

IBM<sup>®</sup> DB2 Universal Database<sup>™</sup>



# Data Recovery and High Availability Guide and Reference

*Version 8.2*



IBM® DB2 Universal Database™



# Data Recovery and High Availability Guide and Reference

*Version 8.2*

Before using this information and the product it supports, be sure to read the general information under *Notices*.

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at [www.ibm.com/shop/publications/order](http://www.ibm.com/shop/publications/order)
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at [www.ibm.com/planetwide](http://www.ibm.com/planetwide)

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 2001 - 2004. All rights reserved.**

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

# Contents

## About this book . . . . . vii

Who should use this book . . . . . vii

How this book is structured. . . . . vii

## Part 1. Data recovery . . . . . 1

### Chapter 1. Developing a good backup and recovery strategy. . . . . 3

Developing a backup and recovery strategy . . . . . 3

Automated backup operations . . . . . 6

Deciding how often to back up . . . . . 6

Storage considerations for recovery. . . . . 8

Keeping related data together. . . . . 9

Using different operating systems . . . . . 9

Crash recovery . . . . . 9

Crash recovery - details . . . . . 11

Recovering damaged table spaces . . . . . 11

Recovering table spaces in recoverable databases 12

Recovering table spaces in non-recoverable

databases . . . . . 12

Reducing the impact of media failure. . . . . 13

Reducing the impact of transaction failure . . . . 15

Recovering from transaction failures in a

partitioned database environment . . . . . 15

Recovering from the failure of a database

partition server . . . . . 18

Recovering indoubt transactions on the host

when DB2 Connect has the DB2 Syncpoint

Manager configured . . . . . 19

Recovering indoubt transactions on the host

when DB2 Connect does not use the DB2

Syncpoint Manager. . . . . 20

Disaster recovery . . . . . 21

Version recovery. . . . . 22

Rollforward recovery . . . . . 22

Incremental backup and recovery . . . . . 25

Incremental backup and recovery - details . . . . 26

Restoring from incremental backup images. . . . 26

Limitations to automatic incremental restore . . 28

Monitoring the progress of backup, restore and

recovery operations. . . . . 30

Understanding recovery logs . . . . . 31

Recovery log details . . . . . 33

Log mirroring . . . . . 33

Reducing logging with the NOT LOGGED

INITIALLY parameter . . . . . 34

Configuration parameters for database logging 35

Managing log files . . . . . 42

Log file allocation and removal. . . . . 44

Managing log files through log archiving . . . . 45

Blocking transactions when the log directory file

is full . . . . . 47

On demand log archive . . . . . 47

Using raw logs . . . . . 48

Including log files with a backup image. . . . . 49

How to prevent losing log files. . . . . 50

Understanding the recovery history file . . . . . 51

Recovery history file - garbage collection . . . . 53

Garbage Collection . . . . . 53

Understanding table space states . . . . . 55

Enhancing recovery performance . . . . . 56

### Chapter 2. Database backup . . . . . 59

Backup overview . . . . . 59

Displaying backup information. . . . . 61

Privileges, authorities, and authorization required to

use backup . . . . . 62

Using backup. . . . . 62

Backing up to tape . . . . . 64

Backing up to named pipes . . . . . 65

BACKUP DATABASE . . . . . 66

db2Backup - Backup database . . . . . 71

Backup sessions - CLP examples . . . . . 78

Optimizing backup performance . . . . . 79

### Chapter 3. Database restore . . . . . 81

Restore overview . . . . . 81

Privileges, authorities, and authorization required to

use restore. . . . . 82

Using restore . . . . . 82

Using incremental restore in a test and production

environment . . . . . 83

Redefining table space containers during a restore

operation (redirected restore) . . . . . 85

Restoring to an existing database . . . . . 86

Restoring to a new database. . . . . 87

RESTORE DATABASE. . . . . 87

db2Restore - Restore database . . . . . 96

Restore sessions - CLP examples . . . . . 107

Optimizing restore performance . . . . . 109

### Chapter 4. Rollforward recovery . . . . . 111

Rollforward overview . . . . . 111

Privileges, authorities, and authorization required

to use rollforward . . . . . 113

Using rollforward . . . . . 113

Rolling forward changes in a table space . . . . 114

Recovering a dropped table . . . . . 118

Recovering data with the load copy location file 120

Synchronizing clocks in a partitioned database

system . . . . . 121

Client/server timestamp conversion. . . . . 122

ROLLFORWARD DATABASE . . . . . 123

db2Rollforward - Rollforward Database . . . . . 131

Rollforward sessions - CLP examples . . . . . 140

### Chapter 5. Database recover . . . . . 145

Recover overview . . . . . 145

Privileges, authorities, and authorization required

to use recover . . . . . 146

Using recover . . . . .	146
Client/server timestamp conversion . . . . .	146
RECOVER DATABASE . . . . .	147
db2Recover - Recover database . . . . .	152

## Part 2. High availability . . . . . 159

### Chapter 6. Introducing high availability and failover support . . . 161

High availability . . . . .	161
High availability through log shipping . . . . .	163
High availability through online split mirror and suspended I/O support . . . . .	164
Online split mirror handling . . . . .	165
Using a split mirror to clone a database . . . . .	165
Using a split mirror as a standby database . . . . .	167
Using a split mirror as a backup image. . . . .	168
Fault monitor facility for UNIX-based systems . . . . .	169
db2fm - DB2 Fault Monitor. . . . .	171

### Chapter 7. High availability disaster recovery (HADR) . . . . . 173

High availability disaster recovery overview . . . . .	173
System requirements for high availability disaster recovery . . . . .	174
Restrictions for high availability disaster recovery . . . . .	175
Database configuration for high availability disaster recovery . . . . .	176
Standby database states in high availability disaster recovery . . . . .	179
Synchronization modes for high availability disaster recovery . . . . .	181
Automatic client reroute and high availability disaster recovery . . . . .	184
Index logging and high availability disaster recovery . . . . .	185
Replicated operations for high availability disaster recovery . . . . .	186
Non-replicated operations for high availability disaster recovery . . . . .	187
Cluster managers and high availability disaster recovery . . . . .	188
Initializing high availability disaster recovery . . . . .	189
START HADR . . . . .	191
db2HADRStart - Start HADR . . . . .	193
Stopping high availability disaster recovery . . . . .	194
STOP HADR . . . . .	196
db2HADRStop - Stop HADR . . . . .	197
Switching database roles in high availability disaster recovery . . . . .	198
Using HADR takeover during a failover . . . . .	199
Reintegrating a database after a takeover operation . . . . .	202
Performing a rolling upgrade in a high availability disaster recovery environment. . . . .	203
TAKEOVER HADR . . . . .	204
db2HADRTakeover - Take Over as Primary Database . . . . .	206

### Chapter 8. Cluster support on AIX . . . 209

High Availability Cluster Multi-Processing support 209

### Chapter 9. Cluster support on the Windows operating system . . . . . 215

Microsoft Cluster Server support . . . . . 215

### Chapter 10. Cluster support for the Solaris Operating Environment. . . . . 219

Cluster support for the Solaris Operating Environment . . . . .	219
Sun Cluster 3.0 support . . . . .	221
VERITAS Cluster Server support . . . . .	224

## Part 3. Appendixes . . . . . 229

### Appendix A. How to read the syntax diagrams . . . . . 231

### Appendix B. Warning, error and completion messages. . . . . 235

### Appendix C. Additional DB2 commands . . . . . 237

System commands. . . . .	237
db2adutl - Managing DB2 objects within TSM . . . . .	237
db2cckbcp - Check Backup . . . . .	243
db2ckrst - Check Incremental Restore Image Sequence . . . . .	247
db2flsn - Find Log Sequence Number . . . . .	249
db2inidb - Initialize a Mirrored Database . . . . .	250
db2mscs - Set up Windows Failover Utility . . . . .	252
db2rfpen - Reset rollforward pending state . . . . .	255
CLP commands . . . . .	255
ARCHIVE LOG . . . . .	255
INITIALIZE TAPE. . . . .	257
LIST HISTORY . . . . .	258
PRUNE HISTORY/LOGFILE . . . . .	260
REWIND TAPE . . . . .	262
SET TAPE POSITION . . . . .	262
UPDATE HISTORY FILE . . . . .	263

### Appendix D. Additional APIs and associated data structures . . . . . 265

db2ArchiveLog - Archive Active Log . . . . .	265
db2HistoryCloseScan - Close History File Scan . . . . .	267
db2HistoryGetEntry - Get Next History File Entry . . . . .	268
db2HistoryOpenScan - Open History File Scan . . . . .	271
db2HistoryUpdate - Update History File . . . . .	275
db2Prune - Prune History File. . . . .	278
db2ReadLogNoConn - Read Log Without a Database Connection . . . . .	280
db2ReadLogNoConnInit - Initialize Read Log Without a Database Connection . . . . .	283
db2ReadLogNoConnTerm - Terminate Read Log Without a Database Connection . . . . .	285
db2ReadLog - Asynchronous Read Log. . . . .	286
db2HistData. . . . .	289
SQLU-LSN . . . . .	293

	<b>Appendix E. Recovery sample programs . . . . .</b>	<b>295</b>		Installing the DB2 Information Center using the DB2 Setup wizard (UNIX) . . . . .	368
	Sample programs with embedded SQL . . . . .	295		Installing the DB2 Information Center using the DB2 Setup wizard (Windows) . . . . .	370
	<b>Appendix F. Tivoli Storage Manager 325</b>			Invoking the DB2 Information Center . . . . .	372
	Configuring a Tivoli Storage Manager client . . . . .	325		Updating the DB2 Information Center installed on your computer or intranet server . . . . .	373
	Considerations for using Tivoli Storage Manager	326		Displaying topics in your preferred language in the DB2 Information Center . . . . .	374
	<b>Appendix G. User exit for database recovery . . . . .</b>	<b>329</b>		DB2 PDF and printed documentation . . . . .	375
	Sample user exit programs . . . . .	329		Core DB2 information . . . . .	375
	Calling format . . . . .	330		Administration information . . . . .	375
	Error handling . . . . .	331		Application development information . . . . .	376
	<b>Appendix H. Backup and restore APIs for vendor products . . . . .</b>	<b>333</b>		Business intelligence information . . . . .	377
	APIs for backup and restore to storage managers	333		DB2 Connect information . . . . .	377
	Operational overview . . . . .	333		Getting started information . . . . .	377
	Operational hints and tips . . . . .	338		Tutorial information . . . . .	378
	Invoking a backup or a restore operation using vendor products . . . . .	338		Optional component information . . . . .	378
	sqluvint - Initialize and Link to Device . . . . .	340		Release notes . . . . .	379
	sqluvget - Reading Data from Device . . . . .	343		Printing DB2 books from PDF files . . . . .	380
	sqluvput - Writing Data to Device . . . . .	344		Ordering printed DB2 books . . . . .	380
	sqluvend - Unlink the Device and Release its Resources . . . . .	346		Invoking contextual help from a DB2 tool . . . . .	381
	sqluvdel - Delete Committed Session . . . . .	348		Invoking message help from the command line processor . . . . .	382
	db2VendorQueryApiVersion - Query Device Supported API Level . . . . .	349		Invoking command help from the command line processor . . . . .	382
	db2VendorGetNextObj - Get Next Object on Device	349		Invoking SQL state help from the command line processor . . . . .	383
	DB2-INFO . . . . .	351		DB2 tutorials . . . . .	383
	VENDOR-INFO . . . . .	353		DB2 troubleshooting information . . . . .	384
	INIT-INPUT . . . . .	354		Accessibility . . . . .	385
	INIT-OUTPUT . . . . .	354		Keyboard input and navigation . . . . .	385
	DATA . . . . .	355		Accessible display . . . . .	385
	RETURN-CODE . . . . .	355		Compatibility with assistive technologies . . . . .	386
	APIs for compressed backups . . . . .	356		Accessible documentation . . . . .	386
	Compression plug-in interface . . . . .	356		Dotted decimal syntax diagrams . . . . .	386
	<b>Appendix I. DB2 Universal Database technical information . . . . .</b>	<b>363</b>		Common Criteria certification of DB2 Universal Database products . . . . .	388
	DB2 documentation and help . . . . .	363		<b>Appendix J. Notices . . . . .</b>	<b>389</b>
	DB2 documentation updates . . . . .	363		Trademarks . . . . .	391
	DB2 Information Center . . . . .	364		<b>Index . . . . .</b>	<b>393</b>
	DB2 Information Center installation scenarios . . . . .	365		<b>Contacting IBM . . . . .</b>	<b>399</b>
				Product information . . . . .	399





---

## About this book

This book provides detailed information about, and shows you how to use, the IBM DB2 Universal Database (UDB) backup, restore, and recovery utilities. The book also explains the importance of high availability, and describes DB2 failover support on several platforms.

---

## Who should use this book

This manual is for database administrators, application programmers, and other DB2 UDB users who are responsible for, or who want to understand, backup, restore, and recovery operations on DB2 database systems.

It is assumed that you are familiar with DB2 Universal Database, Structured Query Language (SQL), and with the operating system environment in which DB2 UDB is running. This manual does not contain instructions for installing DB2, which depend on your operating system.

---

## How this book is structured

The following topics are covered:

### Data Recovery

#### Chapter 1, “Developing a good backup and recovery strategy”

Discusses factors to consider when choosing database and table space recovery methods, including backing up and restoring a database or table space, and using rollforward recovery.

#### Chapter 2, “Database backup”

Describes the DB2 backup utility, used to create backup copies of a database or table spaces.

#### Chapter 3, “Database restore”

Describes the DB2 restore utility, used to rebuild damaged or corrupted databases or table spaces that were previously backed up.

#### Chapter 4, “Rollforward recovery”

Describes the DB2 rollforward utility, used to recover a database by applying transactions that were recorded in the database recovery log files.

#### Chapter 5, “Database recover”

Describes the DB2 recover utility, which performs the necessary restore and rollforward operations to recover a database to a specified time, based on information found in the recovery history file.

### High Availability

#### Chapter 6, “Introducing high availability and failover support”

Presents an overview of the high availability failover support that is provided by DB2.

#### Chapter 7, “High availability disaster recovery (HADR)”

Discusses the concepts and procedures required to set up and manage a high availability disaster recovery (HADR) environment.

**High Availability Cluster Multi-Processing support**

Discusses DB2 support for high availability failover recovery on AIX, which is currently implemented through the Enhanced Scalability (ES) feature of High Availability Cluster Multi-processing (HACMP) for AIX.

**Microsoft Cluster Server support**

Discusses DB2 support for high availability failover recovery on Windows operating systems which is currently implemented through Microsoft Cluster Server (MSCS).

**Chapter 10, "Cluster support for the Solaris Operating Environment"**

Discusses DB2 support for high availability failover recovery in the Solaris Operating Environment, which is currently implemented through Sun Cluster 3.0 (SC3.0) or Veritas Cluster Server (VCS).

**Appendixes****Appendix A, "How to read the syntax diagrams"**

Explains the conventions used in syntax diagrams.

**Appendix B, "Warning, error and completion messages"**

Provides information about interpreting messages generated by the database manager when a warning or error condition has been detected.

**Appendix C, "Additional DB2 commands"**

Describes recovery-related DB2 commands.

**Appendix D, "Additional APIs and associated data structures"**

Describes recovery-related APIs and their data structures.

**Appendix E, "Recovery sample programs"**

Provides the code listing for a sample program containing recovery-related DB2 APIs and embedded SQL calls, and information on how to use them.

**Appendix F, "Tivoli Storage Manager"**

Provides information about the Tivoli Storage Manager (TSM) product, which you can use to manage database or table space backup operations.

**Appendix G, "User exit for database recovery"**

Discusses how user exit programs can be used with database log files, and describes some sample user exit programs.

**Appendix H, "Backup and restore APIs for vendor products"**

Describes the function and use of APIs that enable DB2 to interface with other vendor software.

---

## Part 1. Data recovery



---

## Chapter 1. Developing a good backup and recovery strategy

This section discusses factors to consider when choosing database and table space recovery methods, including backing up and restoring a database or table space, and using rollforward recovery.

The following topics are covered:

- “Developing a backup and recovery strategy”
- “Deciding how often to back up” on page 6
- “Storage considerations for recovery” on page 8
- “Keeping related data together” on page 9
- “Using different operating systems” on page 9
- “Crash recovery” on page 9
- “Disaster recovery” on page 21
- “Version recovery” on page 22
- “Rollforward recovery” on page 22
- “Incremental backup and recovery” on page 25
- “Understanding recovery logs” on page 31
- “Understanding the recovery history file” on page 51
- “Understanding table space states” on page 55
- “Enhancing recovery performance” on page 56

---

### Developing a backup and recovery strategy

A database can become unusable because of hardware or software failure, or both. You may, at one time or another, encounter storage problems, power interruptions, or application failures, and each failure scenario requires a different recovery action. Protect your data against the possibility of loss by having a well rehearsed recovery strategy in place. Some of the questions that you should answer when developing your recovery strategy are:

- Will the database be recoverable?
- How much time can be spent recovering the database?
- How much time will pass between backup operations?
- How much storage space can be allocated for backup copies and archived logs?
- Will table space level backups be sufficient, or will full database backups be necessary?
- Should I configure a standby system, either manually or through high availability disaster recovery (HADR)?

A database recovery strategy should ensure that all information is available when it is required for database recovery. It should include a regular schedule for taking database backups and, in the case of partitioned database systems, include backups when the system is scaled (when database partition servers or nodes are added or dropped). Your overall strategy should also include procedures for recovering command scripts, applications, user-defined functions (UDFs), stored procedure code in operating system libraries, and load copies.

Different recovery methods are discussed in the sections that follow, and you will discover which recovery method is best suited to your business environment.

The concept of a database *backup* is the same as any other data backup: taking a copy of the data and then storing it on a different medium in case of failure or damage to the original. The simplest case of a backup involves shutting down the database to ensure that no further transactions occur, and then simply backing it up. You can then rebuild the database if it becomes damaged or corrupted in some way.

The rebuilding of the database is called *recovery*. *Version recovery* is the restoration of a previous version of the database, using an image that was created during a backup operation. *Rollforward recovery* is the reapplication of transactions recorded in the database log files after a database or a table space backup image has been restored.

*Crash recovery* is the automatic recovery of the database if a failure occurs before all of the changes that are part of one or more units of work (transactions) are completed and committed. This is done by rolling back incomplete transactions and completing committed transactions that were still in memory when the crash occurred.

Recovery log files and the recovery history file are created automatically when a database is created (Figure 1 on page 5). These log files are important if you need to recover data that is lost or damaged.

Each database includes *recovery logs*, which are used to recover from application or system errors. In combination with the database backups, they are used to recover the consistency of the database right up to the point in time when the error occurred.

The *recovery history file* contains a summary of the backup information that can be used to determine recovery options, if all or part of the database must be recovered to a given point in time. It is used to track recovery-related events such as backup and restore operations, among others. This file is located in the database directory.

The *table space change history file*, which is also located in the database directory, contains information that can be used to determine which log files are required for the recovery of a particular table space.

You cannot directly modify the recovery history file or the table space change history file; however, you can delete entries from the files using the the PRUNE HISTORY command. You can also use the *rec\_his\_retentn* database configuration parameter to specify the number of days that these history files will be retained.

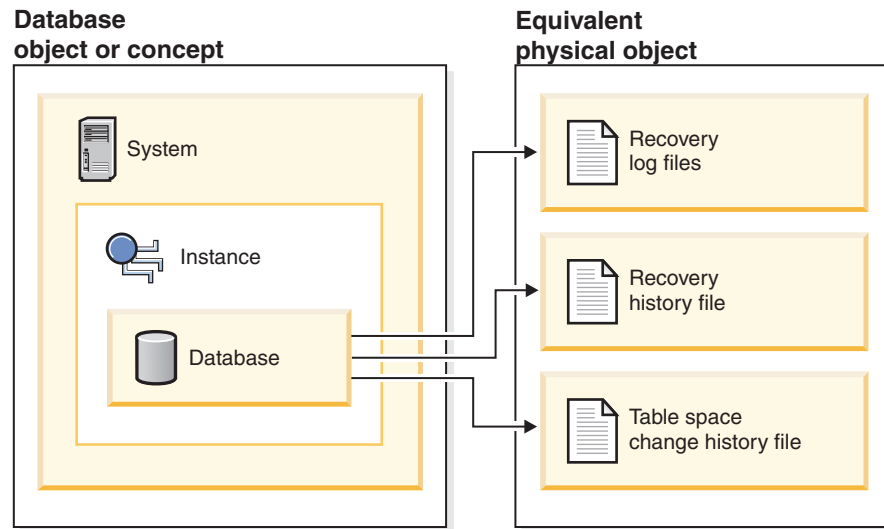


Figure 1. Database recovery files

Data that is easily recreated can be stored in a non-recoverable database. This includes data from an outside source that is used for read-only applications, and tables that are not often updated, for which the small amount of logging does not justify the added complexity of managing log files and rolling forward after a restore operation. *Non-recoverable databases* have the `logarchmeth1` and `logarchmeth2` database configuration parameters set to "OFF". This means that the only logs that are kept are those required for crash recovery. These logs are known as *active logs*, and they contain current transaction data. Version recovery using *offline* backups is the primary means of recovery for a non-recoverable database. (An offline backup means that no other application can use the database when the backup operation is in progress.) Such a database can only be restored offline. It is restored to the state it was in when the backup image was taken and rollforward recovery is not supported.

Data that *cannot* be easily recreated should be stored in a recoverable database. This includes data whose source is destroyed after the data is loaded, data that is manually entered into tables, and data that is modified by application programs or users after it is loaded into the database. *Recoverable databases* have the `logarchmeth1` or `logarchmeth2` database configuration parameters set to a value other than "OFF". Active logs are still available for crash recovery, but you also have the *archived logs*, which contain committed transaction data. Such a database can only be restored offline. It is restored to the state it was in when the backup image was taken. However, with rollforward recovery, you can roll the database *forward* (that is, past the time when the backup image was taken) by using the active and archived logs to either a specific point in time, or to the end of the active logs.

Recoverable database backup operations can be performed either offline or *online* (online meaning that other applications can connect to the database during the backup operation). Online table space restore and rollforward operations are supported only if the database is recoverable. If the database is non-recoverable, database restore and rollforward operations must be performed offline. During an online backup operation, rollforward recovery ensures that *all* table changes are captured and reapplied if that backup is restored.

If you have a recoverable database, you can back up, restore, and roll individual table spaces forward, rather than the entire database. When you back up a table

space online, it is still available for use, and simultaneous updates are recorded in the logs. When you perform an online restore or rollforward operation on a table space, the table space itself is not available for use until the operation completes, but users are not prevented from accessing tables in other table spaces.

## Automated backup operations

Since it can be time-consuming to determine whether and when to run maintenance activities such as backup operations, you can use the Configure Automatic Maintenance wizard to do this for you. With automatic maintenance, you specify your maintenance objectives, including when automatic maintenance can run. DB2 then uses these objectives to determine if the maintenance activities need to be done and then runs only the required maintenance activities during the next available maintenance window (a user-defined time period for the running of automatic maintenance activities).

**Note:** You can still perform manual backup operations when automatic maintenance is configured. DB2 will only perform automatic backup operations if they are required.

### Related concepts:

- “Crash recovery” on page 9
- “Version recovery” on page 22
- “Rollforward recovery” on page 22
- “High availability disaster recovery overview” on page 173
- “Data Links server file backups” in the *DB2 Data Links Manager Administration Guide and Reference*
- “Failure and recovery overview” in the *DB2 Data Links Manager Administration Guide and Reference*

### Related reference:

- “rec\_his\_retentn - Recovery history retention period configuration parameter” in the *Administration Guide: Performance*
- “logarchmeth1 - Primary log archive method configuration parameter” in the *Administration Guide: Performance*
- “DB2 Data Links Manager system setup and backup recommendations” in the *DB2 Data Links Manager Administration Guide and Reference*

---

## Deciding how often to back up

Your recovery plan should allow for regularly scheduled backup operations, because backing up a database requires time and system resources. Your plan may include a combination of full database backups and incremental backup operations.

You should take full database backups regularly, even if you archive the logs (which allows for rollforward recovery). To recover a database, you need a full database backup image that contains all of the table space backup images. Table space backup images are useful for recovering from an isolated disk failure or an application error. In partitioned database environments, you only need to restore the table spaces that have failed. You do not need to restore all of the table spaces or all of the partitions.

You should also consider not overwriting backup images and logs, saving at least two full database backup images and their associated logs as an extra precaution.



If the amount of time needed to apply archived logs when recovering and rolling a very active database forward is a major concern, consider the cost of backing up the database more frequently. This reduces the number of archived logs you need to apply when rolling forward.

You can initiate a backup operation while the database is either *online* or *offline*. If it is online, other applications or processes can connect to the database, as well as read and modify data while the backup operation is running. If the backup operation is running offline, other applications *cannot* connect to the database.

To reduce the amount of time that the database is not available, consider using online backup operations. Online backup operations are supported only if rollforward recovery is enabled. If rollforward recovery is enabled and you have a complete set of recovery logs, you can rebuild the database, should the need arise. You can only use an online backup image for recovery if you have the logs that span the time during which the backup operation was running.

Offline backup operations are faster than online backup operations, since there is no contention for the data files.

The backup utility lets you back up selected table spaces. If you use DMS table spaces, you can store different types of data in their own table spaces to reduce the time required for backup operations. You can keep table data in one table space, long field and LOB data in another table space, and indexes in yet another table space. If you do this and a disk failure occurs, it is likely to affect only one of the table spaces. Restoring or rolling forward one of these table spaces will take less time than it would have taken to restore a single table space containing all of the data.

You can also save time by taking backups of different table spaces at different times, as long as the changes to them are not the same. So, if long field or LOB data is not changed as frequently as the other data, you can back up these table spaces less frequently. If long field and LOB data are not required for recovery, you can also consider not backing up the table space that contains that data. If the LOB data can be reproduced from a separate source, choose the NOT LOGGED option when creating or altering a table to include LOB columns.

**Note:** Consider the following if you keep your long field data, LOB data, and indexes in separate table spaces, but do not back them up together: If you back up a table space that does not contain all of the table data, you cannot perform point-in-time rollforward recovery on that table space. All the table spaces that contain any type of data for a table must be rolled forward simultaneously to the same point in time.

If you reorganize a table, you should back up the affected table spaces after the operation completes. If you have to restore the table spaces, you will not have to roll forward through the data reorganization.

The time required to recover a database is made up of two parts: the time required to complete the restoration of the backup; and, if the database is enabled for forward recovery, the time required to apply the logs during the rollforward operation. When formulating a recovery plan, you should take these recovery costs and their impact on your business operations into account. Testing your overall recovery plan will assist you in determining whether the time required to recover the database is reasonable given your business requirements. Following each test, you may want to increase the frequency with which you take a backup. If

rollforward recovery is part of your strategy, this will reduce the number of logs that are archived between backups and, as a result, reduce the time required to roll the database forward after a restore operation.

**Related concepts:**

- “Developing a backup and recovery strategy” on page 3
- “Incremental backup and recovery” on page 25

**Related reference:**

- Appendix G, “User exit for database recovery,” on page 329
- “Configuration parameters for database logging” on page 35

---

## Storage considerations for recovery

When deciding which recovery method to use, consider the storage space required.

The version recovery method requires space to hold the backup copy of the database and the restored database. The rollforward recovery method requires space to hold the backup copy of the database or table spaces, the restored database, and the archived database logs.

If a table contains long field or large object (LOB) columns, you should consider placing this data into a separate table space. This will affect your storage space considerations, as well as affect your plan for recovery. With a separate table space for long field and LOB data, and knowing the time required to back up long field and LOB data, you may decide to use a recovery plan that only occasionally saves a backup of this table space. You may also choose, when creating or altering a table to include LOB columns, not to log changes to those columns. This will reduce the size of the required log space and the corresponding log archive space.

To prevent media failure from destroying a database and your ability to rebuild it, keep the database backup, the database logs, and the database itself on different devices. For this reason, it is highly recommended that you use the *newlogpath* configuration parameter to put database logs on a separate device once the database is created.

The database logs can use up a large amount of storage. If you plan to use the rollforward recovery method, you must decide how to manage the archived logs. Your choices are the following:

- Specify a log archiving method using the LOGARCHMETH1 or LOGARCHMETH2 configuration parameters.
- Manually copy the logs to a storage device or directory other than the database log path directory after they are no longer in the active set of logs.
- Use a user exit program to copy these logs to another storage device in your environment.

**Related concepts:**

- “Managing log files through log archiving” on page 45

**Related reference:**

- “Configuration parameters for database logging” on page 35
- “logarchmeth1 - Primary log archive method configuration parameter” in the *Administration Guide: Performance*

- “logarchmeth2 - Secondary log archive method configuration parameter” in the *Administration Guide: Performance*

---

## Keeping related data together

As part of your database design, you will know the relationships that exist between tables. These relationships can be expressed at the application level, when transactions update more than one table, or at the database level, where referential integrity exists between tables, or where triggers on one table affect another table. You should consider these relationships when developing a recovery plan. You will want to back up related sets of data together. Such sets can be established at either the table space or the database level. By keeping related sets of data together, you can recover to a point where all of the data is consistent. This is especially important if you want to be able to perform point-in-time rollforward recovery on table spaces.

---

## Using different operating systems

DB2® supports cross-platform backup and restore operations. You can restore databases created on a DB2 Version 8 32-bit Windows® platform to a DB2 Version 8 64-bit Windows platform, or the reverse. You can restore databases created on a DB2 Version 8 32-bit Linux (Intel) platform to a DB2 Version 8 64-bit Linux (Intel) platform, or the reverse. You can restore databases created on DB2 Version 8 AIX®, HP-UX, or the Solaris Operating Environment platforms, in 32-bit or 64-bit, to DB2 Version 8 AIX, HP-UX, or Solaris Operating Environment platforms (32-bit or 64-bit).

You can also restore backup images that were produced on a previous version of DB2 (up to two versions earlier) as long as the word size (32-bit or 64-bit) is the same. Cross-platform restore operations from a backup image created with a previous version of DB2 are not supported. The target system must have the same (or later) version of DB2 as the source system. Restore operations to a down-level system are not supported.

If you must move tables from one operating system to another, and cross-platform backup and restore support is not available in your environment, you can use the **db2move** command, or the export utility followed by the import or the load utility.

### Related reference:

- “db2move - Database Movement Tool Command” in the *Command Reference*
- “EXPORT Command” in the *Command Reference*
- “IMPORT Command” in the *Command Reference*
- “LOAD Command” in the *Command Reference*

---

## Crash recovery

Transactions (or units of work) against a database can be interrupted unexpectedly. If a failure occurs before all of the changes that are part of the unit of work are completed and committed, the database is left in an inconsistent and unusable state. *Crash recovery* is the process by which the database is moved back to a consistent and usable state. This is done by rolling back incomplete transactions and completing committed transactions that were still in memory when the crash

occurred (Figure 2). When a database is in a consistent and usable state, it has attained what is known as a "point of consistency".

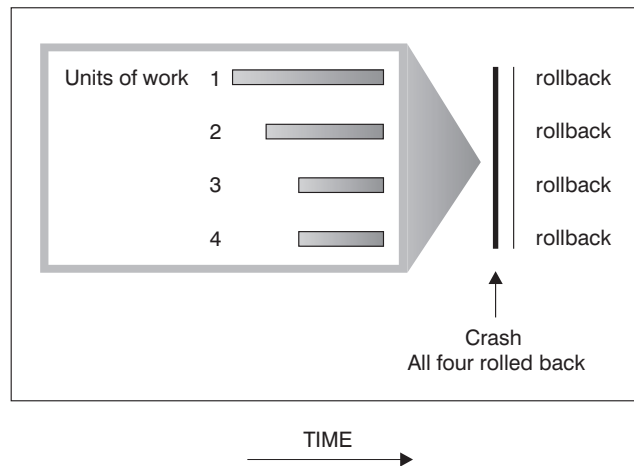


Figure 2. Rolling Back Units of Work (Crash Recovery)

A *transaction failure* results from a severe error or condition that causes the database or the database manager to end abnormally. Partially completed units of work, or UOW that have not been flushed to disk at the time of failure, leave the database in an inconsistent state. Following a transaction failure, the database must be recovered. Conditions that can result in transaction failure include:

- A power failure on the machine, causing the database manager and the database partitions on it to go down
- A hardware failure such as memory corruption, or disk, CPU, or network failure.
- A serious operating system error that causes DB2® to go down

If you want the rollback of incomplete units of work to be done automatically by the database manager, enable the automatic restart (*autorestart*) database configuration parameter by setting it to ON. (This is the default value.) If you do not want automatic restart behavior, set the *autorestart* database configuration parameter to OFF. As a result, you will need to issue the RESTART DATABASE command when a database failure occurs. If the database I/O was suspended before the crash occurred, you must specify the WRITE RESUME option of the RESTART DATABASE command in order for the crash recovery to continue. The administration notification log records when the database restart operation begins.

If crash recovery is applied to a database that is enabled for forward recovery (that is, the *logarchmeth1* configuration parameter is not set to OFF), and an error occurs during crash recovery that is attributable to an individual table space, that table space will be taken offline, and cannot be accessed until it is repaired. Crash recovery continues. At the completion of crash recovery, the other table spaces in the database will be accessible, and connections to the database can be established. However, if the table space that is taken offline is the table space that contains the system catalogs, it must be repaired before any connections will be permitted.

**Related reference:**

- "autorestart - Auto restart enable configuration parameter" in the *Administration Guide: Performance*

- “logarchmeth1 - Primary log archive method configuration parameter” in the *Administration Guide: Performance*

---

## Crash recovery - details

### Recovering damaged table spaces

A damaged table space has one or more containers that cannot be accessed. This is often caused by media problems that are either permanent (for example, a bad disk), or temporary (for example, an offline disk, or an unmounted file system).

If the damaged table space is the system catalog table space, the database cannot be restarted. If the container problems cannot be fixed leaving the original data intact, the only available options are:

- To restore the database
- To restore the catalog table space.

**Notes:**

1. Table space restore is only valid for recoverable databases, because the database must be rolled forward.
2. If you restore the catalog table space, you must perform a rollforward operation to the end of logs.
- 3.

If the damaged table space is *not* the system catalog table space, DB2® attempts to make as much of the database available as possible.

If the damaged table space is the only temporary table space, you should create a new temporary table space as soon as a connection to the database can be made. Once created, the new temporary table space can be used, and normal database operations requiring a temporary table space can resume. You can, if you wish, drop the offline temporary table space. There are special considerations for table reorganization using a system temporary table space:

- If the database or the database manager configuration parameter *indexrec* is set to RESTART, all invalid indexes must be rebuilt during database activation; this includes indexes from a reorganization that crashed during the build phase.
- If there are incomplete reorganization requests in a damaged temporary table space, you may have to set the *indexrec* configuration parameter to ACCESS to avoid restart failures.

**Related tasks:**

- “Recovering table spaces in recoverable databases” on page 12
- “Recovering table spaces in non-recoverable databases” on page 12

**Related reference:**

- “RESTART DATABASE Command” in the *Command Reference*
- “RESTORE DATABASE” on page 87

## Recovering table spaces in recoverable databases

When crash recovery is necessary, a damaged table space will be taken offline and will not be accessible. It will be placed in roll forward pending state. A restart operation will succeed if there are no additional problems.

### Procedure:

To make the damaged table space useable, use one of the procedures that follow:

- Method 1

1. Fix the damaged containers without losing the original data.
2. Complete a table space rollforward operation to the end of the logs.

**Note:** The rollforward operation will first attempt to bring the table space from offline to normal state.

- Method 2

1. Fix the damaged containers with or without losing the original data.
2. Perform a table space restore operation.
3. Complete a table space rollforward operation to the end of the logs or to a point-in-time.

### Related concepts:

- “Recovering damaged table spaces” on page 11

### Related tasks:

- “Recovering table spaces in non-recoverable databases” on page 12

### Related reference:

- “RESTART DATABASE Command” in the *Command Reference*
- “RESTORE DATABASE” on page 87

## Recovering table spaces in non-recoverable databases

Since crash recovery is necessary, and logs are not kept indefinitely, a restart operation can only succeed if you are willing to drop the damaged table spaces. Successful completion of a recovery operation means that the log records necessary to recover the damaged table spaces to a consistent state will be gone; therefore, the only valid action against such table spaces is to drop them.

### Procedure:

To restart a database with damaged table spaces:

1. Invoke an unqualified restart database operation. It will succeed if there are no damaged table spaces. If it fails (SQL0290N), look in the administration notification log file for a complete list of table spaces that are currently damaged.
2. If you are willing to drop all of the damaged table spaces, initiate another restart database operation, listing all of the damaged table spaces under the DROP PENDING TABLESPACES option. If a damaged table space is included in the DROP PENDING TABLESPACES list, the table space is put into drop pending state, and you must drop the table space after the recovery operation is complete.

The restart operation continues without recovering the specified table spaces. If a damaged table space is *not* included in the DROP PENDING TABLESPACES list, the restart database operation fails with SQL0290N.

**Note:** Including a table space name in the DROP PENDING TABLESPACES list does not mean that the table space will be in drop pending state. It will be placed in this state only if the table space is found to be damaged during the restart operation.

3. If the restart database operation is successful, invoke the LIST TABLESPACES command to find out which table spaces are in drop pending state.
4. Issue DROP TABLESPACE statements to drop each of the table spaces that are in drop pending state. Once you have done this, you will be able to reclaim the space that the damaged table spaces were using or recreate the table spaces.
5. If you are unwilling to drop and lose the data in the damaged table spaces, you can:
  - Fix the damaged containers (without losing the original data).
  - Reissue the RESTART DATABASE command.
  - Perform a database restore operation.

**Related concepts:**

- “Recovering damaged table spaces” on page 11

**Related tasks:**

- “Recovering table spaces in recoverable databases” on page 12

**Related reference:**

- “RESTART DATABASE Command” in the *Command Reference*
- “RESTORE DATABASE” on page 87

## Reducing the impact of media failure

To reduce the probability of media failure, and to simplify recovery from this type of failure:

- Mirror or duplicate the disks that hold the data and logs for important databases.
- Use a Redundant Array of Independent Disks (RAID) configuration, such as RAID Level 5.
- In a partitioned database environment, set up a rigorous procedure for handling the data and the logs on the catalog node. Because this node is critical for maintaining the database:
  - Ensure that it resides on a reliable disk
  - Duplicate it
  - Make frequent backups
  - Do not put user data on it.

### Protecting against disk failure

If you are concerned about the possibility of damaged data or logs due to a disk crash, consider the use of some form of disk fault tolerance. Generally, this is accomplished through the use of a *disk array*, which is a set of disks.

A disk array is sometimes referred to simply as a RAID (Redundant Array of Independent Disks). Disk arrays can also be provided through software at the operating system or application level. The point of distinction between hardware and software disk arrays is how CPU processing of input/output (I/O) requests is handled. For hardware disk arrays, I/O activity is managed by disk controllers; for software disk arrays, this is done by the operating system or an application.

**Hardware disk arrays:** In a hardware disk array, multiple disks are used and managed by a disk controller, complete with its own CPU. All of the logic required to manage the disks forming this array is contained on the disk controller; therefore, this implementation is operating system-independent.

There are several types of RAID architecture, differing in function and performance, but only RAID level 1 and level 5 are commonly used today.

RAID level 1 is also known as disk mirroring or duplexing. *Disk mirroring* copies data (a complete file) from one disk to a second disk, using a single disk controller. *Disk duplexing* is similar to disk mirroring, except that disks are attached to a second disk controller (like two SCSI adapters). Data protection is good: Either disk can fail, and data is still accessible from the other disk. With disk duplexing, a disk controller can also fail without compromising data protection. Performance is good, but this implementation requires twice the usual number of disks.

RAID level 5 involves data and parity striping by sectors, across all disks. Parity is interleaved with data, rather than being stored on a dedicated drive. Data protection is good: If any disk fails, the data can still be accessed by using information from the other disks, along with the striped parity information. Read performance is good, but write performance is not. A RAID level 5 configuration requires a minimum of three identical disks. The amount of disk space required for overhead varies with the number of disks in the array. In the case of a RAID level 5 configuration with 5 disks, the space overhead is 20 percent.

When using a RAID (but not a RAID level 0) disk array, a failed disk will not prevent you from accessing data on the array. When hot-pluggable or hot-swappable disks are used in the array, a replacement disk can be swapped with the failed disk while the array is in use. With RAID level 5, if two disks fail at the same time, all data is lost (but the probability of simultaneous disk failures is very small).

You might consider using a RAID level 1 hardware disk array or a software disk array for your logs, because this provides recoverability to the point of failure, and offers good write performance, which is important for logs. To do this, use the *mirrorlogpath* configuration parameter to specify a mirror log path on a RAID level 1 file system. In cases where reliability is critical (because time cannot be lost recovering data following a disk failure), and write performance is not so critical, consider using a RAID level 5 hardware disk array. Alternatively, if write performance is critical, and the cost of additional disk space is not significant, consider a RAID level 1 hardware disk array for your data, as well as for your logs.

For detailed information about the available RAID levels, visit the following web site: [http://www.acnc.com/04\\_01\\_00.html](http://www.acnc.com/04_01_00.html)

**Software disk arrays:** A software disk array accomplishes much the same as does a hardware disk array, but disk traffic is managed by either the operating system, or by an application program running on the server. Like other programs, the



software array must compete for CPU and system resources. This is not a good option for a CPU-constrained system, and it should be remembered that overall disk array performance is dependent on the server's CPU load and capacity.

A typical software disk array provides disk mirroring. Although redundant disks are required, a software disk array is comparatively inexpensive to implement, because costly disk controllers are not required.

**CAUTION:**

**Having the operating system boot drive in the disk array prevents your system from starting if that drive fails. If the drive fails before the disk array is running, the disk array cannot allow access to the drive. A boot drive should be separate from the disk array.**

## Reducing the impact of transaction failure

To reduce the impact of a transaction failure, try to ensure:

- An uninterrupted power supply for each DB2® server
- Adequate disk space for database logs on all partitions
- Reliable communication links among the database partition servers in a partitioned database environment
- Synchronization of the system clocks in a partitioned database environment.

**Related concepts:**

- "Synchronizing clocks in a partitioned database system" on page 121

## Recovering from transaction failures in a partitioned database environment

If a transaction failure occurs in a partitioned database environment, database recovery is usually necessary on both the failed database partition server and any other database partition server that was participating in the transaction:

- Crash recovery occurs on the failed database partition server after the antecedent condition is corrected.
- *Database partition failure recovery* on the other (still active) database partition servers occurs immediately after the failure has been detected.

In a partitioned database environment, the database partition server on which an application is submitted is the coordinator node, and the first agent that works for the application is the coordinator agent. The coordinator agent is responsible for distributing work to other database partition servers, and it keeps track of which ones are involved in the transaction. When the application issues a COMMIT statement for a transaction, the coordinator agent commits the transaction by using the two-phase commit protocol. During the first phase, the coordinator node distributes a PREPARE request to all the other database partition servers that are participating in the transaction. These servers then respond with one of the following:

<b>READ-ONLY</b>	No data change occurred at this server
<b>YES</b>	Data change occurred at this server
<b>NO</b>	Because of an error, the server is not prepared to commit

If one of the servers responds with a NO, the transaction is rolled back. Otherwise, the coordinator node begins the second phase.

During the second phase, the coordinator node writes a COMMIT log record, then distributes a COMMIT request to all the servers that responded with a YES. After all the other database partition servers have committed, they send an acknowledgment of the COMMIT to the coordinator node. The transaction is complete when the coordinator agent has received all COMMIT acknowledgments from all the participating servers. At this point, the coordinator agent writes a FORGET log record.

### **Transaction failure recovery on an active database partition server**

If any database partition server detects that another server is down, all work that is associated with the failed database partition server is stopped:

- If the still active database partition server is the coordinator node for an application, and the application was running on the failed database partition server (and not ready to COMMIT), the coordinator agent is interrupted to do failure recovery. If the coordinator agent is in the second phase of COMMIT processing, SQL0279N is returned to the application, which in turn loses its database connection. Otherwise, the coordinator agent distributes a ROLLBACK request to all other servers participating in the transaction, and SQL1229N is returned to the application.
- If the failed database partition server was the coordinator node for the application, agents that are still working for the application on the active servers are interrupted to do failure recovery. The current transaction is rolled back locally on each server, unless it has been prepared and is waiting for the transaction outcome. In this situation, the transaction is left in doubt on the active database partition servers, and the coordinator node is not aware of this (because it is not available).
- If the application connected to the failed database partition server (before it failed), but neither the local database partition server nor the failed database partition server is the coordinator node, agents working for this application are interrupted. The coordinator node will either send a ROLLBACK or a disconnect message to the other database partition servers. The transaction will only be indoubt on database partition servers that are still active if the coordinator node returns SQL0279.

