

```

BACKEND_AIX=1
BACKEND_LINUX=1
BACKEND_HPUX=1
BACKEND_SOLARIS=1
BACKEND_WINDOWS=1

```

4. Specify the user authentication information for the DB2 Recovery Expert Server.
5. Enter the user name in clear text, and enter the password in encrypted form.
 - a. To generate the encrypted password, run the installation program with one of the following arguments:

Table 10. Generate encrypted passwords during the installation procedure

System	Command
Linux/UNIX	<code>./DB2RecoveryExpertLUW.bin -i console -Decryptedpassword=true</code>
Windows	<code>DB2RecoveryExpertLUW.exe -i console -Decryptedpassword=true</code>

- b. Enter the password in the console dialog and press **Enter**.
 - c. Copy the encrypted password string that is displayed into the response file. For example:

```

SERVER_USERNAME=admin
SERVER_PASSWORD=8BA4C0A899B49FAB

```
 - d. Press **Enter** to terminate the program.
6. Specify the port numbers to be used by the DB2 Recovery Expert Server. Ensure that the selected port numbers do not conflict with other ports in use by the system.

For example:

```

SERVER_HTTP_PORT=13080
SERVER_HTTPS_PORT=13081

```
7. Choose whether to start the server automatically after installation. To start the server automatically, set the value to 1. To leave the server inactive, set the value to 0. For example:

```

START_SERVER=1

```

You can also choose to start the server manually after the installation is complete.

8. After the installation properties response file is ready, run the installation program and specify the path name to the response file by using one of the following commands:

Table 11. Specify the path to the installation properties response file

System	Command
Linux/UNIX	<code>./DB2RecoveryExpertLUW.bin -f <response-file-path></code>
Windows	<code>DB2RecoveryExpertLUW.exe -f <response-file-path></code>

For example:

```
./DB2RecoveryExpertLUW.bin -f re.properties
```

9. Check the results of the installation.

- If the installation is successful, the installation log file is copied to the installation directory and is named `DB2_Recovery_Expert_LUW_InstallLog.txt`.
- If the installation fails, the installation log file is written to the system temporary directory and is named `DB2_Recovery_Expert_LUW_InstallLog_<user>.txt`, where `<user>` is the user name that performed the installation. The system temporary directory is set to `/tmp` for Linux and UNIX systems, or to the value of the `%TEMP%` variable on Windows systems.

10. If the installation failed, correct the problem, and then try again.

What to do next

After successfully completing the installation, perform the following tasks:

- *(Required)* You must start the DB2 Recovery Expert Server before you can use the browser client application. If you do not start the server as part of the installation process, you can start it manually.
- *(Required)* add and configure the managed databases that are to be used with the product.
- Log in to the browser client with one of the following URLs:

The hostname that is used in the URL must be accessible from the system that is running the browser client. If the server and the browser client are in different domains, you must enter the fully qualified hostname in the URL.

– `http://hostname_or_IP_address:port_number/recoveryexpert`

– `https://hostname_or_IP_address:port_number/recoveryexpert`

Default values for port numbers are 13080 for HTTP and 13081 for HTTPS. If you have more than one copy of the DB2 Recovery Expert Server on your computer, to avoid conflicts, you must choose unique port numbers during installation. Make sure that the selected port numbers do not conflict with any other applications on the system.

Related concepts:

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

Installing or upgrading the Command Line Processor (CLP) client by using the installation wizard or the console mode

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the CLP on a computer that does not have a graphical user interface, you can perform the task by using the console mode.

Before you begin

Ensure that your environment meets the product prerequisites, and review product considerations, restrictions, and requirements that are described in the product documentation.

Procedure

Install or upgrade the Command Line Processor (CLP) client by using the installation wizard or the console mode.

You can run the command line processor (CLP) client installation procedure as any user.

1. To start the CLP installation program, run one of the following commands, depending on your operating system environment:

Table 12. CLP installation program start commands

Installation method	System	Command
Installation wizard	AIX	<code>./install.clp.aix.bin</code>
Installation wizard	HP-UX	<code>./install.clp.hpux.bin</code>
Installation wizard	Linux/UNIX	<code>./install.clp.linux.bin</code>
Installation wizard	Solaris	<code>./install.clp.solaris.bin</code>
Installation wizard	Windows	<code>install.clp.windows.exe</code>
Console mode	AIX	<code>./install.clp.aix.bin -i console</code>
Console mode	HP-UX	<code>./install.clp.hpux.bin -i console</code>
Console mode	Linux/UNIX	<code>./install.clp.linux.bin -i console</code>
Console mode	Solaris	<code>./install.clp.solaris.bin -i console</code>
Console mode	Windows	<code>install.clp.windows.exe -i console</code>

2. Complete the pages of the wizard or the questions of the console. The wizard or console indicates whether the installation is successful.

Related concepts:

Chapter 9, “Command line processor (CLP) interface,” on page 139

In addition to the browser client, DB2 Recovery Expert provides a command line processor (CLP) interface that enables you to use the product on machines other than the DB2 server.

“Required privileges and authorities” on page 26

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges, higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects.

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

Related tasks:

“Uninstalling the Command Line Processor (CLP) client by using the uninstallation wizard or the console mode” on page 58

DB2 Recovery Expert provides an uninstallation wizard to uninstall the Command Line Processor (CLP) client. If you want to uninstall the Command Line Processor (CLP) client from a computer that does not have a graphical user interface, you can perform this uninstallation by using the console mode.

Installing a permanent license with IBM DB2 Recovery Expert Activation Kit

Upgrade your trial license to a permanent license with the IBM DB2 Recovery Expert for LUW Activation Kit.

About this task

To install a permanent license, you must be the user who originally installed DB2 Recovery Expert, or have authority equivalent to the user who performed the installation.

Procedure

1. Run one of the following executable files, depending on your operating system environment:

Table 13. Commands for installing a permanent license

Installation method	System	Command
Installation wizard	AIX	<code>./install.license.aix.bin</code>
Installation wizard	HP-UX	<code>./install.license.hpux.bin</code>
Installation wizard	Linux/UNIX	<code>./install.license.linux.bin</code>
Installation wizard	Solaris	<code>./install.license.solaris.bin</code>
Installation wizard	Windows	<code>install.license.windows.exe</code>
Console mode	AIX	<code>./install.license.aix.bin -i console</code>
Console mode	HP-UX	<code>./install.license.hpux.bin -i console</code>
Console mode	Linux/UNIX	<code>./install.license.linux.bin -i console</code>
Console mode	Solaris	<code>./install.license.solaris.bin -i console</code>
Console mode	Windows	<code>install.license.windows.exe -i console</code>

2. Complete the pages of the wizard or the questions of the console. The wizard or console indicates whether the installation is successful.

Related concepts:

“Trial and permanent licenses” on page 39

During the IBM DB2 Recovery Expert for Linux, UNIX, and Windows installation procedure, you are required to install either the permanent license or the trial license. If you install the trial license, the product is fully functional for the 30 day trial license period.

“Required privileges and authorities” on page 26

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges, higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects.

“Installation requirements and recommendations” on page 22

Review the following installation requirements and recommendations before proceeding with the DB2 Recovery Expert installation procedure.

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

“DB2 Recovery Expert upgrade” on page 39
Implemented in DB2 Recovery Expert V3.1.0.1 (Fix Pack 1), the product provides an upgrade path from the previous release.

Backing up product application files

The application installation directory on the DB2 Recovery Expert server includes configuration data.

To prevent loss of product configuration data, perform backups of the application installation directory at suitable intervals.

Configuring remote access to database server machines

Review how to configure remote access to database server machines.

(Windows only) DB2 Recovery Expert uses the IBM Tivoli Remote Execution and Access (RXA) component to access remote computers to install software and perform certain administrative tasks. Some configuration is required to ensure that the software operates correctly.

Many operations that RXA performs require access to resources that are not generally accessible by ordinary user accounts. Therefore, for all target platforms, the account names that you use to log onto remote machines must have administrative privileges on each of the target machines. On UNIX or Linux target machines, it is possible for an ordinary user to connect using `su` or `sudo` to obtain administrative privileges.

Related concepts:

“Installation requirements and recommendations” on page 22

Review the following installation requirements and recommendations before proceeding with the DB2 Recovery Expert installation procedure.

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

Host name resolution and multiple IP addresses for a host name

Use RXA to connect to a target computer system specified by a host name.

RXA uses the Java Runtime Environment (JRE) capability to resolve target addresses. JRE hostname-to-IP resolution is based on the information about the local machine configuration and the network naming services (Domain Name System [DNS]). It is possible that the resolution mechanism returns two or more different IP addresses for one host name. They can be both IPv4 and IPv6 addresses. If more than one IP address is specified for a host, RXA connects only to the first address that is returned by the resolution mechanism. If the address is unavailable, RXA will display a `ConnectionException` with a CTGRI0001E message ID. The user must ensure that the DNS information is correct for their working environment.

It is recommended that you configure local machines and network naming services to return the IP addresses of available hosts. If this is not possible, use an IP address of the target host directly.

Windows targets

Some RXA operations rely on VBScript and Windows Management Instrumentation (WMI) calls to execute scripts on Windows targets. If Windows Scripting Host (WSH) or the WMI service is disabled on the target, or if VBScript is otherwise disabled, some Windows protocol methods will not work.

About this task

If you intend to access Windows targets using the SMB protocol over the NetBIOS, which is determined by `setSMBTransportType()`, port 139 or the port specified by `setNetBIOSPort()`, must not be blocked by firewalls or IP security policies and Enable NetBIOS over TCP/IP must be selected in the Control Panel settings for the machine's network connections properties, as follows:

Windows XP and Windows Server 2003:

Control Panel > Network connections > <some connections> > Properties > Internet Protocol (TCP/IP) (double-click or select this and press properties) > Advanced > WINS > Enable NetBIOS over TCP/IP

Windows Vista and Windows 7:

Control Panel > Network and Internet > Network and Sharing Center > <some connection> > Properties > Internet Protocol Version 4 (TCP/IPv4) (double-click or select this and press properties) > Advanced > WINS > Enable NetBIOS over TCP/IP

Windows Server 2008:

Control Panel > Network and Sharing Center > <some connection> View status > Properties > Internet Protocol Version 4 (TCP/IPv4) (double-click or select this and press Properties) > Advanced > WINS > Enable NetBIOS over TCP/IP

Consult the documentation for your firewall to determine that these ports are not blocked for inbound requests.

To determine if security policies are blocking these ports, click **Start > Settings > Control Panel > Administrative Tools**. Depending on whether your policies are stored locally or in an Active Directory, the next steps are as follows:

Locally stored policies:

Administrative Tools > Local Security Policy > IP Security Policies on Local Computer

Policies stored in Active Directory:

Administrative Tools > Default Domain Security Settings > IP Security Policies on Active Directory

Administrative Tools > Default Domain Controller Security Settings > IP Security Policies on Active Directory

Note: These options are available only if your server is a domain controller.

Examine the IP security policies and edit or remove filters that block the ports listed previously. The following list shows the ports that reserved for NetBIOS. Ensure that all ports currently used by RXA are not blocked.

Note: By default, connections are made using port 445, as is determined by the setting of `setSMBPort()`.

The following table displays the ports reserved for NetBIOS communications.

Table 14. NetBIOS Reserved Ports

Port Number	Use
135	NetBIOS Remote procedure call. At this time, RXA does not use this port.
137	NetBIOS name service.
138	NetBIOS datagram. At this time, RXA does not use this port.
139	NetBIOS session (file/print sharing).
445	CIFS (On XP and Windows 2000).

The target machine must have the Remote Registry service started (which is the default configuration) in order for Remote Execution and Access to connect to the target machine.

RXA requires access to the hidden remote administrative disk share, for access to the system %TEMP% and other directories, to ensure proper operation; access to the Interprocess Communications share (IPC\$) is also required for RXA to access remote registries.

Before you access Interprocess Communications share (IPC\$), make sure the Server service is started (**Control Panel > Administrative Tools > Services > Server**).

See the “Configuring specific versions of Windows” topic to review information about how to perform these operations on different operating systems.

Related concepts:

“Configuring specific versions of Windows”

Review information that is specific to supported versions of the Windows operating system.

Configuring specific versions of Windows

Review information that is specific to supported versions of the Windows operating system.

Windows XP:

Windows XP systems must have Simple File Sharing disabled in order for Remote Execution and Access to operate. Simple Networking forces all logins to authenticate as "guest". A guest login does not have the authorizations necessary for Remote Execution and Access to function.

To disable Simple File Sharing:

1. Start Windows Explorer.
2. Click **Tools > Folder Options**.
3. Select the **View** tab.
4. Scroll through the list of settings until you locate **Use Simple File Sharing**.
5. Remove the check mark next to **Use Simple File Sharing**.
6. Click **Apply**
7. Click **OK**.

Windows XP includes a built-in firewall called the Internet Connection Firewall (ICF). By default, ICF is disabled on Windows XP systems. Windows XP Service Pack 2 comes with the Windows Firewall on by default. If either firewall is enabled on a Windows XP or Vista target, it will block attempted accesses by Remote Execution and Access. On XP Service Pack 2, you can select the **File and Printer Sharing** box in the **Exceptions** tab of the Windows Firewall configuration to allow access.

Windows Server 2003:

No special configuration settings are required for Windows Server 2003.

Windows Server 2008:

On Windows Server 2008, disable the User Account Control if your account is not a domain user account. See the section on Windows Vista to learn how to disable the User Account Control.

Windows Vista:

The User Account Control feature in Windows Vista requires users to perform several steps before RXA applications can communicate with Vista targets.

If you have a domain user account, ensure that the local and the target machine are both members of a Windows domain.

If you are a member of a local administrators group and you use a local user account, complete the following steps to enable performance of administrative tasks on the target machine:

1. Enable the built-in Administrator account and use it to connect. To enable the built-in Administrator account, open the Windows Control Panel and click **Administrative Tools > Local Security Policy > Security Settings > Local Policies > Security Options**. Then double-click on **Accounts: Administrator account status** and select **Enable**.
2. Disable User Account Control if a different Administrator user account is to be used to connect to the Vista target. To disable User Account Control, open the Windows Control Panel and click **Administrative Tools > Local Security Policy > Security Settings > Local Policies > Security Options**. Then double-click on **User Account Control: Run all administrators in Admin Approval Mode** and select **Disable**. Changing this setting requires a system restart.
3. Disable User Account Control when you administer a workstation with a local user account (Security Account Manager user account). Otherwise, you will not connect as a full administrator and will not be able to perform administrative tasks.

To disable the User Account Control, perform the following steps:

- a. Click **Start**.
- b. Click **Run**.
- c. Type `regedit`.
- d. Press **Enter**.
- e. Locate and then click the following registry subkey:
`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System`
- f. If the `LocalAccountTokenFilterPolicy` registry entry does not exist, follow these steps:
 - 1) On the **Edit** menu, point to **New** and click **DWORD Value**.

- 2) Type LocalAccountTokenFilterPolicy and press Enter.
- 3) Right-click **LocalAccountTokenFilterPolicy** and click **Modify**.
- 4) In the **Value** data box, type 1 and click **OK**.
- 5) Restart the computer.

Alternatively, you can modify the registry manually by entering the following command:

```
cmd /c reg add HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\
Policies\system
/v LocalAccountTokenFilterPolicy /t REG_DWORD /d 1 /f
```

Windows 7 and Windows 8:

On Windows 7, the default startup type for the Remote Registry service is "manual." The Remote Registry service must be running to enable RXA.

To check whether the Remote Registry service is enabled and started:

1. Click **Start**.
2. On the **Start Search** box, type `services.msc`.
3. Press **Enter**.
4. When the Microsoft Management Console starts, in the console pane, ensure that the service status is set to **started**. If the service status is not set to **started**, right-click **Remote Registry**, and click **Start**.

To avoid problems with the manual startup, it is advisable to set the Remote Registry service startup type to automatic.

If you want to automatically start the service after the server boot:

1. Right-click **Remote Registry** and select **Properties**.
2. On the **Startup type** option, choose **Automatic**.
3. Click **Apply** > **OK**.

When the system starts up, the Remote Registry will automatically start.

Related tasks:

"Windows targets" on page 54

Some RXA operations rely on VBScript and Windows Management Instrumentation (WMI) calls to execute scripts on Windows targets. If Windows Scripting Host (WSH) or the WMI service is disabled on the target, or if VBScript is otherwise disabled, some Windows protocol methods will not work.

Uninstalling DB2 Recovery Expert

The uninstallation program removes product files that were installed by the installation program.

Procedure

1. To start the uninstallation program, from the `Uninstall` directory of the DB2 Recovery Expert image, run one of the following executable files, depending on your operating system environment:

Table 15. Uninstalling DB2 Recovery Expert by using the uninstallation wizard or the console mode

Uninstallation method	System	Executable file
Uninstallation wizard	Linux/UNIX	<code>./Uninstall</code>

Table 15. Uninstalling DB2 Recovery Expert by using the uninstallation wizard or the console mode (continued)

Uninstallation method	System	Executable file
Uninstallation wizard	Windows	Uninstall.exe
Console mode	Linux/UNIX	./Uninstall -i console
Console mode	Windows	Uninstall.exe -i console

2. Complete the pages of the wizard or the questions of the console. The wizard or console indicates whether the uninstallation was successful.
3. If prompted, the administrator must remove the service IBMDB2RELUWServer:
 - a. Run removeservice.bat.
 - b. Repeat the uninstallation steps.

Related concepts:

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

“DB2 Recovery Expert upgrade” on page 39

Implemented in DB2 Recovery Expert V3.1.0.1 (Fix Pack 1), the product provides an upgrade path from the previous release.

Related tasks:

“Installing or upgrading DB2 Recovery Expert” on page 45

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the product on a computer that does not have a graphical user interface, you can perform the installation by using the console mode.

Uninstalling the Command Line Processor (CLP) client by using the uninstallation wizard or the console mode

DB2 Recovery Expert provides an uninstallation wizard to uninstall the Command Line Processor (CLP) client. If you want to uninstall the Command Line Processor (CLP) client from a computer that does not have a graphical user interface, you can perform this uninstallation by using the console mode.

Procedure

Uninstall the Command Line Processor (CLP) client by using the uninstallation wizard or the console mode:

1. To start the uninstallation program, from the Uninstall directory of the DB2 Recovery Expert image, run one of the following executable files, depending on your operating system environment:

Table 16. Uninstalling the Command Line Processor (CLP) client by using the uninstallation wizard or the console mode

Uninstallation method	System	Executable file
Uninstallation wizard	Linux/UNIX	./Uninstall

Table 16. Uninstalling the Command Line Processor (CLP) client by using the uninstallation wizard or the console mode (continued)

Uninstallation method	System	Executable file
Uninstallation wizard	Windows	Uninstall.exe
Console mode	Linux/UNIX	./Uninstall -i console
Console mode	Windows	Uninstall.exe -i console

2. Complete the pages of the wizard or the questions of the console. The wizard or console indicates whether the uninstallation was successful.

Related concepts:

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

Related tasks:

“Installing or upgrading the Command Line Processor (CLP) client by using the installation wizard or the console mode” on page 50

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the CLP on a computer that does not have a graphical user interface, you can perform the task by using the console mode.

Chapter 4. Configuring DB2 Recovery Expert

Review the necessary steps to administer the DB2 Recovery Expert Server, install or upgrade, and configure the database server components, create datastore repositories for DB2 instances, define local data directories, and modify other application parameters.

Configuration requirements

- Before completing configuration tasks, you must install and configure the DB2 Recovery Expert server.
- You can start, stop, check the status, and modify the configuration for the DB2 Recovery Expert server from the system command line environment.
- Required authorities:
 - To administer the server, you must be the user that originally installed DB2 Recovery Expert or have equivalent authority.
 - The DB2 Recovery Expert administrator can perform configuration tasks by using **Setup** options on the **Open** menu of the DB2 Recovery Expert browser client.
- You cannot use the browser client to perform an upgrade of the database server components if the existing installation was performed using the command line because the installation is not shown in the installation list. The administrator must upgrade these components by using the command line.

Configuration overview

Configuration tasks include

- Starting the DB2 Recovery Expert server. If you did not start the server as part of the installation process, you must start it manually before you can use the browser client application.
- Configuring application security.
- (optional) Installing and configuring the database server components. When you add a managed database for a specific instance, the database server components are deployed at the host (using the instance user name password) for the instance to which the database belongs.
- (optional) Creating datastore repository databases for DB2 instances. The product uses the default datastore for all managed databases unless you create a separate datastore.
- Defining local data directories.

Topics:

- “Administering the DB2 Recovery Expert Server” on page 62
- “Configuring the DB2 Recovery Expert repository database” on page 64
- “Configuring DB2 Recovery Expert security” on page 65
- “Configuring DB2 Recovery Expert Services” on page 66
- “Managing databases used with DB2 Recovery Expert” on page 66
- “Managing DB2 Recovery Expert datastores” on page 71
- “Viewing and upgrading deployed database server components ” on page 75

- “Configuring DB2 Recovery Expert for Log Analysis and remote database processing” on page 78
- “Modifying general settings” on page 79
- “Customizing view preferences” on page 80

Administering the DB2 Recovery Expert Server

To administer the server, you must be the user who originally installed DB2 Recovery Expert, or have equivalent authority. You can start, stop, and display the status of the DB2 Recovery Expert Server at any time with simple commands.

You can also update the configuration file to perform the following tasks:

- Change the administrative user name and password.
- Change DB2 Recovery Expert Server timeout settings.
- Change the port numbers used by the DB2 Recovery Expert Server.

Related concepts:

“Required privileges and authorities” on page 26

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges, higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects.

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

Related tasks:

“Starting and stopping the DB2 Recovery Expert Server”

The DB2 Recovery Expert Server must be running to use the DB2 Recovery Expert browser client application. You can start and stop the DB2 Recovery Expert Server at any time with simple commands.

“Changing the DB2 Recovery Expert Server settings and displaying server status” on page 63

You can update the DB2 Recovery Expert Server configuration to change the administrative user and password, server timeout settings, and port numbers used by the server. You can also display the status of the server.

“Installing or upgrading DB2 Recovery Expert” on page 45

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the product on a computer that does not have a graphical user interface, you can perform the installation by using the console mode.

Starting and stopping the DB2 Recovery Expert Server

The DB2 Recovery Expert Server must be running to use the DB2 Recovery Expert browser client application. You can start and stop the DB2 Recovery Expert Server at any time with simple commands.

About this task

To administer the server, the user must either be the same user that originally installed DB2 Recovery Expert, or have equivalent authority.

The server may take from 30 seconds to a few minutes to start.

Procedure

To start or stop the DB2 Recovery Expert Server, run one of the following commands, depending on your operating system environment:

Table 17. Commands for starting and stopping the server

Task	System	Command
Start the server	Linux/UNIX	<installdir>/bin/start.sh
Start the server	Windows	<installdir>\bin\start.bat
Stop the server	Linux/UNIX	<installdir>/bin/stop.sh
Stop the server	Windows	<installdir>\bin\stop.bat

Related concepts:

“Required privileges and authorities” on page 26

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges, higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects.

“Operational prerequisites” on page 28

Review the operational prerequisites that are necessary to run DB2 Recovery Expert.

“Administering the DB2 Recovery Expert Server” on page 62

To administer the server, you must be the user who originally installed DB2 Recovery Expert, or have equivalent authority. You can start, stop, and display the status of the DB2 Recovery Expert Server at any time with simple commands.

Changing the DB2 Recovery Expert Server settings and displaying server status

You can update the DB2 Recovery Expert Server configuration to change the administrative user and password, server timeout settings, and port numbers used by the server. You can also display the status of the server.

About this task

To administer the server, you must be the user who originally installed DB2 Recovery Expert, or have equivalent authority.

Procedure

1. To change DB2 Recovery Expert Server settings, stop the DB2 Recovery Expert Server.
2. To change the administrative user or password definition, run the following command for your operating system:
 - Linux or UNIX: <installdir>/bin/init.sh <username> <password>
 - Windows: <installdir>\bin\init.bat <username> <password>
3. To change the port numbers used by the DB2 Recovery Expert Server:
 - a. Locate the configuration file for your operating system:
 - Linux or UNIX: <installdir>/Config/dswebserver.properties
 - Windows: <installdir>\Config\dswebserver.properties
 - b. Open the file in a text editor and replace the current port numbers. Ensure that the selected port numbers do not conflict with other ports that are in use by the system.

- port=<HTTP-port-number>
 - https.port=<HTTPS-port-number>
4. To change the timeout settings used by the DB2 Recovery Expert Server:
 - a. Locate the server.xml file for your operating system:
 - Linux or UNIX: <installdir>/wlp/usr/servers/dsweb/server.xml
 - Windows: <installdir>\wlp\usr\servers\dsweb\server.xml
 - b. Open the file in a text editor and modify the following attributes:
 - connectionManager attribute maxIdleTime:

Specifies the amount of time after which an unused or idle connection can be discarded during pool maintenance, if doing so does not reduce the pool below the minimum size. A value of -1 disables this timeout. Specify a positive integer followed by a unit of time, which can be hours (h), minutes (m), or seconds (s). For example, specify 30 seconds as 30s. You can include multiple values in a single entry. For example, 1m30s is equivalent to 90 seconds.

The default value is 30m (30 minutes).
 - httpSession attribute invalidationTimeout:

Specifies the amount of time a session can go unused before it is no longer valid. The default value is 1800 seconds (30 minutes). Consider setting the value to 43200 (12 hours) to cause browser session timeouts to occur during off hours rather than during the next working day. When the session cookie timeout occurs (cookieMaxAge), it affects the whole browser session (all open tabs and dialogs). If an active session was in progress, you can restart the browser client application and reconnect to the interrupted session by using the **Sessions** perspective.
 - httpSession attribute cookieMaxAge:

Specifies the maximum amount of time that a cookie can reside on the client browser. The default value is -1 (no limit).
 5. Start the DB2 Recovery Expert Server.

What to do next

To display the status of the DB2 Recovery Expert Server, run the following command for your operating system:

- Linux or UNIX: <installdir>/bin/status.sh
- Windows: <installdir>\bin\status.bat

Related concepts:

“Administering the DB2 Recovery Expert Server” on page 62

To administer the server, you must be the user who originally installed DB2 Recovery Expert, or have equivalent authority. You can start, stop, and display the status of the DB2 Recovery Expert Server at any time with simple commands.

Configuring the DB2 Recovery Expert repository database

You can define a relational database to serve as the DB2 Recovery Expert repository database, detach a repository database, and create schemas for a repository database.

About this task

The DB2 Recovery Expert repository database stores database connection profiles and server set up information.

Procedure

1. From the **Open** menu, choose **Configuration Repository**.
2. To define the repository database, click **Select Repository Database**.
3. Enter information for the database connection, and then click **Test Connection**.
4. If the test is successful, click **OK** to add the new configuration repository database.

Related concepts:

“Main browser client view” on page 85

The **Welcome** perspective is the main dashboard for **DB2 Recovery Expert**. It displays information about how to get started with the product and references to information resources. The **Open** dropdown menu provides product options.

Configuring DB2 Recovery Expert security

You can configure the DB2 Recovery Expert security settings so that the product authenticates against the DB2 database that you specify.

About this task

The default security method for DB2 Recovery Expert is Web container-managed authentication. With this method, the web application server controls user authentication. Alternatively, you can configure DB2 Recovery Expert to use a DB2 database for authentication.

To configure security, the user must either be the same user that originally installed DB2 Recovery Expert, or have equivalent authority.

Procedure

1. From the **Open** menu, choose **Security Configuration**.
2. Click **Configure Repository Database**.
3. To define the repository database, click **Select Repository Database**.
4. Enter information for the database connection, and then click **Test Connection**.
5. If the test is successful, click **OK** to add the new configuration repository database.
6. Refresh the **Security Configuration** tab to see user and group information.
7. If necessary, grant or revoke privileges to users and groups. Note that only administrators for the DB2 Recovery Expert product should be granted Administrator privileges.

Related concepts:

“Main browser client view” on page 85

The **Welcome** perspective is the main dashboard for **DB2 Recovery Expert**. It displays information about how to get started with the product and references to information resources. The **Open** dropdown menu provides product options.

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more information about the steps.

Configuring DB2 Recovery Expert Services

You can configure the DB2 Recovery Expert logging settings.

About this task

You use the **Services** perspective to configure logging for the application server and control the commands that the server sends to the database server components.

Procedure

1. From the **Open** menu, choose **Services**.
2. In the **Services** column, click a service.
3. Click **Configure**.
4. In the pop-up window for the service that you selected, specify configuration information.
5. Click **OK**.

Related concepts:

“Main browser client view” on page 85

The **Welcome** perspective is the main dashboard for **DB2 Recovery Expert**. It displays information about how to get started with the product and references to information resources. The **Open** dropdown menu provides product options.

Managing databases used with DB2 Recovery Expert

Before you can begin working with DB2 Recovery Expert, you must add and configure the *managed databases* that are to be used with the product. For a managed database, you can create a Schema Level Repository (SLR), run Log Analysis, perform object recovery. A datastore is created by default, and the product uses this datastore by default for all managed databases unless the administrator explicitly creates a separate datastore and sets up the managed database to use the alternate datastore.

You can add, edit properties of, and delete managed databases that have been configured to work with the default datastore or a datastore that you have created. You can manage databases by using the DB2 Recovery Expert browser client or by using the command line. To perform database management tasks, you must be a DB2 Recovery Expert administrator.

Related tasks:

“Installing or upgrading DB2 Recovery Expert” on page 45

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the product on a computer that does not have a graphical user interface, you can perform the installation by using the console mode.

Adding a managed database

You can add a managed database to work with the default datastore or a datastore that you have created.

Before you begin

When the managed database is associated with a remote datastore that resides in another instance, a remote database is cataloged in the same instance as the

managed database so that DB2 Recovery Expert components function correctly.

About this task

Icon buttons in the task bar of the perspectives enable you to perform available tasks. Hover the cursor over each icon button to see its name and function.

Procedure

1. From the **Open** menu, choose **Managed Databases**. The **Managed Databases** perspective shows a list of databases that have been configured to work with the product.
2. Click the **Add** icon. The **Add managed database** window opens.
3. On the **Database connection** step, provide database connection details to configure the managed database for use with the product, and then click **Next**.
4. On the **Datastore selection** step, select a datastore from the list, and then click **Next**.
5. Complete one of the following actions, depending on whether DB2 Recovery Expert database server components have been deployed for the instance of the managed database:
 - If so, the **Finish** button is enabled. Click **Finish** to add the managed database to the list of managed databases.
 - If not, continue with step 6.
6. On the **Deployment details** step, provide directory and other required information to deploy database server components, and then click **Next**.
7. On the **Credentials** step, provide connection information to be used to connect to the host and deploy the database server components, and then click **Next**. Deployment of database server components for the instance starts in the background. You can view information about the deployment in the **Deployment Status** perspective. Note that it can take some time to unpack the necessary files and copy them to the database server.

Related concepts:

“Managing DB2 Recovery Expert datastores” on page 71

A *datastore* is a database that DB2 Recovery Expert uses for storing information about databases that are configured to work with it.

Related tasks:

“Viewing and upgrading deployed database server components ” on page 75

The **Deployment Status** perspective provides a unified view of all deployed database server components for all instances on various hosts. It provides a status for each deployment that you can use to identify database server components that need to be upgraded and issues with a specific set of deployed database server components.

Related reference:

“Database connection step” on page 68

“Credentials” on page 69

When you add a managed database to work with the product, you provide credentials that are to be used to connect to the host. When the system or IT administrator changes the credentials for a DB2 system, the managed database properties must be edited to update the credentials settings for one managed database that is associated with the system of the DB2 instance whose credentials changed. If there are multiple managed databases in the same instance, you only need to do this for one database; the change affects all databases.

“Deployment details step”

The **Deployment details** step provides information about a deployment of database server components, including messages that require attention. The most recent messages are listed first.

Database connection step

Hostname

The host name of the managed database.

Database name

The name of the managed database. You must be able to connect to the database and have the **SELECT** privilege on the following administrative views and table functions to allow DB2 Recovery Expert to determine whether additional hosts exist that are part of a DPF configuration:

- Views:
 - SYSIBM.SYSTABLES
 - SYSIBMADM.ENV_INST_INFO
 - SYSIBM.SYSROUTINES
 - SYSIBMADM.DBMCFG
 - SYSIBMADM.DBCFG
- Table functions:
 - DB_MEMBERS()
 - DB_PARTITIONS()

Database port

DB2 Connection port number for the managed database.

Database username

The user name for the connection user. You can connect to the database by entering this user name.

Database password

The password for the connection user. You can connect to the database by entering this password.

Related tasks:

“Adding a managed database” on page 66

You can add a managed database to work with the default datastore or a datastore that you have created.

Deployment details step

The **Deployment details** step provides information about a deployment of database server components, including messages that require attention. The most recent messages are listed first.

Deployment directory

The path in which the DB2 Recovery Expert database server components are to be installed. This path name must be either a local file system or a shared file system that is visible and writable from the system host name that was specified in the **Database connection** step. Only one path name can be specified. If a DPF installation is performed using non-shared directories, the components are copied to the same path name on each system.

Temporary directory

Temporary files that are used during the deployment are created in the default system temporary directory. If necessary, enter a temporary directory path.

Data directory

The data directory path on the DB2 database server that the product uses to store intermediate data.

Create deployment directory path if necessary

Check this box if you want the product to create the deployment directory that you specified.

Overwrite existing files or components

Check this box if you want the product to replace existing database server component files with the components shown in the **Deployment package** section.

Deployment package

This section displays information about the DB2 Recovery Expert installation package, including the product version, the operating system platform (AIX, HP-UX, Linux, Sun, or Windows), and package timestamp and build.

Related tasks:

“Adding a managed database” on page 66

You can add a managed database to work with the default datastore or a datastore that you have created.

Credentials

When you add a managed database to work with the product, you provide credentials that are to be used to connect to the host. When the system or IT administrator changes the credentials for a DB2 system, the managed database properties must be edited to update the credentials settings for one managed database that is associated with the system of the DB2 instance whose credentials changed. If there are multiple managed databases in the same instance, you only need to do this for one database; the change affects all databases.

Username

The authorized user name to log in to the instance.

Password

The password for the authorized user to log in to the instance.

Additional security privileges for UNIX/Linux systems

To install, upgrade, or uninstall DB2 Recovery Expert on UNIX and Linux systems, specify the instance owner for a single instance installation.

If you did not log in initially as an instance user with authority to run the **sudo** or **su** command, then you must complete a two-part authentication process to change to the identity of the instance user.

Select this check box, and from the **Command type** dropdown list, select **sudo** or **su**.

sudo user

This is the sudo command that is used to provide super user privileges to users. When performing installation and other administrative tasks, specify the instance owner for a deployment.

sudo user's password

The sudo user's password.

Note:

If the `/etc/sudoers` file contains the line `Defaults requiretty` then you must run the `sudo` command from a terminal. Setting this requirement prevents DB2 Recovery Expert from working with `sudo` for database server component installations, datastore creation, and other related operations.

To enable these functions, remove the `Defaults requiretty` line from the `/etc/sudoers` file.

For more information, see the `sudo` manual.

su user

This is the standard `su` command. This is `root` for most UNIX systems. When performing installation and other administrative tasks, specify the instance owner for a deployment. When creating or updating the specified datastore repository, specify the instance owner user name.

su user's password

The `su` user's password.

Related tasks:

"Adding a managed database" on page 66

You can add a managed database to work with the default datastore or a datastore that you have created.

Editing managed database properties

You can edit properties of managed databases that have been configured to work with DB2 Recovery Expert.

Procedure

1. From the **Open** menu, choose **Managed Databases**. The **Managed Databases** perspective shows a list of databases that have been configured to work with the product.
2. If necessary, use filters to limit the databases that are shown in the list:
 - a. In the **Databases** area title bar, click the **Filter** icon (right-pointing arrows).
 - b. In the **Object filters** window, choose to include or exclude objects by system, instance, database, and datastore name, and enter the filtering pattern to be used to identify the objects. The filter syntax is similar to the DB2 LIKE clause; use the percent sign (%) as a wildcard character.
 - c. Click **OK**.
 - d. To clear previously entered filters, click **Clear** and then click **OK**.
3. Select a database from the list. The **Select object information** pane displays property information about the database.
4. To edit a property value, click the ellipsis button (...) in the **Value** cell for that property.
5. After editing properties, click **Apply** to update the database information that is stored in the datastore.

Deleting a managed database from a datastore

Before you begin

Before deleting a managed database, be aware that any deployments that were made when adding the managed database are removed if the managed database that you are deleting is the last one.

Procedure

1. From the **Open** menu, choose **Managed Databases**. The **Managed Databases** perspective shows a list of databases that have been configured to work with the product.
 - a. In the **Databases** area title bar, click the **Filter** icon (right-pointing arrows).
 - b. In the **Object filters** window, choose to include or exclude objects by system, instance, database, and datastore name, and enter the filtering pattern to be used to identify the objects. The filter syntax is similar to the DB2 LIKE clause; use the percent sign (%) as a wildcard character.
 - c. Click **OK**.
 - d. To clear previously entered filters, click **Clear** and then click **OK**.
2. Select a database from the list. The **Select object information** pane displays property information about the database.
3. To delete the database from the list of managed databases, in the title bar, click **Delete**.

Managing DB2 Recovery Expert datastores

A *datastore* is a database that DB2 Recovery Expert uses for storing information about databases that are configured to work with it.

The installation process automatically creates a DB2 database to serve as the repository database, and creates a default datastore and a configuration repository in that database. DB2 Recovery Expert uses the default datastore for all managed databases unless you create a separate datastore. You can create additional datastores in any instance on any host by using the DB2 Recovery Expert browser client or by using the command line.

Note:

When you use a remote datastore, processes that require access to the datastore might run more slowly than if you used a datastore that is in the same instance as the target instance. Performance depends on the strength of the network link between the remote datastore repository and the managed database, and the number of network of round trips that are required by an invocation of the database server components.

Using the **Datastores** perspective, you can

- Create a new datastore.
- Change authorities for a datastore.
- Delete datastores that are no longer needed.

Related concepts:

“Installation roadmap” on page 43

This roadmap lists the high-level steps for installing and configuring IBM DB2 Recovery Expert for Linux, UNIX, and Windows and provides references to more

information about the steps.

“Limited or unsupported DB2 features, environments, and operations” on page 168
DB2 Recovery Expert provides limited or no support for the following DB2 features, environments, and operations.

Related tasks:

“Adding a managed database” on page 66

You can add a managed database to work with the default datastore or a datastore that you have created.

Creating a datastore by using the browser client

You can create additional datastores in a database to meet performance objectives. For example, you can create a datastore in the same instance as a managed database.

Procedure

1. From the **Open** menu, choose **Datastores**. The **Datastores** perspective shows a list of and information about the existing DB2 Recovery Expert datastores.
2. In the toolbar, click the **Add** icon (+).
3. On the **Datastore creation** window, provide information about the new datastore repository.
4. Provide credentials information. The user name must have the authority to connect to the database and create tables.
5. If the database that you specified does not already exist, it must be created in one of the following ways:
 - To ensure that the database is created according to the policies of your site, create the database outside of DB2 Recovery Expert, and then create the datastore.
 - To have DB2 Recovery Expert create the database, select **Create database if necessary**.
6. Click **OK**. The **Datastores** perspective shows the new datastore in the **Datastores** list while creation occurs as a background process. To view the creation status, see the **Selected object information** pane.

Creating a datastore by using the command line tools

DB2 Recovery Expert provides command line tools that you can use as an alternative to the browser client to create a datastore.

About this task

- On Windows platforms, the script can only be used with the system default instance.
- You can run the **installDS** script without any arguments to display a usage message that shows all of the supported options.
- You can run the **installDS** script with options to specify the target instance name or override the default parameter settings.
- Examples are provided at the end of this procedure.

Procedure

1. Using ssh, telnet, or a similar service, log on to the DB2 server as a DB2 instance user or as a user with authority to access the DB2 instance.

2. Change the current working directory to the **installDS** subdirectory that is located in the directory in which the product was installed. (You can skip this step and run the command using a full path name.)
3. Launch the **installDS** script for UNIX or the **installDS.bat** script for Windows and provide the required parameters to create the datastore repository database.
4. Optional: Provide an input DDL file to override the default definitions that the product uses to create the datastore repository database table spaces (ARY*). Input DDL files have the following requirements:
 - You must define user temporary table space ARYUSERTEMPSPACE.
 - You must define two regular table spaces: ARYSPACE and ARYNODESPACE.
 - Table space ARYSPACE must be defined on the catalog partition, for example, by creating it in the catalog partition group (IBMCATGROUP).
 - Table space ARYNODESPACE must be defined on all partitions, for example, by using partition group IBMDEFAULTGROUP or an equivalent.
 - Statements must be terminated by a semicolon (;).
 - Statements must be valid and not produce errors. DB2 Recovery Expert does not perform validation on the input before calling the DB2 processor.
5. To register a manually created datastore with the configuration repository, launch the **registerDS.sh** script for UNIX or the **registerDS.bat** script for Windows. To display the help information for the script, run the **registerDS.sh -help** command.

Example

The following example illustrates how to launch the **installDS** script and provide parameters to create the datastore repository database:

```
./installDS.sh -d datastore -o c
<product_install_dir>/install/installDS.sh -d datastore -o c
```

The following example illustrates how to provide an input DDL file:

```
./installDS.sh -o c -d datastore -tsdf /tmp/arytbsp.ddl
```

The following example shows an input DDL file:

```
CREATE USER TEMPORARY TABLESPACE ARYUSERTEMPSPACE PAGESIZE 8192
MANAGED BY AUTOMATIC STORAGE BUFFERPOOL IBMDEFAULTBP;
CREATE TABLESPACE ARYSPACE IN DATABASE PARTITION GROUP IBMCATGROUP PAGESIZE 8192
MANAGED BY AUTOMATIC STORAGE BUFFERPOOL IBMDEFAULTBP NO FILE SYSTEM CACHING;
CREATE TABLESPACE ARYNODESPACE IN DATABASE PARTITION GROUP IBMDEFAULTGROUP PAGESIZE 8192
MANAGED BY AUTOMATIC STORAGE BUFFERPOOL IBMDEFAULTBP NO FILE SYSTEM CACHING;
```

Related tasks:

“Installing or upgrading database server components from the command line” on page 76

Use the command line to install or upgrade database server components on Linux, UNIX, and Windows systems.

Editing datastore properties

Procedure

1. From the **Open** menu, choose **Datastores**. The **Datastores** perspective shows a list of and information about the existing DB2 Recovery Expert datastore repositories.

2. In the **Datastores** pane, select a datastore. The **Select object information** pane displays property information about the datastore. An ellipsis (...) in the **Value** column for a property indicates that you can edit the value.
3. To edit a property value, click in the **Value** cell.
4. After editing the properties of the selected datastore, to save your changes, click **Apply**. To discard your changes, click **Revert**.
5. To change authorities for a datastore, click the "change authorities" icon and provide the user name and password that were used to create the datastore.

Related concepts:

"Sessions and specifications" on page 88

Review information about the **Sessions** and **Specifications** perspectives in the DB2 Recovery Expert browser client.

Changing database authorities of an existing datastore connection

You can add or revoke authorities on a datastore repository database to allow or disallow user access to DB2 Recovery Expert features.

Before you begin

To grant or revoke authorities to other users, you must have the necessary administrative authority on the database.

Procedure

1. From the **Open** menu, choose **Datastores**. The **Datastores** perspective shows a list of and information about the existing DB2 Recovery Expert datastore repositories.
2. In the **Datastores** pane, select the datastore for which you want to change authorities, and click the "change authorities" button.
3. When prompted, specify the user name and password. The **Change authorities for datastore** window opens.
4. You can take the following actions:
 - To add authorities, click **Add**.
 - To remove an authority that you added, choose an entry from the list, and then click **Remove**.
 - To undo changes, click **Reset**.
 - To remove all grant authorities, click, **Revoke all**.
5. To save authority updates, click **Apply**. To discard your changes, click **Cancel**.

Related concepts:

"Required privileges and authorities" on page 26

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges, higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects.

Related tasks:

"Changing database authorities of an existing datastore connection"

You can add or revoke authorities on a datastore repository database to allow or disallow user access to DB2 Recovery Expert features.

Deleting a datastore

Procedure

1. From the **Open** menu, choose **Datastores**. The **Datastores** perspective shows a list of and information about the existing DB2 Recovery Expert datastore repositories.
2. Select a datastore from the list. The **Select object information** pane displays property information about the datastore.
3. To delete the datastore from the list, click the trash can button above the **Select object information** pane. You can delete the datastore only when all of the managed databases have been removed from the datastore.

Datastore basic information and properties

Datastore name

A user-defined name for the datastore connection.

Hostname

A DB2 server target host name. For physically partitioned instances, the host name must identify the system where the datastore repository catalog partition is defined.

Database name

The datastore repository database name to connect.

Database port

The port of the DB2 instance where the datastore repository database is created.

Username

The user name with credentials to connect to the target host. You can connect to the database by entering this user name.

Password

The password for the connection user. You can connect to the database by entering this password.

Create new database if necessary

Check this box if you want the product to create a new database for the datastore. You must provide basic settings for the new database.

Note: The product can create a new datastore repository. However, to ensure that the database is created according the policies of your site, create the datastore repository outside of the product.

Viewing and upgrading deployed database server components

The **Deployment Status** perspective provides a unified view of all deployed database server components for all instances on various hosts. It provides a status for each deployment that you can use to identify database server components that need to be upgraded and issues with a specific set of deployed database server components.

Procedure

1. From the **Open** menu, choose **Deployment Status**.
2. On the **Deployment Status** perspective, choose from the following actions:

- To view deployment information, including messages that require action, double-click a row, or select a row and in the toolbar, click the **Show deployment details** icon.

The **Deployment details** window opens. Review deployment information, and then click **Close**.

- To upgrade database server components, continue with 3.
3. Perform an upgrade:
 - To upgrade a single deployment, select a row that has a status of **New upgrade available**, and in the toolbar, click the **Upgrade** icon.
 - To upgrade all eligible deployments, in the toolbar, click the **Upgrade all** icon.

The upgrade process runs in the background, and affected managed database unavailable until the upgrade is complete.

Related tasks:

“Adding a managed database” on page 66

You can add a managed database to work with the default datastore or a datastore that you have created.

Related reference:

“Deployment details window” on page 77

The **Deployment details** window is a two-part, read-only window. The upper section displays information about a selected deployment of database server components, as described in this topic. The lower section displays messages about an in-progress deployment of database server components, grouped by host. Deployments that are performed for multiple hosts must complete successfully for every host.

Installing or upgrading database server components from the command line

Use the command line to install or upgrade database server components on Linux, UNIX, and Windows systems.

About this task

If you install the database server components manually by using the command line, no entry is created in the browser client for the host. Without the host entry, you cannot use the browser client to edit the DB2 Recovery Expert configuration for the host. To edit the configuration, you must manually edit the configuration file (`conf/recex.properties`) using a native operating system editor.

In the examples, the variable *vrmm* represents the product version number.

Procedure

1. Open the DB2 Recovery Expert installation package and copy the appropriate zip files for your operating system platform (AIX, HP-UX, Linux, Sun, or Windows) to the machine on which the product will be installed.
2. Create a temporary directory on a file system on the host where you want to install the components.
3. Unzip the server components into the temporary directory that you created.
4. On UNIX or Linux operating systems, set execute access on the scripts in the `install` subdirectory by entering the following command:

```
chmod a+x install/*.sh
```


5. If the components are being installed to work with a DPF instance, install them onto a shared file system that is visible to all of the cooperating systems, or install them separately onto each host to a local file system using the same path for all of the hosts. For example, install to `/opt/revrmm` on each host computer.
6. If you are performing a separate installation, you must run the script once for each host where you want to copy the components.
7. If you are performing a shared installation, you can just perform the installation once onto a shared file system that is accessible to all.
8. Run the installation script to create a target installation directory and copy the files into it:

```
./installPkg.sh -c <targetdir>
```

For example:

```
./installPkg.sh -c /opt/revrmm
```

Note: On Windows, the script is named `installPkg.bat`.

9. Create a datastore repository database to be used with a specific DB2 instance. If you are installing more than one instance, create a datastore for each instance. Use the script `installDS.sh` (or `installDS.bat` on Windows) to create the datastore repository database.

Related concepts:

Chapter 4, “Configuring DB2 Recovery Expert,” on page 61

Review the necessary steps to administer the DB2 Recovery Expert Server, install or upgrade, and configure the database server components, create datastore repositories for DB2 instances, define local data directories, and modify other application parameters.

“Installation requirements and recommendations” on page 22

Review the following installation requirements and recommendations before proceeding with the DB2 Recovery Expert installation procedure.

Related tasks:

“Creating a datastore by using the command line tools” on page 72

DB2 Recovery Expert provides command line tools that you can use as an alternative to the browser client to create a datastore.

Deployment details window

The **Deployment details** window is a two-part, read-only window. The upper section displays information about a selected deployment of database server components, as described in this topic. The lower section displays messages about an in-progress deployment of database server components, grouped by host. Deployments that are performed for multiple hosts must complete successfully for every host.

Deployment version

The version of the installed database server components.

Deployment date

The date on which the database server components were installed.

Hostname

Host name of the database server.

OS type

The operating system platform (AIX, HP-UX, Linux, Sun, or Windows).

Instance name

The name of the database instance.

Type of deployment

The type of installation, such as single system or DPF system.

Path The path in which the DB2 Recovery Expert database server components are installed. This path name must be either a local file system or a shared file system that is visible and writable from the system host name that was specified in the **Database connection** step. Only one path name can be specified. If a DPF installation is performed using non-shared directories, the components are copied to the same path name on each system.

Data directory

The data directory path on the DB2 database server. The product uses this directory to write logs of operations such as Log Analysis, SLR, or Recovery performed against managed databases.

Related tasks:

“Viewing and upgrading deployed database server components ” on page 75
The **Deployment Status** perspective provides a unified view of all deployed database server components for all instances on various hosts. It provides a status for each deployment that you can use to identify database server components that need to be upgraded and issues with a specific set of deployed database server components.

Configuring DB2 Recovery Expert for Log Analysis and remote database processing

Review how to set up DB2 Recovery Expert for Log Analysis and remote database processing.

Before you begin

Be aware of the following prerequisites and requirements:

- Remote database processing is only available for non-DPF databases.
- The remote and local systems must run on the same operating system and type (32 or 64 bit), and run the same version and fix pack levels of DB2.
- The database name must be cataloged under the same name so that DB2 Recovery Expert can find the backups.
- Remote Log Analysis only processes the locally available log files; the product does not read active logs.