Any process (such as an agent or deadlock detector) that attempts to send a request to the failed server is informed that it cannot send the request.

### **Transaction failure recovery on the failed database partition server**

If the transaction failure causes the database manager to end abnormally, you can issue the **db2start** command with the RESTART option to restart the database manager once the database partition has been restarted. If you cannot restart the database partition, you can issue **db2start** to restart the database manager on a different partition.

If the database manager ends abnormally, database partitions on the server may be left in an inconsistent state. To make them usable, crash recovery can be triggered on a database partition server:

- Explicitly, through the RESTART DATABASE command
- Implicitly, through a CONNECT request when the *autorestart* database configuration parameter has been set to ON

Crash recovery reapplies the log records in the active log files to ensure that the effects of all complete transactions are in the database. After the changes have been reapplied, all uncommitted transactions are rolled back locally, *except* for indoubt transactions. There are two types of indoubt transaction in a partitioned database environment:

- On a database partition server that is not the coordinator node, a transaction is in doubt if it is prepared but not yet committed.
- On the coordinator node, a transaction is in doubt if it is committed but not yet logged as complete (that is, the FORGET record is not yet written). This situation occurs when the coordinator agent has not received all the COMMIT acknowledgments from all the servers that worked for the application.

Crash recovery attempts to resolve all the indoubt transactions by doing one of the following. The action that is taken depends on whether the database partition server was the coordinator node for an application:

- If the server that restarted is not the coordinator node for the application, it sends a query message to the coordinator agent to discover the outcome of the transaction.
- If the server that restarted *is* the coordinator node for the application, it sends a message to all the other agents (subordinate agents) that the coordinator agent is still waiting for COMMIT acknowledgments.

It is possible that crash recovery may not be able to resolve all the indoubt transactions (for example, some of the database partition servers may not be available). In this situation, the SQL warning message SQL1061W is returned. Because indoubt transactions hold resources, such as locks and active log space, it is possible to get to a point where no changes can be made to the database because the active log space is being held up by indoubt transactions. For this reason, you should determine whether indoubt transactions remain after crash recovery, and recover all database partition servers that are required to resolve the indoubt transactions as quickly as possible.

If one or more servers that are required to resolve an indoubt transaction cannot be recovered in time, and access is required to database partitions on other servers, you can manually resolve the indoubt transaction by making an heuristic decision. You can use the LIST INDOUBT TRANSACTIONS command to query, commit, and roll back the indoubt transaction on the server.

**Note:** The LIST INDOUBT TRANSACTIONS command is also used in a distributed transaction environment. To distinguish between the two types of indoubt transactions, the *originator* field in the output that is returned by the LIST INDOUBT TRANSACTIONS command displays one of the following:

- DB2<sup>®</sup> Enterprise Server Edition, which indicates that the transaction originated in a partitioned database environment.
- XA, which indicates that the transaction originated in a distributed environment.

### **Identifying the failed database partition server**

When a database partition server fails, the application will typically receive one of the following SQLCODEs. The method for detecting which database manager failed depends on the SQLCODE received:

**SQL0279N**

This SQLCODE is received when a database partition server involved in a transaction is terminated during COMMIT processing.

**SQL1224N**

This SQLCODE is received when the database partition server that failed is the coordinator node for the transaction.

**SQL1229N**

This SQLCODE is received when the database partition server that failed is not the coordinator node for the transaction.

Determining which database partition server failed is a two-step process. The SQLCA associated with SQLCODE SQL1229N contains the node number of the server that detected the error in the sixth array position of the *sqlerrd* field. (The node number that is written for the server corresponds to the node number in the *db2nodes.cfg* file.) On the database partition server that detects the error, a message that indicates the node number of the failed server is written to the administration notification log.

**Note:** If multiple logical nodes are being used on a processor, the failure of one logical node may cause other logical nodes on the same processor to fail.

**Related concepts:**

- “Two-phase commit” in the *Administration Guide: Planning*
- “Error recovery during two-phase commit” in the *Administration Guide: Planning*

**Related tasks:**

- “Manually resolving indoubt transactions” in the *Administration Guide: Planning*

**Related reference:**

- “db2start - Start DB2 Command” in the *Command Reference*
- “LIST INDOUBT TRANSACTIONS Command” in the *Command Reference*

## Recovering from the failure of a database partition server

**Procedure:**

To recover from the failure of a database partition server:

1. Correct the problem that caused the failure.
2. Restart the database manager by issuing the **db2start** command from any database partition server.
3. Restart the database by issuing the **RESTART DATABASE** command on the failed database partition server or servers.

**Related concepts:**

- “Recovering from transaction failures in a partitioned database environment” on page 15

**Related reference:**

- “db2start - Start DB2 Command” in the *Command Reference*
- “RESTART DATABASE Command” in the *Command Reference*

## Recovering indoubt transactions on the host when DB2 Connect has the DB2 Syncpoint Manager configured

If your application has accessed a host or iSeries database server during a transaction, there are some differences in how indoubt transactions are recovered.

To access host or iSeries database servers, DB2 Connect is used. The recovery steps differ if DB2 Connect has the DB2 Syncpoint Manager configured.

### Procedures:

The recovery of indoubt transactions at host or iSeries servers is normally performed automatically by the Transaction Manager (TM) and the DB2 Syncpoint Manager (SPM). An indoubt transaction at a host or iSeries server does not hold any resources at the local DB2 location, but does hold resources at the host or iSeries server as long as the transaction is indoubt at that location. If the administrator of the host or iSeries server determines that a heuristic decision must be made, then the administrator may contact the local DB2 database administrator (for example via telephone) to determine whether to commit or roll back the transaction at the host or iSeries server. If this occurs, the LIST DRDA INDOUBT TRANSACTIONS command can be used to determine the state of the transaction at the DB2 Connect instance. The following steps can be used as a guideline for most situations involving an SNA communications environment.

1. Connect to the SPM as shown below:

```
db2 => connect to db2spm
```

```
Database Connection Information
```

```
Database product      = SPM0500
SQL authorization ID  = CRUS
Local database alias  = DB2SPM
```

2. Issue the LIST DRDA INDOUBT TRANSACTIONS command to display the indoubt transactions known to the SPM. The example below shows one indoubt transaction known to the SPM. The db\_name is the local alias for the host or iSeries server. The partner\_lu is the fully qualified luname of the host or iSeries server. This provides the best identification of the host or iSeries server, and should be provided by the caller from the host or iSeries server. The luwid provides a unique identifier for a transaction and is available at all hosts and iSeries servers. If the transaction in question is displayed, then the uow\_status field can be used to determine the outcome of the transaction if the value is C (commit) or R (rollback). If you issue the LIST DRDA INDOUBT TRANSACTIONS command with the WITH PROMPTING parameter, you can commit, roll back, or forget the transaction interactively.

```
db2 => list drda indoubt transactions
DRDA Indoubt Transactions:
1.db_name: DBAS3   db_alias: DBAS3   role: AR
   uow_status: C  partner_status: I  partner_lu: USIBMSY.SY12DQA
corr_tok: USIBMST.STB3327L
luwid: USIBMST.STB3327.305DFDA5DC00.0001
xid: 53514C2000000017 00000000544D4442 0000000000305DFD A63055E962000000
    00035F
```

3. If an indoubt transaction for the partner\_lu and for the luwid is not displayed, or if the LIST DRDA INDOUBT TRANSACTIONS command returns as follows:

```
db2 => list drda indoubt transactions
SQL1251W No data returned for heuristic query.
```

then the transaction was rolled back.

There is another unlikely but possible situation that may occur. If an indoubt transaction with the proper luwid for the partner\_lu is displayed, but the uow\_status is "I", the SPM doesn't know whether the transaction is to be committed or rolled back. In this situation, you should use the WITH PROMPTING parameter to either commit or roll back the transaction on the DB2 Connect workstation. Then allow DB2 Connect to resynchronize with the host or iSeries server based on the heuristic decision.

**Related tasks:**

- "Recovering indoubt transactions on the host when DB2 Connect does not use the DB2 Syncpoint Manager" on page 20

**Related reference:**

- "db2start - Start DB2 Command" in the *Command Reference*
- "LIST INDOUBT TRANSACTIONS Command" in the *Command Reference*
- "RESTART DATABASE Command" in the *Command Reference*

## Recovering indoubt transactions on the host when DB2 Connect does not use the DB2 Syncpoint Manager

If your application has accessed a host or iSeries database server during a transaction, there are some differences in how indoubt transactions are recovered.

To access host or iSeries database servers, DB2 Connect is used. The recovery steps differ if DB2 Connect has the DB2 Syncpoint Manager configured.

**Procedure:**

Use the information in this section when TCP/IP connectivity is used to update DB2 for z/OS in a multisite update from either DB2 Connect Personal Edition or DB2 Connect Enterprise Server Edition, and the DB2 Syncpoint Manager is not used. The recovery of indoubt transactions in this situation differs from that for indoubt transactions involving the DB2 Syncpoint Manager. When an indoubt transaction occurs in this environment, an alert entry is generated at the client, at the database server, and (or) at the Transaction Manager (TM) database, depending on who detected the problem. The alert entry is placed in the db2alert.log file.

The resynchronization of any indoubt transactions occurs automatically as soon as the TM and the participating databases and their connections are all available again. You should allow automatic resynchronization to occur rather than heuristically force a decision at the database server. If, however, you must do this then use the following steps as a guideline.

**Note:** Because the DB2 Syncpoint Manager is not involved, you cannot use the LIST DRDA INDOUBT TRANSACTIONS command.

1. On the z/OS host, issue the command DISPLAY THREAD TYPE(INDOUBT). From this list identify the transaction that you want to heuristically complete. For details about the DISPLAY command, see the *DB2 for z/OS Command Reference*. The LUWID displayed can be matched to the same luwid at the Transaction Manager Database.
2. Issue the RECOVER THREAD(<LUWID>) ACTION(ABORT|COMMIT) command, depending on what you want to do.

For details about the RECOVER THREAD command, see the *DB2 for z/OS Command Reference*.

**Related tasks:**

- “Recovering indoubt transactions on the host when DB2 Connect has the DB2 Syncpoint Manager configured” on page 19

**Related reference:**

- “LIST INDOUBT TRANSACTIONS Command” in the *Command Reference*

---

## Disaster recovery

The term *disaster recovery* is used to describe the activities that need to be done to restore the database in the event of a fire, earthquake, vandalism, or other catastrophic events. A plan for disaster recovery can include one or more of the following:

- A site to be used in the event of an emergency
- A different machine on which to recover the database
- Off-site storage of database backups and archived logs.

If your plan for disaster recovery is to recover the entire database on another machine, you require at least one full database backup and all the archived logs for the database. You may choose to keep a standby database up to date by applying the logs to it as they are archived. Or, you may choose to keep the database backup and log archives in the standby site, and perform restore and rollforward operations only after a disaster has occurred. (In this case, a recent database backup is clearly desirable.) With a disaster, however, it is generally not possible to recover all of the transactions up to the time of the disaster.

The usefulness of a table space backup for disaster recovery depends on the scope of the failure. Typically, disaster recovery requires that you restore the entire database; therefore, a full database backup should be kept at a standby site. Even if you have a separate backup image of every table space, you cannot use them to recover the database. If the disaster is a damaged disk, a table space backup of each table space on that disk can be used to recover. If you have lost access to a container because of a disk failure (or for any other reason), you can restore the container to a different location.

Another way you can protect your data from partial or complete site failures is to implement the DB2® high availability disaster recovery (HADR) feature. Once it is set up, HADR protects against data loss by replicating data changes from a source database, called the primary, to a target database, called the standby.

DB2 provides you with several options when planning for disaster recovery. Based on your business needs, you might decide to use table space or full database backups as a safeguard against data loss, or you might decide that your environment is better suited to a solution like HADR. Whatever your choice, you should test your recovery procedures in a test environment before implementing them in your production environment.

**Related concepts:**

- “Redefining table space containers during a restore operation (redirected restore)” on page 85
- “High availability disaster recovery overview” on page 173

---

## Version recovery

*Version recovery* is the restoration of a previous version of the database, using an image that was created during a backup operation. You use this recovery method with non-recoverable databases (that is, databases for which you do not have archived logs). You can also use this method with recoverable databases by using the `WITHOUT ROLLING FORWARD` option on the `RESTORE DATABASE` command. A database restore operation will rebuild the entire database using a backup image created earlier. A database backup allows you to restore a database to a state identical to the one at the time that the backup was made. However, every unit of work from the time of the backup to the time of the failure is lost (see Figure 3).




Figure 3. *Version Recovery*. The database is restored from the latest backup image, but all units of work processed between the time of backup and failure are lost.

Using the version recovery method, you must schedule and perform full backups of the database on a regular basis.

In a partitioned database environment, the database is located across many database partition servers (or nodes). You must restore all partitions, and the backup images that you use for the restore database operation must all have been taken at the same time. (Each database partition is backed up and restored separately.) A backup of each database partition taken at the same time is known as a *version backup*.

---

## Rollforward recovery

To use the *rollforward recovery* method, you must have taken a backup of the database and archived the logs (by setting the `logarchmeth1` and `logarchmeth2` configuration parameters to a value other than OFF). Restoring the database and specifying the `WITHOUT ROLLING FORWARD` option is equivalent to using the version recovery method. The database is restored to a state identical to the one at the time that the offline backup image was made. If you restore the database and do *not* specify the `WITHOUT ROLLING FORWARD` option for the restore database operation, the database will be in rollforward pending state at the end of the restore operation. This allows rollforward recovery to take place.



**Note:** The WITHOUT ROLLING FORWARD option cannot be used if the database backup was taken online.

The two types of rollforward recovery to consider are:

- *Database rollforward recovery.* In this type of rollforward recovery, transactions recorded in database logs are applied following the database restore operation (see Figure 4). The database logs record all changes made to the database. This method completes the recovery of the database to its state at a particular point in time, or to its state immediately before the failure (that is, to the end of the active logs.)

In a partitioned database environment, the database is located across many database partitions, and the ROLLFORWARD DATABASE command must be issued on the partition where the catalog tables for the database resides (catalog partition). If you are performing point-in-time rollforward recovery, all database partitions must be rolled forward to ensure that all partitions are at the same level. If you need to restore a single database partition, you can perform rollforward recovery to the end of the logs to bring it up to the same level as the other partitions in the database. Only recovery to the end of the logs can be used if one database partition is being rolled forward. Point-in-time recovery applies to *all* database partitions.

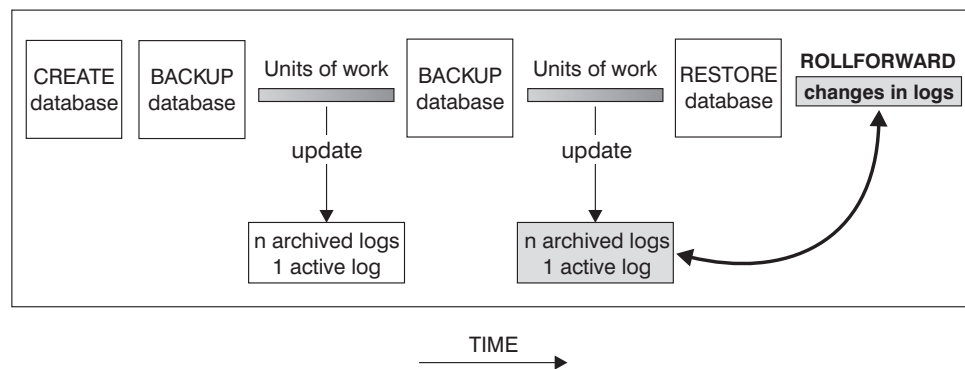


Figure 4. Database Rollforward Recovery. There can be more than one active log in the case of a long-running transaction.

- *Table space rollforward recovery.* If the database is enabled for forward recovery, it is also possible to back up, restore, and roll table spaces forward (see Figure 5 on page 24). To perform a table space restore and rollforward operation, you need a backup image of either the entire database (that is, all of the table spaces), or one or more individual table spaces. You also need the log records that affect the table spaces that are to be recovered. You can roll forward through the logs to one of two points:
  - The end of the logs; or,
  - A particular point in time (called *point-in-time* recovery).

Table space rollforward recovery can be used in the following two situations:

- After a table space restore operation, the table space is always in rollforward pending state, and it must be rolled forward. Invoke the ROLLFORWARD DATABASE command to apply the logs against the table spaces to either a point in time, or to the end of the logs.
- If one or more table spaces are in *rollforward pending* state after crash recovery, first correct the table space problem. In some cases, correcting the table space problem does not involve a restore database operation. For example, a power

loss could leave the table space in rollforward pending state. A restore database operation is not required in this case. Once the problem with the table space is corrected, you can use the ROLLFORWARD DATABASE command to apply the logs against the table spaces to the end of the logs. If the problem is corrected before crash recovery, crash recovery may be sufficient to take the database to a consistent, usable state.

**Note:** If the table space in error contains the system catalog tables, you will not be able to start the database. You must restore the SYSCATSPACE table space, then perform rollforward recovery to the end of the logs.

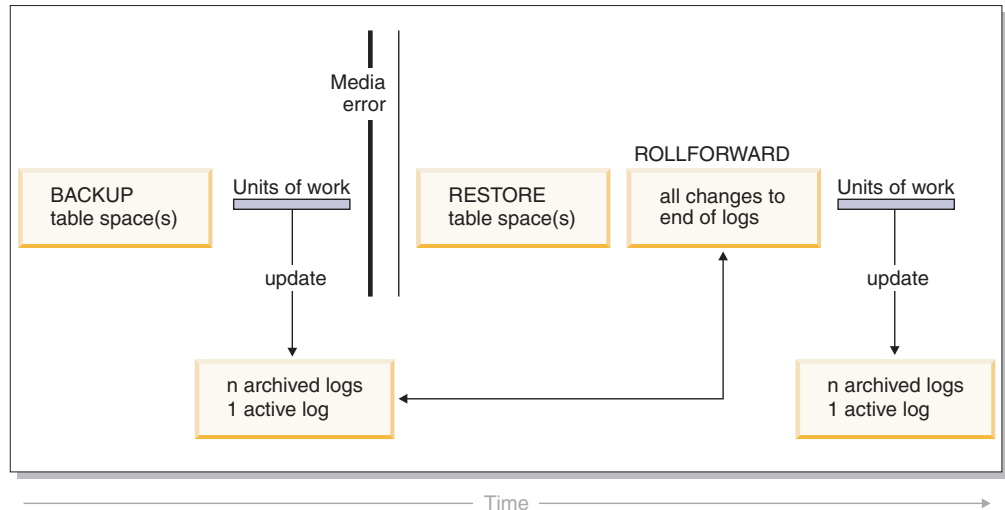


Figure 5. Table Space Rollforward Recovery. There can be more than one active log in the case of a long-running transaction.

In a partitioned database environment, if you are rolling a table space forward *to a point in time*, you do not have to supply the list of nodes (database partitions) on which the table space resides. DB2® submits the rollforward request to all partitions. This means the table space must be restored on all database partitions on which the table space resides.

In a partitioned database environment, if you are rolling a table space forward *to the end of the logs*, you must supply the list of database partitions if you do *not* want to roll the table space forward on all partitions. If you want to roll all table spaces (on all partitions) that are in rollforward pending state forward to the end of the logs, you do not have to supply the list of database partitions. By default, the database rollforward request is sent to all partitions.

**Related concepts:**

- “Understanding recovery logs” on page 31

**Related reference:**

- “ROLLFORWARD DATABASE” on page 123

**Related samples:**

- “dbrecov.out -- HOW TO RECOVER A DATABASE (C)”
- “dbrecov.sqc -- How to recover a database (C)”
- “dbrecov.out -- HOW TO RECOVER A DATABASE (C++)”

- “dbrecov.sqlC -- How to recover a database (C++)”

---

## Incremental backup and recovery

As the size of databases, and particularly warehouses, continues to expand into the terabyte and petabyte range, the time and hardware resources required to back up and recover these databases is also growing substantially. Full database and table space backups are not always the best approach when dealing with large databases, because the storage requirements for multiple copies of such databases are enormous. Consider the following issues:

- When a small percentage of the data in a warehouse changes, it should not be necessary to back up the entire database.
- Appending table spaces to existing databases and then taking only table space backups is risky, because there is no guarantee that nothing outside of the backed up table spaces has changed between table space backups.

To address these issues, DB2<sup>®</sup> provides incremental backup and recovery. An *incremental backup* is a backup image that contains only pages that have been updated since the previous backup was taken. In addition to updated data and index pages, each incremental backup image also contains all of the initial database meta-data (such as database configuration, table space definitions, database history, and so on) that is normally stored in full backup images.

**Note:** If a table space contains long field or large object data and an incremental backup is taken, all of the long field or large object data will be copied into the backup image if any of the pages in that table space have been modified since the previous backup.

Two types of incremental backup are supported:

- *Incremental.* An incremental backup image is a copy of all database data that has changed since the most recent, successful, full backup operation. This is also known as a cumulative backup image, because a series of incremental backups taken over time will each have the contents of the previous incremental backup image. The predecessor of an incremental backup image is always the most recent successful full backup of the same object.
- *Delta.* A delta, or incremental delta, backup image is a copy of all database data that has changed since the last successful backup (full, incremental, or delta) of the table space in question. This is also known as a differential, or non-cumulative, backup image. The predecessor of a delta backup image is the most recent successful backup containing a copy of each of the table spaces in the delta backup image.

The key difference between incremental and delta backup images is their behavior when successive backups are taken of an object that is continually changing over time. Each successive incremental image contains the entire contents of the previous incremental image, plus any data that has changed, or is new, since the previous full backup was produced. Delta backup images contain only the pages that have changed since the previous image of any type was produced.

Combinations of database and table space incremental backups are permitted, in both online and offline modes of operation. Be careful when planning your backup strategy, because combining database and table space incremental backups implies that the predecessor of a database backup (or a table space backup of multiple

table spaces) is not necessarily a single image, but could be a unique set of previous database and table space backups taken at different times.

To rebuild the database or the table space to a consistent state, the recovery process must begin with a consistent image of the entire object (database or table space) to be restored, and must then apply each of the appropriate incremental backup images in the order described below.

To enable the tracking of database updates, DB2 supports a new database configuration parameter, *trackmod*, which can have one of two accepted values:

- NO. Incremental backup is not permitted with this configuration. Database page updates are not tracked or recorded in any way. This is the default value.
- YES. Incremental backup is permitted with this configuration. When update tracking is enabled, the change becomes effective at the first successful connection to the database. Before an incremental backup can be taken on a particular table space, a full backup of that table space is necessary.

For SMS and DMS table spaces, the granularity of this tracking is at the table space level. In table space level tracking, a flag for each table space indicates whether or not there are pages in that table space that need to be backed up. If no pages in a table space need to be backed up, the backup operation can skip that table space altogether.

Although minimal, the tracking of updates to the database can have an impact on the runtime performance of transactions that update or insert data.

**Related tasks:**

- “Restoring from incremental backup images” on page 26

---

## Incremental backup and recovery - details

### Restoring from incremental backup images

**Procedure:**

A restore operation from incremental backup images always consists of the following steps:

1. Identifying the incremental target image.  
Determine the final image to be restored, and request an incremental restore operation from the DB2 restore utility. This image is known as the target image of the incremental restore, because it will be the last image to be restored. The incremental target image is specified using the TAKEN AT parameter in the RESTORE DATABASE command.
2. Restoring the most recent full database or table space image to establish a baseline against which each of the subsequent incremental backup images can be applied.
3. Restoring each of the required full or table space incremental backup images, in the order in which they were produced, on top of the baseline image restored in Step 2.
4. Repeating Step 3 until the target image from Step 1 is read a second time. The target image is accessed twice during a complete incremental restore operation.

During the first access, only initial data is read from the image; none of the user data is read. The complete image is read and processed only during the second access.

The target image of the incremental restore operation must be accessed twice to ensure that the database is initially configured with the correct history, database configuration, and table space definitions for the database that will be created during the restore operation. In cases where a table space has been dropped since the initial full database backup image was taken, the table space data for that image will be read from the backup images but ignored during incremental restore processing.

There are two ways to restore incremental backup images.

- For an automatic incremental restore, the RESTORE command is issued only once specifying the target image to be used. DB2 then uses the database history to determine the remaining required backup images and restores them.
- For a manual incremental restore, the RESTORE command must be issued once for each backup image that needs to be restored (as outlined in the steps above).

### Automatic Incremental Restore Example

To restore a set of incremental backup images using automatic incremental restore, specify the TAKEN AT timestamp option on the RESTORE DATABASE command. Use the time stamp for the last image that you want to restore. For example:

```
db2 restore db sample incremental automatic taken at 20031228152133
```

This will result in the DB2 restore utility performing each of the steps described at the beginning of this section automatically. During the initial phase of processing, the backup image with time stamp 20001228152133 is read, and the restore utility verifies that the database, its history, and the table space definitions exist and are valid.

During the second phase of processing, the database history is queried to build a chain of backup images required to perform the requested restore operation. If, for some reason this is not possible, and DB2 is unable to build a complete chain of required images, the restore operation terminates, and an error message is returned. In this case, an automatic incremental restore will not be possible, and you will have issue the RESTORE DATABASE command with the INCREMENTAL ABORT option. This will cleanup any remaining resources so that you can proceed with a manual incremental restore.

**Note:** It is highly recommended that you not use the FORCE option of the PRUNE HISTORY command. The default operation of this command prevents you from deleting history entries that may be required for recovery from the most recent, full database backup image, but with the FORCE option, it is possible to delete entries that are required for an automatic restore operation.

During the third phase of processing, DB2 will restore each of the remaining backup images in the generated chain. If an error occurs during this phase, you will have to issue the RESTORE DATABASE command with the INCREMENTAL ABORT option to cleanup any remaining resources. You will then have to determine if the error can be resolved before you re-issue the RESTORE command or attempt the manual incremental restore again.

### Manual Incremental Restore Example

To restore a set of incremental backup images, using manual incremental restore, specify the target image using the TAKEN AT *timestamp* option of the RESTORE DATABASE command and follow the steps outlined above. For example:

1. db2 restore database sample incremental taken at <ts>  
  
where:  
    <ts> points to the last incremental backup image (the target image) to be restored
2. db2 restore database sample incremental taken at <ts1>  
  
where:  
    <ts1> points to the initial full database (or table space) image
3. db2 restore database sample incremental taken at <tsX>  
  
where:  
    <tsX> points to each incremental backup image in creation sequence
4. Repeat Step 3, restoring each incremental backup image up to and including image <ts>

If you are performing a database restore operation, and table space backup images have been produced, the table space images must be restored in the chronological order of their backup time stamps.

The **db2ckrst** utility can be used to query the database history and generate a list of backup image time stamps needed for an incremental restore. A simplified restore syntax for a manual incremental restore is also generated. It is recommended that you keep a complete record of backups, and only use this utility as a guide.

**Related concepts:**

- “Incremental backup and recovery” on page 25

**Related reference:**

- “RESTORE DATABASE” on page 87
- “db2ckrst - Check Incremental Restore Image Sequence” on page 247

## Limitations to automatic incremental restore

1. If a table space name has been changed since the backup operation you want to restore from, and you use the new name when you issue a table level restore operation, the required chain of backup images from the database history will not be generated correctly and an error will occur (SQL2571N).

Example:

```
db2 backup db sample -> <ts1>
db2 backup db sample incremental -> <ts2>
db2 rename tablespace from userspace1 to t1
db2 restore db sample tablespace ('t1') incremental automatic taken
at <ts2>
```

```
SQL2571N Automatic incremental restore is unable to proceed.
Reason code: "3".
```

Suggested workaround: Use manual incremental restore.

2. If you drop a database, the database history will be deleted. If you restore the dropped database, the database history will be restored to its state at the time of the restored backup and all history entries after that time will be lost. If you then attempt to perform an automatic incremental restore that would need to

use any of these lost history entries, the RESTORE utility will attempt to restore an incorrect chain of backups and will return an "out of sequence" error (SQL2572N).

Example:

```
db2 backup db sample -> <ts1>
db2 backup db sample incremental -> <ts2>
db2 backup db sample incremental delta -> <ts3>
db2 backup db sample incremental delta -> <ts4>
db2 drop db sample
db2 restore db sample incremental automatic taken at <ts2>
db2 restore db sample incremental automatic taken at <ts4>
```

Suggested workarounds:

- Use manual incremental restore.
  - Restore the history file first from image <ts4> before issuing an automatic incremental restore.
3. If you restore a backup image from one database into another database and then do an incremental (delta) backup, you can no longer use automatic incremental restore to restore this backup image.

Example:

```
db2 create db a
db2 create db b

db2 update db cfg for a using trackmod on

db2 backup db a -> ts1
db2 restore db a taken at ts1 into b

db2 backup db b incremental -> ts2

db2 restore db b incremental automatic taken at ts2
```

SQL2542N No match for a database image file was found based on the source database alias "B" and timestamp "ts1" provided.

Suggested workaround:

- Use manual incremental restore as follows:

```
db2 restore db b incremental taken at ts2
db2 restore db a incremental taken at ts1 into b
db2 restore db b incremental taken at ts2
```
- After the manual restore operation into database B, issue a full database backup to start a new incremental chain

**Related concepts:**

- "Incremental backup and recovery" on page 25

**Related tasks:**

- "Restoring from incremental backup images" on page 26

**Related reference:**

- "RESTORE DATABASE" on page 87

---

## Monitoring the progress of backup, restore and recovery operations

You can use the LIST UTILITIES command to monitor backup, restore, and rollforward operations on a database.

### Procedure:

To use progress monitoring for backup, restore and recovery operations:

- Issue the LIST UTILITIES command and specify the SHOW DETAIL option.

```
list utilities show detail
```

The following is an example of the output for monitoring the performance of an offline database backup operation:

```
LIST UTILITIES SHOW DETAIL

ID                = 2
Type              = BACKUP
Database Name     = SAMPLE
Description       = offline db
Start Time        = 10/30/2003 12:55:31.786115
Throttling:
Priority          = Unthrottled
Progress Monitoring:
Estimated Percentage Complete = 41
  Total Work Units = 20232453 bytes
  Completed Work Units = 230637 bytes
  Start Time      = 10/30/2003 12:55:31.786115
```

For backup operations, an initial estimate of the number of bytes to be processed will be specified. As the backup operation progresses the number of bytes to be processed will be updated. The bytes shown does not correspond to the size of the image and should not be used as an estimate for backup image size. The actual image might be much smaller depending on whether it is an incremental or compressed backup.

For restore operations, no initial estimate will be given. Instead UNKNOWN will be specified. As each buffer is read from the image, the actual amount of bytes read will be updated. For automatic incremental restore operations where multiple images might be restored, the progress will be tracked using phases. Each phase represents an image to be restored from the incremental chain. Initially, only one phase will be indicated. After the first image is restored, the total number of phases will be indicated. As each image is restored the number of phases completed will be updated as will the number of bytes processed.

| For crash recovery and rollforward recovery there will be two phases of progress  
| monitoring: FORWARD and BACKWARD. During the FORWARD phase, log files  
| are read and the log records are applied to the database. For crash recovery, the  
| total amount of work is estimated using the starting log sequence number up to  
| the end of the last log file. For rollforward recovery, when this phase begins  
| UNKNOWN will be specified for the total work estimate. The amount of work  
| processed in bytes will be updated as the process continues.

During the BACKWARD phase, any uncommitted changes applied during the FORWARD phase are rolled back. An estimate for the amount of log data to be processed in bytes will be provided. The amount of work processed in bytes will be updated as the process continues.



**Related concepts:**

- “Crash recovery” on page 9
- “Backup overview” on page 59
- “Restore overview” on page 81
- “Rollforward overview” on page 111

**Related reference:**

- “LIST UTILITIES Command” in the *Command Reference*

---

## Understanding recovery logs

All databases have logs associated with them. These logs keep records of database changes. If a database needs to be restored to a point beyond the last full, offline backup, logs are required to roll the data forward to the point of failure.

There are two types of DB2® logging: *circular*, and *archive*, each provides a different level of recovery capability:

- *Circular logging* is the default behavior when a new database is created. (The *logarchmeth1* and *logarchmeth2* database configuration parameters are set to OFF.) With this type of logging, only full, offline backups of the database are allowed. The database must be offline (inaccessible to users) when a full backup is taken. As the name suggests, circular logging uses a “ring” of online logs to provide recovery from transaction failures and system crashes. The logs are used and retained only to the point of ensuring the integrity of current transactions. Circular logging does not allow you to roll a database forward through transactions performed after the last full backup operation. All changes occurring since the last backup operation are lost. Since this type of restore operation recovers your data to the specific point in time at which a full backup was taken, it is called *version recovery*.

Figure 6 shows that the active log uses a ring of log files when circular logging is active.

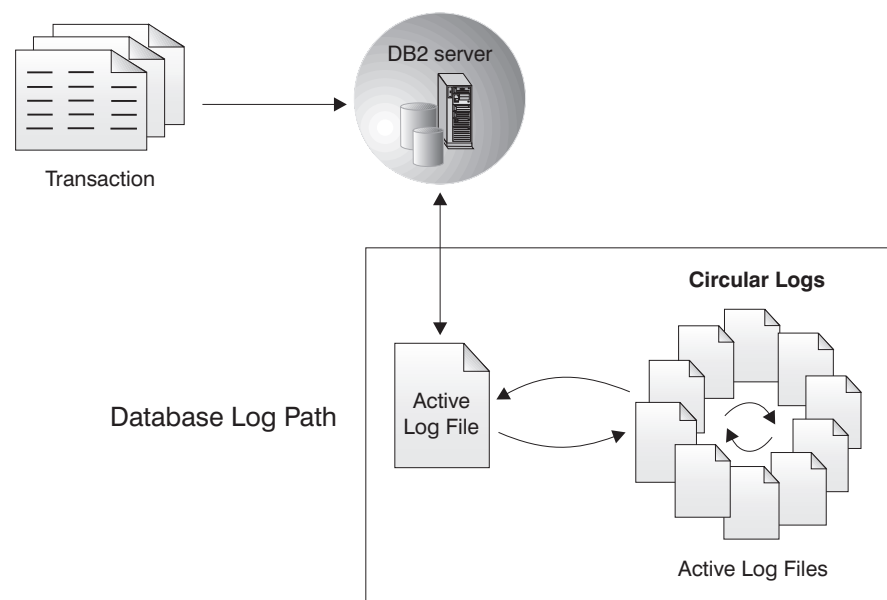



Figure 6. Circular Logging

Active logs are used during crash recovery to prevent a failure (system power or application error) from leaving a database in an inconsistent state. Active logs are located in the database log path directory.

- *Archive* logging is used specifically for rollforward recovery. Archived logs are logs that were active but are no longer required for crash recovery. Use the *logarchmeth1* database configuration parameter to enable archive logging.

The advantage of choosing archive logging is that rollforward recovery can use both archived logs and active logs to rebuild a database either to the end of the logs, or to a specific point in time. The archived log files can be used to recover changes made after the backup was taken. This is different from circular logging where you can only recover to the time of the backup, and all changes made after that are lost.

Taking online backups is only supported if the database is configured for archive logging. During an online backup operation, all activities against the database are logged. When an online backup image is restored, the logs must be rolled forward at least to the point in time at which the backup operation completed. For this to happen, the logs must have been archived and made available when the database is restored. After an online backup is complete, DB2 forces the currently active log to be closed, and as a result, it will be archived. This ensures that your online backup has a complete set of archived logs available for recovery.



Logs are used between backups to track the changes to the databases.

*Figure 7. Active and Archived Database Logs in Rollforward Recovery.* There can be more than one active log in the case of a long-running transaction.

The following database configuration parameters allow you to change where archived logs are stored: The *newlogpath* parameter, and the *logarchmeth1* and *logarchmeth2* parameters. Changing the *newlogpath* parameter also affects where active logs are stored.

To determine which log *extents* in the database log path directory are archived logs, check the value of the *loghead* database configuration parameter. This parameter indicates the lowest numbered log that is active. Those logs with sequence numbers less than *loghead* are archived logs and can be moved. You can check the value of this parameter by using the Control Center; or, by using the command line processor and the GET DATABASE CONFIGURATION command to view the "First active log file". For more information about this configuration parameter, see the *Administration Guide: Performance* book.

**Related concepts:**

- “Log mirroring” on page 33

**Related reference:**

- Appendix G, “User exit for database recovery,” on page 329
- “loghead - First active log file configuration parameter” in the *Administration Guide: Performance*
- “Configuration parameters for database logging” on page 35
- “logarchmeth1 - Primary log archive method configuration parameter” in the *Administration Guide: Performance*

---

## Recovery log details

### Log mirroring

DB2® supports log mirroring at the database level. Mirroring log files helps protect a database from:

- Accidental deletion of an active log
- Data corruption caused by hardware failure

If you are concerned that your active logs may be damaged (as a result of a disk crash), you should consider using the `MIRRORLOGPATH` configuration parameter to specify a secondary path for the database to manage copies of the active log, mirroring the volumes on which the logs are stored.

The `MIRRORLOGPATH` configuration parameter allows the database to write an identical second copy of log files to a different path. It is recommended that you place the secondary log path on a physically separate disk (preferably one that is also on a different disk controller). That way, the disk controller cannot be a single point of failure.

When `MIRRORLOGPATH` is first enabled, it will not actually be used until the next database startup. This is similar to the `NEWLOGPATH` configuration parameter.

If there is an error writing to either the active log path or the mirror log path, the database will mark the failing path as “bad”, write a message to the administration notification log, and write subsequent log records to the remaining “good” log path only. DB2 will not attempt to use the “bad” path again until the current log file is completed. When DB2 needs to open the next log file, it will verify that this path is valid, and if so, will begin to use it. If not, DB2 will not attempt to use the path again until the next log file is accessed for the first time. There is no attempt to synchronize the log paths, but DB2 keeps information about access errors that occur, so that the correct paths are used when log files are archived. If a failure occurs while writing to the remaining “good” path, the database shuts down.

**Related reference:**

- “mirrorlogpath - Mirror log path configuration parameter” in the *Administration Guide: Performance*

## Reducing logging with the NOT LOGGED INITIALLY parameter

If your application creates and populates work tables from master tables, and you are not concerned about the recoverability of these work tables because they can be easily recreated from the master tables, you may want to create the work tables specifying the NOT LOGGED INITIALLY parameter on the CREATE TABLE statement. The advantage of using the NOT LOGGED INITIALLY parameter is that any changes made on the table (including insert, delete, update, or create index operations) in the same unit of work that creates the table will not be logged. This not only reduces the logging that is done, but may also increase the performance of your application. You can achieve the same result for existing tables by using the ALTER TABLE statement with the NOT LOGGED INITIALLY parameter.

### Notes:

1. You can create more than one table with the NOT LOGGED INITIALLY parameter in the same unit of work.
2. Changes to the catalog tables and other user tables are still logged.

Because changes to the table are not logged, you should consider the following when deciding to use the NOT LOGGED INITIALLY table attribute:

- All changes to the table will be flushed out to disk at commit time. This means that the commit may take longer.
- If the NOT LOGGED INITIALLY attribute is activated and an activity occurs that is not logged, the entire unit of work will be rolled back if a statement fails or a ROLLBACK TO SAVEPOINT is executed (SQL1476N).
- You cannot recover these tables when rolling forward. If the rollforward operation encounters a table that was created or altered with the NOT LOGGED INITIALLY option, the table is marked as unavailable. After the database is recovered, any attempt to access the table returns SQL1477N.

**Note:** When a table is created, row locks are held on the catalog tables until a COMMIT is done. To take advantage of the no logging behavior, you must populate the table in the same unit of work in which it is created. This has implications for concurrency.

## Reducing logging with declared temporary tables

If you plan to use declared temporary tables as work tables, note the following:

- Declared temporary tables are not created in the catalogs; therefore locks are not held.
- Logging is not performed against declared temporary tables, even after the first COMMIT.
- Use the ON COMMIT PRESERVE option to keep the rows in the table after a COMMIT; otherwise, all rows will be deleted.
- Only the application that creates the declared temporary table can access that instance of the table.
- The table is implicitly dropped when the application connection to the database is dropped.
- Errors in operation during a unit of work using a declared temporary table do not cause the unit of work to be completely rolled back. However, an error in operation in a statement changing the contents of a declared temporary table

will delete all the rows in that table. A rollback of the unit of work (or a savepoint) will delete all rows in declared temporary tables that were modified in that unit of work (or savepoint).

**Related concepts:**

- “Application processes, concurrency, and recovery” in the *SQL Reference, Volume 1*

**Related tasks:**

- “Creating a table space” in the *Administration Guide: Implementation*

**Related reference:**

- “DECLARE GLOBAL TEMPORARY TABLE statement” in the *SQL Reference, Volume 2*

## Configuration parameters for database logging

### Archive Retry Delay (`archretrydelay`)

Specifies the amount of time (in seconds) to wait between attempts to archive log files after the previous attempt fails. The default value is 20.

### Block on log disk full (`blk_log_dsk_ful`)

This configuration parameter can be set to prevent disk full errors from being generated when DB2 cannot create a new log file in the active log path. Instead, DB2 will attempt to create the log file every five minutes until it succeeds. After each attempt, DB2 will write a message to the administration notification log. The only way to confirm that your application is hanging because of a log disk full condition is to monitor the administration notification log. Until the log file is successfully created, any user application that attempts to update table data will not be able to commit transactions. Read-only queries may not be directly affected; however, if a query needs to access data that is locked by an update request or a data page that is fixed in the buffer pool by the updating application, read-only queries will also appear to hang.

Setting `blk_log_dsk_ful` to YES causes applications to hang when DB2 encounters a log disk full error. You are then able to resolve the error and the transaction can continue. A disk full situation can be resolved by moving old log files to another file system, or by increasing the size of the file system so that hanging applications can complete.

If `blk_log_dsk_ful` is set to NO, a transaction that receives a log disk full error will fail and be rolled back. In some cases, the database will come down if a transaction causes a log disk full error.

### Failover Archive Path (`failarchpath`)

Specifies an alternate directory for the archive log files if the log archive method specified fails. This directory is a temporary storage area for the log files until the log archive method that failed becomes available again at which time the log files will be moved from this directory to the log archive method. By moving the log files to this temporary location, log directory full situations might be avoided. This parameter must be a fully qualified existing directory.

### Log archive method 1 (`logarchmeth1`), log archive method 2 (`logarchmeth2`)

These parameters cause the database manager to archive log files to a location that is not the active log path. If both of these parameters are

specified, each log file is archived twice. This means that you will have two copies of archived log files in two different locations.

Valid values for these parameters include a media type and, in some cases, a target field. Use a colon (:) to separate the values. Valid values are:

**OFF** Specifies that the log archiving method is not to be used. If both *logarchmeth1* and *logarchmeth2* are set to OFF, the database is considered to be using circular logging and will not be rollforward recoverable. This is the default.

#### **LOGRETAIN**

This value can only be used for *logarchmeth1* and is equivalent to setting the *logretain* configuration parameter to RECOVERY. If you specify this value, the *logretain* configuration parameters will automatically be updated.

#### **USEREXIT**

This value is only valid for *logarchmeth1* and is equivalent to setting the *userexit* configuration parameter to ON. If specify this value, the *userexit* configuration parameter will be automatically updated.

**DISK** This value must be followed by a colon(:) and then a fully qualified existing path name where the log files will be archived. For example, if you set *logarchmeth1* to DISK:/u/dbuser/archived\_logs the archive log files will be placed in a directory called /u/dbuser/archived\_logs.

**Note:** If you are archiving to tape, you can use the db2tapemgr utility to store and retrieve log files.

**TSM** If specified without any additional configuration parameters, this value indicates that log files should be archived on the local TSM server using the default management class. If followed by a colon(:) and a TSM management class, the log files will be archived using the specified management class.

#### **VENDOR**

Specifies that a vendor library will be used to archive the log files. This value must be followed by a colon(:) and the name of the library. The APIs provided in the library must use the backup and restore APIs for vendor products.

#### **Notes:**

1. If either *logarchmeth1* or *logarchmeth2* is set to a value other than OFF, the database is configured for rollforward recovery.
2. If you update the *userexit* or *logretain* configuration parameters *logarchmeth1* will automatically be updated and vice versa. However, if you are using either *userexit* or *logretain*, *logarchmeth2* must be set to OFF.

#### **Log archive options 1 (logarchopt1), log archive options 2 (logarchopt2)**

Specifies a string which is passed on to the TSM server or vendor APIs. For TSM, this field is used to allow the database to retrieve logs that were generated on a different TSM node or by a different TSM user. The string must be provided in the following format:

```
"-fromnode=nodename -fromowner=ownername"
```

where nodename is the name of the TSM node that originally archived the log files, and ownername is the name of the TSM user that originally

archived the log files. Each log archive options field corresponds to one of the log archive methods: *logarchopt1* is used with *logarchmeth1*, and *logarchopt2* is used with *logarchmeth2*.

### **Log Buffer (logbufsz)**

This parameter allows you to specify the amount of memory to use as a buffer for log records before writing these records to disk. The log records are written to disk when any one of the following events occurs:

- A transaction commits
- The log buffer becomes full
- Some other internal database manager event occurs.

Increasing the log buffer size results in more efficient input/output (I/O) activity associated with logging, because the log records are written to disk less frequently, and more records are written each time.

### **Log file size (logfilsiz)**

This parameter specifies the size of each configured log, in number of 4-KB pages.

There is a 256-GB logical limit on the total active log space that you can configure. This limit is the result of the upper limit on *logfilsiz*, which is 262144, and the upper limit on (*logprimary* + *logsecond*), which is 256.

The size of the log file has a direct bearing on performance. There is a performance cost for switching from one log to another. So, from a pure performance perspective, the larger the log file size the better. This parameter also indicates the log file size for archiving. In this case, a larger log file is size it not necessarily better, since a larger log file size may increase the chance of failure or cause a delay in log shipping scenarios. When considering active log space, it may be better to have a larger number of smaller log files. For example, if there are 2 very large log files and a transaction starts close to the end of one log file, only half of the log space remains available.

Every time a database is deactivated (all connections to the database are terminated), the log file that is currently being written is truncated. So, if a database is frequently being deactivated, it is better not to choose a large log file size because DB2 will create a large file only to have it truncated. You can use the ACTIVATE DATABASE command to avoid this cost, and having the buffer pool primed will also help with performance.

Assuming that you have an application that keeps the database open to minimize processing time when opening the database, the log file size should be determined by the amount of time it takes to make offline archived log copies.

Minimizing log file loss is also an important consideration when setting the log size. Archiving takes an entire log. If you use a single large log, you increase the time between archiving. If the medium containing the log fails, some transaction information will probably be lost. Decreasing the log size increases the frequency of archiving but can reduce the amount of information loss in case of a media failure since the smaller logs before the one lost can be used.

### **Log retain (logretain)**

This configuration parameter has been replaced by *logarchmeth1*. It is still supported for back-level compatibility.

If `logretain` is set to `RECOVERY`, archived logs are kept in the database log path directory, and the database is considered to be recoverable, meaning that rollforward recovery is enabled.

**Note:** The default value for the `logretain` database configuration parameter does not support rollforward recovery. You must change the value of this parameter if you are going to use rollforward recovery.

### Maximum log per transaction (`max_log`)

This parameter indicates the percentage of primary log space that can be consumed by one transaction.

If the value is set to 0, there is no limit to the percentage of total primary log space that a transaction can consume. If an application violates the `max_log` configuration, the application will be forced to disconnect from the database, the transaction will be rolled back, and error `SQL1224N` will be returned.

You can override this behavior by setting the `DB2_FORCE_APP_ON_MAX_LOG` registry variable to `FALSE`. This will cause transactions that violate the `max_log` configuration to fail and return error `SQL0964N`. The application can still commit the work completed by previous statements in the unit or work, or it can roll the work completed back to undo the unit of work.

**Note:** The following DB2 commands are excluded from the limitation imposed by the `max_log` configuration parameter: `ARCHIVE LOG`, `BACKUP DATABASE`, `LOAD`, `REORG TABLE (online)`, `RESTORE DATABASE`, and `ROLLFORWARD DATABASE`.

### Mirror log path (`mirrorlogpath`)

To protect the logs on the primary log path from disk failure or accidental deletion, you can specify that an identical set of logs be maintained on a secondary (mirror) log path. To do this, change the value of this configuration parameter to point to a different directory. Active logs that are currently stored in the mirrored log path directory are not moved to the new location if the database is configured for rollforward recovery.

Because you can change the log path location, the logs needed for rollforward recovery may exist in different directories. You can change the value of this configuration parameter during a rollforward operation to allow you to access logs in multiple locations.

You must keep track of the location of the logs.

Changes are not applied until the database is in a consistent state. The configuration parameter `database_consistent` returns the status of the database.

To turn this configuration parameter off, set its value to `DEFAULT`.

#### Notes:

1. This configuration parameter is not supported if the primary log path is a raw device.
2. The value specified for this parameter cannot be a raw device.

### New log path (`newlogpath`)

The database logs are initially created in `SQLOGDIR`, which is a subdirectory of the database directory. You can change the location in which active logs and future archived logs are placed by changing the



value of this configuration parameter to point to a different directory or to a device. Active logs that are currently stored in the database log path directory are not moved to the new location if the database is configured for rollforward recovery.

Because you can change the log path location, the logs needed for rollforward recovery may exist in different directories or on different devices. You can change the value of this configuration parameter during a rollforward operation to allow you to access logs in multiple locations.

You must keep track of the location of the logs.

Changes are not applied until the database is in a consistent state. The configuration parameter *database\_consistent* returns the status of the database.

#### **Number of Commits to Group (mincommit)**

This parameter allows you to delay the writing of log records to disk until a minimum number of commits have been performed. This delay can help reduce the database manager overhead associated with writing log records and, as a result, improve performance when you have multiple applications running against a database, and many commits are requested by the applications within a very short period of time.

The grouping of commits occurs only if the value of this parameter is greater than 1, and if the number of applications connected to the database is greater than the value of this parameter. When commit grouping is in effect, application commit requests are held until either one second has elapsed, or the number of commit requests equals the value of this parameter.

#### **Number of archive retries on error (numarchretry)**

Specifies the the number of attempts that will be made to archive log files using the specified log archive method before they are archived to the path specified by the *failarchpath* configuration parameter. This parameter can only be used if the *failarchpath* configuration parameter is set. The default value is 5.

#### **Number log span (num\_log\_span)**

This parameter indicates the number of active log files that an active transaction can span. If the value is set to 0, there is no limit to how many log files one single transaction can span.

If an application violates the *num\_log\_span* configuration, the application will be forced to disconnect from the database and error SQL1224N will be returned.

**Note:** The following DB2 commands are excluded from the limitation imposed by the *num\_log\_span* configuration parameter: ARCHIVE LOG, BACKUP DATABASE, LOAD, REORG TABLE (online), RESTORE DATABASE, and ROLLFORWARD DATABASE.

#### **Overflow log path (overflowlogpath)**

This parameter can be used for several functions, depending on your logging requirements. You can specify a location for DB2 to find log files that are needed for a rollforward operation. It is similar to the OVERFLOW LOG PATH option of the ROLLFORWARD command; however, instead of specifying the OVERFLOW LOG PATH option for every ROLLFORWARD command issued, you can set this configuration parameter once. If both are

used, the OVERFLOW LOG PATH option will overwrite the *overflowlogpath* configuration parameter for that rollforward operation.

If *logsecond* is set to -1, you can specify a directory for DB2 to store active log files retrieved from the archive. (Active log files must be retrieved for rollback operations if they are no longer in the active log path).

If *overflowlogpath* is not specified, DB2 will retrieve the log files into the active log path. By specifying this parameter you can provide additional resource for DB2 to store the retrieved log files. The benefit includes spreading the I/O cost to different disks, and allowing more log files to be stored in the active log path.

For example, if you are using the **db2ReadLog** API for replication, you can use *overflowlogpath* to specify a location for DB2 to search for log files that are needed for this API. If the log file is not found (in either the active log path or the overflow log path) and the database is configured with *userexit* enabled, DB2 will retrieve the log file. You can also use this parameter to specify a directory for DB2 to store the retrieved log files. The benefit comes from reducing the I/O cost on the active log path and allowing more log files to be stored in the active log path.

If you have configured a raw device for the active log path, *overflowlogpath* must be configured if you want to set *logsecond* to -1, or if you want to use the **db2ReadLog** API.

To set *overflowlogpath*, specify a string of up to 242 bytes. The string must point to a path name, and it must be a fully qualified path name, not a relative path name. The path name must be a directory, not a raw device.

**Note:** In a partitioned database environment, the node number is automatically appended to the path. This is done to maintain the uniqueness of the path in multiple logical node configurations.

### Primary logs (logprimary)

This parameter specifies the number of primary logs of size *logfilsz* that will be created.

A primary log, whether empty or full, requires the same amount of disk space. Thus, if you configure more logs than you need, you use disk space unnecessarily. If you configure too few logs, you can encounter a log-full condition. As you select the number of logs to configure, you must consider the size you make each log and whether your application can handle a log-full condition. The total log file size limit on active log space is 256 GB.

If you are enabling an existing database for rollforward recovery, change the number of primary logs to the sum of the number of primary and secondary logs, plus 1. Additional information is logged for LONG VARCHAR and LOB fields in a database enabled for rollforward recovery.

### Secondary logs (logsecond)

This parameter specifies the number of secondary log files that are created and used for recovery, if needed.

If the primary log files become full, secondary log files (of size *logfilsiz*) are allocated, one at a time as needed, up to the maximum number specified by this parameter. If this parameter is set to -1, the database is configured with infinite active log space. There is no limit on the size or number of in-flight transactions running on the database.

**Notes:**

1. Log archiving must be enabled in order to set *logsecond* to -1.
2. If this parameter is set to -1, crash recovery time may be increased since DB2 may need to retrieve archived log files.

**User exit (userexit)**

This configuration parameter has been replaced by *logarchmeth1*. It is still supported for back-level compatibility.

This parameter causes the database manager to call a user exit program for archiving and retrieving logs. The log files are archived in a location that is different from the active log path. If *userexit* is set to 0N, rollforward recovery is enabled.

The data transfer speed of the device you use to store offline archived logs, and the software used to make the copies, must at a minimum match the average rate at which the database manager writes data in the logs. If the transfer speed cannot keep up with new log data being generated, you may run out of disk space if logging activity continues for a sufficiently long period of time. The amount of time it takes to run out of disk space is determined by the amount of free disk space. If this happens, database processing stops.

The data transfer speed is most significant when using tape or an optical medium. Some tape devices require the same amount of time to copy a file, regardless of its size. You must determine the capabilities of your archiving device.

Tape devices have other considerations. The frequency of the archiving request is important. For example, if the time taken to complete any copy operation is five minutes, the log should be large enough to hold five minutes of log data during your peak work load. The tape device may have design limits that restrict the number of operations per day. These factors must be considered when you determine the log size.

**Notes:**

1. This value must be set to 0N to enable infinite active log space.
2. The default value for the *userexit* database configuration parameter does not support rollforward recovery, and must be changed if you are going to use it.

**Related concepts:**

- “Managing log files” on page 42
- “Enhancing recovery performance” on page 56

**Related reference:**

- Appendix G, “User exit for database recovery,” on page 329
- “logfilsiz - Size of log files configuration parameter” in the *Administration Guide: Performance*
- “logprimary - Number of primary log files configuration parameter” in the *Administration Guide: Performance*
- “logsecond - Number of secondary log files configuration parameter” in the *Administration Guide: Performance*
- “logbufsz - Log buffer size configuration parameter” in the *Administration Guide: Performance*

- “mincommit - Number of commits to group configuration parameter” in the *Administration Guide: Performance*
- “newlogpath - Change the database log path configuration parameter” in the *Administration Guide: Performance*
- “logretain - Log retain enable configuration parameter” in the *Administration Guide: Performance*
- “userexit - User exit enable configuration parameter” in the *Administration Guide: Performance*
- “APIs for backup and restore to storage managers” on page 333
- “UPDATE DATABASE CONFIGURATION Command” in the *Command Reference*
- “mirrorlogpath - Mirror log path configuration parameter” in the *Administration Guide: Performance*
- “blk\_log\_dsk\_ful - Block on log disk full configuration parameter” in the *Administration Guide: Performance*
- “overflowlogpath - Overflow log path configuration parameter” in the *Administration Guide: Performance*
- “max\_log - Maximum log per transaction configuration parameter” in the *Administration Guide: Performance*
- “num\_log\_span - Number log span configuration parameter” in the *Administration Guide: Performance*
- “db2tapemgr - Manage Log Files on Tape Command” in the *Command Reference*

## Managing log files

Consider the following when managing database logs:

- The numbering scheme for archived logs starts with S0000000.LOG, and continues through S9999999.LOG, accommodating a potential maximum of 10 million log files. The database manager resets to S0000000.LOG if:
  - A database configuration file is changed to enable rollforward recovery
  - A database configuration file is changed to *disable* rollforward recovery
  - S9999999.LOG has been used.

DB2<sup>®</sup> reuses log file names after restoring a database (with or without rollforward recovery). The database manager ensures that an incorrect log is not applied during rollforward recovery. If DB2 reuses a log file name after a restore operation, the new log files are archived to separate directories so that multiple log files with the same name can be archived. The location of the log files is recorded in the recovery history file so that they can be applied during rollforward recovery. You must ensure that the correct logs are available for rollforward recovery.

When a rollforward operation completes successfully, the last log that was used is truncated, and logging begins with the next sequential log. Any log in the log path directory with a sequence number greater than the last log used for rollforward recovery is re-used. Any entries in the truncated log following the truncation point are overwritten with zeros. Ensure that you make a copy of the logs before invoking the rollforward utility. (You can invoke a user exit program to copy the logs to another location.)

- If a database has not been activated (by way of the ACTIVATE DATABASE command), DB2 truncates the current log file when all applications have disconnected from the database. The next time an application connects to the database, DB2 starts logging to a new log file. If many small log files are being produced on your system, you may want to consider using the ACTIVATE

DATABASE command. This not only saves the overhead of having to initialize the database when applications connect, it also saves the overhead of having to allocate a large log file, truncate it, and then allocate a new large log file.

- An archived log may be associated with two or more different *log sequences* for a database, because log file names are reused (see Figure 8). For example, if you want to recover Backup 2, there are two possible log sequences that could be used. If, during full database recovery, you roll forward to a point in time and stop before reaching the end of the logs, you have created a new log sequence. The two log sequences cannot be combined. If you have an online backup image that spans the first log sequence, you must use this log sequence to complete rollforward recovery.

If you have created a new log sequence after recovery, any table space backup images on the old log sequence are invalid. This is usually recognized at restore time, but the restore utility fails to recognize a table space backup image on an old log sequence if a database restore operation is immediately followed by the table space restore operation. Until the database is actually rolled forward, the log sequence that is to be used is unknown. If the table space is on an old log sequence, it must be “caught” by the table space rollforward operation. A restore operation using an invalid backup image may complete successfully, but the table space rollforward operation for that table space will fail, and the table space will be left in restore pending state.

For example, suppose that a table space-level backup operation, Backup 3, completes between S0000013.LOG and S0000014.LOG in the top log sequence (see Figure 8). If you want to restore and roll forward using the database-level backup image, Backup 2, you will need to roll forward through S0000012.LOG. After this, you could continue to roll forward through either the top log sequence or the (newer) bottom log sequence. If you roll forward through the bottom log sequence, you will not be able to use the table space-level backup image, Backup 3, to perform table space restore and rollforward recovery.

To complete a table space rollforward operation to the end of the logs using the table space-level backup image, Backup 3, you will have to restore the database-level backup image, Backup 2, and then roll forward using the top log sequence. Once the table space-level backup image, Backup 3, has been restored, you can initiate a rollforward operation to the end of the logs.

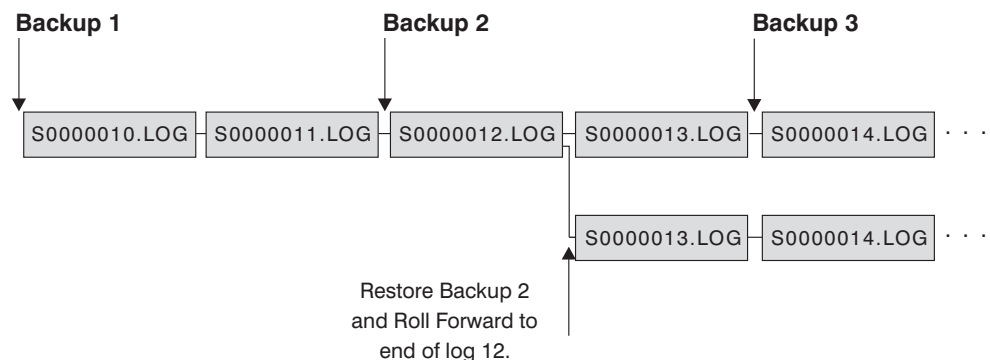


Figure 8. Re-using Log File Names

**Related reference:**

- Appendix G, “User exit for database recovery,” on page 329

## Log file allocation and removal

Log files in the database log directory are never removed if they may be required for crash recovery. However, if you have enabled infinite logging, log files will be deleted once they have been successfully archived. When the *logarchmeth1* database configuration parameter is not set to OFF, a full log file becomes a candidate for removal only after it is no longer required for crash recovery. A log file which is required for crash recovery is called an active log. A log file which is not required for crash recovery is called an archived log.

The process of allocating new log files and removing old log files is dependent on the settings of the *logarchmeth1* database configuration parameter:

### *Logarchmeth1* and *Logarchmeth2* are set to OFF

Circular logging will be used. Rollforward recovery is not supported with circular logging, while crash recovery is.

During circular logging, new log files, other than secondary logs, are not generated and old log files are not deleted. Log files are handled in a circular fashion. That is, when the last log file is full, DB2® begins writing to the first log file.

A log full situation can occur if all of the log files are active and the circular logging process cannot wrap to the first log file. Secondary log files are created when all the primary log files are active and full. Secondary log files are deleted when the database is deactivated or when the space they are using is required for the active log files.

### *Logarchmeth1* is set to LOGRETAIN

Archive logging is used. The database is a recoverable database. Both rollforward recovery and crash recovery are enabled. After you archive the log files, you must delete them from the active log path so that the disk space can be reused for new log files. Each time a log file becomes full, DB2 begins writing records to another log file, and (if the maximum number of primary and secondary logs has not been reached) creates a new log file.

### *Logarchmeth1* is set to a value other than OFF or LOGRETAIN

Archive logging is used. The database is a recoverable database. Both rollforward recovery and crash recovery are enabled. When a log file becomes full, it is automatically archived using the user supplied user exit program.

Log files are usually not deleted. Instead, when a new log file is required and one is not available, an archived log file is renamed and used again. An archived log file, is not deleted or renamed once it has been closed and copied to the log archive directory. DB2 waits until a new log file is needed and then renames the oldest archived log. A log file that has been moved to the database directory during recovery is removed during the recovery process when it is no longer needed. Until DB2 runs out of log space, you will see old log files in the database directory.

If an error occurs when log files are being archived, archiving is suspended for the amount of time specified by the ARCHRETRYDELAY database configuration parameter. You can also use the NUMARCHRETRY database configuration parameter to specify the number of times that DB2 is to try archiving a log file to the primary or secondary archive directory before it tries to archive log files to the failover directory (specified by the FAILARCHPATH database configuration parameter). NUMARCHRETRY is

only used if the FAILARCHPATH database configuration parameter is set. If NUMARCHRETRY is set to 0, DB2 will continuously retry archiving from the primary or the secondary log path.

The easiest way to remove old log files is to restart the database. Once the database is restarted, only new log files and log files that the user exit program failed to archive will be found in the database directory.

When a database is restarted, the minimum number of logs in the database log directory will equal the number of primary logs which can be configured using the *logprimary* database configuration parameter. It is possible for more than the number of primary logs to be found in the log directory. This can occur if the number of empty logs in the log directory at the time the database was shut down, is greater than the value of the *logprimary* configuration parameter at the time the database is restarted. This will happen if the value of the *logprimary* configuration parameter is changed between the database being shut down and restarted, or if secondary logs are allocated and never used.

When a database is restarted, if the number of empty logs is less than the number of primary logs specified by the *logprimary* configuration parameter, additional log files will be allocated to make up the difference. If there are more empty logs than primary logs available in the database directory, the database can be restarted with as many available empty logs as are found in the database directory. After database shutdown, secondary log files that have been created will remain in the active log path when the database is restarted.

## Managing log files through log archiving

The following considerations apply to calling a user exit program for archiving and retrieving log files:

- Setting the database configuration parameter *logarchmeth1* to USEREXIT specifies whether the database manager invokes a user exit program to archive files or to retrieve log files during rollforward recovery of databases. A request to retrieve a log file is made when the rollforward utility needs a log file that is not found in the log path directory.

**Note:** On Windows<sup>®</sup> operating systems, you cannot use a REXX user exit to archive logs.

- When archiving, a log file is passed to the user exit when it is full, even if the log file is still active and is needed for normal processing. This allows copies of the data to be moved away from volatile media as quickly as possible. The log file passed to the user exit is retained in the log path directory until it is no longer needed for normal processing. At this point, the disk space is reused.
- DB2<sup>®</sup> opens a log file in read mode when it starts a user exit program to archive the file. On some platforms, this prevents the user exit program from being able to delete the log file. Other platforms, like AIX<sup>®</sup>, allow processes, including the user exit program, to delete log files. A user exit program should never delete a log file after it is archived, because the file could still be active and needed for crash recovery. DB2 manages disk space reuse when log files are archived.
- When a log file has been archived and it contains no open transactions, DB2 does not delete the file but renames it as the next log file when such a file is needed. This results in a performance gain, because creating a new log file

(instead of renaming the file) causes all pages to be written out to guarantee the disk space. It is more efficient to reuse than to free up and then reacquire the necessary pages on disk.

- DB2 will *not* retrieve log files during crash recovery or rollback unless the *logsecond* database configuration parameter is set to -1.
- Configuring log archiving does not guarantee rollforward recovery to the point of failure, but only attempts to make the failure window smaller. As log files fill, they are queued for the user exit routine. Should the disk containing the log fail before a log file is filled, the data in that log file is lost. Also, since the files are queued for archiving, the disk can fail before all the files are copied, causing any log files in the queue to be lost.
- The configured size of each individual log file has a direct bearing on log archiving. If each log file is very large, a large amount of data can be lost if a disk fails. A database configured with small log files causes the data to be passed to the user exit routine more often.

However, if you are moving the data to a slower device such as tape, you might want to have larger log files to prevent the queue from building up. Using larger log files is also recommended if archiving each file requires substantial overhead, such as rewinding the tape device or establishing a connection to the archive media.

- If *logarchmeth1* is set to USEREXIT, an archive request to the user exit program to occurs each time an active log file is filled. It is possible that an active log file is not full when the last disconnection from the database occurs and the user exit program is also called for a partially filled active log file.

**Note:** To free unused log space, the log file is truncated before it is archived.

- A copy of the log should be made to another physical device so that the offline log file can be used by rollforward recovery if the device containing the log file experiences a media failure. This should not be the same device containing database data files.
- If you have enabled user exit programs and are using a tape drive as a storage device for logs and backup images, you need to ensure that the destination for the backup images and the archived logs is not the same tape drive. Since some log archiving may take place while a backup operation is in progress, an error may occur when the two processes are trying to write to the same tape drive at the same time.
- Locally attached tape drives should not be used to manage log files with a user exit program if you are using any of the following:
  - infinite logging
  - online tablespace level recovery
  - replication
  - the Asynchronous Read Log API (db2ReadLog)
  - high availability disaster recovery (HADR)
- In some cases, if a database is closed before a positive response has been received from a user exit program for an archive request, the database manager will send another request when the database is opened. Thus, a log file may be archived more than once.
- If a user exit program receives a request to archive a file that does not exist (because there were multiple requests to archive and the file was deleted after the first successful archiving operation), or to retrieve a file that does not exist (because it is located in another directory or the end of the logs has been reached), it should ignore this request and pass a successful return code.



- The user exit program should allow for the existence of different log files with the same name after a point in time recovery; it should be written to preserve both log files and to associate those log files with the correct recovery path.
  - If a user exit program is enabled for two or more databases that are using the same tape device to archive log files, and a rollforward operation is taking place on one of the databases, the other database(s) should not be active. If another database tries to archive a log file while the rollforward operation is in progress, the logs required for the rollforward operation may not be found or the new log file archived to the tape device might overwrite the log files previously stored on that tape device.
- To prevent either situation from occurring, you can ensure that no other databases on the node that calls the user exit program are open during the rollforward operation, or write a user exit program to handle this situation.

**Related concepts:**

- “Managing log files” on page 42

## Blocking transactions when the log directory file is full

The *blk\_log\_dsk\_ful* database configuration parameter can be set to prevent “disk full” errors from being generated when DB2® cannot create a new log file in the active log path.

Instead, DB2 attempts to create the log file every five minutes until it succeeds. If a log archiving method is specified, DB2 also checks for the completion of log file archiving. If an archived log file is archived successfully, DB2 can rename the inactive log file to the new log file name and continue. After each attempt, DB2 writes a message to the administration notification log. The only way that you can confirm that your application is hanging because of a log disk full condition is to monitor the administration notification log.

Until the log file is successfully created, any user application that attempts to update table data will not be able to commit transactions. Read-only queries may not be directly affected; however, if a query needs to access data that is locked by an update request, or a data page that is fixed in the buffer pool by the updating application, read-only queries will also appear to hang.

**Related concepts:**

- “Understanding recovery logs” on page 31
- “Managing log files through log archiving” on page 45

## On demand log archive

DB2® supports the closing (and, if enabled, the archiving) of the active log for a recoverable database at any time. This allows you to collect a complete set of log files up to a known point, and then to use these log files to update a standby database.

You can initiate on demand log archiving by invoking the ARCHIVE LOG command, or by calling the **db2ArchiveLog** API.

**Related concepts:**

- “Managing log files through log archiving” on page 45

**Related reference:**

- “ARCHIVE LOG” on page 255
- “Configuration parameters for database logging” on page 35
- “db2ArchiveLog - Archive Active Log” on page 265

## Using raw logs

You can use a raw device for your database log. There are both advantages and disadvantages in doing so.

- The advantages are:
  - You can attach more than 26 physical drives to a system.
  - The file I/O path length is shorter. This may improve performance on your system. You should conduct benchmarks to evaluate if there are measurable benefits for your work load.
- The disadvantages are:
  - The device cannot be shared by other applications; the entire device *must* be assigned to DB2®.
  - The device cannot be operated upon by any operating system utility or third-party tool which would backup or copy from the device.
  - You can easily wipe out the file system on an existing drive if you specify the wrong physical drive number.

You can configure a raw log with the *newlogpath* database configuration parameter. Before doing so, however, consider the advantages and disadvantages listed above, and the additional considerations listed below:

- Only one device is allowed. You can define the device over multiple disks at the operating system level. DB2 will make an operating system call to determine the size of the device in 4-KB pages.

If you use multiple disks, this will provide a larger device, and the striping that results can improve performance by faster I/O throughput.
- DB2 will attempt to write to the last 4-KB page of the device. If the device size is greater than 2 GB, the attempt to write to the last page will fail on operating systems that do not provide support for devices larger than 2 GB. In this situation, DB2 will attempt to use all pages, up to the supported limit.

Information about the size of the device is used to indicate the size of the device (in 4-KB pages) available to DB2 under the support of the operating system. The amount of disk space that DB2 can write to is referred to as the *device-size-available*.

The first 4-KB page of the device is not used by DB2 (this space is generally used by operating system for other purposes.) This means that the total space available to DB2 is  $device-size = device-size-available - 1$ .
- The *logsecond* parameter is not used. DB2 will not allocate secondary logs. The size of active log space is the number of 4-KB pages that result from  $logprimary \times logfilsiz$ .
- Log records are still grouped into log extents, each with a log file size (*logfilsiz*) of 4-KB pages. Log extents are placed in the raw device, one after another. Each extent also consists of an extra two pages for the extent header. This means that the *number of available log extents* the device can support is  $device-size / (logfilsiz + 2)$
- The device must be large enough to support the active log space. That is, the *number of available log extents* must be greater than (or equal to) the value

specified for the *logprimary* configuration parameter. If log archiving is enabled (through the *logarchmeth1* and *logarchmeth2* configuration parameters), ensure that the raw device can contain more logs than the value specified for the *logprimary* configuration parameter. This will compensate for the delay incurred when a log file is being archived.

- If you are using circular logging, the *logprimary* configuration parameter will determine the number of log extents that are written to the device. This may result in unused space on the device.
- If *logarchmeth1* is set to LOGRETAIN, after the *number of available log extents* are all used up, all operations that result in an update will receive a log full error. At this time, you must shut down the database and take an offline backup of it to ensure recoverability. After the database backup operation, the log records written to the device are lost. This means that you cannot use an earlier database backup image to restore the database, then roll it forward. If you take a database backup before the *number of available log extents* are all used up, you can restore and roll the database forward.
- If *logarchmeth1* is set to USEREXIT, the user exit program must be able to read the device and store the archived log as a file.
- DB2 will not retrieve log files from a raw device. Instead, when a log file is needed, DB2 will read the extent headers to determine if the raw device contains the required log file. If the required log file is not found in the raw device, DB2 will search the overflow log path. If the log file is still not found, DB2 will retrieve the log file and place it in the overflow log path. If an overflow log path has not been specified, DB2 will not attempt to retrieve the log file.
- If you have configured a raw device for logging, and are using DataPropagator™ (DPROP), or another application that calls the **db2ReadLog** API, the *overflowlogpath* database configuration parameter must be configured. DB2 may call a user exit program to retrieve the log file and return the log data requested by the **db2ReadLog** API. The retrieved log file will be placed in the path specified by the *overflowlogpath* database configuration parameter.

**Related tasks:**

- “Specifying raw I/O” in the *Administration Guide: Implementation*

**Related reference:**

- “db2ReadLog - Asynchronous Read Log” on page 286
- Appendix F, “Tivoli Storage Manager,” on page 325

## Including log files with a backup image

When performing an online backup operation, you can specify that the log files required to restore and recover a database are included in the backup image. This means that if you need to ship backup images to a disaster recovery site, you do not have to send the log files separately or package them together yourself. Further, you do not have to decide which log files are required to guarantee the consistency of an online backup. This provides some protection against the deletion of log files required for successful recovery.

To make use of this feature specify the INCLUDE LOGS option of the BACKUP DATABASE command. When you specify this option, the backup utility will truncate the currently active log file and copy the necessary set of log extents into the backup image.

To restore the log files from a backup image, use the LOGTARGET option of the RESTORE DATABASE command and specify a fully qualified path that exists on the DB2® server. The restore database utility will then write the log files from the image to the target path. If a log file with the same name already exists in the target path, the restore operation will fail and an error will be returned. If the LOGTARGET option is not specified, no log files will be restored from the backup image.

If the LOGTARGET option is specified and the backup image does not include any log files, an error will be returned before an attempt is made to restore any table space data. The restore operation will also fail if an invalid or read-only path is specified. During a database or table space restore where the LOGTARGET option is specified, if one or more log files cannot be extracted, the restore operation fails and an error is returned.

You can also choose to restore only the log files saved in the backup image. To do this, specify the LOGS option with the LOGTARGET option of the RESTORE DATABASE command. If the restore operation encounters any problems when restoring log files in this mode, the restore operation fails and an error is returned.

During an automatic incremental restore operation, only the logs included in the target image of the restore operation will be retrieved from the backup image. Any logs that are included in intermediate images referenced during the incremental restore process will not be extracted from those backup images. During a manual incremental restore, if you specify a log target directory when restoring a backup image that includes log files, the log files in that backup image will be restored.

**Notes:**

1. This feature is only available on single-partition databases.
2. This feature is not supported for offline backups.
3. When logs are included in an online backup image, the resulting image cannot be restored on releases of DB2 Universal Database™ (UDB) prior to Version 8.2.

**Related tasks:**

- “Restoring from incremental backup images” on page 26

**Related reference:**

- “BACKUP DATABASE” on page 66
- “RESTORE DATABASE” on page 87

## How to prevent losing log files

In situations where you need to drop a database or perform a point-in-time rollforward recovery, it is possible to lose log files that may be required for future recovery operations. In these cases, it is important to make copies of all the logs in the current database log path directory. Consider the following scenarios:

- If you plan to drop a database prior to a restore operation, you need to save the log files in the active log path before issuing the DROP DATABASE command. After the database has been restored, these log files may be required for rollforward recovery because some of them may not have been archived before the database was dropped. Normally, you are not required to drop a database prior to issuing the RESTORE command. However, you may have to drop the database (or drop the database on one partition by specifying the AT NODE option of DROP DATABASE command), because it has been damaged to the

extent that the RESTORE command fails. You may also decide to drop a database prior to the restore operation to give yourself a fresh start.

- If you are rolling a database forward to a specific point in time, log data after the time stamp you specify will be overwritten. If, after you have completed the point-in-time rollforward operation and reconnected to the database, you determine that you actually needed to roll the database forward to a later point in time, you will not be able to because the logs will already have been overwritten. It is possible that the original set of log files may have been archived; however, DB2® may be calling a user exit program to automatically archive the newly generated log files. Depending on how the user exit program is written, this could cause the original set of log files in the archive log directory to be overwritten. Even if both the original and new set of log files exist in the archive log directory (as different versions of the same files), you may have to determine which set of logs should be used for future recovery operations.

**Related concepts:**

- “Understanding recovery logs” on page 31

---

## Understanding the recovery history file

A recovery history file is created with each database and is automatically updated whenever:

- A database or table spaces are backed up
- A database or table spaces are restored
- A database or table spaces are rolled forward
- A table space is created
- A table space is altered
- A table space is quiesced
- A table space is renamed
- A table space is dropped
- A table is loaded
- A table is dropped (when dropped table recovery is enabled)
- A table is reorganized
- On-demand log archiving is invoked
- A new log file is written to (when using recoverable logging)
- A log file is archived (when using recoverable logging)
- A database is recovered

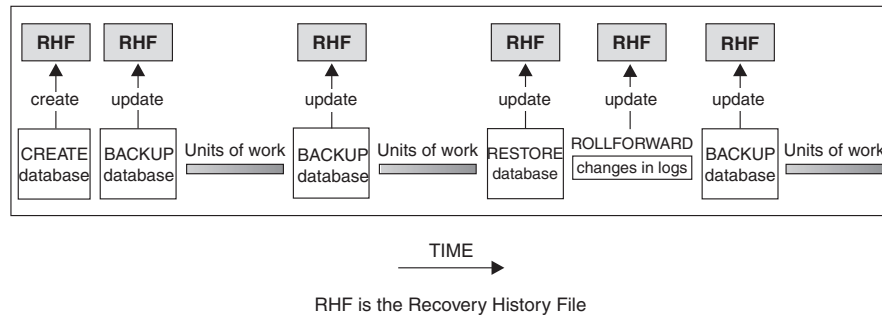


Figure 9. Creating and Updating the Recovery History File

You can use the summarized backup information in this file to recover all or part of a database to a given point in time. The information in the file includes:

- An identification (ID) field to uniquely identify each entry
- The part of the database that was copied and how
- The time the copy was made
- The location of the copy (stating both the device information and the logical way to access the copy)
- The last time a restore operation was done
- The time at which a table space was renamed, showing the previous and the current name of the table space
- The status of a backup operation: active, inactive, expired, or deleted
- The last log sequence number saved by the database backup or processed during a rollforward recovery operation.

To see the entries in the recovery history file, use the LIST HISTORY command.

Every backup operation (database, table space, or incremental) includes a copy of the recovery history file. The recovery history file is associated with the database. Dropping a database deletes the recovery history file. Restoring a database to a new location restores the recovery history file. Restoring does not overwrite the existing recovery history file unless the file that exists on disk has no entries. If that is the case, the database history will be restored from the backup image.

If the current database is unusable or not available, and the associated recovery history file is damaged or deleted, an option on the RESTORE command allows only the recovery history file to be restored. The recovery history file can then be reviewed to provide information on which backup to use to restore the database.

The size of the file is controlled by the `rec_his_retentn` configuration parameter that specifies a retention period (in days) for the entries in the file. Even if the number for this parameter is set to zero (0), the most recent full database backup (plus its restore set) is kept. (The only way to remove this copy is to use the PRUNE with FORCE option.) The retention period has a default value of 366 days. The period can be set to an indefinite number of days by using -1. In this case, explicit pruning of the file is required.

**Related reference:**

- “`rec_his_retentn` - Recovery history retention period configuration parameter” in the *Administration Guide: Performance*
- “LIST HISTORY” on page 258

## Recovery history file - garbage collection

### Garbage Collection

Although you can use the PRUNE HISTORY command at any time to remove entries from the history file, it is recommended that such pruning be left to DB2®. The number of DB2 database backups recorded in the recovery history file is monitored automatically by DB2 *garbage collection*. DB2 garbage collection is invoked:

- After a full, non-incremental database backup operation completes successfully.
- After a database restore operation, where a rollforward operation is not required, completes successfully.
- After a successful database rollforward operation completes successfully.

The configuration parameter *num\_db\_backups* defines how many active full (non-incremental) database backup images are kept. The value of this parameter is used to scan the history file, starting with the last entry.

After every full (non-incremental) database backup operation, the *rec\_his\_retentn* configuration parameter is used to prune expired entries from the history file.

An *active database backup* is one that can be restored and rolled forward using the current logs to recover the current state of the database. An *inactive database backup* is one that, if restored, moves the database back to a previous state.




Figure 10. Active Database Backups. The value of *num\_db\_backups* has been set to four.

All active database backup images that are no longer needed are marked as “expired”. These images are considered to be unnecessary, because more recent backup images are available. All table space backup images and load backup copies that were taken before the database backup image expired are also marked as “expired”.

All database backup images that are marked as “inactive” and that were taken prior to the point at which an expired database backup was taken are also marked as “expired”. All associated inactive table space backup images and load backup copies are also marked as “expired”.

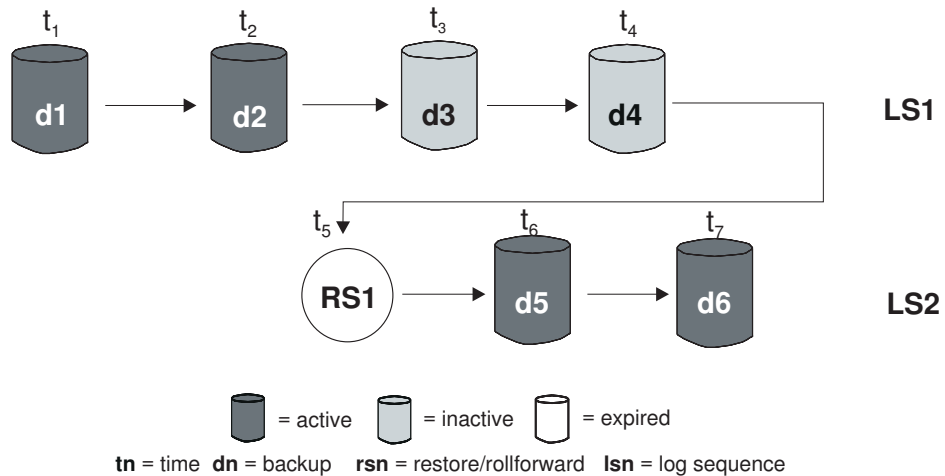


Figure 11. Inactive Database Backups

If an active database backup image is restored, but it is not the most recent database backup recorded in the history file, any subsequent database backup images belonging to the same log sequence are marked as “inactive”.

If an inactive database backup image is restored, any inactive database backups belonging to the current log sequence are marked as “active” again. All active database backup images that are no longer in the current log sequence are marked as “inactive”.

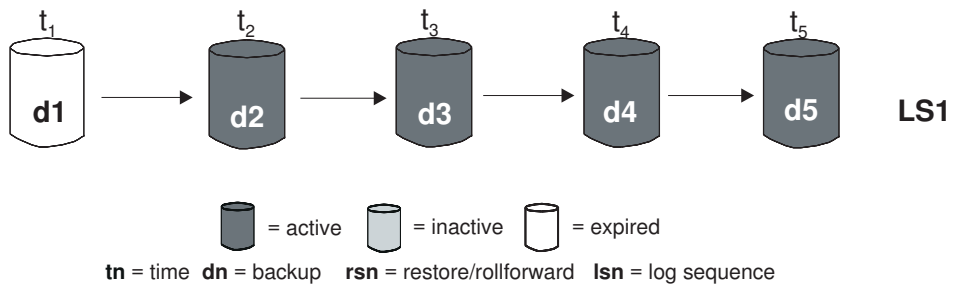


Figure 12. Expired Database Backups

DB2 garbage collection is also responsible for marking the history file entries for a DB2 database or table space backup image as “inactive”, if that backup does not correspond to the current *log sequence*, also called the *current log chain*. The current log sequence is determined by the DB2 database backup image that has been restored, and the log files that have been processed. Once a database backup image is restored, all subsequent database backup images become “inactive”, because the restored image begins a new log chain. (This is true if the backup image was restored without rolling forward. If a rollforward operation has occurred, all database backups that were taken after the break in the log chain are marked as “inactive”. It is conceivable that an older database backup image will have to be restored because the rollforward utility has gone through the log sequence containing a damaged current backup image.)

A table space-level backup image becomes “inactive” if, after it is restored, the current state of the database cannot be reached by applying the current log sequence.



If a backup image contains DATALINK columns, all Data Links servers running the DB2 Data Links Manager are contacted to request garbage collection. DB2 garbage collection then deletes backups of the associated Data Links server files that were contained in the expired backup, but that were unlinked before the next database backup operation.

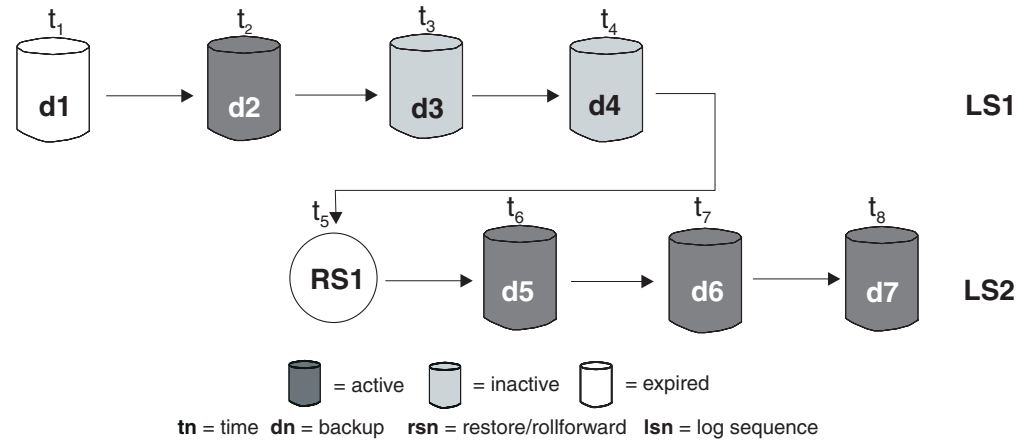


Figure 13. Mixed Active, Inactive, and Expired Database Backups

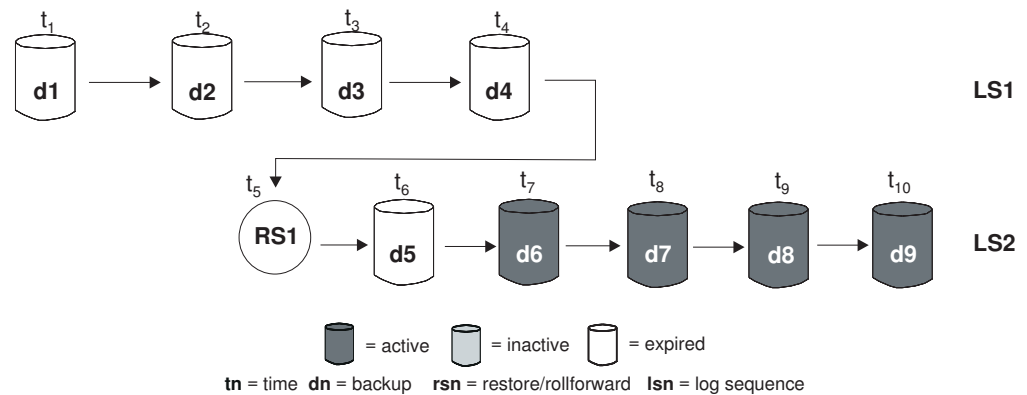


Figure 14. Expired Log Sequence

**Related concepts:**

- “Understanding the recovery history file” on page 51

**Related reference:**

- “PRUNE HISTORY/LOGFILE” on page 260

---

## Understanding table space states

The current status of a table space is reflected by its *state*. The table space states most commonly associated with recovery are:

- *Backup pending*. A table space is put in this state after a point-in-time rollforward operation, or after a load operation with the no copy option. The table space must be backed up before it can be used. (If it is not backed up, the table space cannot be updated, but read-only operations are allowed.)