Procedure

1. Catalog a remote database on a DB2 system (called a *local system*) where the DB2 Recovery Expert database server components are installed. You can complete this step using the DB2 Configuration Assistant or the DB2 Command Line Processor (CLP).
2. Copy the archived log files and backups from the remote host to the local system that you created, or make them available on the local system (for example, via shared NFS).
3. Use either the browser client **Log Analysis** perspective or the DB2 Recovery Expert CLP to run Log Analysis.

Related tasks:

“Connecting to a database to run Log Analysis” on page 106

On the **Location** step in the **Log Analysis** perspective, you can select a local or remote database on which to perform Log Analysis operations.

“Using the CLP to run Log Analysis against remote databases” on page 160
You can use the CLP to run Log Analysis against remote databases.

Modifying general settings

Review how to change or view product-wide settings and configure them for your site's requirements.

Procedure

1. From the **Open** menu, choose **Admin Settings**. The **Admin Settings** perspective shows information such as connections behavior and settings, RXA timeout, and the lifetime for Log Analysis export report files.
2. Use controls that are associated with a particular setting to make changes. For example, to delete all of the existing reports, in the **Export files** area, click **Delete export files**.
3. Choose whether to discard or save changes:
 - To discard changes and to return to original values, click **Restore defaults**. Note that this action resets general settings to default values for all clients.
 - To save changes to the DB2 Recovery Expert server for all clients, click **Apply**.
 - Click "refresh" in the task bar to reload values from DB2 Recovery Expert server and see changes made by other users.

Related concepts:

Chapter 6, “Object History and the SLR,” on page 95

Review Object History and Schema Level Repository (SLR) operations, such as how to view the current SLR range and how to create, rebuild, update, prune, or drop SLR information for user databases.

Administration settings

The **Connections** area contains the following settings

Connection request timeout

Specifies the time limit for the DB2 Recovery Expert Server to establish a connection to a datastore repository when a user running the browser client accesses a DB2 instance.

SSH timeout

Specifies the time limit for attempts to establish an SSH connection.

Idle connection timeout

Specifies the timeframe after which inactive connections between the DB2 Recovery Expert Server and the datastore repository will get disconnected to allow reusing resources. This timeout setting is not related to user inactivity time in the browser client application.

Reap time

Specifies the time interval at which connections will be checked to see if the Idle connection timeout value has been reached. This value should be less than or equal to the Idle connection timeout value.

Maximum connections

Defines the maximum number of connections to a datastore repository database. The maximum number of connections is 1000.

Connection sharing

Controls whether connections to the datastore repository may be shared by

more than one session. Only one action can occur at a time over the connection. If sharing is enabled and one session performs a long running task, other sessions must wait for it to complete. To avoid delays, connection sharing can be disabled. Using separate connections for each session uses more resources, but avoid delays in processing.

The **Remote Execution and Access (RXA) settings** area contains the following settings:

RXA transfer timeout

Specifies the timeframe after which an RXA connection will be disconnected. RXA connections are used to communicate to remote systems to transfer files and run remote commands when installing database server components or when creating a datastore repository.

The **Export files** area contains the following settings:

Lifetime for export files (value in hours)

Specifies the timeframe after which all generated Log Analysis report files will be deleted. You can also delete all of the existing reports by clicking **Delete export files**.

The **Current datastore connections** area contains all of the current active connections to the datastore repository databases and the following fields:

Created on

Displays the date and time when the connection was established.

Uptime

Displays the date and time when the connection will be checked for inactivity.

Database

Displays the datastore repository database name JDBC URL value.

Last IP used

Displays the IP of the client system where the connection was used last.

Customizing view preferences

Customize views throughout the browser client by using the **Preferences** perspective.

Procedure

1. To access the **Preferences** perspective, click **Preferences** in the upper right side of the product title bar.
2. **Number of objects to display within a tree node:** This value controls the maximum number of items that are displayed in a particular tree node/branch. If there are more items in the search result to be displayed, an informational message appears at the end of the selected branch notifying you that the maximum number of items to display has been reached. When this item is selected, a warning message notifies you that the maximum display limit for this branch is reached and that you should narrow your pre-filter search.
3. **Time Representation:** Specify how to display timestamps:
 - **Server time** displays timestamps in the timezone of the DB2 server.
 - **Client time** displays timestamps in the browser client's timezone.

- **UTC time** displays timestamps in Universal Time Coordinated, which is the same as GMT.

Note:

Client time or server timestamp values are displayed in the user interface using the current client or server timezone accordingly. If the timezone observes daylight savings time (DST), the timestamp is displayed using the current daylight savings value.

The time values entered are based on the time zone setting in the client. You must make adjustments for daylight savings time (DST).

If the time representation in the UI has changed from UTC/Client time to Server time or from Server time to UTC/Client time, then all of the values with date/time data will be dynamically changed in accordance with the current DB2 time offset of the active datastore repository for all open **Log Analysis**, **Recovery** and **Object history** tabs.

4. **Status messages to display:** Choose whether to display **All messages**, **Warnings and errors**, or **Errors only**.
5. **Enable support:** Check the box to display the **Support** button in the following locations:
 - **Status** step of the **Object History**, **Log Analysis**, and **Recovery** perspectives
 - **Recovery Plan** and **Dependencies** steps of the **Recovery** perspective. On a physically partitioned instance, the recovery process must complete before you can generate trace files.
 - **SLR Information** step of the **Object History** perspective (if the error occurred after session information became available)

With the **Support** button, you can create a detailed trace that is written to the DB2 Recovery Expert server. Optionally, you can choose to collect database server components trace files for the current session from all hosts and partition locations, and include them in the trace file. You can copy the URL or download the trace file to analyze the operations and provide detailed information for Support.

Chapter 5. DB2 Recovery Expert browser client overview

The DB2 Recovery Expert browser client is a Rich Internet Application (RIA) that runs in any supported web browser.

The DB2 Recovery Expert browser client application provides a graphical user interface for all of the product features, including Log Analysis, Recovery, and administrative tasks. The browser client enables you to run the product from almost any computer, instead of having to log onto the database server machine.

The browser client requires supporting components on other systems. The browser client application servlet is deployed on the DB2 Recovery Expert Server, which communicates to database server components that reside on the database servers. The database server components are provided for supported Linux, UNIX, and Windows environments.

Topics:

- “DB2 Recovery Expert workflow ”
- “Main browser client view” on page 85
- “Database location” on page 84
- “Datastore repository database” on page 84
- “Schema Level Repository (SLR)” on page 85
- “Recovery plans” on page 85
- “Sessions and specifications” on page 88

Related tasks:

“Installing or upgrading DB2 Recovery Expert” on page 45

DB2 Recovery Expert provides an installation wizard to install or upgrade the product. If you want to install or upgrade the product on a computer that does not have a graphical user interface, you can perform the installation by using the console mode.

DB2 Recovery Expert workflow

The product installation is performed by the DB2 Recovery Expert administrator. The administrator also deploys and configures the initial application files on the DB2 Recovery Expert server.

The administrator can then use the DB2 Recovery Expert browser client to define data sources and datastore repositories for DB2 instances and databases that users will access. Database server components must be sent to the various database servers that participate in the configuration, and any remote system configuration changes must be performed to provide storage and other resources that are needed for using the product.

1. After the installation and the initial configuration are complete, users can log in to the DB2 Recovery Expert browser client application by entering one of the following web addresses in a web browser:
 - http://hostname_or_IP_address:port_number/recoveryexpert
 - https://hostname_or_IP_address:port_number/recoveryexpert

The hostname that is used in the URL must be accessible from the system that is running the browser client. If the server and the browser client are in different domains, you must enter the fully qualified hostname in the URL.

The administrator must note the system host name or IP address at the end of the product installation. The default values for the port numbers are 13081 for HTTPS and 13080 for HTTP. If more than one copy of the DB2 Recovery Expert server is installed on one computer, unique port numbers must be assigned during the installation to avoid conflicts.

To ensure proper security, use the https endpoint for DB2 Recovery Expert.

2. If prompted, enter authentication credentials.

Note: You are required to enter information into user interface fields that are marked with an asterisk (*).

3. From the **Welcome** perspective or **Open** menu, choose a task to perform.
4. Select a DB2 system, instance, and database to work with.
5. If prompted, enter authentication credentials to access the product .
6. From the **Open** menu, select the type of tasks to perform, such as Object History, Log Analysis, or Recovery.
7. Specify the necessary options and parameters that are required to perform the task.
8. Save the task specification for later, or submit the task for immediate execution.
9. View the status and upon completion, examine and work with the results of the request.

When working with a database for the first time, the user or administrator must initialize the SLR as a prerequisite for future processing. The user should also perform regular periodic updates to the SLR (using a batch job) so that the history information is up-to-date so that you are prepared to run Log Analysis or Recovery operations.

Database location

To use the product, you must gain access to a specific database in a specific instance that resides on a particular system (machine) or set of systems.

The administrator must identify the databases that are available to use with DB2 Recovery Expert by adding them to a list of managed databases. To perform a DB2 Recovery Expert activity, such as Object History, Log Analysis, or Recovery, you select the database by choosing a predefined managed database from the hierarchy of available databases. This is achieved through the **Location** tab on the various perspectives.

Datastore repository database

DB2 Recovery Expert creates a datastore to track information and status related to a particular database instance.

You can create a datastore by using the **Datastores** perspective in a new database or an existing database on any host. For best performance, the datastore should be created in the same instance in which the managed database resides. The datastore contains stored procedures needed by DB2 Recovery Expert, such as SLR information for databases in the instance, status information for running and completed requests, and result data for Log Analysis reports and generated SQL.

Schema Level Repository (SLR)

A Schema Level Repository (SLR) can be created for each database to track the history of database objects.

The SLR information is initialized from a database backup and is updated from the archive logs, allowing for a continuous record of database history and catalog changes, and allowing recoveries to be performed to any point-in-time, including prior to the time the DB2 Recovery Expert product was installed. The SLR does not need to be updated up to the current time if it is not required to process the requested log range; this eliminates extra log processing.

If needed, Log Analysis can be used immediately, without SLR initialization, to report recent log activity that occurred since the Minimum Recovery Time (MRT) of the table space(s). Reporting prior to the MRT time frame requires creating the SLR to provide historical information about database objects.

Recovery plans

Through the browser client you can specify the point-in-time to recover to and the object name, which can be either existing or dropped.

DB2 Recovery Expert uses the information that you specify through the browser client, and evaluates the available recovery assets to determine which recovery paths are possible to recover the object to the requested point-in-time. DB2 Recovery Expert evaluates the DB2 Recovery Expert SLR entries, available as online and offline DB2 backups, and the DB2 logs which are active and archived. DB2 Recovery Expert will display the possible recovery paths, known as recovery plans, and you can choose the one that you want to execute. Compressed and non-compressed backups are supported.

Main browser client view

The **Welcome** perspective is the main dashboard for **DB2 Recovery Expert**. It displays information about how to get started with the product and references to information resources. The **Open** dropdown menu provides product options.

Open menu options

On the main browser client view, click **Open** and choose a product or set-up option.

DB2 Recovery Expert column options

Object History

The **Object History** (SLR) perspective enables you to create and maintain history information about objects in your database.

Log Analysis

The **Log Analysis** perspective enables you to analyze database archive logs and generate Redo or Undo SQL and reports that can be later exported as text, CSV or XML reports.

Recovery

The **Recovery** perspective enables you to create, save, and perform a recovery plan for a set of database objects.

Sessions

The **Sessions** perspective enables you to track currently running and completed Object History, Log Analysis, and Recovery tasks.

Specifications

The **Specifications** perspective enables you to save a set of user specified values used in Object History, Log Analysis, and Recovery tasks. You can save your current state of processing as a Specification and reload it later to resume processing with the same settings that you entered up to that point of work. Saved Specifications are always stored permanently in the datastore database and are available in the **Specifications** perspective.

Setup column

The **Setup** section of the **Open** menu provides options that enable you to get up and running with DB2 Recovery Expert. Most options requires administrator privileges.

Managed Databases

Choose this option to add, edit, and delete the databases with which the product works.

Datastores

Choose this option to create new or upgrade existing datastore repository databases, and review a list of current datastore connections and controls that are available to create, change or modify these datastore connections. You can also change authorities and change the product's settings.

A default datastore is created in the configuration repository when the DB2 Recovery Expert server is first started and is always available. You can create additional datastores to meet performance objectives. For example, you can create a datastore in the same instance as a managed database for improved performance.

Deployment Status

Choose this option to view the status of all deployed database server components. You can check deployment progress, whether an upgrade is required, and other conditions related to a set of deployed database server components.

Admin Settings

Choose this option to modify various options of the product, such as the datastore connection parameters, the Remote Execution and Access (RXA) protocol transfer timeout, and the control lifetime of the Log Analysis export files. You can also see the current active datastore connections on this step.

Configuration Repository

Choose this option to define a default repository for DB2 Recovery Expert related information, such as installed versions and connections.

Security Configuration

Choose this option to change the security level that DB2 Recovery Expert uses for user authentication.

Services

Choose this option to enable or disable logging and other Optim™ Data Studio related services.

Logs Choose this option to view the Optim Data Studio logs and the DB2

Recovery Expert logs to resolve problems. To view information about a specific log file, select the row of the file and then click **View Log**.

Welcome

The **Welcome** perspective displays general information about the current dashboard and how to customize the current view.

Related concepts:

Chapter 6, “Object History and the SLR,” on page 95

Review Object History and Schema Level Repository (SLR) operations, such as how to view the current SLR range and how to create, rebuild, update, prune, or drop SLR information for user databases.

Chapter 7, “Log Analysis,” on page 105

On the **Log Analysis** perspective of the browser client, you can generate reports, perform Redo or Undo SQL, apply filters and options for Log Analysis processing, and work with reports.

Chapter 8, “Recovery,” on page 123

On the **Recovery** perspective, you can recover data and dropped objects for a database.

“Sessions and specifications” on page 88

Review information about the **Sessions** and **Specifications** perspectives in the DB2 Recovery Expert browser client.

Related tasks:

“Configuring the DB2 Recovery Expert repository database” on page 64

You can define a relational database to serve as the DB2 Recovery Expert repository database, detach a repository database, and create schemas for a repository database.

“Configuring DB2 Recovery Expert security” on page 65

You can configure the DB2 Recovery Expert security settings so that the product authenticates against the DB2 database that you specify.

“Configuring DB2 Recovery Expert Services” on page 66

You can configure the DB2 Recovery Expert logging settings.

Standard action buttons

Several standard actions are available on the various DB2 Recovery Expert perspectives.

The standard actions that are available on the various DB2 Recovery Expert perspectives are as follows:

Cancel

Cancels a running operation.

Save Saves the current parameter settings as a Specification for future reference.

Refresh

Saves the current parameter settings as a Specification for future reference.

Run Performs the task related to the current perspective, such as running Log Analysis, performing a Recovery, or other operations.

New Tab button

The product perspectives, such as **Object History**, **Log Analysis** and **Recovery**, contain a **New Tab** button that enables you to create a new view of the current

perspective. This allows the user to start a new sequence of actions accessible through the same main application panel, while retaining the current action state.

Each time a new view is created, a new button or tab is created on the panel for accessing the view. You can create as many new views (sessions) as necessary, and you can switch among the views by selecting the button related to the view with which you want to work. For example, a user could create a Log Analysis view for database DATA1, another Log Analysis view for database DATA2, and a third Log Analysis view for database DATA3. Each session view can be in a different state; for example, the first view could be actively performing Log Analysis, while the second view could be in the middle of specifying filtering options, and the third view could be in the process of reviewing the results from a requested report. When multiple views are open, the currently displayed view can be closed by selecting the X (close) icon at the right of the panel.

Sessions and specifications

Review information about the **Sessions** and **Specifications** perspectives in the DB2 Recovery Expert browser client.

Topics:

- “Sessions”
- “Specifications” on page 89
- “Saving sessions or specifications” on page 89
- “Searching for saved sessions or specifications” on page 90

Related concepts:

“Main browser client view” on page 85

The **Welcome** perspective is the main dashboard for **DB2 Recovery Expert**. It displays information about how to get started with the product and references to information resources. The **Open** dropdown menu provides product options.

Related tasks:

“Editing datastore properties” on page 73

Related reference:

“DB2 Recovery Expert architecture and process flow” on page 9
Review the DB2 Recovery Expert architecture and process flow.

Sessions

On the **Sessions** perspective, you can track currently running and completed **Object History**, **Log Analysis**, and **Recovery** tasks.

When you run a task, a temporary session object for that task is automatically created and stored in the datastore repository database. You can permanently save a session in the **Sessions** perspective. Use a session to track currently running or finished tasks. For long running tasks, you can disconnect and later reconnect to the session to retrieve the results, display or export reports, and generate Redo or Undo SQL.

Each session has the following attributes:

- **Session ID:** A unique, numerical ID associated with the session.
- **Owner:** The instance user that owns the current session.
- **Database:** The database name for which the selected task was run.
- **Task:** One of the following tasks for which this session was created:

- Object History
- Log Analysis
- Recovery
- Operation: The operation that was performed:
 - Object history: Create, Rebuild, Update, Prune, or Drop.
 - Log Analysis: Summary or Full.
 - Recovery: Dependencies, Generate, or Execute.
- Start time: The start time for the task associated with this session.
- End time: The end time for the task associated with this session.
- Expires in: This field is only used for non-saved sessions. If the session is saved, the **Expires in** value is not used.
- Saved: Whether the session is saved or not saved. If the session is not saved, the session will be deleted after the session's lifetime has expired.
- Status: A message that describes the status of the session, such as running, completed, or failed.
- % finished: The percentage finished displays the percentage of the sessions work that has completed.
- Message: A message that provides a detailed explanation about why the session has failed. Or, if the session has not failed, it will say that the session has completed successfully.
- Comment: A user-defined description given to the session at creation.

Specifications

On the **Specifications** perspective, you can save a set of user-specified values for use in Object History, Log Analysis, and Recovery tasks.

You can save the current state of processing as a specification and reload it later to resume processing with the same settings that you had entered up to that point of work. Saved specifications are stored in the datastore repository database, and are available in the **Specifications** perspective.

When loading a saved Log Analysis, Recovery, or Object History specification, the previously saved settings and values are retrieved for that task so that a run with those parameters can be executed at any time. The loaded specification will update based on the current state of the system when it is loaded.

Each specification has the following attributes:

- Owner: The instance user that owns the current specification.
- Name: The user-defined name for the saved specification.
- Database: The database name for which the selected task was run.
- Type: One of the following tasks for which the specification was created:
 - Object History
 - Log Analysis
 - Recovery
- Creation time: The time at which the specification was created.
- Comment: A user-defined description given to the specification at creation.

Saving sessions or specifications

You can save session result information, and specification settings and task values.

About this task

Saving a session enables you to display or export reports and generate Redo or Undo SQL after running Log Analysis. Saving a specification enables you to run the same task later with the specified parameters.

Procedure

To save the current session or specification:

1. Click **Save**.
2. On the **Save mode** menu, select **Save specification only** or **Save session**.
 - **Save specification only** is available while the session is processing and after the session has completed. You can review a saved specification in the **Specifications** perspective.
 - **Save session** is only available if the session has been executed or is still in progress. You can review a saved session in the **Sessions** perspective. The saved session provides error messages and reports.
3. In the **Name** field, specify an alphanumeric value for the session or specification.
4. In the **Comment** field, type a comment for the saved session or specification.
5. Click **OK**.

Results

The session or specification is saved in the current datastore repository database.

Related tasks:

“Searching for saved sessions or specifications”

You can search for a saved session or specification by specifying one or more attributes as search criteria.

Searching for saved sessions or specifications

You can search for a saved session or specification by specifying one or more attributes as search criteria.

Procedure

1. On the left side of the **Sessions** or **Specifications** perspective, specify search criteria.
 - For a session, you can specify one or more of the following attributes:
 - Location
 - Database
 - Task
 - Owner
 - Status
 - Session ID
 - For a specification, you can specify one or more of the following attributes:
 - Location
 - Database
 - Task
 - Owner
 - Name

The search results are displayed on the right side of the window. If security is enabled on the DB2 Recovery Expert server side, then only search results for the instance user that you used to run the associated task are displayed.

2. Click a session or specification to load it.

Results

For a session, the saved settings and task results are retrieved. For a specification, the saved settings and task values are retrieved. The loaded specification is updated based on the current state of the system when it is loaded.

Related tasks:

“Saving sessions or specifications” on page 89

You can save session result information, and specification settings and task values.

Tips for using the browser client

This topic provides basic how-to information for new Recovery Expert users and tips for experienced users.

Task bar buttons

Icon buttons in the task bar of the perspectives enable you to perform available tasks. Hover the cursor over each icon button to see its name and function.

Selecting objects for processing

A plus sign (+) beside an object indicates that the object is part of a tree that you can expand and collapse. To expand the selected tree root element, click the plus sign. To collapse the selected tree root element, click the minus (-) sign.

To select multiple objects, press and hold **Shift** or **Ctrl**, and then click the objects that you want to select.

You can drag objects from an **Available** panel and drop them into a **Selected** panel.

Keyboard shortcuts

You can use keyboard shortcuts to perform all DB2 Recovery Expert functions.

The following table shows the keyboard shortcut keys that are supported by DB2 Recovery Expert.

Table 18. Keyboard shortcut keys

Keyboard shortcut keys	Action
Tab key	Press the Tab key to move the keyboard focus to the next user interface element. That element will be highlighted by the focus rectangle.
Shift+Tab keys	Press the Shift+Tab keys together to move the keyboard focus to the previous user interface element.

Table 18. Keyboard shortcut keys (continued)

Keyboard shortcut keys	Action
Up and Down arrow keys	Press the Up and Down arrow keys to move the selection or focus to the next or previous group item, such as, groups of buttons, tables, groups of radio buttons and combination box lists.
Space bar key	Press the Space bar key to activate a button.
Page Up and Page Down, Home and End, Left and Right arrow keys	Press the Page Up, Page Down, Home and End keys to scroll through the current list of tables and tree controls. Press the Left and Right arrow keys to expand and collapse the selected tree root element.
Enter and Escape keys	Press the Enter key on a dialog window to confirm the OK user interface icon. Press the Escape key to cancel and close the dialog window. Press the Enter or Escape key to close a message window.
Ctrl+C, Ctrl+X and Ctrl+V keys	Press the Ctrl+C keys together to copy text. Press the Ctrl+X keys together to cut text. Press the Ctrl+V keys together to paste text.
Ctrl+Down arrow keys	Press the Ctrl+Down arrow keys together to open a list menu.
Delete key	Press the Delete key to close a selected tab at the tab bar.

Tips for using patterns to filter objects

You can use a combination of a pattern and exclusion. For example, you can select pattern Q.ST% for tables to be processed, and then exclude table Q.STAFF from that list.

You can select different kinds of objects at the same time. The interaction between different types of objects is specified as AND. The interaction between objects of the same kind is specified as OR. For example, if you select tables PARTS, SALES, and PROJECT, and table space USERSPACE1, Log Analysis will only report on tables named PARTS, SALES, or PROJECT that exist in USERSPACE1.

Example: To exclude the catalog-related schemas and only include schema names that begin with "IBM":

1. On the **Objects** step, enter IBM% as the schema filter value.
2. Select **Include**.
3. On the **Filters** step, check **Exclude System and Tool objects**.
4. Click the refresh icon.

To see the objects that conform to a pattern, you can expand that pattern in the selected objects area.

Related concepts:

“Accessibility features” on page 14

Accessibility features help a user who has a physical disability, such as restricted

mobility or limited vision, to use a software product successfully.

Chapter 6. Object History and the SLR

Review Object History and Schema Level Repository (SLR) operations, such as how to view the current SLR range and how to create, rebuild, update, prune, or drop SLR information for user databases.

Object History prerequisites and considerations

Before you can perform Object History or SLR operations, complete the following installation and configuration tasks:

- (optional) Add a datastore repository by using the **Datastores** perspective.
- Add a managed database by using the **Managed Databases** perspective.
- When backup images are located in an IBM Tivoli Storage Manager (TSM) or vendor storage manager, the VENDOROPT database configuration item must be used to pass an OPTION value for reading backup images, if required.

DB2 Recovery Expert uses the SLR for Log Analysis when Log Analysis is not used in Minimum Recovery Time (MRT) mode, and to perform object recoveries. The SLR should be built and kept up-to-date for all user databases that must be processed in the future. Each user database has its own SLR, but all SLRs are stored in the datastore repository database. The datastore repository database must be created for each instance, and therefore, each datastore will cover only those user databases that are defined within that instance.

Note: Archive logging must be enabled to use DB2 Recovery Expert. If archive logging is disabled in DB2 and then re-enabled, before you run Log Analysis, you must re-create the SLR from a new starting point that is later than the most-recent enabling of logging.

Object History overview

The typical workflow to create or rebuild the SLR is as follows:

1. Select an existing user database that requires the SLR to be created.
2. Select an existing backup that contains system catalog information (for example, SYSIBM tables). This represents the start point of the future SLR.
3. Choose the end point of the future SLR. By default, the end point is set to the current state; however, it can be easily adjusted to a required period.
4. Click **Run** to initiate the SLR create and update phase.

You use the following steps in the **Object History** perspective to perform operations related to the SLR:

- **Overview** step. Review a description of the **Object History** perspective and the operations that can be performed from this perspective.
- **Location** step. Review the list of registered databases in the **List of databases** panel. Select a recoverable database and view detailed information about that database in the **Selected database information** panel.
- **SLR information** step. Review information such as whether the SLR is built, database registration timestamp, and start and end information for the current history range.

Note: For Database Partitioning Facility (DPF) environments, this information is applicable to the catalog partition only; however, the SLR is also built on non-catalog partitions.

- **Options** step. Review additional log directories and backup locations for SLR operations. This step is optional, and is not necessary in most cases. The options can be useful when backups and logs are moved to another directory.
- **Operations** step. Specify the necessary operations to be performed on the SLR, including creating, rebuilding, updating, pruning, and dropping the SLR.
- **Status** step. Review information about the current processing of the SLR.

Topics:

- “Connecting to a database to run the Object History tool”
- “Reviewing SLR Information”
- “Specifying additional locations to search for log files and backups” on page 97
- “Performing SLR operations” on page 98
- “Creating the SLR for a database” on page 98
- “Rebuilding the SLR” on page 100
- “Updating the SLR” on page 101
- “Pruning SLR information” on page 102
- “Dropping the SLR” on page 102
- “Reviewing the SLR status” on page 102
- “Updating the SLR from logs” on page 103

Connecting to a database to run the Object History tool

On the **Location** step in the **Object History** perspective, you can select a registered database for creating and maintaining Object History.

About this task

You might be required to log into the datastore repository database and the **SLR Information** step, depending on the connection properties that the administrator defined for the datastore repository database.

Procedure

1. To see the list of managed databases, expand the tree in the **List of databases** panel.
2. Select a recoverable database from the tree. Information about this database is displayed in the **Selected database information** panel on the right.
3. If prompted, specify a user name and password to connect to the database. If you are performing operations on a remote database, the user name and password are those for the remote database (not the local databases, and not the datastore).
4. Proceed to the **SLR Information** step. If prompted, log into the **SLR Information** step.

Reviewing SLR Information

The **SLR Information** step displays information about the current state of the Schema Level Repository (SLR) for the selected database.

About this task

If the SLR has not been created for the database, the status states that information. For databases for which the SLR has been created, you can see the current state of the SLR and its various parameters, such as consistency and time range.

Procedure

1. Log into the database if required by the datastore connection properties.
2. To get a snapshot of SLR tables without performing an SLR operation that is not required, click **Support**. If **Support** is not displayed, you can enable it in the **Preferences** perspective.
3. When you have finished reviewing the SLR Information, proceed to the **Options** step.

Specifying additional locations to search for log files and backups

On the **Options** step, you can specify additional locations to search for log files and disk backup files that are not located in the default location recorded in the DB2 database history. By default, DB2 Recovery Expert uses backups and logs that DB2 records in the Recovery History File (RHF).

About this task

Only a backup from the catalog partition will be used to create the SLR with the initial data. No backups from non-catalog partitions will be processed or accessed.

When performing operations on a remote database, you must add at least one location for logs. If the SLR has not been built or must be rebuilt, then you must also specify at least one backup location.

Procedure

1. Specify additional locations to search for backups and logs:
 - a. On the **Log directories** or **Backup directories** panel, click the plus sign icon (+), and then browse to the required location. For Database Partitioning Facility (DPF) environments, to specify a location for a specific database partition, select the partition in the **Database Partition** menu.
 - b. Click **OK**.
 - c. Complete optional tasks as needed:
 - To modify the added location, select the location and then click the edit icon.
 - To delete one or more locations, select the location and then click **Remove** or **Remove all**.
2. Proceed to the **Operations** step.

Related concepts:

“Operational prerequisites” on page 28

Review the operational prerequisites that are necessary to run DB2 Recovery Expert.

“Considerations for DPF environments” on page 171

Review considerations for using DB2 Recovery Expert in a Database Partition Facility (DPF) environment.

Performing SLR operations

On the **Operations** step, you can perform operations against a selected database.

The following SLR operations are available:

- Create
- Rebuild
- Update
- Prune
- Drop

Alternatively, you can perform SLR operations from the Command Line Processor. For more information, see the section about the CLP interface.

Related concepts:

“Object History features and benefits” on page 5

Object History offers several features and benefits that aid in the recovery process.

“Considerations for DPF environments” on page 171

Review considerations for using DB2 Recovery Expert in a Database Partition Facility (DPF) environment.

Related tasks:

“Creating the SLR for a database”

To get started using IBM DB2 Recovery Expert for Linux, UNIX and Windows, create the Schema Level Repository (SLR) on the **Object History** perspective.

“Rebuilding the SLR” on page 100

Review how to rebuild the Schema Level Repository (SLR) for any user database, if the SLR is not in a consistent state, by using the browser client.

“Updating the SLR” on page 101

Review how to update an existing Schema Level Repository (SLR) for any user database by using the browser client.

“Pruning SLR information” on page 102

Review how to prune the existing Schema Level Repository (SLR) for any user database by using the browser client.

“Dropping the SLR” on page 102

Review how to drop the Schema Level Repository (SLR) for any user database, if the SLR is not in a consistent state, by using the browser client.

“Reviewing the SLR status” on page 102

The **Status** step provides the progress and completion status of the selected operation to modify the SLR, and enables you to provide feedback regarding the operation, or cancel a pending operation.

Creating the SLR for a database

To get started using IBM DB2 Recovery Expert for Linux, UNIX and Windows, create the Schema Level Repository (SLR) on the **Object History** perspective.

About this task

- An SLR is required for most object recoveries and for using Log Analysis beyond the MRT range.
- The amount of time that is required to create the SLR can vary between a few seconds to several hours, and depends on the size of the selected database and the timeframe that the SLR span.

- The SLR is initialized using a backup of the System Catalog. The backup can be an online or offline, and full or table space level. After the SLR has been initialized with the data from the backup image, the logged operations against the System Catalog tables are processed and recorded in the SLR.
- Only a backup of the catalog partition will be used to create the SLR with the initial data. No backups from non-catalog partitions will be processed or accessed.
- When the value of the DB2_COMPATIBILITY_VECTOR registry variable is set to ORA (Oracle compatibility), warning messages might be issued that are related to tables ARYSYSTABLES, ARYSYSINDEXES, ARYSYSDATAPARTITIONS, and ARYSYSPLAN. These warnings occur because the database was created in Oracle compatibility mode, and fields in the user database differ from those in the datastore database. You can ignore these messages.
- The OPTION parameter of RESTORE is not available when reading backup images. Instead, set the VENDOROPT database configuration item, if required.

Archive logging considerations

- DB2 Recovery Expert requires archive logging to be turned on for user databases. One of the log archive method database configuration parameters (logarchmeth1 or logarchmeth2) must be set to one of the following values:
 - DISK
 - TSM if you are using the Tivoli® Storage Manager product
 - VENDOR if you are using a supported storage product other than TSM.

The SLR can be built from any full database or table space backup that contains the table space SYSCATSPACE and that was taken after roll-forward recovery was enabled for the database.

- In a Database Partitioning Facility (DPF) environment, the first backup after enabling archive logging should not be used to build the SLR.
- If archive logging is turned on or off in DB2, you must re-create the SLR when you run Log Analysis.

Procedure

1. Navigate to the **Object History** perspective and review the information that is displayed on the **Overview** step.
2. Proceed to the **Location** step and select a registered database from the **List of databases** panel.
3. On the **Operations** step, select **Create**.
Create is preselected by default for databases that do not have the SLR built. For databases that already have the SLR built, **Create** is disabled and the default operation is **Update**.
4. From the **Backup images** table, select a backup to use as a starting point for the SLR.
5. From the **End point** dropdown list, choose one of the following SLR end point ranges:
 - **End of logs:** The SLR will contain information about database objects starting from the selected backup until the end of the available logs. This option is selected by default.
 - **Specified time:** The SLR will contain information about database objects starting from the selected backup until the time that you specify.

- **Recent time:** The SLR will contain information about database objects starting from the selected backup until the **days, hours, or minutes** that you enter, as calculated by using your computer clock setting.

6. Click **Run**.

Related concepts:

“Performing SLR operations” on page 98

On the **Operations** step, you can perform operations against a selected database.

Related tasks:

“Updating the SLR” on page 101

Review how to update an existing Schema Level Repository (SLR) for any user database by using the browser client.

Rebuilding the SLR

Review how to rebuild the Schema Level Repository (SLR) for any user database, if the SLR is not in a consistent state, by using the browser client.

About this task

- The SLR must be created before it can be rebuilt.
- The amount of time that is required to rebuild the SLR can vary between a few seconds to several hours, and depends on the size of the selected database and the timeframe that is spanned by the SLR.
- The **OPTION** parameter of **RESTORE** is not available when reading backup images. Instead, set the **VENDOROPT** database configuration item, if required.

Archive logging considerations

- DB2 Recovery Expert requires archive logging to be turned on for user databases. One of the log archive method database configuration parameters (**logarchmeth1** or **logarchmeth2**) must be set to one of the following values:
 - **DISK**
 - **TSM** if you are using the Tivoli® Storage Manager product
 - **VENDOR** if you are using a supported storage product other than TSM.

The SLR can be built from any full database or table space backup that contains the table space **SYSCATSPACE** and that was taken after roll-forward recovery was enabled for the database.

- In a Database Partitioning Facility (DPF) environment, the first backup after enabling archive logging should not be used to build the SLR.
- If archive logging is turned on or off in DB2, you must re-create the SLR when you run Log Analysis.

Procedure

1. On the **Operations** step, select **Rebuild**.
2. From the **Backup images** table, select a backup to use as a starting point for the SLR.
3. From the **End point** dropdown list, choose one of the following SLR end point ranges:
 - **End of logs:** The SLR will contain information about database objects starting from the selected backup until the end of the available logs. This option is selected by default.
 - **Specified time:** The SLR will contain information about database objects starting from the selected backup until the time that you specify.

- **Recent time:** The SLR will contain information about database objects starting from the selected backup until the **days, hours, or minutes** that you enter, as calculated by using your computer clock setting.

4. Click **Run**.

Related concepts:

“Performing SLR operations” on page 98

On the **Operations** step, you can perform operations against a selected database.

Updating the SLR

Review how to update an existing Schema Level Repository (SLR) for any user database by using the browser client.

About this task

- The SLR must be created before it can be updated.
- Updating the SLR processes the logs starting from the last update log position. The operations against the System Catalog tables are recorded in the SLR.
- Keep the SLR in an updated state that covers the range of the most valuable activity so that Log Analysis or Recovery can be performed when it is required.
- The amount of time that is required to update the SLR can vary between a few seconds to several hours, and depends on the size of the selected database and the timeframe that is spanned by the SLR.

Procedure

1. On the **Operations** step, select **Update**. **Update** is selected by default if the SLR is already built and is in a consistent state.
2. Choose one of the following SLR end points:
 - **End of logs** updates an SLR that contains information about database objects starting from the selected backup until the end of the available logs.
By default, the end point is set to the end of logs; therefore, it covers all of the existing activity from the point of the last SLR update. If required, you can set the SLR end point to a specific time, specific LSN/LRI/LSO, or a recent time.
 - **Timestamp** updates an SLR that contains information about database objects starting from the selected backup until the timestamp that you specify.

Note: Log file and LSN/LRI/LSO values are only applicable to the catalog partition.

LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

3. Click **Run**. The SLR process begins, and the **Status** step becomes active.

Related concepts:

“Performing SLR operations” on page 98

On the **Operations** step, you can perform operations against a selected database.

Related tasks:

“Creating the SLR for a database” on page 98

To get started using IBM DB2 Recovery Expert for Linux, UNIX and Windows, create the Schema Level Repository (SLR) on the **Object History** perspective.

Pruning SLR information

Review how to prune the existing Schema Level Repository (SLR) for any user database by using the browser client.

About this task

- The SLR for the database must be created and consistent before you can perform prune operations.
- Pruning the SLR removes old data that is no longer required. You can prune the SLR to control the SLR range. When the old history is no longer needed, you can adjust the starting point of the SLR by issuing the SLR prune command.

Procedure

1. On the **Operations** step, select **Prune**.
2. Enter a timestamp.
3. Click **Run**. The SLR process begins, and the **Status** step becomes active.

Related concepts:

“Performing SLR operations” on page 98

On the **Operations** step, you can perform operations against a selected database.

Dropping the SLR

Review how to drop the Schema Level Repository (SLR) for any user database, if the SLR is not in a consistent state, by using the browser client.

About this task

- If the SLR is not built, you cannot drop the SLR.
- Dropping the SLR might be required if datastore repository database maintenance must be performed, or if the SLR for a specific database is no longer required because the database is no longer used.
- This operation permanently deletes all SLR definitions, data, and history for the specified database.
- If the SLR was previously built for any user database, and that database is later restored from a backup, then the SLR for the database is no longer valid. The SLR must be rebuilt from a new backup. If the database was restored to the same name, you can create the SLR from the same backup that was used for the restore. Otherwise, you must create the SLR from a new backup that was taken after the restore was performed.

Procedure

1. On the **Operations** step, select **Drop**.
2. Click **Run**. The SLR process begins, and the **Status** step becomes active.

Related concepts:

“Performing SLR operations” on page 98

On the **Operations** step, you can perform operations against a selected database.

Reviewing the SLR status

The **Status** step provides the progress and completion status of the selected operation to modify the SLR, and enables you to provide feedback regarding the operation, or cancel a pending operation.

About this task

- If the SLR was previously built for any user database, and that database is later restored from a backup, then the SLR for the database is no longer valid. The SLR must be rebuilt from a new backup. If the database was restored to the same name, you can create the SLR from the same backup that was used for the restore. Otherwise, you must create the SLR from a new backup that was taken after the restore was performed.
- The **Status** step is the same throughout the product for **Log Analysis**, **Recovery**, and **Object History** operations.
- The header area of the window displays the following information:
 - Operation: The operation that you are currently performing.
 - Session ID: The ID of the submitted task.
 - Start time and end time of the operation.
 - Status: The progress of the task, the result of the task, and whether the operation was successful.
- While the operation is in progress, a time indicator is displayed near the lower left corner.
- The **Status** area displays the time, database partition, and messages related to the operation.

Procedure

You can perform the following optional tasks:

1. To cancel any pending operation, click **Cancel**.
2. To choose which kinds of error messages are displayed, from the dropdown list, select one of the following options:
 - **All messages**
 - **Warnings and errors**
 - **Errors only**
3. To copy messages to a clipboard so that you can paste them into a text editor, click the copy icon.
4. To download a compressed file that contains the messages, and either view the compressed file directly, or save the compressed file to a local directory, click the download icon.
5. To create a detailed trace that can be used to analyze the operation and provide detailed information for Software Support:
 - a. Click **Support**. If **Support** is not displayed, you can enable it in the **Preferences** perspective.
 - b. Select options and then click **OK**. The **Download trace file** window opens. You can view the file and save it.

Related concepts:

“Performing SLR operations” on page 98

On the **Operations** step, you can perform operations against a selected database.

Updating the SLR from logs

You can bring the information in the Schema Level Repository (SLR) up to date by processing the recent database recovery logs.

About this task

The SLR information is initialized from a database backup and is updated from the recovery logs to

- provide a continuous record of database history and catalog changes.
- allow recoveries to be performed to any point in time, including before the time that the product was installed.

Typically, the SLR should be built and kept up-to-date for all of the user databases that need to be processed in the future. Each user database has its own SLR, but all of them are stored in the data store databases that are created within the same DB2 instance.

The SLR does not need to be updated up to the current time if that is not required to process the requested log range. This feature eliminates extra log processing.

To eliminate the need for the SLR update process to read backup images to process masked update log records, enable DATA CAPTURE CHANGES (DCC) for all System Catalog tables.

Log Analysis considerations:

- If needed, Log Analysis can be used immediately, without SLR initialization, to report archive and active log activity that occurred within the Minimum Recovery Time (MRT) range. Reporting before the MRT time frame requires initializing the SLR to provide historical information about database objects.
- Log Analysis determines the set of required log files from the SLR time frame in combination with user-supplied report parameters.

Procedure

Update the SLR by processing recent database recovery logs. This procedure is similar to rebuilding the SLR. For more information, see “Rebuilding the SLR” on page 100.

1. Select an existing user database.
2. Select an existing backup that contains system catalog information (for example, SYSIBM tables) to represent the start point of the SLR.
3. Choose the end point of the SLR.
4. Click **Run**.

Chapter 7. Log Analysis

On the **Log Analysis** perspective of the browser client, you can generate reports, perform Redo or Undo SQL, apply filters and options for Log Analysis processing, and work with reports.

Log Analysis prerequisites and considerations

- Before performing Log Analysis operations, review the following information:
 - (optional) Add a datastore repository by using the **Datastores** perspective.
 - Add a managed database by using the **Managed Databases** perspective.
- Though not required for all Log Analysis operations, some operations require that you first create the Schema Level Repository (SLR) for the database on which you want to use the Log Analysis tool. For example, you must create the SLR if you want to use Log Analysis to report on the log activity that is beyond the minimum recovery time (MRT) range.
- When backup images are located in an IBM Tivoli Storage Manager (TSM) or vendor storage manager, the VENDOROPT database configuration item must be used to pass an OPTION value for reading backup images, if required.

Log Analysis overview

The typical workflow for a Log Analysis operation is as follows:

1. Access the **Log Analysis** perspective: from the **Open** click **Log Analysis**.
2. Select an existing user database on which to run the Log Analysis tool.
3. Select the type of report (**Summary**, **Detail**, or **Quiet Time**), the database partitions to run against, the start and end time of the log activity, and provide filters and a set of objects to be included in the report.
4. Click **Run**.

You use the following steps in the **Log Analysis** perspective to perform Log Analysis operations:

- **Overview** step. Review a description of the operations that can be performed on the **Log Analysis** perspective.
- **Location** step. Review the registered databases and detailed information about the selected database.
- **Details** step. Specify the report type and mode, the time range that should be covered by the report, whether to generate Redo or Undo SQL, and the database partitions on which to perform Log Analysis.
- **Objects** step. Filter and only produce the report for objects that will be selected in this step.
- **Filters** step. Specify operations to perform (combinations of inserts, updates, and deletes), the type of transactions that should be reported, and the additional filters necessary for the report output.
- **Options** step. Review additional log directories and backup locations for Log Analysis. This step is optional, and is not necessary in most cases. The options can be useful when backups and logs are moved to another directory.

You can set up some performance options for the detailed phases of Log Analysis, such as specifying whether LONG or LOB data should be processed,

whether access to backups is allowed, and whether Log Analysis can access the current DB2 pages to undo some of the operations.

- **Status** step. Review information about the current processing of Log Analysis.
- **Result** step. After the Log Analysis task completes successfully, review Log Analysis statistics, and compile the generated report for further review. You can also provide post filters and choose a different order to create report activity.

Topics:

- “Connecting to a database to run Log Analysis”
- “Providing report details”
- “Selecting objects for Log Analysis reporting” on page 108
- “Providing filters for Log Analysis” on page 111
- “Specifying additional options for Log Analysis processing” on page 112
- “Reviewing Log Analysis status” on page 113
- “Working with Log Analysis results” on page 114
- “Configuring DB2 Recovery Expert for Log Analysis and remote database processing” on page 78

Connecting to a database to run Log Analysis

On the **Location** step in the **Log Analysis** perspective, you can select a local or remote database on which to perform Log Analysis operations.

About this task

You might be required to log into the database, depending on the connection properties that the administrator defined for the datastore connection.

Procedure

1. To see the existing systems, instances, and databases, expand the tree in the **List of databases** panel.
2. Select a recoverable database from the tree.
Information about the database is displayed in the **Selected database information** panel in the **Parameter** and **Value** columns.
3. If prompted, specify a user name and password to connect to the database. If you are performing remote Log Analysis, the user name and password are those for the remote database (not the local databases, and not the datastore).
4. Proceed to the **Details** step.

Related tasks:

“Working with Log Analysis results” on page 114

The **Result** step consists of the **Statistics**, **Post filters**, and **Report** views. You can switch between the sub-views by clicking the appropriate title in the top bar.

“Configuring DB2 Recovery Expert for Log Analysis and remote database processing” on page 78

Review how to set up DB2 Recovery Expert for Log Analysis and remote database processing.

Providing report details

On the **Details** step, you can specify several settings that affect the Log Analysis output.

Procedure

1. Report type: Choose from the following report types:

- **Summary:** Contains only the list of transactions and a list of operations.
- **Detail:** Includes details about all Log Analysis operations and transactions.
- **Quiet time:** Displays periods during which no active transactions were modifying a set of table spaces and tables on the selected database.

For the most accurate results, activate application information logging. For more information, see the section on capturing AUTHID, application name, and ID information in the operational prerequisites topic.

2. Report mode: Choose from the following report modes:

- **Available log range activity:** If the SLR has been created, choose this mode to indicate that the SLR will be used. (If the SLR has not been created, this mode is not available.)
- **Minimum Recovery Timestamp activity:** If the SLR has *not* been created or is not to be used, choose this mode.

Note that the minimum recovery timestamp value is unique for every table space on every database partition. When you report on multiple table spaces in **Minimum Recovery Timestamp** mode, Log Analysis uses the most recent minimum recovery timestamp value on each database partition for all of the table spaces that are included in the report. This action can lead to omission from the report of activity for one table space because the activity occurred before the minimum recovery timestamp value of another table space. To avoid this outcome, specify a timestamp for the report in the specific table space to include in the report.

3. Type of SQL:

After the report has been generated, you can specify options to control the SQL output:

- **None:** No SQL is generated during the report.
- **Undo SQL:** Available for **Detail** report type only. Produces SQL statements to reverse the changes made by the original database log operations.
- **Redo SQL:** Available for **Detail** report type only. Produces SQL statements to reapply the changes made by the original database log operations, and optionally includes DDL statements to replicate changes that were made to the database catalog.

4. Start and end points of the report range:

You can specify an exact time or a relative time for which the Log Analysis report should be generated:

- **Earliest available:** Uses the SLR start range.
- **Specified time:** Uses a start or end time that you enter.
- **Recent time:** Specifies a relative time that is calculated using your computer clock setting.
- **Specified LSN/LSO:** Available for single-database-partition analysis only. Specifies a log sequence number (LSN) or log stream offset (LSO) value. The LSN or LSO value is a 16-character hexadecimal number that represents the byte offset of the log record from the beginning of the database log.
- **Specified LRI:** Available when running DB2 10.1 and later. Specifies a log record identifier. The LRI value is a 34-character hexadecimal number. The first two characters indicate the type of log record, the next 16 characters indicate the unique log flush sequence number (LFS), and the final 16 characters indicate the associated LSN.

Note: DB2 Recovery Expert supports only type 1 LRIs.

- **Latest available:** Uses the current time as the end point.

For more information about timestamps, see the topic about operational prerequisites.

5. **Log files range:** Specify exact boundaries for Log Analysis processing. Log Analysis will not include log files from outside the range that you specify. This option is only available for non-DPF databases.
6. **Database Partitions to launch Log Analysis on:** This option is available in Database Partitioning Facility (DPF) environments only. Choose from the following options:
 - **All database partitions:** Perform Log Analysis across all partitions.
 - **Catalog partition only:** Perform Log Analysis on the catalog partition only.
 - **Selected database partitions below:** Perform Log Analysis on partitions that you select from the list. To select multiple database partitions, press and hold **Ctrl** and click each partition.
7. Proceed to the **Objects** step.

Related concepts:

“Operational prerequisites” on page 28

Review the operational prerequisites that are necessary to run DB2 Recovery Expert.

Related tasks:

“Working with Log Analysis results” on page 114

The **Result** step consists of the **Statistics**, **Post filters**, and **Report** views. You can switch between the sub-views by clicking the appropriate title in the top bar.

Selecting objects for Log Analysis reporting

On the **Objects** step, select objects for which Log Analysis will provide report information. By default, Log Analysis processes all of the objects in the list. To override the default, select objects for Log Analysis to process or to exclude from processing.

About this task

- You can specify filters for the following objects for Log Analysis:
 - Database partition groups
 - Schemas
 - Tables
 - Table spaces
- Selecting objects for processing:

You drill down to the object that you want to select. The objects appear in different colors based on where the object exists:

 - **Black:** In both the Schema Level Repository (SLR) and the DB2 system catalog.
 - **Red:** In the SLR only (for example, when the object is dropped).
 - **Blue:** In the DB2 system catalog only.

DB2 Recovery Expert internal objects are not available for selection when the datastore is installed in the user database.
- DB2 Recovery Expert does not display views, MQT, and aliases in the table node of the object tree.

- Changing the DB2 objects (display) filters does not affect the objects that appear in the **Selected objects** list. If objects are later removed from the **Selected objects** list, then their appearance in the **Available objects** list depends on whether the current filter restricts them from being displayed.
- If you do not select an object, and you proceed to another step in the **Log Analysis** perspective, Log Analysis processes *all objects* when generating the report.
- For objects that have schemas, DB2 Recovery Expert has two levels of hierarchy, as shown in the following example:

```
Tables
  -Schema1
    Table1
    Table2
  +Schema2
  -Schema3
    Table3
    Table4
```

- You can hide the following catalog and tool objects:
 - Table spaces matching pattern SYS% and ARY%.
 - Schema matching pattern ARY%.