- *Restore pending.* A table space is put in this state if a rollforward operation on that table space is cancelled, or if a rollforward operation on that table space encounters an unrecoverable error, in which case the table space must be restored and rolled forward again.
- *Rollforward-in-progress.* A table space is put in this state when a rollforward operation on that table space is in progress. Once the rollforward operation completes successfully, the table space is no longer in rollforward-in-progress state. The table space can also be taken out of this state if the rollforward operation is cancelled.
- *Rollforward pending.* A table space is put in this state after it is restored, or following an input/output (I/O) error. After it is restored, the table space can be rolled forward to the end of the logs or to a point in time. Following an I/O error, the table space must be rolled forward to the end of the logs.

---

## Enhancing recovery performance

The following should be considered when thinking about recovery performance:

- You can improve performance for databases that are frequently updated by placing the logs on a separate device. In the case of an online transaction processing (OLTP) environment, often more I/O is needed to write data to the logs than to store a row of data. Placing the logs on a separate device will minimize the disk arm movement that is required to move between a log and the database files.

You should also consider what other files are on the disk. For example, moving the logs to the disk used for system paging in a system that has insufficient real memory will defeat your tuning efforts.

DB2<sup>®</sup> automatically attempts to minimize the time it takes to complete a backup or restore operation by choosing an optimal value for the number of buffers, the buffer size and the parallelism settings. The values are based on the amount of utility heap memory available, the number of processors available and the database configuration.

- To reduce the amount of time required to complete a restore operation:
  - Adjust the restore buffer size. The buffer size must be a multiple of the buffer size that was used during the backup operation.
  - Increase the number of buffers.

If you use multiple buffers and I/O media devices, you should use at least twice as many buffers as media devices to ensure that they do not have to wait for data. The size of the buffers used will also contribute to the performance of the restore operation. The ideal restore buffer size should be a multiple of the extent size for the table spaces.

If you have multiple table spaces with different extent sizes, specify a value that is a multiple of the largest extent size.

The *minimum* recommended number of buffers is the number of media devices plus the number specified for the PARALLELISM option.

- Use multiple source devices.
- Set the PARALLELISM option for the restore operation to be at least one (1) greater than the number of source devices.
- If a table contains large amounts of long field and LOB data, restoring it could be very time consuming. If the database is enabled for rollforward recovery, the RESTORE command provides the capability to restore selected table spaces. If the long field and LOB data is critical to your business, restoring these table spaces should be considered against the time required to complete the backup

task for these table spaces. By storing long field and LOB data in separate table spaces, the time required to complete the restore operation can be reduced by choosing not to restore the table spaces containing the long field and LOB data. If the LOB data can be reproduced from a separate source, choose the NOT LOGGED option when creating or altering a table to include LOB columns. If you choose not to restore the table spaces that contain long field and LOB data, but you need to restore the table spaces that contain the table, you must roll forward to the end of the logs so that all table spaces that contain table data are consistent.

**Note:** If you back up a table space that contains table data without the associated long or LOB fields, you cannot perform point-in-time rollforward recovery on that table space. All the table spaces for a table must be rolled forward simultaneously to the same point in time.

- The following apply for both backup and restore operations:
  - Multiple I/O buffers and devices should be used.
  - Allocate at least twice as many buffers as devices being used.
  - Do not overload the I/O device controller bandwidth.
  - Use more buffers of smaller size rather than a few large buffers.
  - Tune the number and the size of the buffers according to the system resources.
  - Use of the PARALLELISM option
- DB2<sup>®</sup> uses multiple agents to perform both crash recovery and database rollforward recovery. You can expect better performance during these operations, particularly on symmetric multi-processor (SMP) machines; using multiple agents during database recovery takes advantage of the extra CPUs that are available on SMP machines.

The agent type introduced by parallel recovery is db2agnsc. DB2 chooses the number of agents to be used for database recovery based on the number of CPUs on the machine.

DB2 distributes log records to these agents so that they can be reapplied concurrently, where appropriate. For example, the processing of log records associated with insert, delete, update, add key, and delete key operations can be parallelized in this way. Because the log records are parallelized at the page level (log records on the same data page are processed by the same agent), performance is enhanced, even if all the work was done on one table.

**Related concepts:**

- “Optimizing backup performance” on page 79
- “Restore overview” on page 81



---

## Chapter 2. Database backup

This section describes the DB2 UDB backup utility, which is used to create backup copies of a database or table spaces.

The following topics are covered:

- “Backup overview”
- “Privileges, authorities, and authorization required to use backup” on page 62
- “Using backup” on page 62
- “Backing up to tape” on page 64
- “Backing up to named pipes” on page 65
- “BACKUP DATABASE” on page 66
- “db2Backup - Backup database” on page 71
- “Backup sessions - CLP examples” on page 78
- “Optimizing backup performance” on page 79

---

### Backup overview

The simplest form of the DB2<sup>®</sup> BACKUP DATABASE command requires only that you specify the alias name of the database that you want to back up. For example:

```
db2 backup db sample
```

If the command completes successfully, you will have acquired a new backup image that is located in the path or the directory from which the command was issued. It is located in this directory because the command in this example does not explicitly specify a target location for the backup image. On Windows<sup>®</sup> operating systems, for example, this command (when issued from the root directory) creates an image that appears in a directory listing as follows:

```
Directory of D:\SAMPLE.0\DB2\NODE0000\CATN0000\20010320
```

```
03/20/2001 12:26p      <DIR>      .
03/20/2001 12:26p      <DIR>      ..
03/20/2001 12:27p              12,615,680 122644.001
```

**Note:** If the DB2 client and server are not located on the same system, DB2 will determine which directory is the current working directory on the client machine and use that as the backup target directory on the server. For this reason, it is recommended that you specify a target directory for the backup image.

Backup images are created at the target location specified when you invoke the backup utility. This location can be:

- A directory (for backups to disk or diskette)
- A device (for backups to tape)
- A Tivoli<sup>®</sup> Storage Manager (TSM) server
- Another vendor’s server

The recovery history file is updated automatically with summary information whenever you invoke a database backup operation. This file is created in the same directory as the database configuration file.

On UNIX<sup>®</sup> based systems, file names for backup images created on disk consist of a concatenation of several elements, separated by periods:

```
DB_alias.Type.Inst_name.NODEnnnn.CATNnnnn.timestamp.Seq_num
```

For example:

```
STAFF.0.DB201.NODE0000.CATN0000.19950922120112.001
```

On Windows operating systems, a four-level subdirectory tree is used:

```
DB_alias.Type\Inst_name\NODEnnnn\CATNnnnn\yyyymmdd\hhmmss.Seq_num
```

For example:

```
SAMPLE.0\DB2\NODE0000\CATN0000\20010320\122644.001
```

<b>Database alias</b>	A 1- to 8-character database alias name that was specified when the backup utility was invoked.
<b>Type</b>	Type of backup operation, where: 0 represents a full database-level backup, 3 represents a table space-level backup, and 4 represents a backup image generated by the LOAD...COPY TO command.
<b>Instance name</b>	A 1- to 8-character name of the current instance that is taken from the <b>DB2INSTANCE</b> environment variable.
<b>Node number</b>	The node number. In non-partitioned database systems, this is always NODE0000. In partitioned database systems, it is NODExxxx, where xxxx is the number assigned to the node in the db2nodes.cfg file.
<b>Catalog node number</b>	The node number of the catalog node for the database. In non-partitioned database systems, this is always CATN0000. In partitioned database systems, it is CATNxxxx, where xxxx is the number assigned to the node in the db2nodes.cfg file.
<b>Time stamp</b>	A 14-character representation of the date and time at which the backup operation was performed. The time stamp is in the form <i>yyyymmddhhmmss</i> , where: <ul style="list-style-type: none"><li>• <i>yyyy</i> represents the year (1995 to 9999)</li><li>• <i>mm</i> represents the month (01 to 12)</li><li>• <i>dd</i> represents the day of the month (01 to 31)</li><li>• <i>hh</i> represents the hour (00 to 23)</li><li>• <i>nn</i> represents the minutes (00 to 59)</li><li>• <i>ss</i> represents the seconds (00 to 59)</li></ul>
<b>Sequence number</b>	A 3-digit number used as a file extension.

When a backup image is written to tape:

- File names are not created, but the information described above is stored in the backup header for verification purposes.

- A tape device must be available through the standard operating system interface. On a large partitioned database system, however, it may not be practical to have a tape device dedicated to each database partition server. You can connect the tape devices to one or more TSM servers, so that access to these tape devices is provided to each database partition server.
- On a partitioned database system, you can also use products that provide virtual tape device functions, such as REELlibrarian 4.2 or CLIO/S. You can use these products to access the tape device connected to other nodes (database partition servers) through a pseudo tape device. Access to the remote tape device is provided transparently, and the pseudo tape device can be accessed through the standard operating system interface.

You cannot back up a database that is in an unusable state, except when that database is in backup pending state. If any table space is in an abnormal state, you cannot back up the database or that table space, unless it is in backup pending state.

Concurrent backup operations on the same table space are not permitted. Once a backup operation has been initiated on a table space, any subsequent attempts will fail (SQL2048).

If a database or a table space is in a partially restored state because a system crash occurred during the restore operation, you must successfully restore the database or the table space before you can back it up.

A backup operation will fail if a list of the table spaces to be backed up contains the name of a temporary table space.

The backup utility provides concurrency control for multiple processes that are making backup copies of different databases. This concurrency control keeps the backup target devices open until all the backup operations have ended. If an error occurs during a backup operation, and an open container cannot be closed, other backup operations targeting the same drive may receive access errors. To correct such access errors, you must terminate the backup operation that caused the error and disconnect from the target device. If you are using the backup utility for concurrent backup operations to tape, ensure that the processes do not target the same tape.

## Displaying backup information

You can use **db2ckbkp** to display information about existing backup images. This utility allows you to:

- Test the integrity of a backup image and determine whether or not it can be restored.
- Display information that is stored in the backup header.
- Display information about the objects and the log file header in the backup image.

### Related concepts:

- “Developing a backup and recovery strategy” on page 3
- “Understanding the recovery history file” on page 51
- “Including log files with a backup image” on page 49
- “Automatic maintenance” in the *Administration Guide: Planning*

**Related reference:**

- Appendix F, “Tivoli Storage Manager,” on page 325

---

## Privileges, authorities, and authorization required to use backup

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges and higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects. Users can access only those objects for which they have the appropriate authorization; that is, the required privilege or authority.

You must have SYSADM, SYSCTRL, or SYSMAINT authority to use the backup utility.

---

## Using backup

**Prerequisites:**

You should not be connected to the database that is to be backed up: the backup database utility automatically establishes a connection to the specified database, and this connection is terminated at the completion of the backup operation. If you are connected to a database that is to be backup up, you will be disconnected when the BACKUP DATABASE command is issued and the backup operation will proceed.

The database can be local or remote. The backup image remains on the database server, unless you are using a storage management product such as Tivoli Storage Manager (TSM).

On a partitioned database system, database partitions are backed up individually. The operation is local to the database partition server on which you invoke the utility. You can, however, issue **db2\_all** from one of the database partition servers in the instance to invoke the backup utility on a list of servers, which you identify by node number. (Use the LIST NODES command to identify the nodes, or database partition servers, that have user tables on them.) If you do this, you must back up the catalog node first, then back up the other database partitions. You can also use the Command Editor to back up database partitions. Because this approach does not support rollforward recovery, back up the database residing on these nodes regularly. You should also keep a copy of the db2nodes.cfg file with any backup copies you take, as protection against possible damage to this file.

On a distributed request system, backup operations apply to the distributed request database and the metadata stored in the database catalog (wrappers, servers, nicknames, and so on). Data source objects (tables and views) are not backed up, unless they are stored in the distributed request database.

If a database was created with a previous release of the database manager, and the database has not been migrated, you must migrate the database before you can back it up.

**Restrictions:**

The following restrictions apply to the backup utility:



- A table space backup operation and a table space restore operation cannot be run at the same time, even if different table spaces are involved.
- If you want to be able to do rollforward recovery in a partitioned database environment, you must regularly back up the database on the list of nodes, and you must have at least one backup image of the rest of the nodes in the system (even those that do not contain user data for that database). Two situations require the backed-up image of a database partition at a database partition server that does not contain user data for the database:
  - You added a database partition server to the database system after taking the last backup, and you need to do forward recovery on this database partition server.
  - Point-in-time recovery is used, which requires that all database partitions in the system are in rollforward pending state.
  - Online backup operations for DMS table spaces are incompatible with the following operations:
    - load
    - reorganization (online and offline)
    - drop table space
    - table truncation
    - index creation
    - not logged initially (used with the CREATE TABLE and ALTER TABLE statements)

**Procedure:**

The backup utility can be invoked through the command line processor (CLP), the Backup Database notebook or Wizard in the Control Center, or the **db2Backup** application programming interface (API).

Following is an example of the BACKUP DATABASE command issued through the CLP:

```
db2 backup database sample to c:\DB2Backups
```

To open the Backup Database notebook or wizard:

1. From the Control Center, expand the object tree until you find the Databases folder.
2. Click on the Databases folder. Any existing databases are displayed in the pane on the right side of the window (the contents pane).
3. Click the right mouse button on the database you want in the contents pane, and select Backup Database or Backup Database Using Wizard from the pop-up menu. The Backup Database notebook or the Backup Database wizard opens.

Detailed information is provided through the online help facility within the Control Center.

**Related concepts:**

- “Administrative APIs in Embedded SQL or DB2 CLI Programs” in the *Application Development Guide: Programming Client Applications*
- “Introducing the plug-in architecture for the Control Center” in the *Administration Guide: Implementation*

**Related tasks:**

- “Migrating databases” in the *Quick Beginnings for DB2 Servers*

**Related reference:**

- “LIST DBPARTITIONNUMS Command” in the *Command Reference*
- “db2Backup - Backup database” on page 71

## Backing up to tape

When you back up your database or table space, you must correctly set your block size and your buffer size. This is particularly true if you are using a variable block size (on AIX®, for example, if the block size has been set to zero).

There is a restriction on the number of fixed block sizes that can be used when backing up. This restriction exists because DB2® writes out the backup image header as a 4-KB block. The only fixed block sizes DB2 supports are 512, 1024, 2048, and 4096 bytes. If you are using a fixed block size, you can specify any backup buffer size. However, you may find that your backup operation will not complete successfully if the fixed block size is not one of the sizes that DB2 supports.

If your database is large, using a fixed block size means that your backup operations might take a long time to complete. You may want to consider using a variable block size.

**Note:** Use of a variable block size is currently *not* supported. If you must use this option, ensure that you have well tested procedures in place that enable you to recover successfully, using backup images that were created with a variable block size.

When using a variable block size, you must specify a backup buffer size that is less than or equal to the maximum limit for the tape devices that you are using. For optimal performance, the buffer size must be equal to the maximum block size limit of the device being used.

Before a tape device can be used on a Windows® operating system, the following command must be issued:

```
db2 initialize tape on <device> using <blksize>
```

Where:

**<device>**

is a valid tape device name. The default on Windows operating systems is \\.\TAPE0.

**<blksize>**

is the blocking factor for the tape. It must be a factor or multiple of 4096. The default value is the default block size for the device.

Restoring from a backup image with variable block size may return an error. If this happens, you may need to rewrite the image using an appropriate block size. Following is an example on AIX:

```
tctl -b 0 -Bn -f /dev/rmt0 read > backup_filename.file
dd if=backup_filename.file of=/dev/rmt0 obs=4096 conv=sync
```

The backup image is dumped to a file called backup\_filename.file. The **dd** command dumps the image back onto tape, using a block size of 4096.

There is a problem with this approach if the image is too large to dump to a file. One possible solution is to use the **dd** command to dump the image from one tape device to another. This will work as long as the image does not span more than one tape. When using two tape devices, the **dd** command is:

```
dd if=/dev/rmt1 of=/dev/rmt0 obs=4096
```

If using two tape devices is not possible, you may be able to dump the image to a raw device using the **dd** command, and then to dump the image from the raw device to tape. The problem with this approach is that the **dd** command *must* keep track of the number of blocks dumped to the raw device. This number must be specified when the image is moved back to tape. If the **dd** command is used to dump the image from the raw device to tape, the command dumps the entire contents of the raw device to tape. The **dd** utility cannot determine how much of the raw device is used to hold the image.

When using the backup utility, you will need to know the maximum block size limit for your tape devices. Here are some examples:

Device	Attachment	Block Size Limit	DB2 Buffer Size Limit (in 4-KB pages)
8 mm	scsi	131,072	32
3420	s370	65,536	16
3480	s370	65,536	16
3490	s370	65,536	16
3490E	s370	65,536	16
7332 (4 mm) <sup>1</sup>	scsi	262,144	64
3490e	scsi	262,144	64
3590 <sup>2</sup>	scsi	2,097,152	512
3570 (magstar MP)		262,144	64

**Notes:**

1. The 7332 does not implement a block size limit. 256 KB is simply a suggested value. Block size limit is imposed by the parent adapter.
2. While the 3590 does support a 2-MB block size, you could experiment with lower values (like 256 KB), provided the performance is adequate for your needs.
3. For information about your device limit, check your device documentation or consult with the device vendor.

---

## Backing up to named pipes

Support is now available for database backup to (and database restore from) local named pipes on UNIX based systems.

**Prerequisites:**

Both the writer and the reader of the named pipe must be on the same machine. The pipe must exist and be located on a local file system. Because the named pipe is treated as a local device, there is no need to specify that the target is a named pipe.

**Procedure:**

Following is an AIX example:

1. Create a named pipe:  
`mkfifo /u/dmcinnis/mypipe`
2. If this backup image is going to be used by the restore utility, the restore operation must be invoked *before* the backup operation, so that it will not miss any data:  
`db2 restore db sample into mynewdb from /u/dmcinnis/mypipe`
3. Use this pipe as the target for a database backup operation:  
`db2 backup db sample to /u/dmcinnis/mypipe`

**Related tasks:**

- “Using backup” on page 62

**Related reference:**

- “BACKUP DATABASE” on page 66
- “RESTORE DATABASE” on page 87

---

## BACKUP DATABASE

Creates a backup copy of a database or a table space.

**Scope:**

This command only affects the database partition on which it is executed.

**Authorization:**

One of the following:


- *sysadm*
- *sysctrl*
- *sysmaint*

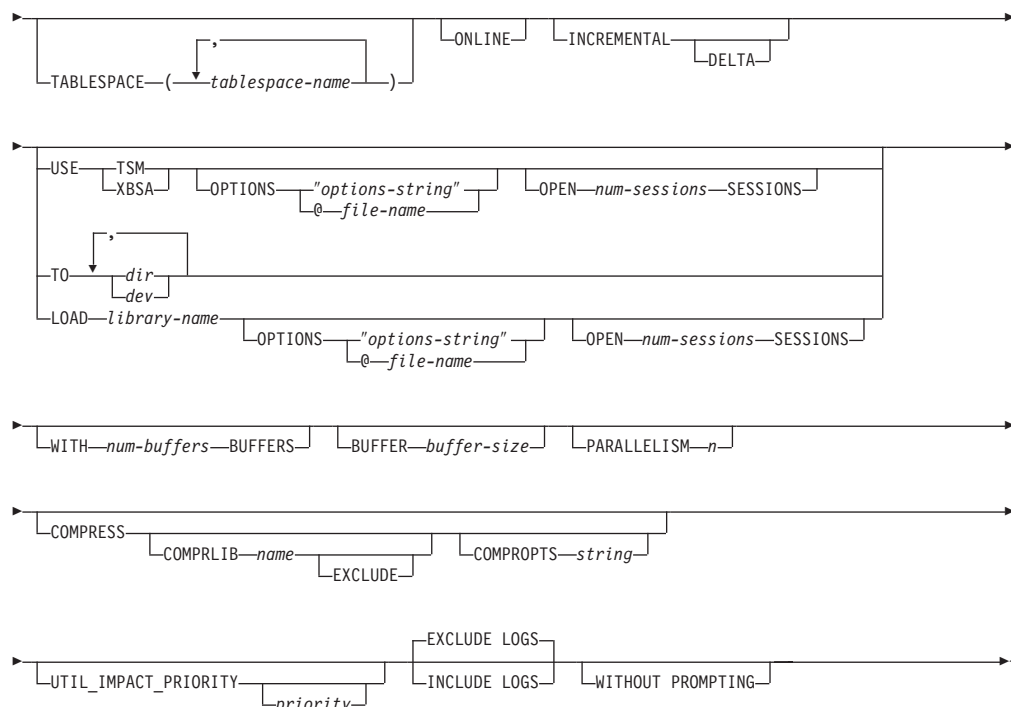
**Required connection:**

Database. This command automatically establishes a connection to the specified database.

**Note:** If a connection to the specified database already exists, that connection will be terminated and a new connection established specifically for the backup operation. The connection is terminated at the completion of the backup operation.

**Command syntax:**





**Command parameters:**

**DATABASE database-alias**

Specifies the alias of the database to back up.

**USER username**

Identifies the user name under which to back up the database.

**USING password**

The password used to authenticate the user name. If the password is omitted, the user is prompted to enter it.

**TABLESPACE tablespace-name**

A list of names used to specify the table spaces to be backed up.

**ONLINE**

Specifies online backup. The default is offline 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 and S (Share) locks on LOB data in SMS table spaces.

**INCREMENTAL**

Specifies a cumulative (incremental) backup image. An incremental backup image is a copy of all database data that has changed since the most recent successful, full backup operation.

**DELTA**

Specifies a non-cumulative (delta) backup image. A delta backup image is a copy of all database data that has changed since the most recent successful backup operation of any type.

**USE TSM**

Specifies that the backup is to use Tivoli Storage Manager output.

**USE XBSA**

Specifies that the XBSA interface is to be used. Backup Services APIs

## BACKUP DATABASE

(XBSA) are an open application programming interface for applications or facilities needing data storage management for backup or archiving purposes.

### OPTIONS

*"options-string"*

Specifies options to be used for the backup operation. The string will be passed to the vendor support library, for example TSM, exactly as it was entered, without the quotes.

**Note:** Specifying this option overrides the value specified by the VENDOROPT database configuration parameter.

*@file-name*

Specifies that the options to be used for the backup operation are contained in a file located on the DB2 server. The string will be passed to the vendor support library, for example TSM. The file must be a fully qualified file name.

### OPEN num-sessions SESSIONS

The number of I/O sessions to be created between DB2 and TSM or another backup vendor product.

**Note:** This parameter has no effect when backing up to tape, disk, or other local device.

### TO dir/dev

A list of directory or tape device names. The full path on which the directory resides must be specified. If USE TSM, TO, and LOAD are omitted, the default target directory for the backup image is the current working directory of the client computer. This target directory or device must exist on the database server. This parameter may be repeated to specify the target directories and devices that the backup image will span. If more than one target is specified (target1, target2, and target3, for example), target1 will be opened first. The media header and special files (including the configuration file, table space table, and history file) are placed in target1. All remaining targets are opened, and are then used in parallel during the backup operation. Because there is no general tape support on Windows operating systems, each type of tape device requires a unique device driver. To back up to the FAT file system on Windows operating systems, users must conform to the 8.3 naming restriction.

Use of tape devices or floppy disks may generate messages and prompts for user input. Valid response options are:

- c** Continue. Continue using the device that generated the warning message (for example, when a new tape has been mounted)
- d** Device terminate. Stop using *only* the device that generated the warning message (for example, when there are no more tapes)
- t** Terminate. Abort the backup operation.

If the tape system does not support the ability to uniquely reference a backup image, it is recommended that multiple backup copies of the same database not be kept on the same tape.

### LOAD library-name

The name of the shared library (DLL on Windows operating systems) containing the vendor backup and restore I/O functions to be used. It can

contain the full path. If the full path is not given, it will default to the path on which the user exit program resides.

**WITH num-buffers BUFFERS**

The number of buffers to be used. DB2 will automatically choose an optimal value for this parameter unless you explicitly enter a value. However, when creating a backup to multiple locations, a larger number of buffers may be used to improve performance.

**BUFFER buffer-size**

The size, in 4 KB pages, of the buffer used when building the backup image. DB2 will automatically choose an optimal value for this parameter unless you explicitly enter a value. The minimum value for this parameter is 8 pages.

If using tape with variable block size, reduce the buffer size to within the range that the tape device supports. Otherwise, the backup operation may succeed, but the resulting image may not be recoverable.

When using tape devices on SCO UnixWare 7, specify a buffer size of 16.

With most versions of Linux, using DB2's default buffer size for backup operations to a SCSI tape device results in error SQL2025N, reason code 75. To prevent the overflow of Linux internal SCSI buffers, use this formula:

$$\text{bufferpages} \leq \text{ST\_MAX\_BUFFERS} * \text{ST\_BUFFER\_BLOCKS} / 4$$

where *bufferpages* is the value you want to use with the BUFFER parameter, and ST\_MAX\_BUFFERS and ST\_BUFFER\_BLOCKS are defined in the Linux kernel under the drivers/scsi directory.

**PARALLELISM n**

Determines the number of table spaces which can be read in parallel by the backup utility. DB2 will automatically choose an optimal value for this parameter unless you explicitly enter a value.

**UTIL\_IMPACT\_PRIORITY *priority***

Specifies that the backup will run in throttled mode, with the priority specified. Throttling allows you to regulate the performance impact of the backup operation. Priority can be any number between 1 and 100, with 1 representing the lowest priority, and 100 representing the highest priority. If the UTIL\_IMPACT\_PRIORITY keyword is specified with no priority, the backup will run with the default priority of 50. If UTIL\_IMPACT\_PRIORITY is not specified, the backup will run in unthrottled mode. An impact policy must be defined by setting the *util\_impact\_lim* configuration parameter for a backup to run in throttled mode.

**COMPRESS**

Indicates that the backup is to be compressed.

**COMPRLIB *name***

Indicates the name of the library to be used to perform the compression. The name must be a fully qualified path referring to a file on the server. If this parameter is not specified, the default DB2 compression library will be used. If the specified library cannot be loaded, the backup will fail.

**EXCLUDE**

Indicates that the compression library will not be stored in the backup image.

## BACKUP DATABASE

### COMPROPTS *string*

Describes a block of binary data that will be passed to the initialization routine in the compression library. DB2 will pass this string directly from the client to the server, so any issues of byte reversal or code page conversion will have to be handled by the compression library. If the first character of the data block is '@', the remainder of the data will be interpreted by DB2 as the name of a file residing on the server. DB2 will then replace the contents of string with the contents of this file and will pass this new value to the initialization routine instead. The maximum length for *string* is 1024 bytes.

### EXCLUDE LOGS

Specifies that the backup image should not include any log files.

**Note:** When performing an offline backup operation, logs are excluded whether or not this option is specified.

### INCLUDE LOGS

Specifies that the backup image should include the range of log files required to restore and roll forward this image to some consistent point in time. This option is not valid for an offline backup.

### WITHOUT PROMPTING

Specifies that the backup will run unattended, and that any actions which normally require user intervention will return an error message.

### Examples:

1. In the following example, the database WSDB is defined on all 4 partitions, numbered 0 through 3. The path /dev3/backup is accessible from all partitions. Partition 0 is the catalog partition, and needs to be backed-up separately since this is an offline backup. To perform an offline backup of all the WSDB database partitions to /dev3/backup, issue the following commands from one of the database partitions:

```
db2_all ' <<+0< db2 BACKUP DATABASE wsdb TO /dev3/backup '  
db2_all ' | <<-0< db2 BACKUP DATABASE wsdb TO /dev3/backup '
```

In the second command, the db2\_all utility will issue the same backup command to each database partition in turn (except partition 0). All four database partition backup images will be stored in the /dev3/backup directory.

2. In the following example database SAMPLE is backed up to a TSM server using two concurrent TSM client sessions. DB2 calculates the optimal buffer size for this environment.

```
db2 backup database sample use tsm open 2 sessions with 4 buffers
```

3. In the following example, a table space-level backup of table spaces (syscatspace, userspace1) of database payroll is done to tapes.

```
db2 backup database payroll tablespace (syscatspace, userspace1) to  
/dev/rmt0, /dev/rmt1 with 8 buffers without prompting
```

4. The USE TSM OPTIONS keywords can be used to specify the TSM information to use for the backup operation. The following example shows how to use the USE TSM OPTIONS keywords to specify a fully qualified file name:

```
db2 backup db sample use TSM options @/u/dmcinnis/myoptions.txt
```

The file myoptions.txt contains the following information: -fromnode=bar  
-fromowner=dmcinnis



5. Following is a sample weekly incremental backup strategy for a recoverable database. It includes a weekly full database backup operation, a daily non-cumulative (delta) backup operation, and a mid-week cumulative (incremental) backup operation:

```
(Sun) db2 backup db sample use tsm
(Mon) db2 backup db sample online incremental delta use tsm
(Tue) db2 backup db sample online incremental delta use tsm
(Wed) db2 backup db sample online incremental use tsm
(Thu) db2 backup db sample online incremental delta use tsm
(Fri) db2 backup db sample online incremental delta use tsm
(Sat) db2 backup db sample online incremental use tsm
```

6. In the following example, three identical target directories are specified for a backup operation on database SAMPLE. You might want to do this if the target file system is made up of multiple physical disks.

```
db2 backup database sample to /dev3/backup, /dev3/backup, /dev3/backup
```

The data will be concurrently backed up to the three target directories, and three backup images will be generated with extensions .001, .002, and .003.

**Related reference:**

- “RESTORE DATABASE” on page 87
- “ROLLFORWARD DATABASE” on page 123

---

## db2Backup - Backup database

Creates a backup copy of a database or a table space.

**Scope:**

This API only affects the database partition on which it is executed.

**Authorization:**

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

**Required connection:**

Database. This API automatically establishes a connection to the specified database.

The connection will be terminated upon the completion of the backup.

**API include file:**

*db2ApiDf.h*

**C API syntax:**

```
/* File: db2ApiDf.h */
/* API: db2Backup */
/* ... */
SQL_API_RC SQL_API_FN
db2Backup (
    db2UInt32    versionNumber,
```

## db2Backup - Backup database

```
void          *pDB2BackupStruct,
struct sqlca *pSqlca);

typedef SQL_STRUCTURE db2BackupStruct
{
    char                *piDBAlias;
    char                oApplicationId[SQLU_APPLID_LEN+1];
    char                oTimestamp[SQLU_TIME_STAMP_LEN+1];
    struct db2TablespaceStruct *piTablespaceList;
    struct db2MediaListStruct *piMediaList;
    char                *piUsername;
    char                *piPassword;
    void                *piVendorOptions;
    db2Uint32           iVendorOptionsSize;
    db2Uint32           oBackupSize;
    db2Uint32           iCallerAction;
    db2Uint32           iBufferSize;
    db2Uint32           iNumBuffers;
    db2Uint32           iParallelism;
    db2Uint32           iOptions;
    db2Uint32           iUtilImpactPriority;
    char                *piComprLibrary;
    void                *piComprOptions;
    db2Uint32           iComprOptionsSize;
} db2BackupStruct;

typedef SQL_STRUCTURE db2TablespaceStruct
{
    char                **tablespaces;
    db2Uint32           numTablespaces;
} db2TablespaceStruct;

typedef SQL_STRUCTURE db2MediaListStruct
{
    char                **locations;
    db2Uint32           numLocations;
    char                locationType;
} db2MediaListStruct;
/* ... */
```

### Generic API syntax:

```
/* File: db2ApiDf.h */
/* API: db2Backup */
/* ... */
SQL_API_RC SQL_API_FN
db2gBackup (
    db2Uint32           versionNumber,
    void                *pDB2gBackupStruct,
    struct sqlca *pSqlca);

typedef SQL_STRUCTURE db2gBackupStruct
{
    char                *piDBAlias;
    db2Uint32           iDBAliasLen;
    char                *poApplicationId;
    db2Uint32           iApplicationIdLen;
    char                *poTimestamp;
    db2Uint32           iTimestampLen;
    struct db2gTablespaceStruct *piTablespaceList;
    struct db2gMediaListStruct *piMediaList;
    char                *piUsername;
    db2Uint32           iUsernameLen;
    char                *piPassword;
    db2Uint32           iPasswordLen;
    void                *piVendorOptions;
    db2Uint32           iVendorOptionsSize;
```

```

db2UInt32          oBackupSize;
db2UInt32          iCallerAction;
db2UInt32          iBufferSize;
db2UInt32          iNumBuffers;
db2UInt32          iParallelism;
db2UInt32          iOptions;
db2UInt32          iUtilImpactPriority;
char               *piComprLibrary;
db2UInt32          iComprLibraryLen;
void               *piComprOptions;
db2UInt32          iComprOptionsSize;
} db2gBackupStruct;

typedef SQL_STRUCTURE db2gTablespaceStruct
{
    struct db2Char          *tablespaces;
    db2UInt32              numTablespaces;
} db2gTablespaceStruct;

typedef SQL_STRUCTURE db2gMediaListStruct
{
    struct db2Char          *locations;
    db2UInt32              numLocations;
    char                   locationType;
} db2gMediaListStruct;

typedef SQL_STRUCTURE db2Char
{
    char                   *pioData;
    db2UInt32              iLength;
    db2UInt32              oLength;
} db2Char;
/* ... */

```

**API parameters:****versionNumber**

Input. Specifies the version and release level of the structure passed as the second parameter *pDB2BackupStruct*.

**pDB2BackupStruct**

Input. A pointer to the *db2BackupStruct* structure.

**pSqlca**

Output. A pointer to the *sqlca* structure.

**piDBAlias**

Input. A string containing the database alias (as cataloged in the system database directory) of the database to back up.

**iDBAliasLen**

Input. A 4-byte unsigned integer representing the length in bytes of the database alias.

**oApplicationId**

Output. The API will return a string identifying the agent servicing the application. Can be used to obtain information about the progress of the backup operation using the database monitor.

**poApplicationId**

Output. Supply a buffer of length `SQLU_APPLID_LEN+1` (defined in `sqlutil.h`). The API will return a string identifying the agent servicing the application. Can be used to obtain information about the progress of the backup operation using the database monitor.

## db2Backup - Backup database

### **iApplicationIdLen**

Input. A 4-byte unsigned integer representing the length in bytes of the *poApplicationId* buffer. Should be equal to `SQLU_APPLID_LEN+1` (defined in `sqlutil.h`).

### **oTimestamp**

Output. The API will return the time stamp of the backup image

### **poTimestamp**

Output. Supply a buffer of length `SQLU_TIME_STAMP_LEN+1` (defined in `sqlutil.h`). The API will return the time stamp of the backup image.

### **iTimestampLen**

Input. A 4-byte unsigned integer representing the length in bytes of the *poTimestamp* buffer. Should be equal to `SQLU_TIME_STAMP_LEN+1` (defined in `sqlutil.h`).

### **piTablespaceList**

Input. List of table spaces to be backed up. Required for table space level backup only. Must be NULL for a database level backup. See structure `DB2TablespaceStruct`.

### **piMediaList**

Input. This structure allows the caller to specify the destination for the backup operation. The information provided depends on the value of the *locationType* parameter. The valid values for *locationType* (defined in `sqlutil.h`) are:

#### **SQLU\_LOCAL\_MEDIA**

Local devices (a combination of tapes, disks, or diskettes).

#### **SQLU\_TSM\_MEDIA**

TSM. If the locations pointer is set to NULL, the TSM shared library provided with DB2 is used. If a different version of the TSM shared library is desired, use `SQLU_OTHER_MEDIA` and provide the shared library name.

#### **SQLU\_OTHER\_MEDIA**

Vendor product. Provide the shared library name in the locations field.

#### **SQLU\_USER\_EXIT**

User exit. No additional input is required (only available when server is on OS/2).

For more information, see the *db2MediaListStruct* structure .

### **piUsername**

Input. A string containing the user name to be used when attempting a connection. Can be NULL.

### **iUsernameLen**

Input. A 4-byte unsigned integer representing the length in bytes of the user name. Set to zero if no user name is provided.

### **piPassword**

Input. A string containing the password to be used with the user name. Can be NULL.

### **iPasswordLen**

Input. A 4-byte unsigned integer representing the length in bytes of the password. Set to zero if no password is provided.

**piVendorOptions**

Input. Used to pass information from the application to the vendor functions. This data structure must be flat; that is, no level of indirection is supported. Note that byte-reversal is not done, and code page is not checked for this data.

**iVendorOptionsSize**

Input. The length of the *piVendorOptions* field, which cannot exceed 65535 bytes.

**oBackupSize**

Output. Size of the backup image (in MB).

**iCallerAction**

Input. Specifies action to be taken. Valid values (defined in *db2ApiDf.h*) are:

**DB2BACKUP\_BACKUP**

Start the backup.

**DB2BACKUP\_NOINTERRUPT**

Start the backup. Specifies that the backup will run unattended, and that scenarios which normally require user intervention will either be attempted without first returning to the caller, or will generate an error. Use this caller action, for example, if it is known that all of the media required for the backup have been mounted, and utility prompts are not desired.

**DB2BACKUP\_CONTINUE**

Continue the backup after the user has performed some action requested by the utility (mount a new tape, for example).

**DB2BACKUP\_TERMINATE**

Terminate the backup after the user has failed to perform some action requested by the utility.

**DB2BACKUP\_DEVICE\_TERMINATE**

Remove a particular device from the list of devices used by backup. When a particular medium is full, backup will return a warning to the caller (while continuing to process using the remaining devices). Call backup again with this caller action to remove the device which generated the warning from the list of devices being used.

**DB2BACKUP\_PARM\_CHK**

Used to validate parameters without performing a backup. This option does not terminate the database connection after the call returns. After successful return of this call, it is expected that the user will issue a call with *SQLUB\_CONTINUE* to proceed with the action.

**DB2BACKUP\_PARM\_CHK\_ONLY**

Used to validate parameters without performing a backup. Before this call returns, the database connection established by this call is terminated, and no subsequent call is required.

**iBufferSize**

Input. Backup buffer size in 4 KB allocation units (pages). Minimum is 8 units.

## db2Backup - Backup database

### iNumBuffers

Input. Specifies number of backup buffers to be used. Minimum is 2. Maximum is limited by memory.

### iParallelism

Input. Degree of parallelism (number of buffer manipulators). Minimum is 1. Maximum is 1024.

### iOptions

Input. A bitmap of backup properties. The options are to be combined using the bitwise OR operator to produce a value for *iOptions*. Valid values (defined in *db2ApiDf.h*) are:

#### DB2BACKUP\_OFFLINE

Offline gives an exclusive connection to the database.

#### DB2BACKUP\_ONLINE

Online allows database access by other applications while the backup operation occurs.

**Note:** An online backup operation may appear to hang if users are holding locks on SMS LOB data.

#### DB2BACKUP\_DB

Full database backup.

#### DB2BACKUP\_TABLESPACE

Table space level backup. For a table space level backup, provide a list of table spaces in the *piTablespaceList* parameter.

#### DB2BACKUP\_INCREMENTAL

Specifies a cumulative (incremental) backup image. An incremental backup image is a copy of all database data that has changed since the most recent successful, full backup operation.

#### DB2BACKUP\_DELTA

Specifies a noncumulative (delta) backup image. A delta backup image is a copy of all database data that has changed since the most recent successful backup operation of any type.

#### DB2BACKUP\_COMPRESS

Specifies that the backup should be compressed.

#### DB2BACKUP\_INCLUDE\_COMPR\_LIB

Specifies that the library used for compressing the backup should be included in the backup image.

#### DB2BACKUP\_EXCLUDE\_COMPR\_LIB

Specifies that the library used for compressing the backup should be not included in the backup image.

#### DB2BACKUP\_INCLUDE\_LOGS

Specifies that the backup image should also include the range of log files required to restore and roll forward this image to some consistent point in time. This option is not valid for an offline backup or a multi-partition backup.

#### DB2BACKUP\_EXCLUDE\_LOGS

Specifies that the backup image should not include any log files.

**Note:** When performing an offline backup operation, logs are excluded whether or not this option is specified.

**iUtilImpactPriority**

Specifies the priority value that will be used during a backup. The priority value can be any number between 0 and 100, with 0 representing unthrottled and 100 representing the highest priority.

**piComprLibrary**

Input. Indicates the name of the external library to be used to perform compression of the backup image. The name must be a fully-qualified path referring to a file on the server. If the value is a null pointer or a pointer to an empty string, DB2 will use the default library for compression. If the specified library is not found, the backup will fail.

**piComprLibraryLen**

Input. A four-byte unsigned integer representing the length in bytes of the name of the library specified in piComprLibrary. Set to zero if no library name is given.

**piComprOptions**

Input. Describes a block of binary data that will be passed to the initialization routine in the compression library. DB2 will pass this string directly from the client to the server, so any issues of byte-reversal or code-page conversion will have to be handled by the compression library. If the first character of the data block is '@', the remainder of the data will be interpreted by DB2 as the name of a file residing on the server. DB2 will then replace the contents of piComprOptions and iComprOptionsSize with the contents and size of this file respectively and will pass these new values to the initialization routine instead.

**iComprOptionsSize**

Input. A four-byte unsigned integer representing the size of the block of data passed as piComprOptions. iComprOptionsSize shall be zero if and only if piComprOptions is a null pointer.

**tablespaces**

A pointer to the list of table spaces to be backed up. For C, the list is null-terminated strings. In the generic case, it is a list of *db2Char* structures.

**numTablespaces**

Number of entries in the *tablespaces* parameter.

**locations**

A pointer to the list of media locations. For C, the list is null-terminated strings. In the generic case, it is a list of *db2Char* structures.

**numLocations**

The number of entries in the *locations* parameter.

**locationType**

A character indicated the media type. Valid values (defined in *sqlutil.h*) are:

**SQLU\_LOCAL\_MEDIA**

Local devices (tapes, disks, diskettes, or named pipes).

**SQLU\_TSM\_MEDIA**

Tivoli Storage Manager.

**SQLU\_OTHER\_MEDIA**

Vendor library.

**SQLU\_USER\_EXIT**

User exit (only available when the server is on OS/2).

## db2Backup - Backup database

### pioData

A pointer to the character data buffer.

### iLength

Input. The size of the *pioData* buffer.

### oLength

Output. Reserved for future use.

### Related reference:

- “sqlmgdb - Migrate Database” in the *Administrative API Reference*
- “db2Rollforward - Rollforward Database” on page 131
- “SQLCA” in the *Administrative API Reference*
- “db2Restore - Restore database” on page 96

### Related samples:

- “dbrecov.sqc -- How to recover a database (C)”
- “dbrecov.sqC -- How to recover a database (C++)”

---

## Backup sessions - CLP examples

### Example 1

In the following example database SAMPLE is backed up to a TSM server using 2 concurrent TSM client sessions. The backup utility will use 4 buffers. The optimal size of the buffers, in 4 KB pages, is automatically calculated based on the amount of memory and the number of target devices that are available. The parallelism setting is also automatically calculated and is based on the number of processors available and the number of table spaces to be backed up.

```
db2 backup database sample use tsm open 2 sessions with 4 buffers
```

```
db2 backup database payroll tablespace (syscatspace, userspace1) to  
/dev/rmt0, /dev/rmt1 with 8 buffers without prompting
```

### Example 2

Following is a sample weekly incremental backup strategy for a recoverable database. It includes a weekly full database backup operation, a daily non-cumulative (delta) backup operation, and a mid-week cumulative (incremental) backup operation:

```
(Sun) db2 backup db kdr use tsm  
(Mon) db2 backup db kdr online incremental delta use tsm  
(Tue) db2 backup db kdr online incremental delta use tsm  
(Wed) db2 backup db kdr online incremental use tsm  
(Thu) db2 backup db kdr online incremental delta use tsm  
(Fri) db2 backup db kdr online incremental delta use tsm  
(Sat) db2 backup db kdr online incremental use tsm
```

### Example 3

To initiate a backup operation to a tape device in a Windows environment, issue:

```
db2 backup database sample to \\.\tape0
```

### Related tasks:

- “Using backup” on page 62



---

## Optimizing backup performance

When you perform a backup operation, DB2® will automatically choose an optimal value for the number of buffers, the buffer size and the parallelism settings. The values will be based on the amount of utility heap memory available, the number of processors available, and the database configuration. The objective is to minimize the time it takes to complete a backup operation. Unless you explicitly enter a value for the following BACKUP DATABASE command parameters, DB2 will select one for them:

- WITH num-buffers BUFFERS
- PARALLELISM n
- BUFFER buffer-size

The values specified by the database manager configuration parameters BACKBUFSZ and RESTBUFSZ are ignored. If you want to use these values, you must explicitly specify a buffer size when you issue the BACKUP DATABASE command.

You can also choose to do any of the following to reduce the amount of time required to complete a backup operation:

- Specify table space backup.  
You can back up (and subsequently recover) part of a database by using the TABLESPACE option on the BACKUP DATABASE command. This facilitates the management of table data, indexes, and long field or large object (LOB) data in separate table spaces.
- Increase the value of the PARALLELISM parameter on the BACKUP DATABASE command so that it reflects the number of table spaces being backed up.

The PARALLELISM parameter defines the number of processes or threads that are started to read data from the database and to compress data during a compressed backup operation. Each process or thread is assigned to a specific table space, so there is no benefit to specifying a value for the PARALLELISM parameter that is larger than the number of table spaces being backed up. When it finishes backing up this table space, it requests another. Note, however, that each process or thread requires both memory and CPU overhead.

- Increase the backup buffer size.  
The ideal backup buffer size is a multiple of the table space extent size plus one page. If you have multiple table spaces with different extent sizes, specify a value that is a common multiple of the extent sizes plus one page.
- Increase the number of buffers.  
Use at least twice as many buffers as backup targets (or sessions) to ensure that the backup target devices do not have to wait for data.
- Use multiple target devices.

### Related concepts:

- “Backup overview” on page 59

### Related tasks:

- “Using backup” on page 62



---

## Chapter 3. Database restore

This section describes the DB2 UDB restore utility, which is used to rebuild damaged or corrupted databases or table spaces that were previously backed up.

The following topics are covered:

- “Restore overview”
- “Privileges, authorities, and authorization required to use restore” on page 82
- “Using restore” on page 82
- “Using incremental restore in a test and production environment” on page 83
- “Redefining table space containers during a restore operation (redirected restore)” on page 85
- “Restoring to an existing database” on page 86
- “Restoring to a new database” on page 87
- “RESTORE DATABASE” on page 87
- “db2Restore - Restore database” on page 96
- “Restore sessions - CLP examples” on page 107
- “Optimizing restore performance” on page 109

---

### Restore overview

The simplest form of the DB2® RESTORE DATABASE command requires only that you specify the alias name of the database that you want to restore. For example:

```
db2 restore db sample
```

| In this example, because the SAMPLE database exists and will be replaced when  
| the RESTORE DATABASE command is issued, the following message is returned:

```
| SQL2539W Warning! Restoring to an existing database that is the same as  
| the backup image database. The database files will be deleted.  
| Do you want to continue ? (y/n)
```

| If you specify *y*, the restore operation should complete successfully.

A database restore operation requires an exclusive connection: that is, no applications can be running against the database when the operation starts, and the restore utility prevents other applications from accessing the database until the restore operation completes successfully. A table space restore operation, however, can be done online.

A table space is not usable until the restore operation (followed by rollforward recovery) completes successfully.

If you have tables that span more than one table space, you should back up and restore the set of table spaces together.

When doing a partial or subset restore operation, you can use either a table space-level backup image, or a full database-level backup image and choose one or more table spaces from that image. All the log files associated with these table spaces from the time that the backup image was created must exist.

**Related concepts:**

- “Including log files with a backup image” on page 49

---

## Privileges, authorities, and authorization required to use restore

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges and higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects. Users can access only those objects for which they have the appropriate authorization; that is, the required privilege or authority.

You must have SYSADM, SYSCTRL, or SYSMAINT authority to restore to an *existing* database from a full database backup. To restore to a *new* database, you must have SYSADM or SYSCTRL authority.

---

## Using restore

**Prerequisites:**

When restoring to an *existing* database, you should not be connected to the database that is to be restored: the restore utility automatically establishes a connection to the specified database, and this connection is terminated at the completion of the restore operation. When restoring to a *new* database, an instance attachment is required to create the database. When restoring to a *new remote* database, you must first attach to the instance where the new database will reside. Then, create the new database, specifying the code page and the territory of the server.

The database can be local or remote.

**Restrictions:**

The following restrictions apply to the restore utility:

- You can only use the restore utility if the database has been previously backed up using the DB2 backup utility.
- A database restore operation cannot be started while the rollforward process is running.
- You can restore a table space only if the table space currently exists, and if it is the same table space; “same” means that the table space was not dropped and then recreated between the backup and the restore operation.)
- You cannot restore a table space-level backup to a new database.
- You cannot perform an online table space-level restore operation involving the system catalog tables.

**Procedure:**

The restore utility can be invoked through the command line processor (CLP), the Restore Database notebook or wizard in the Control Center, or the **db2Restore** application programming interface (API).

Following is an example of the RESTORE DATABASE command issued through the CLP:

```
db2 restore db sample from D:\DB2Backups taken at 20010320122644
```

To open the Restore Database notebook or wizard:

1. From the Control Center, expand the object tree until you find the Databases folder.
2. Click on the Databases folder. Any existing databases are displayed in the pane on the right side of the window (the contents pane).
3. Click the right mouse button on the database you want in the contents pane, and select Restore Database or Restore Database Using Wizard from the pop-up menu. The Restore Database notebook or the Restore Database wizard opens.

Detailed information is provided through the online help facility within the Control Center.

**Related concepts:**

- “Administrative APIs in Embedded SQL or DB2 CLI Programs” in the *Application Development Guide: Programming Client Applications*
- “Introducing the plug-in architecture for the Control Center” in the *Administration Guide: Implementation*

**Related reference:**

- “db2Restore - Restore database” on page 96

---

## Using incremental restore in a test and production environment

Once a production database is enabled for incremental backup and recovery, you can use an incremental or delta backup image to create or refresh a test database. You can do this by using either manual or automatic incremental restore. To restore the backup image from the production database to the test database, use the `INTO target-database-alias` option on the `RESTORE DATABASE` command. For example, in a production database with the following backup images:

```
backup db prod
Backup successful. The timestamp for this backup image is : <ts1>
```

```
backup db prod incremental
Backup successful. The timestamp for this backup image is : <ts2>
```

an example of a manual incremental restore would be:

```
restore db prod incremental taken at <ts2> into test without
prompting
DB20000I The RESTORE DATABASE command completed successfully.
```

```
restore db prod incremental taken at <ts1> into test without
prompting
DB20000I The RESTORE DATABASE command completed successfully.
```

```
restore db prod incremental taken at <ts2> into test without
prompting
DB20000I The RESTORE DATABASE command completed successfully.
```

If the database TEST already exists, the restore operation will overwrite any data that is already there. If the database TEST does not exist, the restore utility will create it and then populate it with the data from the backup images.

Since automatic incremental restore operations are dependent on the database history, the restore steps change slightly based on whether or not the test database

exists. To perform an automatic incremental restore to the database TEST, its history must contain the backup image history for database PROD. The database history for the backup image will replace any database history that already exists for database TEST if:

- the database TEST does not exist when the RESTORE DATABASE command is issued, or
- the database TEST exists when the RESTORE DATABASE command is issued, and the database TEST history contains no records.

The following example shows an automatic incremental restore to database TEST which does not exist:

```
restore db prod incremental automatic taken at <ts2> into test without
prompting
DB20000I The RESTORE DATABASE command completed successfully.
```

The restore utility will create the TEST database and populate it.

If the database TEST does exist and the database history is not empty, you must drop the database before the automatic incremental restore operation as follows:

```
drop db test
DB20000I The DROP DATABASE command completed successfully.

restore db prod incremental automatic taken at <ts2> into test without
prompting
DB20000I The RESTORE DATABASE command completed successfully.
```

If you do not want to drop the database, you can issue the PRUNE HISTORY command using a timestamp far into the future and the WITH FORCE OPTION parameter before issuing the RESTORE DATABASE command:

```
connect to test
Database Connection Information

Database server          = <server id>
SQL authorization ID    = <id>
Local database alias    = TEST

prune history 9999 with force option
DB20000I The PRUNE command completed successfully.

connect reset
DB20000I The SQL command completed successfully.
restore db prod incremental automatic taken at <ts2> into test without
prompting
SQL2540W Restore is successful, however a warning "2539" was
encountered during Database Restore while processing in No
Interrupt mode.
```

In this case, the RESTORE DATABASE COMMAND will act in the same manner as when the database TEST did not exist.

If the database TEST does exist and the database history is empty, you do not have to drop the database TEST before the automatic incremental restore operation:

```
restore db prod incremental automatic taken at <ts2> into test without
prompting
SQL2540W Restore is successful, however a warning "2539" was
encountered during Database Restore while processing in No
Interrupt mode.
```

You can continue taking incremental or delta backups of the test database without first taking a full database backup. However, if you ever need to restore one of the incremental or delta images you will have to perform a manual incremental restore. This is because automatic incremental restore operations require that each of the backup images restored during an automatic incremental restore be created from the same database alias.

If you make a full database backup of the test database after you complete the restore operation using the production backup image, you can take incremental or delta backups and can restore them using either manual or automatic mode.

**Related concepts:**

- “Incremental backup and recovery” on page 25

**Related reference:**

- “BACKUP DATABASE” on page 66
- “RESTORE DATABASE” on page 87
- “LIST HISTORY” on page 258

---

## Redefining table space containers during a restore operation (redirected restore)

During a database backup operation, a record is kept of all the table space containers associated with the table spaces that are being backed up. During a restore operation, all containers listed in the backup image are checked to determine if they exist and if they are accessible. If one or more of these containers is inaccessible because of media failure (or for any other reason), the restore operation will fail. A successful restore operation in this case requires redirection to different containers. DB2® supports adding, changing, or removing table space containers.

You can redefine table space containers by invoking the RESTORE DATABASE command and specifying the REDIRECT parameter, or by using the Containers page of the Restore Database notebook in the Control Center. The process for invoking a redirected restore of an incremental backup image is similar to the process for a non-incremental backup image: Call the RESTORE DATABASE command with the REDIRECT parameter and specify the backup image from which the database should be incrementally restored.

During a redirected restore operation, directory and file containers are automatically created if they do not already exist. The database manager does not automatically create device containers.

Container redirection provides considerable flexibility for managing table space containers. For example, even though adding containers to SMS table spaces is not supported, you could accomplish this by specifying an additional container when invoking a redirected restore operation.

The following example, show how to perform a redirected restore on database SAMPLE:

```
db2 restore db sample redirect without prompting
SQL1277N Restore has detected that one or more table space containers are
inaccessible, or has set their state to 'storage must be defined'.
DB20000I The RESTORE DATABASE command completed successfully.
```

```
db2 set tablespace containers for 2 using (path 'userspace1.0', path
'userspace1.1')
DB20000I The SET TABLESPACE CONTAINERS command completed successfully.
```

```
db2 restore db sample continue
DB20000I The RESTORE DATABASE command completed successfully.
```

**Related reference:**

- “RESTORE DATABASE” on page 87
- “Restore sessions - CLP examples” on page 107

**Related samples:**

- “dbrecov.out -- HOW TO RECOVER A DATABASE (C)”
- “dbrecov.sqc -- How to recover a database (C)”
- “dbrecov.out -- HOW TO RECOVER A DATABASE (C++)”
- “dbrecov.sqC -- How to recover a database (C++)”

---

## Restoring to an existing database

You can restore a full database backup image to an existing database. The backup image may differ from the existing database in its alias name, its database name, or its database seed.

A database *seed* is a unique identifier for a database that does not change during the life of the database. The seed is assigned by the database manager when the database is created. DB2® always uses the seed from the backup image.

When restoring to an existing database, the restore utility:

- Deletes table, index, and long field data from the existing database, and replaces it with data from the backup image.
- Replaces table entries for each table space being restored.
- Retains the recovery history file, unless it is damaged or has no entries. If the recovery history file is damaged or contains no entries, the database manager copies the file from the backup image.
- Retains the authentication type for the existing database.
- Retains the database directories for the existing database. The directories define where the database resides, and how it is cataloged.
- Compares the database seeds. If the seeds are different:
  - Deletes the logs associated with the existing database.
  - Copies the database configuration file from the backup image.
  - Sets NEWLOGPATH to the value of the *logpath* database configuration parameter if NEWLOGPATH was specified on the RESTORE DATABASE command.

If the database seeds are the same:

- Deletes the logs if the image is for a non-recoverable database.
- Retains the current database configuration file, unless the file has been corrupted, in which case the file is copied from the backup image.
- Sets NEWLOGPATH to the value of the *logpath* database configuration parameter if NEWLOGPATH was specified on the RESTORE DATABASE command; otherwise, copies the current log path to the database



configuration file. Validates the log path: If the path cannot be used by the database, changes the database configuration to use the default log path.

---

## Restoring to a new database

You can create a new database and then restore a full database backup image to it. If you do not create a new database, the restore utility will create one.

When restoring to a new database, the restore utility:

- Creates a new database, using the database alias name that was specified through the target database alias parameter. (If a target database alias was not specified, the restore utility creates the database with an alias that is the same as that specified through the source database alias parameter.)
- Restores the database configuration file from the backup image.
- Sets NEWLOGPATH to the value of the *logpath* database configuration parameter if NEWLOGPATH was specified on the RESTORE DATABASE command. Validates the log path: If the path cannot be used by the database, changes the database configuration to use the default log path.
- Restores the authentication type from the backup image.
- Restores the comments from the database directories in the backup image.
- Restores the recovery history file for the database.

---

## RESTORE DATABASE

Rebuilds a damaged or corrupted database that has been backed up using the DB2 backup utility. The restored database is in the same state it was in when the backup copy was made. This utility can also overwrite a database with a different image, or restore to a new database.

You can restore databases created on a DB2 Version 8 32-bit Windows platform to a DB2 Version 8 64-bit Windows platform, or the reverse. You can restore databases created on a DB2 Version 8 32-bit Linux (Intel) platform to a DB2 Version 8 64-bit Linux (Intel) platform, or the reverse. You can restore databases created on DB2 Version 8 AIX, HP-UX, or the Solaris Operating Environment platforms, in 32-bit or 64-bit, to DB2 Version 8 AIX, HP-UX, or Solaris Operating Environment platforms (32-bit or 64-bit).

The restore utility can also be used to restore backup images that were produced on a previous version of DB2 (up to two versions earlier) as long as the word size (32-bit or 64-bit) is the same. Cross-platform restore operations from a backup image created with a previous version of DB2 are not supported. If a migration is required, it will be invoked automatically at the end of the restore operation.

If, at the time of the backup operation, the database was enabled for rollforward recovery, the database can be brought to the state it was in prior to the occurrence of the damage or corruption by invoking the rollforward utility after successful completion of a restore operation.

This utility can also restore a table space level backup.

Incremental and delta images cannot be restored when there is a difference in operating systems or word size (32-bit or 64-bit).

## RESTORE DATABASE

Following a successful restore from one environment to a different environment, no incremental or delta backups are allowed until a non-incremental backup is taken. (This is not a limitation following a restore to the same environment.)

Even with a successful restore from one environment to a different environment, there are some considerations: packages must be rebound before use (using the BIND command, the REBIND command, or the db2rbind utility); SQL procedures must be dropped and recreated; and all external libraries must be rebuilt on the new platform. (These are not considerations when restoring to the same environment.)

### Scope:

This command only affects the node on which it is executed.

### Authorization:

To restore to an existing database, one of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

To restore to a new database, one of the following:

- *sysadm*
- *sysctrl*

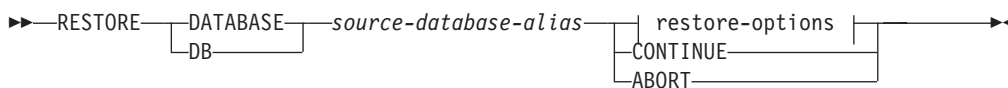
### Required connection:

The required connection will vary based on the type of restore action you wish to do:

- Database, to restore to an existing database. This command automatically establishes an exclusive connection to the specified database.
- Instance and database, to restore to a new database. The instance attachment is required to create the database.

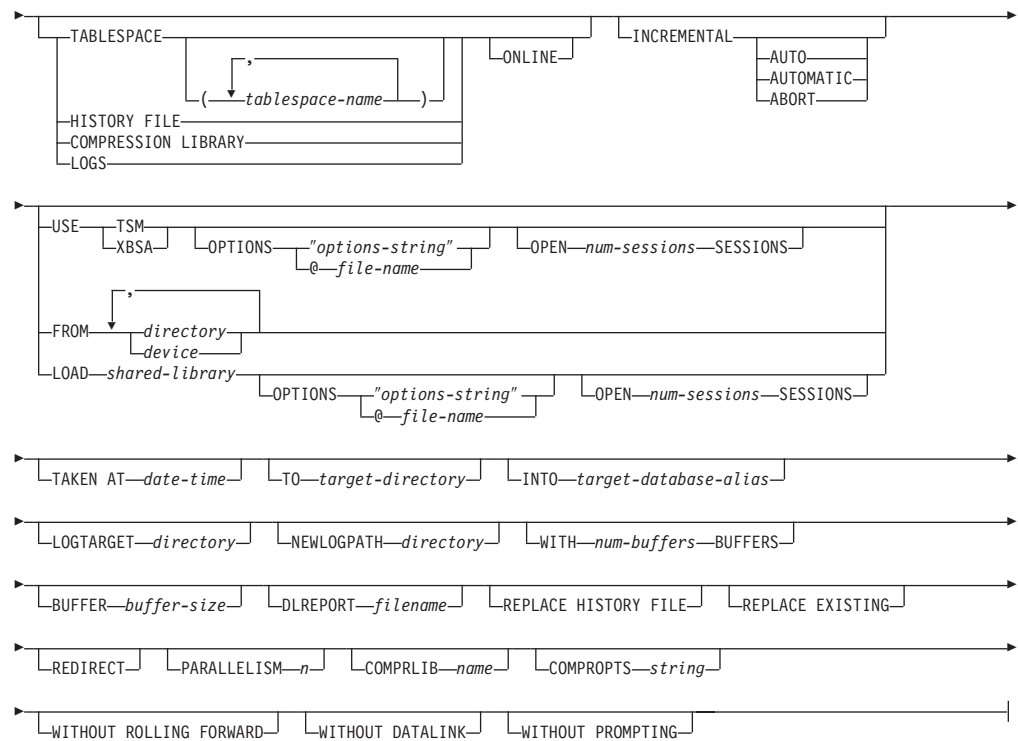
To restore to a new database at an instance different from the current instance, it is necessary to first attach to the instance where the new database will reside. The new instance may be local or remote. The current instance is defined by the value of the DB2INSTANCE environment variable.

### Command syntax:



### restore-options:





## Command parameters:

### DATABASE *source-database-alias*

Alias of the source database from which the backup was taken.

### CONTINUE

Specifies that the containers have been redefined, and that the final step in a redirected restore operation should be performed.

### ABORT

This parameter:

- Stops a redirected restore operation. This is useful when an error has occurred that requires one or more steps to be repeated. After RESTORE DATABASE with the ABORT option has been issued, each step of a redirected restore operation must be repeated, including RESTORE DATABASE with the REDIRECT option.
- Terminates an incremental restore operation before completion.

### USER *username*

Identifies the user name under which the database is to be restored.

### USING *password*

The password used to authenticate the user name. If the password is omitted, the user is prompted to enter it.

### TABLESPACE *tablespace-name*

A list of names used to specify the table spaces that are to be restored.

### ONLINE

This keyword, applicable only when performing a table space-level restore operation, is specified to allow a backup image to be restored online. This means that other agents can connect to the database while the backup image is being restored, and that the data in other table spaces will be available while the specified table spaces are being restored.

## RESTORE DATABASE

### HISTORY FILE

This keyword is specified to restore only the history file from the backup image.

### COMPRESSION LIBRARY

This keyword is specified to restore only the compression library from the backup image. If the object exists in the backup image, it will be restored into the database directory. If the object does not exist in the backup image, the restore operation will fail.

**LOGS** This keyword is specified to restore only the set of log files contained in the backup image. If the backup image does not contain any log files, the restore operation will fail. If this option is specified, the LOGTARGET option must also be specified.

### INCREMENTAL

Without additional parameters, INCREMENTAL specifies a manual cumulative restore operation. During manual restore the user must issue each restore command manually for each image involved in the restore. Do so according to the following order: last, first, second, third and so on up to and including the last image.

### INCREMENTAL AUTOMATIC/AUTO

Specifies an automatic cumulative restore operation.

### INCREMENTAL ABORT

Specifies abortion of an in-progress manual cumulative restore operation.

### USE TSM

Specifies that the database is to be restored from TSM-managed output.

### OPTIONS

*"options-string"*

Specifies options to be used for the restore operation. The string will be passed to the vendor support library, for example TSM, exactly as it was entered, without the quotes.

**Note:** Specifying this option overrides the value specified by the VENDOROPT database configuration parameter.

*@file-name*

Specifies that the options to be used for the restore operation are contained in a file located on the DB2 server. The string will be passed to the vendor support library, for example TSM. The file must be a fully qualified file name.

### OPEN *num-sessions* SESSIONS

Specifies the number of I/O sessions that are to be used with TSM or the vendor product.

### USE XBSA

Specifies that the XBSA interface is to be used. Backup Services APIs (XBSA) are an open application programming interface for applications or facilities needing data storage management for backup or archiving purposes.

### FROM *directory/device*

The fully qualified path name of the directory or device on which the backup image resides. If USE TSM, FROM, and LOAD are omitted, the default value is the current working directory of the client machine. This target directory or device must exist on the database server.

On Windows operating systems, the specified directory must not be a DB2-generated directory. For example, given the following commands:

```
db2 backup database sample to c:\backup
db2 restore database sample from c:\backup
```

Using these commands, DB2 generates subdirectories under the c:\backup directory to allow more than one backup to be placed in the specified top level directory. The DB2-generated subdirectories should be ignored. To specify precisely which backup image to restore, use the TAKEN AT parameter. There may be several backup images stored on the same path.

If several items are specified, and the last item is a tape device, the user is prompted for another tape. Valid response options are:

- c** Continue. Continue using the device that generated the warning message (for example, continue when a new tape has been mounted).
- d** Device terminate. Stop using *only* the device that generated the warning message (for example, terminate when there are no more tapes).
- t** Terminate. Abort the restore operation after the user has failed to perform some action requested by the utility.

### **LOAD** *shared-library*

The name of the shared library (DLL on Windows operating systems) containing the vendor backup and restore I/O functions to be used. The name can contain a full path. If the full path is not given, the value defaults to the path on which the user exit program resides.

### **TAKEN AT** *date-time*

The time stamp of the database backup image. The time stamp is displayed after successful completion of a backup operation, and is part of the path name for the backup image. It is specified in the form *yyyymmddhhmmss*. A partial time stamp can also be specified. For example, if two different backup images with time stamps 20021001010101 and 20021002010101 exist, specifying 20021002 causes the image with time stamp 20021002010101 to be used. If a value for this parameter is not specified, there must be only one backup image on the source media.

### **TO** *target-directory*

The target database directory. This parameter is ignored if the utility is restoring to an existing database. The drive and directory that you specify must be local.

**Note:** On Windows operating systems, when using this parameter, specify only the drive letter. If you specify a path, an error is returned.

### **INTO** *target-database-alias*

The target database alias. If the target database does not exist, it is created.

When you restore a database backup to an existing database, the restored database inherits the alias and database name of the existing database. When you restore a database backup to a nonexistent database, the new database is created with the alias and database name that you specify. This new database name must be unique on the system where you restore it.

### **LOGTARGET** *directory*

The absolute path name of an existing directory on the database server, to

## RESTORE DATABASE

be used as the target directory for extracting log files from a backup image. If this option is specified, any log files contained within the backup image will be extracted into the target directory. If this option is not specified, log files contained within a backup image will not be extracted. To extract only the log files from the backup image, specify the LOGS option.

### **NEWLOGPATH** *directory*

The absolute pathname of a directory that will be used for active log files after the restore operation. This parameter has the same function as the *newlogpath* database configuration parameter, except that its effect is limited to the restore operation in which it is specified. The parameter can be used when the log path in the backup image is not suitable for use after the restore operation; for example, when the path is no longer valid, or is being used by a different database.

### **WITH** *num-buffers* **BUFFERS**

The number of buffers to be used. DB2 will automatically choose an optimal value for this parameter unless you explicitly enter a value. A larger number of buffers can be used to improve performance when multiple sources are being read from, or if the value of PARALLELISM has been increased.

### **BUFFER** *buffer-size*

The size, in pages, of the buffer used for the restore operation. DB2 will automatically choose an optimal value for this parameter unless you explicitly enter a value. The minimum value for this parameter is 8 pages.

The restore buffer size must be a positive integer multiple of the backup buffer size specified during the backup operation. If an incorrect buffer size is specified, the buffers are allocated to be of the smallest acceptable size.

### **DLREPORT** *filename*

The file name, if specified, must be specified as an absolute path. Reports the files that become unlinked, as a result of a fast reconcile, during a restore operation. This option is only to be used if the table being restored has a DATALINK column type and linked files.

### **REPLACE HISTORY FILE**

Specifies that the restore operation should replace the history file on disk with the history file from the backup image.

### **REPLACE EXISTING**

If a database with the same alias as the target database alias already exists, this parameter specifies that the restore utility is to replace the existing database with the restored database. This is useful for scripts that invoke the restore utility, because the command line processor will not prompt the user to verify deletion of an existing database. If the WITHOUT PROMPTING parameter is specified, it is not necessary to specify REPLACE EXISTING, but in this case, the operation will fail if events occur that normally require user intervention.

### **REDIRECT**

Specifies a redirected restore operation. To complete a redirected restore operation, this command should be followed by one or more SET TABLESPACE CONTAINERS commands, and then by a RESTORE DATABASE command with the CONTINUE option.

**Note:** All commands associated with a single redirected restore operation must be invoked from the same window or CLP session.

**WITHOUT ROLLING FORWARD**

Specifies that the database is not to be put in rollforward pending state after it has been successfully restored.

If, following a successful restore operation, the database is in rollforward pending state, the ROLLFORWARD command must be invoked before the database can be used again.

If this option is specified when restoring from an online backup image, error SQL2537N will be returned.

**WITHOUT DATALINK**

Specifies that any tables with DATALINK columns are to be put in DataLink\_Reconcile\_Pending (DRP) state, and that no reconciliation of linked files is to be performed.

**PARALLELISM *n***

Specifies the number of buffer manipulators that are to be spawned during the restore operation. DB2 will automatically choose an optimal value for this parameter unless you explicitly enter a value.

**COMPRLIB *name***

Indicates the name of the library to be used to perform the decompression. The name must be a fully qualified path referring to a file on the server. If this parameter is not specified, DB2 will attempt to use the library stored in the image. If the backup was not compressed, the value of this parameter will be ignored. If the specified library cannot be loaded, the restore will fail.

**COMPROPTS *string***

Describes a block of binary data that will be passed to the initialization routine in the decompression library. DB2 will pass this string directly from the client to the server, so any issues of byte reversal or code page conversion will have to be handled by the decompression library. If the first character of the data block is '@', the remainder of the data will be interpreted by DB2 as the name of a file residing on the server. DB2 will then replace the contents of *string* with the contents of this file and will pass this new value to the initialization routine instead. The maximum length for string is 1024 bytes.

**WITHOUT PROMPTING**

Specifies that the restore operation is to run unattended. Actions that normally require user intervention will return an error message. When using a removable media device, such as tape or diskette, the user is prompted when the device ends, even if this option is specified.

**Examples:**

1. In the following example, the database WSDB is defined on all 4 partitions, numbered 0 through 3. The path /dev3/backup is accessible from all partitions. The following offline backup images are available from /dev3/backup:

```
wsdb.0.db2inst1.NODE0000.CATN0000.20020331234149.001
wsdb.0.db2inst1.NODE0001.CATN0000.20020331234427.001
wsdb.0.db2inst1.NODE0002.CATN0000.20020331234828.001
wsdb.0.db2inst1.NODE0003.CATN0000.20020331235235.001
```

To restore the catalog partition first, then all other database partitions of the WSDB database from the /dev3/backup directory, issue the following commands from one of the database partitions:

## RESTORE DATABASE

```
db2_all '<<+0< db2 RESTORE DATABASE wsdb FROM /dev3/backup
TAKEN AT 20020331234149
  INTO wsdb REPLACE EXISTING'
db2_all '<<+1< db2 RESTORE DATABASE wsdb FROM /dev3/backup
TAKEN AT 20020331234427
  INTO wsdb REPLACE EXISTING'
db2_all '<<+2< db2 RESTORE DATABASE wsdb FROM /dev3/backup
TAKEN AT 20020331234828
  INTO wsdb REPLACE EXISTING'
db2_all '<<+3< db2 RESTORE DATABASE wsdb FROM /dev3/backup
TAKEN AT 20020331235235
  INTO wsdb REPLACE EXISTING'
```

The db2\_all utility issues the restore command to each specified database partition.

- Following is a typical redirected restore scenario for a database whose alias is MYDB:

- Issue a RESTORE DATABASE command with the REDIRECT option.

```
db2 restore db mydb replace existing redirect
```

After successful completion of step 1, and before completing step 3, the restore operation can be aborted by issuing:

```
db2 restore db mydb abort
```

- Issue a SET TABLESPACE CONTAINERS command for each table space whose containers must be redefined. For example:

```
db2 set tablespace containers for 5 using
(file 'f:\ts3con1' 20000, file 'f:\ts3con2' 20000)
```

To verify that the containers of the restored database are the ones specified in this step, issue the LIST TABLESPACE CONTAINERS command.

- After successful completion of steps 1 and 2, issue:

```
db2 restore db mydb continue
```

This is the final step of the redirected restore operation.

- If step 3 fails, or if the restore operation has been aborted, the redirected restore can be restarted, beginning at step 1.
- Following is a sample weekly incremental backup strategy for a recoverable database. It includes a weekly full database backup operation, a daily non-cumulative (delta) backup operation, and a mid-week cumulative (incremental) backup operation:

```
(Sun) backup db mydb use tsm
(Mon) backup db mydb online incremental delta use tsm
(Tue) backup db mydb online incremental delta use tsm
(Wed) backup db mydb online incremental use tsm
(Thu) backup db mydb online incremental delta use tsm
(Fri) backup db mydb online incremental delta use tsm
(Sat) backup db mydb online incremental use tsm
```

For an automatic database restore of the images created on Friday morning, issue:

```
restore db mydb incremental automatic taken at (Fri)
```

For a manual database restore of the images created on Friday morning, issue:



```

restore db mydb incremental taken at (Fri)
restore db mydb incremental taken at (Sun)
restore db mydb incremental taken at (Wed)
restore db mydb incremental taken at (Thu)
restore db mydb incremental taken at (Fri)

```

4. To produce a backup image, which includes logs, for transportation to a remote site:

```
backup db sample online to /dev3/backup include logs
```

To restore that backup image, supply a LOGTARGET path and specify this path during ROLLFORWARD:

```
restore db sample from /dev3/backup logtarget /dev3/logs
rollforward db sample to end of logs and stop overflow log path /dev3/logs
```

5. To retrieve only the log files from a backup image that includes logs:

```
restore db sample logs from /dev3/backup logtarget /dev3/logs
```

6. The USE TSM OPTIONS keywords can be used to specify the TSM information to use for the restore operation. On Windows platforms, omit the -fromowner option.

- Specifying a delimited string:

```
db2 restore db sample use TSM options "-fromnode bar -fromowner dmcinnis"
```

- Specifying a fully qualified file:

```
db2 restore db sample use TSM options @/u/dmcinnis/myoptions.txt
```

The file myoptions.txt contains the following information: -fromnode=bar -fromowner=dmcinnis

#### Usage notes:

- A RESTORE DATABASE command of the form db2 restore db <name> will perform a full database restore with a database image and will perform a table space restore of the table spaces found in a table space image. Any RESTORE DATABASE command of the form db2 restore db <name> tablespace will perform a table space restore of the table spaces found in the image. Any RESTORE DATABASE command in which a list of table spaces is provided will perform a restore of whatever table spaces are explicitly listed.
- Following the restore of an online backup, you must perform a roll forward recovery.
- If a backup image is compressed, DB2 will detect this and automatically decompress the data before restoring it. If a library is specified on the db2Restore API, it will be used for decompressing the data. Otherwise, if a library that is stored in the backup image will be used. Otherwise, the data cannot be decompressed, so the restore will fail.
- If the compression library is to be restored from a backup image (either explicitly by specifying the DB2RESTORE\_COMPR\_LIB restore type or implicitly by performing a normal restore of a compressed backup), the restore operation must be done on the same platform and operating system that the backup was taken on. If the platform the backup was taken on is not the same as the platform that the restore is being done on, the restore operation will fail, even if DB2 normally supports cross-platform restores involving the two systems.
- 
- To restore log files from a backup image which contains them, the LOGTARGET option must be specified, providing a fully qualified and valid path which exists on the DB2 server. If those conditions are satisfied, the restore utility will write the log files from the image to the target path. If a LOGTARGET is specified

## RESTORE DATABASE

during a restore of a backup image which does not include logs, the restore will return an error before attempting to restore any table space data. A restore will also fail with an error if an invalid, or read-only, LOGTARGET path is specified.

- If any log files exist in the LOGTARGET path at the time the RESTORE DATABASE command is issued, a warning prompt will be returned to user. This warning will not be returned if WITHOUT PROMPTING is specified.
- During a restore operation where a LOGTARGET is specified, if any log file can not be extracted, for any reason, the restore will fail and return an error. If any of the log files being extracted from the backup image have the same name as an existing file already in the LOGTARGET path, the restore operation will fail and an error will be returned. The restore database utility will not overwrite existing log files in the LOGTARGET directory.
- It is also possible to restore only the saved log set from a backup image. To indicate that only the log files are to be restored, specify the LOGS option in addition to the LOGTARGET path. Specifying the LOGS option without a LOGTARGET path will result in an error. If any problem occurs while restoring log files in this mode of operation, the restore operation will terminate immediately and an error will be returned.
- During an automatic incremental restore operation, only the logs included in the target image of the restore operation will be retrieved from the backup image. Any logs included in intermediate images referenced during the incremental restore process will not be extracted from those intermediate backup images. During a manual incremental restore operation, the LOGTARGET path should only be specified with the final restore command to be issued.

### Related reference:

- “BACKUP DATABASE” on page 66
- “ROLLFORWARD DATABASE” on page 123
- “db2move - Database Movement Tool Command” in the *Command Reference*

---

## db2Restore - Restore database

Rebuilds a damaged or corrupted database that has been backed up using db2Backup - Backup Database. The restored database is in the same state it was in when the backup copy was made. This utility can also restore to a database with a name different from the database name in the backup image (in addition to being able to restore to a new database).

This utility can also be used to restore DB2 databases created in the two previous releases.

This utility can also restore from a table space level backup, or restore table spaces from within a database backup image.

### Scope:

This API only affects the database partition from which it is called.

### Authorization:

To restore to an existing database, one of the following:

- *sysadm*
- *sysctrl*

- *sysmaint*

To restore to a new database, one of the following:

- *sysadm*
- *sysctrl*

#### Required connection:

Database, to restore to an existing database. This API automatically establishes a connection to the specified database and will release the connection when the restore operation finishes.

Instance and database, to restore to a new database. The instance attachment is required to create the database.

To restore to a new database at an instance different from the current instance (as defined by the value of the DB2INSTANCE environment variable), it is necessary to first attach to the instance where the new database will reside.

#### API include file:

*db2ApiDf.h*

#### C API syntax:

```

/* File: db2ApiDf.h */
/* API: db2Restore */
/* ... */
SQL_API_RC SQL_API_FN
db2Restore (
    db2UInt32      versionNumber,
    void          *pDB2RestoreStruct,
    struct sqlca *pSqlca);
/* ... */

typedef SQL_STRUCTURE db2RestoreStruct
{
    char          *piSourceDBAlias;
    char          *piTargetDBAlias;
    char          oApplicationId[SQLU_APPLID_LEN+1];
    char          *piTimestamp;
    char          *piTargetDBPath;
    char          *piReportFile;
    struct db2TablespaceStruct *piTablespaceList;
    struct db2MediaListStruct *piMediaList;
    char          *piUsername;
    char          *piPassword;
    char          *piNewLogPath;
    void          *piVendorOptions;
    db2UInt32     iVendorOptionsSize;
    db2UInt32     iParallelism;
    db2UInt32     iBufferSize;
    db2UInt32     iNumBuffers;
    db2UInt32     iCallerAction;
    db2UInt32     iOptions;
    char          *piComprLibrary;
    void          *piComprOptions;
    db2UInt32     iComprOptionsSize;
    char          *piLogTarget;
} db2RestoreStruct;

typedef SQL_STRUCTURE db2TablespaceStruct

```

## db2Restore - Restore database

```
{
|
|         char                **tablespaces;
|         db2UInt32           numTablespaces;
|     } db2TablespaceStruct;
|
|     typedef SQL_STRUCTURE db2MediaListStruct
|     {
|         char                **locations;
|         db2UInt32           numLocations;
|         char                locationType;
|     } db2MediaListStruct;
|     /* ... */
|
|     Generic API syntax:
|
|     /* File: db2ApiDf.h */
|     /* API: db2gRestore */
|     /* ... */
|     SQL_API_RC SQL_API_FN
|     db2gRestore (
|         db2UInt32           versionNumber,
|         void                *pDB2gRestoreStruct,
|         struct sqlca *pSqlca);
|
|     typedef SQL_STRUCTURE db2gRestoreStruct
|     {
|         char                *piSourceDBAlias;
|         db2UInt32           iSourceDBAliasLen;
|         char                *piTargetDBAlias;
|         db2UInt32           iTargetDBAliasLen;
|         char                *poApplicationId;
|         db2UInt32           iApplicationIdLen;
|         char                *piTimestamp;
|         db2UInt32           iTimestampLen;
|         char                *piTargetDBPath;
|         db2UInt32           iTargetDBPathLen;
|         char                *piReportFile;
|         db2UInt32           iReportFileLen;
|         struct db2gTablespaceStruct *piTablespaceList;
|         struct db2gMediaListStruct *piMediaList;
|         char                *piUsername;
|         db2UInt32           iUsernameLen;
|         char                *piPassword;
|         db2UInt32           iPasswordLen;
|         char                *piNewLogPath;
|         db2UInt32           iNewLogPathLen;
|         void                *piVendorOptions;
|         db2UInt32           iVendorOptionsSize;
|         db2UInt32           iParallelism;
|         db2UInt32           iBufferSize;
|         db2UInt32           iNumBuffers;
|         db2UInt32           iCallerAction;
|         db2UInt32           iOptions;
|         char                *piComprLibrary;
|         db2UInt32           iComprLibraryLen;
|         void                *piComprOptions;
|         db2UInt32           iComprOptionsSize;
|         char                *piLogTarget;
|         db2UInt32           iLogTargetLen;
|     } db2gRestoreStruct;
|
|     typedef SQL_STRUCTURE db2gTablespaceStruct
|     {
|         struct db2Char        *tablespaces;
|         db2UInt32           numTablespaces;
|     } db2gTablespaceStruct;
```

```

typedef SQL_STRUCTURE db2gMediaListStruct
{
    struct db2Char          *locations;
    db2UInt32              numLocations;
    char                   locationType;
} db2gMediaListStruct;

typedef SQL_STRUCTURE db2Char
{
    char                   *pioData;
    db2UInt32              iLength;
    db2UInt32              oLength;
} db2Char;
/* ... */

```

**API parameters:****versionNumber**

Input. Specifies the version and release level of the structure passed as the second parameter *pDB2RestoreStruct*.

**pDB2RestoreStruct**

Input. A pointer to the *db2RestoreStruct* structure

**pSqlca**

Output. A pointer to the *sqlca* structure.

**piSourceDBAlias**

Input. A string containing the database alias of the source database backup image.

**iSourceDBAliasLen**

Input. A 4-byte unsigned integer representing the length in bytes of the source database alias.

**piTargetDBAlias**

Input. A string containing the target database alias. If this parameter is null, the *piSourceDBAlias* will be used.

**iTargetDBAliasLen**

Input. A 4-byte unsigned integer representing the length in bytes of the target database alias.

**oApplicationId**

Output. The API will return a string identifying the agent servicing the application. Can be used to obtain information about the progress of the backup operation using the database monitor.

**poApplicationId**

Output. Supply a buffer of length `SQLU_APPLID_LEN+1` (defined in `sqlutil`). The API will return a string identifying the agent servicing the application. Can be used to obtain information about the progress of the backup operation using the database monitor.

**iApplicationIdLen**

Input. A 4-byte unsigned integer representing the length in bytes of the `poApplicationId` buffer. Should be equal to `SQLU_APPLID_LEN+1` (defined in `sqlutil`).

**piTimestamp**

Input. A string representing the timestamp of the backup image. This field is optional if there is only one backup image in the source specified.

## db2Restore - Restore database

### **iTimestampLen**

Input. A 4-byte unsigned integer representing the length in bytes of the piTimestamp buffer.

### **piTargetDBPath**

Input. A string containing the relative or fully qualified name of the target database directory on the server. Used if a new database is to be created for the restored backup; otherwise not used.

### **piReportFile**

Input. The file name, if specified, must be fully qualified. The datalinks files that become unlinked during restore (as a result of a fast reconcile) will be reported.

### **iReportFileLen**

Input. A 4-byte unsigned integer representing the length in bytes of the piReportFile buffer.

### **piTablespaceList**

Input. List of table spaces to be restored. Used when restoring a subset of table spaces from a database or table space backup image. See the *DB2TablespaceStruct* structure. The following restrictions apply:

- The database must be recoverable; that is, log retain or user exits must be enabled.
- The database being restored to must be the same database that was used to create the backup image. That is, table spaces can not be added to a database through the table space restore function.
- The rollforward utility will ensure that table spaces restored in a partitioned database environment are synchronized with any other database partition containing the same table spaces. If a table space restore operation is requested and the *piTablespaceList* is NULL, the restore utility will attempt to restore all of the table spaces in the backup image.

When restoring a table space that has been renamed since it was backed up, the new table space name must be used in the restore command. If the old table space name is used, it will not be found.

### **piMediaList**

Input. Source media for the backup image. The information provided depends on the value of the locationType field. The valid values for locationType (defined in sqlutil) are:

#### **SQLU\_LOCAL\_MEDIA**

Local devices (a combination of tapes, disks, or diskettes).

#### **SQLU\_XBSA\_MEDIA**

XBSA interface. Backup Services APIs (XBSA) are an open application programming interface for applications or facilities needing data storage management for backup or archiving purposes.

#### **SQLU\_TSM\_MEDIA**

TSM. If the locations pointer is set to NULL, the TSM shared library provided with DB2 is used. If a different version of the TSM shared library is desired, use SQLU\_OTHER\_MEDIA and provide the shared library name.

**SQLU\_OTHER\_MEDIA**

Vendor product. Provide the shared library name in the locations field.

**SQLU\_USER\_EXIT**

User exit. No additional input is required (only available when server is on OS/2).

**piUsername**

Input. A string containing the user name to be used when attempting a connection. Can be NULL.

**iUsernameLen**

Input. A 4-byte unsigned integer representing the length in bytes of piUsername. Set to zero if no user name is provided.

**piPassword**

Input. A string containing the password to be used with the user name. Can be NULL.

**iPasswordLen**

Input. A 4-byte unsigned integer representing the length in bytes of piPassword. Set to zero if no password is provided.

**piNewLogPath**

Input. A string representing the path to be used for logging after the restore has completed. If this field is null the default log path will be used.

**iNewLogPathLen**

Input. A 4-byte unsigned integer representing the length in bytes of piNewLogPath.

**piVendorOptions**

Input. Used to pass information from the application to the vendor functions. This data structure must be flat; that is, no level of indirection is supported. Note that byte-reversal is not done, and the code page is not checked for this data.

**iVendorOptionsSize**

Input. The length of the piVendorOptions, which cannot exceed 65535 bytes.

**iParallelism**

Input. Degree of parallelism (number of buffer manipulators). Minimum is 1. Maximum is 1024.

**iBufferSize**

Input. Backup buffer size in 4 KB allocation units (pages). Minimum is 8 units. The size entered for a restore must be equal to or an integer multiple of the buffer size used to produce the backup image.

**iNumBuffers**

Input. Specifies number of restore buffers to be used.

**iCallerAction**

Input. Specifies action to be taken. Valid values (defined in db2ApiDf) are:

**DB2RESTORE\_RESTORE**

Start the restore operation.