Procedure

1. From the **Available objects** panel, select one or more objects or patterns, and then click **Include** or click **Exclude**.
2. In the **Pattern value** field, specify the filtering pattern to be used to identify objects.
3. To filter the list of available objects:
 - a. Click **Object filters**.
 - b. Enter the filtering pattern to be used to identify those objects. The syntax is similar to the DB2 LIKE clause; use % and _ symbols to specify required filters. To include all of the currently available DB2 objects, enter the default filter % or leave the field blank.

To limit the number of objects that are displayed in a branch, you can filter each level by schema and then by name. The schema and the name filters are independent; that is, all tables under each schema are filtered with the name filter. Specify a filter value in following format: *schema_name.object_name*. If you do not specify either a schema or an object filter value, DB2 Recovery Expert uses the wildcard value "%". The following examples illustrate schema and object filtering.

Table 19. Schema and object filtering include examples

Operation	Filter value	Schema filter	Object filter	Explanation
Include	(no value)	%	%	Include all objects in all schemas.
Include	A%	%	A%	Include all objects starting with "A" in all schemas.

Table 19. Schema and object filtering include examples (continued)

Operation	Filter value	Schema filter	Object filter	Explanation
Include	D%.A%	D%	A%	Include all objects starting with "A" where the schema name begins with "D".

Table 20. Schema and object filtering exclude examples

Operation	Filter value	Schema filter	Object filter	Explanation
Exclude	D%.%	D%	NOT %	Exclude all objects where the schema name begins with "D".
Exclude	%	%	NOT %	Exclude all objects in all schemas.
Exclude	M%.A%	M%	NOT A%	Exclude all objects starting with "A" where the schema name begins with "M".
Exclude	D97A1.C%	D97A1	NOT C%	Exclude all objects starting with "C" where the schema name is "D97A1".

For more information about filtering objects, see the tips in "About this task."

c. Click **OK**.

Excluded objects are identified in the **Available objects** list.

4. To view further details, in **Object properties** in the **Selected objects** panel, click the object. You can see whether the object exists only in the SLR, only in the database, and whether the object definition between the SLR and its current state is different.
5. Click **OK**.
6. Proceed with the **Filters** step.

Related tasks:

"Working with Log Analysis results" on page 114

The **Result** step consists of the **Statistics**, **Post filters**, and **Report** views. You can switch between the sub-views by clicking the appropriate title in the top bar.

Related reference:

"Tips for using the browser client" on page 91

This topic provides basic how-to information for new Recovery Expert users and tips for experienced users.

Providing filters for Log Analysis

On the **Filters** step, you can specify filters for Log Analysis reports.

About this task

Log Analysis provides the following predefined filters:

- **Operations** to process all operations (inserts, updates and deletes).
- **Transactions** to process only committed transactions.
- **Exclude System and Tool objects** to filter out SYS% and ARY% objects.

Restriction: With column-organized tables, a system table that is not in a system table space is not excluded.

- **Additional filters** to filter by Application Name, Auth ID, object ID, and so on.

The report for processing partial, uncommitted, and rolled backed transactions contains only summary information. Log Analysis cannot generate details on such transactions.

You can modify the existing filters and also add a new set of filters to run on Log Analysis.

Procedure

To set additional filters for the activity in the Log Analysis report:

1. In the **Operations** panel, choose one or more operations (Inserts, Updates, or Deletes).
2. In the **Transactions** panel, choose one or more transactions. The report for processing partial, uncommitted, and rolled backed transactions contains summary information only. Log Analysis cannot generate details on such transactions.
 - (Required) **Committed** transactions are those for which Log Analysis can locate a starting log record and an ending commit log record within the set of logs that are processed. Log Analysis always reports committed transactions.
 - **Partial** transactions are those for which Log Analysis can find an ending commit log record, but cannot locate the start of the transaction within the set of logs that are processed.
 - **Uncommitted** transactions are partial or full transactions for which Log Analysis cannot locate an ending commit log record within the set of logs that are processed.
 - **Rolled back** transactions are uncommitted transactions that were canceled and rolled back to their starting point. This option is available in DB2 Recovery Expert for Linux, UNIX, and Windows Version 3.1 and later. In earlier versions, Log Analysis reported rolled back transactions.
3. In the **Additional filters** area:
 - a. In the **Operation** column, choose whether to include or exclude types that match the filter. For the following filter types, you can specify wildcards by selecting **Include like** or **Exclude like**:
 - **Application name**
 - **Application IDs**
 - **Auth ID**

- b. From the **Type** dropdown list, select an additional filter type to filter log activity by the value that you specify in the **Value** column.
 - c. In the **Value** column, specify the name or ID by which to filter log activity.

If log records contain mixed activity, for example, if some of the log records have application information and some do not, the **Application names** (appnames), **Application IDs** (appids), and **Auth IDs** (authids) filters can be set to {empty} where application information was not set.

For **Transactions ID**, specify a transaction ID or range of transaction IDs using the format ID1..ID2. A transaction ID is a hexadecimal value, as follows:

 - Local Transaction IDs have a length between 1 and 12 hexadecimal characters (0-9, A-F).
 - Global Transaction IDs have a length between 13 and 40 hexadecimal characters.
 - XA Transaction IDs have a length between 41 and 140 hexadecimal characters.

For **Table ID**, specify TID:FID pairs, which identify tables.
 - d. Complete the following optional steps as needed:
 - To add new filters to the **Additional filters** panel, click the plus sign icon (+).
 - To delete filters from the **Additional filters** panel, select a filter and click the remove or remove all icon.
4. Click **Next** to proceed to the **Options** step.

Related tasks:

“Working with Log Analysis results” on page 114

The **Result** step consists of the **Statistics**, **Post filters**, and **Report** views. You can switch between the sub-views by clicking the appropriate title in the top bar.

Related reference:

“Tips for using the browser client” on page 91

This topic provides basic how-to information for new Recovery Expert users and tips for experienced users.

Specifying additional options for Log Analysis processing

On the **Options** step, you can specify additional options for Log Analysis processing, such as detailed report processing options, and disk backups and logs that are not located in the default location recorded in the DB2 database history. By default, DB2 Recovery Expert uses backups and logs that DB2 records in the Recovery History File (RHF).

About this task

The amount of time that is required to run a Log Analysis can vary from a few seconds to several hours, and depends on several factors, such as the size of the selected database, how many DB2 objects are included within the Log Analysis report, and the kinds of transactions and operations to be included in the report.

When performing operations on a remote database, if you selected the Summary report, you must add at least one backup location. For the Detail report, you must specify at least one backup location unless you check **Don't access current database when processing masked updates**.

To enable the **Detailed report processing options** panel of the **Options** step, you must set the report type to **Detail** in the **Details** step.

Procedure

1. On the **Detailed report processing options** panel, set the following options, as needed:
 - When you do not need Log Analysis to report LONG and LOB data, check the **Skip processing for LONG and LOB column data** check box.
 - When your database has DATA CAPTURE CHANGES (DCC) enabled or your isolation level is set to currently committed, improve performance by checking the **Don't access backups when processing masked updates** check box.
 - To skip accessing current databases pages when performing an undo from the current state, check the **Don't access current database when processing masked updates** check box.
2. Specify additional locations to search for backups and logs:
 - a. On the **Log directories** or **Backup directories** panel, click the plus sign icon (+), and then browse to the required location. For Database Partitioning Facility (DPF) environments, to specify a location for a specific database partition, select the partition in the **Database Partition** menu.
 - b. Click **OK**.
 - c. Complete optional tasks as needed:
 - To modify the added location, select the location and then click the edit icon.
 - To delete one or more locations, select the location and then click **Remove** or **Remove all**.
3. Click **Save**.
4. Click **Run** or **Next**.

Related concepts:

“Considerations for DPF environments” on page 171

Review considerations for using DB2 Recovery Expert in a Database Partition Facility (DPF) environment.

Related tasks:

“Working with Log Analysis results” on page 114

The **Result** step consists of the **Statistics**, **Post filters**, and **Report** views. You can switch between the sub-views by clicking the appropriate title in the top bar.

Reviewing Log Analysis status

When all of the required options and filters are specified, you can submit a Log Analysis task for execution.

About this task

The amount of time that is required to run a Log Analysis task can vary between a few seconds to several hours depending on the size of the selected database, how many DB2 objects are included within the Log Analysis report, the kinds of transactions and operations to be included in the report, in addition to other factors.

When the Log Analysis task is submitted, you are automatically redirected to the **Status** step.

- The **Status** step is the same throughout the product for **Log Analysis**, **Recovery**, and **Object History** operations.
- The header area of the window displays the following information:
 - Operation: The operation that you are currently performing.
 - Session ID: The ID of the submitted task.
 - Start time and end time of the operation.
 - Status: The progress of the task, the result of the task, and whether the operation was successful.
- While the operation is in progress, a time indicator is displayed near the lower left corner.
- The **Status** area displays the time, database partition, and messages related to the operation.

Procedure

You can perform the following optional tasks:

1. To cancel any pending operation, click **Cancel**.
2. To choose which kinds of error messages are displayed, from the dropdown list, select one of the following options:
 - **All messages**
 - **Warnings and errors**
 - **Errors only**
3. To copy messages to a clipboard so that you can paste them into a text editor, click the copy icon.
4. To download a compressed file that contains the messages, and either view the compressed file directly, or save the compressed file to a local directory, click the download icon.
5. To create a detailed trace that can be used to analyze the operation and provide detailed information for Software Support:
 - a. Click **Support**. If **Support** is not displayed, you can enable it in the **Preferences** perspective.
 - b. Select options and then click **OK**. The **Download trace file** window opens. You can view the file and save it.

Working with Log Analysis results

The **Result** step consists of the **Statistics**, **Post filters**, and **Report** views. You can switch between the sub-views by clicking the appropriate title in the top bar.

About this task

On the Log Analysis **Result** step, you can create a detailed trace for troubleshooting, review statistics, apply filters to the Log Analysis report, generate the report, and apply and run Redo or Undo SQL that was generated by Log Analysis.

Procedure

1. To create a detailed trace that can be used to analyze the operation and provide detailed information for Software Support, click **Support**. If **Support** is not displayed, you can enable it in the **Preferences** perspective.

Statistics view

Review the overall statistics and a graph of Log Analysis activity against your database for a given time frame, for objects, and for report filters.

Note: If different versions of a table with the same name have existed at different times during the life of a database, the **Inserts**, **Updates**, and **Deletes** statistics for each version of the table that were found during a Log Analysis report are combined for a single table name entry.

For example, you cannot view the first 10 transactions of the first version of Table X, nor the first 25 transactions of a different version of Table X that existed at a later point in time. The output in the **Statistics** panel displays only 35 actions for Table X, because Log Analysis does not track the statistics separately for different versions of a table with the same name.

This issue occurs only with the **Statistics** panel. The detailed reports and SQL correctly associate activities of different versions of objects with the same name.

2. Perform one or more of the following actions:

- Switch the current statistics information by clicking the appropriate grouping title.
- Select multiple items to see activity represented in the graph objects by holding **Shift** and clicking the objects.
- Suppress or activate totals in the graph by clicking the **Show totals** check box.
- Activate stack display by clicking the **Stacked** check box.

Post filters view

Apply filters to the Log Analysis report to see only a selected subset of the information in the **Details** step. You can select different kinds of post filters at the same time.

3. Perform one or more of the following actions:

- Specify start and end points for the report range.
- Reset the time input to the start or end point by clicking the blue reset icon.
- Set the quiet time threshold to limit reporting to quiet times of a minimum length. The default value, 00:30:00, indicates that only quiet times of 30 minutes or longer are included in the Quiet time report.
- Use the **Selected report post filters** list to select objects. By default, if the **Selected report post filters** list contains no objects, then the Log Analysis report includes all objects. When you select different types of objects in the **Selected report post filters** list, those objects are applied using the Boolean operator OR within the post filter, and the Log Analysis report shows the result.
- Create column-level post filters for a selected table. For more information, see “Creating column-level filters” on page 118:

Report view

Review transaction results for the report. Displayed information depends on the selections that you made on the **Details** step:

- If you chose a Quiet Time report, the view displays periods during which no active transactions were modifying the selected object.
- If you chose a Detail report, log operation details are included.

- If you requested SQL generation, Redo or Undo SQL details are included.
 - The **Transactions** table displays information according to the selection you make in the grouping control to the left of the table.
4. Perform one or more of the following actions:
 - Specify the information that is displayed in the **Transactions** table by clicking the **Customize columns** icon.
 - Display a list of the related operations for a selected transaction.
 - Modify the information that is displayed in the **Transaction columns** and the **Operation columns** by clicking **Display**. These columns become unavailable when appropriate information is not present in the report.
 - Specify whether details are displayed vertically or horizontally by clicking the **Show details vertically** check box. Showing details horizontally reduces the performance of the Flash player; however, it is not recommended if the column count is greater than 10.
 - When a transaction is selected and operations are shown in the **Operations** table, change the display position in the list of available data by using the scroll bars and navigational buttons.
 - When operations are shown in the **Operations** table, display transaction information that occurred before and after the operation was executed by selecting an operation and clicking **Details**.
 - When operation details are displayed, retrieve and save values for data types that are too large to display directly in the browser client by clicking **Download**. LOB, LONG, GRAPHIC, VARGRAPHIC, VARCHAR and CHAR data types might be too large to display directly in the browser client. Column values up to 1500 bytes are displayed in place.
 - If Redo SQL or Undo SQL was generated as part of the report, display the SQL details for the operation by clicking **Redo SQL** or **Undo SQL**.
 5. To set properties for the report that you generate, click **Export > Export report**.
 6. To set properties for the SQL that you export, click **Export > Export SQL**.
 7. To generate the report, click **Generate**.

Note: If you are using Internet Explorer 8 and the export fails, you might need to enable automatic file downloads.

8. On the **Download report** window, you can perform the following actions:
 - Copy the web address at which you can view the report by clicking **Copy URL**.
 - Choose to open or save a zip file that contains the report by clicking **Download**.
 - Close the window by clicking **Close**.

The code set of the exported report or SQL file is UTF-8. Comments in the report provide the following information:

- Connection information that was specified in the **Location** step.
 - Report range options that were specified in the **Details** step.
 - Objects and filters that were specified in the **Objects** step and the **Filters** step.
 - Additional options that were specified in the **Options** step.
9. *Before* applying Redo or Undo SQL that was generated by Log Analysis, review the following precautions and considerations:
 - When you click **Run SQL**, the SQL is executed against the *current* database.

- Ensure that the selected SQL is valid to perform against the current database tables in their current state.
 - Performing Redo or Undo SQL is not the same as performing a recovery. A recovery takes into account the state of the database at the selected recovery time, including the interaction of related database objects. The Redo or Undo SQL that Log Analysis generates represents database activity that occurred at a specific point in time; however, preceding or subsequent database activity might affect the performance of the generated SQL statements.
 - DDL operations are not included in the Run SQL processing. The SQL is run in the current database state. Any subsequent object definition or data changes in other transactions might affect the validity of the generated SQL in the current database.
 - The SQL to be performed contains transactions that you selected on the **Report details** view before you clicked **Apply SQL**. You can select a single transaction, a set of transactions, or all available transactions in the report. You cannot select specific operations; only the SQL for complete transactions as generated by the report can be performed. These transactions might not represent the complete original transactions if you selected filters before generating the report. For example, it is possible to choose to generate a report only for DELETE operations. In this case, the report and generated transaction SQL will not include other operations (INSERTs or UPDATEs) that were part of the original transactions.
 - The **Post Filter** settings on the Log Analysis **Results** view do not affect the Redo or Undo SQL that will be performed by **Run SQL** processing. Post filters only affect the report details as displayed in the browser client or exported to a file. The SQL that is applied by the **Run SQL** option represents the complete transaction contents as generated by the Log Analysis report settings.
10. To specify settings for running the generated SQL, click **Run SQL**.
 - To perform the SQL using the original commit scope as read from the DB2 logs, choose **Preserve commit scope**.
 - To enter a user specified commit scope value, choose **Override commit scope**, and then enter the number of operations to perform between commit statements.
 11. To perform the Redo or Undo SQL, click **Run**, or to return to the Log Analysis **Report** view, click **Cancel**.

Related tasks:

“Customizing view preferences” on page 80

Customize views throughout the browser client by using the **Preferences** perspective.

“Connecting to a database to run Log Analysis” on page 106

On the **Location** step in the **Log Analysis** perspective, you can select a local or remote database on which to perform Log Analysis operations.

“Providing report details” on page 106

On the **Details** step, you can specify several settings that affect the Log Analysis output.

“Selecting objects for Log Analysis reporting” on page 108

On the **Objects** step, select objects for which Log Analysis will provide report information. By default, Log Analysis processes all of the objects in the list. To override the default, select objects for Log Analysis to process or to exclude from processing.

“Providing filters for Log Analysis” on page 111

On the **Filters** step, you can specify filters for Log Analysis reports.

“Specifying additional options for Log Analysis processing” on page 112

On the **Options** step, you can specify additional options for Log Analysis processing, such as detailed report processing options, and disk backups and logs that are not located in the default location recorded in the DB2 database history. By default, DB2 Recovery Expert uses backups and logs that DB2 records in the Recovery History File (RHF).

Related reference:

“Log Analysis Result step field definitions” on page 119

Review results of the Log Analysis operation on the **Result** step.

Creating column-level filters

On the Log Analysis **Result** step **Post Filters** view, you can define advanced filters for report type **Detail** to allow specification of filtering down to the column level.

About this task

Column-level filters are available for the **Detail** report type. You select the report type on the **Details** step.

The **Column filter for table <tableName>** window enables you to include one or more column post filters for a table that you choose.

Correct ordering of defined predicates is crucial to building the correct column post filter. You can reorder predicates in the **SQL column predicates** table by selecting a row and clicking the arrow icons to move the row up or down.

The **SQL output for column post filter** section is a read-only text area that displays the full SQL statement for the currently defined column post filter that will be applied when the Log Analysis report is constructed.

Column post filters are included in exported Log Analysis reports, and post-filter settings are included in saved sessions and specifications.

The following restrictions apply to column post filters:

- You cannot filter columns if the column data type changed over the Log Analysis range.
- (*DB2 version 9.5 only*) REAL and DOUBLE data type columns are not supported.

Procedure

To create a column-level filter:

1. In the **Available report post filters** area, expand the **Columns** root node to display a list of individual tables.
2. Select the table with which you want to work, and expand it to display the **Column filter** node.
3. Drag the **Column filter** node to the **Selected report post filters** area. The **Column filter for table <tableName>** window opens.
4. Click the plus sign (+) icon to add a predicate row to the **SQL column predicates** area.
5. In the row for the predicate that you want to define, choose or enter the following information:

- Left and right parentheses.
- The name of the column in the table that you selected.
- The SQL operator.

Choosing either the **Changed** or **Unchanged** operator means that you want to filter all rows in which the specified column value has changed or has not changed, respectively. Therefore, **Column Value** is not available for these operators.

- The value in the column.
- The condition.

6. Click OK.

Example

To define the filter ((A>a OR B<b) AND (C=c OR D Like d)) OR (E<= e AND F>=f) for a table, list the following predicates in the order shown:

```
Left Parenthesis: ((, Name: A, Operator: Greater than, Value: a, Right Parenthesis: ), Condition: OR
Left Parenthesis: (, Name: B, Operator: Less than, Value: b, Right Parenthesis: ), Condition: AND
Left Parenthesis: (, Name: C, Operator: Equals, Value: c, Right Parenthesis: ), Condition: OR
Left Parenthesis: (, Name: D, Operator: Like, Value: d, Right Parenthesis: )), Condition: OR
Left Parenthesis: (, Name: E, Operator: Less than or equal, Value: e, Right Parenthesis: ), Condition: AND
Left Parenthesis: (, Name: F, Operator: Greater than or equal, Value: f, Right Parenthesis: ), Condition:
```

What to do next

Optionally complete the following tasks in the **Selected report post filters** panel:

- View column post filters for a selected table.
- Modify an existing column post filter: Click the edit icon.
- Delete an existing column post filter: Click the remove icon.

Log Analysis Result step field definitions

Review results of the Log Analysis operation on the **Result** step.

Table 21. Alphabetical list of Log Analysis Result step fields and descriptions

Field	Description
Action	Indicates the specified operation (inserts, deletes, or updates).
APPID List	List of application IDs that were used for report filtering.
Application ID	Groups the activity for all application IDs. (This group is only shown when Log Analysis has run and the information about application IDs was present.)
Application name	Groups the activity for all application names. (This group is only shown when Log Analysis has run and the information about application names was present.)
Apply post filters	Apply post filters when generating the SQL output.
APPNAME List	List of application names that were used for report filtering.
Auth ID	Groups the activity for all authorization IDs. (This group is only shown when Log Analysis has run and the information about authorization IDs was present.)
AUTHID List	List of authorization IDs that were used for report filtering.
Backup Directory	List of additional backup directories that were specified.
Client codepage	The code page value that was used by a DB2 application that produced a transaction in the logs, with the application log record that was processed by the Log Analysis.

Table 21. Alphabetical list of Log Analysis Result step fields and descriptions (continued)

Field	Description
Client DB alias	The client database alias name as it was recorded in the application log record that was processed by Log Analysis.
Commit scope	Specify how many operations should be performed between commits. Commit statements will be inserted into the SQL output at the specified interval. The Commit scope defaults to 0, which uses the original transaction commit point.
Database	Database on which the report was performed.
Database Partition	Number of the database partition on which the operation was executed.
Database Partition Groups	Displays the SQL activity that is associated with the appropriate database partition groups. On the report, lists the included database partition groups as specified in Database Partitions to launch Log Analysis on .
Database Partition List	Database partitions on which Log Analysis was performed.
datastore repository	Name of the datastore repository.
Date	Display transactions according to the date. On the report, the date and time of the transaction commit log record.
Display SQL	Display the SQL statements that are performed.
End LSN/LSO/LRI	Ending log sequence number or log stream offset of the report range, if specified.
End date	End date for the transaction period.
End time	End time for the transaction period.
End timestamp	End date and time for the transaction period.
Global transaction ID	The transaction ID that is unique across all database partitions in the database; it embeds the local transaction ID value.
Include DDL	Applies to Redo SQL only. Log Analysis generated Undo SQL does not provide the ability to undo database catalog changes. The Run SQL action does not perform DDL operations.
Include LONG	Include LONG and LOB data in the SQL output.
Include SQL comments	Include comments in the SQL output.
Local transaction ID	Displays transaction IDs.
Log directory	List of additional log directories that were specified.
Logbounds	Starting and ending log numbers, as specified in Log files range .
LRI	Log record identifier for the operation.
LSN/LSO	Log sequence number or log stream offset for the operation.
MRT	Displayed if the chosen report mode was Minimum Recovery Timestamp activity .
NOBACKUP	Displayed if Don't access backups when processing masked updates was selected.
NOLONG	Displayed if Skip processing for LONG and LOB column data was selected.

Table 21. Alphabetical list of Log Analysis Result step fields and descriptions (continued)

Field	Description
NOPAGE	Displayed if Don't access current database when processing masked updates was selected.
Operations	Displays the number of Inserts, Updates and Deletes .
Quiet Time threshold	Set the quiet time threshold to limit reporting to quiet times of a minimum length. The default value, 00:30:00, indicates that only quiet times of 30 minutes or longer are included in the Quiet time report.
Redo Direction SQL or Undo Direction SQL	Type of generated SQL (Redo or Undo).
Report format	Select the format in which you want the report output to be displayed.
Report type	Select the amount of detail to include: summary information only, details for transactions, or both.
Report verbosity	Select whether to include statistics only, or to include all available details as specified by the Report type setting.
Sequence number	The value in the transaction that shows when the transaction (unit of work) ended. This value together with the application ID can be used to uniquely identify a transaction.
Start date	Start date for the transaction period.
Start LSN/LSO/LRI	Starting log sequence number or log stream offset of the report range, if specified.
Start time	Start time for the transaction period.
Start timestamp	Start date and time for the transaction period.
Status	Transaction status (Committed, Uncommitted, Partial, or Rolled back).
Stop on error	Stops performing SQL if an error is encountered.
Table List	List of tables that were used to generate the report.
Tables	Groups the activity for all tables.
Tablespaces	Groups the activity for all table spaces. On the report, the list of table spaces that were used to generate the report.
TID FID Map	List of table space ID (TID) values or table space and table ID (TID:FID) paired values that were used for report filtering.
Time	The point in time at which the transaction was committed.
Totals	Contains information about inserts, updates, and deletes. The total of these rows is displayed in the Total row at the bottom of the table.
Transaction IDs	List of transaction IDs that were used for report filtering.
Transaction start LRI	Start time of the log record identifier for the transaction period.
Transaction LRI	List of log record identifiers that were used for report filtering.
Transaction LSN/LSO	Log sequence number or log stream offset of the commit log record.
Transaction start LSN	Start time of the sequence number or log stream offset for the transaction period.
Transaction state	Types of transactions included into the report (committed, uncommitted, rolled back, or partial).
User	Database user ID.

Table 21. Alphabetical list of Log Analysis Result step fields and descriptions (continued)

Field	Description
XA transaction ID	Written only by XA applications; the transaction ID that is unique across all databases in the instance.

Related tasks:

“Working with Log Analysis results” on page 114

The **Result** step consists of the **Statistics**, **Post filters**, and **Report** views. You can switch between the sub-views by clicking the appropriate title in the top bar.

Chapter 8. Recovery

On the **Recovery** perspective, you can recover data and dropped objects for a database.

Recovery overview

The typical workflow for a Recovery operation is as follows:

1. Access the **Recovery** perspective (from the **Open**, click **Recovery**).
2. From the **List of databases** panel, select an existing user database on which to perform recovery operations.
3. Select one or more database objects or patterns to be recovered. Although you can select objects of different types, the data will be restored only if objects of the same type are selected.
4. Specify the point in time of recovery by selecting **Timestamp** or **End of logs**.
5. Specify additional locations to search for log files and backup images.
6. View dependent objects and specify related tables for recovery.
7. Review the available recovery plans and select a plan to perform the recovery of database objects that were previously specified.
8. Review the progress and completion status of the recovery.
9. Click **Run**.

You use the following steps in the **Recovery** perspective:

- **Overview** step. Review a description of the operations that can be performed on the **Recovery** perspective.
- **Location** step. Review the registered databases and detailed information about a selected database.
- **Objects** step. Select one or more database objects to be recovered.
- **Point in time** step. Define the time to which an object is recovered.
- **Options** step. Specify object definitions for recovery and additional locations to search for log files and backup files if they are not located in the default location recorded in the DB2 database history.
- **Dependencies** step. View dependent objects and specify related tables for recovery.
- **Recovery plan** step. Create and review the available recovery plans and select the one that you want to perform.
- **Status** step. Review the progress and completion status of the recovery request.

Topics:

- “Selecting a database for recovering objects” on page 124
- “Selecting objects for recovery” on page 124
- “Selecting a point in time for recovery” on page 126
- “Specifying additional recovery options” on page 126
- “Viewing dependent objects and specifying related tables for recovery” on page 127
- “Selecting a recovery plan” on page 129
- “Reviewing recovery status” on page 137

Selecting a database for recovering objects

On the **Location** step, you can select a registered database from the **List of databases** panel to perform recovery operations.

About this task

You might be required to log into the database, depending on the connection properties that the administrator defined for the datastore connection.

Although you can select objects of different types, the data will be restored only if you select objects of the same type.

DB2 10.5 FP1 only: For a database to be available for recovery, NLS_STRING_UNITS must be set to NULL or SYSTEM, and STRING_UNITS must be set to SYSTEM.

Procedure

To select a database for recovering objects:

1. To see the managed systems, instances, and databases, expand the tree in the **List of databases** panel.
2. Select a recoverable database from the tree. Information about the database is displayed in the **Selected database information** panel on the right.
3. If prompted, specify a user name and password to connect to the database. This requirement is defined by the connection properties of the datastore repository database that the admin user has defined.
4. Proceed to the **Objects** step.

Related concepts:

“Limited or unsupported DB2 features, environments, and operations” on page 168
DB2 Recovery Expert provides limited or no support for the following DB2 features, environments, and operations.

Selecting objects for recovery

On the **Objects** step, select database objects to be recovered. By default, Recovery processing excludes all of the objects from processing. To override the default, select objects for recovery.

About this task

DB2 Recovery Expert supports recovery of the following non-catalog-based DB2 object types:

- Recovery includes definitions for the following objects:
 - buffer pools
 - event monitors
 - functions
 - database partition groups
 - schemas
 - sequences
 - SQL procedures
 - user types

- variables
- Recovery includes definitions and data for the following objects:
 - tables
 - table spaces
 - storage groups
- Selecting objects for processing:

You drill down to the object that you want to select. The objects appear in different colors based on where the object exists:

 - Black: In both the Schema Level Repository (SLR) and the DB2 system catalog.
 - Red: In the SLR only (for example, when the object is dropped).
 - Blue: In the DB2 system catalog only.
- Certain DB2 system and Recovery Expert internal objects, such as schemas, table spaces, and buffer pools, are not included in this list because such objects are not recoverable. These objects will not appear in the list, even when you use the filter "%" or leave the field blank.
- For objects that have schemas, DB2 Recovery Expert has two levels of hierarchy, as shown in the following example:

```
Tables
  -Schema1
    Table1
    Table2
  +Schema2
  -Schema3
    Table3
    Table4
```

Procedure

1. From the **Available objects** panel, click the plus sign to expand the tree.
2. Select objects for recovery using one of the following actions:
 - Select one or more available objects and drag them to the **Selected objects** panel.
 - To include a group of objects, select **Pattern**, and then click **Include**. In the **Pattern** window, specify filter values. To see the objects list, expand the **Pattern** group in the **Selected objects** panel.
3. To filter the list of available objects:
 - a. Click **Object filters**.
 - b. Click **Include**.
 - c. Enter the filtering pattern to be used to identify the objects. The syntax is similar to the DB2 LIKE clause; use % and _ symbols to specify required filters. To include all of the currently available DB2 objects, enter the default filter % or leave the field blank.

To limit the number of objects that are displayed in a branch, you can filter each level by schema and then by name. The schema and the name filters are independent; that is, all tables under each schema are filtered with the name filter. Specify a filter value in following format: *schema_name.object_name*. If you do not specify either a schema or an object filter value, DB2 Recovery Expert uses the wildcard value "%". The following examples illustrate schema and object filtering.

Table 22. Schema and object filtering include examples

Operation	Filter value	Schema filter	Object filter	Explanation
Include	(no value)	%	%	Include all objects in all schemas.
Include	A%	%	A%	Include all objects starting with "A" in all schemas.
Include	D%.A%	D%	A%	Include all objects starting with "A" where the schema name begins with "D".

- d. Click **OK**.
4. To view further details in the **Object properties** panel, in the **Selected objects** panel, click the object.
5. Proceed with the **Point in time** step.

Selecting a point in time for recovery

Use the **Point in time** step to define the time to which an object is recovered.

Procedure

1. Specify the point in time of recovery by using the option that best suits your situation.
 - In the following situations, select **End of logs** to recover the selected database objects:
 - To perform a table space recovery, and the table space already exists.
 - To fix a problem where the log data was correct but the database contents became physically corrupted.

The product restores data and processes the recovery logs through the end of the available log data.
 - In the following situations, select **Timestamp**:
 - To recover to a previous point in time.
 - When data becomes logically inconsistent and you can identify a prior timestamp when the data was consistent.
2. Specify a time or timestamp as follows:
 - (*Nonpartitioned database environment*) Click **Select backup time**, and then select or specify a backup time.
 - (*Partitioned database environment*) Specify the end timestamp of the partition backup in a set of backups that ended last or later.

Specifying additional recovery options

On the **Options** step, you can specify whether to recover object definitions only. You can also specify additional locations to search for log files and backup images if they are not located in the default location that is recorded in the DB2 database history.

About this task

When you advance to the **Recovery plan** step, one or more recovery plans will be generated using the supplied resources and point-in-time information. Some time is required to generate the recovery plans. You can then choose a recovery plan to run to perform the recovery.

Procedure

1. If necessary, select the **Recover object definition only (data will not be restored)** check box. By default, the product recovers object definitions and data for the selected objects.
2. If necessary, select the **Recover all privileges for the restored database objects** check box. Choose this option to ensure that users have the same level of access to the database object after it is recovered. The product supports the following types of privileges:
 - table
 - table space
 - view
 - index
 - schema
 - variable
 - module
 - procedure
 - function
 - sequence
3. Click the plus sign (+) in the **Log directories** or **Backup directories** panels, and browse to the required location. For Database Partitioning Facility (DPF) environments, to specify the location for all database partitions (when the location is the same), click **All Database Partitions**. To specify a location for a specific database partition, select it in the **Database Partition** menu.

Note: The first path specified for all database partitions in the log directories list is used as the default overflow log path. If the recovery requires running the DB2 ROLLFORWARD command, and no path is specified, the command fails if any database partition overflow log paths are specified for the ROLLFORWARD command.

4. Click **OK**.
5. Complete optional tasks as follows, or proceed to the **Dependencies** step.
 - To modify the added location, select the location and then click the edit icon.
 - To delete one or more locations, select the locations and then click **Remove** or **Remove all**.

Related concepts:

“Considerations for DPF environments” on page 171

Review considerations for using DB2 Recovery Expert in a Database Partition Facility (DPF) environment.

Viewing dependent objects and specifying related tables for recovery

On the **Dependencies** step, you can make sure that the SLR exists and is up to date, and review dependent objects and specify related tables for recovery.

Before you begin

To perform a recovery, the SLR must exist, and it must include the point in time (PIT). You can update the SLR from this step, if necessary. However, if the SLR does not exist, you must create it in the **Object History** perspective. When you create the SLR, specify the base backup and the end point.

The **Update SLR** button is displayed when the SLR is obsolete. For example, **Update SLR** appears if the selected point in time is earlier than the SLR end point. If this occurs, you must perform an update SLR operation.

About this task

- After checking the SLR, DB2 Recovery Expert generates the list of dependent objects, which are related to the objects that you selected on the **Objects** step. The amount of time required to generate a list of dependent objects is based on the number of objects that you selected for recovery, and the number of dependencies that the selected objects have. If many dependencies exist, generating the list might require a significant amount of time.
- Tabs on the **Dependencies** step of the **Recovery** perspective show relationships to the selected objects, as follows:
 - The **Related tables** tab displays all of the tables that are directly or indirectly related. Use action buttons to choose objects for the recovery plan. (Hover the cursor over each button to show button names.)
 - The **Dependencies** tab displays all other objects (aliases, foreign keys, indexes, table spaces, views) that are directly related.
- By default, all dependent objects are recovered along with objects that you select. If you do not want to recover some of the related tables, you can exclude them from the recovery.

Note: Only tables can be excluded. If there are no related tables, only the **Dependencies** option is enabled and the **Related tables** option is disabled.

Procedure

1. If the SLR is not up to date, click **Update SLR**.

The **Status** area displays the **Time**, **Database Partition**, and **Message** of the operation.
2. If necessary, you can create a detailed trace that can be used to analyze the operation and provide detailed information for Software Support by clicking **Support**. If **Support** is not displayed, you can enable it in the **Preferences** perspective.
3. On the **Status** pane, specify the type of messages to be displayed by selecting from the menu **All messages**, **Warnings and errors**, or **Errors only**.

To copy the messages to a clipboard so that you can paste them into a text editor, click **Copy status messages to clipboard**.

To download a compressed file that contains the specified messages and either view the compressed file directly, or save the compressed file to a local directory, click **Download status messages**.
4. Click **Continue**.

The **Dependencies** perspective displays the dependencies and the related tables that have been selected.
5. On the **Dependencies** tab, you can review a list of all of the dependent objects.
6. On the **Related tables** tab, you can

- Move tables from the **Excluded related tables** pane to the **Selected related tables** pane by selecting the table and clicking **Include** or **Include all**.
- If related tables exist, select related tables to include or exclude for the recovery.
- Drag tables to the appropriate pane as necessary.
- Remove tables from the **Selected related tables** pane by clicking the **Remove** or **Remove all** icon.

What to do next

The Recovery plan generation is executed as a session. To save the session for later use, click **Save**.

Note the following considerations for saved sessions:

- User selection of related tables, and the event presence of them, is not known at the moment of dependencies generation. Therefore, this information is not recorded within the Session. If you want to save this, you must generate the Recovery plan and then save the Session.
- If a session was saved on the **Dependencies** step and this step contains related table information, then any changes made to the related tables will not be saved.
- When you click **Refresh**, the state of the SLR is checked. If the state of the SLR is stable, then it reloads the dependencies; however, the session is not regenerated. To regenerate the dependencies, you can change the PIT or select different objects.

Selecting a recovery plan

The Recovery plan is generated when you navigate to the **Recovery plan** step for the first time.

About this task

The Recovery plan will be regenerated if your selection of objects or the point in time (PIT) has changed, or if you click **Regenerate**. If you click **Regenerate**, the session named **Generate** becomes available on the **Sessions** perspective, and you can view this session later. If a session was saved on the **Recovery Plan** step, then any excluded related tables from the **Dependencies** step will be saved.

DB2 Recovery Expert only uses backups that are recorded in the database Recovery History File (RHF) when generating recovery scenarios.

The amount of time that it takes to generate Recovery plans depends on factors such as the number of backups, the number of objects for recovery, and the SLR range.

Procedure

Select a plan to perform the recovery of database objects:

1. From the **Active scenario** menu, select a recovery option. Each scenario is a set of recovery steps that will be performed on a specific partition. Only certain recovery options that apply to your recovery scenario are displayed in the **Active scenario** menu.

After you select a scenario, all of the steps for the selected scenario are displayed on the **Scenario steps** panel.

The table on the **Scenario steps** panel contains the following columns:

- **Step** is the ID number of the step.
- **Description** is a brief description of the step.
- **Database partition number** is the ID number of the partition where this step will be performed.

It is possible that some of the steps with the same name will be listed in the scenario that will be performed on several or all database partitions.

2. On the **Scenario steps** panel, select a scenario step.

The details of the scenario step appear on the **Selected step details** panel. The **Selected step details** panel consists of the following columns. The ... icon is displayed if a field can be edited.

- **Option** displays the field name.
- **Value** displays the value of the field.

3. (Optional) Redirect table space containers. Complete the following steps to define table space containers for the restored table space on the target system:
 - a. On the **Scenario steps** panel, select **Create object**.

Note: More than one step might show **Create object** in the Description column. Select the step and then look for **Table space parameters** in the **Selected step details** table Option column.

- b. On the **Table space parameters** row in the **Selected step details** panel, click the ... button. The **Storage parameters** window opens.
- c. Define table space containers for the restored table space on the target system. Click the plus (+) icon to add rows to the table for additional locations.

Note: The following considerations apply to containers:

- If one or more of the containers is inaccessible for any reason, the restore fails. To restore in such a case, redirecting of table space containers is supported during the restore.
 - For a container in a DMS table space, **Number of pages** must be greater than 0.
- d. Click **Resize parameters** and define resize parameters for the recovered table space.
 - e. Click **Save**.

Note: Be aware of the following points:

- If an error occurs later during the recovery at the **Create Object** step, the recovery is stopped. You must return to the **Recovery plan** step, regenerate the Recovery plan, enter a new, correct path into the table space container location, and then run the recovery again.
 - In Database Partitioning Facility (DPF) environments, you must define containers for each database partition.
4. (Optional) To edit scenario step details, complete the following steps:
 - On the **Selected step details** panel, to edit the value, click on the field for simple data types, such as numbers, text, or list choices.
 - Click the ... button to open the dialog for complex data types, such as table space containers and backups.
 5. (Optional) Perform additional tasks as needed:
 - To undo your changes, click **Restore**. The selected step details will be replaced with the default values.

- To undo all of your changes for all tasks, click **Refresh**.
- To forcibly regenerate the Recovery plan, click **Regenerate**. This task might be useful if external conditions have changed or if you want to make sure that the Recovery plan is up to date.
- To specify which kinds of messages are displayed, on the **Status messages** panel, select **All messages**, **Warnings and errors**, or **Errors only**.
- To copy the messages that you have specified in the menu to a clipboard so that you can paste them into a text editor, click the copy icon.
- To download a compressed file that contains the specified messages and either view the compressed file directly, or save the compressed file to a local directory, click the download icon.

6. Click **Run**.

Related reference:

“Recovery plan definitions”

This topic provides a reference to the scenarios and scenario steps that are available as part of a recovery plan.

Recovery plan definitions

This topic provides a reference to the scenarios and scenario steps that are available as part of a recovery plan.

Recovery scenarios

Translated table space restore from backup image and Redo SQL

This recovery scenario is offered by DB2 Recovery Expert when one or more table spaces need to be recovered. The recovery scenario will re-create table space DDLs if these table spaces were dropped. The DB2 Restore API call will be initiated after table spaces were re-created. The restore API will work with the DB2 Recovery Expert component to restore table space data. The **Rollforward** operation will be performed after the restore, and the Log Analysis task will be performed to generate a **Redo SQL**, which will follow the **Rollforward**.

Translated table space restore from backup image and Undo SQL

This recovery scenario is offered by DB2 Recovery Expert when one or more table spaces need to be recovered. The recovery scenario will re-create table space DDLs if these table spaces were dropped. The **DB2 Restore API** call will be initiated after table spaces were re-created. The **Restore API** will work with the DB2 Recovery Expert component to restore table space data. The **Rollforward** operation will be performed after the restore, and the Log Analysis task will be performed to generate an **Undo SQL** after the **Rollforward**.

Extract data from backup image and generate Redo SQL

This recovery scenario will be offered by DB2 Recovery Expert when recovering tables. This recovery scenario extracts the data from a backup image in a loadable format and is followed by a Log Analysis task to generate **Redo SQL**.

Note: During this scenario, any constraints that prevent DB2 Recovery Expert from recovering the selected objects will be dropped and then re-created after the recovery. For example, if the selected tables for recovery contain primary or foreign key constraints (Referential Integrity) these constraints will be dropped before the recovery and re-created after the recovery.

Re-create object definitions from the SLR and generate Redo SQL

This recovery scenario is offered by DB2 Recovery Expert when recovering either tables or table spaces. The recovery will re-create all object definitions from the SLR, and the subsequent Log Analysis task will generate and then apply **Redo SQL**.

Extract data from backup image and generate Undo SQL

This recovery scenario will be offered by DB2 Recovery Expert when recovering tables. This recovery scenario extracts the data from a backup image in a loadable format and is followed by a Log Analysis task to generate **Undo SQL**.

Note: During this scenario, any constraints that prevent DB2 Recovery Expert from recovering the selected objects will be dropped and then re-created after the recovery. For example, if the selected tables for the recovery contain primary or foreign key constraints (Referential Integrity), these constraints will be dropped before the recovery and re-created after the recovery.

Generate Undo SQL from current state

This recovery scenario will be offered by DB2 Recovery Expert when recovering tables or table spaces, and is only offered when the SLR end-point time is greater than the Minimum Recovery Timestamp of table spaces that need to be recovered. The Log Analysis task will be initiated to generate **Undo SQL** and then applied.

Note: This scenario will not be offered for tables or table spaces that contain LONG or LOB data, extended row size, or any constraints that prohibit the recovery.

Restore table space using DB2 RESTORE and ROLLFORWARD

This recovery scenario will be offered by DB2 Recovery Expert when one or more table spaces need to be recovered and these table spaces are not dropped. This recovery scenario will invoke the DB2 **Restore** command, which will be followed by the **Roll forward** command.

Re-create object definitions from SLR (DDL only; no data will be restored)

This recovery scenario will be offered by DB2 Recovery Expert when no object data (such as tables or table spaces) will be recovered. This recovery scenario will only re-create object definitions using DDL that is produced from information that is recorded in the SLR. Any associated table or table space data will not be restored.

Scenario steps

The following steps are available on the **Scenario steps** panel:

Drop existing objects

This step removes objects, which are listed in the **Selected step details** panel, from the recovery object list, including tables (except in UNDO scenarios), table spaces, indexes, aliases, views, monitors, triggers, procedures, functions, modules, and sequences. When no objects must be dropped, you can exclude this step.

Create objects

This step creates objects that are absent in the database but must be recovered, and that were removed in the **Drop objects** step. Tables will be created without check constraints and foreign keys. When no objects must be created, you can exclude this step.

The **Selected step details** panel shows

- the list of objects (not mutable) that will be created
- whether to prepare database objects to apply SQL
- table space type (AUTO or DMS)

For DMS table spaces, you can

- add or change the table space location container (the location type can be file or device)
- set the number of pages for the table space container

Click **Resize parameters** to access the following additional options:

- **Autoresize:** Choose whether to enable the auto-resize capability of a DMS or AUTO table space. Auto-resizable table spaces automatically increase in size when they become full.
- **Initial size:** This option is only valid for AUTO table spaces. Specify the initial size, per database partition, of an automatic storage table space. Specify an integer value followed by K (for kilobytes), M (for megabytes), or G (for gigabytes). The value must be at least 48 K.

Note: The actual value that is used might be slightly smaller than that specified because the database manager strives to maintain a consistent size across containers in the table space. Also, if the table space is auto-resizable and the initial size is not large enough to contain meta data that must be added to the new table space, the database manager will continue to extend the table space until there is enough space. If an initial size is not specified, the database manager determines an appropriate value.

- **Increment by:** Specify the amount, per database partition, by which a table space that is enabled for auto-resize will automatically be increased when the table space is full and a request for space has been made. Specify an integer value followed by K (for kilobytes), M (for megabytes), or G (for gigabytes).

Note: The actual value that is used might be slightly different than that specified because the database manager strives to maintain consistent growth across containers in the table space. If the table space is auto-resizable, but the amount by which to increment is not specified, the database manager determines an appropriate value.

- **Maximum size:** Specify the maximum size, per database partition, to which auto-resizable table spaces can automatically be increased. If the table space is auto-resizable, but the maximum size is not specified, the default is None. Specify an integer followed by K (for kilobytes), M (for megabytes), or G (for gigabytes). Specify None to indicate that the table space is to be allowed to grow to file system capacity, or to the maximum table space size.

Note: The actual value that is used might be slightly smaller than that specified because the database manager strives to maintain consistent growth across containers in the table space.

- **Caching:** Specify whether I/O operations are to be cached at the file system level. If no value is specified, the default is
 - **File system caching** for JFS on AIX, Linux System z[®], all non-VxFS file systems on Solaris, HP-UX, SMS temporary table space files on all

platforms, and all LOB and large data objects. **File system caching** specifies that all I/O operations in the target table space are to be cached at the file system level.

- **No file system caching** on all other platforms and file system types. **No file system caching** specifies that all I/O operations are to bypass the file system level cache.

Run recovery components

This step runs the recovery components. When running on an instance with a single partition, this step has no options. When running on an instance with multiple partitions, the **Selected step details** panel shows the following option:

- **Database Partition:** The database partition number on which the recovery will be run.

Validate data

This step validates options that were specified in the **Options** step for a specific database partition.

The **Selected step details** panel shows the following options:

- **Backup locations:** The specified backup.
- **Database:** The specified target database name.
- **Log locations:** The path to the specified log file.
- **Backup vendor options configuration parameter:** Optionally specify a value to override the VENDOROPT database configuration item for reading backup images.
- **Primary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT1 database configuration item for reading log files from LOGARCHMETH1.
- **Secondary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT2 database configuration item for reading log files from LOGARCHMETH2.

Synchronize recovery plan results

This step is required for multiple-partition databases to make sure that tasks are completed on all partitions. This step has no options.

Extract

This step extracts table data from a backup.

The **Selected step details** panel shows the following options:

- **Backup timestamp:** Timestamp of the backup from which data is extracted.
- **Database:** The target database name.
- **Extract map:** The TID:FID for tables.
- **Keep export file:** Whether the export file is saved to the session directory (Yes) or the temporary directory (No).
- **Backup vendor options configuration parameter:** Optionally specify a value to override the VENDOROPT database configuration item for reading backup images.
- **Primary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT1 database configuration item for reading log files from LOGARCHMETH1.

- **Secondary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT2 database configuration item for reading log files from LOGARCHMETH2.

Load This step loads table data into the database.

The **Selected step details** panel shows the following options:

- **Save mode:** Choose **Save load copy** to copy the loading data. Choose **Make backup** to perform a backup of the table spaces and the tables that are loaded by the backup step.
- **Table:** The name of the DEL file with data.

Backup

This step performs a table space backup if, in the Load step, the value of **Save mode** is **Make backup**.

The **Selected step details** panel shows the following options:

- **Buffer size:** The size, in 4 KB pages, of the buffer that is used when building the backup image. If this option has no value, then DB2 will automatically choose an optimal value.
- **Media header timestamp**
- **Number of buffers** to be used for the backup.
- **Online:** Whether to perform an online backup. Online backups are only available for databases configured with **logretain** or **userexit** enabled. During an online backup, DB2 obtains IN (Intent None) locks on all tables existing in SMS table spaces as they are processed. S (share locks) are no longer held on LOB data in SMS table spaces during online backup.
- **Parallelism:** Determines the number of table spaces that the backup utility can read in parallel. If no value is provided, DB2 will automatically choose an optimal value in the range 0 through 999999.
- **Table spaces:** (not mutable) The table spaces to be backed up.
- **Backups:** (mutable) The type of backup, as follows:
 - **Disk location:** You can add, edit, or delete backup locations.
 - **TSM:** You can set the number of I/O sessions (1 through 1000) to be created between DB2 and the backup product. When the number of sessions is also specified, choose the path to the backup product library.
 - **Vendor** (to specify another backup vendor's product): You can set the number of I/O sessions (1 through 1000) to be created between DB2 and the backup product.
 - **Backup vendor options configuration parameter:** Optionally specify a value to override the VENDOROPT database configuration item for reading backup images.

Translate table space

This step re-creates the specified table space from a backup.

- The **Selected step details** panel shows the backup timestamp (not mutable), the name of the database that contains the table space to be translated (not mutable), and the translation map.
- **Backup vendor options configuration parameter:** Optionally specify a value to override the VENDOROPT database configuration item for reading backup images.

Rollforward

This step performs a roll forward action for the specified table space of a given database.

The **Selected step details** panel shows the caller action, the database, and the table space for which the action applies.

Generate redo SQL

This step generates redo SQL for the specified database and TID:FID filter from begin to end time; it has the following options:

- **Backup vendor options configuration parameter:** Optionally specify a value to override the VENDOROPT database configuration item for reading backup images.
- **Primary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT1 database configuration item for reading log files from LOGARCHMETH1.
- **Secondary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT2 database configuration item for reading log files from LOGARCHMETH2.

Generate undo SQL

This step generates undo SQL for the specified database and TID:FID filter from begin to end time; it has the following options:

- **Backup vendor options configuration parameter:** Optionally specify a value to override the VENDOROPT database configuration item for reading backup images.
- **Primary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT1 database configuration item for reading log files from LOGARCHMETH1.
- **Secondary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT2 database configuration item for reading log files from LOGARCHMETH2.