**DB2RESTORE\_NOINTERRUPT**

Start the restore. Specifies that the restore will run unattended, and that scenarios which normally require user intervention will either be attempted without first returning to the caller, or will generate

## db2Restore - Restore database

an error. Use this caller action, for example, if it is known that all of the media required for the restore have been mounted, and utility prompts are not desired.

### **DB2RESTORE\_CONTINUE**

Continue the restore after the user has performed some action requested by the utility (mount a new tape, for example).

### **DB2RESTORE\_TERMINATE**

Terminate the restore after the user has failed to perform some action requested by the utility.

### **DB2RESTORE\_DEVICE\_TERMINATE**

Remove a particular device from the list of devices used by restore. When a particular device has exhausted its input, restore will return a warning to the caller. Call restore again with this caller action to remove the device which generated the warning from the list of devices being used.

### **DB2RESTORE\_PARM\_CHK**

Used to validate parameters without performing a restore. This option does not terminate the database connection after the call returns. After successful return of this call, it is expected that the user will issue a call with DB2RESTORE\_CONTINUE to proceed with the action.

### **DB2RESTORE\_PARM\_CHK\_ONLY**

Used to validate parameters without performing a restore. Before this call returns, the database connection established by this call is terminated, and no subsequent call is required.

### **DB2RESTORE\_TERMINATE\_INCRE**

Terminate an incremental restore operation before completion.

### **DB2RESTORE\_RESTORE\_STORDEF**

Initial call. Table space container redefinition requested.

### **DB2RESTORE\_STORDEF\_NOINTERRUPT**

Initial call. The restore will run uninterrupted. Table space container redefinition requested.

## **iOptions**

Input. A bitmap of restore properties. The options are to be combined using the bitwise OR operator to produce a value for *iOptions*. Valid values (defined in db2ApiDf) are:

### **DB2RESTORE\_OFFLINE**

Perform an offline restore operation.

### **DB2RESTORE\_ONLINE**

Perform an online restore operation.

### **DB2RESTORE\_DB**

Restore all table spaces in the database. This must be run offline

### **DB2RESTORE\_TABLESPACE**

Restore only the table spaces listed in the *piTablespaceList* parameter from the backup image. This can be online or offline.

### **DB2RESTORE\_HISTORY**

Restore only the history file.



**DB2RESTORE\_COMPR\_LIB**

Indicates that the compression library is to be restored. This option cannot be used simultaneously with any other restore type. If the object exists in the backup image, it will be restored into the database directory. If the object does not exist in the backup image, the restore operation will fail.

**DB2RESTORE\_LOGS**

Specify to restore only the set of log files contained in the backup image. If the backup image did not include log files, the restore operation will fail. If this option is specified, the *piLogTarget* parameter must also be supplied.

**DB2RESTORE\_INCREMENTAL**

Perform a manual cumulative restore operation.

**DB2RESTORE\_AUTOMATIC**

Perform an automatic cumulative (incremental) restore operation. Must be specified with DB2RESTORE\_INCREMENTAL.

**DB2RESTORE\_DATALINK**

Perform reconciliation operations. Tables with a defined DATALINK column must have RECOVERY YES option specified.

**DB2RESTORE\_NODATALINK**

Do not perform reconciliation operations. Tables with DATALINK columns are placed into DataLink\_Roconcile\_pending (DRP) state. Tables with a defined DATALINK column must have the RECOVERY YES option specified.

**DB2RESTORE\_ROLLFWD**

Place the database in rollforward pending state after it has been successfully restored.

**DB2RESTORE\_NOROLLFWD**

Do not place the database in rollforward pending state after it has been successfully restored. This cannot be specified for backups taken online or for table space level restores. If, following a successful restore, the database is in roll-forward pending state, db2Rollforward - Rollforward Database must be executed before the database can be used.

**piComprLibrary**

Input. Indicates the name of the external library to be used to perform decompression of the backup image if the image is compressed. The name must be a fully-qualified path referring to a file on the server. If the value is a null pointer or a pointer to an empty string, DB2 will attempt to use the library stored in the image. If the backup was not compressed, the value of this parameter will be ignored. If the specified library is not found, the restore will fail.

**piComprLibraryLen**

Input. A four-byte unsigned integer representing the length in bytes of the name of the library specified in piComprLibrary. Set to zero if no library name is given.

**piComprOptions**

Input. Describes a block of binary data that will be passed to the initialization routine in the decompression library. DB2 will pass this string directly from the client to the server, so any issues of byte-reversal or code-page conversion will have to be handled by the compression library.

## db2Restore - Restore database

If the first character of the data block is '@', the remainder of the data will be interpreted by DB2 as the name of a file residing on the server. DB2 will then replace the contents of `piComprOptions` and `iComprOptionsSize` with the contents and size of this file respectively and will pass these new values to the initialization routine instead.

### **iComprOptionsSize**

Input. A four-byte unsigned integer representing the size of the block of data passed as `piComprOptions`. `iComprOptionsSize` shall be zero if and only if `piComprOptions` is a null pointer.

### **piLogTarget**

Input. The absolute path of an existing directory on the database server to be used as the target directory for extracting log files from a backup image. If this parameter is specified, any log files included in the backup image will be extracted into the target directory. If this parameter is not specified, log files included in the backup image will not be extracted. To extract only the log files from the backup image, use the `DB2RESTORE_LOGS` parameter.

### **iLogTargetLen**

Input. A four-byte unsigned integer representing the length, in bytes, of the path in `piLogTarget`.

### **tablespaces**

A pointer to the list of table spaces to be backed up. For C, the list is null-terminated strings. In the generic case, it is a list of `db2Char` structures.

### **numTablespaces**

Number of entries in the `tablespaces` parameter.

### **locations**

A pointer to the list of media locations. For C, the list is null-terminated strings. In the generic case, it is a list of `db2Char` structures.

### **numLocations**

The number of entries in the `locations` parameter.

### **locationType**

A character indicated the media type. Valid values (defined in `sqlutil`) are:

#### **SQLU\_LOCAL\_MEDIA**

Local devices(tapes, disks, diskettes, or named pipes).

#### **SQLU\_XBSA\_MEDIA**

XBSA interface.

#### **SQLU\_TSM\_MEDIA**

Tivoli Storage Manager.

#### **SQLU\_OTHER\_MEDIA**

Vendor library.

#### **SQLU\_USER\_EXIT**

User exit (only available when the server is on OS/2).

### **pioData**

A pointer to the character data buffer.

### **iLength**

Input. The size of the `pioData` buffer

### **oLength**

Output. Reserved for future use.

### Usage notes:

For offline restore, this utility connects to the database in exclusive mode. The utility fails if any application, including the calling application, is already connected to the database that is being restored. In addition, the request will fail if the restore utility is being used to perform the restore, and any application, including the calling application, is already connected to any database on the same workstation. If the connect is successful, the API locks out other applications until the restore is completed.

The current database configuration file will not be replaced by the backup copy unless it is unusable. If the file is replaced, a warning message is returned.

The database or table space must have been backed up using db2Backup - Backup Database.

If the caller action is DB2RESTORE\_NOINTERRUPT, the restore continues without prompting the application. If the caller action is DB2RESTORE\_RESTORE, and the utility is restoring to an existing database, the utility returns control to the application with a message requesting some user interaction. After handling the user interaction, the application calls RESTORE DATABASE again, with the caller action set to indicate whether processing is to continue (DB2RESTORE\_CONTINUE) or terminate (DB2RESTORE\_TERMINATE) on the subsequent call. The utility finishes processing, and returns an SQLCODE in the *sqlca*.

To close a device when finished, set the caller action to DB2RESTORE\_DEVICE\_TERMINATE. If, for example, a user is restoring from 3 tape volumes using 2 tape devices, and one of the tapes has been restored, the application obtains control from the API with an SQLCODE indicating end of tape. The application can prompt the user to mount another tape, and if the user indicates "no more", return to the API with caller action SQLUD\_DEVICE\_TERMINATE to signal end of the media device. The device driver will be terminated, but the rest of the devices involved in the restore will continue to have their input processed until all segments of the restore set have been restored (the number of segments in the restore set is placed on the last media device during the backup process). This caller action can be used with devices other than tape (vendor supported devices).

To perform a parameter check before returning to the application, set caller action to DB2RESTORE\_PARM\_CHK.

Set caller action to DB2RESTORE\_RESTORE\_STORDEF when performing a redirected restore; used in conjunction with sqlbstsc - Set Tablespace Containers.

If a system failure occurs during a critical stage of restoring a database, the user will not be able to successfully connect to the database until a successful restore is performed. This condition will be detected when the connection is attempted, and an error message is returned. If the backed-up database is not configured for roll-forward recovery, and there is a usable current configuration file with either of these parameters enabled, following the restore, the user will be required to either take a new backup of the database, or disable the log retain and user exit parameters before connecting to the database.

Although the restored database will not be dropped (unless restoring to a nonexistent database), if the restore fails, it will not be usable.

## db2Restore - Restore database

If the restore type specifies that the history file on the backup is to be restored, it will be restored over the existing history file for the database, effectively erasing any changes made to the history file after the backup that is being restored. If this is undesirable, restore the history file to a new or test database so that its contents can be viewed without destroying any updates that have taken place.

If, at the time of the backup operation, the database was enabled for roll forward recovery, the database can be brought to the state it was in prior to the occurrence of the damage or corruption by issuing db2Rollforward after successful execution of db2Restore. If the database is recoverable, it will default to roll forward pending state after the completion of the restore.

If the database backup image is taken offline, and the caller does not want to roll forward the database after the restore, the DB2RESTORE\_NOROLLFWD option can be used for the restore. This results in the database being useable immediately after the restore. If the backup image is taken online, the caller must roll forward through the corresponding log records at the completion of the restore.

To restore log files from a backup image which contains them, the LOGTARGET option must be specified, providing a fully qualified and valid path which exists on the DB2 server. If those conditions are satisfied, the restore utility will write the log files from the image to the target path. If a LOGTARGET is specified during a restore of a backup image which does not include logs, the restore will return an error before attempting to restore any table space data. A restore will also fail with an error if an invalid, or read-only, LOGTARGET path is specified.

If any log files exist in the LOGTARGET path at the time the Restore command is issued, a warning prompt will be returned to user. This warning will not be returned if WITHOUT PROMPTING is specified.

During a restore where a LOGTARGET is specified, if any log file can not be extracted, for any reason, the restore will fail and return an error. If any of the log files being extracted from the backup image have the same name as an existing file already in the LOGTARGET path, the restore operation will fail and an error will be returned. The restore utility will not overwrite existing log files in the LOGTARGET directory.

It is also possible to restore only the saved log set from a backup image. To indicate that only the log files are to be restored, specify the LOGS option in addition to the LOGTARGET path. Specifying the LOGS option without a LOGTARGET path will result in an error. If any problem occurs while restoring log files in this mode of operation, the restore will terminate immediately and an error will be returned.

During an automatic incremental restore, only the logs included in the target image of the restore operation will be retrieved from the backup image. Any logs included in intermediate images referenced during the incremental restore process will not be extracted from those intermediate backup images. During a manual incremental restore, the LOGTARGET path should only be specified with the final restore command to be issued.

If a backup is compressed, DB2 will detect this and automatically decompress the data before restoring it. If a library is specified on the db2Restore API, it will be used for decompressing the data. Otherwise, if a library that is stored in the backup image will be used. Otherwise, the data cannot be decompressed, so the restore will fail.

If the compression library is to be restored from a backup image (either explicitly by specifying the DB2RESTORE\_COMPR\_LIB restore type or implicitly by performing a normal restore of a compressed backup), the restore operation must be done on the same platform and operating system that the backup was taken on. If the platform the backup was taken on is not the same as the platform that the restore is being done on, the restore operation will fail, even if DB2 normally supports cross-platform restores involving the two systems.

**Related reference:**

- “sqlmgdb - Migrate Database” in the *Administrative API Reference*
- “db2Rollforward - Rollforward Database” on page 131
- “SQLCA” in the *Administrative API Reference*
- “db2Backup - Backup database” on page 71
- “db2CfgGet - Get Configuration Parameters” in the *Administrative API Reference*

**Related samples:**

- “dbrecov.sqc -- How to recover a database (C)”
- “dbrecov.sqC -- How to recover a database (C++)”

---

## Restore sessions - CLP examples

**Example 1**

Following is a typical non-incremental redirected restore scenario for a database whose alias is MYDB:

1. Issue a RESTORE DATABASE command with the REDIRECT option.
 

```
db2 restore db mydb replace existing redirect
```
2. Issue a SET TABLESPACE CONTAINERS command for each table space whose containers you want to redefine. For example, in a Windows environment:
 

```
db2 set tablespace containers for 5 using
(file 'f:\ts3con1'20000, file 'f:\ts3con2'20000)
```

To verify that the containers of the restored database are the ones specified in this step, issue the LIST TABLESPACE CONTAINERS command for every table space whose container locations are being redefined.

3. After successful completion of steps 1 and 2, issue:
 

```
db2 restore db mydb continue
```

This is the final step of the redirected restore operation.
4. If step 3 fails, or if the restore operation has been aborted, the redirected restore can be restarted, beginning at step 1.

**Notes:**

1. After successful completion of step 1, and before completing step 3, the restore operation can be aborted by issuing:
 

```
db2 restore db mydb abort
```
2. If step 3 fails, or if the restore operation has been aborted, the redirected restore can be restarted, beginning at step 1.

**Example 2**

Following is a typical manual incremental redirected restore scenario for a database whose alias is MYDB and has the following backup images:

```
backup db mydb
Backup successful. The timestamp for this backup image is : <ts1>
```

```
backup db mydb incremental
Backup successful. The timestamp for this backup image is : <ts2>
```

1. Issue a RESTORE DATABASE command with the INCREMENTAL and REDIRECT options.  

```
db2 restore db mydb incremental taken at <ts2> replace existing redirect
```
2. Issue a SET TABLESPACE CONTAINERS command for each table space whose containers must be redefined. For example, in a Windows environment:  

```
db2 set tablespace containers for 5 using
(file 'f:\ts3con1'20000, file 'f:\ts3con2'20000)
```

To verify that the containers of the restored database are the ones specified in this step, issue the LIST TABLESPACE CONTAINERS command.

3. After successful completion of steps 1 and 2, issue:  

```
db2 restore db mydb continue
```
4. The remaining incremental restore commands can now be issued as follows:  

```
db2 restore db mydb incremental taken at <ts1>
db2 restore db mydb incremental taken at <ts2>
```

This is the final step of the redirected restore operation.

#### Notes:

1. After successful completion of step 1, and before completing step 3, the restore operation can be aborted by issuing:  

```
db2 restore db mydb abort
```
2. After successful completion of step 3, and before issuing all the required commands in step 4, the restore operation can be aborted by issuing:  

```
db2 restore db mydb incremental abort
```
3. If step 3 fails, or if the restore operation has been aborted, the redirected restore can be restarted, beginning at step 1.
4. If either restore command fails in step 4, the failing command can be reissued to continue the restore process.

#### Example 3

Following is a typical automatic incremental redirected restore scenario for the same database:

1. Issue a RESTORE DATABASE command with the INCREMENTAL AUTOMATIC and REDIRECT options.  

```
db2 restore db mydb incremental automatic taken at <ts2>
replace existing redirect
```
2. Issue a SET TABLESPACE CONTAINERS command for each table space whose containers must be redefined. For example, in a Windows environment:  

```
db2 set tablespace containers for 5 using
(file 'f:\ts3con1'20000, file 'f:\ts3con2'20000)
```

To verify that the containers of the restored database are the ones specified in this step, issue the LIST TABLESPACE CONTAINERS command.

3. After successful completion of steps 1 and 2, issue:  

```
db2 restore db mydb continue
```

This is the final step of the redirected restore operation.

**Notes:**

1. After successful completion of step 1, and before completing step 3, the restore operation can be aborted by issuing:  
`db2 restore db mydb abort`
2. If step 3 fails, or if the restore operation has been aborted, the redirected restore can be restarted, beginning at step 1 after issuing:  
`db2 restore db mydb incremental abort`

**Related reference:**

- “RESTORE DATABASE” on page 87
- “LIST TABLESPACE CONTAINERS Command” in the *Command Reference*
- “SET TABLESPACE CONTAINERS Command” in the *Command Reference*

---

## Optimizing restore performance

When you perform a restore operation, DB2 will automatically choose an optimal value for the number of buffers, the buffer size and the parallelism settings. The values will be based on the amount of utility heap memory available, the number of processors available and the database configuration. The objective is to minimize the time it takes to complete a restore operation. Unless you explicitly enter a value for the following RESTORE DATABASE command parameters, DB2 will select one for them:

- WITH num-buffers BUFFERS
- PARALLELISM n
- BUFFER buffer-size

For restore operations, a multiple of the buffer size used by the backup operation will always be used. The values specified by the database manager configuration parameters BACKBUFSZ and RESTBUFSZ are ignored. If you want to use these values, you must explicitly specify a buffer size when you issue the RESTORE DATABASE command.

You can also choose to do any of the following to reduce the amount of time required to complete a restore operation:

- Increase the restore buffer size.  
The restore buffer size must be a positive integer multiple of the backup buffer size specified during the backup operation. If an incorrect buffer size is specified, the buffers allocated will be the smallest acceptable size.
- Increase the number of buffers.  
The value you specify must be a multiple of the number of pages that you specified for the backup buffer. The minimum number of pages is 8.
- Increase the value of the PARALLELISM parameter.  
This will increase the number of buffer manipulators (BM) that will be used to write to the database during the restore operation.

**Related concepts:**

- “Restore overview” on page 81

**Related tasks:**

- “Using restore” on page 82





---

## Chapter 4. Rollforward recovery

This section describes the DB2 UDB rollforward utility, which is used to recover a database by applying transactions that were recorded in the database recovery log files.

The following topics are covered:

- “Rollforward overview”
- “Privileges, authorities, and authorization required to use rollforward” on page 113
- “Using rollforward” on page 113
- “Rolling forward changes in a table space” on page 114
- “Recovering a dropped table” on page 118
- “Recovering data with the load copy location file” on page 120
- “Synchronizing clocks in a partitioned database system” on page 121
- “Client/server timestamp conversion” on page 122
- “ROLLFORWARD DATABASE” on page 123
- “db2Rollforward - Rollforward Database” on page 131
- “Rollforward sessions - CLP examples” on page 140

---

### Rollforward overview

The simplest form of the DB2<sup>®</sup> ROLLFORWARD DATABASE command requires only that you specify the alias name of the database that you want to rollforward recover. For example:

```
db2 rollforward db sample to end of logs and stop
```

In this example, the command returns:

```
Rollforward Status

Input database alias           = sample
Number of nodes have returned status = 1

Node number                    = 0
Rollforward status             = not pending
Next log file to be read       =
Log files processed             = -
Last committed transaction     = 2001-03-11-02.39.48.000000
```

```
DB20000I The ROLLFORWARD command completed successfully.
```

The general approach to rollforward recovery involves:

1. Invoking the rollforward utility without the STOP option.
2. Invoking the rollforward utility with the QUERY STATUS option

If you specify recovery to the end of the logs, the QUERY STATUS option can indicate that one or more log files is missing, if the returned point in time is earlier than you expect.

If you specify point-in-time recovery, the QUERY STATUS option will help you to ensure that the rollforward operation has completed at the correct point.

3. Invoking the rollforward utility with the STOP option. After the operation stops, it is not possible to roll additional changes forward.

A database must be restored successfully (using the restore utility) before it can be rolled forward, but a table space does not. A table space may be temporarily put in rollforward pending state, but not require a restore operation to undo it (following a power interruption, for example).

When the rollforward utility is invoked:

- If the database is in rollforward pending state, the database is rolled forward. If table spaces are also in rollforward pending state, you must invoke the rollforward utility again after the database rollforward operation completes to roll the table spaces forward.
- If the database is *not* in rollforward pending state, but table spaces in the database *are* in rollforward pending state:
  - If you specify a list of table spaces, only those table spaces are rolled forward.
  - If you do not specify a list of table spaces, all table spaces that are in rollforward pending state are rolled forward.

A database rollforward operation runs offline. The database is not available for use until the rollforward operation completes successfully, and the operation cannot complete unless the STOP option was specified when the utility was invoked.

A table space rollforward operation can run offline. The database is not available for use until the rollforward operation completes successfully. This occurs if the end of the logs is reached, or if the STOP option was specified when the utility was invoked.

You can perform an *online* rollforward operation on table spaces, as long as SYSCATSPACE is not included. When you perform an online rollforward operation on a table space, the table space is not available for use, but the other table spaces in the database *are* available.

When you first create a database, it is enabled for circular logging only. This means that logs are reused, rather than being saved or archived. With circular logging, rollforward recovery is not possible: only crash recovery or version recovery can be done. Archived logs document changes to a database that occur after a backup was taken. You enable log archiving (and rollforward recovery) by setting the *logarchmeth1* database configuration parameter to a value other than its default of OFF. When you set *logarchmeth1* to a value other than OFF, the database is placed in backup pending state, and you must take an offline backup of the database before it can be used again.

**Note:** Entries will be made in the recovery history file for each log file that is used in a rollforward operation.

**Related concepts:**

- “Recovering data with the load copy location file” on page 120
- “Understanding recovery logs” on page 31

**Related reference:**

- “ROLLFORWARD DATABASE” on page 123
- “Configuration parameters for database logging” on page 35

- “logarchmeth1 - Primary log archive method configuration parameter” in the *Administration Guide: Performance*

---

## Privileges, authorities, and authorization required to use rollforward

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges and higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects. Users can access only those objects for which they have the appropriate authorization; that is, the required privilege or authority.

You must have SYSADM, SYSCTRL, or SYSMAINT authority to use the rollforward utility.

---

## Using rollforward

### Prerequisites:

You should not be connected to the database that is to be rollforward recovered: the rollforward utility automatically establishes a connection to the specified database, and this connection is terminated at the completion of the rollforward operation.

Do not restore table spaces without cancelling a rollforward operation that is in progress; otherwise, you may have a table space set in which some table spaces are in rollforward in progress state, and some table spaces are in rollforward pending state. A rollforward operation that is in progress will only operate on the table spaces that are in rollforward in progress state.

The database can be local or remote.

### Restrictions:

The following restrictions apply to the rollforward utility:

- You can only invoke one rollforward operation at a time. If there are many table spaces to recover, you can specify all of them in the same operation.
- If you have renamed a table space following the most recent backup operation, ensure that you use the new name when rolling the table space forward. The previous table space name will not be recognized.
- You cannot cancel a rollforward operation that is running. You can only cancel a rollforward operation that has completed, but for which the STOP option has not been specified, or a rollforward operation that has failed before completing.
- You cannot *continue* a table space rollforward operation to a point in time, specifying a time stamp that is less than the previous one. If a point in time is not specified, the previous one is used. You can initiate a rollforward operation to a point in time by just specifying STOP, but this is only allowed if the table spaces involved were all restored from the same offline backup image. In this case, no log processing is required. If you start another rollforward operation with a different table space list before the in-progress rollforward operation is either completed or cancelled, an error message (SQL4908) is returned. Invoke the LIST TABLESPACES command on all nodes to determine which table spaces

are currently being rolled forward (rollforward in progress state), and which table spaces are ready to be rolled forward (rollforward pending state). You have three options:

- Finish the in-progress rollforward operation on all table spaces.
  - Finish the in-progress rollforward operation on a subset of table spaces. (This may not be possible if the rollforward operation is to continue to a specific point in time, which requires the participation of all nodes.)
  - Cancel the in-progress rollforward operation.
- In a partitioned database environment, the rollforward utility must be invoked from the catalog node of the database.

**Procedure:**

The rollforward utility can be invoked through the command line processor (CLP), the Rollforward Database notebook in the Control Center, or the **db2Rollforward** application programming interface (API).

Following is an example of the ROLLFORWARD DATABASE command issued through the CLP:

```
db2 rollforward db sample to end of logs and stop
```

To open the Rollforward Database notebook:

1. From the Control Center, expand the object tree until you find the Databases folder.
2. Click on the Databases folder. Any existing databases are displayed in the pane on the right side of the window (the contents pane).
3. Click the right mouse button on the database you want in the contents pane, and select Rollforward from the pop-up menu. The Rollforward Database notebook opens.

Detailed information is provided through the online help facility within the Control Center.

**Related concepts:**

- “Administrative APIs in Embedded SQL or DB2 CLI Programs” in the *Application Development Guide: Programming Client Applications*
- “Introducing the plug-in architecture for the Control Center” in the *Administration Guide: Implementation*

**Related reference:**

- “db2Rollforward - Rollforward Database” on page 131

---

## Rolling forward changes in a table space

If the database is enabled for forward recovery, you have the option of backing up, restoring, and rolling forward table spaces instead of the entire database. You may want to implement a recovery strategy for individual table spaces because this can save time: it takes less time to recover a portion of the database than it does to recover the entire database. For example, if a disk is bad, and it contains only one table space, that table space can be restored and rolled forward without having to recover the entire database, and without impacting user access to the rest of the database, unless the damaged table space contains the system catalog tables; in this situation, you cannot connect to the database. (The system catalog table space can

be restored independently if a table space-level backup image containing the system catalog table space is available.) Table space-level backups also allow you to back up critical parts of the database more frequently than other parts, and requires less time than backing up the entire database.

After a table space is restored, it is always in rollforward pending state. To make the table space usable, you must perform rollforward recovery on it. In most cases, you have the option of rolling forward to the end of the logs, or rolling forward to a point in time. You cannot, however, roll table spaces containing system catalog tables forward to a point in time. These table spaces must be rolled forward to the end of the logs to ensure that all table spaces in the database remain consistent.

When a table space is rolled forward, DB2® will process all log files even if they do not contain log records that affect that table space. To skip the log files known not to contain any log records affecting the table space, set the `DB2_COLLECT_TS_REC_INFO` registry variable to `ON`. To ensure that the information required for skipping log files is collected, the registry variable must be set before the log files are created and used.

The table space change history file (`DB2TSCHG.HIS`), located in the database directory, keeps track of which logs should be processed for each table space. You can view the contents of this file using the `db2logsForRfwd` utility, and delete entries from it using the `PRUNE HISTORY` command. During a database restore operation, `DB2TSCHG.HIS` is restored from the backup image and then brought up to date during the database rollforward operation. If no information is available for a log file, it is treated as though it is required for the recovery of every table space.

Since information for each log file is flushed to disk after the log becomes inactive, this information can be lost as a result of a crash. To compensate for this, if a recovery operation begins in the middle of a log file, the entire log is treated as though it contains modifications to every table space in the system. After this, the active logs will be processed and the information for them will be rebuilt. If information for older or archived log files is lost in a crash situation and no information for them exists in the data file, they will be treated as though they contain modifications for every table space during the table space recovery operation.

Before rolling a table space forward, invoke the `LIST TABLESPACES SHOW DETAIL` command. This command returns the *minimum recovery time*, which is the earliest point in time to which the table space can be rolled forward. The minimum recovery time is updated when data definition language (DDL) statements are run against the table space, or against tables in the table space. The table space must be rolled forward to at least the minimum recovery time, so that it becomes synchronized with the information in the system catalog tables. If recovering more than one table space, the table spaces must be rolled forward to at least the highest minimum recovery time of all the table spaces being recovered. In a partitioned database environment, issue the `LIST TABLESPACES SHOW DETAIL` command on all partitions. The table spaces must be rolled forward to at least the highest minimum recovery time of all the table spaces on all partitions.

If you are rolling table spaces forward to a point in time, and a table is contained in multiple table spaces, all of these table spaces must be rolled forward simultaneously. If, for example, the table data is contained in one table space, and the index for the table is contained in another table space, you must roll both table spaces forward simultaneously to the same point in time.

If the data and the long objects in a table are in separate table spaces, and the long object data has been reorganized, the table spaces for both the data and the long objects must be restored and rolled forward together. You should take a backup of the affected table spaces after the table is reorganized.

If you want to roll a table space forward to a point in time, and a table in the table space is either:

- An underlying table for a materialized query or staging table that is in another table space
- A materialized query or staging table for a table in another table space

You should roll both table spaces forward to the same point in time. If you do not, the materialized query or staging table is placed in check pending state at the end of the rollforward operation. The materialized query table will need to be fully refreshed, and the staging table will be marked as incomplete.

If you want to roll a table space forward to a point in time, and a table in the table space participates in a referential integrity relationship with another table that is contained in another table space, you should roll both table spaces forward simultaneously to the same point in time. If you do not, the child table in the referential integrity relationship will be placed in check pending state at the end of the rollforward operation. When the child table is later checked for constraint violations, a check on the entire table is required. If any of the following tables exist, they will also be placed in check pending state with the child table:

- Any descendent materialized query tables for the child table
- Any descendent staging tables for the child table
- Any descendent foreign key tables of the child table

These tables will require full processing to bring them out of the check pending state. If you roll both table spaces forward simultaneously, the constraint will remain active at the end of the point-in-time rollforward operation.

Ensure that a point-in-time table space rollforward operation does not cause a transaction to be rolled back in some table spaces, and committed in others. This can happen if:

- A point-in-time rollforward operation is performed on a subset of the table spaces that were updated by a transaction, and that point in time precedes the time at which the transaction was committed.
- Any table contained in the table space being rolled forward to a point in time has an associated trigger, or is updated by a trigger that affects table spaces other than the one that is being rolled forward.

The solution is to find a suitable point in time that will prevent this from happening.

You can issue the `QUIESCE TABLESPACES FOR TABLE` command to create a transaction-consistent point in time for rolling table spaces forward. The quiesce request (in share, intent to update, or exclusive mode) waits (through locking) for all running transactions against those table spaces to complete, and blocks new requests. When the quiesce request is granted, the table spaces are in a consistent state. To determine a suitable time to stop the rollforward operation, you can look in the recovery history file to find quiesce points, and check whether they occur after the minimum recovery time.

After a table space point-in-time rollforward operation completes, the table space is put in backup pending state. You must take a backup of the table space, because

all updates made to it between the point in time to which you rolled forward and the current time have been removed. You can no longer roll the table space forward to the current time from a previous database- or table space-level backup image. The following example shows why the table space-level backup image is required, and how it is used. (To make the table space available, you can either back up the entire database, the table space that is in backup pending state, or a set of table spaces that includes the table space that is in backup pending state.)

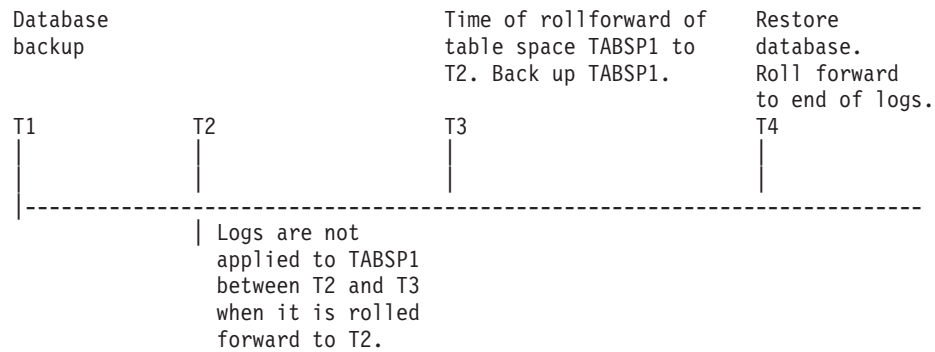


Figure 15. Table Space Backup Requirement

In the preceding example, the database is backed up at time T1. Then, at time T3, table space TABSP1 is rolled forward to a specific point in time (T2). The table space is backed up after time T3. Because the table space is in backup pending state, this backup operation is mandatory. The time stamp of the table space backup image is after time T3, but the table space is at time T2. Log records from between T2 and T3 are not applied to TABSP1. At time T4, the database is restored, using the backup image created at T1, and rolled forward to the end of the logs. Table space TABSP1 is put in restore pending state at time T3, because the database manager assumes that operations were performed on TABSP1 between T3 and T4 without the log changes between T2 and T3 having been applied to the table space. If these log changes were in fact applied as part of the rollforward operation against the database, this assumption would be incorrect. The table space-level backup that must be taken after the table space is rolled forward to a point in time allows you to roll that table space forward past a previous point-in-time rollforward operation (T3 in the example).

Assuming that you want to recover table space TABSP1 to T4, you would restore the table space from a backup image that was taken after T3 (either the required backup, or a later one), then roll TABSP1 forward to the end of the logs.

In the preceding example, the most efficient way of restoring the database to time T4 would be to perform the required steps in the following order:

1. Restore the database.
2. Restore the table space.
3. Roll the database forward.
4. Roll the table space forward.

Because you restore the table space before rolling the database forward, resource is not used to apply log records to the table space when the database is rolled forward.

If you cannot find the TABSP1 backup image that follows time T3, or you want to restore TABSP1 to T3 (or earlier), you can:

- Roll the table space forward to T3. You do not need to restore the table space again, because it was restored from the database backup image.
- Restore the table space again, using the database backup taken at time T1, then roll the table space forward to a time that precedes time T3.
- Drop the table space.

In a partitioned database environment:

- You must simultaneously roll all parts of a table space forward to the same point in time at the same time. This ensures that the table space is consistent across database partitions.
- If some database partitions are in rollforward pending state, and on other database partitions, some table spaces are in rollforward pending state (but the database partitions are not), you must first roll the database partitions forward, and then roll the table spaces forward.
- If you intend to roll a table space forward to the end of the logs, you do not have to restore it at each database partition; you only need to restore it at the database partitions that require recovery. If you intend to roll a table space forward to a point in time, however, you must restore it at each database partition.

**Related concepts:**

- “Recovering data with the load copy location file” on page 120

**Related reference:**

- “ROLLFORWARD DATABASE” on page 123

## Recovering a dropped table

You may occasionally drop a table whose data you still need. If this is the case, you should consider making your critical tables recoverable following a drop table operation.

You could recover the table data by invoking a database restore operation, followed by a database rollforward operation to a point in time before the table was dropped. This may be time-consuming if the database is large, and your data will be unavailable during recovery.

The dropped table recovery feature lets you recover your dropped table data using table space-level restore and rollforward operations. This will be faster than database-level recovery, and your database will remain available to users.

**Prerequisites:**

For a dropped table to be recoverable, the table space in which the table resides must have the DROPPED TABLE RECOVERY option turned on. This can be done during table space creation, or by invoking the ALTER TABLESPACE statement. The DROPPED TABLE RECOVERY option is table space-specific and limited to regular table spaces. To determine if a table space is enabled for dropped table recovery, you can query the DROP\_RECOVERY column in the SYSCAT.TABLESPACES catalog table. Dropped table recovery is enabled by default for newly created data table spaces.

When a DROP TABLE statement is run against a table whose table space is enabled for dropped table recovery, an additional entry (identifying the dropped



table) is made in the log files. An entry is also made in the recovery history file, containing information that can be used to recreate the table.

### Restrictions:

There are some restrictions on the type of data that is recoverable from a dropped table. It is not possible to recover:

- Large object (LOB) or long field data. The DROPPED TABLE RECOVERY option is not supported for large table spaces. If you attempt to recover a dropped table that contains LOB or LONG VARCHAR columns, these columns will be set to NULL in the generated export file. The DROPPED TABLE RECOVERY option can only be used for regular table spaces, not for temporary or large table spaces.
- The metadata associated with row types. (The data is recovered, but not the metadata.) The data in the hierarchy table of the typed table will be recovered. This data may contain more information than appeared in the typed table that was dropped.

If the data being recovered is of the GRAPHIC or VARGRAPHIC data type, it may include more than one code page. In order to recover this data, you need to specify the USEGRAPHICCODEPAGE file type modifier of the IMPORT or LOAD commands. In this case, using the LOAD command to recover the data will increase the performance of the recovery operation.

### Procedure:

Only one dropped table can be recovered at a time. You can recover a dropped table by doing the following:

1. Identify the dropped table by invoking the LIST HISTORY DROPPED TABLE command. The dropped table ID is listed in the Backup ID column.
2. Restore a database- or table space-level backup image taken before the table was dropped.
3. Create an export directory to which files containing the table data are to be written. This directory must either be accessible to all database partitions, or exist on each partition. Subdirectories under this export directory are created automatically by each database partition. These subdirectories are named NODE $nnnn$ , where  $nnnn$  represents the database partition or node number. Data files containing the dropped table data as it existed on each database partition are exported to a lower subdirectory called data. For example, \export\_directory\NODE0000\data.
4. Roll forward to a point in time after the table was dropped, using the RECOVER DROPPED TABLE option on the ROLLFORWARD DATABASE command. Alternatively, roll forward to the end of the logs, so that updates to other tables in the table space or database are not lost.
5. Recreate the table using the CREATE TABLE statement from the recovery history file.
6. Import the table data that was exported during the rollforward operation into the table.

The names of linked files associated with DATALINK columns *can* be recovered. After importing the table data, the table should be reconciled with the DB2 Data Links Manager. Backup images of the files may or may not be restored by the DB2 Data Links Manager, depending on whether garbage collection has already deleted them.

**Related reference:**

- “ALTER TABLESPACE statement” in the *SQL Reference, Volume 2*
- “CREATE TABLE statement” in the *SQL Reference, Volume 2*
- “ROLLFORWARD DATABASE” on page 123
- “LIST HISTORY” on page 258

---

## Recovering data with the load copy location file

The DB2LOADREC registry variable is used to identify the file with the load copy location information. This file is used during rollforward recovery to locate the load copy. It has information about:

- Media type
- Number of media devices to be used
- Location of the load copy generated during a table load operation
- File name of the load copy, if applicable

If the location file does not exist, or no matching entry is found in the file, the information from the log record is used.

The information in the file might be overwritten before rollforward recovery takes place.

**Notes:**

1. In a partitioned database environment, the DB2LOADREC registry variable must be set for all the database partition servers using the **db2set** command.
2. In a partitioned database environment, the load copy file must exist at each database partition server, and the file name (including the path) must be the same.
3. If an entry in the file identified by the DB2LOADREC registry variable is not valid, the old load copy location file is used to provide information to replace the invalid entry.

The following information is provided in the location file. The first five parameters must have valid values, and are used to identify the load copy. The entire structure is repeated for each load copy recorded. For example:

TIMestamp	19950725182542	* Time stamp generated at load time
SCHema	PAYROLL	* Schema of table loaded
TABlename	EMPLOYEES	* Table name
DATabasename	DBT	* Database name
DB2instance	toronto	* DB2INSTANCE
BUFFernumber	NULL	* Number of buffers to be used for recovery
SESSionnumber	NULL	* Number of sessions to be used for recovery
TYPeofmedia	L	* Type of media - L for local device A for TSM 0 for other vendors
LOCationnumber	3	* Number of locations
ENTry	/u/toronto/dbt.payroll.employees.001	
ENT	/u/toronto/dbt.payroll.employees.002	
ENT	/dev/rmt0	
TIM	19950725192054	
SCH	PAYROLL	
TAB	DEPT	
DAT	DBT	
DB2®	toronto	

BUF	NULL
SES	NULL
TYP	A
TIM	19940325192054
SCH	PAYROLL
TAB	DEPT
DAT	DBT
DB2	toronto
BUF	NULL
SES	NULL
TYP	0
SHRlib	/@sys/lib/backup_vendor.a

**Notes:**

1. The first three characters in each keyword are significant. All keywords are required in the specified order. Blank lines are not accepted.
2. The time stamp is in the form *yyyymmddhhmmss*.
3. All fields are mandatory, except for BUF and SES, which can be NULL. If SES is NULL, the value specified by the *numloadrecses* configuration parameter is used. If BUF is NULL, the default value is SES+2.
4. If even one of the entries in the location file is invalid, the previous load copy location file is used to provide those values.
5. The media type can be local device (L for tape, disk or diskettes), TSM (A), or other vendor (0). If the type is L, the number of locations, followed by the location entries, is required. If the type is A, no further input is required. If the type is 0, the shared library name is required.
6. The SHRlib parameter points to a library that has a function to store the load copy data.
7. If you invoke a load operation, specifying the COPY NO or the NONRECOVERABLE option, and do not take a backup copy of the database or affected table spaces after the operation completes, you cannot restore the database or table spaces to a point in time that follows the load operation. That is, you cannot use rollforward recovery to rebuild the database or table spaces to the state they were in following the load operation. You can only restore the database or table spaces to a point in time that precedes the load operation.

If you want to use a particular load copy, you can use the recovery history file for the database to determine the time stamp for that specific load operation. In a partitioned database environment, the recovery history file is local to each database partition.

**Related reference:**

- Appendix F, “Tivoli Storage Manager,” on page 325

---

## Synchronizing clocks in a partitioned database system

You should maintain relatively synchronized system clocks across the database partition servers to ensure smooth database operations and unlimited forward recoverability. Time differences among the database partition servers, plus any potential operational and communications delays for a transaction should be less than the value specified for the *max\_time\_diff* (maximum time difference among nodes) database manager configuration parameter.

To ensure that the log record time stamps reflect the sequence of transactions in a partitioned database system, DB2<sup>®</sup> uses the system clock on each machine as the basis for the time stamps in the log records. If, however, the system clock is set

ahead, the log clock is automatically set ahead with it. Although the system clock can be set back, the clock for the logs cannot, and remains at the *same* advanced time until the system clock matches this time. The clocks are then in synchrony. The implication of this is that a short term system clock error on a database node can have a long lasting effect on the time stamps of database logs.

For example, assume that the system clock on database partition server A is mistakenly set to November 7, 1999 when the year is 1997, and assume that the mistake is corrected *after* an update transaction is committed in the partition at that database partition server. If the database is in continual use, and is regularly updated over time, any point between November 7, 1997 and November 7, 1999 is virtually unreachable through rollforward recovery. When the COMMIT on database partition server A completes, the time stamp in the database log is set to 1999, and the log clock remains at November 7, 1999 until the system clock matches this time. If you attempt to roll forward to a point in time within this time frame, the operation will stop at the first time stamp that is beyond the specified stop point, which is November 7, 1997.

Although DB2 cannot control updates to the system clock, the *max\_time\_diff* database manager configuration parameter reduces the chances of this type of problem occurring:

- The configurable values for this parameter range from 1 minute to 24 hours.
- When the first connection request is made to a non-catalog node, the database partition server sends its time to the catalog node for the database. The catalog node then checks that the time on the node requesting the connection, and its own time are within the range specified by the *max\_time\_diff* parameter. If this range is exceeded, the connection is refused.
- An update transaction that involves more than two database partition servers in the database must verify that the clocks on the participating database partition servers are in synchrony before the update can be committed. If two or more database partition servers have a time difference that exceeds the limit allowed by *max\_time\_diff*, the transaction is rolled back to prevent the incorrect time from being propagated to other database partition servers.

**Related reference:**

- “*max\_time\_diff* - Maximum time difference among nodes configuration parameter” in the *Administration Guide: Performance*

---

## Client/server timestamp conversion

This section explains the generation of timestamps in a client/server environment:

- If you specify a local time for a rollforward operation, all messages returned will also be in local time.

**Note:** All times are converted on the server and (in partitioned database environments) on the catalog node.

- The timestamp string is converted to GMT on the server, so the time represents the server’s time zone, not the client’s. If the client is in a different time zone from the server, the server’s local time should be used.
- If the timestamp string is close to the time change due to daylight savings time, it is important to know whether the stop time is before or after the time change so that it is specified correctly.

**Related concepts:**

- “Rollforward overview” on page 111
- “Synchronizing clocks in a partitioned database system” on page 121

## ROLLFORWARD DATABASE

Recovers a database by applying transactions recorded in the database log files. Invoked after a database or a table space backup image has been restored, or if any table spaces have been taken offline by the database due to a media error. The database must be recoverable (that is, the *logarchmeth1* or *logarchmeth2* database configuration parameters must be set to a value other than OFF) before the database can be recovered with rollforward recovery.

### Scope:

In a partitioned database environment, this command can only be invoked from the catalog partition. A database or table space rollforward operation to a specified point in time affects all partitions that are listed in the *db2nodes.cfg* file. A database or table space rollforward operation to the end of logs affects the partitions that are specified. If no partitions are specified, it affects all partitions that are listed in the *db2nodes.cfg* file; if rollforward recovery is not needed on a particular partition, that partition is ignored.

### Authorization:

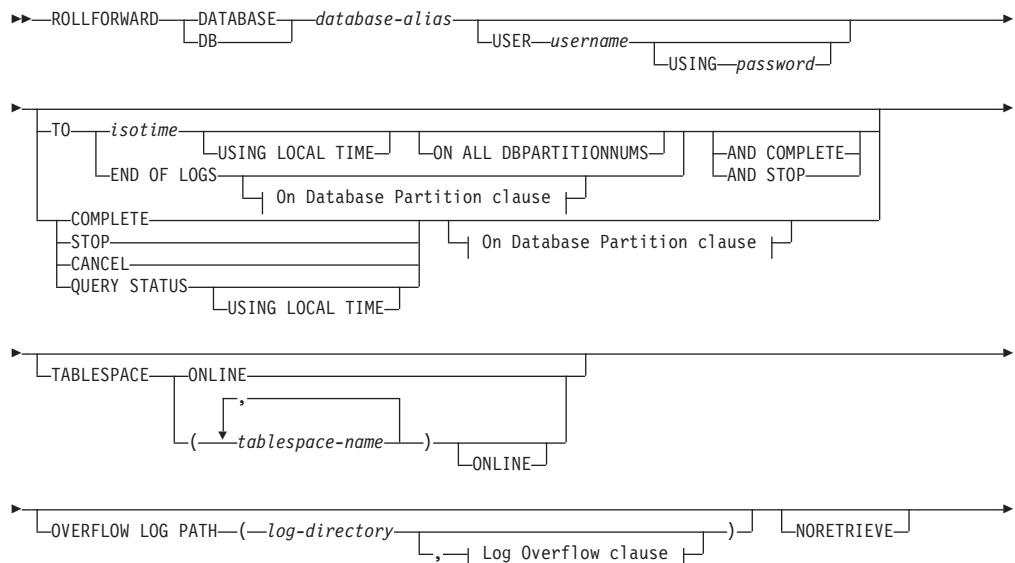
One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

### Required connection:

None. This command establishes a database connection.

### Command syntax:



## ROLLFORWARD DATABASE

RECOVER DROPPED TABLE—*drop-table-id*—TO—*export-directory*

### On Database Partition clause:

ON Database Partition List clause  
ALL DBPARTITIONNUMS EXCEPT Database Partition List clause

### Database Partition List clause:

DBPARTITIONNUM ( DBPARTITIONNUMS

,  
db-partition-number1 TO db-partition-number2 )

### Log Overflow clause:

,  
log-directory ON DBPARTITIONNUM db-partition-number1

### Command parameters:

#### DATABASE *database-alias*

The alias of the database that is to be rollforward recovered.

#### USER *username*

The user name under which the database is to be rollforward recovered.

#### USING *password*

The password used to authenticate the user name. If the password is omitted, the user is prompted to enter it.

#### TO

##### *isotime*

The point in time to which all committed transactions are to be rolled forward (including the transaction committed precisely at that time, as well as all transactions committed previously).

This value is specified as a time stamp, a 7-part character string that identifies a combined date and time. The format is *yyyy-mm-dd-hh.mm.ss.nnnnnn* (year, month, day, hour, minutes, seconds, microseconds), expressed in Coordinated Universal Time (UTC). UTC helps to avoid having the same time stamp associated with different logs (because of a change in time associated with daylight savings time, for example). The time stamp in a backup image is based on the local time at which the backup operation started. The CURRENT TIMEZONE special register specifies the difference between UTC and local time at the application server. The difference is represented by a time duration (a decimal number in which the first two digits represent the number of hours, the next two digits represent the number of minutes, and the last two

digits represent the number of seconds). Subtracting CURRENT TIMEZONE from a local time converts that local time to UTC.

### USING LOCAL TIME

Allows the user to rollforward to a point in time that is the user's local time rather than GMT time. This makes it easier for users to rollforward to a specific point in time on their local machines, and eliminates potential user errors due to the translation of local to GMT time.

#### Notes:

1. If the user specifies a local time for rollforward, all messages returned to the user will also be in local time. Note that all times are converted on the server, and in partitioned database environments, on the catalog database partition.
2. The timestamp string is converted to GMT on the server, so the time is local to the server's time zone, not the client's. If the client is in one time zone and the server in another, the server's local time should be used. This is different from the local time option from the Control Center, which is local to the client.
3. If the timestamp string is close to the time change of the clock due to daylight savings, it is important to know if the stop time is before or after the clock change, and specify it correctly.

### END OF LOGS

Specifies that all committed transactions from all online archive log files listed in the database configuration parameter *logpath* are to be applied.

### ALL DBPARTITIONNUMS

Specifies that transactions are to be rolled forward on all partitions specified in the *db2nodes.cfg* file. This is the default if a database partition clause is not specified.

### EXCEPT

Specifies that transactions are to be rolled forward on all partitions specified in the *db2nodes.cfg* file, except those specified in the database partition list.

### ON DBPARTITIONNUM / ON DBPARTITIONNUMS

Roll the database forward on a set of database partitions.

#### **db-partition-number1**

Specifies a database partition number in the database partition list.

#### **db-partition-number2**

Specifies the second database partition number, so that all partitions from *db-partition-number1* up to and including *db-partition-number2* are included in the database partition list.

### COMPLETE / STOP

Stops the rolling forward of log records, and completes the rollforward recovery process by rolling back any incomplete transactions and turning off the rollforward pending state of the database. This allows access to the database or table spaces that are being rolled forward. These keywords are equivalent; specify one or the other, but not both. The keyword AND permits specification of multiple operations at once; for example, `db2 rollforward db sample to end of logs and complete`.

## ROLLFORWARD DATABASE

**Note:** When rolling table spaces forward to a point in time, the table spaces are placed in backup pending state.

### CANCEL

Cancels the rollforward recovery operation. This puts the database or one or more table spaces on all partitions on which forward recovery has been started in restore pending state:

- If a *database* rollforward operation is not in progress (that is, the database is in rollforward pending state), this option puts the database in restore pending state.
- If a *table space* rollforward operation is not in progress (that is, the table spaces are in rollforward pending state), a table space list must be specified. All table spaces in the list are put in restore pending state.
- If a table space rollforward operation *is* in progress (that is, at least one table space is in rollforward in progress state), all table spaces that are in rollforward in progress state are put in restore pending state. If a table space list is specified, it must include all table spaces that are in rollforward in progress state. All table spaces on the list are put in restore pending state.
- If rolling forward to a point in time, any table space name that is passed in is ignored, and all table spaces that are in rollforward in progress state are put in restore pending state.
- If rolling forward to the end of the logs with a table space list, only the table spaces listed are put in restore pending state.

This option cannot be used to cancel a rollforward operation *that is actually running*. It can only be used to cancel a rollforward operation that is in progress but not actually running at the time. A rollforward operation can be in progress but not running if:

- It terminated abnormally.
- The STOP option was not specified.
- An error caused it to fail. Some errors, such as rolling forward through a non-recoverable load operation, can put a table space into restore pending state.

**Note:** Use this option with caution, and only if the rollforward operation that is in progress cannot be completed because some of the table spaces have been put in rollforward pending state or in restore pending state. When in doubt, use the LIST TABLESPACES command to identify the table spaces that are in rollforward in progress state, or in rollforward pending state.

### QUERY STATUS

Lists the log files that the database manager has rolled forward, the next archive file required, and the time stamp (in CUT) of the last committed transaction since rollforward processing began. In a partitioned database environment, this status information is returned for each partition. The information returned contains the following fields:

#### Database partition number

#### Rollforward status

Status can be: database or table space rollforward pending, database or table space rollforward in progress, database or table space rollforward processing STOP, or not pending.



### Next log file to be read

A string containing the name of the next required log file. In a partitioned database environment, use this information if the rollforward utility fails with a return code indicating that a log file is missing or that a log information mismatch has occurred.

### Log files processed

A string containing the names of processed log files that are no longer needed for recovery, and that can be removed from the directory. If, for example, the oldest uncommitted transaction starts in log file *x*, the range of obsolete log files will not include *x*; the range ends at *x* - 1.

### Last committed transaction

A string containing a time stamp in ISO format (*yyyy-mm-dd-hh.mm.ss*). This time stamp marks the last transaction committed after the completion of rollforward recovery. The time stamp applies to the database. For table space rollforward recovery, it is the time stamp of the last transaction committed to the database.

**Note:** QUERY STATUS is the default value if the TO, STOP, COMPLETE, or CANCEL clauses are omitted. If TO, STOP, or COMPLETE was specified, status information is displayed if the command has completed successfully. If individual table spaces are specified, they are ignored; the status request does not apply only to specified table spaces.

## TABLESPACE

This keyword is specified for table space-level rollforward recovery.

### tablespace-name

Mandatory for table space-level rollforward recovery to a point in time. Allows a subset of table spaces to be specified for rollforward recovery to the end of the logs. In a partitioned database environment, each table space in the list does not have to exist at each partition that is rolling forward. If it *does* exist, it must be in the correct state.

## ONLINE

This keyword is specified to allow table space-level rollforward recovery to be done online. This means that other agents are allowed to connect while rollforward recovery is in progress.

## OVERFLOW LOG PATH log-directory

Specifies an alternate log path to be searched for archived logs during recovery. Use this parameter if log files were moved to a location other than that specified by the *logpath* database configuration parameter. In a partitioned database environment, this is the (fully qualified) default overflow log path *for all partitions*. A relative overflow log path can be specified for single-partition databases.

**Note:** The OVERFLOW LOG PATH command parameter will overwrite the value (if any) of the database configuration parameter OVERFLOWLOGPATH.

## log-directory ON DBPARTITIONNUM

In a partitioned database environment, allows a different log path to override the default overflow log path for a specific partition.

## ROLLFORWARD DATABASE

### NORETRIEVE

Allows the user to control which log files to be rolled forward on the standby machine by allowing the user to disable the retrieval of archived logs. The benefits of this are:

- By controlling the logfiles to be rolled forward, one can ensure that the standby machine is X hours behind the production machine, to prevent the user affecting both systems.
- If the standby system does not have access to archive (eg. if TSM is the archive, it only allows the original machine to retrieve the files)
- It might also be possible that while the production system is archiving a file, the standby system is retrieving the same file, and it might then get an incomplete log file. Noretrieve would solve this problem.

### RECOVER DROPPED TABLE **drop-table-id**

Recovers a dropped table during the rollforward operation. The table ID can be obtained using the LIST HISTORY command.

### TO **export-directory**

Specifies a directory to which files containing the table data are to be written. The directory must be accessible to all database partitions.

### Examples:

#### Example 1

The ROLLFORWARD DATABASE command permits specification of multiple operations at once, each being separated with the keyword AND. For example, to roll forward to the end of logs, and complete, the separate commands:

```
db2 rollforward db sample to end of logs
db2 rollforward db sample complete
```

can be combined as follows:

```
db2 rollforward db sample to end of logs and complete
```

Although the two are equivalent, it is recommended that such operations be done in two steps. It is important to verify that the rollforward operation has progressed as expected, before stopping it and possibly missing logs. This is especially important if a bad log is found during rollforward recovery, and the bad log is interpreted to mean the "end of logs". In such cases, an undamaged backup copy of that log could be used to continue the rollforward operation through more logs.

#### Example 2

Roll forward to the end of the logs (two table spaces have been restored):

```
db2 rollforward db sample to end of logs
db2 rollforward db sample to end of logs and stop
```

These two statements are equivalent. Neither AND STOP or AND COMPLETE is needed for table space rollforward recovery to the end of the logs. Table space names are not required. If not specified, all table spaces requiring rollforward recovery will be included. If only a subset of these table spaces is to be rolled forward, their names must be specified.

#### Example 3

After three table spaces have been restored, roll one forward to the end of the logs, and the other two to a point in time, both to be done online:

```
db2 rollforward db sample to end of logs tablespace(TBS1) online
```

```
db2 rollforward db sample to 1998-04-03-14.21.56.245378 and stop  
tablespace(TBS2, TBS3) online
```

Note that two rollforward operations cannot be run concurrently. The second command can only be invoked after the first rollforward operation completes successfully.

### Example 4

After restoring the database, roll forward to a point in time, using `OVERFLOW LOG PATH` to specify the directory where the user exit saves archived logs:

```
db2 rollforward db sample to 1998-04-03-14.21.56.245378 and stop  
overflow log path (/logs)
```

### Example 5 (partitioned database environments)

There are three database partitions: 0, 1, and 2. Table space TBS1 is defined on all partitions, and table space TBS2 is defined on partitions 0 and 2. After restoring the database on database partition 1, and TBS1 on database partitions 0 and 2, roll the database forward on database partition 1:

```
db2 rollforward db sample to end of logs and stop
```

This returns warning SQL1271 ("Database is recovered but one or more table spaces are off-line on database partition(s) 0 and 2.").

```
db2 rollforward db sample to end of logs
```

This rolls TBS1 forward on database partitions 0 and 2. The clause `TABLESPACE(TBS1)` is optional in this case.

### Example 6 (partitioned database environments)

After restoring table space TBS1 on database partitions 0 and 2 only, roll TBS1 forward on database partitions 0 and 2:

```
db2 rollforward db sample to end of logs
```

Database partition 1 is ignored.

```
db2 rollforward db sample to end of logs tablespace(TBS1)
```

This fails, because TBS1 is not ready for rollforward recovery on database partition 1. Reports SQL4906N.

```
db2 rollforward db sample to end of logs on dbpartitionnums (0, 2)  
tablespace(TBS1)
```

This completes successfully.

```
db2 rollforward db sample to 1998-04-03-14.21.56.245378 and stop  
tablespace(TBS1)
```

This fails, because TBS1 is not ready for rollforward recovery on database partition 1; all pieces must be rolled forward together.

## ROLLFORWARD DATABASE

**Note:** With table space rollforward to a point in time, the database partition clause is not accepted. The rollforward operation must take place on all the database partitions on which the table space resides.

After restoring TBS1 on database partition 1:

```
db2 rollforward db sample to 1998-04-03-14.21.56.245378 and stop
tablespace(TBS1)
```

This completes successfully.

### Example 7 (partitioned database environment)

After restoring a table space on all database partitions, roll forward to PIT2, but do not specify AND STOP. The rollforward operation is still in progress. Cancel and roll forward to PIT1:

```
db2 rollforward db sample to pit2 tablespace(TBS1)
db2 rollforward db sample cancel tablespace(TBS1)
```

```
** restore TBS1 on all database partitions **
```

```
db2 rollforward db sample to pit1 tablespace(TBS1)
db2 rollforward db sample stop tablespace(TBS1)
```

### Example 8 (partitioned database environments)

Rollforward recover a table space that resides on eight database partitions (3 to 10) listed in the db2nodes.cfg file:

```
db2 rollforward database dwtest to end of logs tablespace (tssprodt)
```

This operation to the end of logs (not point in time) completes successfully. The database partitions on which the table space resides do not have to be specified. The utility defaults to the db2nodes.cfg file.

### Example 9 (partitioned database environment)

Rollforward recover six small table spaces that reside on a single-partition database partition group (on database partition 6):

```
db2 rollforward database dwtest to end of logs on dbpartitionnum (6)
tablespace(tsstore, tssbuyer, tsstime, tsswhse, tsslscat, tssvendor)
```

This operation to the end of logs (not point in time) completes successfully.

### Usage notes:

If restoring from an image that was created during an online backup operation, the specified point in time for the rollforward operation must be later than the time at which the online backup operation completed. If the rollforward operation is stopped before it passes this point, the database is left in rollforward pending state. If a table space is in the process of being rolled forward, it is left in rollforward in progress state.

If one or more table spaces is being rolled forward to a point in time, the rollforward operation must continue at least to the minimum recovery time, which is the last update to the system catalogs for this table space or its tables. The minimum recovery time (in Coordinated Universal Time, or UTC) for a table space can be retrieved using the LIST TABLESPACES SHOW DETAIL command.

Rolling databases forward may require a load recovery using tape devices. If prompted for another tape, the user can respond with one of the following:

- c** Continue. Continue using the device that generated the warning message (for example, when a new tape has been mounted)
- d** Device terminate. Stop using the device that generated the warning message (for example, when there are no more tapes)
- t** Terminate. Terminate all devices.

If the rollforward utility cannot find the next log that it needs, the log name is returned in the SQLCA, and rollforward recovery stops. If no more logs are available, use the STOP option to terminate rollforward recovery. Incomplete transactions are rolled back to ensure that the database or table space is left in a consistent state.

### Compatibilities:

For compatibility with versions earlier than Version 8:

- The keyword NODE can be substituted for DBPARTITIONNUM.
- The keyword NODES can be substituted for DBPARTITIONNUMS.

### Related reference:

- “BACKUP DATABASE” on page 66
- “RESTORE DATABASE” on page 87

---

## db2Rollforward - Rollforward Database

Recovers a database by applying transactions recorded in the database log files. Called after a database or a table space backup has been restored, or if any table spaces have been taken offline by the database due to a media error. The database must be recoverable (that is, the *logarchmeth1* database configuration parameter must be set to on) before the database can be recovered with rollforward recovery.

### Scope:

In a partitioned database environment, this API can only be called from the catalog partition. A database or table space rollforward call specifying a point-in-time affects all database partition servers that are listed in the *db2nodes.cfg* file. A database or table space rollforward call specifying end of logs affects the database partition servers that are specified. If no database partition servers are specified, it affects all database partition servers that are listed in the *db2nodes.cfg* file; if no roll forward is needed on a particular database partition server, that database partition server is ignored.

### Authorization:

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

### Required connection:

## db2Rollforward - Rollforward Database

None. This API establishes a database connection.

### API include file:

*db2ApiDf.h*

### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2Rollforward */
/* ... */
SQL_API_RC SQL_API_FN
db2Rollforward (
    db2UInt32 versionNumber,
    void *pDB2RollforwardStruct,
    struct sqlca *pSqlca);

typedef SQL_STRUCTURE db2RollforwardStruct
{
    struct db2RfwdInputStruct *piRfwdInput;
    struct db2RfwdOutputStruct *poRfwdOutput;
} db2RollforwardStruct;

typedef SQL_STRUCTURE db2RfwdInputStruct
{
    sqluint32          iVersion;
    char               *piDbAlias;
    db2UInt32          iCallerAction;
    char               *piStopTime;
    char               *piUserName;
    char               *piPassword;
    char               *piOverflowLogPath;
    db2UInt32          iNumChngLgOvrflw;
    struct sqlurf_newlogpath *piChngLogOvrflw;
    db2UInt32          iConnectMode;
    struct sqlu_tablespace_bkrst_list *piTablespaceList;
    db2int32           iAllNodeFlag;
    db2int32           iNumNodes;
    SQL_PDB_NODE_TYPE *piNodeList;
    db2int32           iNumNodeInfo;
    char               *piDroppedTblID;
    char               *piExportDir;
    db2UInt32          iRollforwardFlags;
} db2RfwdInputStruct;

typedef SQL_STRUCTURE db2RfwdOutputStruct
{
    char               *poApplicationId;
    sqlint32           *poNumReplies;
    struct sqlurf_info *poNodeInfo;
} db2RfwdOutputStruct;

typedef SQL_STRUCTURE sqlurf_newlogpath
{
    SQL_PDB_NODE_TYPE nodenum;
    unsigned short     pathlen;
    char               logpath[SQL_LOGPATH_SZ+SQL_LOGFILE_NAME_SZ+1];
} sqlurf_newlogpath;

typedef SQL_STRUCTURE sqlu_tablespace_bkrst_list
{
    long               num_entry;
    struct sqlu_tablespace_entry *tablespace;
} sqlu_tablespace_bkrst_list;

typedef SQL_STRUCTURE sqlu_tablespace_entry
```

```

{
    sqluint32          reserve_len;
    char              tablespace_entry[SQLU_MAX_TBS_NAME_LEN+1];
    char              filler[1];
} sqlu_tablespace_entry;

typedef SQL_STRUCTURE sqlurf_info
{
    SQL_PDB_NODE_TYPE nodenum;
    sqlint32          state;
    unsigned char     nextarclog[SQLUM_ARCHIVE_FILE_LEN+1];
    unsigned char     firstarcdel[SQLUM_ARCHIVE_FILE_LEN+1];
    unsigned char     lastarcdel[SQLUM_ARCHIVE_FILE_LEN+1];
    unsigned char     lastcommit[SQLUM_TIMESTAMP_LEN+1];
} sqlurf_info;
/* ... */

```

**Generic API syntax:**

```

/* File: db2ApiDf.h */
/* API: db2Rollforward */
/* ... */
SQL_API_RC SQL_API_FN
db2gRollforward (
    db2Uint32 versionNumber,
    void *pDB2gRollforwardStruct,
    struct sqlca *pSqlca);

typedef SQL_STRUCTURE db2gRollforwardStruct
{
    struct db2gRfwdInputStruct *piRfwdInput;
    struct db2RfwdOutputStruct *poRfwdOutput;
} db2gRollforwardStruct;

SQL_STRUCTURE db2gRfwdInputStruct
{
    db2Uint32          iDbAliasLen;
    db2Uint32          iStopTimeLen;
    db2Uint32          iUserNameLen;
    db2Uint32          iPasswordLen;
    db2Uint32          iOvrflwLogPathLen;
    db2Uint32          iDroppedTblIDLen;
    db2Uint32          iExportDirLen;
    sqluint32          iVersion;
    char               *piDbAlias;
    db2Uint32          iCallerAction;
    char               *piStopTime;
    char               *piUserName;
    char               *piPassword;
    char               *piOverflowLogPath;
    db2Uint32          iNumChngLgOvrflw;
    struct sqlurf_newlogpath *piChngLogOvrflw;
    db2Uint32          iConnectMode;
    struct sqlu_tablespace_bkrst_list *piTablespaceList;
    db2int32           iAllNodeFlag;
    db2int32           iNumNodes;
    SQL_PDB_NODE_TYPE *piNodeList;
    db2int32           iNumNodeInfo;
    char               *piDroppedTblID;
    char               *piExportDir;
    db2Uint32          iRollforwardFlags;
} db2gRfwdInputStruct;

typedef SQL_STRUCTURE db2RfwdOutputStruct
{
    char               *poApplicationId;
    sqlint32           poNumReplies;
}

```

## db2Rollforward - Rollforward Database

```
    struct sqlurf_info          *poNodeInfo;
} db2RfwdOutputStruct;

typedef SQL_STRUCTURE sqlurf_newlogpath
{
    SQL_PDB_NODE_TYPE          nodenum;
    unsigned short             pathlen;
    char                       logpath[SQL_LOGPATH_SZ+SQL_LOGFILE_NAME_SZ+1];
} sqlurf_newlogpath;

typedef SQL_STRUCTURE sqlu_tablespace_bkrst_list
{
    long                       num_entry;
    struct sqlu_tablespace_entry *tablespace;
} sqlu_tablespace_bkrst_list;

typedef SQL_STRUCTURE sqlu_tablespace_entry
{
    sqluint32                  reserve_len;
    char                       tablespace_entry[SQLU_MAX_TBS_NAME_LEN+1];
    char                       filler[1];
} sqlu_tablespace_entry;

typedef SQL_STRUCTURE sqlurf_info
{
    SQL_PDB_NODE_TYPE          nodenum;
    sqlint32                   state;
    unsigned char               nextarclog[SQLUM_ARCHIVE_FILE_LEN+1];
    unsigned char               firstarcdel[SQLUM_ARCHIVE_FILE_LEN+1];
    unsigned char               lastarcdel[SQLUM_ARCHIVE_FILE_LEN+1];
    unsigned char               lastcommit[SQLUM_TIMESTAMP_LEN+1];
} sqlurf_info;
/* ... */
```

### API parameters:

#### versionNumber

Input. Specifies the version and release level of the structure passed as the second parameter.

#### pDB2RollforwardStruct

Input. A pointer to the *db2RollforwardStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### piRfwdInput

Input. A pointer to the *db2RfwdInputStruct* structure.

#### poRfwdOutput

Output. A pointer to the *db2RfwdOutputStruct* structure.

#### iDbAliasLen

Input. Specifies the length in bytes of the database alias.

#### iStopTimeLen

Input. Specifies the length in bytes of the stop time parameter. Set to zero if no stop time is provided.

#### iUserNameLen

Input. Specifies the length in bytes of the user name. Set to zero if no user name is provided.

#### iPasswordLen

Input. Specifies the length in bytes of the password. Set to zero if no password is provided.



**iOverflowLogPathLen**

Input. Specifies the length in bytes of the overflow log path. Set to zero if no overflow log path is provided.

**iVersion**

Input. The version ID of the rollforward parameters. It is defined as SQLUM\_RFWD\_VERSION.

**piDbAlias**

Input. A string containing the database alias. This is the alias that is cataloged in the system database directory.

**iCallerAction**

Input. Specifies action to be taken. Valid values (defined in db2ApiDf.h) are:

**DB2ROLLFORWARD\_ROLLFWD**

Rollforward to the point in time specified by *piStopTime*. For database rollforward, the database is left in *rollforward-pending* state. For table space rollforward to a point in time, the table spaces are left in *rollforward-in-progress* state.

**DB2ROLLFORWARD\_STOP**

End roll-forward recovery. No new log records are processed and uncommitted transactions are backed out. The *rollforward-pending* state of the database or table spaces is turned off. Synonym is DB2ROLLFORWARD\_RFWD\_COMPLETE.

**DB2ROLLFORWARD\_RFWD\_STOP**

Rollforward to the point in time specified by *piStopTime*, and end roll-forward recovery. The *rollforward-pending* state of the database or table spaces is turned off. Synonym is DB2ROLLFORWARD\_RFWD\_COMPLETE.

**DB2ROLLFORWARD\_QUERY**

Query values for *nextarclog*, *firstarclog*, *lastarclog*, and *lastcommit*. Return database status and a node number.

**DB2ROLLFORWARD\_PARM\_CHECK**

Validate parameters without performing the roll forward.

**DB2ROLLFORWARD\_CANCEL**

Cancel the rollforward operation that is currently running. The database or table space are put in recovery pending state.

**Note:** This option cannot be used while the rollforward is actually running. It can be used if the rollforward is paused (that is, waiting for a STOP), or if a system failure occurred during the rollforward. It should be used with caution.

Rolling databases forward may require a load recovery using tape devices. The rollforward API will return with a warning message if user intervention on a device is required. The API can be called again with one of the following three caller actions:

**DB2ROLLFORWARD\_LOADREC\_CONT**

Continue using the device that generated the warning message (for example, when a new tape has been mounted).

**DB2ROLLFORWARD\_DEVICE\_TERM**

Stop using the device that generated the warning message (for example, when there are no more tapes).

### DB2ROLLFORWARD\_LOAD\_REC\_TERM

Terminate all devices being used by load recovery.

### piStopTime

Input. A character string containing a time stamp in ISO format. Database recovery will stop when this time stamp is exceeded. Specify `SQLUM_INFINITY_TIMESTAMP` to roll forward as far as possible. May be NULL for `DB2ROLLFORWARD_QUERY`, `DB2ROLLFORWARD_PARM_CHECK`, and any of the load recovery (`DB2ROLLFORWARD_LOADREC_xxx`) caller actions.

### piUserName

Input. A string containing the user name of the application. May be NULL.

### piPassword

Input. A string containing the password of the supplied user name (if any). May be NULL.

### piOverflowLogPath

Input. This parameter is used to specify an alternate log path to be used. In addition to the active log files, archived log files need to be moved (by the user) into the *logpath* before they can be used by this utility. This can be a problem if the user does not have sufficient space in the *logpath*. The overflow log path is provided for this reason. During roll-forward recovery, the required log files are searched, first in the *logpath*, and then in the overflow log path. The log files needed for table space roll-forward recovery can be brought into either the *logpath* or the overflow log path. If the caller does not specify an overflow log path, the default value is the *logpath*. In a partitioned database environment, the overflow log path must be a valid, fully qualified path; the default path is the default overflow log path for each node. In a single-partition database environment, the overflow log path can be relative if the server is local.

### iNumChngLgOvrflw

Input. Partitioned database environments only. The number of changed overflow log paths. These new log paths override the default overflow log path for the specified database partition server only.

### piChngLogOvrflw

Input. Partitioned database environments only. A pointer to a structure containing the fully qualified names of changed overflow log paths. These new log paths override the default overflow log path for the specified database partition server only.

### iConnectMode

Input. Valid values (defined in `db2ApiDf.h`) are:

#### DB2ROLLFORWARD\_OFFLINE

Offline roll forward. This value must be specified for database roll-forward recovery.

#### DB2ROLLFORWARD\_ONLINE

Online roll forward.

### piTablespaceList

Input. A pointer to a structure containing the names of the table spaces to be rolled forward to the end-of-logs or to a specific point in time. If not specified, the table spaces needing rollforward will be selected.

### iAllNodeFlag

Input. Partitioned database environments only. Indicates whether the

rollforward operation is to be applied to all database partition servers defined in `db2nodes.cfg`. Valid values are:

### **DB2\_NODE\_LIST**

Apply to database partition servers in a list that is passed in *piNodeList*.

### **DB2\_ALL\_NODES**

Apply to all database partition servers. *piNodeList* should be NULL. This is the default value.

### **DB2\_ALL\_EXCEPT**

Apply to all database partition servers except those in a list that is passed in *piNodeList*.

### **DB2\_CAT\_NODE\_ONLY**

Apply to the catalog partition only. *piNodeList* should be NULL.

### **iNumNodes**

Input. Specifies the number of database partition servers in the *piNodeList* array.

### **piNodeList**

Input. A pointer to an array of database partition server numbers on which to perform the roll-forward recovery.

### **iNumNodeInfo**

Input. Defines the size of the output parameter *poNodeInfo*, which must be large enough to hold status information from each database partition that is being rolled forward. In a single-partition database environment, this parameter should be set to 1. The value of this parameter should be same as the number of database partition servers for which this API is being called.

### **piDroppedTblID**

Input. A string containing the ID of the dropped table whose recovery is being attempted.

### **piExportDir**

Input. The directory into which the dropped table data will be exported.

### **RollforwardFlags**

Input. Specifies the rollforward flags. Valid values (defined in `db2ApiDf.h`):

#### **DB2ROLLFORWARD\_EMPTY\_FLAG**

No flags specified.

#### **DB2ROLLFORWARD\_LOCAL\_TIME**

Allows the user to rollforward to a point in time that is the user's local time rather than GMT time. This makes it easier for users to rollforward to a specific point in time on their local machines, and eliminates potential user errors due to the translation of local to GMT time.

#### **DB2ROLLFORWARD\_NO\_RETRIEVE**

Controls which log files to be rolled forward on the standby machine by allowing the user to disable the retrieval of archived logs. By controlling the log files to be rolled forward, one can ensure that the standby machine is X hours behind the production machine, to prevent the user affecting both systems. This option is useful if the standby system does not have access to archive, for example, if TSM is the archive, it only allows the original machine

## db2Rollforward - Rollforward Database

to retrieve the files. It will also remove the possibility that the standby system would retrieve an incomplete log file while the production system is archiving a file and the standby system is retrieving the same file.

### **poApplicationId**

Output. The application ID.

### **poNumReplies**

Output. The number of replies received.

### **poNodeInfo**

Output. Database partition reply information.

### **nodenum**

Node number.

### **pathlen**

The length of the new logpath.

### **logpath**

The new overflow log path.

### **num\_entry**

Number of entries in the list pointed to by the *tablespace* field.

### **tablespace**

A pointer to the *sqlu\_tablepsace\_entry* structure.

### **reserve\_len**

Length of the character string provided in the *tablespace\_entry* field. For languages other than C.

### **tablespace\_entry**

Table space name.

### **state**

State information.

### **nextarclog**

A buffer to hold the returned name of the next archived log file required. If a caller action other than DB2ROLLFORWARD\_QUERY is supplied, the value returned in this field indicates that an error occurred when accessing the file. Possible causes are:

- The file was not found in the database log directory, nor on the path specified by the overflow log path parameter
- The log archiving method failed to return the archived file.

### **firstarcdel**

A buffer to hold the returned name of the first archived log file no longer needed for recovery. This file, and all files up to and including lastarcdel, can be moved to make room on the disk.

For example, if the values returned in firstarcdel and lastarcdel are S0000001.LOG and S0000005.LOG, the following log files can be moved:

- S0000001.LOG
- S0000002.LOG
- S0000003.LOG
- S0000004.LOG
- S0000005.LOG

### **lastarcdel**

A buffer to hold the returned name of the last archived log file that can be removed from the database log directory.

### **lastcommit**

A string containing a time stamp in ISO format. This value represents the time stamp of the last committed transaction after the rollforward operation terminates.

### **Usage notes:**

The database manager uses the information stored in the archived and the active log files to reconstruct the transactions performed on the database since its last backup.

The action performed when this API is called depends on the *rollforward\_pending* flag of the database prior to the call. This can be queried using `db2CfgGet - Get Configuration Parameters`. The *rollforward\_pending* flag is set to `DATABASE` if the database is in roll-forward pending state. It is set to `TABLESPACE` if one or more table spaces are in `SQLB_ROLLFORWARD_PENDING` or `SQLB_ROLLFORWARD_IN_PROGRESS` state. The *rollforward\_pending* flag is set to `NO` if neither the database nor any of the table spaces needs to be rolled forward.

If the database is in roll-forward pending state when this API is called, the database will be rolled forward. Table spaces are returned to normal state after a successful database roll-forward, unless an abnormal state causes one or more table spaces to go offline. If the *rollforward\_pending* flag is set to `TABLESPACE`, only those table spaces that are in roll-forward pending state, or those table spaces requested by name, will be rolled forward.

**Note:** If table space rollforward terminates abnormally, table spaces that were being rolled forward will be put in `SQLB_ROLLFORWARD_IN_PROGRESS` state. In the next invocation of `ROLLFORWARD DATABASE`, only those table spaces in `SQLB_ROLLFORWARD_IN_PROGRESS` state will be processed. If the set of selected table space names does not include all table spaces that are in `SQLB_ROLLFORWARD_IN_PROGRESS` state, the table spaces that are not required will be put into `SQLB_RESTORE_PENDING` state.

If the database is not in roll-forward pending state and no point in time is specified, any table spaces that are in rollforward-in-progress state will be rolled forward to the end of logs. If no table spaces are in rollforward-in-progress state, any table spaces that are in rollforward pending state will be rolled forward to the end of logs.

This API reads the log files, beginning with the log file that is matched with the backup image. The name of this log file can be determined by calling this API with a caller action of `DB2ROLLFORWARD_QUERY` before rolling forward any log files.

The transactions contained in the log files are reapplied to the database. The log is processed as far forward in time as information is available, or until the time specified by the stop time parameter.

Recovery stops when any one of the following events occurs:

- No more log files are found
- A time stamp in the log file exceeds the completion time stamp specified by the stop time parameter

## db2Rollforward - Rollforward Database

- An error occurs while reading the log file.

Some transactions might not be recovered. The value returned in *lascommit* indicates the time stamp of the last committed transaction that was applied to the database.

If the need for database recovery was caused by application or human error, the user may want to provide a time stamp value in *piStopTime*, indicating that recovery should be stopped before the time of the error. This applies only to full database roll-forward recovery, and to table space rollforward to a point in time. It also permits recovery to be stopped before a log read error occurs, determined during an earlier failed attempt to recover.

When the *rollforward\_recovery* flag is set to DATABASE, the database is not available for use until roll-forward recovery is terminated. Termination is accomplished by calling the API with a caller action of DB2ROLLFORWARD\_STOP or DB2ROLLFORWARD\_RFWRD\_STOP to bring the database out of roll-forward pending state. If the *rollforward\_recovery* flag is TABLESPACE, the database is available for use. However, the table spaces in SQLB\_ROLLFORWARD\_PENDING and SQLB\_ROLLFORWARD\_IN\_PROGRESS states will not be available until the API is called to perform table space roll-forward recovery. If rolling forward table spaces to a point in time, the table spaces are placed in backup pending state after a successful rollforward.

When the *RollforwardFlags* option is set to DB2ROLLFORWARD\_LOCAL\_TIME, all messages returned to the user will also be in local time. All times are converted on the server, and on the catalog partition, if it is a partitioned database environment. The timestamp string is converted to GMT on the server, so the time is local to the server's time zone, not the client's. If the client is in one time zone and the server in another, the server's local time should be used. This is different from the local time option from the Control Center, which is local to the client. If the timestamp string is close to the time change of the clock due to daylight savings, it is important to know if the stop time is before or after the clock change, and specify it correctly.

### Related reference:

- "SQLCA" in the *Administrative API Reference*
- "db2Restore - Restore database" on page 96

### Related samples:

- "dbrecov.sqc -- How to recover a database (C)"
- "dbrecov.sqC -- How to recover a database (C++)"

---

## Rollforward sessions - CLP examples

### Example 1

The ROLLFORWARD DATABASE command permits specification of multiple operations at once, each being separated with the keyword AND. For example, to roll forward to the end of logs, and complete, the separate commands:

```
db2 rollforward db sample to end of logs
db2 rollforward db sample complete
```

can be combined as follows:

```
db2 rollforward db sample to end of logs and complete
```

Although the two are equivalent, it is recommended that such operations be done in two steps. It is important to verify that the rollforward operation has progressed as expected, before stopping it and possibly missing logs. This is especially important if a bad log is found during rollforward recovery, and the bad log is interpreted to mean the “end of logs”. In such cases, an undamaged backup copy of that log could be used to continue the rollforward operation through more logs.

### Example 2

Roll forward to the end of the logs (two table spaces have been restored):

```
db2 rollforward db sample to end of logs
db2 rollforward db sample to end of logs and stop
```

These two statements are equivalent. Neither AND STOP or AND COMPLETE is needed for table space rollforward recovery to the end of the logs. Table space names are not required. If not specified, all table spaces requiring rollforward recovery will be included. If only a subset of these table spaces is to be rolled forward, their names must be specified.

### Example 3

After three table spaces have been restored, roll one forward to the end of the logs, and the other two to a point in time, both to be done online:

```
db2 rollforward db sample to end of logs tablespace(TBS1) online
db2 rollforward db sample to 1998-04-03-14.21.56.245378 and stop
tablespace(TBS2, TBS3) online
```

Note that two rollforward operations cannot be run concurrently. The second command can only be invoked after the first rollforward operation completes successfully.

### Example 4

After restoring the database, roll forward to a point in time, using OVERFLOW LOG PATH to specify the directory where the user exit saves archived logs:

```
db2 rollforward db sample to 1998-04-03-14.21.56.245378 and stop
overflow log path (/logs)
```

### Example 5 (partitioned database environments)

There are three nodes: 0, 1, and 2. Table space TBS1 is defined on all nodes, and table space TBS2 is defined on nodes 0 and 2. After restoring the database on node 1, and TBS1 on nodes 0 and 2, roll the database forward on node 1:

```
db2 rollforward db sample to end of logs and stop
```

This returns warning SQL1271 (“Database is recovered but one or more table spaces are offline on node(s) 0 and 2.”).

```
db2 rollforward db sample to end of logs
```

This rolls TBS1 forward on nodes 0 and 2. The clause TABLESPACE(TBS1) is optional in this case.

### Example 6 (partitioned database environments)

After restoring table space TBS1 on nodes 0 and 2 only, roll TBS1 forward on nodes 0 and 2:

```
db2 rollforward db sample to end of logs
```

Node 1 is ignored.

```
db2 rollforward db sample to end of logs tablespace(TBS1)
```

This fails, because TBS1 is not ready for rollforward recovery on node 1. Reports SQL4906N.

```
db2 rollforward db sample to end of logs on nodes (0, 2) tablespace(TBS1)
```

This completes successfully.