Apply SQL

This step applies generated SQL to the target database.

The **Selected step details** panel shows the following options:

- **Commit scope:** (mutable) The number of rows (0 through 999999) that are grouped together in a commit scope.
- **Direction:** (not mutable) The type of the applied SQL, either **redosql** or **undosql**.
- **Stop on error:** (mutable) Specifies whether the action stops when an error occurs.

Update SLR history

This step executes SLR update. It is inserted at the end of the scenario if the scenario includes the **Create objects** step. It is inserted at the beginning of the **Generate Undo SQL from current state** recovery scenario.

If no data was entered on the **Options** step, the **Selected step details** panel shows no options. Otherwise, the **Selected step details** panel shows the following options:

- **Backup locations:** The path to the backup that was specified in the **Options** step.
- **Log locations:** The path to the log file that was specified in the **Options** step.

- **Backup vendor options configuration parameter:** Optionally specify a value to override the VENDOROPT database configuration item for reading backup images.
- **Primary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT1 database configuration item for reading log files from LOGARCHMETH1.
- **Secondary log archive method configuration parameter:** Optionally specify a value to override the LOGARCHOPT2 database configuration item for reading log files from LOGARCHMETH2.

Related tasks:

“Selecting a recovery plan” on page 129

The Recovery plan is generated when you navigate to the **Recovery plan** step for the first time.

Reviewing recovery status

The **Status** step displays the progress and completion status of the recovery.

About this task

- The **Status** step is the same throughout the product for **Log Analysis**, **Recovery**, and **Object History** operations.
- The header area of the window displays the following information:
 - Operation: The operation that you are currently performing.
 - Session ID: The ID of the submitted task.
 - Start time and end time of the operation.
 - Status: The progress of the task, the result of the task, and whether the operation was successful.
- While the operation is in progress, a time indicator is displayed near the lower left corner.
- The **Status** area displays the time, database partition, and messages related to the operation.

Best practice: Take a full backup or a table space backup of the recovered objects after the recovery has completed. For a table space backup, include the catalog table space (SYSCATSPACE) and all user table spaces that were affected by the recovery. This backup will create a new recovery point for use by future recoveries. If a backup is not taken, it might not be possible to recover the same objects in the future by using certain recovery plans.

Procedure

You can perform the following optional tasks:

1. To cancel any pending operation, click **Cancel**.
2. To choose which kinds of error messages are displayed, from the dropdown list, select one of the following options:
 - **All messages**
 - **Warnings and errors**
 - **Errors only**
3. To copy messages to a clipboard so that you can paste them into a text editor, click the copy icon.

4. To download a compressed file that contains the messages, and either view the compressed file directly, or save the compressed file to a local directory, click the download icon.
5. To create a detailed trace that can be used to analyze the operation and provide detailed information for Software Support:
 - a. Click **Support**. If **Support** is not displayed, you can enable it in the **Preferences** perspective.
 - b. Select options and then click **OK**. The **Download trace file** window opens. You can view the file and save it.



Chapter 9. Command line processor (CLP) interface

In addition to the browser client, DB2 Recovery Expert provides a command line processor (CLP) interface that enables you to use the product on machines other than the DB2 server.

You must perform a separate installation procedure to install the CLP client. DB2 Recovery Expert versions 3.1 and later use a distributed product architecture. This architecture does not support the command-line executables of earlier versions of the product.

Topics:

- “Java Command Line Processor (CLP) overview”
- “Configuring the CLP” on page 140
- “Command descriptions and usage details” on page 141
- “Using the CLP to run Log Analysis against remote databases” on page 160

Java Command Line Processor (CLP) overview

DB2 Recovery Expert includes a Java CLP client that can be installed onto any supported client system to connect to the configured remote DB2 servers and access specific DB2 Recovery Expert features from a command-line environment.

The Java CLP can be installed for shared use or for use by an individual user. The user who installs the client can choose the installation location. The installation process installs a suitable version of Java for use by the CLP.

Supported features include

- Log Analysis
- Schema Level Repository (SLR) operations
- database object data extraction from backups

(Full recovery processing is not supported when you use the CLP.)

The Java CLP uses a natural language style of syntax that is similar to the DB2 Command Line Processor. The Java CLP runs in interactive shell mode and batch mode.

Interactive shell mode

In interactive shell mode, each command is entered one at a time after starting the shell.

Batch mode

In batch mode, a text input file is created to contain the sequence of commands to run, and the command input file name is passed to the Java CLP for processing. Java CLP batch mode processing can be scheduled to run a specific set of commands at a specific time using the operating system scheduling interface. Using the CLP in batch mode makes it easy to run daily Log Analysis audit reports or perform regularly scheduled SLR updates on databases of interest. Most processing and data storage occurs on the remote database server.

Some commands create a session ID number that must be referenced by following related commands. For example, you can run a Log Analysis session to produce report data, and then reference that session's results to export the generated report or SQL to a file. The @ substitution symbol can be used to reference the last session ID that was created by the previous **run** command. This is useful in command input files where the previous session number cannot be known in advance.

Running the command

The main Java CLP command format is as follows:

- Linux and UNIX:
aryclp.sh [-f commandfile] [-h] [-s] [-v] [-V] [-w codepage]
- Windows:
aryclp.bat [-f commandfile] [-h] [-s] [-v] [-V] [-w codepage]

Interactive shell mode: To start a command processor shell, run the command without arguments.

Batch mode: To run a set of commands in batch mode, use the `-f` option and supply a command input file. Commands are read from the command input file until the end of the file is reached; the input file does not need to include a **quit** command.

Using command options

- f commandfile**
Run commands from specified command input file path name.
- h** Display command help messages.
- s** Stop on error when processing command input file.
- v** Display Java CLP version information.
- V** Run in verbose mode.
- w codepage**
Run using specified character code page.

Return status

When you run a set of commands in batch mode by using the `-f commandfile` option, the command return status is set as follows:

- 0** All commands were successful.
- 1** One or more commands in the input file failed.

You can test this status from the calling command shell.

Configuring the CLP

The configuration file, `aryclp.properties`, is located in the CLP installation directory tree under the `conf` subdirectory. You can edit the file to identify the location of the `datastores.xml` file and to initialize other settings.

Before you begin

- Do not edit the `datastores.xml` file directly.
- You must have access to create files in the installation directory and to modify the `datastores.xml` file to allow the `add datastore connection` command to update the file.

- Optionally, you can make the `aryclp.properties` and `datastores.xml` files read-only to prevent accidental changes after the datastore definitions are available.

Following is an example of the `aryclp.properties` file:

```
Default contents of the conf/aryclp.properties file:
datastore.path={INSTALL_DIR}/conf/datastores.xml
log.path={USER_HOME_DIR}/aryclp.log
log.level=1
time.representation=server
```

In the examples, the variable `vrmm` represents the product version number.

Procedure

1. Set the property `datastore.path` to an absolute path name of the file containing the datastore connection definitions. The default setting creates the file in the installation directory by referencing the `INSTALL_DIR` predefined configuration variable.

In the following example, the Java CLP components were installed in the directory `/opt/revrmm/clp`:

```
datastore.path=/opt/revrmm/clp/datastores.xml
```

2. Add definitions for all of the datastore connections that are available to the command processor to the `datastores.xml` file. You can copy the `datastores.xml` file from another Java CLP environment if the definitions have previously been established, or you can add custom definitions to the local configuration by using the `add datastore connection` command.
3. Add the bin directory under the installation directory path name that contains the main Java CLP command to the command shell `PATH` variable. For example:

```
Linux/UNIX: PATH=/data/RE41/CLP/bin:$PATH
```

```
Windows: set PATH=c:\data\RE41\CLP\bin;%PATH%
```

If the directory is not available using the search path, you can run the command by using a relative or absolute path.

4. Set `log.path`. The `log.path` property identifies the location for the diagnostic trace file. The `USER_HOME_DIR` predefined configuration variable causes it to be created in the user's home directory.
5. Set `log.level`. The `log.level` property controls the log verbosity level. The default setting of "1" sets normal logging detail. Customer Support might suggest other settings to assist in problem diagnosis.
6. Set `time.representation`. The `time.representation` property controls how timestamp values used by the CLP will be interpreted. Values are as follows:
 - **client**: Interpret timestamps in terms of the client system where the CLP commands are run.
 - **server** (default): Interpret timestamps in terms of the time the target DB2 instance that is being accessed.
 - **utc**: Interpret timestamps as UTC (Coordinated Universal Time).

Command descriptions and usage details

Commands consist of one or a few required words followed by additional keywords, options, and input parameters. This section provides a summary of commands organized by type, and details about using each command.

Getting Help for a command

To display a help summary in command interactive shell mode, type `help` or `?` (question mark). To display additional details about a command, type `help` or `?` (question mark) followed by the specific command name. Type all of the words that make up the command name as shown in the command list. For example:

```
help export report
? export report
```

Using special characters

The following special characters are supported:

- Use the `@` symbol to substitute the session ID that was created by the previous `run` command.
- Use the `#` character as a comment character in the command file to add descriptive text.

Using filters and patterns

To exclude objects that match a specific name, enter an exclamation point as the first character before the name.

If a list of more than one name is provided, an exclamation point must be used before each name in the list.

Command summary

The following table provides a summary of each command. The section after the table provides details about each command, organized alphabetically by command name.

Table 23. Command summary

Command name	Command type	Description
<code>sync with configuration repository</code>	Connection related	Synchronizes managed database information from the application server repository
<code>list datastores</code>	Connection related	List available datastore connections
<code>reset connection</code>	Connection related	Reset the current datastore connection
<code>set current datastore with label</code>	Connection related	Set the current datastore connection
<code>show current connection</code>	Connection related	Display the current datastore connection
<code>apply SQL</code>	Action	Perform Redo/Undo SQL generated by Log Analysis
<code>export report</code>	Action	Export Log Analysis report results
<code>export SQL</code>	Action	Export Redo/Undo SQL generated by Log Analysis
<code>run ddlmaker</code>	Action	Produce DDL from the SLR for specified objects

Table 23. Command summary (continued)

Command name	Command type	Description
run la	Action	Run Log Analysis
run ox	Action	Run Object Extract tool
run slr	Action	Run commands to create or update the SLR
get slr info	Status and session	Display SLR information for database
save session	Status and session	Preserve the specified session ID so that the related data is not deleted. The @ symbol can be used to substitute the session ID that was created by the previous run command.
unsave session	Status and session	Release the specified saved session ID so that the session data can be deleted and the ID can be reused by a new session.
help or ?	Other	Display help inside the command shell
quit or q	Other	Exit the command shell
#	Other	Comment character used in command input file

Command and option details

Each command accepts a set of additional keywords, options, and input parameters that define the details of the command.

Words in the usage descriptions that are shown in plain text must be entered exactly as shown. Words enclosed inside angle brackets (such as <connection name>) represent replaceable parameters that are defined in each Description section.

Options inside square brackets (such as [show detail]) can be omitted or specified in any order on the command line.

sync with configuration repository

Usage:

```
sync with configuration repository database <dbname>
host <hostname> port <port> [user <username> password <password>]
```

Description:

This command synchronizes the managed database information from the configuration repository. Issue this command to load information about datastores, managed databases, deployments, credentials, and so on from your configuration repository into the local XML file.

Options:

dbname

Database name of the DB2 database.

host Hostname of the DB2 database server.

port DB2 connection port number.

user User name that can connect to the host.

password

Password used with the specified user name.

Example:

```
sync with configuration repository database test1  
host skala port 62005 user d95a1 password pass1
```

apply SQL

Usage:

```
apply SQL for database <database> from session <session_id>  
[user <user>] [password <password>] [stoponerror]  
[commitscope <numoperations>]
```

Description:

Perform the Redo or Undo SQL statements that were generated by a previous Log Analysis command that was run with the option to generate Redo or Undo SQL. The supplied session ID identifies the Log Analysis session from, which the SQL will be obtained. This will usually be related to the immediately previous command.

Options:

database

User database name to which the SQL will be applied.

numoperations

The number of operations that will occur between commits in the exported SQL file. Specify 0 (zero) to preserve the commit frequency used by the original transactions.

password

Password used with the specified user name.

If the password value that you specify matches a product keyword, then you must enclose the value in quotation marks (").

session_id

Identifies the related Log Analysis session that generated the SQL that will be applied. The @ symbol can be used to substitute the session ID that was created by the previous run command.

stoponerror

Indicates to stop if an error is encountered when applying the SQL.

user User name that can connect to the DB2 instance (optional). If not specified, authentication parameters will be requested at time of connection.

Example:

```
apply SQL for database SAMPLE from session @ stoponerror
```

export report

Usage:

```
export report from session <session_id>  
into <directory> as <text | xml | csv>  
level <summary | details | full | quiet>  
[threshold <time>] [statisticsonly]
```

Description:

Export a report generated by a previous Log Analysis command to a file on the local client system.

Options:

Directory

Directory or file name in which to store the output file on the client system where the command is run. Make sure that the base directory exists, and that you have permission to write to the specified location. When you specify a directory, the output is stored in one of the following default file names for exports:

- txt (report.txt)
- xml (report.xml)
- csv (ary.csv)

When you specify a file name without an extension, the default extension for that export type is added (.txt, .csv, or .xml).

session_id

Identifies the related Log Analysis session that generated the SQL that will be applied. Use the @ symbol to substitute the session ID that was created by the previous run command.

summary | details | full | quiet

These options control the content of the report file.

summary

Statistics and transaction summary only.

details

Transaction details only.

full Statistics, summary, and transaction details.

statisticsonly

Report only the statistics totals.

quiet Report on periods during which no active transactions were modifying a set of table spaces and tables on the selected database.

threshold

Specify the minimum length of quiet times to be included in the quiet time report.

Use the format `[[hh.]mm.]ss or #<s|m|h>`, where hh and h = hours, mm and m = minutes, and s and ss = seconds. For example, 30m for 30 minutes. Hours and minutes are optional, and default to 0 if not specified.

The default value, 00:30:00, indicates that only quiet times of 30 minutes or longer are included in the report. The minimum valid threshold is 1 second, 00:00:01, or 1s.

text | xml | csv

These options controls the format of the output for the exported report.

text Human readable text format.

xml XML format.

csv Comma separated value file format.

Example:

```
export report from session @ into /tmp/export1 as text level summary
```

export SQL

Usage:

```
export SQL from session <session_id> into <path>  
[commitscope <numoperations>]  
[includeddl] [includelong] [withcomments]
```

Description:

Export the Redo or Undo SQL statements generated by a previous Log Analysis command to a file on the local client system.

Options:

path Directory or file name in which to store the output file on the client system where the command is run. Make sure that the base directory exists, and that you have permission to write to the specified location. When you specify a directory, the output is stored in ary.sql the default file name for SQL exports.

When you specify a file name without an extension, the default extension for SQL exports, .sql, is added.

includeddl

Include DDL statements that occur with transactions in the generated Redo SQL output. DDL statements are not available when generating Undo SQL.

includelong

Include LONG and LOB column data with the exported output. If specified, the long data is written to files in the specified output directory. When using this option, it is recommended to place the output file in a new, empty directory, because many additional files might be created for the LOB data.

numoperations

The number of operations that will occur between commits in the exported SQL. Specify 0 (zero) to preserve the commit frequency used by the original transactions.

session_id

Identifies the related Log Analysis session that generated the SQL that will be applied. The @ symbol can be used to substitute the session ID that was created by the previous **run** command.

withcomments

Include transaction comments in the output file.

Example:

```
export SQL from session @ into /tmp/exportSQL commitscope 1000 includedd1
```

get slr info**Usage:**

```
get slr info for database <jdbc_url> [user <user>]
[password <password>]
```

Description:

Display the Schema Level Repository (SLR) status information for the specified database. This information shows whether the SLR exists for the database, and if so, shows details about the current state of the SLR.

Options:**database**

User database for which to process the SLR. Specify a valid JDBC URL in the format:

```
jdbc:db2://<hostname>:<port>/<dbname>
```

user

User name that can connect to the DB2 instance (optional). If not specified, authentication parameters will be requested at time of connection.

password

Password used with the specified user name.

Example:

```
get slr info for database jdbc:db2://verta:62022/TC5543
user user1 password passwd1
```

Example output:

```
SLR status: Created
Database registration timestamp: 2010/09/27 14:30:25
SLR last update timestamp: 2010/09/27 14:31:00
SLR start log: S0000000.LOG
SLR end log: S0000000.LOG
SLR start timestamp: 2010/09/27 14:30:25
SLR end timestamp: 2010/09/27 14:31:00
SLR start LSN: 00000000010B30EC
SLR end LSN: 00000000013882F5
```

help**Usage:**

```
help [<command_name>]
? [<command_name>]
```

Description:

In command interactive shell mode, type help or ? (question mark) to display a help summary. Type help or ? (question mark) followed by a specific command name to display additional details about the command. Type all the words that make up the command name as it is shown in the command list.

Options:**command_name**

Name of command to display help about.

Examples:

? Display general summary help.

help export report

Displays help about the specified command.

list datastores

Usage:

list datastores [show detail]

Description:

List the set of currently defined datastore connections that are available for use by the Java Command Line Processor environment.

Options:

show detail

Display details about each connection.

quit

Usage: quit

q

Description:

The commands **quit** and **q** will exit from the command shell. It is permitted, but not required, in the batch input file.

Options:

None.

reset connection

Usage: reset connection

Description

Reset (disconnect) the current datastore connection. This is required before a new connection can be established to a different datastore.

Options:

None

run ddlmaker

Usage:

run ddlmaker for database <database> objecttype <table | tablespace | procedure | bufferpool | eventmonitor | dbpartitiongroup | schema | datatype | stogroup | function | module | sequence | variable> name [schema.]<name> serverfile <file> from <time> [specificname] [user <user>] [password <password>]

Description:

Generate DDL statements for the specified database object.

Options:

database

User database name that will be processed.

file

The directory or file path name where the output will be written on the database server system. In physically partitioned environments, the file will be located on the host computer where the datastore connection was defined.

If an existing directory name is specified, the DDL output will be created in this directory using a file name consisting of the schema, table or other object name provided on the command. Otherwise, the specified output file will be created, if necessary, and the DDL statements will be written to the file.

objecttype

Values can be:

table | tablespace | procedure | bufferpool |
eventmonitor | dbpartitiongroup | schema | datatype |
stogroup | function | module | sequence | variable

If you want the product to generate DDL for an object's privileges, then specify one of the following privilege type values by using the `objecttype` option:

TABLEAUTH | TABLESPACEAUTH | INDEXAUTH |
VARIABLEAUTH | MODULEAUTH | PROCEDUREAUTH |
FUNCTIONAUTH | SCHEMAAUTH | SEQUENCEAUTH

Privilege types `PROCEDUREAUTH` and `FUNCTIONAUTH` require the addition of the `specificname` option and the specific name of the procedure or function. `ddlmaker` does not generate routine privileges by routine name; the specific name must be used.

password

Password used with the specified user name.

If the password value that you specify matches a product keyword, then you must enclose the value in quotation marks ("").

[schema].name

Identifies the name of the database object. Specify the optional schema name if needed to identify a table.

specificname

The specific name of the `CREATE PROCEDURE (SQL)` statement.

To generate DDL for privilege types `PROCEDUREAUTH` and `FUNCTIONAUTH`, specify option `specificname` and the specific name of the procedure or function.

time

Timestamp that identifies the version of the object to generate DDL for in the format:

[utc]YYYYMMDDHHMMSS[.uuuuuu]

For more information, see "Timestamp considerations" on page 159.

user

User name that can connect to the DB2 instance (optional). If not specified, authentication parameters will be requested at time of connection.

Example:

```
run ddlmaker for database SAMPLE objecttype  
table name DB2INST1.ORG serverfile /tmp/ddl1.txt from 20100625123916
```

run 1a

Usage:

```
run la for database <jdbc_url> report {summary | full |
quiet | details for session <session_id>
[user <user>] [password <password>] [from <begintime>]
[to <endtime>] [operations <i | u | d>] [mode M] [updateslr]
[savesession] [nobackup] [sql <undo | redo>] [startlog <lognumber>]
[endlog <lognumber>] [partitions <all|"1 2 3 ...">] [nolong]
[nopageaccess] [transactions ...] [<tables...| tablespaces ...
| authids ... | dbpartitiongroups ... | appids... | appnames ...
| tidfid ...>] [backups <dirlist>] [logs <dirlist>]
```

Description:

Run the Log Analysis command against the specified database with options set to control the report and optional SQL that will be generated. For the database, specify a valid JDBC URL in the following format:

```
jdbc:db2://<hostname>:<port>/<dbname>
```

Options:

appids [list]

Specify list of application ids to include in the report, included in quotes, separated by spaces.

Note: If log records contain mixed activity, for example, if some of the log records have application information and some do not, the Application names filter (appnames), the Application IDs filter (appids) and the Authorization IDs filter (authids) can be set to {empty} where application information was not set.

appnames [list]

Specify list of application names to include in the report, which should be included in quotes and separated by spaces.

Note: If log records contain mixed activity, for example, if some of the log records have application information and some do not, the Application names filter (appnames), the Application IDs filter (appids) and the Authorization IDs filter (authids) can be set to {empty} where application information was not set.

authids [list]

Specify list of Authorization IDs to include in the report, included in quotes, separated by spaces.

Note: If log records contain mixed activity, for example, if some of the log records have application information and some do not, the Application names filter (appnames), the Application IDs filter (appids) and the Authorization IDs filter (authids) can be set to {empty} where application information was not set.

backups <dirlist>

Specify an optional list of alternate directories for locating backups.

For more information, see “Specifying additional locations to search for backups and log files” on page 158.

begintime

Begin time of the report range, in the format:
[utc]YYYYMMDDHHMMSS[.uuuuuu]

For more information, see “Timestamp considerations” on page 159.

database

User database name to, which the SQL will be applied.
Specify a valid JDBC URL in the format:
jdbc:db2://<hostname>:<port>/<dbname>

endlog <lognumber>

Constrain processing to logs up to the indicated ending log number.

endtime

End time of the report range, specified in the same format as the begintime option.

logs <dirlist>

Specify optional list of alternate directories for locating logs.

For more information, see “Specifying additional locations to search for backups and log files” on page 158.

mode M

Report mode option:

M Minimum Recovery Time (MRT);

SLR is not required within this range. The default report mode is Normal, which requires the SLR to exist if the report range extends prior to the MRT value.

nobackup

Specifies to not access backup images for processing of masked update operations. By default, backups are accessed when needed.

nolong Omit LONG and LOB column data from processing.

nopageaccess

Do not access the current database pages during masked update processing. By default, the current database can be accessed when needed.

operations [i][d][u]

Specified operations to include:

i inserts
d deletes
u updates

You can specify one or more values; for example:
operations id

All operations are included if the options is not specified.

dbpartitiongroups [list]

Specify list of database partition groups to include in the report, included in quotes, separated by spaces.

partitions {all | "1 2 3 ..." }

Specify which database partitions to run Log Analysis on in a partitioned database environment. By default, Log Analysis will run on all database partitions. To process all database partitions explicitly, specify "all". Multiple database partition numbers can be specified by separating them with commas or spaces. If spaces are used, the entire option string must be surrounded by quotes.

password

Password used with the specified user name.

If the password value that you specify matches a product keyword, then you must enclose the value in quotation marks (").

savesession

This parameter preserves the session data until you run the `unsavesession` command to release the session ID.

session_id

Identifies the related Log Analysis session that generated the SQL that will be applied. The @ symbol can be used to substitute the session ID that was created by the previous run command.

startlog <lognumber>

Constrain the processing to the indicated starting log number and following logs.

summary | full | quiet | details for session <session_id>

These options control the content of the report that will be generated:

summary

Statistics and transaction summary only.

full Statistics, summary, and transaction details.

quiet Periods during which no transactions occurred.

details for session <session_id>

Transaction details based on previously run summary or full report.

sql undo | redo

Generate Redo or Undo SQL during the report processing. This option requires setting the report type to details.

tables [list]

Specify a list of tables to include in the report, included in quotes, separated by spaces. The table names have the form:

[schema.]name

where the schema prefix is optional.

tablespaces [list]

Specify a list of table spaces to include in the report, included in quotes, separated by spaces.

tidfid [list]

Specify a list of table space ID (TID) values or table space and table ID (TID:FID) paired values to identify tables or table spaces to include or exclude from the Log Analysis report. A TID is an integer value that identifies a table space. A TID:FID pair is a pair of integer values separated by a colon character (":") that identify a specific table as recorded in the database catalog. The list of IDs can contain a list of TIDs, a list of TID:FID pairs, or a mixture of TID and TID:FID values. By default, specified values indicate that the table or table space should be included in the Log Analysis report. To indicate that a table or table space should be excluded from the report, enter an exclamation point character ("!") immediately preceding the TID or TID:FID value. Use care when mixing IDs to both include and exclude objects because it is possible to construct a list of non-intersecting values that will result in an empty report.

Multiple TID and TID:FID pairs can be specified separated by commas or spaces. If spaces are used, the entire option string must be surrounded by quotes.

Examples:

2:5 Include the table with TID:FID value 2:5 (table space ID 2, file ID 5).

2,15,16 Include table spaces with TID values 2, 15, and 16.

"2 4:6 4:7"
Include table space with TID 2, and tables identified by TID:FID values 4:6 and 4:7.

!0 Exclude table space with TID value of zero (the database catalog table space).

"2 3 !2:16"
Include activity for table space TIDs 2 and 3, but exclude table with TID:FID 2:16.

transactions [list]

Specify a list of transaction IDs (Local Transaction IDs, Global Transaction IDs, and XA Transaction IDs) to include in the report, included in quotes, separated by spaces. To specify a range of IDs, use the following format: ID1..ID2, ID1.., ..ID2

updateslr

Specifies that the SLR should be updated as part of running the Log Analysis command. The SLR is not updated by default.

user User name that can connect to the DB2 instance (optional). If not specified, authentication parameters will be requested at time of connection.

Examples:

```
run la for database jdbc:db2://  
verta.abccompany.com:62055/ASDB report summary from  
20100905000000 to 20100905235959 updateslr
```

```
run la for database jdbc:db2://  
verta.abccompany.com:62055/ASDB report detail from  
20100905000000 to 20100905235959 updateslr sql redo  
partitions all tables "db2inst1.employee  
db2inst1.department"
```

run ox

Usage:

```
run ox for database <jdbc_url> from <backuptime> format <del | ixf>  
<tidfid ... | tables ...> serverdir <serverdirectory> [user <user>]  
[password <password>] [nolong] [partition <partition>] [offline]  
[lobsinfile] [backups <dirlist>] [logs <dirlist>]
```

Description:

Run the **Object Extract** command against the specified database to extract table data from a specified backup image into an output file located on the database server. This data file can then be loaded into a database with a matching table structure using the **DB2 LOAD** command.

For the database, specify a valid JDBC URL in the following format:

```
jdbc:db2://<hostname>:<port>/<dbname>
```

Options:

backups <dirlist>

Specify optional list of alternate directories for locating backups.

For more information, see "Specifying additional locations to search for backups and log files" on page 158.

backuptime

Identifies the backup image timestamp in the format:

```
YYYYMMDDHHMMSS
```

This timestamp is always in terms of the database server time zone, so it matches the filename of the actual backup image.

database

User database name that will be processed. Specify a valid JDBC URL in the format:

```
jdbc:db2://<hostname>:<port>/<dbname>
```

format del | ixf

Specifies if the output file will be created in DEL or IXF format.

lobsinfile

Include LOB data in the output file.

logs <dirlist>

Specify optional list of alternate directories for locating logs.

For more information, see “Specifying additional locations to search for backups and log files” on page 158.

no long Omit LONG and LOB column data from processing.

offline

Indicates to process an online backup in offline mode. Database logs are not processed and any in-flight transactions are not taken into account when extracting the data, therefore the data is not guaranteed to be consistent when using this option. This option is not recommended, but might be useful in certain cases where the user knows that the data in the backup was in a static state or when it is considered acceptable to retrieve the data in this way.

This option has no effect when processing an offline backup.

partition <partition>

Specifies the database partition number to process.

Note: If the option "partition" of "run ox" is not set, then "run ox" will be started on the catalog partition of the target database.

For Database Partitioning Facility (DPF) environments, this option is mandatory. If the database partition does not exist, then DB2 Recovery Expert Command Line Processor (CLP) displays the following error message:

```
ARY4235E: The partition "<db-partition-number>" does not exist.
```

For non-DPF databases, this option is optional and the default catalog partition number is used.

password

Password used with the specified user name.

If the password value that you specify matches a product keyword, then you must enclose the value in quotation marks (").

serverdirectory

Specifies path name on the database server where the extracted data will be written. The user must ensure that they have permission to write into the specified directory.

tables [list]

Specify a list of tables for extracting data, included in quotes, separated by spaces. The table names have the form:

```
[schema.]name
```

where the schema prefix is optional.

tidfid [list]

Specify list of TID:FID pairs for extracting table data. TID is the catalog table space ID. FID is the table File ID.

Format is: TID:FID

Multiple TID:FID pairs can be specified separated by commas or spaces. If spaces are used, the entire option string must be surrounded by quotes.

user User name that can connect to the DB2 instance (optional). If not specified, authentication parameters will be requested at time of connection.

Example:

```
| run ox for database jdbc:db2://verta.abccompany.com:62055/ASDB from 20100905142316
| partition 15 format del tidfid 2:15
| serverdir /data/export/user1
```

run slr

Usage:

```
| run slr for database <jdbc_url> perform
| <create | update | rebuild | drop | prune>
| [user <user>] [password <password>]
| [from <backuptime>]
| [to <endtime>] [backups <dirlist>]
| [logs <dirlist>]
```

Description:

Run the specified command to create or maintain the Schema Level Repository (SLR) for the specified database. The SLR contains history about the database objects over the lifetime of the database within a specific starting and ending point in the recovery logs.

For the database, specify a valid JDBC URL in the following format:

```
| jdbc:db2://<hostname>:<port>/<dbname>
```

Options:

backups <dirlist>

Specify an optional list of alternate directories for locating backups.

For more information, see “Specifying additional locations to search for backups and log files” on page 158.

backuptime

Identifies the backup image timestamp in the following format:

```
YYYYMMDDHHMMSS
```

This timestamp is always in terms of the database server time zone, so it matches the filename of the actual backup image.

create Create new SLR from specified database backup image.

database

User database name to process the SLR for. Specify a valid JDBC URL in the format:

```
| jdbc:db2://<hostname>:<port>/<dbname>
```

drop Drop SLR.

endtime

Timestamp that identifies the timestamp in the database logs to update the SLR to:

```
[utc]YYYYMMDDHHMMSS[.uuuuuu]
```


For more information, see “Timestamp considerations” on page 159.

logs <dirlist>

Specify optional list of alternate directories for locating logs.

For more information, see “Specifying additional locations to search for backups and log files” on page 158.

password

Password used with the specified user name.

If the password value that you specify matches a product keyword, then you must enclose the value in quotation marks (“”).

prune Prune old SLR information before specified timestamp.

rebuild

Re-create existing SLR using specified database backup.

update Update existing SLR to specified end point.

user User name that can connect to the DB2 instance (optional). If not specified, authentication parameters will be requested at time of connection.

Examples:

The following examples create the SLR from a specified backup, and update it until the end of logs (default behavior when end time option is not specified).

```
run slr for database jdbc:db2://verta.abccompany.com:62055/ASDB  
perform create from 20100905142316
```

```
run slr for database jdbc:db2://verta.abccompany.com:62055/ASDB perform update
```

save session

Usage:

```
save session <session_id> [comment <comment>]
```

Description:

Preserve the specified session ID so that the related data is not deleted. The @ symbol can be used to substitute the session ID that was created by the previous run command.

Option:

comment

Short user description for saved session.

session_id

Identifies the related Log Analysis session.

set current datastore with label <label_name>

Usage:

```
set current datastore with label <label_name>  
[user <user> password <password>]
```

Description:

Set the current datastore connection to define the DB2 instance that the following commands will work with.

Options:

| **label_name**
 | User-selected name for the new connection.

| **password**
 | Password used with the specified user name.

| **user** User name that can connect to the DB2 instance (optional).
 | If not specified, authentication parameters will be
 | requested at time of connection.

Example:

alias_on_user1_skala user user1 password passwd1

show current connection

Usage: show current connection

Description:

Display the name of the current datastore connection.

Options:

None.

unsave session

Usage: unsave session <session_id>

Description:

Release the specified saved session ID so that the session data can be deleted and the ID can be reused by a new session.

Option:

session_id

Identifies the related Log Analysis session.

Specifying additional locations to search for backups and log files

DB2 Recovery Expert locates backups and log files by using information that DB2 records in the Recovery History File (RHF). If the information is not available, or if the backups or log files have been moved from their original location, you can specify an alternate disk directory or list of directories that DB2 Recovery Expert can search for them.

Use one of the following formats to specify alternate disk directories for backups or log files:

Format 1

<filetype> <dirlist>

For backups:

- If <dirlist> contains multiple path names, separate them with commas, and enclose them with single quotation marks so that the argument is treated as a single string.
- DB2 Recovery Expert searches the <dirlist> directory for each database partition.

For logs, DB2 Recovery Expert searches as follows:

- On a single partition instance, in the <dirlist> directory.

- On a multi-partition instance, for each partition, as follows:
 - DB2 10.1 and later:


```
<base path> / <instance> / <db name> / <NODE0XXX> / <LOGSTREAM0XXX> / <current chain>
```
 - DB2 version 9.7 and earlier:


```
<base path> / <instance> / <db name> / <NODE0XXX> / <current chain>
```

<base path> is the value that was provided in <dirlist>; it must be a fully qualified directory name.

Examples of format 1:

```
backups /data/backups
```

```
backups "/data/backups1,/data/backups2"
```

```
logs /data/logs
```

Format 2

```
<filetype> partitionnum1:<pathlist1> partitionnum2:<pathlist2>...
```

In this format, <partitionnumN> is a database partition number. If you want the specified path to be used on all partitions, set <partitionnumN> to 1000. (This is a reserved special number that can be used for such cases.)

In <pathlistN>, separate multiple path names for a specific database partition with commas. If the path name contains spaces, then enclose it in single quotation marks so that the argument is treated as a single string.

Note: DB2 Recovery Expert does not support specifying the path with and without the prefix at the same time.

Examples of format 2:

```
backups "0:/tmp/dir0 1:/tmp/dir1 2:'/tmp/dir 2a','/tmp/dir 2b','/tmp/dir 2c'"
```

```
logs "0:/data/logs1 1:/data/logs2"
```

Timestamp considerations

Command option timestamp parameters are usually in the following format:
[utc]YYYYMMDDHHMMSS[.uuuuuu]

The default time zone is controlled by the Java CLP `time.representation` property that is preset to be the same as the database server. To indicate that the time is specified in Universal Coordinated Time (UTC), specify the optional prefix `utc`. The optional suffix `.uuuuuu` is a period followed by a 6-digit microsecond sequence value. You can omit it if a high-resolution timestamp is not required.

Note: The backup timestamp is given in the format `YYYYMMDDHHMMSS` when a specific backup image is identified by using its timestamp for creating the SLR or extracting object data. This exception exists because the timestamp string will match the actual file name that is used by DB2, and the timestamp is interpreted in the time zone of the database server.

Using the CLP to run Log Analysis against remote databases

You can use the CLP to run Log Analysis against remote databases.

Before you begin

Configure the product for Log Analysis and remote database processing.

You might be required to log into the database, depending on the connection properties that the administrator defined for the datastore connection. The user name and password that are required for Schema Level Repository (SLR) operations and Log Analysis are for the remote database (not the local databases, and not the datastore). You can provide datastore credentials by issuing the `set current connection` command, if required.

You must specify fully qualified paths for the location of log files. When locating log files, the product does not add the instance name, database name, database partition, or log chain to the path.

Procedure

1. Launch the DB2 Recovery Expert CLP.
2. Add and activate a datastore connection (issue the `add datastore connection` command, and then the `set current connection` command).

```
add datastore connection <connectionName> to hostname <hostName> for
database <databaseName> port <portNumber> user <remoteUserID>
password <remotePassword>

set current connection <connectionName> [user <remoteUserID>
password <remotePassword>]
```
3. To create the Schema Level Repository (SLR), issue the following command:

```
run slr for database <databaseName> perform <create | update | ...>
[user <remoteUserID>] [password <remotePassword>]
from <backupImageTS> backups <localBackupsPath>
logs <localLogsPath>
```
4. To run Log Analysis, issue the following command:

```
run la for database <databaseName> report <summary | full>
[user <remoteUserID>] [password <remotePassword>]
backups <localBackupsPath> logs <localLogsPath>
```

Related tasks:

“Configuring DB2 Recovery Expert for Log Analysis and remote database processing” on page 78

Review how to set up DB2 Recovery Expert for Log Analysis and remote database processing.

Chapter 10. Considerations and restrictions



This section lists the considerations and restrictions for installing, configuring, and using DB2 Recovery Expert.

The restrictions listed in this section apply to DB2 Recovery Expert for enhanced recovery capabilities, which include dropped object recovery, object versioning, and log analysis. These restrictions do not apply when DB2 Recovery Expert is using standard DB2 recovery operations.

Recovery using DB2 standard utilities is supported for existing objects to the end of logs or a point in time after the most recent minimum recovery timestamp of all table spaces to recover.

Recommendations for using DB2 Recovery Expert:

- If you set the DB2 DATA CAPTURE CHANGES (DCC) option for schema or tables, log activity detail reporting will run significantly faster for those objects.
- DB2 versions 9.7 and later support currently committed semantics with the cursor stability (CS) isolation level to eliminate lock contention. This option causes additional row information to be logged. To significantly improve Log Analysis performance when processing UPDATES, set the currently committed option for any DB2 9.7 or later databases for which you can recover objects or use Log Analysis. The option is set by default on newly created databases, and can be enabled for other databases by setting the **CUR_COMMIT** database configuration parameter.

Topics:

- “Operational limitations”
- “Recovery, Log Analysis, and SQL limitations and considerations” on page 163
- “Limited or unsupported DB2 features, environments, and operations” on page 168
- “Unsupported DDL statements” on page 170
- “HADR considerations” on page 172
- “SET INTEGRITY statement considerations” on page 178
- “MDC table restrictions” on page 173
- “IDENTITY column considerations” on page 173
- “Limitations for tables with LOB and LONG columns and extended row size” on page 175
- “Materialized query table restrictions and considerations” on page 175
- “Log chain restrictions and considerations” on page 177

Operational limitations

Concurrency limitations

DB2 Recovery Expert does not support running multiple concurrent operations for the same database (such as two recoveries, or a recovery and log analysis, or two log analysis processes). For example, do not start recovering a table, and at the same time, from another window, start running log analysis for the same database.

The two processes share the same data areas and temporary files; one process will overwrite the information of the other process. Recovery Expert does not prevent you from attempting to run these kinds of concurrent operations, but will fail to perform either operation successfully.

Note: DB2 Recovery Expert *does* support running concurrent operations on different DB2 instances or different databases within the same instance.

Date and time limitations

Date and time fixed formats

The date and time formats that DB2 Recovery Expert uses are displayed in a fixed format, and are not adjusted based on user region settings.

Effects of changing the system clock

DB2 Recovery Expert for Linux, UNIX, and Windows relies on DB2 changes to the system clock. When major time shifts occur, such as the change to or from Daylight Saving Time (DST), consider the following situations:

- If you execute point-in-time recovery, you must be aware of any significant time shifts.
- Function definitions include the time and date that they were created in the form of a timestamp. At function invocation, DB2 attempts to resolve the function definition. As part of the function resolution, the timestamp value that was logged in the function definition at create-time is checked. If you move the system clock back to a time before the functions were created, DB2 will not resolve references to those functions. In DB2 Recovery Expert for Linux, UNIX, and Windows this could result in the inability to run tasks and in receiving error SQL0440N.

For more information, see <http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=%2Fcom.ibm.db2.luw.admin.ha.doc%2Fdoc%2Ft0054835.html>.

Uploading files using HTTPS connection from the web client

When working with DB2 Recovery Expert using a secure HTTPS connection, uploading files from the browser client using the Firefox, Safari or Chrome web browser is restricted. Microsoft Internet Explorer is supported and is recommended for these situations. Alternatively, you can use an http connection to perform an upload only.

User locale when running DB2 utilities must match DB2 instance user locale

DB2 Recovery Expert reads the Recovery History File (RHF) using the DB2 instance user's environment on Linux and UNIX. If DB2 utilities are run using a different locale, then the entries written in the RHF will use the user's locale. This action causes the RHF to contain entries with differing locales that DB2 Recovery Expert cannot read.

Using database aliases to locate backups

When backups are taken for a database, they are recorded using the name of the database alias (name) that was used with the backup command. If DB2 Recovery

Expert connects to an alternate database alias to perform log analysis or a recovery, unless backups exist for the database under the selected alias name, DB2 Recovery Expert cannot locate the backups.

Databases cataloged using local protocol

DB2 Recovery Expert does not support the use of a database cataloged in an instance that references a database that exists in another instance on the same host using the local protocol. As a workaround, use the TCP/IP protocol to catalog the database and support it as a remote database.

Microsoft Windows installation limitation

On Windows operating systems, Recovery Expert supports only one installation at a time.

SLR build error

When two databases with the same name or alias are associated with a single datastore, an error might occur after the SLR is built. Messages ARY4001E / SQL0811N indicate this issue. The SLR build is successful, but an SQL error is produced when closing the outstanding command line sessions internally. The SLR should be usable after this error occurs. This error might also occur while using the command line processor. To avoid this error, keep databases with the same name in separate datastores.

Related concepts:

“Limited or unsupported DB2 features, environments, and operations” on page 168
DB2 Recovery Expert provides limited or no support for the following DB2 features, environments, and operations.

“Recovery, Log Analysis, and SQL limitations and considerations”
Review the following restrictions and considerations.

Recovery, Log Analysis, and SQL limitations and considerations

Review the following restrictions and considerations.

Specifying log files ranges in MRT mode

You can specify exact boundaries for Log Analysis processing when you use a Minimum Recovery Time (MRT) report mode. Review the following information when specifying a log files range.

When no log range is specified:

- If all logs are accessible within the MRT range, the Log Analysis process reads all archive logs starting with the MRT, and then reads active logs.
- If archive logs are missing within the MRT range, the Log Analysis process reads only the archive logs that are numbered higher than the last missed log, and then reads active logs.
- If the newest archive log is inaccessible or is out of the MRT range, the Log Analysis process reads only active logs.

When only the end log is specified:

- If all logs are accessible within the MRT range and the specified log, the Log Analysis process reads only the archive logs starting at MRT and ending with the specified log.
- If archive logs are missing between the MRT range and the specified log, the Log Analysis process reads only the archive logs that are numbered higher than the last missed log through the specified log.
- If the specified log is inaccessible, the Log Analysis process displays an error indicating the inaccessible log file.

Only the start log is specified:

- If archive logs are missing within the MRT range, the Log Analysis process displays an error indicating the missing log file.
- If all logs are accessible within the MRT range and:
 - The specified log is out of MRT range, the Log Analysis process displays an error indicating the specified log file.
 - The specified log is within the MRT range, the Log Analysis process reads all archive logs starting the specified log, and then reads active logs.

When both the start and end log are specified:

- If archive logs are missing within the MRT range (or the start log if it is numbered higher) and the end log, the Log Analysis process displays an error indicating the missing log files.
- If no archive logs are missing within the MRT range (or the start log if it is numbered higher) and the end log, the Log Analysis process reads only these archive logs.

Note: When the end log is specified, the Log Analysis process does not read active logs, even if the specified log is active.

Translated table space recovery limitation

When a translated table space recovery is performed, any subsequent database recovery that is performed using a backup image that was taken before that recovery will fail at the roll-forward phase that follows the database recovery. This issue occurs because the DML activity that DB2 will try to roll forward will contain information about row identifiers (RIDs) that do not exist in the current database after the database recovery. During the roll forward, DB2 might report the following error: SQL1042C An unexpected system error occurred. SQLSTATE=58004.

As a workaround, recover all table spaces except those that were used in the translated table space recovery, and then recover later table spaces using the backup that DB2 Recovery Expert creates after the translated table space recovery.

Note: This workaround will not recover the system catalog. Therefore, any dropped or re-created objects or table spaces that reuse TIDs/FIDs cannot be recovered because of conflicting TID/FID values.

ATTACH and DETACH operations might cause Recovery processing to fail

Operations that attach or detach data partitions on a partitioned table by using the ATTACH PARTITION or DETACH PARTITION clause on the ALTER TABLE statement can cause a recovery to fail. For example, if a data partition was detached and still exists, the recovery will replay the DETACH operation. The

replay of the DETACH operation will fail because the detached data partition currently exists. Recovery Expert is not able to determine that the detached data partition was at one time part of the data partition table because, at the Recovery point in time, no relationship exists between the data partitioned table and the detached data partition.

DDL activity immediately before an offline backup

When performing a table space recovery using the translated table space restore method, a problem might occur if DDL activity that affects tables that are being recovered occurs during the same second as the starting time of the backup. In this case, DB2 Recovery Expert might be unable to determine whether the DDL activity was performed before or after the backup, resulting in incorrect recovery execution with unpredictable results.

ROW CHANGE TIMESTAMP consideration

Consider the following information if you have tables with `TIMESTAMP NOT NULL GENERATED ALWAYS FOR EACH ROW ON UPDATE AS ROW CHANGE TIMESTAMP`.

The `ROW CHANGE TIMESTAMP` column is a special column used for optimistic locking. The value of the column is always generated by DB2. The column is used to generate a unique value to determine whether a row has been updated since the last time it was queried. For this column type, it is valid to allow the recovery to not restore the exact original value of the column. Using `DEFAULT` to allow DB2 to generate a new value for the column is valid in this case.

Compression dictionary age and renamed tables

DB2 Recovery Expert cannot locate the compression dictionary for renamed tables when the dictionary has been reset more recently than the Schema Level Repository (SLR) coverage range. As a result, DB2 Recovery Expert cannot process detailed row data during Log Analysis.

Log Analysis might quiesce in share mode

You can use Log Analysis as follows:

- to generate a detailed activity report
- to generate Redo or Undo SQL statements
- as part of the Recovery process, when it is necessary to generate Redo or Undo SQL to recover the objects

When Log Analysis is running, it might be necessary to temporarily quiesce table spaces that are related to the log activity that is being processed. The quiesce action occurs when masked update processing occurs and incomplete information is logged. In this situation, Log Analysis must obtain the original row image from either a backup image or a table space container data page.

When obtaining data pages from table space containers, Log Analysis must quiesce the table space in share mode (execute the DB2 command **QUIESCE TABLESPACES FOR TABLE** with the `SHARE` option), and then flush all of the table space pages from the buffer pools. After the table space is flushed to disk, the current data page that is required is read from the table space container, and the table space quiesce is released.

You can disable this processing by specifying the `nopagesaccess` option or by checking the **Don't access current database when processing masked updates** option in the **Options** step of the **Log Analysis** perspective. Taking this action forces exclusive use of backup images to obtain row images for update processing. It is not possible to turn off table space container access for recovery processing.

Log Analysis sessions save only the last related SQL execution

Only the last executed SQL execution session is saved with its related Log Analysis session. Each time that a Run SQL operation is performed, it overwrites the previous SQL session. This is also true for multiple users; if two users load a Log Analysis session and perform the Run SQL operation, only the last execution is saved.

Tables defined with adaptive compression require the DATA CAPTURE CHANGES clause

(DB2 10.1 and later only) Log Analysis detail and SQL generation for logged operations of adaptively compressed tables is not possible because the compression dictionary is not available. Enabling DATA CAPTURE CHANGES (DCC) logs the operations without compression, and allows Log Analysis and Recovery using Redo/Undo SQL of tables that use adaptive compression. You must enable DCC before performing the operations, and you can enable DCC at the schema or table level.

UPDATE operations on partitioned data

When you use the Log Analysis tool to generate Redo or Undo SQL for tables that use the following data organizational schemes, UPDATE operations might be generated as a pair of DELETE and INSERT operations:

- database partitioning (DISTRIBUTE BY)
- table partitioning (PARTITION BY)
- multidimensional clustering (ORGANIZE BY)

This action will occur if any distribution, partitioning, or clustering key column is changed in a way that causes the storage location of the row to change. In this situation, DB2 logs the UPDATE operation as a DELETE of the row from the previous location and an INSERT of the updated row in the new location. This limitation should not affect the result of applying the SQL.

Database initiated operations

DB2 log records do not identify database-initiated operations (such as Referential Integrity and Triggers) separately from user-initiated actions. Log Analysis generates SQL for all operations that are recorded in the database logs as normal DML statements. Execution of the generated SQL might fail or produce invalid data.

Log Analysis transaction count consideration

If rolled back transactions are not reported, you must update the Schema Level Repository (SLR) with a later end range that includes at least one transaction that was committed after the rolled back transactions. Taking this action enables Log Analysis to report on rolled back transactions that were previously missing from the report.

Log Analysis restriction in pureScale environment

When you use DB2 version 10.1 and later, you can run local Log Analysis processing in a pureScale environment. Remote Log Analysis is not supported in a pureScale environment.

Quiet Time and Transaction Start Time accuracy issue

The reported transaction start time and quiet time ranges do not reflect the actual transaction start time. DB2 Recovery Expert obtains the start time from an informational log record that is written at the start of a transaction. DB2 does not update the recorded start time before writing the log record. The start time value is updated for any other activity in the database that uses DB2 virtual time, such as commit, rollback, and online backup complete. Because the transaction start time value is not updated, Quiet Time reports will not include all of the quiet times, and will report shorter quiet time ranges than have actually occurred. Log Analysis reports might report a transaction start time that is the same as the last transaction commit time, even though time might have passed between the last commit and the start of the new transaction.

The following DB2 APARs correct the transaction start time value:

- DB2 V10 Fix Pack 1: APAR IC87105
- DB2 V9.7 Fix Pack 8: APAR IC88046
- DB2 V9.5 Fix Pack 11: APAR IC88047

Related concepts:

“Operational limitations” on page 161

“HADR considerations” on page 172

If your site uses DB2 high availability disaster recovery (HADR), review the HADR considerations for DB2 Recovery Expert processing.

“Limited or unsupported DB2 features, environments, and operations” on page 168
DB2 Recovery Expert provides limited or no support for the following DB2 features, environments, and operations.

“Unsupported DDL statements” on page 170

Review the DDL restrictions for DB2 Recovery Expert.

“Log Analysis features and benefits” on page 4

Log Analysis offers several features and benefits that aid in the recovery process.

“DB2 Recovery Expert features and benefits” on page 3

DB2 Recovery Expert offers several unique and significant features that you can use to improve your DB2 environment.

SLR considerations

SLR creation might fail if archive log files are moved

Creating the schema level repository (SLR) might fail when all of the following conditions exist. When these actions occur, DB2 Recovery Expert might be unable to locate the active log files.

- Archive log files are moved from the default location immediately before you attempt to create the SLR.
- DB2 is archiving log files.

To avoid this issue, use one of the following recommendations:

- Before you create the SLR, wait for the DB2 archive process to finish.
- While you create the SLR, leave at least one archive log file in the default location.
- Before you create the SLR, run following command to ensure that current active logs are archived and the latest logs are moved to the default directory:
db2 "ARCHIVE LOG FOR DATABASE db_name ON ALL DBPARTITIONNUMS"

SLR build error

When two databases with the same name or alias are associated with a single datastore, an error might occur after the SLR is built. Messages ARY4001E / SQL0811N indicate this issue. The SLR build is successful, but an SQL error is produced when closing the outstanding command line sessions internally. The SLR should be usable after this error occurs. This error might also occur while using the command line processor. To avoid this error, keep databases with the same name in separate datastores.

Limited or unsupported DB2 features, environments, and operations

DB2 Recovery Expert provides limited or no support for the following DB2 features, environments, and operations.

DB2 Version 10.5 limitations

DDL-only recovery is fully supported for all new features. Limitations apply to support for other features, as follows.

- Index on expression is not supported for recovery. If encountered, the product ignores the index and issues a warning message.
- To recover column organized tables, you must use the DB2 RESTORE/ROLLFORWARD Recovery plan.
- For databases that contain the DB2 Recovery Expert configuration repository and datastore repository, the value of the database configuration parameter **dft_table_org** must be ROW. The value COLUMN is not supported.
- If encountered during Log Analysis processing, the product ignores column-organized tables and issues a warning message.

Database partition group changes

Table space recovery does not support database partition group changes between the time of the backup and the recovery point in time (PIT).

Recovery between different DB2 releases

DB2 Recovery Expert does not support the recovery of objects between different DB2 releases. DB2 objects must be recovered into the same DB2 version as the original backup. DB2 9.5 is the minimum database that is supported.

Recovery during redistribution

Recovering data for a table space or table while adding or removing partitions is not supported. For Log Analysis update processing to succeed, a backup of the affected table spaces is required following redistribution.

User-defined structured types

DB2 Recover Expert support of user-defined structured types (UDST) has the following limitations:

- DB2 Recover Expert does not support tables or table spaces that contain UDSTs.
- UDST definitions (without the data) can be recovered from the browser client Object Selection window.
- Table space recovery from offline backups is supported for table spaces that contain tables that are dependent on UDSTs as long as the recovery path does not require running Log Analysis to generate Redo or Undo SQL. The browser client does not display recovery paths that involve tables that are dependent on UDSTs.
- DB2 Recovery Expert does not support recovery from an online backup through the browser client application for table spaces that contain tables with UDST columns.

LBAC, RCAC, and XML

If any tables that are selected for recovery have label based access control (LBAC), row and column access control (RCAC), or XML columns, the only recovery path is DB2 RESTORE and ROLLFORWARD. Existing table space recovery to any point in time is supported using DB2 RESTORE and ROLLFORWARD. DB2 Recovery Expert does not offer recovery paths that lead to SQL generation, and issues a message stating that the recovery paths do not include Undo/Redo SQL because the table contains columns that use LBAC, RCAC, or XML.

XML data type and related DDL

Dropped object recovery and Log Analysis do not support tables that contain XML columns. Object versioning does not generate XML column data types.

Global Temporary Tables

Global Temporary Tables are not supported for recovery.

Federated databases

DB2 Recovery Expert does not support federated databases. The definitions are not versioned and DDL generation is not available.

Workload Manager

DB2 Recovery Expert does not support Workload Manager. The definitions are not versioned and DDL generation is not available.

Trusted Context

DB2 Recovery Expert does not support Trusted Context. The definitions are not versioned and DDL generation is not available.

Advanced Copy Services (ACS)

DB2 Recovery Expert does not support ACS snapshot backups.

Order of recovery for dependent objects

DB2 Recovery Expert is aware of dependent objects. Except for noted restrictions, when recovering selected objects, the product also recovers the related and dependent objects. For example, if dropped Table A is selected for recovery, and the table space in which it resides was also dropped, the product first recovers the dependent table space, and then recovers the table. For more information, see “Viewing dependent objects and specifying related tables for recovery” on page 127.

Translated table space recovery for table spaces that are part of a partitioned table definition

A partitioned table is defined by using the PARTITION BY [RANGE] clause on the CREATE TABLE statement. The table can be defined across multiple table spaces. The translated table space recovery method is not supported for table spaces that are part of a partitioned table definition; however, those table spaces can be recovered by using other methods that DB2 Recovery Expert provides.

Recovery of statistics

DB2 Recovery Expert does not recover statistics for dropped objects along with the data. To set appropriate statistics for the recovered objects, run the RUNSTATS utility.

Range clustered tables (ORGANIZE BY KEY SEQUENCE)

- DDL generation does not support the ORGANIZE BY KEY SEQUENCE clause.
- Log Analysis does not process log records specific to range clustered tables.

Related concepts:

“Operational limitations” on page 161

“Recovery, Log Analysis, and SQL limitations and considerations” on page 163
Review the following restrictions and considerations.

Unsupported DDL statements

Review the DDL restrictions for DB2 Recovery Expert.

DB2 Recovery Expert does not support DDL statements as described in this topic.

- CREATE INDEX statements with the following clauses:
 - XMLPATTERN
 - INDEX EXTENSION
- CREATE/ALTER TABLE statements as follows:
 - For the following objects:
 - staging tables
 - materialized query tables (MQTs)
 - user defined structured types (UDST)
 - With the following clauses:
 - ORGANIZE BY KEY SEQUENCE
 - DISTRIBUTE BY REPLICATION
 - SECURITY POLICY

- OPTIONS
- NOT LOGGED INITIALLY
- ALTER TABLE statements with the following clauses:
 - PCTFREE
 - LOCKSIZE
 - APPEND ON/OFF
 - VOLATILE/NOT VOLATILE
 - LOG INDEX BUILD
 - ADD/DROP SECURITY POLICY
- CREATE/ALTER/DROP statements with the following clauses:
 - FUNCTION MAPPING
 - SERVER
 - TYPE MAPPING
 - USER MAPPING
 - WRAPPER
 - AUDIT POLICY
- CREATE TYPE statements for the following types:
 - array
 - cursor
 - structured
- Label Based Access Control (LBAC) statements
- Workload manager statements

Related concepts:

“Recovery, Log Analysis, and SQL limitations and considerations” on page 163
 Review the following restrictions and considerations.

Considerations for DPF environments

Review considerations for using DB2 Recovery Expert in a Database Partition Facility (DPF) environment.

- Do not enter a database-partition-specific ending of a log path because this will be automatically constructed.
- If the specified backup disk location is the same on all database partitions, then to specify just one backup location for all partitions, select **All Database Partitions** and provide the path.
- The log location that is specified for all partitions will be used as a base path. The actual log directory location will be appended with the partition and log chain number values.

Examples of the log directory location appended with the partition and log chain number values are as follows:

For DB2 10.1 and later:

```
on partition XXX: <base path> / <instance> / <db name> /
<NODE0XXX> / <LOGSTREAM0XXX> / <current chain>;
```

For earlier versions of DB2:

```
on partition XXX: <base path> / <instance> / <db name> /
<NODE0XXX> / <current chain>;
```

HADR considerations

If your site uses DB2 high availability disaster recovery (HADR), review the HADR considerations for DB2 Recovery Expert processing.

To support a user database that has HADR enabled, the datastore repository must also have HADR enabled. When the datastore repository is created within the user database, the HADR setup for user database is used. When the datastore repository is created in a *separate* database from the user database, HADR must be set up for the database containing the datastore repository.

Considerations for Schema Level Repository (SLR) operations

- Normally, the SLR minimizes the logging impact of processing by turning off LOB and table operation logging when updating. When BLOCKNONLOGGED is enabled in a HADR environment, increased logging will occur because all operations will be logged.
- SLR Update processing uses a cache of system catalog row images for log reconstruction. This information is stored on disk in the DATA_DIR path. For future SLR updates to succeed, the DB2 Recovery Expert database server components and working directory DATA_DIR must be located on a shared file system that is accessible on the standby server when a failover occurs.

Considerations for Recovery processing

- During Recovery processing, the original NOT LOGGED clause for LOB columns will not be included in generated DDL if BLOCKNONLOGGED is currently enabled. Including the NOT LOGGED clause would cause the recovery to fail. As a consequence, if BLOCKNONLOGGED is disabled in the future, the LOB column will still be logged.
- DDL that is generated by using Log Analysis or the command-line DDL generator will include the NOT LOGGED clause, regardless of the BLOCKNONLOGGED setting. Running this DDL would fail if the database has BLOCKNONLOGGED enabled. DB2 Recovery Expert generates DDL in several places. DB2 will change all current NOT LOGGED columns, such as LOB columns, to LOGGED if BLOCKNONLOGGED is enabled after the columns are defined.
- Recovery processing uses the LOAD utility when recovering dropped tables and recovering a table back to a prior definition that requires dropping the current table. In HADR, when LOAD is used, the load file must be available on the secondary server. When BLOCKNONLOGGED is enabled, LOAD utility operations must include the COPY clause. In this configuration DB2 Recovery Expert will use the COPY YES clause of the LOAD utility. The copy location will be the same location as the backup from which the table data was exported.

Related concepts:

“Recovery, Log Analysis, and SQL limitations and considerations” on page 163
Review the following restrictions and considerations.

“Log chain restrictions and considerations” on page 177

DB2 Recovery Expert cannot process logs across log chain boundaries. Review the following considerations that are associated with this restriction.

IDENTITY column considerations

RESTART value of IDENTITY columns might be incorrect after a recovery

IDENTITY columns using the CACHE clause might leave holes in the sequence during normal application processing. Log Analysis Redo or Undo SQL might generate different values when applied than the original application logic.

The RESTART value of the IDENTITY column might not be correct after the recovery has completed. Before you generate new identity values, make sure that the restart value is correct.

Redo and Undo SQL cannot replicate original row values

When Log Analysis generates Redo or Undo SQL for table columns defined as GENERATED AS IDENTITY, the Redo SQL that replicates INSERT operations, and the Undo SQL that rolls back DELETE operations cannot replicate the complete original row value. This is because INSERT produces a new IDENTITY column value. When you perform SQL on a table containing IDENTITY columns, unexpected results can occur depending on how the IDENTITY column is used within the database. Subsequent SQL statements that depend on the value of the IDENTITY column might also be affected. This is especially true if the IDENTITY column is part of the primary or unique key, because primary or unique keys are used in the WHERE clause of SQL DELETE and UPDATE statements that the Log Analysis tool generates.

Index values limitations

Review the following limitations for index values.

Log Analysis cannot mimic application program logic

The Log Analysis tool cannot mimic application program logic. For example, if an application program exists that does UPDATE (or DELETE) statements that are WHERE CURRENT OF CURSOR, the DB2 log merely records the fact that the update occurred to a particular row. SQL that DB2 Recovery Expert generates cannot be performed against such a row; the WHERE clause only determines the rows against which the UPDATE (or DELETE) will be performed. This can cause problems against non-uniquely keyed tables. Such an UPDATE (or DELETE) operation, when redone, might update many rows.

Tables without a primary or unique key

SQL generation for tables without a primary or unique key will be incorrect for operations that match multiple rows. Log Analysis is not able to generate a predicate that uniquely identifies the row that was modified by the logged operation. For cursor based updates and deletes, it is not possible to generate SQL that matches the application logic.

MDC table restrictions

Recovery restrictions for MDC tables

- During object recoveries that involve the generation of Redo or Undo SQL, DB2 Recovery Expert might be unable to generate complete information for DELETE

operations against multidimensional clustering (MDC) tables when MDC rollout logging is enabled. MDC tables are defined using the ORGANIZE BY DIMENSIONS clause. In DB2 version 9.5 and later, optimized logging is enabled by default. DB2 writes additional log information that allows reporting the correct number of DELETES that occurred during the MDC rollout. During the summary phase of Log Analysis, the Log Analysis tool might be unable to show the exact number of deletes for MDC tables. This issue does not affect the detailed report or SQL generation phases because in those phases, the Log Analysis tool will read the backup and can undo or redo the deletes accordingly. The complete set of DELETE operations will always be reported in the detailed report and in any generated SQL.

- In some cases when MDC rollout logging occurs, DB2 Recovery Expert might not be able to report the complete details or generate SQL statements for the related DELETE operations. This can occur when DB2 Recovery Expert cannot reconstruct the DELETE operations from the logs in the forward direction from a previous backup. For example, when an intervening LOAD, REORG, or ALTER TABLE event occurred between the backup and the MDC rollout event. Because MDC rollout events do not log the complete details of the deleted data, DB2 Recovery Expert cannot reconstruct the DELETES in the reverse direction. If this situation occurs, the detailed report will indicate that a DELETE occurred, but it will not contain the operation details, and corresponding SQL DELETE statements will not be generated.

Log Analysis restrictions for MDC tables

The DB2 Recovery Expert Log Analysis tool might be unable to generate complete information for DELETE operations against multidimensional clustering (MDC) tables when MDC rollout logging is enabled. MDC tables are defined using the ORGANIZE BY DIMENSIONS clause. Optimized logging is enabled by default in DB2 version 9.5 and later.

In DB2 version 9.5 and later, DB2 writes additional log information that allows reporting the correct number of DELETES that occurred during the MDC rollout.

This issue does not affect the detailed report or SQL generation phase. The complete set of DELETE operations will always be reported in the detailed report and in any generated SQL.

In some cases when MDC rollout logging occurs, Recovery Expert might not be able to report the complete details or generate SQL statements for the related DELETE operations. This can occur when it is not possible to reconstruct the DELETE operations from the logs in the forward direction from a previous backup, such as when there is an intervening LOAD, REORG, or ALTER TABLE event between the backup and the MDC rollout event. Because MDC rollout events do not log the complete details of the deleted data, it is not possible to reconstruct the DELETES in the reverse direction. If this situation occurs, the detailed report will indicate that a DELETE occurred, but it will not contain the operation details, and corresponding SQL DELETE statements will not be generated.

Limitations for tables with LOB and LONG columns and extended row size

Restriction for LOB and LONG columns in separate DMS table spaces

DB2 Recovery Expert does not support recovering dropped tables when the table data and the LOB or LONG column data for the table are located in separate backups. This issue can arise if separate table space backups are performed for the table spaces that contain the related data.

Backup requirement for SQL generation

This requirement applies to tables that contain LOB or LONG columns or extended row size. You must perform a full database or table space backup that includes the affected table space following the LOAD. The backup image establishes a new sequence point for SQL generation following the LOAD to ensure that correct SQL is generated for UPDATES and DELETES after a DB2 LOAD.

DB2 Recovery Expert performs a DB2 LOAD as part of the table recovery procedure. To avoid leaving the table space in a backup pending state, Recovery Expert either performs the LOAD with the COPY option (default), or you can select an option to perform a table space backup after the data is loaded.

If you do not perform a backup after a LOAD, Log Analysis might be unable to reconstruct details for UPDATE and DELETE operations that are subsequently performed on the affected table. In this case, Log Analysis will report "NOT AVAILABLE" for LOB or LONG column or extended row size values in the detailed report and in any generated SQL statements that reference those columns.

Materialized query table restrictions and considerations

Review the following restrictions and considerations for materialized query tables (MQT).

MQTs are not supported for recovery

The following restrictions apply to MQTs during recovery:

- MQTs are not listed as tables that can be recovered. However, they are listed as dependencies for tables that are selected for recovery.
- For table space recovery to a point in time (PIT) that is after the Minimum Recovery Timestamp (MRT), the recovery is supported because DB2 RESTORE and ROLLFORWARD processing is used. However, the recovery might fail for dropped table or table space recovery or recoveries involving Undo SQL when MQTs are present.
- Recoveries that require defining the MQT are not supported because DDL generation of MQTs is not supported.
- Dropped table space recovery is not supported if it contains a dependent MQT. In this case, the user must recover each table individually, or select a group of tables that avoids the MQT dependencies.
- For Undo SQL recovery, an MQT can cause the recovery to fail if the base table must be altered in a way that is not allowed when an MQT is present. Most Undo SQL recoveries will work if the base table does not include any columns that are defined with the GENERATED ALWAYS clause. To allow an Undo SQL recovery of a table with a GENERATED ALWAYS column to succeed, the dependent MQT must be manually

dropped. For all Undo SQL recoveries, the dependent MQTs will not be processed as part of the recovery plan.

- The **Recovery** step will warn the user of the presence of unsupported MQTs during recovery plan generation. The dependent MQT will not be processed during plan execution. The dependent MQT must be manually refreshed after the recovery of the base table.
- Dropped table recovery will warn the user of the presence of unsupported MQTs during recovery generation. The dependent MQTs will not be processed during recovery.

The following message will be added for both Undo SQL and dropped object recoveries:

```
ARY0063W: Unsupported MQT {0}.{1} in recovery plan {2} will not be recovered.
```

Explanation:

Materialized Query Tables (MQT) are not supported in recoveries other than table space restore. Processing of the MQT will be skipped during recovery processing.

User Response:

For the recovery to complete the MQT may need dropped. To drop the MQT relationship to the base table run:

```
ALTER TABLE {0}.{1} DROP MATERIALIZED QUERY
```

The message parameters are as follows:

- 0 - MQT Schema
- 1 - MQT Name

Note: The removal of the dependent MQT from the Recovery plan can result in other dependent objects being dropped and not recovered. For example, if a function is dependent on the MQT and another table that is being recovered, the function will be dropped to allow the table to be recovered. The function cannot be recovered due to the dependency on the MQT. The function can be recovered in a separate step after the dependent MQT has been manually defined.

MQTs and staging tables are not supported by DDL generation during dropped object recovery

A dropped MQT or staging table is recovered as a regular table. The user must select the base, MQT, and staging table spaces for recovery.

Log Analysis is unable to determine if operations against an MQT or staging table were database or user initiated. Invalid SQL statements can be generated during recovery and SQL generation.

MQTs are not refreshed during recovery

When performing a recovery that involves a materialized query table (MQT), complete the following tasks:

1. Alter the table to drop the MQT.
2. Recover the base table.
3. Add the MQT and refresh the data.

Table spaces containing MQTs can only be recovered using the DB2 RESTORE API command under certain conditions

A table space that does *not* contain an MQT can be recovered using the following methods:

- Executing the DB2 RESTORE API command.
For more information, see DB2 RESTORE API.

- Executing the translated table space algorithm.
For more information, see translated table space algorithm.
- Recover individual tables in the table space.

If a table space *does* contain an MQT, the only available method to recover information is to execute the DB2 RESTORE API command. You cannot recover information from a table space that contains an MQT by executing the translated table space algorithm or by recovering individual table spaces. If you cannot enter the DB2 RESTORE API command, information cannot be recovered from the table space.

Log chain restrictions and considerations

DB2 Recovery Expert cannot process logs across log chain boundaries. Review the following considerations that are associated with this restriction.

HADR environments

In a high availability disaster recovery (HADR) environment, when a takeover occurs, a new log chain is started on the new primary system. DB2 Recovery Expert cannot process new log chains so that the information from previous chains is preserved for use with the new chains. See also “HADR considerations” on page 172.

Schema Level Repository (SLR)

- DB2 Recovery Expert cannot use the same backup image to create the SLR that was used for a database restore if the database was restored under a different name.
- DB2 Recovery Expert cannot create the SLR from any backup before the restore point, and cannot update the SLR to a point that is not in the current log chain.
- An existing SLR cannot be retained for a database that was restored. After restoring the database, the SLR must be re-created.

Log Analysis

The Log Analysis tool can only use the latest current log chain. The tool cannot cross the database restore point because it cannot determine when the current log chain started or read logs from the previous log before the beginning of the current log chain.

For example, assume that the following log chains existed for a database:
C0000000/S0000000.LOG, C0000000/S0000001.LOG, C0000000/S0000002.LOG,
C0000000/S0000003.LOG, C0000001/S0000002.LOG, C0000001/S0000003.LOG

The current log chain is C0000001, with logs S0000002.LOG and S0000003.LOG. After the last log on chain C0000001 is read, the log chain is decremented, and the logs are processed.

Related concepts:

“HADR considerations” on page 172

If your site uses DB2 high availability disaster recovery (HADR), review the HADR considerations for DB2 Recovery Expert processing.

SET INTEGRITY statement considerations

Processing of SET INTEGRITY statements during recovery

During recovery, the processing of SET INTEGRITY statements might require the creation of exception tables. DB2 Recovery Expert for Linux, UNIX, and Windows defines these tables using the original table name. The table name format is `<schema>.ARY<timestamp><counter><table_name>`, where variables are defined as follows:

- `<schema>` is the original table schema
- `<timestamp>` is the original transaction timestamp
- `<counter>` is a sequential value to ensure a unique table name
- `<table_name>` is the original table name

Recovering a table in a SET INTEGRITY pending state

If a table was in the SET INTEGRITY pending state at the time that it was originally backed up, the table data is recovered to a newly defined table with the same name that is not in the SET INTEGRITY pending state, and conflicting data is written to an exception table named with the suffix "_exception", which is appended to the original table name. Because DB2 does not know the correct state of the data, you must ensure that the data is valid after recovery and account for any data that was written to the exception table when the data was loaded. A warning message is issued instructing you to verify that the data is correct. Refer to the DB2 documentation for more information about exception tables that are used with the LOAD command and the SET INTEGRITY statement.

Log Analysis support for SET INTEGRITY statements

Log Analysis is unable to generate complete SET INTEGRITY statements when creating Redo SQL. In this case, a comment about the SET INTEGRITY statement will be included in the SQL output and a warning message is issued that instructs the user to check the SQL file to verify that it is correct. If necessary, the user can edit the SQL and supply a complete SET INTEGRITY statement with the appropriate clauses.

Chapter 11. Troubleshooting

These troubleshooting topics are designed to provide you with information about DB2 Recovery Expert errors and issues. They also describe documentation that you might need to provide to IBM Software Support for assistance.

Topics:

- “Information required to analyze errors”
- “Information required to analyze performance issues” on page 180
- “Gathering diagnostic information” on page 182

Information required to analyze errors

Before you report a problem with DB2 Recovery Expert for Linux, UNIX, and Windows to IBM Software Support, you need to gather the appropriate diagnostic information.

All errors

Provide the following information for all DB2 Recovery Expert for Linux, UNIX, and Windows errors:

- A clear description of the problem and the steps that are required to re-create the problem
- All messages that were issued as a result of the problem
- Screen captures of pertinent windows
- The name and version of your browser
- The operating system of the client, the DB2 Recovery Expert for Linux, UNIX, and Windows server, and the DB2 server
- The complete version of the DB2 Recovery Expert console component as shown in the browser client "About" window
- The version of DB2 and the type and version of the operating system that you are using

Specific errors

You might need to provide additional information based on the error message that you received or if requested by Support. For information about locating the DB2 Recovery Expert server log files, see the topic about accessing server log files. Provide additional information related to product messages as follows:

Messages ARY000E through ARY2000E

Database server components produce messages in this range, usually during Log Analysis, Recovery, and creating or updating SLR.

Gather the following additional information:

- initial trace results, including all tables that are selected by default
- database server components trace files

Messages ARY4000E through ARY4999E

Application server, DB2 Recovery Expert server components produce messages in this range. These errors usually occur when

- JDBC connectivity issues exist between the DB2 Recovery Expert server and DB2
- Database server components and the datastore repository installation are incorrect, incomplete, or out-of-sync
- Incorrect information was specified by a user (for example, an invalid user ID or password for connecting to DB2)
- Recovery Expert files and repository were improperly configured or configuration was corrupted
- System errors occur due to lack of space or memory
- RXA component or network issues occur
- A product code problem exists

Gather the following additional information:

- ARY-MP log file
- DS_System log file
- hostlogs log files

Messages ARY5000 through ARY8000

The browser UI produces messages in this range. These errors usually occur when a user enters incorrect information (such as a value that is out of range), when the network connection between the browser and the DB2 Recovery Expert server is interrupted, or when a product code problem exists. No additional information is required for messages in this range.

Messages ARY8000E through ARY8999

Product scripts produce messages in this range. These errors usually occur during installation, uninstallation, or verification of the database server components, or during datastore creation, and occur because

- Incorrect information was provided by the user (such as an invalid user name or password).
- The user does not have the privileges required to perform the operation.
- Insufficient disk space was available or system errors occurred on one of the DB2 systems.
- RXA/network related issues occurred.
- A product code problem exists.

Gather the following additional information:

- ARY-MP log file from the DB2 Recovery Expert server.
- screen capture showing the error message
- printed output from the process

Related tasks:

“Accessing server log files” on page 183

You can access log files on the DB2 Recovery Expert server or from the product user interface in a browser window.

Information required to analyze performance issues

When DB2 Recovery Expert tasks such as SLR operations, Log Analysis, or Recovery take longer to complete than expected, IBM Software Support might require information to analyze the cause of the issue.

All performance issues

IBM Software Support might request the following information:

- Output from the **db2support** utility, a DB2 tool that collects information about a specified user database.
- ARY-MP log file. This log file contains information about the task that was submitted for execution and its parameters.
- Screen captures. Include the **Status** step with the most recent status messages.

SLR issues

When an issue occurs during SLR operations such as create, rebuild, or update, the issue probably occurred while the product was reading a backup image or DB2 logs. Support will need the following information:

- Size and location of the backup that was specified for the SLR operation
- Archived log file information as follows:
 - log file location
 - whether the log files are located on disk, TSM, or vendor storage
 - total size of archived log files that were to be processed, starting from the first log file that corresponds to the end log file of the specified backup image
- If the SLR was previously built, then a screen capture of **SLR Info** step is required.

Log Analysis issues

When an issue occurs during Log Analysis operation, Support will need the following information:

- If Log Analysis was launched in MRT mode, an output of the following command for the specified database, which will print MRT timestamp values for all table spaces in your database: **db2 list tablespaces show detail**
- If Log Analysis was launched in regular mode, a screen capture of the SLR range (This information is available in the Object History perspective on the **SLR Info** step for the selected database).
- archived log file information as follows:
 - log file location
 - whether the log files are located on disk, TSM, or vendor storage
 - total size of archived log files that were to be processed, starting from the first log file that corresponds to the end log file of the specified backup image

Recovery issues

When an issue occurs during Recovery operations, Support will need the following information:

- Trace output from the Recovery plan.
- Size and location of the backup that was used for the recovery.
- If the recovery involves Log Analysis, archived log file information as follows:
 - log file location
 - whether the log files are located on disk, TSM, or vendor storage
 - total size of archived log files that were to be processed, starting from the first log file that corresponds to the end log file of the specified backup image

Gathering diagnostic information

Before you report a problem to IBM Software Support, you need to gather the appropriate diagnostic information.

Automatically generating trace files

DB2 Recovery Expert can automatically collect the trace files for a current session from all hosts and partition locations, and then include them in the file that you download to send to Support. If you cannot use the automatic process, you can manually located database server components information.

Procedure

1. In the window where the error occurred, click **Support > Database Server components**.
2. In the **Select tables to trace** window, select one or more of the options as directed by Support.
3. If you have a Problem Management Report (PMR) number from Support, you can enter it in the **PMR Number** field. The default name of the trace file is arytrace.zip. If you provide a PRM number, it is included in the file name.

Related tasks:

“Locating database server components information”

You can manually activate tracing and generate trace files for the database server components:

Locating database server components information

You can manually activate tracing and generate trace files for the database server components:

Procedure

1. In the product user interface, click **Open > Admin > Datastores**.
2. Select the datastore connection that you were using when the problem occurred, and then click **Settings**.
3. From the **Log Level** dropdown menu, choose Debug and note the DATA_DIR property setting, and then click **Apply**.
4. Navigate to the perspective where the error occurred (such as Object history, Log Analysis, or Recovery), and then retry the actions that led to the error. Make note of the session ID that is used.
5. When the error occurs, navigate to the DATA_DIR location on the DB2 server where you performed the action, and then navigate to the following subdirectory:
 - a. DB2 instance name (for example, db2inst1)
 - b. database partition (for example, NODE0000)
 - c. database (for example, SATURN)
 - d. sessions
 - e. numerical session ID (for example, if your session ID is 5, there will be directory "5") For example: /data/opt/IBM/DB2RELUW/db2inst1/NODE0000/SATURN/sessions/5/
6. Zip the entire directory to send to Support.
7. If your database is in a DPF environment, zip the directory for each database partition (that is, navigate to all partitions NODExxxx). The database and session ID must be the same.

8. From the **Log Level** dropdown menu, choose **Error**, and then click **Apply**.

Related tasks:

“Automatically generating trace files” on page 182

DB2 Recovery Expert can automatically collect the trace files for a current session from all hosts and partition locations, and then include them in the file that you download to send to Support. If you cannot use the automatic process, you can manually located database server components information.

“Generating output from the db2support utility”

Generating output from the db2support utility

Procedure

1. Issue the following command:

```
db2support -d <database alias> -cl 0
```

2. Optional: For additional information about the **db2support** utility, issue the following command:

```
db2support -h
```

Related tasks:

“Locating database server components information” on page 182

You can manually activate tracing and generate trace files for the database server components:

Accessing server log files

You can access log files on the DB2 Recovery Expert server or from the product user interface in a browser window.

Related reference:

“Information required to analyze errors” on page 179

Before you report a problem with DB2 Recovery Expert for Linux, UNIX, and Windows to IBM Software Support, you need to gather the appropriate diagnostic information.

Accessing log files on the server

Procedure

1. Log in to your DB2 Recovery Expert server system directly (via ssh, ftp, telnet, or Windows log in).
2. Change your current directory to the location where the DB2 Recovery Expert server was installed. This is the location that was specified during the initial product installation.
3. Locate log files in the following directories in the product installation tree. More than one of each type of log file might exist if log rotation was enabled for the server.
 - server log files in directory logs
The ARY-MP log files are named ARY-MP.0, ARY-MP.1, and so on. The DS_System log files are named DS_System.0, DS_System.1, and so on.
 - hostlog files in directory logs/ARY/hostlogs
4. Zip the entire logs directory and send the compressed file to Support.

Accessing log files in a browser window Procedure

1. Click **Open > Logs**.
2. Select ARY-MP and then click **View log**.
3. Click **View in browser**.
4. Save the contents of the displayed file.
5. Repeat for the DS_System log file.

DB2 Recovery Expert messages

This topic documents the messages that DB2 Recovery Expert issues. Messages are presented in ascending numerical order.

Message parameters that are represented with variables such as {0} and {1} are replaced with relevant text strings when the message is displayed.

ARY0001I **Applying SQL for transaction ID {0}, operation ID {1}, from database partition {2}.**

Explanation: The product is performing the SQL apply phase for the specified transaction, operation, and database partition.

Response: No action is required.

ARY0001W **Operation at LSN/LRI/LSO {0} column {1}, value is truncated.**

Explanation: The value of the indicated column for the operation at the indicated LSN/LRI/LSO has been truncated. This occurs during SQL generation using the CSV format. The SQL statements in the CSV file format include text values as character constant values. The column value exceeded the maximum allowed length of a character constant, currently 32677.

Response: Verify that the generated SQL is acceptable.

ARY0001E **Processing stopped due to lack of available memory.**

Explanation: A memory allocation error occurred.

Response: Make sure that system has enough memory resources, and then repeat the command.

ARY0002I **Applying SQL: {0}{1}{2}{3}**

Explanation: The specified SQL statement is being executed.

Response: No action is required.

ARY0002W **Table {0} has unsupported structured type, ignored.**

Explanation: The table specified has a column with a user defined structured data type.

Response: Table with columns of user defined structured types are not supported.

ARY0002E **Message not found: {0}**

Explanation: The error message with the specified key could not be found.

Response: Refer to the product documentation to make sure that the application is being launched correctly. If you continue to receive this message after attempting to correct the error, contact IBM Software Support.

ARY0003I **Total processed : {0}**

Explanation: The specified SQL processing occurred.

Response: No action is required.

ARY0003E **Internal product error.**

Explanation: An internal product error was encountered when processing a message.

Response: Contact IBM Software Support.

ARY0004I **Succeeded : {0}**

Explanation: The specified number of successful SQL operations completed.

Response: No action is required.

ARY0004W **Unsupported log record found at LSN/LRI/LSO {0}, possible data loss.**

Explanation: Log Analysis processing encountered an unsupported log record at the identified log sequence number, log record identifier, or log stream offset. For DB2 9.7 and earlier, the value is an LSN; for DB2 10.1 and later, it is an LRI (if length is 34 hexadecimal

characters) or an LSO (if length is 16 hexadecimal characters).

Response: Log Analysis processing ignores log records that it does not understand. This message is issued to warn you of possible data loss. If you cannot determine the cause of the error, contact IBM Software Support.

ARY0004E Unable to access resources.

Explanation: The database user name was not specified, but the password was specified.

Response: Specify database user name.

ARY0005I Failed : {0}.

Explanation: The specified number of failed SQL operations occurred.

Response: No action is required.

ARY0005E Unable to format message: {0}.

Explanation: An error was encountered when formatting a product message.

Response: Make sure that the product is being launched correctly. If you continue to receive this message after attempting to correct the error, contact IBM Software Support.

ARY0006I Warnings : {0}.

Explanation: The specified number of warnings occurred during the SQL apply phase.

Response: No action is required.

ARY0006W Inconsistency detected between reported DELETE operations count {0} and actual deleted rows count {1}.

Explanation: The number of DELETE operations due to purging of cluster of MDC table is not equal to the number of actually deleted rows from the MDC table. This might result in the inclusion of dummy operations without details in the report, or some DELETE operations might not be reported.

Response: Contact IBM Software Support.

ARY0007W Data from table TID {0}, FID {1} is truncated.

Explanation: This message occurs because of a limitation in the product. Data is extracted for a table in IXF mode and the converted data exceeds the maximum defined column size for a character column in the table. The process of extracting character data in IXF mode converts table data from the backup image into UTF-8 format when writing it to the IXF file. This can cause data that originally fit in a column when

stored in its original code set to expand past the defined column size after conversion. The likelihood of the problem occurring depends on the original table code page definition and the amount and type of data stored in the columns. It is more likely to occur when the columns are "mostly full." This limitation does not affect recoveries that are performed by using the browser client because table data is processed in DEL format. The problem only occurs if db2ox is used to extract table data in IXF mode and the necessary conditions are met.

Response: If the error occurs, the problem can be avoided by extracting the data in DEL format instead.

ARY0008I Total completed transactions : {0}.

Explanation: The specified number of committed transactions occurred in the SQL apply phase.

Response: No action is required.

ARY0008W No more files.

Explanation: While processing a specific set of log files, specified with the -P parameter, the last log file was encountered. Additional log files are required to complete processing.

Response: Expand the log file range that is specified with the -P parameter to include additional log files.

ARY0009I Total rolled back transactions : {0}.

Explanation: The specified number of rolled back transactions occurred during the SQL apply phase.

Response: No action is required.

ARY0010I Applying SQL statement number {0} on line {1}.

Explanation: The specified statement number/line number from the SQL file is being executed.

Response: No action is required.

ARY0011I Ignored : {0}.

Explanation: The specified number of ignored SQL operations occurred.

Response: No action is required.

ARY0012W For table with TID={0}, FID={1} contents of row {2}, column(s) {3} have been substituted with null/empty values.

Explanation: The table row that is being recovered was inserted before the specified column that is being added to the table, and does not contain a value for the column.

Response: The table row being recovered was inserted

before the specified column that is being added to the table, and does not contain a value for the column.

ARY0013W Incomplete table space recovery, unable to translate indexes for TID={0} FID={1}.

Explanation: The table that was recovered contained more than 300 indexes. The product was not able to translate the internal index IDs. This might result in incorrect query results when the query path uses an index.

Response: Run REORG INDEXES ALL FOR TABLE <table name> for the table that is identified by the TID and FID.

ARY0015W Unable to recover table structure, table name: {0}, LSN/LRI/LSO: {1}.

Explanation: This is an internal error. During a Log Analysis task (either during a Log Analysis task that is initiated from the Log Analysis perspective or a Log Analysis task that is initiated from the Recovery perspective), the Log Analysis tool was unable to obtain the table structure {0} at the specified LSN/LRI/LSO {1}. Because of that condition, Log Analysis was not able to show the details of an operation that affected that table. LSN/LRI/LSO depends on the version of DB2 that you are using. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Contact IBM Software Support.

ARY0016W Unable to recover table row, table name: {0}, LSN/LRI/LSO: {1}.

Explanation: During a Log Analysis task (either during a Log Analysis task that is initiated from the Log Analysis perspective or a Log Analysis task that is initiated from the Recovery perspective) the Log Analysis tool encountered an operation for table {0} at LSN/LRI/LSO {1} that is not recoverable. The operation that Log Analysis attempted to execute does not have the details in the logs to show the details or generate SQL. This could be the normal condition or it could be an internal error. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Contact IBM Software Support.

ARY0017W One or more LONG columns not logged, table name: {0}, LSN/LRI/LSO: {1}.

Explanation: During the Log Analysis operation, some or all LOB data are not logged. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: No action is required.

ARY0018W Unable to recover one or more LONG columns, table name: {0}, LSN/LRI/LSO: {1}.

Explanation: During Log Analysis, the product was unable to get the details for some or all LOB or LONG data for the specified table and specified LSN/LRI/LSO. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Contact IBM Software Support.

ARY0021W Log activity is detected for table with TID = {0}, FID = {1}.

Explanation: This message is displayed during backup reading, and indicates that log activity for the specified table with TID and FID has occurred.

Response: Contact IBM Software Support.

ARY0023W At least one DDL operation was missed during SQL generating. Internal code: {0}.

Explanation: During DDL generation or Recovery processing, at least one DDL statement was not generated properly.

Response: Verify that the generated DDL is accurate. If the DDL is accurate, then no action is required.

ARY0024W Fail occurred during DDL operation generating. Internal code: {0}.

Explanation: This warning message is generated by Log Analysis processing during DDL generation, and indicates that unusual DDL activity was observed.

Response: This could be a normal situation; however, in the case of incorrect results, contact IBM Software Support.

ARY0025W Warning processing SQL statement {0}{1}{2}{3}.

Explanation: This warning message is displayed during the SQL apply phase, and contains the actual warning message.

Response: See the warning message that was returned by DB2 and follow the suggested response.

ARY0026W During rollforward, message: {0} was received.

Explanation: This warning message is displayed during Recovery processing, and contains the actual warning message.

Response: See the warning message that was returned by DB2 and follow the suggested response.

ARY0027W Data were inserted into table {0} in database {1} with warning: {2}.

Explanation: This warning message is displayed during DB2 LOAD or IMPORT calls for the specified table and contains the actual warning message.

Response: See the warning message that was returned by DB2 and follow the suggested response.

ARY0028W LONG data for table {0};{1} has not been found in backup {2}

Explanation: This warning is processed during backup reading and signifies that LOB or LONG data for specified table is not found in specified backup.

Response: This could be a normal situation; however, in the case of an incorrect result, contact IBM Software Support.

ARY0029W {0}

Explanation: This warning message is displayed when Recovery Expert called DB2 API.

Response: Observe DB2 warning message and try to use one of the suggested response .

ARY0030W SLR is not built, log analysis is performed in MRT mode

Explanation: Warning message that are displayed from Log Analysis that signifies that Log Analysis is being run in minimum recovery timestamp mode.

Response: No action is required.

ARY0031W One or more LONG(s) in DML statement for table {0} have been substituted with null/empty value(s).

Explanation: This warning message is displayed during SLR or Log Analysis operations and signifies that one or more LOB or LONG values in a DML statement for specified table was substituted with NULL or empty value.

Response: This could be a normal situation, however in case of incorrect result contact IBM Software Support.

ARY0032W DML statement for table {0} is too long.

Explanation: This warning message is displayed during SLR or Log Analysis operations and signifies that the SQL that needs to be applied is too long.

Response: Contact IBM Software Support.

ARY0033W SLR update processing stopped due to REORG on catalog table FID {0} at LSN/LRI/LSO: {1}

Explanation: During SLR operations, a REORG on the specified catalog table was found that stopped processing. There is no backup for the system catalog data that can be used to properly reconstruct the data. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Use a later backup to re-create the SLR, or contact IBM Software Support.

ARY0034W Log Analysis generates UPDATE or DELETE statement for the table having no unique column. Table name: {0}, TID/FID: {1}/{2}.

Explanation: This warning message is displayed during Log Analysis processing on specified table that signifies that this table does not have primary or unique indexes that can be used in the WHERE clause.

Response: No action is required.

ARY0037W Recovery History File entry not found for table space {0}.

Explanation: During recovery generation, the RHF was not found for the specified table space.

Response: This could be a normal situation, and in this case, additional information for the specified table space (such as table space containers) must be provided on the Recovery Plan step. (see "Selecting a recovery plan" on page 129).

ARY0038W Recovery History File (RHF) entry for table space {0} parsed with error.

Explanation: During recovery generation, the RHF was not found or is corrupted for the specified table space.

Response: This could be a normal situation, and in this case, additional information for the specified table space (such as table space containers) must be provided on the Recovery Plan step (see “Selecting a recovery plan” on page 129).

ARY0039W Unable to generate USING clause for table space {0}.

Explanation: The table space container details were not found. This can occur for a dropped table space without a CREATE or DROP entry in the Recovery History File (RHF).

Response: Provide the table space container details for the table space.

ARY0040W {0} TABLE detected for {1} at LSN/LRI/LSO/Timestamp: {2}

Explanation: The specified operation for the specified table was detected but is not supported. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: No action is required.

ARY0041W Log Analysis end setting is not within the defined SLR range.

Explanation: During Log Analysis processing, either the specified End LSN/LRI/LSO or End Timestamp value is outside of the SLR range. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Modify the value for End LSN/LRI/LSO or End Time stamp.

ARY0042W The process number of open files limit is {0}, which is less than the minimum recommended value of {1} files.

Explanation: The specified value for the minimum number of open files should be equal to or greater than the recommended minimum value.

Response: Modify the specified value to meet the minimum requirement.

ARY0043W The process maximum file size limit is {0} bytes, which is less than the minimum recommended value of {1} GB.

Explanation: The specified value for the maximum file size limit should be equal to or greater than the recommended minimum value.

Response: Modify the specified value to meet the minimum requirement.

ARY0044W The process data segment size limit is {0} bytes, which is less than the minimum recommended value of {1} GB.

Explanation: The specified value for the process data segment size limit should be equal to or greater than the recommended minimum value.

Response: Modify the specified value to meet the minimum requirement.

ARY0045W The process virtual memory limit is {0} bytes, which is less than the minimum recommended value of {1} GB.

Explanation: The specified value for the process virtual memory limit should be equal to or greater than the recommended minimum value.

Response: Modify the specified value to meet the minimum requirement.

ARY0046W Unable to recover one or more variable length columns from extended row, table name: {0}, LSN/LRI/LSO: {1}.

Explanation: During Log Analysis, the product was unable to get the details for some or all variable string data for the specified table and specified LSN/LRI/LSO. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Contact IBM Software Support.

ARY0047W The backup that contains the table having TID/FID: {0}/{1} is in check pending state.

Explanation: During Recovery processing, the product found that the table with the specified TID/FID is in check pending state.

Response: If required, remove check pending state by using a SET INTEGRITY statement. Otherwise, no action is required.

ARY0048W One or more scenarios were not generated because no source backup was found.

Explanation: During Recovery processing, DB2 Recovery Expert could not find backup assets.

Response: If the backup assets are present and should be used for recovery, then use the Options step to specify backup locations, and then regenerate the recovery plan. Otherwise, no action is required.

ARY0050W Column organized table {0} is not supported.

Explanation: Specified table has column organized structure that is not supported.

Response: Column organized tables are currently not supported.

ARY0051E Operating system error has occurred: {0}.

Explanation: The operating system has returned an error during processing. Check the message for a description of the error that occurred.

Response: Correct the operating system error and then try the operation again.

ARY0054W Cannot calculate exact value for restart identity sequence of table: *table_name* .

Explanation: This warning message is displayed when Log Analysis is unable to calculate the exact value for tables with identity columns.

Response: No action is required.

ARY0061W Recovery might fail due to missing log file at {0}.

Explanation: This warning message is displayed during the Recovery that signifies that missing log file was found that could prevent the Recovery.

Response: If Recovery fails find missing log files and replace it into the proper location and repeat the Recovery.

ARY0062W Recovery might fail due to missing backup image file on {0}.

Explanation: This warning message is displayed during Recovery that signifies that backup image is missing and it could possibly prohibit the Recovery.

Response: If Recovery is unsuccessful, find missing backup provided for the Recovery Expert and repeat the Recovery.

ARY0063W Unsupported MQT {0}.{1} will not be recovered.

Explanation: This warning message is displayed during Recovery Generation and signifies that specified MQT object will not be recovered.

Response: No action is required.

ARY0065W Aryox translation map empty

Explanation: This warning message is displayed during table extraction, and indicates that the extraction map was not specified.

Response: Provide the extraction map and repeat the command.

ARY0066W RUNSTATS command against table {0} failed with error: {1}.

Explanation: After the SLR is created, RUNSTATS is run on all SLR tables to ensure that they perform well for subsequent queries. An attempt to run the RUNSTATS command against the specified table failed with the reported underlying error message. This warning message does not mean that the SLR creation failed, but performance of these SLR tables might be affected if the SLR command is not run manually after correcting the issue reported by the underlying error message.

Response: Check the underlying error for more information about the RUNSTATS command failure. Run the RUNSTATS command manually to ensure that the performance of queries against the table is not degraded.

ARY0070W Unable to generate SET INTEGRITY statement at LSN/LRI/LSO {0}

Explanation: During Log Analysis, the SET INTEGRITY statement cannot be generated for the required DDL at the specified LSN/LRI/LSO. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: After the recovery, manually execute the SET INTEGRITY command to verify that the data is correct.

ARY0071W NOT LOGGED INITIALLY encountered in transaction {0} for table {1}.{2} TID/FID {3}.

Explanation: The table was created or altered with the NOT LOGGED INITIALLY clause, and therefore, table activity information is missing. Because information is missing, the result of log analysis or extraction from an

online backup might be wrong. Correct results cannot be obtained from old logs.

Response: No action is required. To avoid this problem in the future, do not use the NOT LOGGED INITIALLY clause.

ARY0075W ATTACH operation to table {0} might cause recovery processing to fail.

Explanation: The ATTACH operation is detected for the specified partitioned table during the recovery plan generation. Depending on the specific activity, some or all recovery scenarios may not be able to detect and properly handle this situation.

Response: No action is required.

ARY0076W DETACH operation to table {0} might cause recovery processing to fail.

Explanation: This warning message is displayed when the DETACH operation is detected for the specified partitioned table during recovery plan generation.

Response: Depending on the specific activity it might be possible that some or all recovery scenarios will not be able to detect and properly handle this situation.

ARY0077W Object {0} was removed from the current scenario.

Explanation: During recovery plan generation, the specified object was restricted and will not be recovered for the specified scenario.

Response: No action is required.

ARY0078W During the backup, the product received a message: {0}.

Explanation: During Recovery processing in the DB2 BACKUP operation, the specified DB2 message was received.

Response: If the recovery finishes successfully, then no action is required. Otherwise, see the DB2 message and follow the suggested response. If required, repeat the recovery.

ARY0079W During the restore, the product received message: {0}.

Explanation: During Recovery processing in the DB2 RESTORE operation, the specified DB2 message was received.

Response: If the recovery finishes successfully, then no action is required. Otherwise, see the DB2 message and follow the suggested response. If required, repeat the recovery.

ARY0080W During the export, the product received message: {0}.

Explanation: During SLR operations or Recovery processing in the DB2 EXPORT operation, the specified DB2 message was received.

Response: If the operation finished successfully, then no action is required. Otherwise, see the DB2 message, follow the suggested response, and then repeat the operation.

ARY0081W Timestamp '{0}' for roll forward point in time is later than time in logs.

Explanation: During Recovery processing in the DB2 ROLLFORWARD operation, roll forward reached the end of the logs. The timestamp that was specified for the rollforward is greater than the time that is recorded in the logs.

Response: Issue the command DB2 LIST TABLESPACES SHOW DETAIL to determine which table spaces are in the rollforward pending state. Issue the command ROLLFORWARD DATABASE database_name TO END OF LOGS AND COMPLETE TABLESPACE(tablespace_name). Separate multiple table space names with commas.

ARY0082W Report end point {0} might be not reached, processing stopped at end log {1}.

Explanation: The specified end log range value was specified with the -H option, and the report end point was specified with the -e option. Log Analysis processing will stop at the end log range value that was specified with the -H option. Log Analysis might not reach the report end point.

Response: Increase the end value that is specified for the log files range so that Log Analysis processing can reach the end of the time range that was specified for the report.

ARY0083W Unable to read data from backup taken at {0}, reason code {1}.

Explanation: An error occurred while the product was reading a required backup. Possible values for reason code include:

- 1 - No such file
- 5 - Access denied
- 9 - Too many open files
- 10 - No space left on device

Response: Make sure that the required backup exists and is accessible. Also make sure that enough product working storage space exists to cache the information that is required from the backup image. For example, when processing compressed tables, the compression dictionary and any needed row images will need to be

cached to disk for processing.

ARY0085W Specified additional log directory '{0}' is skipped; log files cannot be accessed in that location.

Explanation: The specified additional log directory is either empty, does not exist, or cannot be accessed due to conflicting owner or permission settings.

Response: Verify that the additional log directory is correct and that the log files contained in it are readable by the DB2 instance user. Make any necessary changes, and then try the request again.

ARY0086W Specified end log value {0} is greater than or equal to first active log file {1}.

Explanation: The specified end log boundary for Log Analysis processing is greater than or equal to the first active log file. All active logs will be processed by the Log Analysis tool.

Response: If only archive log needs to be processed, then correct end log boundary to be less than first active log file and resubmit the Log Analysis task.

ARY0102E Cannot open backup directory: {0}

Explanation: DB2 Recovery Expert is unable to open the backup directory identified in the message. The backup location directory being added does not have the same shared library name as the existing backup locations.

Response: Set the backup location to specify backups taken with the same shared library and then try again.

ARY0103E Specified backup image does not exist: {0}

Explanation: A timestamp has been entered that does not point to the supported backup image, or the backup file does not exist in the Recovery History File (RHF) or on the user-specified directory.

Response: Provide the correct timestamp and specify the directory that contains the backup file, if its location differs from the RHF (see "Specifying additional locations to search for log files and backups" on page 97).

ARY0104E Target backup image does not exist

Explanation: The target backup image does not exist.

Response: Verify that a backup for the object being recovered exists for the target database and then try again.

ARY0105E Unexpected error reading data from backup taken at {0}, error code {1}.

Explanation: An unexpected operating system error occurred while reading the backup.

Response: Correct the operating system error and then try the operation again.

ARY0106E Unable to locate a backup for SLR creation.

Explanation: A backup image was not found for creation of the SLR. For remote database support, the recovery history file (RHF) is not used, and a local path is required.

Response: Make sure that a backup of the system catalog is available with a valid location in the RHF or that a valid path to the backup is provided.

ARY0108E Unable to locate database archive log file number <number> .

Explanation: DB2 Recovery Expert is unable to locate the specified archive log file because it cannot be found in the location that is recorded in the database Recovery History File (RHF), or in the currently configured archive log location.

For more information, see the section "Archive log files must be accessible using DB2 configuration" in "Operational prerequisites" on page 28.

Response: Perform one of the following operations:

1. Ensure that the log file is located in the specified locations.
2. Update the database RHF entry for the log using the DB2 UPDATE HISTORY command to point to the correct location of the file.
3. Copy any missing log files to an alternate location and instruct DB2 Recovery Expert to look in this location for the files.

ARY0109E Database is not properly configured. To use the product rollforward recovery must be enabled for the database.

Explanation: The database is not properly configured to work with DB2 Recovery Expert.

Response: For a database to work with DB2 Recovery Expert, it must be configured for rollforward recovery. This can be done by setting LOGARCHMETH1 to LOGRETAIN, USEREXIT, DISK, TSM, or VENDOR.

ARY0110E Invalid SQL generation mode {0}.

Explanation: The Log Analysis SQL generation options were not set correctly, as specified in the message.

Response: Verify that the values provided for the Log

ARY0112E • ARY0133E

Analysis -q option are r or u and the optional -f value specifies a valid file name, and then try again.

ARY0112E Invalid commit scope {0}.

Explanation: The value that was provide for the Log Analysis commit scope is not valid.

Response: Set the Log Analysis commit scope to a numeric value greater than 0 and then try again.

ARY0113E Unable to open file "{0}": {1}

Explanation: DB2 Recovery Expert was unable to open the file identified in the message.

Response: Correct the error specified in the message and then try again.

ARY0114E Cannot read file {0}, {1}

Explanation: DB2 Recovery Expert was unable to read the file identified in the message.

Response: Correct the error specified in the message and then try again.

ARY0115E Cannot get file statistics

Explanation: DB2 Recovery Expert was unable to get file statistics.

Response: Correct the error specified in the message and then try again.

ARY0116E The start date and time must be before the end date and time.

Explanation: You have specified a start date and time that occur after the end date and time.

Response: Correct the entry and then try again.

ARY0119E Invalid begin timestamp {0}.

Explanation: The timestamp that is specified in the message is not valid.

Response: Set the begin timestamp to a valid timestamp and then try again.

ARY0120E Specified begin time is greater than the current time

Explanation: The begin time specified is after the current time.

Response: Set the begin time to a value that is less than or equal to the current time and then try again.

ARY0123E Invalid end timestamp {0}

Explanation: The timestamp that is specified in the message is not valid.

Response: Set the end timestamp to a valid timestamp and then try again.

ARY0125E Unable to determine component ID

Explanation: DB2 Recovery Expert was unable to determine the component ID for a log record.

Response: Contact IBM Software Support.

ARY0126E Unable to determine function ID

Explanation: DB2 Recovery Expert was unable to determine the function ID for a log record.

Response: Contact IBM Software Support.

ARY0127E Unable to determine table space ID

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0128E Unable to determine table ID

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0129E Unable to determine action.

Explanation: DB2 Recovery Expert was unable to determine the action type for a log record.

Response: Contact IBM Software Support.

ARY0130E Action {0} is not valid.

Explanation: The action code specified in the message is not valid. Valid action codes for transactions types are I, U, or D.

Response: Set the action code to I, U, or D and then try again.

ARY0132E FID {0} is not valid.

Explanation: The file ID (FID) specified in the message is not valid.

Response: Set the FID to a valid numeric value and then try again.

ARY0133E TID {0} not valid.

Explanation: The table space ID (TID) specified in the message is not valid.

Response: Set the TID to a valid numeric value and then try again.

ARY0135E Database partition number {0} is not valid.

Explanation: The database partition number specified in the message is not valid.

Response: Set the database partition to a valid numeric value and then try again.

ARY0136E Bad Local Transaction ID {0}

Explanation: The Local Transaction ID specified in the message is not valid.

Response: Set the value for the Transaction ID filter to a value that is 12 characters in length and then try again.

ARY0137E Bad content level {0}

Explanation: The value provide for the content level is not valid. Valid levels are 0 - maximum content, 1 - high content, 2 - medium content, 3 - minimum content, 4 - silent mode.

Response: Set the value for the content level to 0, 1, 2, 3, or 4, and then try again.

ARY0138E Content level {0} is not valid.

Explanation: The value that was provide for the content level is not valid. Valid levels are 0 - maximum content, 1 - high content, 2 - medium content, 3 - minimum content, 4 - silent mode.

Response: Set the value for the content level to 0, 1, 2, 3, or 4, and then try again.

ARY0141E Required parameter missing

Explanation: The block buffer parameter was not provided.

Response: Contact IBM Software Support.

ARY0142E Table is not valid; unable to recover.

Explanation: The table that was specified for recovery is not a valid table. DB2 Recovery Expert cannot recover it.

Response: Verify that the table exists in the backup image that is being used for the recovery and then try again.

ARY0144E Incorrect value for masked update processing option. Must set to do processing in backward or forward direction.

Explanation: You specified an incorrect value for the -G option which identifies the direction for masked update reconstruction.

Response: For the -G option, you must specify either b for roll back or f for roll forward.

ARY0145E Unable to reconstruct masked update details for table TID = <TID>, FID = <FID>. Local Transaction ID: <transaction ID> No reconstruction path available

Explanation: DB2 Recovery Expert was unable to reconstruct the details of a masked update operation for the transaction and table identified in the message.

Response: If the operation details are required, ensure that the necessary backup or database resources are available for processing by Log Analysis.

ARY0146E Unable to reconstruct update details by rollforward

Explanation: DB2 Recovery Expert was unable to reconstruct the update details using the roll forward process.

Response: Attempt to reconstruct the update details using a roll back process.

ARY0147E Unable to reconstruct update details by rollback

Explanation: DB2 Recovery Expert was unable to reconstruct the update details using the roll back process.

Response: Attempt to reconstruct the update details using a roll forward process.

ARY0149E Translation map is not valid.

Explanation: The translation map specified for the -X option is not valid. The syntax for this option is: -X "OLDTID1:OLDFID1[-NEWTID1:NEWFID1 [OLDTID2:OLDFID2[-NEWTID2:NEWFID2]] ...]" For table extract mode only the old TID:FID pairs are required. For table space recovery the new TID:FID pair is required for translation.

Response: Correct the translation map value and then try again.

ARY0158E Cannot continue without connection

Explanation: DB2 Recovery Expert cannot continue without a valid connection to a DB2 database.

Response: Connect to the DB2 database that you want to work with and then try again.

ARY0159E Cannot continue without table space ID

Explanation: DB2 Recovery Expert cannot continue without the table space ID.

Response: Specify the table space ID and then try again.

ARY0161E Cannot write table

Explanation: An error was encountered when writing to a table data page.

Response: Verify that the current user has write permissions for the default working directory, or the table output directory if using table extract mode, and then try again.

ARY0162E Cannot read table

Explanation: An error was encountered when reading a table data page.

Response: Verify that the current user has read permissions for the underlying database files and then try again.

ARY0165E Unable to determine table name

Explanation: An error occurred when attempting to determine table name.

Response: Verify that the table owner and name are correct and then try again.

ARY0171E Cannot read backup DMS table space

Explanation: An error was encountered when reading DMS table space data from a backup image.

Response: Verify that the table space that is being processed exists in the backup image that is being used, and then try again. If you continue to receive this message after attempting to correct the error, contact IBM Software Support.

ARY0172E Recovery History Scan is not open

Explanation: The Recovery History File (RHF) has not been opened for read access.

Response: Contact IBM Software Support.

ARY0174E Unable to locate the TID/FID for the specified table in the backup image

Explanation: The specified table name does not exist in the backup image.

Response: Verify that the specified table name exists in the backup image and then try again.

ARY0175E Tablespace recovery: The source and target backup images are not compatible

Explanation: For one or more of the following parameters, the source backup image is not the same as the target backup image:

- buffer size
- session count
- compress

Response: Verify that the parameters for the source and target backup images are the same, and then try again.

ARY0176E Log file IDs region is not valid.

Explanation: The values that were specified for the set of log files to be processed are not valid. The syntax for the -P argument is: { [n][A] | [n]:[m|A]} where 'n' and 'm' represent archive log numbers and 'A' is a constant that means active logs.

Response: Correct the values for the set of log files to be processed and then try again.

ARY0177E Database logging configuration was recently changed. Database might be in backup pending state. Disconnect all applications and try again.

Explanation: The database logging configuration was recently changed and has invalidated the current connection information.

Response: Disconnect from the database and then try again.

ARY0178E End LSN/LRI/LSO {0} must be greater than start LSN/LRI/LSO {1}.

Explanation: The LSN/LRI/LSO end value that was specified is less than the start value that was specified.

Response: Specify correct start and end values, and then repeat the command.

ARY0179E Range begin and end values are not valid: start LSN/LRI/LSO is greater than end LSN/LRI/LSO.

Explanation: During Log Analysis processing, the specified Start LSN/LRI/LSO value is larger than the End LSN/LRI/LSO. LSN/LRI/LSO depends on your

version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Specify correct Start and/or End LSN/LRI/LSO and repeat the command.

ARY0180E **Begin and end points must have the same type.**

Explanation: During Log Analysis, the beginning and ending point value types (LSN or LRI) are different.

Response: Modify the beginning and ending point values to have the same type (LSN or LRI), and repeat the command.

ARY0181E **Internal error in file {0} line {1}**

Explanation: An internal DB2 Recovery Expert error has been encountered in the file specified in the message.

Response: Record the file and line number and contact IBM Software Support.

ARY0182E **Wrong export format: {0}**

Explanation: The export format specified in the message is not valid. The only supported format for table extract is IXF.

Response: Set the value for the -e option to IXF and then try again.

ARY0183E **Wrong argument(s)**

Explanation: The arguments used for the tool being run are not valid.

Response: Refer to the product documentation for the supported arguments for the tool that is being run, and then try again.

ARY0184E **Start and End timestamp/LSN/LRI/LSO must be of the same type.**

Explanation: An error was encountered when attempting to issue a base checkpoint for log archive metadata.

Response: Verify that the current user has write permission for the log archive metadata directory and then try again.

ARY0186E **Attempt to create sysstorages for non-catalog database partition without having data from catalog database partition**

Explanation: Log Analysis is being run on a

non-catalog database partition with the update archive log meta data option set, but no data exists for the catalog database partition.

Response: Run Log Analysis on the catalog database partition with the update archive log meta data option set and then rerun on the non-catalog database partition.

ARY0188E **Data block is too small**

Explanation: This is an internal Recovery Expert error.

Response: Contact IBM Software Support.

ARY0189E **Data block is too big**

Explanation: This is an internal Recovery Expert error.

Response: Contact IBM Software Support.

ARY0191E **Size of block is not valid**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0192E **Specified key is not valid: {0}.**

Explanation: An incorrect parameter was specified.

Response: Remove the incorrect parameter and repeat the command.

ARY0193E **Path name is too long: {0}**

Explanation: Specified file or path name is too long.

Response: Move or modify the path name to shorten it, and repeat the command.

ARY0194E **Cannot open configuration file**

Explanation: The Recovery Expert configuration file cannot be opened.

Response: Verify that the configuration file exists and repeat the command.

ARY0198E **Unsupported LFH version: {0}**

Explanation: New version of log file header has been detected that is not supported.

Response: Contact IBM Software Support.

ARY0199E **Database with the name: '{0}' does not exist.**

Explanation: Specified user database does not exist.

Response: Verify that the database name is correct and repeat the command.

ARY0200I Generate scenario based on {0}

Explanation: This message is for informational purposes only.

Response: No action is required.

ARY0201I There is no CREATE TABLE in SLR to recover the table: {0} from logs

Explanation: In the SLR create table, DDL is not found for the specified table.

Response: No action is required.

ARY0201E Archive log file {0} seems corrupt

Explanation: Specified archive log file is corrupted.

Response: Verify that the specified archive log file is correct or replace it with the backup version of this file.

ARY0203E Inconsistency in archivelog at file '{0}'.

Explanation: The sequence of the archive log file is incorrect.

Response: Verify that log files that are provided for Log Analysis are sequential and that there is no missing log files. Repeat the command.

ARY0204E Buffer with size {0} is too large. Log file '#{1}' seems corrupt.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0206E Logfile '{0}' open error '#{1}': {2}

Explanation: Specified log file is missing.

Response: Locate and replace the log file and repeat the command.

ARY0207E Logfile '{0}' page read error '#{1}': {2}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0208E File name while retrieving: {0} is not valid.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0209E File I/O error: {0}

Explanation: I/O internal error message

Response: Contact IBM Software Support.

ARY0210I No backup resources are available to recover data for the table: {0} by extract/load/redo-SQL between {1} and {2}.

Explanation: The product could not find a backup image from which to extract data to recover the specified table.

Response: If the backup is required for this scenario, and is available, then specify this backup for Recovery Expert and regenerate the Recovery Plan.

ARY0210E Not attached

Explanation: Current active database is not attached.

Response: Contact IBM Software Support.

ARY0211I No backup resources are available to recover data for the table: {0} by extract/load/redo-SQL on database partitions: {1}.

Explanation: Backup images are missing for the specified table on the specified database partitions.

Response: If backup images exist for the specified database partition, specify the backups to Recovery Expert and regenerate the recovery plan.

ARY0212I No backup resources are available to recover data for the table: {0} by extract/load/redo-SQL on database partitions: {1} due to some restrictions.

Explanation: During recovery plan generation, no backups were found for the specified table for the specified scenario because of the specified restrictions.

Response: If disk backups are available for the specified table and are not listed in the Recovery History File (RHF), and should be used for the recovery, then specify these disk backup locations and regenerate the recovery plan (see "Specifying additional recovery options" on page 126).

ARY0212E Unknown backup otype: {0}

Explanation: Unsupported backup operation type was found.

Response: Contact IBM Software Support.

ARY0213I No backup resources are available to recover data for the table: {0} by extract/load/redo-SQL due to some table modifications.

Explanation: The product could not find the backup image from which to extract data to recover the specified table.

Response: If the backup is required for this scenario, and is available, then specify this backup for Recovery Expert and regenerate the Recovery Plan.

ARY0213E Backup image filename {0} is not valid.

Explanation: The product detected a backup image file name that is incorrect or not valid.

Response: Verify that the backup image file name is named correctly and repeat the command.

ARY0220I Cannot recover the table: {0} that does not exist by recovery from current table method.

Explanation: The table does not exist in the current catalog.

Response: No action is required.

ARY0221I Cannot recover the table: {0} from current table. Between PIT and current table there are some table geometry changes at {1}.

Explanation: The recovery point in time (PIT) is earlier than the table creation time.

Response: No action is required.

ARY0222I Cannot recover the object: {0} by Undo from current method. There are some restrictions created by recover of previous objects.

Explanation: The specified object cannot be recovered by UNDO SQL.

Response: No action is required.

ARY0223I Recovery of table {0} with LONG data or extended row size is not available for this scenario.

Explanation: The specified table contains LONG data and this table cannot be recovered by this scenario.

Response: No action is required.

ARY0224I Cannot recover the table: {0} that is not accessible by Recovery from current table method.

Explanation: The specified table is contained in the table space that is in rollforward or backup pending state.

Response: Verify that the table space is accessible, and then regenerate the Recovery plan.

ARY0224E Unable to allocate buffer size {0}

Explanation: This is a memory allocation error.

Response: Make sure that DB2 system has enough memory resources and repeat the command.

ARY0225I The recovery plan was not generated because the SLR information for the object {0} is not current.

Explanation: The SLR information for the specified object was modified after the last SLR update.

Response: Update the SLR, and then regenerate the recovery plan.

ARY0228E Failed to create directory {0}.

Explanation: Recovery Expert could not create the specified directory because not enough free space is available on the DB2 system, or privileges are not set correctly.

Response: Correct the disk space or privileges issue, and then repeat the command.

ARY0229E Failed to create export file '{0}' of type '{2}' in '{1}'.

Explanation: During data extraction for a backup image, the specified export file could not be created in the specified directory due to insufficient privileges or insufficient disk space on the DB2 system.

Response: Correct the privileges or disk space issues and then repeat the command.

ARY0230I No backup resources are available to translate the table space: {0} between {1} and {2}.

Explanation: The product could not find a backup image for the specified table space for the specified time frame.

Response: If this scenario needs to be generated and the backup images for the specified time frame are available, then specify these backup images for Recovery Expert and regenerate the recovery plan.

ARY0230E Failed to open table space with TID = {0}

Explanation: Table space data entry was not found in the backup image for the specified table space.

Response: Contact IBM Software Support.

ARY0231I Backups for table space: {0} translation do not contain any tables.

Explanation: Backups for the specified table space do not contain table data.

Response: No action is required.

ARY0231E Calculated OTR Extent Number {0} more than OTR Extent map entries {1}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0232I Cannot recover table space: {0} due to a part of partitioned table {1} in that table space or dependent table spaces.

Explanation: The specified table space cannot be recovered because of the specified partition table partitions that reside in other table spaces.

Response: If the specified table space needs to be recovered, then specify all of the dependent table spaces where the specified partitioned table resides and regenerate the Recovery plan.

ARY0232E Backup taken at {2} on partition {3} does not contain object with FID = {1} in the table space with TID = {0}.

Explanation: The table with the specified ID does not exist in the specified backup.

Response: Verify that the table is actually contained in the backup. If not, remove that table from the extraction object set.

ARY0233I Restore of table space: {0} is not supported with HADR.

Explanation: In a high availability disaster recovery (HADR) environment, table space RESTORE is restricted by DB2. Recoveries that use table space RESTORE are not valid recovery paths.

Response: No action is required.

ARY0233E Unable to allocate memory OTR buffer size {0}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0234E Error verifying OT page against OT EMP page

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0235E Data object in OTR does not exist

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0238E Error verifying data EMP page against the Object Table Record

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0239E Error verifying If EMP page against Object Table Record

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0240I No backup resources are available to recover the table space: {0} by API between {1} and {2}.

Explanation: The product could not find a backup image to recover the specified table space by using DB2 API for the specified time frame.

Response: If specified table space needs to be recovered using this scenario and backup images exists for specified time frame then specify those backup image locations and regenerate the Recovery Plan.

ARY0240E Error verifying lob EMP page against Object Table Record

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0241I Cannot recover table space {0} by API, it was dropped at {1}.

Explanation: The specified table space cannot be recovered by DB2 API recovery scenario because it was dropped at the specified time stamp.

Response: No action is required.

ARY0242E Error getting long data

Explanation: Error message that is displayed during processing of LOB or LONG data and signifies that LOB or LONG data cannot be undone by UNDO SQL.

Response: No action is required.

ARY0243E Error accessing file {0}

Explanation: DB2 Recovery Expert encountered an error when accessing the file that is identified in the message.

Response: Make sure the log files are in the expected location and accessible.

ARY0244E Error verifying DAT0 page against Object Table Record

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0245E Unable to open table null page {0}, size {1} from backup

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0246E Table with TID = {0}, FID = {1} not found in backup image

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0250I No backup resources are available to recover the data for table: {0} by extract/load/undo-SQL between {1} and {2}.

Explanation: The product could not find backup images to recover the specified table during the specified time frame.

Response: If this scenario needs to be generated and backup images are available for specified time frame then specify these backup locations and regenerate the Recovery Plan.

ARY0250E Unable to access table TID={0}, FID={1}

Explanation: Table descriptor was not found in either source or template backup images.

Response: Contact IBM Software Support.

ARY0251I No backup resources are available to recover the data for table: {0} by extract/load/undo-SQL on database partitions: {1}.

Explanation: During recovery generation, backup images were not found to recover the specified table on the specified database partitions.

Response: If backup images for the specified database partitions exist, and this scenario must be generated, then specify the backups for the recovery and regenerate the recovery plan.

ARY0251E Tablespace recovery: Source table (tid={0}, fid={1}) is not compatible with Target table (tid={2}, fid={3})

Explanation: During TSR recovery, table descriptions between the source and template backups were found to be incompatible.

Response: Contact IBM Software Support.

ARY0252I No backup resources are available to recover the data for table: {0} by extract/load/undo-SQL on database partitions: {1} due to some restrictions.

Explanation: During recovery generation, backup images were not found to recover the specified table on the specified database partitions.

Response: No action is required.

ARY0252E Tablespace recovery: Indexes of source table (tid={0}, fid={1}) are not compatible with indexes of target table (tid={2}, fid={3}).

Explanation: During recovery generation and before translation, the index sets of the source and template backup images were found to be different. Translation cannot occur and execution of the recovery scenario stops.

Response: Try recovering this database using a different recovery scenario.

ARY0253I No backup resources are available to recover the data for table: {0} by extract/load/undo-SQL due to some table modifications.

Explanation: No backup images are available to recover the specified table because the table definition was changed.

Response: No action is required.

ARY0255E Recovery is not possible, because user defaults were added to source table (tid={0}, fid={1})

Explanation: During TSR recovery, the user default data were found in the source backup.

Response: Try recovering the table space using different recovery scenario.

ARY0258E Unable to read database configuration file: {0}

Explanation: Recovery Expert was not able to find the database configuration file or privileges are not sufficient.

ARY0261E • ARY0287E

Response: Make sure that the configuration file exists and the privileges are set correctly, and then repeat the command.

ARY0261E Cannot create file in working directory: {0}

Explanation: Files cannot be created in specified working directory because of insufficient disk space or insufficient privileges.

Response: Correct the disk space or privileges issue and then repeat the command.

ARY0262E Cannot create file in datacache directory: {0}

Explanation: The cache file with backup data cannot be created in the specified directory because of insufficient disk space or insufficient privileges.

Response: Correct the disk space or privileges issue and then repeat the command.

ARY0264E Tablespace recovery: Unable to open option file {0}: {1}

Explanation: During TSR recovery, the specified option file could not be opened because of insufficient disk space or insufficient privileges.

Response: Correct the disk space or privileges issue and then repeat the command.

ARY0265E Tablespace recovery: Code page different

Explanation: During TSR recovery, the source backup code page and template backup code page were not compatible.

Response: Contact IBM Software Support.

ARY0267E Tablespace recovery: Source and latest backup are equal

Explanation: During TSR recovery, no template backup was available to perform the recovery.

Response: Contact IBM Software Support.

ARY0268E Could not determine tablespace info

Explanation: During table space recovery, no information was available for the specified table space name.

Response: Contact IBM Software Support.

ARY0270E Cannot translate table spaces

Explanation: During TSR recovery, the table count for the specified table space did not match between the source and the template backups.

Response: Contact IBM Software Support.

ARY0271E Tablespace recovery: Table space count is zero

Explanation: During TSR recovery, no table spaces were specified for the recovery.

Response: Specify table spaces that need to be recovered and then repeat the command.

ARY0272E Not BPS Page

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0273E Page unsupported

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0281E Numeric pattern ({0}) is not valid.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0283E Cannot create file in directory: '{0}'.

Explanation: Files could not be created in the specified directory because of insufficient disk space or insufficient privileges.

Response: Correct the disk space or privileges issue and then repeat the command.

ARY0284E Cannot read pattern from file '{0}'

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0286E FileStorage '{0}': new size for storage: {1} is not valid.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0287E FileStorage '{0}': block size read error.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0288I Recovery of table {0} restricted due to absence of primary key or unique constraint.

Explanation: The table without a primary key or unique constraint has been detected during the plan generation and it cannot be recovered.

Response: If this table needs to be recovered, the only way to recover it is to recover the whole table space where this table resides by using the DB2 API Recovery scenario.

ARY0288E FileStorage '{0}': error reading record offset {1} from index file.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0289I Recovery of table {0} restricted due to RESTORE.

Explanation: The RESTORE operation has been detected during the plan generation for a particular table and one or more recovery scenarios cannot be performed.

Response: If this table must be recovered, then the recovery PIT must be modified so that DB2 Recovery Expert can recover this table. For example, the PIT can be changed to a point before the RESTORE operation was initiated.

ARY0289E FileStorage '{0}': file open error: {1}.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0290I Cannot recover at least one object by scenario: {0}

Explanation: This information message is displayed during recovery, and indicates that at least one object cannot be recovered by the specified scenario.

Response: No action is required.

ARY0290E FileStorage '{0}': Header read error.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0291I Recovery of MQT table {0} is restricted.

Explanation: The MQT table is detected for the recovery and it cannot be recovered. DB2 Recovery Expert will not re-create such an object.

Response: If this MQT needs to be recovered, the only way to recover it is to recover the whole table space

where this MQT resides by using the DB2 API Recovery scenario.

ARY0291E FileStorage '{0}': Read-only storage.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0292I Recovery of table {0} having XML column is restricted.

Explanation: A table with an XML column has been detected. DB2 Recovery Expert does not support tables with XML columns.

Response: If this table must be recovered, the whole table space must be recovered where this table resides by using the DB2 API Recovery scenario.

ARY0292E FileStorage '{0}': RestoreState found mismatch between file size ({1}) and DataEnd field ({2}).

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0293I Recovery of table {0} having SECURITY POLICY is restricted.

Explanation: An LBAC object is detected for the recovery and it cannot be recovered.

Response: If this object must be recovered then the whole table space must be recovered by using the DB2 API Recovery scenario.

ARY0293E Archive Log Inventory Data (ALID) does not match the product release and needs to be rebuilt.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0294I Recovery of table {0} having UDST column is restricted.

Explanation: A table with a User Defined Structured Type (UDST) column is detected for the recovery and it cannot be recovered.

Response: If this table must be recovered, then the whole table space where this table resides must be recovered by using DB2 API Recovery scenario.

ARY0295I Recovery of table {0} restricted due to REORG TABLE.

Explanation: The REORG operation has been detected during the plan generation for a particular table and one, or all, recovery scenarios cannot be performed.

Response: If this table must be recovered then the recovery PIT must be modified so that DB2 Recovery Expert can recover this table. For example, the PIT can be changed to a point before the REORG operation was initiated.

ARY0296I Recovery of table {0} restricted due to LOAD TABLE.

Explanation: The LOAD TABLE operation has been detected during the plan generation for a particular table and one, or all, recovery scenarios cannot be performed.

Response: If this table needs to be recovered then the recovery PIT must be modified so that DB2 Recovery Expert can recover this table. For example, the PIT can be changed to a point before the LOAD operation was initiated.

ARY0297I Recovery of table {0} restricted due to TRUNCATE TABLE.

Explanation: This message issues for scenarios:

1. Translated TBS restore from offline backup and rollforward/rollback
2. Extract data from backup image and generate redo/undo SQL
3. Re-create object definitions from SLR and generate redo SQL
4. Generate undo SQL from current state

Response: If this table needs to be recovered then the recovery PIT must be modified so that DB2 Recovery Expert can recover this table. For example, the PIT can be changed to a point after the TRUNCATE operation was performed.

ARY0297E Storage '{0}': illegal Storage copy-constructor called.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0298I Recovery of COLUMN-ORGANIZED table {0} is restricted.

Explanation: The COLUMN-ORGANIZED table is detected for the recovery and it cannot be recovered. DB2 Recovery Expert will not re-create such an object.

Response: If this table needs to be recovered, the only way to recover it is to recover the whole table space where this table resides by using the DB2 API Recovery scenario.

ARY0298E Storage '{0}': illegal Storage copy-operator called.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0299E Storage '{0}': Index={1} out of bounds [0..{2}]

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0300E Storage '{0}': must be opened in read/write mode.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0301I Begin LSN/LRI/LSO for database partition {0} set to {1}

Explanation: During Log Analysis in MRT mode, the start point will be set to the specified LSN/LRI/LSO. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: No action is required.

ARY0301E Storage '{0}': must be opened.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0302I Accessing backup image taken at {0} to obtain row images for logged operation reconstruction.

Explanation: Incomplete information was logged for one or more DML operations. Some operations that log incomplete information are UPDATE, MDC Block Delete, and LOB delete. To reconstruct the complete operation details, DB2 Recovery Expert for Linux, UNIX, and Windows will read a row image, LOB data, or both from a backup image. Reading a backup image will affect processing time.

Response: No action is required.

Note: To improve Log Analysis detail processing performance for future logged operations, complete the following tasks:

- Set the DB2 DATA CAPTURE CHANGES (DCC) option for schema or tables to reduce the time required for log activity detail reporting for those objects.

- In DB2 versions 9.7 and later, use currently committed semantics for cursor stability transaction isolation.

ARY0302E Storage::Open({0}, {1}) called with empty name argument

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0303I Accessing containers for table space {0} to obtain row images for logged operation reconstruction.

Explanation: Incomplete information was logged for one or more DML operations. Some operations that log incomplete information are UPDATE, MDC Block Delete, and LOB delete. To reconstruct the complete operation details, DB2 Recovery Expert for Linux, UNIX, and Windows will read a row image from the current database table space containers. Reading the current database requires that each table space read be quiesced in share mode.

Response: No action is required.

Note: To improve Log Analysis detail processing performance for future logged operations, complete the following tasks:

- Set the DB2 DATA CAPTURE CHANGES (DCC) option for schema or tables that are being processed.
- In DB2 versions 9.7 and later, use currently committed semantics for cursor stability transaction isolation.

ARY0305I Minimum Recovery Timestamp unknown.

Explanation: The database was made recoverable after a table space was created, and no database activity has updated the minimum recovery timestamp (MRT) or LSN/LRI/LSO value. In this case, the MRT and LSN are zero. Log Analysis will start reading logs from log extent zero when the database was first made recoverable. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Provide a value for the begin timestamp of the report.

ARY0306I Begin LSN/LRI/LSO for database partition {0} set to zero.

Explanation: The database was made recoverable after a table space was created, and no database activity has updated the minimum recovery timestamp (MRT) or

LSN/LRI/LSO values. In this case, the MRT and LSN/LRI/LSO are zero. Log Analysis will start reading logs from log extent zero when the database was first made recoverable. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Provide a value for the beginning timestamp of the report.

ARY0312E Missing required parameter: database alias

Explanation: The database alias was not specified.

Response: Specify the database alias and then repeat the command.

ARY0313E Database partition number must be numeric

Explanation: A specified database partition number is not valid.

Response: Specify the correct partition number and repeat the command.

ARY0315E Database partition number must be between 0 and 999

Explanation: A specified database partition number is not valid.

Response: Specify the correct partition number and repeat the command.

ARY0318E Password option requires username option

Explanation: Datastore user name is not specified.

Response: Specify the datastore user name and repeat the command.

ARY0321E Database alias must be specified

Explanation: Incorrect or empty database alias is specified.

Response: Specify valid database alias and repeat the command.

ARY0323E Datastore database alias must be specified.

Explanation: Datastore alias is not specified.

Response: Specify datastore database alias and repeat the command.

ARY0328E **Must set all of the values for: object type, schema and name, and point in time**

Explanation: Object name or schema is not specified.

Response: Provide all object names and schemas and repeat the command.

ARY0329E **Backup was not found. Specify backup timestamp and optional backup location for the desired backup image.**

Explanation: Specified backup time stamp is not found or is not valid.

Response: Specify the correct timestamp and/or optional backup location and repeat the command.

ARY0330E **Password option requires username option**

Explanation: User name for database alias must be specified.

Response: Provide database alias user name and repeat the command.

ARY0342E **Translation map or list of tables and table spaces must be specified**

Explanation: Empty translation map or list of tables and table spaces is detected.

Response: Provide list of tables and table spaces and repeat the command.

ARY0347E **Remote database connection requires specification of a log file directory**

Explanation: Specifying additional log path is required for remote database for Log Analysis.

Response: Specify additional log path and repeat the command.

ARY0348E {0}

Explanation: Generic DB2 error message.

Response: Observe the DB2 error message and follow DB2 response.

ARY0349E **Not connected**

Explanation: Connection to user database is not established.

Response: Contact IBM Software Support.

ARY0350E **Cannot change database partition name of existing database: {0}**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0352E **Unable to open backup image: {0}**

Explanation: Specify backup image cannot be open.

Response: Contact IBM Software Support.

ARY0355E **Error accessing file "{0}": {1}**

Explanation: Error message that is displayed when Recovery Expert failed to access specified file. I/O error message is displayed.

Response: Make sure that disk space on DB2 system is available and privileges are correct. Repeat the command.

ARY0356E **Error determining the current working directory: {1}**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0357E **Parse Error, string not well-formed, check value**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0359E **Unable to open converter using codeset "{0}", open status = {1}**

Explanation: Conversion table for specified code set is not available.

Response: Contact IBM Software Support.

ARY0360E **Data manager internal error {0}**

Explanation: Internal error message

Response: Contact IBM Software Support.

ARY0370E **Invalid SQLCA object.**

Explanation: The specified SQLCA object is invalid.

Response: Contact IBM Software Support.

ARY0371E **Character conversion error: {0}.**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0373E Object Translation Tool cannot connect to remote database

Explanation: Connection to the remote database is not established and not supported.

Response: Specify local backup image locations and repeat the command.

ARY0382E Cannot determine DB2 instance user name. DB2INSTANCE variable is not set.

Explanation: DB2 instance user name cannot be obtained.

Response: Set DB2INSTANCE variable accordingly and repeat the command.

ARY0387E Invalid DB2 instance user: cannot be root.

Explanation: Recovery Expert components are launched as root user.

Response: Relaunch the command as an instance user or any other non-root user.

ARY0394E Cannot find {0} in your PATH.

Explanation: Specified file is not found in the PATH environment variable.

Response: Modify your PATH environment variable and repeat the command.

ARY0403E You cannot use this program to run itself.

Explanation: The aryrun command is specified twice.

Response: Modify your command line and repeat the command.

ARY0415E Error: Unable to find program {0} using the current PATH definition.

Explanation: The product cannot find the specified application.

Response: Make sure that application is available and repeat the command.

ARY0416E Product library directory does not exist or is not accessible: {0}

Explanation: The installation directory is missing or the permissions are not set correctly.

Response: Correct the installation or reinstall the components to reset the directory contents. If the problem persists, contact IBM Software Support.

ARY0417E DB2 instance default library path does not exist or is not accessible: {0}.

Explanation: The directory is missing or the permissions are not set correctly.

Response: Contact IBM Software Support.

ARY0419E Product configuration file is not secure: {0}

Explanation: Installation directory permissions is not set correctly.

Response: Contact IBM Software Support.

ARY0426E Login user must be in the DB2 instance administration group: {0}.

Explanation: Installation directory permissions is not set correctly.

Response: Contact IBM Software Support.

ARY0427E Product installation path does not exist or is not accessible: {0}

Explanation: Installation directory does not exist or the permissions are not set correctly.

Response: Contact IBM Software Support.

ARY0435E Program is not accessible or secure: {0}

Explanation: Installation directory permissions is not set correctly.

Response: Contact IBM Software Support.

ARY0442E Error: missing option to set working directory path.

Explanation: Option that specifies working path is missing.

Response: Specify working directory and repeat the command.

ARY0443E Unable to open page from a cache

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0444E Error: Unknown transaction.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0446E **Error: Process ALTER TABLE record not supported yet.**

Explanation: Alter table DDL is found during the online backup processing and is not supported.

Response: No action is required.

ARY0447E **Error: failed to load compression library.**

Explanation: Compression library cannot be loaded.

Response: Contact IBM Software Support.

ARY0450E **Cannot obtain correct structure of table {0} from SLR database at LSN/LRI/LSO: {1}.**

Explanation: During recovery execution or data extraction from a backup image, the table structure definition was not found in the SLR for the specified table at the specified LSN/LRI/LSO. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Contact IBM Software Support.

ARY0453E **Unable to retrieve logfile '{0}'.**

Explanation: Specified log file cannot be retrieved.

Response: Make sure that specified log file is available and repeat the command.

ARY0454E **Error: Hole in archivlogs on file: {0}**

Explanation: Specified missing log is detected.

Response: Replace the missing log and repeat the command.

ARY0455E **Online backup is not supported for offline table space recovery.**

Explanation: Online backup is detected for TSR recovery that is not supported.

Response: No action is required.

ARY0456E **Nothing for dump - first active log number 0**

Explanation: No log activity is detected for specified database.

Response: No action is required.

ARY0457E **Table ({0},{1}) with XML columns is not supported.**

Explanation: XML data is detected and is not supported.

Response: No action is required.

ARY0459E **ESCAPE character is empty or too long in option '{0}'.**

Explanation: ESCAPE character is not specified or is too long in specified option.

Response: Modify the ESCAPE character and repeat the command.

ARY0461E **Unclosed quote in option '{0}': {1}**

Explanation: Missing quote is detected.

Response: Check all arguments for any unclosed quotes, modify them, and repeat the command.

ARY0466E **The database alias is too long.**

Explanation: Database alias provided is too long.

Response: Modify database alias and repeat the command.

ARY0469E **Archive log number is set but Schema Level Repository is not available.**

Explanation: The archive log number that was specified is either beyond the SLR range or the SLR is not created for the database where Log Analysis is launched.

Response: Ensure that the SLR is created and that it covers the archived logs that were specified for the Log Analysis task.

ARY0470E **RCC memory exception: {0}.**

Explanation: There is not enough memory available on the DB2 server where the DB2 Recovery Expert database components are working.

Response: Ensure that there is enough free memory for the DB2 Recovery Expert database server components and then repeat the command.

ARY0472E **Unexpected character in the position {0}.**

Explanation: Incorrect translation map at specified position is encountered.

Response: Modify translation map and repeat the command.

ARY0473E **Unable find template for input, position from {0}**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0474E **XML syntax error**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0475E **Invalid XML format of recovery plan.**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0476E **Recovery information not found for object '{0}' at specified PIT {1}.**

Explanation: Recovery is not possible for specified object as specified PIT.

Response: Modify PIT or change object that needs to be recovered and regenerate the Recovery Plan.

ARY0477E **Detailed report call for base session ID**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0478E **List of objects is incorrect: "{0}".**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0479E **Option {0} must be specified.**

Explanation: Specified option is missing.

Response: Provide missing option and repeat the command.

ARY0480E **Wrong value '{0}' for option {1}.**

Explanation: Incorrect specified value is observed for specified option.

Response: Correct value for specified option and repeat the command.

ARY0481E **Parameter for ARRAY INSERT too large**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0485E **No recovery plan available for given objects.**

Explanation: Recovery plan was not selected for the specified objects.

Response: Observe the PIT and selected objects, provide additional information and regenerate the Recovery Plan.

ARY0486E **Log Analysis begin setting is not within the defined SLR range**

Explanation: Begin time that is specified for Log Analysis is less than the SLR start time stamp.

Response: Modify begin time and repeat the command.

ARY0489E **Unknown object type '{0}'.**

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0490E **Internal parsing error.**

Explanation: Internal parsing error.

Response: Contact IBM Software Support.

ARY0491E **Comma in last position.**

Explanation: Comma was found in one of the parameter.

Response: Review past parameters, correct and repeat the command.

ARY0492E **Translation map or list of tables and table spaces must be specified.**

Explanation: Empty translation map or list of tables and table spaces is detected.

Response: Provide list of tables and table spaces and repeat the command.

ARY0493E **Specified start and/or end archive log file is not in SLR range: {0};{1}.**

Explanation: Specified logs for processing Log Analysis are not covered by SLR range.

Response: Modify the specified log range and repeat the command.

ARY0494E **Specified end archive log file has the start LSN/LRI/LSO: {0} greater than SLR end LSN/LRI/LSO: {1}.**

Explanation: Specified end archive log file LSN/LRI/LSO is greater than SLR end LSN/LRI/LSO. LSN/LRI/LSO depends on your version of DB2. For

ARY0495E • ARY0507E

DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Modify end LSN/LRI/LSO and repeat the command.

ARY0495E Cannot read info about end archive log file: {0} from DB2TSCHG.HIS file.

Explanation: Log information is not available in the DB2 table space change history file.

Response: No action is required.

ARY0496E Environment not attached to database

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0497E Unable to locate appropriate backup after REORG of table *table_name* .

Explanation: A REORG of the system catalog table was detected during SLR CREATE/UPDATE operation. To complete this operation, a backup image of SYSCATSPACE is required. The backup must be taken after this REORG operation, but before the next REORG operation of the same table.

Response: Take a new backup of SYSCATSPACE and repeat the SLR operation. Otherwise, limit the end point of the operation with the time of the REORG, or choose another starting backup.

ARY0498E Unable to start component, rc = {0} {1}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0499E Recovery task {0} finished with error.

Explanation: Recovery task has finished with specified error message.

Response: Observe the specified error message and follow its response.

ARY0500E Error obtained value for ARY_APP_DIR from table SYSTOOLS.ARY_PROPERTY, rc = {0} {1}

Explanation: Recovery Expert installation directory variable in the datastore database does not point to the installed components. Installation might be corrupted.

Response: Observe the specified installation location and reinstall database server components if required.

ARY0501E Datastore was not configured correctly because value for ARY_APP_DIR is not set, rc = {0} {1}

Explanation: This error is displayed when installation or the datastore creation is corrupted.

Response: Re-create datastore database first or reinstall database server components.

ARY0502E Aryrun could not start, rc = {0} {1}

Explanation: The Recovery Expert launcher component did not start.

Response: Contact IBM Software Support.

ARY0503E Launcher was started successfully but failed after, rc = {0} {1}

Explanation: The Recovery Expert launcher component was started, but then immediately stopped.

Response: Contact IBM Software Support.

ARY0504E Stored procedure was aborted, rc = {0} {1}

Explanation: The Recovery Expert stored procedure ended abnormally.

Response: Contact IBM Software Support.

ARY0505E Failed to launch Recovery Expert database server component on database partition {0}.

Explanation: A DB2 Recovery Expert database server component could not be launched during the "Timeout for launching server components" interval for database partition {0}. The cause of the problem might be that the operating system cannot start a new process, or that other running sessions do not allow concurrent changes in the datastore database.

Response: Check that the operating system can start new processes or check the list of the ARY sessions which are working at that time.

ARY0506E Aryrun was not found in ARY_APP_DIR, rc = {0} {1}

Explanation: The Recovery Expert launcher component was not found in the installation directory.

Response: Contact IBM Software Support.

| **ARY0507E** Abandoned {0} on host {1} because it did not start successfully within {2} seconds.

| **Explanation:** The database server component that is specified in the message did not start and report its

| status within the timeout value that is configured in the
| server, and was abandoned.

| **Response:** Check the host to see whether the database
| server component is still running or is stalled. It can be
| cancelled if it has come online, and this might indicate
| that the timeout value needs to be adjusted to
| accommodate for network latency. The timeout value
| can be adjusted by executing the following SQL
| statement against the datastore: UPDATE
| SYSTOOLS.ARY_PROPERTY SET VALUE =
| 'timeout_seconds' WHERE KEY =
| 'LAUNCHING_TIME'.

ARY0508E Component is not running on expected database partition. Current : {0}, expected: {1}.

Explanation: One of the database server components is not running on the expected database partition.

Response: Contact IBM Software Support.

ARY0511E Error in database '{1}': {0}.

Explanation: A DB2 database error has occurred.

Response: See the displayed DB2 error message and follow the suggested response.

ARY0512E Unable to get buffer size for template backup.

Explanation: The buffer size of the template backup that is required for TSR recovery cannot be obtained.

Response: Regenerate the recovery plan and then repeat the recovery. If the error persists, contact IBM Software Support.

ARY0514E An error has occurred while decompressing table row.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0515E Log file IDs region can not be set with negative value(s).

Explanation: Incorrect or negative values were specified for the log range for Log Analysis.

Response: Correct the log range values and complete the command.

ARY0516E Unable to determine column value for logged operation at *lsn* of transaction *transaction_id* committed at *timestamp* on table *table_name* due to a REORG.

Explanation: During an SLR operation, the operation identified by *lsn* against table *table_name* cannot be

reconstructed due to a REORG operation.

Response: Specify a backup that was taken after *timestamp* as the starting point of the SLR.

ARY0517E On LSN/LRI/LSO {0} table {1} have ONLINE REORG, process stopped.

Explanation: During SLR operations, an online REORG for the specified table and specified LSN/LRI/LSO was detected and is not supported. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: No action is required.

ARY0518E Data Manager {0} I/O error: {1}

Explanation: Internal I/O error message in the data manager.

Response: Contact IBM Software Support.

ARY0519E Incorrect end character in "relative" timestamp format

Explanation: During timestamp value parsing, the product detected an incorrect character.

Response: Correct the timestamp value and then repeat the command.

ARY0520E Invalid relative timestamp format does not contain any digits

Explanation: During timestamp value parsing, the product detected an incorrect character

Response: Correct the timestamp value and then repeat the command.

ARY0521E No data is found in the backup for specified table {0}

Explanation: The product found no data in the backup for the specified table.

Response: If this error message is displayed during object extraction command then chose a different backup that contains data for the specified table and repeat the command. Otherwise contact IBM Software Support.

ARY0522E Option format is mandatory

Explanation: No options were specified for the extraction command.

Response: Specify parameters for extraction and then repeat the command.

ARY0524E Table dictionary has not been gathered. Compressed rows cannot be restored.

Explanation: During Log Analysis processing, the product could not obtain the table dictionary to decompress the data.

Response:

ARY0525E Unable to perform SLR Update due to duplicate row on System Catalog table {0} following REORG.

Explanation: The data for System Catalog table *parameter_name* was reorganized and contains duplicate rows. SLR Update processing was not able to locate a row image to reconstruct an operation.

Response: The SLR will need to be rebuilt using a backup image taken after the table was reorganized.

ARY0526E Page dictionary has not been gathered. Compressed rows can not be restored.

Explanation: During Log Analysis processing, the product could not obtain the page dictionary to decompress the data.

Response:

ARY0600E Schema Level Repository (SLR) operation must be specified.

Explanation: The operation parameter is missing.

Response: Specify the operation parameter and then repeat the command.

ARY0601E Invalid Schema Level Repository (SLR) command: '{0}'.

Explanation: A specified operation for the SLR component is not valid.

Response: Correct the SLR operation parameter value and then repeat the command.

ARY0602E Schema Level Repository (SLR) exists, but '{0}' command was issued.

Explanation: For an existing SLR, the specified command was issued.

Response: Modify the operation parameter value and then repeat the command.

ARY0603E Schema Level Repository (SLR) does not exist, but '{0}' command was issued.

Explanation: The SLR has not been created.

Response: Create the SLR, and then repeat the command.

ARY0606E Schema Level Repository metadata incorrect.

Explanation: The datastore database is missing Recovery Expert tables or stored procedures.

Response: Re-create the datastore database and then repeat the command.

ARY0608E The end timestamp must be greater than the start timestamp.

Explanation: The end timestamp value is less than the start timestamp value.

Response: Modify the end timestamp value and then repeat the command.

ARY0609E The end log range value {0} must not be less than the start log range value {1}.

Explanation: Both the start and end log range values were specified with the -H option, and the end log range value is less than the start value.

Response: Take one of the following actions: Specify no log range values; modify the end log range value so that it is less or equal to the start log range value; specify only one log range value (either start or end).

ARY0610E Timestamp or LSN/LRI/LSO must be provided for pruning.

Explanation: During SLR operations, an incorrect or invalid timestamp value was provided for pruning of the SLR. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Specify the valid timestamp or LSN/LRI/LSO and repeat the command.

ARY0614E Schema Level Repository does not exist.

Explanation: The SLR has not been created.

Response: Create the SLR and then repeat the command.

ARY0616E Previous operations failed, you must rebuild Schema Level Repository

Explanation: SLR tables are missing.

Response: Rebuild the datastore database and then repeat the command.

ARY0617E Schema Level Repository table count is incorrect; you must rebuild Schema Level Repository

Explanation: SLR tables are missing.

Response: Rebuild the datastore database and repeat the command.

ARY0619E Schema Level Repository on catalog partition does not contain required transactions; SLR must be updated on catalog partition.

Explanation: During SLR operations, the SLR must be updated on the catalog partition first.

Response: Issue an SLR Update command.

ARY0620E Schema Level Repository is inconsistent, but '{0}' command was issued; SLR must be dropped and re-created.

Explanation: The SLR is inconsistent.

Response: Re-create the SLR and repeat the command if required.

ARY0621E Database id undefined.

Explanation: The database ID in the SLR table is undefined.

Response: Drop the SLR for this database and re-create it.

ARY0622E Schema Level Repository range error - end LSN/LRI/LSO {0} less than begin LSN/LRI/LSO {1}.

Explanation: During SLR operations, the specified End LSN/LRI/LSO is less than the Begin LSN. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Modify the End LSN/LRI/LSO and repeat the command.

ARY0623E Schema Level Repository range error - end lfid less than begin lfid

Explanation: The specified end log file is less than the start log file.

Response: Modify the end log file and then repeat the command.

ARY0624E Schema Level Repository range error - end timestamp less than begin timestamp

Explanation: The specified end timestamp is less than the begin timestamp.

Response: Modify the end timestamp value and then repeat the command.

ARY0626E Datastore version and application version not compatible

Explanation: The datastore version and DB2 database server components version do not match.

Response: Check whether the datastore database or database server components were reinstalled or modified recently and reinstall or modify the other component so that the version is the same.

ARY0627E Unable to retrieve datastore information

Explanation: SLR tables are missing.

Response: Re-create the datastore database and repeat the command.

ARY0628E Schema Level Repository has inconsistent state

Explanation: The SLR is in an inconsistent state.

Response: Rebuild the SLR and then repeat the Log Analysis command.

ARY0629E Backup for creating Schema Level Repository must contain database catalog

Explanation: The backup image that was provided for the SLR Create command does not contain the system catalog tables.

Response: Specify the backup image that contains the system catalog tables for the command that is used.

ARY0630E LSN/LRI/LSO values cannot be specified for Log Analysis begin or end arguments when running on multiple partitions.

Explanation: During the SLR operation, a database ID in the SLR table is undefined.

Response: Drop the SLR for this database and re-create it.

ARY0632E An invalid value '{0}' was specified for the parameter '{1}'.

Explanation: The parameter value that was used for the specified command is not valid.

Response: Correct the specified value and then repeat the command.

ARY0633E Prune command issued, but timestamp not specified

Explanation: A timestamp value was not specified for the prune operation.

Response: Specify a prune timestamp value and then repeat the command.

ARY0637E Only one database partition can be processed with this request

Explanation: Several database partitions were specified for Log Analysis processing, but only one partition is allowed with this command.

Response: Specify the desired database partition and repeat the command.

ARY0638E Log reading met reorg on rollback reading on LSN/LRI/LSO: {0}

Explanation: During SLR operations, a REORG operation was found during ROLLBACK to reconstruct system catalog tables at the specified LSN/LRI/LSO; the SLR cannot be updated further. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Use the appropriate backup to re-create the SLR from this point, or contact IBM Software Support.

ARY0639E Backup directory does not exist

Explanation: The specified backup directory does not exist.

Response: Modify the backup directory parameter and repeat the command.

ARY0640E Ranges backup does not exist

Explanation: The SLR backup does not exist.

Response: Contact IBM Software Support.

ARY0642E Unable to open table space info file: {0}.

Explanation: The product could not open the specified table space information file.

Response: Contact IBM Software Support.

ARY0643E Unable to read table space info for TID: {0}.

Explanation: Recovery Expert could not read the specified table space information file.

Response: Contact IBM Software Support.

ARY0644E The current log chain does not match initial log chain.

Explanation: The current log chain number does not match the initial log chain due to possible restore of the database.

Response: Re-create the SLR for the specified database using any backup that was created after the restore point.

ARY0701E Table name length is too short for extract suffix.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0702E Duplicate catalog partition.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0703E Schema Level Repository has not been created for the database.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0704E The database was re-created but Schema Level Repository is obsolete. Only drop/rebuild operations against Schema Level Repository are available.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0705E One of the SLR operation (create/drop/update/rebuild/prune) is being performed at the moment.

Explanation: Only one SLR operation (create/drop/update/rebuild/prune) can be performed simultaneously.

Response: Wait until SLR operation is finished and try again.

ARY0706E Invalid database ID.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0707E The SLR cannot be accessed while it is being updated by another process. Try the request again after the SLR update is complete.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0708E The SLR cannot be modified while it is being accessed by another process. Try the request again when the SLR is no longer being accessed by other processes.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0709E Invalid session ID.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0710E Session ID not available.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0711E General error occurred when closing session.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0712E Write session does not exist, unable to close write session.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0713E Read session does not exist, unable to close read session.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0714E Incorrect tables count in SLR database.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0715E Incorrect views count in SLR database.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0716E Incorrect procedures count in SLR database.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0717E Sessions limit exceeded.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0718E Key was not found in properties.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0719E Running sessions limit exceeded.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0720E Java procedures count is incorrect.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0721E Database or database partition was not found in SLR.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0722E Incorrect session time.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0723E For spectator session, session ID is undefined.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0724E • ARY0740E

ARY0724E Duplicate database id and key in ARY_PROPERTY table.

Explanation: An internal error occurred.

Response: Contact IBM Software Support.

ARY0725E Internal error - parent session is called, but session number is undefined.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0726E Parent session is undefined.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0727E Parent session is not open.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0728E Parent session was not found.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0729E Unable to close parent session, child sessions are still running.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0730E Parent session does not exist, unable to close parent session.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0731E Parent session does not exist, unable to create new session.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0732E Unable to drop session - session tasks active.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0733E DDL operations number is less than markers number.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0734E Attempt to obtain the unknown constant from ARY_CONSTANT table.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0735E Session wait is canceled.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0736E Unable to reuse session - this session is currently in use.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0737E Timestamp specified for prune is greater than the end timestamp of the SLR range.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0738E The timestamp value provided for the prune operation is outside of the SLR range.

Explanation: The product cannot perform the prune operation with the timestamp value provided; however, SLR status is not affected.

Response: Specify a timestamp value that is between the SLR start and end timestamps.

ARY0739E Unable attach to session with type that is not 'U'.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0740E Specified table is not in the compile list.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0741E Unable to start detailed report for empty session.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0742E Only one type of task is allowed in a session.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0743E Previous summary or full report is not generated, unable to do detailed report.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0744E Session type incorrect.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0745E Unable to start recovery plan execution - session is empty.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0746E Recovery plan execution is failed.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0747E Timestamp specified for prune is less than the start timestamp of the SLR range.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0748E {0} specified for prune is less than the start {1} of the SLR range.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0749E SLR processing error, internal code: 1.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0750E SLR processing error, internal code: 2.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0751E SLR processing error, internal code: 2.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0752E SLR processing error, internal code: 3.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0753E SLR processing error, internal code: 4.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0754E SLR processing error, internal code: 5.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0755E Unable to set borders on prune - not in transaction range

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0756E Tablespace count in SLR database incorrect

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0757E Unable to create GET_SQL view - empty components

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0758E Invalid scenario id: {0}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0759E Backup records do not exist in SLR

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0760E SLR in use

Explanation: The SLR is in use, which prohibits the requested operation.

Response: After all clients have stopped using the SLR, repeat the command.

ARY0761E Processing stopped due to error in prior task. Check messages for details.

Explanation: Tasks on one or more database partitions completed with errors.

Response: For more information, see previous messages.

ARY0801E Invalid UTIL request.

Explanation: A utility was not valid.

Response: Contact IBM Software Support.

ARY0802E Database alias must be set for the requested operation.

Explanation: This error message is displayed when the database alias is missing for the utility component.

Response: Contact IBM Software Support.

ARY0803E Utility operation type must be specified.

Explanation: This error message is displayed when operation is not specified for the utility component.

Response: Contact IBM Software Support.

ARY0805E Datastore username and/or password is specified, but datastore option is missing.

Explanation: This error message is displayed when the datastore database is not specified.

Response: Specify the datastore database and repeat the command.

ARY0806E Invalid transaction state: {0}.

Explanation: An invalid transaction parameter value was specified for Log Analysis.

Response: Modify the transaction parameter value and repeat the command.

ARY0807E Database partition number, username, and password must be set to get remote database configuration.

Explanation: The database partition number, user name, or password is missing. The information is required for the remote database operation.

Response: Specify the missing parameters and repeat the command.

ARY0808E Options '{0}' and '{1}' cannot be specified together.

Explanation: This error message is displayed when two specified options cannot be used together for Log Analysis.

Response: Modify your command line accordingly and repeat the command.

ARY0809E Starting or ending log numbers cannot be specified when Minimum Recovery Time (MRT) processing mode is selected.

Explanation: This error message is displayed when starting or ending log numbers are specified as a filter for Log Analysis to be run in the minimum recovery timestamp processing mode.

Response: Remove log number filters and repeat the command.

ARY0811E SQL file name or source SQL session ID was not provided .

Explanation: This error message is displayed when the SQL file name or the session ID is not provided for SQL generation.

Response: Specify the missing parameter and repeat the command.

ARY0812E The session {0} contains no right SQL: {1}.

Explanation: This error message is displayed when there is no SQL to generate for the specified session.

Response: No action is required.

ARY0813E Error when applying SQL: {0}{1}{2}{3}.

Explanation: This error message is displayed when an error occurred during applying SQL statement. Partition, transaction ID, operation ID, and count are displayed.

Response: No action is required.

ARY0814E Unsupported type of a host variable

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0815E Unexpected end of sql-statement.

Explanation: An unexpected end of SQL statement was observed.

Response: Contact IBM Software Support.

ARY0816E Specified start log range value {0} for processing is less than MRT minimum log file {1}.

Explanation: The start log range value that was specified with the -H option is less than the minimum recovery timestamp (MRT) minimum log file. The Log Analysis feature cannot process log files that are earlier than the MRT.

Response: Take one of the following actions: Specify no log range values; modify the start log range value so that it is greater than or equal to the MRT minimum log file, and less than or equal to the end log range value (if you specified one).

ARY0817E Specified end log range value {0} for processing is less than minimum recovery timestamp (MRT) minimum log file {1}.

Explanation: The end log range value that was specified with the -H option is less than MRT minimum log file. The Log Analysis feature cannot process log files that are earlier than the MRT.

Response: Take one of the following actions: Specify no log range values; modify the end log range value so that it is greater than or equal to both the MRT minimum log file and the start log range value (if you specified one).

ARY0818E Report starting point {0} is before specified log range {1}.

Explanation: The start log range value was specified with the -H option, the report starting point was specified with the -b option, and the specified report starting point is earlier than the start log range value. The Log Analysis feature cannot process log files that are outside of the log range.

Response: Take one of the following actions: Specify neither the -H nor the -b option; specify only an end log range value with the -H option that is greater than or equal to the report starting point; Specify either the -H option or the -b option, but not both; modify either the report starting point(-b), the start log range value(-H), or both so that the first value is greater than or equal to the second value.

ARY0819E Report starting point {0} is after specified log range {1}.

Explanation: The end log range value was specified with -H option, the report starting point was specified with the -b option, and the specified report starting point is later than the end log range value. The Log Analysis feature cannot process log files that are outside of the log range.

Response: Take one of the following actions: Specify neither the -H nor the -b option; specify only a start log range value with the -H option that is less than or equal to the starting report point; specify either the -H option or the -b option, but not both; modify either the report starting point (-b), the end log range value (-H), or both so that the first value is less than or equal to the second value.

ARY0830E Value for option '{0}' is out of the SLR range.

Explanation: This error message is displayed when the PIT that is specified for DDL generation component is outside of the SLR range.

Response: Modify the PIT for the DDL generation and repeat the command.

ARY0831E Object specified by options '{0}' and '{1}' does not exist at specified PIT.

Explanation: This error message is displayed when specified object for DDL generation does not exist at specified PIT.

Response: Modify object or the PIT value and repeat the command.

ARY0832E Base directory path name must be specified for 'browse' option.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0833E List of database partitions must be specified for retrieving backup information from the Recovery History File (RHF).

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0834E Pending state for DB {0} on database partition: {1}

Explanation: A user database was detected in backup-, restore-, or rollforward-pending state.

Response: Bring the database to the active state and repeat the command.

ARY0835E Incorrect file db2nodes.cfg format

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0836E No LFH file was found.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0838E No GLFH file was found.

Explanation: The global LFH file was not found in the user database.

Response: Verify that the database is not physically damaged and that LFH files are available, then rerun the command. If the error persists, contact IBM Software Support.

ARY0900E 3DES error: encrypt/decrypt error, code: {0}.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0901E 3DES error: there is not enough room to completely encrypt/decrypt, source/destination length: {0}/{1}.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0902E 3DES error: Invalid 3DES-encrypted data was encountered.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0903E 3DES error: Invalid 3DES-encrypted data was encountered (pad = {0}).

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0904E Unexpected character in the position {0} in log paths

Explanation: This error message is displayed when an unexpected character in specified position in the user-provided log path value.

Response: Verify and correct log path value parameter and repeat the command.

ARY0905E Unexpected character in the position {0} in backup paths

Explanation: This error message is displayed when an unexpected character in the specified position was found in the provided user backup path locations.

Response: Verify and correct user backup path locations and repeat the command.

ARY0906E Dependencies not found in the SLR.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY0910E Incompatible remote and local database version. Remote instance: DB2 {0}, local instance: DB2 {1}.

Explanation: During Log Analysis or SLR operations against a remote user database, it was discovered that the remote and local DB2 versions on which the product database server components are installed do not match.

Response: Catalog the remote database on a system with matching versions of DB2 where remote processing needs to be performed, and copy DB2 assets (logs and backups) that are required for processing. Then install the database server components on the system and repeat the required processing.

ARY0911E Incompatible remote and local instance environment. Remote instance: OS {0}, Platform: {1}, local instance: OS {2}, Platform: {3}.

Explanation: During Log Analysis or SLR operations against a remote user database, it was discovered that the remote OS platform and local OS platform on which the product database server components are installed are not compatible.

Response: Catalog the remote database on a system that is compatible with the remote database, and copy DB2 assets (logs and backups) that are required for processing. Then install the database server components on the system and repeat the required processing.

ARY0912E Partitioned remote database is not supported.

Explanation: During Log Analysis or SLR operations against a remote user database, it was discovered that the remote database is a partitioned database (DPF), and is not supported.

Response: No action is required.

ARY4000W Property "log.path" in aryclp.properties points on inaccessible file.

Explanation: This warning message is displayed when the specified log path directory in the properties file contains inaccessible or nonexistent directory.

Response: Modify log path in the properties file and repeat the command.

ARY4000E During creating datastore session SQL error occurred: {0}

Explanation: This error signifies that specified DB2 SQL error has occurred.

Response: See the displayed DB2 error message and follow the suggested response.

ARY4001W Property "log.level" in aryclp.properties is not valid.

Explanation: This warning message is displayed when the specified log level property in the properties file is not valid.

Response: Modify the log level value to a desired level in the properties file and repeat the command.

ARY4001E During closing inactive sessions for database SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error is observed while closing inactive session.

Response: Observe the displayed DB2 SQL error message and follow its response.

ARY4002E While obtaining status information about datastore session, SQL error occurred: {0}

Explanation: While the product was obtaining information about the datastore session, a DB2 SQL error was observed.

Response: Observe the displayed DB2 error message and follow its response.

ARY4003W Property "time.representation" in aryclp.properties has not been set, defaulting to local time.

Explanation: This warning message signifies that time representation parameter in the properties file has not been set and the default time representation will be set to local time.

Response: If the default time representation needs to be changed then modify time representation parameter in the properties file and repeat the command.

ARY4003E The product version {0} is not compatible with the datastore repository version {1}.

Explanation: The displayed version of the installed product is not compatible with the displayed version of the datastore repository.

Response: If the datastore repository version is older, use the product Admin perspective to upgrade it. Otherwise, upgrade the product installation.

ARY4004W Property "time.representation" value of {0} in aryclp.properties does not match a valid setting (utc or server or local), defaulting to local time.

Explanation: This warning message is displayed when time representation value in the properties file is not valid and defaults to the local time.

Response: Modify time representation parameter in the properties file and repeat the command if required.

ARY4004E During registering database {0} in Datastore SQL error occurred: {1}

Explanation: This error message signifies that DB2 SQL error message occurred during registering of specified user database.

Response: Observe the displayed DB2 error message and follow its response.

ARY4005E XML file {0} was not found.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4006E During parsing XML file {0} error occurred: {1}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4007E Datastore with name "{0}" already exists.

Explanation: This error message signifies that existing datastore database name was found in the repository database.

Response: Remove existing datastore database item from the repository database and repeat the command.

ARY4008E Format of ARY data XML file "{0}" is incorrect.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4009E During transformation XML file "{0}" string error occurred: {1}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4010E While checking the existence of session {0} SQL error occurred: {1}.

Explanation: While the product was checking the existence of the session, a DB2 SQL error occurred.

Response: See the displayed DB2 error message and follow the suggested response.

ARY4011E While obtaining information about the state of session {0}, an SQL error occurred: {1}.

Explanation: While the product was obtaining information about the state of the specified session, a DB2 SQL error occurred.

Response: See the displayed DB2 error message and follow the suggested response.

ARY4012E Directory "{0}" either does not exist or is not permitted.

Explanation: The specified directory does not exist or does not have proper privileges for reading.

Response: Modify privileges for the specified directory or specify another directory, and then repeat the command.

ARY4013E Required data about installed product was not found.

Explanation: This error message is displayed when the information about installed components is not found.

Response: Contact IBM Software Support.

ARY4014E During obtaining operation details SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error is observed during obtaining of Log Analysis details.

Response: Observe the displayed DB2 error message and follow its response. If error is persistent, contact IBM Software Support.

ARY4015E Required database partition number {0} does not exist.

Explanation: The database partition number that was specified for browsing is not valid.

Response: Contact IBM Software Support.

ARY4016E During working with "aryutil" proxy table SQL error occurred: {0}

Explanation: This error message signifies DB2 error message is observed while working with the utility component.

Response: See the displayed DB2 error message and follow the suggested response. If the error persists, then contact IBM Software Support.

ARY4017E "{0}" does not specify path to backup image.

Explanation: This error message signifies that specified user path is not valid backup image path name.

Response: Contact IBM Software Support.

ARY4018W Property "ssh.timeout" in aryclp.properties is not valid.

Explanation: The specified SSH timeout property in the properties file is not valid.

Response: Modify the SSH timeout value in the properties file to a valid time (in seconds), and then repeat the command.

ARY4019E During starts process via Java SP SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error message is observed during process launching.

Response: See the displayed DB2 error message and follow the suggested response. If the error persists, contact IBM Software Support.

ARY4020E Database partition number must be integer between 0 and 999.

Explanation: The specified database partition number is not valid.

Response: Correct the database partition number value and repeat the command.

ARY4021E During reading datastore database property table SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error is observed while reading datastore database property table.

Response: See the displayed DB2 error message and follow its suggested response. If the error persists, contact IBM Software Support.

Note: It is necessary to install DB2 Enterprise Server Edition 9.7 Fix Pack 3a and higher or APAR IC67948, if you receive error message ARY4021E or ARY4121E.

ARY4022E Internal error. Class "{0}" is used incorrectly.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4023E Target database birthday does not match. You can change Log Analysis mode to MRT or rebuild the SLR to process earlier logs.

Explanation: This error message signifies that during Log Analysis it was detected that the specified user database was re-created but SLR information was built for previous version of the user database.

Response: Rebuild SLR and repeat the command or modify Log Analysis mode to launch Log Analysis in MRT mode.

ARY4024E Internal error. Required class-executor was not found by alias "{0}".

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4025E Error occurred while decoding base64.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4026E Error occurred while decoding data.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4028E One scenario of the plan has already been applied. The plan must be regenerated before it is run again.

Explanation: One scenario of the plan has already been applied.

Response: The plan must be regenerated before being run again because the database geometry could have changed.

ARY4029E While obtaining information about plan sessions, an SQL error occurred: {0}

Explanation: This error message indicates that DB2 error message is observed while obtaining information about the recovery plan.

Response: Observe the displayed DB2 error message and follow its response. If the error message persists,

contact IBM Software Support.

ARY4030E Internal error. XML request is not valid.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4031E Operation does not have details because it is terminal operation for delete of MDC table.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4032E SQL error occurred while obtaining datastore database authorities: {0}

Explanation: This error message signifies that DB2 error message is observed during obtaining authorization information about the datastore database.

Response: See the displayed DB2 SQL error message and follow the suggested response. If the error persists, contact IBM Software Support.

ARY4033E Datastore database is corrupted or version of server/client part of DB2 Recovery Expert is not compatible with version of the DB2 Recovery Expert Datastore.

Explanation: This error message signifies that there is version mismatch between datastore database and DB2 Recovery Expert server components or datastore database is corrupted.

Response: Re-create the datastore database and repeat the command.

ARY4034E Database "{0}" is in backup pending state.

Explanation: This error message signifies that specified database is in backup pending state.

Response: Bring specified database to available state and repeat the command.

ARY4035E Database "{0}" is in rollforward pending state.

Explanation: This error message signifies that specified database is in rollforward pending state.

Response: Bring specified database to available state and repeat the command.

ARY4036E Database "{0}" is in restore pending state.

Explanation: This error message signifies that specified database is in restore pending state.

Response: Bring specified database to available state and repeat the command.

ARY4037E Unable to create a database connection, an SQL error occurred: {0}

Explanation: A SQL error occurred when trying to connect to the selected database. This might have occurred because the authorization credentials entered were incorrect or not valid.

Response: Review the SQL error and correct the problem based on the error message. If the SQL error indicates a character conversion problem, then ensure that the authorization credentials follow the rules required by the operating system of the database server. For example, on Linux and UNIX platforms, characters for user names must be lowercase a through z, 0 through 9, and _ (underscore) for names not starting with 0 through 9.

ARY4038E Blank password is not allowed.

Explanation: This error message signifies that no password was provided for DB2 database.

Response: Specify password and repeat the command.

ARY4039E Invalid userID or password.

Explanation: This error message signifies that invalid username or password was provided for connection to DB2 database.

Response: Modify username and/or password and repeat the command.

ARY4040E During opening database connection unexpected error occurred: {0}

Explanation: This error message signifies that unexpected error message is observed and is displayed.

Response: Contact IBM Software Support.

ARY4041E Internal error. Input parameter "time point" is incorrect.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4042E JDBC connection cannot be established because connection settings do not allow it. Try to change connection settings and repeat task on new tab.

Explanation: This error message signifies that JDBC connection cannot be established.

Response: Try to change connection settings and repeat task on new tab.

ARY4043E During obtaining information about list of sessions SQL error occurred: {0}

Explanation: This error message signifies that DB2 error message is observed during obtaining list of sessions.

Response: Observe the displayed DB2 error message and follow its response. If error persists contact IBM Software Support.

ARY4044E Internal error. Type of SLR process is incorrect.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4045E During obtaining list of databases registered in datastore SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error message is observed while obtaining list of registered databases.

Response: Observe the displayed SQL error message and follow its response. If error messages persists contact IBM Software Support.

ARY4046E Unable to retrieve report details for transaction ID {0} LSN/LRI/LSO {1} and database partition number {2}.

Explanation: Log Analysis report details are not recovered and not available for the specified transaction and LSN/LRI/LSO. A blocking activity (such as a LOAD or REORG operation) prevented the Log Analysis tool from accessing a backup image or current database to reconstruct the details of an operation. It is also possible that no backup image is available between two REORGs or LOADs. In this case, Log Analysis cannot recover some details for the time frame between these two blocking activities. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: If a backup image is available for the specified database partition for the time frame of the

transaction, then specify the location for this backup. If accessing backups or reading database pages is prohibited, then allow the Log Analysis tool to use backups and access the current database, and repeat the command.

ARY4047E Unable to retrieve generated SQL for transaction ID {0} LSN/LRI/LSO {1} and database partition number {2}.

Explanation: During Log Analysis, SQL was not generated. LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Contact IBM Software Support.

ARY4048E Operation does not have SQL because it is terminal operation for delete of MDC table.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4049E During obtaining string SQL statement SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error message is observed during obtaining generated SQL statement.

Response: Observe the displayed DB2 error message. If error persists, contact IBM Software Support.

ARY4050E IRepositoryContributor with ID {0} was not found.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4051E Error occurred while opening connection to ARY repository. Error message is: {0}.

Explanation: Internal error message.

Response: Contact IBM Software Support.

| **ARY4052E** An SQL error occurred while working
| with the product repository database:
| {0}.

| **Explanation:** The specified SQL error occurred while
| attempting to access the configuration repository. This
| might occur if an incorrect database was specified for
| the configuration repository.

| **Response:** Verify that the database options that are
| specified are for the configuration repository database.

ARY4053E IRepositoryMgrService was not found.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4054E During parsing, an XML string error occurred: {0}.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4055E During validation, an XML string error occurred: {0}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4056E During processing remote command error occurred: {0}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4057E User "{0}" does not have authority to execute datastore database procedures.

Explanation: This error message signifies that specified user does not have privileges to execute datastore database stored procedures.

Response: If specified user needs to have permissions then grant permissions for specified user and repeat the command.

ARY4058E During obtaining a last "Run SQL" session SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error message is observed during obtaining RUN SQL information.

Response: Observe the displayed DB2 error message and follow its response. If error persists contact IBM Software Support.

| **ARY4059E** While opening SSH connection to host
| "{0}" error occurred: {1}.

| **Explanation:** An error occurred while attempting to
| establish SSH connections.

| **Response:** If it is timeout problem then go to
| datastore settings and change SSH timeout parameter
| to the greater one.

ARY4060E Cannot create file "{0}" for datastore definitions.

Explanation: This error message signifies that specified file cannot be created to store datastore connection definitions.

Response: Make sure that there is enough free disk space and privileges are correct. Repeat the command.

ARY4061E File "{0}" for datastore definitions must be readable and writable.

Explanation: This error message signifies that specified file for datastore definitions is not accessible.

Response: Modify privileges for specified file and make sure that the specified file exists and repeat the command.

ARY4062E Unable to find required property "{0}" in file aryclp.properties.

Explanation: This error message signifies that displayed property is not found in the properties file.

Response: Contact IBM Software Support.

ARY4063E Command Line Processor configuration file "{0}" cannot be found.

Explanation: This error message signifies that properties file is not found.

Response: Replace missing properties file and repeat the command.

ARY4064E I/O error occurred while reading aryclp.properties file. Error message is: "{0}"

Explanation: This error message signifies that I/O error occurred while reading the properties file. The observed error message is displayed.

Response: Follow the displayed error message response and repeat the command.

ARY4065E A connection to target host "{0}" cannot be created.

Explanation: Connection attempt to the target machine was denied. This has occurred because an incorrect hostname or incorrect authorization credentials were entered.

Response: It is necessary to verify that the hostname and credentials were entered correctly and it is also necessary to verify that this computer allows remote access.

ARY4066E Datastore "{0}" name was not found.

Explanation: This error message signifies that specified datastore name is not found in the repository.

Response: Make sure that the datastore database entry is valid and exists and repeat the command

ARY4067E During creating Log Analysis report file in CSV format SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error is observed while creating Log Analysis CSV report file.

Response: Observe the displayed DB2 error message and follow its response.

ARY4068E During obtaining general information of Log Analysis report SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error is observed while obtaining information about Log Analysis SQL report.

Response: Observe the displayed DB2 error message and follow its response.

ARY4069E During obtaining part of Log Analysis summary report SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error is observed while obtaining Log Analysis summary report.

Response: Observe the displayed DB2 error message and follow its response.

ARY4070E During generating Log Analysis formatted text report SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error is observed while creating Log Analysis text report.

Response: Observe the displayed DB2 error message and follow its response.

ARY4071E During generating Log Analysis XML SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error is observed while creating Log Analysis XML report file.

Response: Observe the displayed DB2 error message and follow its response.

ARY4072E During applying recovery scenario parameters SQL error occurred: {0}

Explanation: This error message signifies that DB2 SQL error is observed while applying recovery scenario parameters.

Response: Observe the displayed DB2 error message and follow its response.

ARY4073E During obtaining recovery scenario information SQL error occurred: {0}

Explanation: This error message is displayed during obtaining recovery scenario information and signifies that the displayed DB2 SQL error message is observed.

Response: Follow that DB2 error message response.

ARY4074E Specified DB2 Recovery Expert session does not exist, because it was re-created by another user.

Explanation: This error message is displayed when the DB2 Recovery Expert session was reused by another user.

Response: No action is required.

ARY4075E During updating DB2 Recovery Expert session SQL error occurred: {0}

Explanation: This error message is displayed during updating DB2 Recovery Expert session and signifies that DB2 SQL error is observed.

Response: Follow the displayed DB2 error message response and repeat the command.

ARY4076E Internal error. Incorrect type of generated log analysis report.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4077E During obtaining session options SQL error occurred: {0}

Explanation: This error message is displayed during obtaining session information and signifies that the displayed DB2 SQL error message is observed.

Response: Follow the displayed DB2 error message response and repeat the command.

ARY4078E During creating trace files SQL error occurred: {0}

Explanation: This error message is displayed while creating trace files and signifies that the displayed DB2 error message is observed.

Response: Follow the displayed DB2 SQL error

message response and repeat the command.

ARY4079E Incorrect type for grouping report objects.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4080E During obtaining report "{0}" SQL error occurred: {1}

Explanation: This error message is displayed during obtaining specified report objects and signifies that the displayed DB2 SQL error message has been encountered.

Response: Follow the displayed DB2 error message response and repeat the command.

ARY4081E SQL error occurred while setting datastore database authorities: {0}

Explanation: This error message is displayed while setting datastore database privileges and signifies that the displayed SQL error has been observed.

Response: Follow the displayed error message response and repeat the command.

ARY4082E Error occurred while encoding data.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4083E Unhandled error occurred while redirecting to /arympservice.

Explanation: Internal error message.

Response: This internal error message could mean that one or several server components failed to run. Try restarting Recovery Expert server and if the error persists then contact IBM Software Support.

ARY4085E SQL error occurred while obtaining list of specifications: {0}

Explanation: This error message is displayed during obtaining list of specifications and signifies that the displayed DB2 SQL error message is observed.

Response: Follow the displayed DB2 error message response and repeat the command.

ARY4086E SQL error occurred while obtaining information about JDBC connection: {0}

Explanation: While obtaining information about the JDBC connection, the product observed the displayed DB2 SQL error message.

Response: For information about the SQL error, see

ARY4087E • ARY4101E

the IBM documentation for error codes issued by the IBM Data Server Driver for JDBC and SQLJ. Follow the IBM recommended response, and then repeat the command.

ARY4087E DB2 version is not supported.

Explanation: DB2 version of the user database is not supported.

Response: No action is required.

ARY4088E Connection to datastore database must be created for catalog partition.
Reconnect to hostname "{0}".

Explanation: The connection to the datastore database that was specified on the Admin perspective is using non-catalog partition hostname.

Response: Modify current datastore connection and change the hostname to the displayed hostname, and then repeat the command.

ARY4089E While obtaining current timestamp via JDBC connection, SQL error occurred: {0}

Explanation: While obtaining the current timestamp value, a DB2 SQL error message was observed.

Response: Follow the displayed DB2 SQL error message response and then repeat the command.

ARY4090E Internal unexpected error. Internal thread was interrupted.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4092E Path to DB2 Recovery Expert components is not configured.

Explanation: DB2 Recovery Expert components were not installed properly or the datastore database is corrupted.

Response: Verify that the DB2 Recovery Expert components are installed properly and verify that the datastore database exists and is not corrupted.

ARY4093E DB2 Recovery Expert component "aryrun" does not exist or is not available in the directory pointed to by ARY_APP_DIR in table SYSTOOLS.ARY_PROPERTY in the datastore database.

Explanation: DB2 Recovery Expert components were not installed properly or the datastore database is corrupted.

Response: Verify that the DB2 Recovery Expert components are installed properly and verify that the datastore database exists and is not corrupted.

ARY4094E Internal error. Incorrect filter object type.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4095E Error occurred while starting DB2 Recovery Expert component "aryutil".
Error message is: {0}.

Explanation: Internal error message

Response: Contact IBM Software Support.

ARY4096E DB2 Recovery Expert component "aryrun" finished with error: {0}.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4097E Datastore database is corrupted or version of server/client part of DB2 Recovery Expert is not compatible with version of datastore of DB2 Recovery Expert.

Explanation: The datastore database is corrupted or the version of the DB2 Recovery Expert components and the datastore database do not match.

Response: Verify that the installed components are installed properly and the datastore database is not corrupted.

ARY4099E Internal server error. Process that runs DB2 Recovery Expert components used incorrectly.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4100E Specified DB2 Recovery Expert session does not exist.

Explanation: The chosen DB2 Recovery Expert session could not be opened because it no longer exists.

Response: No action is required.

ARY4101E Error occurred while parsing string timestamp.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4102E Internal error. LA report filter type is incorrect.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4103E During obtaining LA statistics for a object SQL error occurred: {0}

Explanation: While obtaining information about Log Analysis statistics, the displayed SQL error message was observed.

Response: Follow the displayed DB2 error message response and repeat the command.

ARY4104E Internal error. Incorrect input parameters for sending ZIP file.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4105E File for loading does not exist.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4107E Specification with input name "{0}" already exist.

Explanation: The specification name already exists.

Response: Modify the specification name and repeat the command.

ARY4108E While registering specification in a session SQL error occurred: {0}

Explanation: While saving a session, a DB2 SQL error was observed.

Response: Follow the displayed DB2 error message response and repeat the command.

ARY4109E During reading session task name SQL error occurred: {0}

Explanation: While reading the session information, the displayed SQL error message was observed.

Response: Follow the displayed DB2 error message response and repeat the command.

ARY4110E Network error occurred.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4111E Unknown hostname "{0}".

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4112E Hostname is not reachable within timeout {0}.

Explanation: This error message is displayed when the hostname is not available for specified timeout period.

Response: Verify that your systems are available in the network and repeat the command.

ARY4113E During exporting LOB from detail report SQL error occurred: {0}

Explanation: This error message is displayed while exporting LOB data from the Log Analysis report and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4114E During exporting LOB from detail report I/O error occurred: {0}

Explanation: This error message is displayed while exporting LOB data from the Log Analysis report and signifies that I/O error is observed.

Response: Contact IBM Software Support.

ARY4115E During unsaving session SQL error occurred: {0}

Explanation: This error message is displayed while marking session as unsaved and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4116E RXA connection error occurred: {0}

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4117E Platform "{0}" is not supported.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4118E Internal error. DB2 Recovery Expert Web Service input ID is incorrect.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4119E Internal error. Class-executor "{0}" cannot be loaded.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4120E Internal error. Class-executor does not inherit required class.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4121E During obtaining information from session log SQL error occurred: {0}

Explanation: This error message is displayed during obtained information from the session log and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

Note: It is necessary to install DB2 Enterprise Server Edition 9.7 Fix Pack 3a and higher or APAR IC67948, if you receive error message ARY4021E or ARY4121E.

ARY4122E During deleting specification SQL error occurred: {0}

Explanation: This error message is displayed while deleting specification and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4123E During obtaining timezone of datastore database SQL error occurred: {0}

Explanation: This error message is displayed during obtaining information about the datastore database and signifies that the SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4124E Internal error. Incorrect input language parameter.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4125E During applying datastore properties SQL error occurred: {0}

Explanation: This error message is displayed while applying new datastore properties and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4126E During obtaining datastore properties SQL error occurred: {0}

Explanation: This error message is displayed while retrieving datastore properties and signifies that the DB2 SQL error message is observed.

Response: Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4127E Cannot create directory into work path.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4128E During clearing previous session data SQL error occurred: {0}

Explanation: This error message is displayed while clearing previous session data and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4129E Datastore XML file "{0}" is corrupted.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4130E I/O error occurred while compressing contents of a directory.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4131E SQL error occurred while obtaining SLR range: {0}

Explanation: This error message is displayed while obtaining the SLR range and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4132E SQL error occurred while obtaining database storage type from SLR: {0}

Explanation: This error message is displayed while obtaining database storage type and signifies that the DB2 SQL error message is observed.

Response: Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4134E During obtaining SLR state for registered in datastore target database SQL error occurred: {0}

Explanation: This error message is displayed while obtaining SLR state and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4135E During obtaining SLR status for the database from datastore SQL error occurred: {0}

Explanation: This error message is displayed while obtaining SLR status and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4136E During canceling agent session SQL error occurred: {0}

Explanation: This error message is displayed while canceling agent session and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4137E The session's time expired.

Explanation: This error message is displayed when the open session has expired.

Response: No action is required.

ARY4138E Directory "{0}" can not be found on target machine.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4139E During obtaining list of transaction SQL error occurred: {0}

Explanation: This error message is displayed while obtaining list of transactions and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4140E During obtaining list of initial recovery objects SQL error occurred: {0}

Explanation: This error message is displayed while obtaining list of recovery objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error

message response and repeat the command.

ARY4141E During obtaining list of depend recovery objects SQL error occurred: {0}

Explanation: This error message is displayed while obtaining list of dependent recovery objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4142E During saving new data to recex.properties I/O error occurred: {0}

Explanation: This error message is displayed when the I/O error is observed while saving the properties file.

Response: Observe the I/O error and contact IBM Software Support.

ARY4143E During reading recex.properties file I/O error occurred: {0}

Explanation: This error message is displayed while reading the properties file and signifies that the I/O error is observed during reading the properties file.

Response: Observe the I/O error message and contact IBM Software Support.

ARY4146E Unable to launch the DB2 Recovery Expert component "aryrun" due to version incompatibility. Required version is {0} or later.

Explanation: The specified database server component could not be launched because the installed version is not compatible with the application server.

Response: Upgrade the database server components to the specified version or a later version. The installed version of database server components can be found on the Deployment Status perspective.

ARY4147E During saving specification SQL error occurred: {0}

Explanation: This error message is displayed while saving specification and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4148E During saving session SQL error occurred: {0}

Explanation: This error message is displayed while saving the session and signifies that the DB2 SQL error message is observed.

ARY4149E • ARY4163E

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4149E Datastore database is corrupted. SLR table space count is incorrect.

Explanation: This error message is displayed when the datastore database is corrupted.

Response: Re-create the datastore database.

ARY4150E Datastore database is corrupted. Java procedure count is incorrect.

Explanation: This error message is displayed when the I/O error is observed while saving the properties file.

Response: Observe the I/O error and contact IBM Software Support.

ARY4151E Datastore database is corrupted. SLR table count is incorrect.

Explanation: This error message is displayed when the datastore database is corrupted.

Response: Re-create the datastore database.

ARY4152E Datastore database is corrupted. SLR views count is incorrect.

Explanation: This error message is displayed when the datastore database is corrupted.

Response: Re-create the datastore database.

ARY4153E Datastore database is corrupted. SLR SQL procedure count incorrect.

Explanation: This error message is displayed when the datastore database is corrupted.

Response: Re-create the datastore database.

ARY4154E Datastore database is corrupted or chosen database is not datastore.

Explanation: This error message is displayed when the datastore database is corrupted.

Response: Re-create the datastore database.

ARY4155E SQL error occurred while checking datastore database: {0}

Explanation: This error message is displayed while checking the datastore database and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4156E Connection "{0}" is active. Reset connection and try again.

Explanation: This error message is displayed when the specified connection is in use.

Response: Repeat your command after the specified connection is not in use.

ARY4157E Configuration error. Product location environment is not set.

Explanation: This error message is displayed when the product installation is corrupted or incomplete.

Response: Reinstall DB2 Recovery Expert server component.

ARY4158E Configuration error. Product location directory "{0}" does not exist or not accessible.

Explanation: This error message is displayed when the product installation is corrupted or incomplete.

Response: Reinstall DB2 Recovery Expert server component.

ARY4159E System error. Error occurred during working a system process.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4160E Error occurred during executing a system process.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4161E Internal error. Output format from "arylum" is not valid.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4162E Internal error occurred while applying license.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4163E Internal error occurred while validating license.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4164E No license for this product is available.

Explanation: The license for the product has expired or is missing.

Response: Add the license for the product, or contact IBM Software Support.

ARY4165E No license for this product is available: enrollment certificate is missing.

Explanation: The License for the product has expired or is missing.

Response: Add the license for the product, or contact IBM Software Support.

ARY4166E Trial license has expired.

Explanation: License for the DB2 Recovery Expert has expired or missing.

Response: Add the license for the product, or contact IBM Software Support.

ARY4167E License problem: {0}.

Explanation: Internal license problem.

Response: Contact IBM Software Support.

ARY4168W Application ID {2} terminated unexpectedly and disconnected from session {0} for database {1} .

Explanation: One of the tasks identified by application-id for the session identified by session-id has disconnected unexpectedly from the session.

Response: Examine the session messages to identify what process may have terminated. Verify that there is sufficient disk space and that any process limits for memory and open files are set to recommended values. If any corrective action does not solve the problem, contact IBM software support.

ARY4169E An SQL error occurred while obtaining list of sequences from target database: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4170E An SQL error occurred while obtaining list of sequences from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the

DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4171E An SQL error occurred while obtaining list of variables from target database: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4172E An SQL error occurred while obtaining list of variables from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4173E An SQL error occurred while obtaining list of modules from target database: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4174E An SQL error occurred while obtaining list of modules from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4175E An SQL error occurred while obtaining list of UDTs from target database: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4176E An SQL error occurred while obtaining list of UDTs from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the

ARY4177E • ARY4189E

DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4177E An SQL error occurred while obtaining list of SQL functions from target database: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4178E An SQL error occurred while obtaining list of SQL functions from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4179E An SQL error occurred while obtaining list of SQL procedures from target database: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4180E An SQL error occurred while obtaining list of SQL procedures from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining specified list of objects and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4181E A Schema Level Repository error occurred while canceling related objects: {0}

Explanation: This error message is displayed during the cancel operation and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed error message response and repeat the command.

ARY4182E The specified PIT, {0}, occurs before the SLR history coverage starting point, {1}.

Explanation: The specified PIT is not covered by the SLR beginning time.

Response: If the PIT that was provided is correct then the SLR Rebuild operation is required to cover the PIT for the required recovery.

ARY4183E Internal error. Incorrect servlet input parameters.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4184E An SQL error occurred while obtaining operations list for transactions: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4185E Unable to start a new session, because the maximum of {0} have been started.

Explanation: Unable to start a new session, because the maximum of {0} have been started.

Response: No action is required.

ARY4186E Unable to start a new session, because the maximum of {0} are running.

Explanation: Unable to start a new session, because the maximum of {0} are running.

Response: No action is required.

ARY4187E DB2 JDBC driver is not loaded.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4188E Server location "{0}" does not exist, is invalid or is not accessible.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4189E An SQL error occurred while obtaining list of buffer pools from target database: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4190E An SQL error occurred while obtaining list of buffer pools from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4191E An SQL error occurred while obtaining list of event monitors from target database: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4192E An SQL error occurred while obtaining list of event monitors from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4193E An SQL error occurred while obtaining list of database partition groups from target database: {0}

Explanation: While obtaining the list of specified objects, a DB2 SQL error message was observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4194E An SQL error occurred while obtaining list of database partition groups from Schema Level Repository: {0}

Explanation: While obtaining the list of specified objects, a DB2 SQL error message was observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4195E An SQL error occurred while obtaining list of schemas from target database: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4196E An SQL error occurred while obtaining list of schemas from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4197E An SQL error occurred while obtaining list of tables from target database: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4198E An SQL error occurred while obtaining list of tables from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4199E An SQL error occurred while obtaining list of table spaces from target database: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4200E An SQL error occurred while obtaining list of table spaces from Schema Level Repository: {0}

Explanation: This error message is displayed during obtaining the list of specified objects and signifies that DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4201E An SQL error occurred while obtaining database ID from SLR by name: {0}

Explanation: This error message is displayed when obtaining database ID from SLR and signifies that the DB2 error message is observed.

ARY4202E • ARY4216E

Response: Follow the displayed DB2 error message response and repeat the command.

ARY4202E JDBC URL "{0}" is incorrect.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4203E Directory "{0}" has no DB2 server components for {1}.

Explanation: The specified directory does not contain DB2 server components for the required OS.

Response: Verify that the specified directory contains DB2 server component for the required OS.

ARY4204E I/O error opening file for printing.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4205E Option "{0}" does not exist.

Explanation: Internal error message

Response: Contact IBM Software Support.

ARY4206E Option "{0}" must have value.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4207E Internal error occurred while parsing option "{0}".

Explanation: Internal error message

Response: Contact IBM Software Support.

ARY4208E Session ID that is specified does not point to Log Analysis session.

Explanation: This error message is displayed when the session ID that was specified point to a session that does not have the requested information.

Response: Contact IBM Software Support.

ARY4209E Session ID specifies session that is not completed.

Explanation: This error message is displayed when the session ID that was specified point to a session that does not have the requested information.

Response: Contact IBM Software Support.

ARY4210E Session ID that is specified does not point to Log Analysis session with SQL.

Explanation: This error message is displayed when the session ID that was specified point to a session that does not have the requested information.

Response: Contact IBM Software Support.

ARY4211E Session ID that is specified does not point to Log Analysis session with required report type.

Explanation: This error message is displayed when the session ID that was specified point to a session that does not have the requested information.

Response: Contact IBM Software Support.

ARY4212E Connection to datastore database is corrupted due to error message: {0}.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4213E For executing this command you need set current connection to datastore.

Explanation: This error message is displayed when the datastore connection is not established.

Response: Establish datastore connection first and then repeat the command.

ARY4214E SQL error occurred while closing connection. Error message: {0}.

Explanation: This error message is displayed while closing the connection and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 error message response and repeat the command.

ARY4215E I/O error occurred while reading file "{0}".

Explanation: This error message is displayed and signifies that I/O error occurred while reading the specified file.

Response: Contact IBM Software Support.

ARY4216E File "{0}" was not found.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4217E I/O error occurred: {0}

Explanation: This error message is displayed and signifies that I/O error is observed.

Response: Contact IBM Software Support.

ARY4218E Internal error. Executor class cannot be created.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4219E User name for login to datastore is incorrect.

Explanation: This error message is displayed when the user name that was provided for the datastore connection is incorrect.

Response: Correct the username and repeat the command.

ARY4220E Client server session is expired.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4222E No web session data.

Explanation: No web session data.

Response: No action is required.

ARY4223E User name for target database is incorrect.

Explanation: This error message is displayed when the user name for the target database is not provided or empty.

Response: Provide user name and repeat the command.

ARY4225E WSID is incorrect. Update web session is required.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4226E An SQL error occurred while obtaining information about target database partitions: {0}

Explanation: While obtaining information about the specified database partitions, a DB2 SQL error message was observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4228E Unsupported symbol with code "{0}" has been found at position "{1}".

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4229E The following filename format is not supported: "{0}".

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4230E An SQL error occurred while updating the existing specification: {0}

Explanation: This error message is displayed during updating the specification information and signifies that the DB2 SQL error message is observed.

Response: Follow the displayed DB2 SQL error message response and repeat the command.

ARY4231E Internal unexpected error occurred: {0}

Explanation: Internal unexpected error occurred on DB2 Recovery Expert server side.

Response: Collect all available logs and sent them to developers.

ARY4232E Remote path {0} does not exist on the target system "{1}".

Explanation: This error message is displayed when the specified remote pass is not available on the target system.

Response: Make sure that the path is correct and repeat the command.

ARY4233E Remote path {0} cannot be accessed on the target system "{1}" by the user.

Explanation: This error message is displayed when the specified remote path cannot be accessed on the target system.

Response: Modify the path or its privileges and repeat the command.

ARY4234E For Database Partitioning Facility (DPF) instances an option "partition" is mandatory for "run OX" command.

Explanation: This error message is displayed for OX command and signifies that the partition must be specified.

Response: Specify partition option and repeat the command.

ARY4235E • ARY4258E

ARY4235E The database partition "{0}" does not exist.

Explanation: This error message is displayed for OX command and signifies that the specified database partition is not valid.

Response: Correct the database partition number and repeat the command.

ARY4236E The Schema Level Repository must exist in order to use the "table list" option.

Explanation: This error message is displayed for OX command and signifies that table list option is available only when SLR is built.

Response: Correct your command line accordingly and repeat the command.

ARY4237E Internal error. Incorrect format of XML output of "aryutil".

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4238E Format of JDBC options is incorrect. Format must be "key1=value1;key2=value2;...".

Explanation: Specified JDBC options for datastore connection is incorrect.

Response: Correct the JDBC options for the current datastore connection and repeat the command.

ARY4241E Unsupported codepage "{0}".

Explanation: This error message signifies that an unsupported character code page was set.

Response: Modify the character code page and repeat the command.

ARY4243E Internal error occurred while LSO/PSO conversion.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4244E Internal error. Incorrect LRI format.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4245E Internal error. Error occurred while retrieving version information from manifest.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4248E Version of the product datastore database is unsupported for upgrade.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4249E SQL error {0} occurred while updating datastore database.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4250E Internal error. Error {0} occurred while opening SQL batch file.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4251E Internal error. Wrong SQL batch file name: {0}.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4252E Internal error. Error {0} occurred while obtaining list of SQL batch files.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4253E Internal error. SQL batch files for datastore upgrade were not found.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4254E Internal error. Error {0} occurred while updating datastore database.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4258E An SQL error occurred while the product was obtaining the list of storage groups from target database: {0}.

Explanation: This error message indicates that a DB2 SQL error was observed while the product was reading information about storage groups from the user database.

Response: See the displayed DB2 error message and follow the suggested response. If the error persists, contact IBM Software Support.

ARY4259E An SQL error occurred while the product was obtaining the list of storage groups from Schema Level Repository: {0}.

Explanation: A DB2 SQL error was observed while the product was reading information about storage groups from the datastore database.

Response: See the displayed DB2 error message and follow the suggested response. If the error persists, contact IBM Software Support.

ARY4260E An SQL error occurred while processing the report transactions.

Explanation: The reported transaction information is corrupted.

Response: Check the ARY-MP log file in the Logs perspective for error details. Check that the datastore database is available, and then rerun the report to collect the transaction information.

ARY4265E Directory "{0}" does not exist.

Explanation: This error message indicates that the specified directory is not accessible on the server system.

Response: Try the command again. If the error persists, contact IBM Software Support.

ARY4267E Processing is disabled for database {0} because database is not recoverable.

Explanation: The database specified is not recoverable and therefore cannot be processed.

Response: Make sure that the database selected is a recoverable database.

ARY4270E User name or password for remote database is incorrect.

Explanation: This error message signifies that invalid username or password was provided for connection to database.

Response: Modify username and/or password and repeat the command

ARY4271E During reading BL trace file {0} I/O error occurred: {1}

Explanation: This error message is displayed and signifies that I/O error occurred while reading the trace file.

Response: Observe the I/O error message and contact IBM Software Support.

ARY4272E Internal error occurred while obtaining information about shared installations from product repository database.

Explanation: Internal error message.

Response: Contact IBM Software Support.

ARY4273E Invalid option specified, threshold is only supported for Quiet Time reports.

Explanation: The report type is not Quiet Time report and a quiet time threshold was specified.

Response: Correct the report type or remove the quiet time threshold parameter.

ARY4274E Invalid quiet time threshold "{0}".

Explanation: The value provided for the quiet time threshold is not valid. The minimum allowed threshold value is 1 second.

Response: Provide a threshold value in the format of: hh.mm.ss where hh is hours, mm is minutes, and ss is seconds or n<slm|h> where n is a positive integer value.

ARY4275E While obtaining the list of schemas from the schema level repository, an SQL error occurred.

Explanation: While the product was obtaining the list of schemas for the specified objects from the schema level repository, an SQL error occurred.

Response: Contact IBM Software Support.

ARY4276E An SQL error occurred while obtaining of the list of schemas from the target database.

Explanation: While the product was obtaining the list of schemas for the specified objects from the target database, an SQL error occurred.

Response: Contact IBM Software Support.

ARY4277E Cannot establish the ssh connection to "{0}".

Explanation: The error occurred while attempting to connect a socket to a remote address.

Response: Contact IBM Software Support.

ARY4278E Command "{0}" is no longer supported. Use the "{1}" command instead.

Explanation: The specified command is not supported for the current release.

Response: Use the suggested command.

| **ARY4279E** **Cannot get credentials from repository file.**

| **Explanation:** The product cannot get the required credentials from the XML repository file.

| **Response:** Use the sync command to synchronize managed database information from the application server repository. If the problem persists, verify that you specified the correct configuration repository.

| **ARY4280E** **Cannot get datastore from repository file.**

| **Explanation:** The product cannot get the required datastore from the XML repository file.

| **Response:** Use the sync command to synchronize managed database information from the application server repository. If the problem persists, verify that you specified the correct configuration repository.

| **ARY4281E** **Database name must be valid JDBC URL.**

| **Explanation:** The database name must be a valid JDBC URL, as shown in the following example: "jdbc:db2://hostname:port/DBNAME"

| **Response:** Specify a valid JDBC URL for the database name.

ARY5000W **You should allow local storage to work correctly.**

Explanation: The amount of cache allotted in the Flash Player is insufficient to store the necessary Flex objects.

Response: Use the Flash Player Settings Manager to allocate at least 10 KB of cache so that the necessary objects can be stored.

| **ARY5000E** **Error occurred while installing the product.**

| **Explanation:** An error occurred during installation of the Recovery Expert database server components.

| **Response:** Check the status window for the source of the error, correct the error, and then try again.

ARY5001E **Error occurred while configuring datastore database.**

Explanation: An error occurred during configuration of the datastore database.

Response: Check the status screen for the source of the error, correct the error, and then try again.

ARY5002W **You have selected a point in time that is in the future. Recovery could not be done for this time.**

Explanation: The point in time (PIT) value that was selected for the recovery is in future and is invalid.

Response: Modify recovery PIT to the point in time where object(s) existed in the past.

ARY5002E **Error while writing local value.**

Explanation: This error can occur if the local information storage for objects of this domain are permanently disallowed.

Response: Ensure that local information storage is allowed for objects of this domain.

ARY5003E **Processing is disabled for database {0} because database is in backup pending state.**

Explanation: The database specified is in a backup pending state and therefore cannot be processed.

Response: After the database backup is complete, try the operation again.

ARY5004E **Processing is disabled for database {0} because database is in ROLLFORWARD pending state.**

Explanation: A database in ROLLFORWARD pending state cannot be processed.

Response: After the database ROLLFORWARD is complete, try the operation again.

ARY5005E **Processing is disabled for database {0} because database is in restore pending state.**

Explanation: The database specified is in a restore pending state and therefore cannot be processed.

Response: After the database restore is complete, try the operation again.

ARY5006E **Processing is disabled for database {0} because database is a remote database.**

Explanation: The database specified is a remote database and therefore cannot be processed.

Response: Ensure that the database selected is a local database.

ARY5007E Processing is disabled for database {0} because database is not recoverable.

Explanation: The database specified is not recoverable and therefore cannot be processed.

Response: Ensure that the database selected is a recoverable database.

ARY5008E Processing is not supported for the DB2 Recovery Expert datastore database.

Explanation: The database specified is a DB2 Recovery Expert datastore database and therefore cannot be processed.

Response: Ensure that the database selected is not a DB2 Recovery Expert datastore database.

ARY5009E You must fill report window.

Explanation: The Report step in the Log Analysis perspective has not been filled.

Response: Complete the necessary fields within the Report step, and then try the operation again.

ARY5010E The selected filter contains an empty or incorrect value.

Explanation: The value of the selected filter in the Filters step of the Log Analysis perspective is incorrect.

Response: Correct the filter value for the selected filter and then try the operation again.

ARY5011E At least one item must be selected for Operations.

Explanation: At least one type of database operation must be selected in the Filters step of the Log Analysis perspective.

Response: Select one or more database operation types and then try the operation again.

ARY5012E Network communication error.

Explanation: A possible network communication error occurred.

Response: Make sure that the network communication between the browser client and Recovery Expert Server is restored, and then try the operation again.

ARY5013E Required parameter "{0}" is not filled properly on scenario step {1}.

Explanation: The Recovery plan task parameter displayed in the message was not properly set up in the selected scenario.

Response: Ensure that the Recovery plan task

parameter is properly set within the selected scenario, then repeat this operation.

ARY5014E Invalid XML response.

Explanation: The XML response from the Recovery Expert Server is invalid.

Response: Contact IBM Software Support.

ARY5015E Start point should be less or equal to end point and greater or equal to SLR start point.

Explanation: In the Report step of the Log Analysis perspective, the chosen start time was either greater than the end time or less than the SLR start time.

Response: Ensure that the chosen start time is less than or equal to the end time and that it is greater than or equal to the SLR start time.

ARY5016E Start log number should be less or equal to end log number.

Explanation: The Start log number is greater than the End log number.

Response: Ensure that the Start log number is less than or equal to End log number.

ARY5018E The specified LSN/LRI/LSO is not a valid 16-character hexadecimal value.

Explanation: LSN/LRI/LSO depends on your version of DB2. For DB2 9.7 and earlier, the value is a log sequence number; for DB2 10.1 and later, it is a log record identifier (if length is 34 hexadecimal characters) or a log stream offset (if length is 16 hexadecimal characters).

Response: Enter a correct LSN/LRI/LSO value.

ARY5019E Invalid file path input.

Explanation: The current file input path is not valid.

Response: Enter a correct file input path.

ARY5020E The session has been expired due to long idle. The tab will be closed.

Explanation: The current tab will close due to expiration of the current session.

Response: Try to login again and reopen the tab.

ARY5021E An error occurred while uploading the file: {0}.

Explanation: The specified file could not be uploaded to the Recovery Expert server.

Response: Make sure that the Recovery Expert server

ARY5022E • ARY5031E

is running and that there are no network communication issues between the browser client and server, and then try again.

ARY5022E Unexpected server response: {0}

Explanation: Internal product error.

Response: Repeat the last operation. If the error persists, contact IBM Software Support.

ARY5023E The session could not be opened because it has been expired.

Explanation: This error messages signifies that the session that was selected to be open has expired and no longer exist.

Response: No action is required.

ARY5024E Start point should be greater or equal to SLR start point.

Explanation: This error message signifies that the start point timestamp value, that is specified for Log Analysis, is less than the SLR starting point.

Response: Modify the Log Analysis starting point and repeat the command.

ARY5025E Start log number should be greater or equal to SLR start log number.

Explanation: This error message signifies that the start log number value, that is specified for Log Analysis, is less than the SLR start log number.

Response: Modify the Log Analysis start log number and repeat the command.

ARY5026E Select at least one database partition.

Explanation: The database partitions where Log Analysis needs to be launched are not specified.

Response: Select the database partitions where the Log Analysis task must be executed, and repeat the command.

ARY5027W There are currently too many DB2 objects to display. Only the first {0} objects of the selected type are displayed. You can further refine your search by entering more restrictive values in the DB2 object display filter.

Explanation: This warning message is displayed when the total number of objects to be displayed exceeds the maximum number of objects that is specified.

Response: If object(s) that you must work with are not displayed, specify a different filter to cover your objects and/or increase the maximum number of objects to be

displayed and repeat the command.

ARY5027E Required parameter "{0}" is not filled properly for table space "{1}" on scenario step {2}.

Explanation: This error message signifies that the specified parameter is not set for the specified table space for the specified step.

Response: Provide the required information for the specified step and repeat the command.

ARY5028E The script has been executed for too long because of too many objects. Set object filters or reduce the number of objects to display.

Explanation: This error message is displayed when the Recovery Expert client takes more than 15 seconds to perform the required command.

Response: This could be a normal situation. Observe whether the Recovery Expert client works correctly and if there are problems and then specify additional filters to filter the number of objects.

ARY5029E A connection to the target host "{0}" could not be established. Would you like to try testing the connections to all specified hosts again?

Explanation: Whether to perform the connectivity check on all hosts is determined by the **Test connections to all specified hosts** check box on the **Credentials** step. DB2 Recovery Expert attempted to perform RXA connections to all hosts, and at least one connection failed.

Response: Click **Yes** to attempt connection testing again.

ARY5030E The selected installation did not pass the verification.

Explanation: During the **Verify install** step, DB2 Recovery Expert determines whether required components exist, and whether any components have been modified (for example, applying a fix) since the initial installation.

Response: Review the status messages to determine the required action. Correct any problems, and if necessary, reinstall the components.

ARY5031E A connection to the target host {0} could not be established due to invalid user credentials.

Explanation: DB2 Recovery Expert performed a connections test before uninstalling the database server components. The connection failed because the user

credentials that were entered are not valid.

Response: Using authorized credentials, try the uninstallation again.

ARY5032E The upgrade for the specified database server components is not viable because the selected version is older than the currently installed version.

Explanation: DB2 Recovery Expert does not allow the installation of database server components that are older than the currently installed version.

Response: To continue with the upgrade, select a newer version of the database server components.

ARY5033E The specified LSO is not a valid 16-character hexadecimal value.

Explanation: An invalid log stream offset (LSO) has been entered.

Response: Enter a correct log stream offset.

ARY5034E The specified LRI is not a valid 34-character hexadecimal value.

Explanation: An invalid log record identifier (LRI) has been entered.

Response: Enter a correct log record identifier.

ARY5035E The start and end point types are not consistent with each other.

Explanation: The specified start and end points are of different types. Log Analysis processing requires that they be of the same type.

Response: Modify the start and end points so that both fields are specified using a timestamp, both fields are specified by an LSN, LSO or LRI, or at least one field is specified using the "Earliest available"/"Latest available" selection.

ARY5036E The minimum number of pages specified for a container within a DMS table space must be 1 or greater.

Explanation: The minimum size of a DMS table space is five extents. Attempting to create a table space smaller than five extents will result in an error (SQL1422N). Three extents in the table space are reserved for overhead. At least two extents are required to store any user table data. (These extents are required for the regular data for one table, and not for any index, long field or large object data, which require their own extents.)

Response: Modify "Number of pages" column cell so that this value is greater than 0.

ARY5037E The Schema Level Repository (SLR) needs to be up-to-date in order to obtain the object dependencies and proceed with the recovery. Would you like to update the SLR?

Explanation: The Schema Level Repository (SLR) needs to be up-to-date in order to obtain the object dependencies and proceed with the recovery.

Response: Click **Yes** to update the SLR now, **No** to ignore this message.

ARY5038E A corresponding right parenthesis should follow the left one.

Explanation: The closing parenthesis is missing. Both opening and closing parentheses are required.

Response: Provided both opening and closing parentheses, and then retry the action.

ARY5039E The left and right parentheses count should be equal.

Explanation: The number of opening parentheses does not match the number of closing parentheses.

Response: Provide an equal number of opening and closing parentheses, and then retry the action.

ARY5040E A column name should be selected for row {0}.

Explanation: The required column name field selection has not been made at the specified row.

Response: Select a valid column name.

ARY5041E An operator should be selected for row {0}.

Explanation: The required operator field selection has not been made at the specified row.

Response: Select a valid operator.

ARY5042E A column value must be provided for row {0}.

Explanation: The required column value field is missing at the specified row.

Response: Specify a column value.

ARY5043E The identity of the current user is identical to the specified user.

Explanation: The current user and the user identity to which you requested to change are the same.

Response: Use `su` or `sudo` to switch to a different user identity than the current user.

ARY5044E When creating a datastore, it is not valid to switch to the "root" user.

Explanation: A non-root user identity that has access to the target instance/database is required.

Response: Use su or sudo to switch to a non-root user identity that has access to the target instance/database.

ARY5045E Database {0} cannot be selected because the database property 'National language string units' must be set to 'SYSTEM' instead of '{1}'.

Explanation: The current setting of the database property 'National language string units' makes it unusable with the product.

Response: Set 'National language string units' to the value SYSTEM.

ARY5046E Database {0} cannot be selected because the database property 'String units' must be set to SYSTEM instead of '{1}'.

Explanation: The current setting of the database property 'String units' makes it unusable with the product.

Response: Set 'String units' to the value SYSTEM.

ARY5047E Database {0} cannot be selected because the database property 'Default table organization' must be set to ROW instead of '{1}'.

Explanation: The current setting of the database property 'Default table organization' makes it unusable with the product.

Response: Set 'Default table organization' to the value ROW.

ARY5048E Datastore cannot be updated as it has been updated by another client

Explanation: The attempt to update the datastore information failed as it has been updated by another user

Response: Refresh and retry the update

ARY5049E Unable to find datastore in the Configuration Repository

Explanation: An attempt to find the datastore in the configuration repository failed.

Response: Reload the list of datastores and attempt the operation again

ARY5050E Unable to add datastore to the configuration repository as it already exists.

Explanation: An attempt to define a new datastore failed because a datastore with the same properties already exists

Response: Reload the list of datastores to load the datastore that already exists

ARY5051E Managed database {0} (host: {1}, port: {2}) already exists in the Configuration Repository

Explanation: Unable to add the managed database to the configuration repository as it already exists

Response: Reload list of managed databases to retrieve the latest list of managed databases that are already configured.

ARY5052E Invalid user ID or password.

Explanation: An invalid user name or password was provided for connection to the hostname.

Response: Specify a valid user name and password and repeat the command.

ARY5053E DB2 Recovery Expert encountered the following error: {0}

Explanation: Failure to properly add the {0} datastore.

Response: Remove this incorrectly added datastore, then try creating this datastore again by referring to the User's Guide for the correct settings. Alternatively, you can click [here](#) to make any corrections to your previous datastore profile.

ARY5054E Datastore database with specified properties does not exist

Explanation: Update of datastore properties cannot be proceeded because connection to the database with specified parameters has been failed

Response: Check datastore properties for validity and update the appropriate properties

ARY5055E Datastore with database name {0} already exists in configuration repository.

Explanation: An attempt was made to add a new datastore with the specified name failed because a datastore with that name already exists.

Response: Reload the list of datastores and add the new datastore with a new name.

ARY5056E SQL error occurred while populating datastore: {0}

Explanation: A SQL error occurred during the population of the datastore and it could not be populated properly

Response: Identify the SQL error and attempt to add the datastore again

ARY5057E Cannot remove datastore {0} because one or more managed databases refer to it.

Explanation: An attempt to remove a datastore failed because a managed database exists in the configuration repository and it points to the datastore being removed

Response: Remove managed databases connected to datastore and retry

ARY5058E Database cannot be updated as it has been updated by another client.

Explanation: Database has been updated by another client.

Response: Refresh and retry an update

ARY5059E DB2 Recovery Expert encountered an error; managed database not added.

Explanation: Failure to add the {0} managed database. For more information, see deployment details on the Deployments perspective.

Response: Remove the incorrectly added database, then try adding the managed database again. For the correct settings, see the product documentation.

ARY5060E Invalid alternate user ID or password.

Explanation: An invalid alternate user name or password was provided for execution.

Response: Specify a valid user name and password and repeat the command.

ARY5061E Credentials cannot be updated as it has been updated by another client.

Explanation: Credentials have been updated by another client.

Response: Refresh and retry an update

ARY5062E Alternate credentials cannot be updated as it has been updated by another client.

Explanation: Alternate credentials have been updated by another client.

Response: Refresh and retry an update

ARY5063W Unable to remove database from datastore.

Explanation: The product cannot connect to the datastore database to remove the database. This error might occur if the database has been dropped, or is in a pending or inconsistent state. Datastore content will remain unchanged.

Response: No action is required.

ARY5064E Deployment cannot be upgraded as it has been updated by another client.

Explanation: Deployment has been updated by another client.

Response: Refresh and retry an upgrade

ARY5065E Managed database {0} is not usable with Recovery Expert because {1} is set to {2}.

Explanation: Recovery Expert cannot use the managed database because the database is in pending state or its configuration prevents log reading.

Response: Correct the specified database configuration parameter.

ARY5066E Invalid user ID or password.

Explanation: An invalid user name or password was provided for connection to the instance.

Response: Specify a valid user name and password, and then repeat the command.

ARY5067E Instance has a failed deployment.

Explanation: The instance on which the new managed database is being added has a failed deployment.

Response: Reinstall the failed deployment by removing all existing managed databases in the instance.

ARY5068E Database configuration parameter {0} is set to {1}. The value must be equal to or greater than {2}.

Explanation: The specified parameter value is invalid.

Response: Specify a valid value for the configuration parameter.

ARY5069E Label of current datastore is '{0}'. Set current datastore with label '{1}' to use specified managed database.

Explanation: The specified managed database belongs to a different datastore than the current one.

Response: Set the datastore for the specified managed database as the current datastore.

ARY6018W Closing the current dialog will not result in termination of the process that is currently running in the background. By closing the dialog, any informational or warning messages resulting from the background process will not be displayed. Are you sure you want to close the dialog?

Explanation: Closing the wizard during the installation, uninstallation, upgrade, or installation verification process will not end the actual process. The process will continue to run in the background.

Response: Click **Yes** to close the wizard, or **No** to leave it open.

ARY6019W The installation you have selected already exists on the specified instance. Are you sure you want to continue and overwrite the existing installation with this new one?

Explanation: DB2 Recovery Expert does not allow two identical installations on the same instance.

Response: To overwrite the existing installation, click **Yes**. To close the dialog without overwriting the existing installation, click **No**.

ARY6021W There are one or more unsaved configurations. Would you like to save them?

Explanation: There are one or more unsaved configurations. Would you like to save them?

Response: Click **Yes** to save all changes made to any unsaved configuration(s) **Yes**. Click **No** to discard all changes made to any unsaved configuration(s) **No**. Click **Cancel** to close this warning dialog **Cancel**.

ARY8001W Closing any existing connections to the database.

Explanation: The user requested to drop the target database, but some connections to the database still exist. These connections will be closed.

Response: No action required.

ARY8001E Required input parameter database name is not set.

Explanation: Required input parameter database name is not set.

Response: Add the required input parameter database name.

ARY8002I Setting current database partition to "{0}".

Explanation: Displayed during the datastore create operation to identify the current database partition.

Response: No action is required.

ARY8002E No users or groups were specified.

Explanation: At least one user or group name must be specified when granting or revoking authorities.

Response: Try the request again and specify at least one user or group name.

ARY8003W Directory does not exist; unable to delete "{0}".

Explanation: The specified directory does not exist; it cannot be deleted.

Response: No action required.

ARY8003E \$HOME/sqllib/db2profile does not exist.

Explanation: \$HOME/sqllib/db2profile does not exist.

Response: No action is required.

ARY8004W Warning: "{0}" argument ignored.

Explanation: The specified command argument was ignored.

Response: Verify that the request was run with the correct arguments. If necessary, reissue the request with the correct arguments.

ARY8005W Messages suppressed for: {0}

Explanation: The message is used when running in verbose mode to show a list of message tags for which messages would normally be hidden from the user.

Response: No action required.

ARY8005E Unable to set client partition parameters.

Explanation: The script was unable to set the specified DB2 client connect and attach parameters.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8006W The configured active log size of "{0}" 4KB units is less than the recommended default setting of "{1}" 4KB units.

Explanation: The log size may not be sufficient to perform some operations on the datastore database.

Response: You can adjust the log configuration parameters LOGPRIMARY, LOGSECOND, and LOGFILSIZ to increase the size of the active log.

ARY8006E Unable to connect to database "{0}".

Explanation: The script was unable to connect to the specified database.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8007W The database configuration parameter "{0}" is set to "{1}", which is less than the recommended value of "{2}".

Explanation: The specified database configuration parameter value is less than the indicated recommended value. This may affect product operation.

Response: Increase the configuration parameter to a value equal to or greater than the recommended value.

ARY8007E Unable to connect to database "{0}" to determine partition list.

Explanation: The script was unable to connect to the specified database.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8008W Installation location is empty or corrupted. Check it manually.

Explanation: The database server files cannot be uninstalled using the standard uninstallation procedure because the installation location is empty or missing expected files.

Response: Delete the product files manually on any host computers where they were installed.

ARY8008E Unable to determine list of database partitions.

Explanation: The script was unable to determine the list of partition numbers for the database.

Response: Contact IBM Software Support.

ARY8009I Deleting datastore definitions from existing database "{0}".

Explanation: The existing datastore objects are being dropped.

Response: No action is required.

ARY8009W Datastore database "{0}" is defined using codeset "{1}". DB2 Recovery Expert can only process user databases that match this codeset definition.

Explanation: The specified datastore database is defined using the indicated codeset. This means that DB2 Recovery Expert can only process user databases that match the specified codeset. Users cannot connect to other databases.

Response: To work with databases that are defined with a different codeset, create the datastore in a new UTF-8 database or in an existing database that is defined to use the UTF-8 codeset.

ARY8009E A problem was detected while validating the contents of the datastore database "{0}". It is necessary to re-create the datastore. If the problem persists, contact IBM Software Support.

Explanation: The datastore database contents could not be validated.

Response: Re-create the datastore database. If the problem persists, contact IBM Software Support.

ARY8010W You have requested to drop the existing database *parameter_name*. All of the data in the database will be lost. Type "1" to continue, or "0" to stop.

Explanation: Dropping and then re-creating the DB2 Recovery Expert datastore database might result in unwanted loss of data. If the datastore was created in an existing user database, and you drop the database, you will lose all of the user data. However, if the datastore was created in its own database, dropping that database will only affect DB2 Recovery Expert.

Response: Verify that you really intend to drop the containing database. If you chose to create the DB2 Recovery Expert datastore in an existing database, all of the datastore definitions will be automatically dropped and created again without dropping the whole database. This is the recommended method unless there is a specific reason to drop and re-create the whole database.

ARY8010E Error setting client connect database partition number to "{0}".

Explanation: The script was unable to set the specified DB2 client connect partition number.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8011W Unable to get value for database configuration parameter *parameter_name* .

Explanation: The script was unable to obtain the specified configuration information. Processing continues; however, other problems might occur.

Response: Check the database to determine why it is not possible to obtain the configuration details, and correct the problem, if possible.

ARY8011E Unable to connect to database "{0}" on partition "{1}".

Explanation: The script was unable to connect to the database on the specified partition number.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8012E Error refreshing Java classes for database "{0}" on partition "{1}".

Explanation: The script was unable to refresh the Java class definitions on all the database partitions.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8013W Unable to get value for database manager configuration parameter *parameter_name* .

Explanation: The script was unable to obtain the specified configuration information. Processing continues; however, other problems might occur.

Response: Check the database manager configuration to determine why it is not possible to obtain the configuration details, and correct the problem, if possible.

ARY8013E Error setting current database partition to "{0}".

Explanation: The script was unable to set the connect and attach parameters for the specified DB2 database partition number.

Response: See previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8014W The database manager configuration parameter "{0}" is set to "{1}", which is less than the recommended value of "{2}".

Explanation: The specified database manager configuration parameter value is less than the indicated

recommended value. This might affect product operation.

Response: Increase the configuration parameter to a value equal to or greater than the recommended value.

ARY8014E Unable to attach to instance "{0}" on partition "{1}".

Explanation: The script was unable to attach to the specified DB2 instance using the specified partition number.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8015W An error occurred while deleting file "{0}".

Explanation: The script was unable to delete the specified file.

Response: Review the previous errors messages and ensure that the user has sufficient authority to perform the uninstallation.

ARY8015E Unexpected argument "{0}".

Explanation: The specified argument was not expected by the script.

Response: Correct the error and reissue the command.

ARY8016W Errors occurred while removing product files.

Explanation: One or more errors occurred during uninstallation.

Response: Review the errors and take appropriate action. If necessary, delete the product files and directories manually.

ARY8016E Required input parameter is not set for "{0}" argument.

Explanation: The argument requires a parameter that was not specified on the command line.

Response: Reissue the command and specify the required parameter.

ARY8017W An error occurred while deleting directory "{0}".

Explanation: The script was unable to delete the specified directory.

Response: Review the errors and take appropriate action. If necessary, delete the product files and directories manually.

ARY8017E Specified DB2 instance name "{0}" does not match current DB2 instance context "{1}".

Explanation: The script must be run for the current DB2 context. It cannot specify a different DB2 instance name.

Response: Correct the command line or establish the correct DB2 context, and reissue the command.

ARY8018W An error occurred while deleting product installation directory "{0}".

Explanation: The specified directory could not be deleted.

Response: Review the previous errors messages and ensure that the user has sufficient authority to perform the uninstallation.

ARY8018E Unable to determine home directory for user "{0}".

Explanation: The script was unable to determine the home directory for the specified DB2 instance user.

Response: Contact IBM Software Support.

ARY8019W Errors occurred while removing product directories.

Explanation: The specified directories could not be deleted.

Response: Review the previous errors messages and ensure that the user has sufficient authority to perform the uninstallation.

ARY8019E The DB2 instance context is not set.

Explanation: The DB2 instance context could not be established for running the script.

Response: Contact IBM Software Support.

ARY8020W Errors occurred during product uninstallation.

Explanation: One or more errors occurred while uninstalling the product.

Response: Review the errors and take appropriate action. If necessary, delete the product files and directories manually.

ARY8020E DB2 profile "{0}" does not exist or is not readable.

Explanation: The script was unable to establish the DB2 instance context because the DB2 profile file does not exist or is not readable at the indicated path.

Response: Contact IBM Software Support.

ARY8021W Errors occurred while removing product files.

Explanation: An error occurred while deleting product files from the installation directory.

Response: Review the error messages to determine the cause of the error. If necessary, delete the product files and directories manually.

ARY8021E Datastore database "{0}" already exists.

Explanation: The specified database already exists.

Response: To reinitialize the datastore, repeat the request with the option set to create the datastore in an existing database, or, ensure that no user data will be lost, and then drop and re-create the database that contains the datastore definitions. To create the datastore, either drop the existing database and reissue the request, or reissue the request with the options set to re-create the database or to use an existing database.

ARY8022E An error occurred while enabling the Self Tuning Memory Manager (STMM).

Explanation: The Self Tuning Memory Manager could not be enabled for the database.

Response: Refer to the previous DB2 error message to determine why the command failed. If the problem can be corrected, the Self Tuning Memory Manager can be enabled manually.

ARY8023E Unable to drop database "{0}". Resolve the problem and try again.

Explanation: The database could not be dropped.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible or drop the database manually, and reissue the command.

ARY8024E Unable to set database catalog partition to "{0}".

Explanation: The script was unable to set the database catalog partition to the specified partition number when creating the new database.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8025E Unable to attach to instance "{0}".

Explanation: The script was unable to attach to the specified DB2 instance.

Response: Refer to the previous DB2 error message to determine why the command failed.

ARY8026E Unable to create datastore database "{0}" on specified catalog partition "{1}".

Explanation: The script was unable to create the database on the specified catalog partition number.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8027E Unable to create datastore database "{0}".

Explanation: The script was unable to create the database.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8028E Error setting client database partition parameters.

Explanation: The script was unable to set the specified DB2 client connect and attach parameters.

Response: See previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8029E Unable to connect to database "{0}" on catalog partition.

Explanation: The script was unable to connect to the specified database on the catalog partition.

Response: Refer to the previous DB2 error message to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8030E A problem occurred while creating the datastore environment. Resolve the problem and try again.

Explanation: An unexpected error occurred while establishing the datastore database settings.

Response: Refer to the previous error messages to determine why the command failed. Correct the problem if possible and reissue the command.

ARY8031E Unable to create the datastore database. The DB2_COMPATIBILITY_VECTOR registry definition "{0}" is not compatible with the database. The "{1}" setting causes a conflict. To avoid this problem, you can temporarily remove the conflicting definition or reset the value to 0, restart the DB2 instance, and repeat the request to create the datastore database. You can restore the original DB2_COMPATIBILITY_VECTOR definition after the database has been created.

Explanation: The current DB2_COMPATIBILITY_VECTOR settings are not compatible with the datastore database. One or more of the option settings may cause a problem when the database is created.

Response: Follow the directions specified in the error message. You can remove the conflicting settings from the vector to avoid the problem, and restore the original settings after the database is created.

ARY8032E Installation target directory pathname must be specified.

Explanation: A required parameter was not specified.

Response: Reissue the command and specify the installation target directory location.

ARY8033E The product cannot be installed.

Explanation: It was not possible to install the product with the settings specified by the user.

Response: Review the preceding error messages, correct the problem, and then try the installation again.

ARY8034E Installation target directory "{0}" is same as source directory.

Explanation: The product cannot be installed into the same directory tree where the installation source files are.

Response: Choose a different target directory for the installation.

ARY8035E Specified installation directory "{0}" does not exist. Repeat the request and either specify a valid installation directory path or specify the option to create the target directory.

Explanation: The installation target directory does not exist, but the installation did not request to create it.

Response: Specify an installation directory that exists or specify the option to create the target directory and repeat the request.

ARY8036E Unable to create installation target directory "{0}", status="{1}".

Explanation: The installation target directory could not be created.

Response: Check the target directory path and review the error status. Correct the problem and then try the installation again.

ARY8037E 1 file in the installation directory was not updated.

Explanation: One file was not overwritten during installation.

Response: No action required.

ARY8038E {0} files in the installation directory were not updated.

Explanation: One or more file were not overwritten during installation.

Response: No action required.

ARY8039E Unable to run "{0}" to set access on product files.

Explanation: It was not possible to run the specified command to set file permissions on the installed files.

Response: Check that the command file exists and has the correct permissions. Correct the problem and repeat the install, or run the command directly to set the correct permissions on the installed files.

ARY8040E Error occurred while installing product files; status="{0}". Verify that the target filesystem has sufficient free space.

Explanation: The indicated error occurred while installing the product files.

Response: Review the error status and any previous error messages. If necessary, correct the problem and repeat the installation request.

ARY8041E Invalid installation directory: "{0}" The product should not be installed using that path name. Correct the installation if necessary and specify a valid installation directory path.

Explanation: An invalid target directory path was used for installing the product.

Response: Choose a different installation directory path and repeat the request.

ARY8042E Required permissions not set on directory: "{0}". Current settings are: {1}

Explanation: Incorrect permission settings were detected on the specified product directory.

Response: Run the setAccessscript to fix the product owner and permission settings, or change the permissions manually.

ARY8043E Required owner not set to 'root' on file: "{0}".

Explanation: The specified file must be owned by the root user.

Response: Run the setAccessscript to fix the product owner and permission settings, or change the permissions manually.

ARY8044E Required permissions not set on file: "{0}". Current settings are: {1}

Explanation: Incorrect permission settings were detected on the specified product file.

Response: Run the setAccessscript to fix the product owner and permission settings, or change the permissions manually.

ARY8045E Invalid validate mode: "{0}".

Explanation: Unsupported option for validating the installation.

Response: Repeat the request after correcting the command option.

ARY8046E Specified installation directory path "{0}" is not a directory. Specify a valid installation directory path and then try again.

Explanation: If the specified installation path exists, it must be a directory.

Response: Choose a different installation location and then repeat the request.

ARY8047E Specified installation directory path "{0}" does not exist or is not a directory. Specify a valid installation directory path and try again.

Explanation: There is a problem with the specified installation directory path. If the specified installation path exists, it must be a directory.

Response: Choose a different installation location and repeat the request.

ARY8048E **Verification failed.**

Explanation: The verification of the installed product files failed.

Response: Review the previous messages to determine what is wrong. If necessary, reinstall the product to fix the problem.

ARY8049E **Specified inventory file "{0}" does not exist or is not readable. Specify a valid path name for the inventory file and try again.**

Explanation: The product inventory file cannot be read.

Response: Determine the correct location of the inventory file and repeat the request if possible.

ARY8050E **Specified validation program "{0}" does not exist or is not executable. Specify a valid path name for the validation program and try again.**

Explanation: The specified product validation program cannot be run.

Response: Determine the correct location of the validation program, correct any problems, and repeat the request if possible.

ARY8051E **Unable to change to directory "{0}". Verify that the installation directory is correct.**

Explanation: The installation script cannot change to the specified directory.

Response: Verify the correct installation directory and repeat the request after correcting any problems.

ARY8052E **Product validation program "{0}" does not exist or is not executable.**

Explanation: The specified product validation program cannot be run.

Response: Determine the correct location of the validation program, correct any problems, and repeat the request if possible.

ARY8053E **Product inventory file does not exist or is not readable: "{0}". If the validation script was not run from the product installation directory, rerun the script and specify the installation directory path name as an argument.**

Explanation: The product inventory file cannot be read.

Response: Determine the correct location of the

inventory file and repeat the request if possible.

ARY8054E **Product directory not found in the installation directory: "{0}"**

Explanation: An expected product directory was missing from the installation tree. The installation may be corrupted.

Response: Check the installation files. Reinstall the product if necessary to correct the problem.

ARY8055E **Product file not found in the installation directory: "{0}"**

Explanation: An expected product file was missing from the installation tree. The installation may be corrupted.

Response: Check the installation files. Reinstall the product if necessary to correct the problem.

ARY8056E **Unable to validate file signature using command: "{0}".**

Explanation: A problem occurred when validating the file signature.

Response: Check the installation files. Reinstall the product if necessary to correct the problem.

ARY8057E **Product file signature does not match expected value for file: "{0}".**

Explanation: The specified file signature does not match the expected value. The file may have been changed after the original installation. This may not be a problem if the file was updated to fix a product problem.

Response: Verify if the changed file is valid. Reinstall the product if necessary to correct the problem.

ARY8058E **Specified argument "{0}" is not a directory or does not exist. Specify a valid installation directory path and try again.**

Explanation: If the specified installation path exists, it must be a directory.

Response: Choose a different installation location and repeat the request.

ARY8059E **The product files cannot be updated.**

Explanation: The product files permission settings cannot be updated due to a previously reported problem.

Response: Review the error messages to determine the cause of the problem. Reinstall the product if necessary to correct the problem.

ARY8060E Unable to perform path safety checks. The verification script "{0}" does not exist or is not executable.

Explanation: A required product file is missing. The product ownership and permissions cannot be updated.

Response: Check the installation files. Reinstall the product if necessary to correct the problem.

ARY8061E Error setting owner on installation directory "{0}".

Explanation: The script was unable to update the owner setting on the indicated directory.

Response: Review the error messages and verify that the user running the command has sufficient authority to modify ownership on the directory.

ARY8062E Error setting permissions on installation directory "{0}".

Explanation: The script was unable to update the permission settings on the indicated directory.

Response: Review the error messages and verify that the user running the command has sufficient authority to modify permissions on the directory.

ARY8063E Error setting owner on files in product directory "{0}".

Explanation: The script was unable to update the owner setting on files in the indicated directory.

Response: Review the error messages and verify that the user running the command has sufficient authority to modify ownership on the directory.

ARY8064E Error setting permissions in product directory "{0}".

Explanation: The script was unable to update the permission settings on the contents of the indicated directory.

Response: Review the error messages and verify that the user running the command has sufficient authority to modify permissions in the directory.

ARY8065E Error setting permissions on files in product directory "{0}".

Explanation: The script was unable to update the permission settings on the contents of the indicated directory.

Response: Review the error messages and verify that the user running the command has sufficient authority to modify permissions in the directory.

ARY8066E Error setting permissions on product file "{0}".

Explanation: The script was unable to update the permission settings on the specified file.

Response: Review the error messages and verify that the user running the command has sufficient authority to modify permissions on the specified file.

ARY8068E Path name of directory to be uninstalled must be specified.

Explanation: A required argument was omitted.

Response: Specify the name of the directory for uninstalling the product files.

ARY8069E The product cannot be uninstalled.

Explanation: It was not possible to uninstall the product with the settings specified by the user.

Response: Review the preceding error messages, correct the problem, and then try the request again.

ARY8070E Product inventory file "{0}" does not exist or is not readable. Specify a valid installation directory path name as an argument. If the uninstallation script is not run from the product installation directory, rerun the script and specify the installation directory path name as an argument. The product cannot be uninstalled.

Explanation: The product could not be uninstalled. The installation directory path may be incorrect, or the installation directory contents may be incomplete or corrupted.

Response: Verify that the installation directory path is correct. If necessary, delete the product files and directories manually.

ARY8077E Errors occurred while performing GRANT commands.

Explanation: Errors occurred while trying to grant authorities to the specifies users or groups.

Response: Review the previous error messages and verify that the user has sufficient authority to grant the requested authorities.

ARY8078E Unknown operation code "{0}".

Explanation: Unexpected operation value for creating the datastore.

Response: Contact IBM Software Support.

ARY8079E Close any existing connections to the database "{0}".

Explanation: Unable to process the request due to outstanding connections to the database.

Response: Close any existing connections and try the request again.

ARY8080E Error occurred while installing product files.

Explanation: An error occurred while installing the product files.

Response: Review the previous error messages and reissue the request after correcting the problem.

ARY8081E Unable to create target directory "{0}", status="{1}".

Explanation: The script was unable to create the indicated target directory. Possible reasons include lack of disk space, insufficient user authority, or an incorrect location path.

Response: Review the error messages to determine the cause of the error. Reissue the request after correcting the problem.

ARY8082E Error occurred while copying files to "{0}".

Explanation: An error occurred while installing the product files. Possible reasons include lack of disk space, insufficient user authority, or an incorrect location path.

Response: Review the error status and any previous error messages. Correct the problem and reissue the request.

ARY8083E Missing inventory file path name for "{0}" argument.

Explanation: A required value is missing for the specified argument.

Response: Reissue the request and supply the missing value.

ARY8084E Specified installation directory path "{0}" does not exist.

Explanation: The specified directory does not exist.

Response: Determine the correct installation directory path and reissue the command.

ARY8085E The uninstallation script cannot be run to delete its own directory tree.

Explanation: On Windows platforms, the uninstallation script cannot delete itself, so it cannot be run from within the product installation directory tree.

Response: Copy the uninstallation script to another directory outside of the installation directory tree. Run the script to uninstall the product, then delete the uninstallation script.

ARY8086E Error occurred while removing the product installation directory.

Explanation: An error occurred while deleting the product installation directory.

Response: Review the error messages to determine the cause of the error. If necessary, delete the product files and directories manually.

ARY8088E Invalid operation value "{0}".

Explanation: Invalid value specified for the operation argument.

Response: Reissue the request after correcting the command arguments.

ARY8089E Required input parameter is not set for "{0}" argument.

Explanation: A required value is missing for the specified argument.

Response: Reissue the request after correcting the command arguments.

ARY8090E Missing temporary directory path name for "{0}" argument.

Explanation: A valid temporary directory path must be specified for the specified argument.

Response: Reissue the request after correcting the command arguments.

ARY8091E Specified temporary directory path "{0}" does not exist. Specify a valid directory path name and try again.

Explanation: The specified temporary directory does not exist.

Response: Reissue the request and specify a valid temporary directory path.

ARY8092E The specified temporary directory path name "{0}" is invalid. Specify a valid temporary directory path that exists and is not located within the installation tree to be uninstalled.

Explanation: The specified temporary directory pathname "{0}" is invalid. Specify a valid temporary directory path that exists and is not located within the installation tree to be uninstalled.

Response: Reissue the request and specify a valid temporary directory path.

ARY8093I File "{0}" was not deleted.

Explanation: The specified file could not be deleted.

Response: Review and error messages to determine the cause of the problem. If necessary, delete the product files and directories manually.

ARY8093E Unable to copy file "{0}" to specified temporary directory "{1}".

Explanation: The specified file could not be copied to the indicated temporary directory. Possible reasons include lack of disk space, insufficient user authority, or an incorrect location path.

Response: Review the error messages to determine the cause of the error. Correct the problem and reissue the request.

ARY8094I Directory "{0}" was not deleted.

Explanation: The specified directory could not be deleted.

Response: Review and error messages to determine the cause of the problem. If necessary, delete the product files and directories manually.

ARY8095I DB2 Recovery Expert datastore creation timestamp (UTC):{0}.

Explanation: This message identifies the date and time when the datastore was created.

Response: No action required.

ARY8096I DB2 Recovery Expert datastore version: {0}

Explanation: This displays diagnostic version information about the datastore.

Response: No action required.

ARY8096E The datastore cannot be administered by the root user. You must perform datastore administration tasks as the DB2 instance owner or as another user that has authority to access the desired DB2 instance."

Explanation: Datastore tasks cannot be performed as the root user.

Response: Run the datastore request as some other user, such as the DB2 instance owner or another user that has the authority to access the desired DB2 instance.

ARY8097I SLR migrated successfully for instance {0}, partition {1}, database {2}.

Explanation: Migration of the SLR persistent data storage is successful for the specified database in the specified instance.

Response: No action required.

ARY8097E Unable to establish DB2 context for user "{0}". Verify that the specified instance name or current user matches a valid DB2 instance.

Explanation: It was not possible to establish the environment context for the specified DB2 instance.

Response: Review any previous error messages and verify that the correct DB2 instance name was specified. Correct any problem and reissue the command.

ARY8098E Invalid argument combination: -r can only be specified with -o g

Explanation: The specified arguments cannot be used together.

Response: Reissue the request after correcting the command arguments.

ARY8099E An error occurred while updating datastore authority for ID "{0}".

Explanation: It was not possible to update the authorities on the datastore database for the specified user ID.

Response: Verify that the issuing user has sufficient authority to modify the database authorities for the specified user. Correct the problem and reissue the request.

ARY8100E Duplicate argument "{0}" specified.

Explanation: A command argument was specified twice.

Response: Reissue the command after resolving the problem.

ARY8101E Script "{0}" finished with error

Explanation: The specified script finished with error status.

Response: Review the previous error messages to determine the cause of the problem. Reissue the command after resolving any problems.

ARY8102E Command "{0}" finished with error

Explanation: The specified command finished with error status.

Response: Review the previous error messages to determine the cause of the problem. Reissue the command after resolving any problems.

ARY8103E Unable to create the datastore in existing database "{0}" defined using codeset "{1}". The database codeset must be "UTF-8".

Explanation: The database could not be created because the code set is not UTF-8.

Response: Create the datastore in a new database, or in a different database that uses the UTF-8 code set.

ARY8104E Unable to install DB2 Recovery Expert with user ID "{0}". DB2 Recovery Expert must be installed by the "root" user.

Explanation: The product cannot be installed by a non-root user.

Response: Reissue the command after resolving the problem. Perform the installation while running as the root user or as another user with root authority. The "sudo" and "su" commands can be used to achieve root authority if necessary.

ARY8105E Unable to create the datastore in existing database "{0}" that has configuration parameter "{1}" set to "{2}". The parameter must be set to "{3}".

Explanation: The datastore cannot be created in the specified database that has the indicated configuration parameter setting.

Response: Change the configuration parameter or create the datastore in a different database.

ARY8106E Invalid argument combination: "{0}" may only be specified with "{1}".

Explanation: The specified command argument can only be used with one of the listed command arguments.

Response: Correct the command line and perform the operation again.

ARY8107E A problem was detected while upgrading the datastore definitions in database "{0}". It is necessary to re-create the datastore. If the problem persists, contact IBM Software Support.

Explanation: The datastore definitions in the specified database cannot be upgraded.

Response: Reinitialize the datastore. Previous product data that is stored in the datastore tables will be lost.

ARY8123E Unable to read table space definition file "{0}".

Explanation: The input file containing table space definitions cannot be opened or read using the provided path name.

Response: Verify that the path name to the DDL input file is correct and that the file is readable. Correct the path name if necessary and perform the operation again.

ARY8124E Interactive mode argument value "{0}" is not valid.

Explanation: The specified argument value is not valid.

Response: Try the operation again with one of the following argument values: 1 to run in interactive mode; 0 to disable prompt messages

ARY8125E The DB2 Recovery Expert datastore cannot be created in database "{0}" that has the "dft_table_org" configuration parameter value set to "{1}". The value must be set to "ROW" for DB2 to create row organized tables by default. Choose or create a different database that has compatible settings, and then try the request again.

Explanation: The DB2 Recovery Expert datastore does not support using column organized tables.

Response: Create the datastore in a database that is set to create row organized tables by default. If you choose to create the datastore in an existing database and you change the "dft_table_org" database configuration parameter to "ROW", do not alter the parameter afterwards. The "dft_table_org" parameter must remain set to "ROW" as long as DB2 Recovery Expert is used with that datastore database.

ARY8126E DB2 Recovery Expert cannot create the datastore database "{0}" because the "DB2_WORKLOAD" registry variable is set to "ANALYTICS". This setting initializes the database configuration parameter "dft_table_org" in new databases so that DB2 will create column organized tables by default. Choose or create a different database that has compatible settings, and then try the request again.

Explanation: The DB2 Recovery Expert datastore does not support using column organized tables.

Response: Create the datastore in a database that is set to create row organized tables by default. If you choose to create the datastore in an existing database and you change the "dft_table_org" database configuration parameter value to "ROW", do not alter the parameter afterwards. The "dft_table_org" parameter must remain set to "ROW" as long as DB2 Recovery Expert is used with that datastore database.

ARY8127W Unable to migrate the SLR persistent data storage for instance {0}, partition {1}, database {2}, data directory {3}. Recreate the SLR for this database.

Explanation: During the migration of the SLR, persistent storage failure has occurred. Migration is not possible for the specified database in the specified instance.

Response: Remove the contents of the SLR directory in the specified data directory and recreate the SLR.

Notices

This information was developed for products and services offered in the U.S.A. If the products or services are targeted for offering outside of the U.S.A., please consult your local IP Law department.

This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to: IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510,
Japan

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

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

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

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

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

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003 U.S.A.

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© Copyright IBM Corp. 2014 All rights reserved.

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

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

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

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

Index

A

- accessibility 14
- administering the DB2 Recovery Expert Server 62
- application program logic
 - Log Analysis cannot mimic 163
- architecture 9
- authorities 26

B

- backing up
 - product application files 53
- backup files 97
- backups 142
- browser client overview 83

C

- changing authorities
 - existing datastore connection 74
- CLP 6
- CLP (command line processor) 139
- CLP configuration 140
- command information 142
- command line processor (CLP) 139
- command set 142
- comments
 - methods for providing x
- configuration 61
- configuration file recex.properties
 - Java heap size requirements 28
- configuration repository 6
- Configuration tasks
 - terms and field definitions 68, 69, 75, 77, 79
- configuring
 - remote access to database server machines 53
- configuring Linux systems 40, 41
- configuring UNIX systems 40, 41
- considerations 161
- considerations for DPF environments 171
- conventions
 - highlighting ix
- creating 98
 - datastore repository database using the command line 72
- currently committed semantics 28, 161

D

- DATA CAPTURE CHANGES 28, 104, 161
- database location 84
- database objects 129
- Database Partitioning Facility (DPF) 5
- database server components 6

- database server components status 75
- database server machines 53
- datastore 72
- datastore connection 71, 74
- datastore repository 6
- datastore repository database 84
- datastore repository database using the command line 72
- DB2
 - support 168
- DB2 Recovery Expert
 - settings 79
- DB2 Recovery Expert repository database 65
- DB2 Recovery Expert security 65, 66
- DB2 Recovery Expert Server
 - changing settings 63
- DB2 support limitations 168
- delete
 - existing datastore connection 71
- dependent objects 128
- deployment
 - dropped object restore 6
- details 107
- diagnostic information
 - gathering 179, 181, 182
- documentation 14
- documentation changes 1
- DPF 5
- DPF environments
 - considerations 171

F

- feature interaction with DB2 12
- feature interaction with host 12
- features and benefits 3, 5
 - Log Analysis 4
 - Recovery 4
- feedback
 - methods for providing x

G

- gathering information 182, 183
- glossary 6

H

- hardware 17
- high availability disaster recovery (HADR) 172
- highlighting conventions ix
- host name resolution 53

I

- index values limitations
 - floating point numbers 173

- index values limitations (*continued*)
 - Log Analysis cannot mimic
 - application program logic 173
 - tables without a primary or unique key 173
- installation 43
 - requirements and recommendations 23
- installation requirements 17
- installation roadmap 43
- installing
 - Command Line Processor client 50, 62
 - database server components 76
 - DB2 Recovery Expert 45
 - permanent license 52

J

- Java CLP configuration 140
- Java Command Line Processor (CLP) 6, 139

L

- license
 - installing 52
- licenses 39
- Log Analysis 105, 113
 - connecting to a database 106
 - creating column-level filters 118
 - filters 111
 - objects 108
 - Report 119
 - results 114
 - specifying options 112
- Log Analysis cannot mimic application program logic 163
- Log Analysis restrictions and considerations
 - quiesce in share mode
 - log chain restrictions 163
 - SET INTEGRITY statement 163
- Log Analysis tool 6
- log chain restrictions 177
- log files 97
- logs 142
- LRI 6
- LSN 6
- LSO 6

M

- main browser client view 85
- managed database 6, 66
- masked update reconstruction 6
- materialized query tables (MQT) 175
- messages
 - methods for accessing ix
- metadata repository 6
- Minimum Recovery Time (MRT) 6

MQT restrictions and considerations 175
MRT 6
multiple IP addresses for a host
 name 53
multiple object recovery 5, 6

N

new tab button 88

O

Object History 5, 6, 95, 96
operating systems
 Windows 55
 Windows 2008 55
 Windows 7 55
 Windows Server 2003 55
 Windows Vista 55
operational prerequisites 28
overview 1

P

Point-in-time object recovery 6
preferences 80
preparing environment 17
privileges 26
problems
 diagnostic information about 179,
 181, 182
process flow 9

Q

Quiet Time report 6

R

rebuild
 SLR (Schema Level Repository) 100
recovery 124
 backup images 127
 log files 127
 selecting a database object or
 pattern 124
 selecting a point in time 126
 user-defined structured types 168
Recovery 123
Recovery History File (RHF) 28, 142
Recovery plan 6, 131
Recovery plans 85
Recovery scenario 6
recovery status 137
related tables 128
remote access 53
Remote database processing 78, 160
Remote repository 6
report 107
restrictions 161

S

Schema Level Repository (SLR) 6, 85,
95, 97, 98
 dropping 102
 operations 98
 pruning 102
 rebuild 100
 reviewing status 103
 updating from logs 104
service information ix
Session 6
sessions 88, 90
SET INTEGRITY statement
 Log Analysis restrictions and
 considerations
 restrictions for multidimensional
 clustering (MDC) tables 163
share mode
 Log Analysis restrictions and
 considerations 163
SLR (Schema Level Repository) 85, 95,
97
 dropping 102
 operations 98
 pruning 102
 reviewing status 103
 updating from logs 104
SLR range
 updating 101
software 17
special characters 142
Specification 6
specifications 88, 89, 90
standard action buttons 87
summary of changes 1
support
 required information 179, 181, 182
support information ix
system temporary table space page
 size 163

T

table space
 recovery restriction 168
temporary disk space 28
terminology 6
Tips for using Recovery Expert 91
Translated table space recovery 6
Troubleshooting 182, 183

U

UDST (User-Defined Structured Types)
 table recovery 168
 table space recovery
 online backup restriction 168
uninstalling
 Command Line Processor client 58
uninstalling
 DB2 Recovery Expert 57
upgrade 39, 43
upgrading
 database server components 76
User-Defined Structured Types (UDST)
 table recovery 168

User-Defined Structured Types (UDST)
(*continued*)
 table space recovery
 online backup restriction 168

W

what's new 1
Windows 2008 55
Windows 7 55
Windows Server 2003 55
Windows targets 54
Windows Vista 55
Windows XP 55
workflow 83



Product Number: 5724-H93

Printed in USA

SC19-3992-02