```
db2 rollforward db sample to 1998-04-03-14.21.56.245378 and stop  
tablespace(TBS1)
```

This fails, because TBS1 is not ready for rollforward recovery on node 1; all pieces must be rolled forward together.

**Note:** With table space rollforward to a point in time, the node clause is not accepted. The rollforward operation must take place on all the nodes on which the table space resides.

After restoring TBS1 on node 1:

```
db2 rollforward db sample to 1998-04-03-14.21.56.245378 and stop  
tablespace(TBS1)
```

This completes successfully.

### Example 7 (partitioned database environments)

After restoring a table space on all nodes, roll forward to PIT2, but do not specify AND STOP. The rollforward operation is still in progress. Cancel and roll forward to PIT1:

```
db2 rollforward db sample to pit2 tablespace(TBS1)  
db2 rollforward db sample cancel tablespace(TBS1)
```

```
** restore TBS1 on all nodes **
```

```
db2 rollforward db sample to pit1 tablespace(TBS1)  
db2 rollforward db sample stop tablespace(TBS1)
```

### Example 8 (partitioned database environments)

Rollforward recover a table space that resides on eight nodes (3 to 10) listed in the db2nodes.cfg file:

```
db2 rollforward database dwtest to end of logs tablespace (tssprodt)
```

This operation to the end of logs (not point in time) completes successfully. The nodes on which the table space resides do not have to be specified. The utility defaults to the db2nodes.cfg file.

### Example 9 (partitioned database environments)

Rollforward recover six small table spaces that reside on a single node database partition group (on node 6):



```
db2 rollforward database dwtest to end of logs on node (6)
    tablespace(tsstore, tssbuyer, tsstime, tsswhse, tsslscat, tssvendor)
```

This operation to the end of logs (not point in time) completes successfully.



---

## Chapter 5. Database recover

This section describes the DB2 UDB recover utility, which performs the necessary restore and rollforward operations to recover a database to a specified time, based on information found in the recovery history file.

The following topics are covered:

- “Recover overview”
- “Privileges, authorities, and authorization required to use recover” on page 146
- “Using recover” on page 146
- “Client/server timestamp conversion” on page 146
- “RECOVER DATABASE” on page 147
- “db2Recover - Recover database” on page 152

---

### Recover overview

The recover utility performs the necessary restore and rollforward operations to recover a database to a specified time, based on information found in the recovery history file. When you use this utility, you specify that the database be recovered to a point-in-time or to the end of the log files. The utility will then select the best suitable backup image and perform the recovery operations.

The recover utility does not support the following RESTORE DATABASE command options:

- TABLESPACE tablespace-name. Table space restore operations are not supported.
- INCREMENTAL. Incremental restore operations are not supported.
- OPEN num-sessions SESSIONS. You cannot indicate the number of I/O sessions that are to be used with TSM or another vendor product.
- BUFFER buffer-size. You cannot set the size of the buffer used for the restore operation.
- DLREPORT filename. You cannot specify a file name for reporting files that become unlinked.
- WITHOUT ROLLING FORWARD. You cannot specify that the database is not to be placed in rollforward pending state after a successful restore operation.
- WITHOUT DATALINK. You cannot specify that no reconciliation of linked files is to be performed.
- PARALLELISM n. You cannot indicate the degree of parallelism for the restore operation.
- WITHOUT PROMPTING. You cannot specify that a restore operation is to run unattended

#### Related concepts:

- “Client/server timestamp conversion” on page 122

#### Related tasks:

- “Using recover” on page 146

#### Related reference:

- “RESTORE DATABASE” on page 87
- “Configuration parameters for database logging” on page 35

---

## Privileges, authorities, and authorization required to use recover

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges and higher-level database manager maintenance and utility operations. Together, these act to control access to the database manager and its database objects. Users can access only those objects for which they have the appropriate authorization; that is, the required privilege or authority.

You must have SYSADM, SYSCTRL, or SYSMAINT authority to use the recover utility.

---

## Using recover

### Prerequisites:

You should not be connected to the database that is to be recovered: the recover database utility automatically establishes a connection to the specified database, and this connection is terminated at the completion of the recover operation.

The database can be local or remote.

### Procedure:

The recover utility can be invoked through the command line processor (CLP) or the **db2Recover** application programming interface (API).

The following example shows how to use the RECOVER DATABASE command through the CLP:

```
db2 recover db sample
```

**Note:** In a partitioned database environment, the recover utility must be invoked from the catalog node of the database.

### Related concepts:

- “Recover overview” on page 145

### Related reference:

- “db2Recover - Recover database” on page 152
- “RECOVER DATABASE” on page 147

---

## Client/server timestamp conversion

This section explains the generation of timestamps in a client/server environment:

- If you specify a local time for a rollforward operation, all messages returned will also be in local time.

**Note:** All times are converted on the server and (in partitioned database environments) on the catalog node.

- The timestamp string is converted to GMT on the server, so the time represents the server's time zone, not the client's. If the client is in a different time zone from the server, the server's local time should be used.
- If the timestamp string is close to the time change due to daylight savings time, it is important to know whether the stop time is before or after the time change so that it is specified correctly.

**Related concepts:**

- "Rollforward overview" on page 111
- "Synchronizing clocks in a partitioned database system" on page 121

## RECOVER DATABASE

Restores and rolls forward a database to a particular point in time or to the end of the logs.

**Scope:**

In a partitioned database environment, this command can only be invoked from the catalog partition. A database recover operation to a specified point in time affects all partitions that are listed in the `db2nodes.cfg` file. A database recover operation to the end of logs affects the partitions that are specified. If no partitions are specified, it affects all partitions that are listed in the `db2nodes.cfg` file.

**Authorization:**

To recover an existing database, one of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

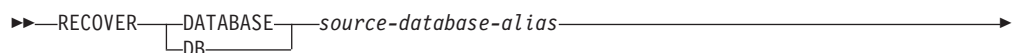
To recover to a new database, one of the following:

- *sysadm*
- *sysctrl*

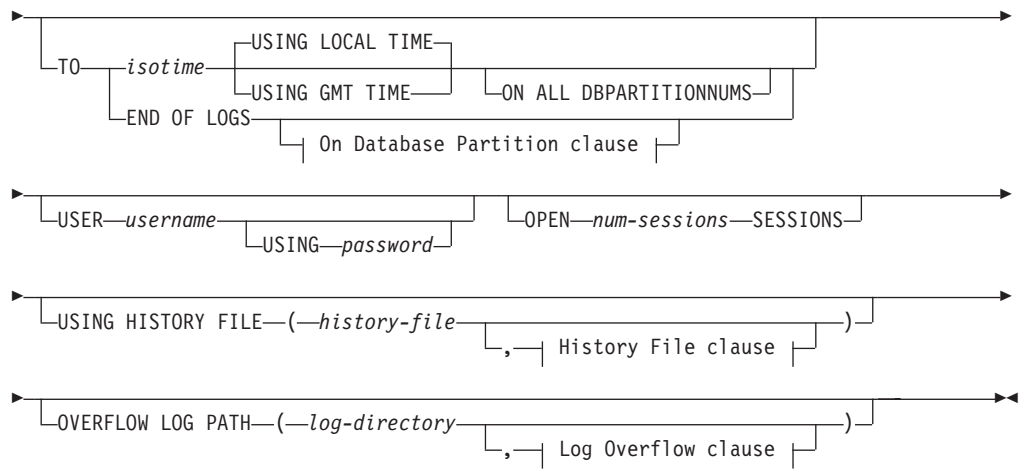
**Required connection:**

To recover an existing database, a database connection is required. This command automatically establishes a connection to the specified database and will release the connection when the recover operation finishes. To recover to a new database, an instance attachment and a database connection are required. The instance attachment is required to create the database.


**Command syntax:**




## RECOVER DATABASE




### On Database Partition clause:




### Database Partition List clause:



### Log Overflow clause:



### History File clause:



### Command parameters:

#### DATABASE *database-alias*

The alias of the database that is to be recovered.

#### USER *username*

The user name under which the database is to be recovered.

#### USING *password*

The password used to authenticate the user name. If the password is omitted, the user is prompted to enter it.

#### TO

*isotime* The point in time to which all committed transactions are to be recovered (including the transaction committed precisely at that time, as well as all transactions committed previously).

This value is specified as a time stamp, a 7-part character string that identifies a combined date and time. The format is *yyyy-mm-dd-hh.mm.ss.nnnnnn* (year, month, day, hour, minutes, seconds, microseconds), expressed in Coordinated Universal Time (UTC). UTC helps to avoid having the same time stamp associated with different logs (because of a change in time associated with daylight savings time, for example). The time stamp in a backup image is based on the local time at which the backup operation started. The CURRENT TIMEZONE special register specifies the difference between UTC and local time at the application server. The difference is represented by a time duration (a decimal number in which the first two digits represent the number of hours, the next two digits represent the number of minutes, and the last two digits represent the number of seconds). Subtracting CURRENT TIMEZONE from a local time converts that local time to UTC.

#### USING LOCAL TIME

Specifies the point in time to which to recover. This option allows the user to recover to a point in time that is the user's local time rather than GMT time. This makes it easier for users to recover to a specific point in time on their local machines, and eliminates potential user errors due to the translation of local to GMT time.

##### Notes:

1. If the user specifies a local time for recovery, all messages returned to the user will also be in local time. Note that all times are converted on the server, and in partitioned database environments, on the catalog database partition.
2. The timestamp string is converted to GMT on the server, so the time is local to the server's time zone, not the client's. If the client is in one time zone and the server in another, the server's local time should be used. This is different from the local time option from the Control Center, which is local to the client.
3. If the timestamp string is close to the time change of the clock due to daylight saving time, it is important to know if the stop time is before or after the clock change, and specify it correctly.

#### USING GMT TIME

Specifies the point in time to which to recover.

#### END OF LOGS

Specifies that all committed transactions from all online archive log files listed in the database configuration parameter *logpath* are to be applied.

#### ON ALL DBPARTITIONNUMS

Specifies that transactions are to be rolled forward on all partitions specified in the *db2nodes.cfg* file. This is the default if a database partition clause is not specified.

#### EXCEPT

Specifies that transactions are to be rolled forward on all partitions specified in the *db2nodes.cfg* file, except those specified in the database partition list.

#### ON DBPARTITIONNUM / ON DBPARTITIONNUMS

Roll the database forward on a set of database partitions.

## RECOVER DATABASE

*db-partition-number1*

Specifies a database partition number in the database partition list.

*db-partition-number2*

Specifies the second database partition number, so that all partitions from *db-partition-number1* up to and including *db-partition-number2* are included in the database partition list.

**OPEN** *num-sessions* **SESSIONS**

Specifies the number of I/O sessions that are to be used with Tivoli Storage Manager (TSM) or the vendor product.

**USING HISTORY FILE** *history-file*

*history-file* **ON DBPARTITIONNUM**

In a partitioned database environment, allows a different history file

**OVERFLOW LOG PATH** *log-directory*

Specifies an alternate log path to be searched for archived logs during recovery. Use this parameter if log files were moved to a location other than that specified by the *logpath* database configuration parameter. In a partitioned database environment, this is the (fully qualified) default overflow log path *for all partitions*. A relative overflow log path can be specified for single-partition databases.

**Note:** The OVERFLOW LOG PATH command parameter will overwrite the value (if any) of the database configuration parameter *overflowlogpath*.

*log-directory* **ON DBPARTITIONNUM**

In a partitioned database environment, allows a different log path to override the default overflow log path for a specific partition.

### Examples:

In a single-partition database environment, where the database being recovered currently exists, and so the most recent version of the history file is available in the *dftdbpath*:

1. To use the latest backup image and rollforward to the end of logs using all default values:  

```
RECOVER DB SAMPLE
```
2. To recover the database to a PIT, issue the following. The most recent image that can be used will be restored, and logs applied until the PIT is reached.  

```
RECOVER DB SAMPLE TO 2001-12-31-04:00:00
```
3. To recover the database using a saved version of the history file, issue the following. For example, if the user needs to recover to an extremely old PIT which is no longer contained in the current history file, the user will have to provide a version of the history file from this time period. If the user has saved a history file from this time period, this version can be used to drive the recover.

```
RECOVER DB SAMPLE TO 1999-12-31-04:00:00
USING HISTORY FILE (/home/user/old1999files/db2rhist.asc)
```

In a single-partition database environment, where the database being recovered does not exist, you must use the USING HISTORY FILE clause to point to a history file.

1. If you have not made any backups of the history file, so that the only version available is the copy in the backup image, the recommendation is to issue a



RESTORE followed by a ROLLFORWARD. However, to use RECOVER, you would first have to extract the history file from the image to some location, for example /home/user/oldfiles/db2rhist.asc, and then issue this command. (This version of the history file does not contain any information about log files that are required for rollforward, so this history file is not useful for RECOVER.)

```
RECOVER DB SAMPLE TO END OF LOGS
USING HISTORY FILE (/home/user/fromimage/db2rhist.asc)
```

2. If you have been making periodic or frequent backup copies of the history, the USING HISTORY clause should be used to point to this version of the history file. If the file is /home/user/myfiles/db2rhist.asc, issue the command:

```
RECOVER DB SAMPLE TO PIT
USING HISTORY FILE (/home/user/myfiles/db2rhist.asc)
```

(In this case, you can use any copy of the history file, not necessarily the latest, as long as it contains a backup taken before the point-in-time (PIT) requested.)

In a partitioned database environment, where the database exists on all database partitions, and the latest history file is available on *dftdbpath* on all database partitions:

1. To recover the database to a PIT on all nodes. DB2 will verify that the PIT is reachable on all nodes before starting any restore operations.

```
RECOVER DB SAMPLE TO 2001:12:31:04:00:00
```

2. To recover the database to this PIT on all nodes. DB2 will verify that the PIT is reachable on all nodes before starting any restore operations. The RECOVER operation on each node is identical to a single-partition RECOVER.

```
RECOVER DB SAMPLE TO END OF LOGS
```

3. Even though the most recent version of the history file is in the *dftdbpath*, you might want to use several specific history files. Unless otherwise specified, each partition will use the history file found locally at /home/user/oldfiles/db2rhist.asc. The exceptions are nodes 2 and 4. Node 2 will use: /home/user/node2files/db2rhist.asc, and node 4 will use: /home/user/node4files/db2rhist.asc.

```
RECOVER DB SAMPLE TO 1999:12:31:04:00:00
USING HISTORY FILE (/home/user/oldfiles/db2rhist.asc,
/home/user/node2files/db2rhist.asc ON DBPARTITIONNUM 2,
/home/user/node4files/db2rhist.asc ON DBPARTITIONNUM 4)
```

4. It is possible to recover a subset of nodes instead of all nodes, however a PIT RECOVER can not be done in this case, the recover must be done to EOL.

```
RECOVER DB SAMPLE TO END OF LOGS ON DBPARTITIONNUMS(2 TO 4, 7, 9)
```

In a partitioned database environment, where the database does not exist:

1. If you have not made any backups of the history file, so that the only version available is the copy in the backup image, the recommendation is to issue a RESTORE followed by a ROLLFORWARD. However, to use RECOVER, you would first have to extract the history file from the image to some location, for example, /home/user/oldfiles/db2rhist.asc, and then issue this command. (This version of the history file does not contain any information about log files that are required for rollforward, so this history file is not useful for the recover.)

```
RECOVER DB SAMPLE TO PIT
USING HISTORY FILE (/home/user/fromimage/db2rhist.asc)
```

## RECOVER DATABASE

2. If you have been making periodic or frequent backup copies of the history, the USING HISTORY clause should be used to point to this version of the history file. If the file is /home/user/myfiles/db2rhist.asc, you can issue the following command:

```
RECOVER DB SAMPLE TO END OF LOGS
USING HISTORY FILE (/home/user/myfiles/db2rhist.asc)
```

### Usage notes:

- Recovering a database might require a load recovery using tape devices. If prompted for another tape, the user can respond with one of the following:
  - c** Continue. Continue using the device that generated the warning message (for example, when a new tape has been mounted).
  - d** Device terminate. Stop using the device that generated the warning message (for example, when there are no more tapes).
  - t** Terminate. Terminate all devices.
- If there is a failure during the restore portion of the recover operation, you can reissue the RECOVER DATABASE command. If the restore operation was successful, but there was an error during the rollforward operation, you can issue a ROLLFORWARD DATABASE command, since it is not necessary (and it is time-consuming) to redo the entire recover operation.
- In a partitioned database environment, if there is an error during the restore portion of the recover operation, it is possible that it is only an error on a single database partition. Instead of reissuing the RECOVER DATABASE command, which restores the database on all database partitions, it is more efficient to issue a RESTORE DATABASE command for the database partition that failed, followed by a ROLLFORWARD DATABASE command.

---

## db2Recover - Recover database

Restores and rolls forward a database to a particular point in time or to the end of the logs.

### Scope:

In a partitioned database environment, this API can only be called from the catalog partition. If no database partition servers are specified, it affects all database partition servers that are listed in the db2nodes.cfg file. If a point in time is specified, the API affects all database partitions.

### Authorization:

To recover an existing database, one of the following:

- sysadm*
- sysctrl*
- sysmaint*

To recover to a new database, one of the following:

- sysadm*
- sysctrl*

### Required connection:

To recover an existing database, a database connection is required. This API automatically establishes a connection to the specified database and will release the connection when the recover operation finishes. Instance and database, to recover to a new database. The instance attachment is required to create the database.

#### API include file:

*db2ApiDf.h*

#### C API syntax:

```

/* File: db2ApiDf.h */
/* API: db2Recover */
/* ... */
SQL_API_RC SQL_API_FN
db2Recover (
    db2UInt32 versionNumber,
    void * pDB2RecovStruct,
    struct sqlca * pSqlca);

typedef SQL_STRUCTURE db2RecoverStruct
{
    char                *piSourceDBAlias;
    char                *piUsername;
    char                *piPassword;
    db2UInt32           iRecoverCallerAction;
    db2UInt32           iOptions;
    sqlint32            *poNumReplies;
    struct sqlurfl_info *poNodeInfo;
    char                *piStopTime;
    char                *piOverflowLogPath;
    db2UInt32           iNumChngLgOvrflw;
    struct sqlurfl_newlogpath *piChngLogOvrflw;
    db2int32            iAllNodeFlag;
    db2int32            iNumNodes;
    SQL_PDB_NODE_TYPE  *piNodeList;
    db2int32            iNumNodeInfo;
    db2UInt32           iRollforwardFlags;
    char                *piHistoryFile;
    db2UInt32           iNumChngHistoryFile;
    struct sqlu_histFile *piChngHistoryFile;
} db2RecoverStruct;
/* ... */

```

#### Generic API syntax:

```

/* File: db2ApiDf.h */
/* API: db2gRecover */
/* ... */
SQL_API_RC SQL_API_FN
db2gRecover (
    db2UInt32 versionNumber,
    void * pDB2gRecoverStruct,
    struct sqlca * pSqlca);

typedef SQL_STRUCTURE db2gRecoverStruct
{
    char                *piSourceDBAlias;
    db2UInt32           iSourceDBAliasLen;
    char                *piUserName;
    db2UInt32           iUserNameLen;
    char                *piPassword;
    db2UInt32           iPasswordLen;
    db2UInt32           iRecoverCallerAction;
    db2UInt32           iOptions;
    sqlint32            *poNumReplies;
}

```

## db2Recover - Recover database

```
|      struct sqlurf_info      *poNodeInfo;
|      char                   *piStopTime;
|      db2Uint32              iStopTimeLen;
|      char                   *piOverflowLogPath;
|      db2Uint32              iOverflowLogPathLen;
|      db2Uint32              iNumChngLgOvrflw;
|      struct sqlurf_newlogpath *piChngLogOvrflw;
|      db2int32               iAllNodeFlag;
|      db2int32               iNumNodes;
|      SQL_PDB_NODE_TYPE      *piNodeList;
|      db2int32               iNumNodeInfo;
|      db2Uint32              iRollforwardFlags;
|      char                   *piHistoryFile;
|      db2Uint32              iHistoryFileLen;
|      db2Uint32              iNumChngHistoryFile;
|      struct sqlu_histFile   *piChngHistoryFile;
|  } db2gRecoverStruct;
|  /* ... */
```

### API parameters:

#### versionNumber

Input. Specifies the version and release level of the structure passed as the second parameter *pDB2RecoverStruct*.

#### pDB2RecoverStruct

Input. A pointer to the *db2RecoverStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### piSourceDBAlias

Input. A string containing the database alias of the database to be recovered.

#### iSourceDBAliasLen

Length in bytes of *piSourceDBAlias*.

#### piUserName

Input. A string containing the user name to be used when attempting a connection. Can be NULL.

#### iUserNameLen

Length in bytes of *piUsername*.

#### piPassword

Input. A string containing the password to be used with the user name. Can be NULL.

#### iPasswordLen

Length in bytes of *piPassword*.

#### iRecoverCallerAction

Input. Valid values are:

##### DB2RESTORE\_NOINTERRUPT

Starts the restore operation. Specifies that the restore will run unattended, and that scenarios that normally require user intervention will either be attempted without first returning to the caller, or will generate an error. Use this caller action, for example, if it is known that all of the media required for the restore have been mounted, and utility prompts are not desired.

**DB2RESTORE\_CONTINUE**

Continues the restore operation after the user has performed some action requested by the utility (mounting a new tape, for example).

**DB2RESTORE\_TERMINATE**

Terminates the restore operation after the user has failed to perform some action requested by the utility.

**DB2RESTORE\_DEVICE\_TERMINATE**

Removes a particular device from the list of devices used by the restore operation. When a particular device has exhausted its input, restore will return a warning to the caller. Call the restore utility again with this caller action to remove the device that generated the warning from the list of devices being used.

**DB2RESTORE\_PARM\_CHK**

Used to validate parameters without performing a restore operation. This option does not terminate the database connection after the call returns. After successful return of this call, it is expected that the user will issue a call with **DB2RESTORE\_CONTINUE** to proceed with the action.

**DB2RESTORE\_PARM\_CHK\_ONLY**

Used to validate parameters without performing a restore operation. Before this call returns, the database connection established by this call is terminated, and no subsequent call is required.

**DB2RESTORE\_TERMINATE\_INCRE**

Terminates an incremental restore operation before completion.

**DB2ROLLFORWARD\_LOADREC\_CONT**

Continue using the device that generated the warning message (for example, when a new tape has been mounted).

**DB2ROLLFORWARD\_DEVICE\_TERM**

Stop using the device that generated the warning message (for example, when there are no more tapes).

**DB2ROLLFORWARD\_LOAD\_REC\_TERM**

Terminate all devices being used by load recovery.

**iOptions**

Input. Valid values are:

**DB2RECOVER\_EMPTY\_FLAG**

No flags specified.

**DB2RECOVER\_LOCAL\_TIME**

Indicates that the value specified for the stop time by *piStopTime* is in local time, not GMT. This is the default setting.

**DB2RECOVER\_GMT\_TIME**

This flag indicates that the value specified for the stop time by *piStopTime* is in GMT (Greenwich Mean Time).

**poNumReplies**

Output. The number of replies received.

**poNodeInfo**

Output. Database partition reply information.

## db2Recover - Recover database

### **piStopTime**

Input. A character string containing a time stamp in ISO format. Database recovery will stop when this time stamp is exceeded. Specify SQLUM\_INFINITY\_TIMESTAMP to roll forward as far as possible. May be NULL for DB2ROLLFORWARD\_QUERY, DB2ROLLFORWARD\_PARM\_CHECK, and any of the load recovery (DB2ROLLFORWARD\_LOADREC\_) caller actions.

### **iStopTimeLen**

Length in bytes of *piStopTime*.

### **piOverflowLogPath**

Input. This parameter is used to specify an alternate log path to be used. In addition to the active log files, archived log files need to be moved (by the user) into the location specified by the *logpath* configuration parameter before they can be used by this utility. This can be a problem if the user does not have sufficient space in the log path. The overflow log path is provided for this reason. During roll-forward recovery, the required log files are searched, first in the log path, and then in the overflow log path. The log files needed for table space rollforward recovery can be brought into either the log path or the overflow log path. If the caller does not specify an overflow log path, the default value is the log path. In a partitioned database environment, the overflow log path must be a valid, fully qualified path; the default path is the default overflow log path for each database partition. In a single-partition database environment, the overflow log path can be relative if the server is local.

### **iOverflowLogPathLen**

Length in bytes of *piOverflowLogPath*.

### **iNumChngLgOvrflw**

Input. Partitioned database environments only. The number of changed overflow log paths. These new log paths override the default overflow log path for the specified database partition server only.

### **piChngLogOvrflw**

Input. Partitioned database environments only. A pointer to a structure containing the fully qualified names of changed overflow log paths. These new log paths override the default overflow log path for the specified database partition server only.

### **iAllNodeFlag**

Input. Partitioned database environments only. Indicates whether the rollforward operation is to be applied to all database partition servers defined in *db2nodes.cfg*. Valid values are:

#### **DB2\_NODE\_LIST**

Apply to database partition servers in a list that is passed in *piNodeList*.

#### **DB2\_ALL\_NODES**

Apply to all database partition servers. *piNodeList* should be NULL. This is the default value.

#### **DB2\_ALL\_EXCEPT**

Apply to all database partition servers except those in a list that is passed in *piNodeList*.

#### **DB2\_CAT\_NODE\_ONLY**

Apply to the catalog partition only. *piNodeList* should be NULL.

**iNumNodes**

Input. Specifies the number of database partition servers in the *piNodeList* array.

**piNodeList**

Input. A pointer to an array of database partition server numbers on which to perform the rollforward recovery.

**iNumNodeInfo**

Input. Defines the size of the output parameter *poNodeInfo*, which must be large enough to hold status information from each database partition that is being rolled forward. In a single-partition database environment, this parameter should be set to 1. The value of this parameter should be the same as the number of database partition servers for which this API is being called.

**RollforwardFlags**

Input. Specifies the rollforward flags. Valid values (defined in *db2ApiDf.h*):

**DB2ROLLFORWARD\_EMPTY\_FLAG**

No flags specified.

**DB2ROLLFORWARD\_LOCAL\_TIME**

Allows the user to roll forward to a point in time that is the user's local time rather than GMT time. This makes it easier for users to roll forward to a specific point in time on their local machines, and eliminates potential user errors due to the translation of local to GMT time.

**piHistoryFile**

History file.

**iHistoryFileLen**

Length in bytes of *piHistoryFile*.

**iNumChngHistoryFile**

Number of history files in list.

**piChngHistoryFile**

List of history files.

**Usage notes:****Related reference:**

- "RECOVER DATABASE" on page 147





---

## **Part 2. High availability**



---

## Chapter 6. Introducing high availability and failover support

Successful e-businesses depend on the uninterrupted availability of transaction processing systems, which in turn are driven by database management systems, such as DB2, that must be available 24 hours a day and 7 days a week (“24 x 7”). This section discusses the following:

- “High availability”
- “High availability through online split mirror and suspended I/O support” on page 164
- “Fault monitor facility for UNIX-based systems” on page 169
- “db2fm - DB2 Fault Monitor” on page 171

---

### High availability

*High availability* (HA) is the term that is used to describe systems that run and are available to customers more or less all the time. For this to occur:

- Transactions must be processed efficiently, without appreciable performance degradations (or even loss of availability) during peak operating periods. In a partitioned database environment, DB2<sup>®</sup> can take advantage of both *intrapartition* and *interpartition* parallelism to process transactions efficiently. *Intrapartition parallelism* can be used in an SMP environment to process the various components of a complex SQL statement simultaneously. *Interpartition parallelism* in a partitioned database environment, on the other hand, refers to the simultaneous processing of a query on all participating nodes; each node processes a subset of the rows in the table.
- Systems must be able to recover quickly when hardware or software failures occur, or when disaster strikes. DB2 has an advanced continuous checkpointing system and a parallel recovery capability that allow for extremely fast crash recovery.  
The ability to recover quickly can also depend on having a proven backup and recovery strategy in place.
- Software that powers the enterprise databases must be continuously running and available for transaction processing. To keep the database manager running, you must ensure that another database manager can take over if it fails. This is called *failover*. *Failover* capability allows for the automatic transfer of workload from one system to another when there is hardware failure.

You can implement failover capability through the high availability disaster recovery (HADR) database replication feature. HADR is a high availability solution that protects against data loss by replicating changes from a source database, called the primary database, to a target database, called the standby database.

Failover protection can also be achieved by keeping a copy of your database on another machine that is perpetually rolling the log files forward. *Log shipping* is the process of copying whole log files to a standby machine, either from an archive device, or through a user exit program running against the primary database. With this approach, the primary database is restored to the standby machine, using either the DB2 restore utility or the split mirror function. You can also use suspended I/O support to quickly initialize the new database. The standby database on the standby machine continuously rolls the log files forward. If the

primary database fails, any remaining log files are copied over to the standby machine. After a rollforward to the end of the logs and stop operation, all clients are reconnected to the standby database on the standby machine.

Failover support can also be provided through platform-specific software that you can add to your system. For example:

- Tivoli<sup>®</sup> System Automation for Linux.  
For detailed information about Tivoli System Automation, see the whitepaper entitled “Highly Available DB2 Universal Database™ using Tivoli System Automation for Linux”, which is available from the “DB2 UDB and DB2 Connect™ Online Support” web site (<http://www.ibm.com/software/data/pubs/papers/>).
- High Availability Cluster Multi-Processing, Enhanced Scalability, for AIX<sup>®</sup>.  
For detailed information about HACMP/ES, see the white paper entitled “IBM<sup>®</sup> DB2 Universal Database Enterprise Edition for AIX and HACMP/ES”, which is available from the “DB2 UDB and DB2 Connect Online Support” web site (<http://www.ibm.com/software/data/pubs/papers/>).
- Microsoft<sup>®</sup> Cluster Server, for Windows<sup>®</sup> operating systems.  
For information about Microsoft Cluster Server see the following white paper which is available from the “DB2 UDB and DB2 Connect Online Support” web site (<http://www.ibm.com/software/data/pubs/papers/>): “Implementing IBM DB2 Universal Database V8.1 Enterprise Server Edition with Microsoft Cluster Server”.
- Sun<sup>™</sup> Cluster, or VERITAS Cluster Server, for the Solaris<sup>™</sup> Operating Environment.  
For information about Sun Cluster, see the white paper entitled “DB2 Universal Database and High Availability on Sun Cluster 3.X”, which is available from the “DB2 UDB and DB2 Connect Online Support” web site (<http://www.ibm.com/software/data/pubs/papers/>). For information about VERITAS Cluster Server, see the white paper entitled “DB2 UDB and High Availability with VERITAS Cluster Server”, which is available from the “IBM Support and downloads” Web site (<http://www.ibm.com/support/docview.wss?uid=swg21045033>).
- Multi-Computer/ServiceGuard, for Hewlett-Packard.  
For detailed information about HP MC/ServiceGuard, see the white paper which discusses IBM DB2 ESE V8.1 with HP MC/ServiceGuard High Availability Software, which is available from the “IBM DB2 Information Management Products for HP” web site (<http://www-306.ibm.com/software/data/hp/>).

Failover strategies are usually based on clusters of systems. A *cluster* is a group of connected systems that work together as a single system. Each physical machine within a cluster contains one or more logical nodes. Clustering allows servers to back each other up when failures occur, by picking up the workload of the failed server.

IP address takeover (or IP takeover) is the ability to transfer a server IP address from one machine to another when a server goes down; to a client application, the two machines appear at different times to be the same server.

Failover software may use *heartbeat monitoring* or *keepalive packets* between systems to confirm availability. Heartbeat monitoring involves system services that

maintain constant communication between all the nodes in a cluster. If a heartbeat is not detected, failover to a backup system starts. End users are usually not aware that a system has failed.

The two most common failover strategies on the market are known as *idle standby* and *mutual takeover*, although the configurations associated with these terms may also be associated with different terms that depend on the vendor:

#### **Idle Standby**

In this configuration, one system is used to run a DB2 instance, and the second system is “idle”, or in standby mode, ready to take over the instance if there is an operating system or hardware failure involving the first system. Overall system performance is not impacted, because the standby system is idle until needed.

#### **Mutual Takeover**

In this configuration, each system is the designated backup for another system. Overall system performance may be impacted, because the backup system must do extra work following a failover: it must do its own work plus the work that was being done by the failed system.

Failover strategies can be used to failover an instance, a partition, or multiple logical nodes.

#### **Related concepts:**

- “Parallelism” in the *Administration Guide: Planning*
- “Developing a backup and recovery strategy” on page 3
- “High availability through online split mirror and suspended I/O support” on page 164
- “Cluster support for the Solaris Operating Environment” on page 219
- “Sun Cluster 3.0 support” on page 221
- “VERITAS Cluster Server support” on page 224
- “Microsoft Cluster Server support” on page 215
- “High Availability Cluster Multi-Processing support” on page 209
- “High availability disaster recovery overview” on page 173

---

## **High availability through log shipping**

Log shipping is the process of copying whole log files to a standby machine either from an archive device, or through a user exit program running against the primary database. The standby database is continuously rolling forward through the log files produced by the production machine. When the production machine fails, a failover occurs and the following takes place:

- The remaining logs are transferred over to the standby machine.
- The standby database rolls forward to the end of the logs and stops.
- The clients reconnect to the standby database and resume operations.

The standby machine has its own resources (i.e., disks), but must have the same physical and logical definitions as the production database. When using this approach the primary database is restored to the standby machine, by using restore utility or the split mirror function.

To ensure that you are able to recover your database in a disaster recovery situation consider the following:

- The archive location should be geographically separate from the primary site.
- Remotely mirror the log at the standby database site
- Use a synchronous mirror for no loss support. You can do this using modern disk subsystems such as ESS and EMC, or another remote mirroring technology. NVRAM cache (both local and remote) is also recommended to minimize the performance impact of a disaster recovery situation.

**Notes:**

1. When the standby database processes a log record indicating that an index rebuild took place on the primary database, the indexes on the standby server are not automatically rebuilt. The index will be rebuilt on the standby server either at the first connection to the database, or at the first attempt to access the index after the standby server is taken out of rollforward pending state. It is recommended that the standby server be resynchronized with the primary server if any indexes on the primary server are rebuilt.
2. If the load utility is run on the primary database with the COPY YES option specified, the standby database must have access to the copy image.
3. If the load utility is run on the primary database with the COPY NO option specified, the standby database should be resynchronized, otherwise the table space will be placed in restore pending state.
4. There are two ways to initialize a standby machine:
  - a. By restoring to it from a backup image.
  - b. By creating a split mirror of the production system and issuing the **db2inidb** command with the STANDBY option.

Only after the standby machine has been initialized can you issue the ROLLFORWARD command on the standby system.

5. Operations that are not logged will not be replayed on the standby database. As a result, it is recommended that you re-sync the standby database after such operations. You can do this through online split mirror and suspended I/O support.

**Related concepts:**

- “High availability through online split mirror and suspended I/O support” on page 164
- “High availability disaster recovery overview” on page 173

**Related tasks:**

- “Using a split mirror as a standby database” on page 167

**Related reference:**

- Appendix G, “User exit for database recovery,” on page 329

---

## High availability through online split mirror and suspended I/O support

*Suspended I/O* supports continuous system availability by providing a full implementation for online split mirror handling; that is, splitting a mirror without shutting down the database. A *split mirror* is an “instantaneous” copy of the database that can be made by mirroring the disks containing the data, and splitting the mirror when a copy is required. *Disk mirroring* is the process of writing all of your data to two separate hard disks; one is the mirror of the other. *Splitting* a mirror is the process of separating the primary and secondary copies of the database.

If you would rather not back up a large database using the DB2® backup utility, you can make copies from a mirrored image by using suspended I/O and the split mirror function. This approach also:

- Eliminates backup operation overhead from the production machine
- Represents a fast way to clone systems
- Represents a fast implementation of idle standby failover. There is no initial restore operation, and if a rollforward operation proves to be too slow, or encounters errors, reinitialization is very fast.

The **db2inidb** command initializes the split mirror so that it can be used:

- As a clone database
- As a standby database
- As a backup image

This command can only be issued against a split mirror, and it must be run before the split mirror can be used.

In a partitioned database environment, you do not have to suspend I/O writes on all partitions simultaneously. You can suspend a subset of one or more partitions to create split mirrors for performing offline backups. If the catalog node is included in the subset, it must be the last partition to be suspended.

In a partitioned database environment, the **db2inidb** command must be run on every partition before the split image from any of the partitions can be used. The tool can be run on all partitions simultaneously using the **db2\_all** command. If; however, you are using the RELOCATE USING option, you cannot use the **db2\_all** command to run **db2inidb** on all of the partitions simultaneously. A separate configuration file must be supplied for each partition, that includes the NODENUM value of the partition being changed. For example, if the name of a database is being changed, every partition will be affected and the **db2relocatedb** command must be run with a separate configuration file on each partition. If containers belonging to a single database partition are being moved, the **db2relocatedb** command only needs to be run once on that partition.

**Note:** Ensure that the split mirror contains all containers and directories which comprise the database, including the volume directory.

**Related reference:**

- “db2inidb - Initialize a Mirrored Database” on page 250

---

## Online split mirror handling

### Using a split mirror to clone a database

Use the following procedure to create a clone database. Although you can write to clone databases, they are generally used for read-only activities such as running reports.

**Restrictions:**

You cannot back up a cloned database, restore the backup image on the original system, or roll forward through log files produced on the original system.

**Procedure:**

To clone a database, follow these steps:

1. Suspend I/O on the primary database:

```
db2 set write suspend for database
```

2. Use appropriate operating system-level commands to split the mirror or mirrors from the primary database.

**Note:** Ensure that you copy the entire database directory including the volume directory. You must also copy the log directory and any container directories that exist outside the database directory.

3. Resume I/O on the primary database:

```
db2 set write resume for database
```

4. Catalog the mirrored database on the secondary system.

**Note:** By default, a mirrored database cannot exist on the same system as the primary database. It must be located on a secondary system that has the same directory structure and uses the same instance name as the primary database. If the mirrored database must exist on the same system as the primary database, you can use the **db2relocatedb** utility or the RELOCATE USING option of the **db2inidb** command to accomplish this.

5. Start the database instance on the secondary system:

```
db2start
```

6. Initialize the mirrored database on the secondary system:

```
db2inidb database_alias as snapshot
```

If required, specify the RELOCATE USING option of the db2inidb command to relocate the clone database:

```
db2inidb database_alias as snapshot relocate using relocatedbcfg.txt
```

where the relocatedbcfg.txt file contains the information required to relocate the database.

**Notes:**

- a. This command will roll back transactions that are in flight when the split occurs, and start a new log chain sequence so that any logs from the primary database cannot be replayed on the cloned database.
- b. The database directory (including the volume directory), the log directory, and the container directories must be moved to the desired location before you use the RELOCATE USING option.

**Related concepts:**

- “High availability through online split mirror and suspended I/O support” on page 164

**Related tasks:**

- “Using a split mirror as a standby database” on page 167
- “Using a split mirror as a backup image” on page 168

**Related reference:**

- “db2relocatedb - Relocate Database Command” in the *Command Reference*
- “SET WRITE Command” in the *Command Reference*



## Using a split mirror as a standby database

Use the following procedure to create a split mirror of a database for use as a standby database. If a failure occurs on the primary database and crash recovery is necessary, you can use the standby database to take over for the primary database.

### Procedure:

To use a split mirror as a standby database, follow these steps:

1. Suspend I/O on the primary database:  
`db2 set write suspend for database`
2. Use appropriate operating system-level commands to split the mirror or mirrors from the primary database.

**Note:** Ensure that you copy the entire database directory including the volume directory. You must also copy the log directory and any container directories that exist outside the database directory.

3. Resume I/O on the primary database:  
`db2 set write resume for database`
4. Catalog the mirrored database on the secondary system.

**Note:** By default, a mirrored database cannot exist on the same system as the primary database. It must be located on a secondary system that has the same directory structure and uses the same instance name as the primary database. If the mirrored database must exist on the same system as the primary database, you can use the **db2relocatedb** utility or the RELOCATE USING option of the **db2inidb** command to accomplish this.

5. Start the database instance on the secondary system:  
`db2start`
6. Initialize the mirrored database on the secondary system by placing it in rollforward pending state:  
`db2inidb database_alias as standby`

If required, specify the RELOCATE USING option of the db2inidb command to relocate the standby database:

```
db2inidb database_alias as standby relocate using relocatedbcfg.txt
```

where the relocatedbcfg.txt file contains the information required to relocate the database.

### Notes:

- a. If you have only DMS table spaces (database managed space), you can take a full database backup to offload the overhead of taking a backup on the production database.
  - b. The database directory (including the volume directory), the log directory, and the container directories must be moved to the desired location before you use the RELOCATE USING option.
7. Set up a user exit program to retrieve the log files from the primary system.
  8. Roll the database forward to the end of the logs or to a point-in-time.
  9. Continue retrieving log files, and rolling the database forward through the logs until you reach the end of the logs or the point-in-time required for the standby database.

10. To bring the standby database online issue the ROLLFORWARD command with the STOP option specified.

**Note:** The logs from the primary database cannot be applied to the mirrored database once it has been taken out of rollforward pending state.

**Related concepts:**

- “High availability through online split mirror and suspended I/O support” on page 164

**Related tasks:**

- “Using a split mirror to clone a database” on page 165
- “Using a split mirror as a backup image” on page 168

**Related reference:**

- “db2inidb - Initialize a Mirrored Database” on page 250
- “db2relocatedb - Relocate Database Command” in the *Command Reference*
- “SET WRITE Command” in the *Command Reference*

## Using a split mirror as a backup image

Use the following procedure to create a split mirror of a primary database for use as a backup image. This procedure can be used instead of performing backup database operations on the primary database.

**Procedure:**

To use a split mirror as a “backup image”, follow these steps:

1. Suspend I/O on the primary database:  
`db2 set write suspend for database`
2. Use appropriate operating system-level commands to split the mirror or mirrors from the primary database.
3. Resume I/O on the primary database:  
`db2 set write resume for database`
4. A failure occurs on the primary system, necessitating a restore from backup.
5. Stop the primary database instance:  
`db2stop`
6. Use operating system-level commands to copy the split-off data over the primary system. **Do not copy the split-off log files**, because the primary logs will be needed for rollforward recovery.
7. Start the primary database instance:  
`db2start`
8. Initialize the primary database:  
`db2inidb database_alias as mirror`
9. Roll the primary database forward to the end of the logs or to a point-in-time and stop.

**Related concepts:**

- “High availability through online split mirror and suspended I/O support” on page 164

**Related tasks:**

- “Using a split mirror to clone a database” on page 165
- “Using a split mirror as a standby database” on page 167

**Related reference:**

- “db2inidb - Initialize a Mirrored Database” on page 250

---

## Fault monitor facility for UNIX-based systems

On UNIX<sup>®</sup> based systems, the Fault Monitor Facility improves the availability of non-clustered DB2<sup>®</sup> environments through a sequence of processes that work together to ensure that DB2 is running. That is, the *init* daemen monitors the Fault Monitor Coordinator (FMC), the FMC monitors the fault monitors and the fault monitors monitor DB2.

The Fault Monitor Coordinator (FMC) is the process of the Fault Monitor Facility that is started at the UNIX boot sequence. The *init* daemon starts the FMC and will restart it if it terminates abnormally. The FMC starts one fault monitor for each DB2 instance. Each fault monitor runs as a daemon process and has the same user privileges as the DB2 instance. Once a fault monitor is started, it will be monitored to make sure it does not exit prematurely. If a fault monitor fails, it will be restarted by the FMC. Each fault monitor will, in turn, be responsible for monitoring one DB2 instance. If the DB2 instance exits prematurely, the fault monitor will restart it.

**Notes:**

1. If you are using a high availability clustering product (i.e., HACMP or MSCS), the fault monitor facility must be turned off since the instance startup and shut down is controlled by the clustering product.
2. The fault monitor will only become inactive if the **db2stop** command is issued. If a DB2 instance is shut down in any other way, the fault monitor will start it up again.

**Fault Monitor Registry File**

A fault monitor registry file is created for every instance on each physical machine when the fault monitor daemon is started. The values in this file specify the behavior of the fault monitors. The file can be found in the `/sqllib/` directory and is called `fm.<machine_name>.reg`. This file can be altered using the **db2fm** command. The entries are as follows:

```
FM_ON = no
FM_ACTIVE = yes
START_TIMEOUT = 600
STOP_TIMEOUT = 600
STATUS_TIMEOUT = 20
STATUS_INTERVAL = 20
RESTART_RETRIES = 3
ACTION_RETRIES = 3
NOTIFY_ADDRESS = <instance_name>@<machine_name>
```

where:

**FM\_ON**

Specifies whether or not the fault monitor should be started. If the value is set to NO, the fault monitor daemon will not be started, or will be turned off if it had already been started. The default value is NO.

### **FM\_ACTIVE**

Specifies whether or not the fault monitor is active. The fault monitor will only take action if both FM\_ON and FM\_ACTIVE are set to YES. If FM\_ON is set to YES and FM\_ACTIVE is set to NO, the fault monitor daemon will be started, but it will not be active. That means that it will not try to bring DB2 back online if it shuts down. The default value is YES.

### **START\_TIMEOUT**

Specifies the amount of time within which the fault monitor must start the service it is monitoring. The default value is 600 seconds.

### **STOP\_TIMEOUT**

Specifies the amount of time within which the fault monitor must bring down the service it is monitoring. The default value is 600 seconds.

### **STATUS\_TIMEOUT**

Specifies the amount of time within which the fault monitor must get the status of the service it is monitoring. The default value is 20 seconds.

### **STATUS\_INTERVAL**

Specifies the minimum time between two consecutive calls to obtain the status of the service that is being monitored. The default value is 20 seconds.

### **RESTART\_RETRIES**

Specifies the number of times the fault monitor will try to obtain the status of the service being monitored after a failed attempt. Once this number is reached the fault monitor will take action to bring the service back online. The default value is 3.

### **ACTION\_RETRIES**

Specifies the number of times the fault monitor will attempt to bring the service back online. The default value is 3.

### **NOTIFY\_ADDRESS**

Specifies the e-mail address to which the fault monitor will send notification messages. The default is <instance\_name>@<machine\_name>

This file can be altered using the **db2fm** command. For example:

To update the START\_TIMEOUT value to 100 seconds for instance DB2INST1, issue:

```
db2fm -i db2inst1 -T 100
```

To update the STOP\_TIMEOUT value to 200 seconds for instance DB2INST1, issue:

```
db2fm -i db2inst1 -T /200
```

To update the START\_TIMEOUT value to 100 seconds and the STOP\_TIMEOUT value to 200 seconds for instance DB2INST1, issue:

```
db2fm -i db2inst1 -T 100/200
```

To turn on fault monitoring for instance DB2INST1, issue:

```
db2fm -i db2inst1 -f yes
```

To turn off fault monitoring for instance DB2INST1, issue:

```
db2fm -i db2inst1 -f no
```

**Note:** If the fault monitor registry file does not exist, the default values will be used.

**Related reference:**

- “db2fm - DB2 Fault Monitor” on page 171

---

## db2fm - DB2 Fault Monitor

Controls the DB2 fault monitor daemon. You can use **db2fm** to configure the fault monitor.

This command is only available on UNIX-based platforms.

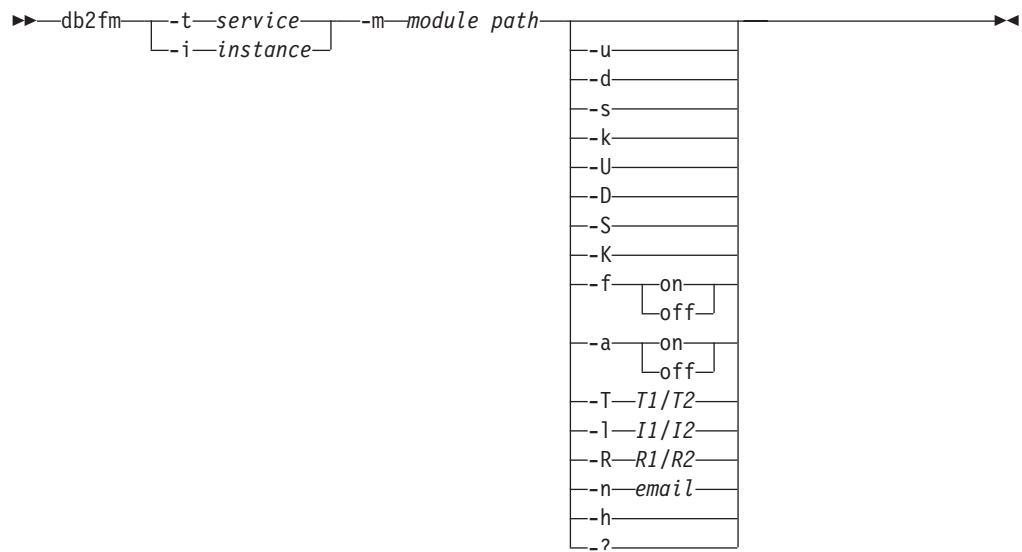
**Authorization:**

Authorization over the instance against which you are running the command.

**Required connection:**

None.

**Command syntax:**



**Command parameters:**

**-m** *module-path*

Defines the full path of the fault monitor shared library for the product being monitored. The default is \$INSTANCEHOME/sql/lib/libdb2gcf.

**-t** *service*

Gives the unique text descriptor for a service.

**-i** *instance*

Defines the instance of the service.

**-u** Brings the service up.

**-U** Brings the fault monitor daemon up.

**-d** Brings the service down.

**-D** Brings the fault monitor daemon down.

## db2fm - DB2 Fault Monitor

- k Kills the service.
- K Kills the fault monitor daemon.
- s Returns the status of the service.
- S Returns the status of the fault monitor daemon.

**Note:** the status of the service or fault monitor can be one of the following

- Not properly installed,
- INSTALLED PROPERLY but NOT ALIVE,
- ALIVE but NOT AVAILABLE (maintenance),
- AVAILABLE, or
- UNKNOWN

- f *on | off*  
Turns fault monitor on or off.

**Note:** If this option is set off, the fault monitor daemon will not be started, or the daemon will exit if it was running.

- a *on | off*  
Activates or deactivate fault monitoring.

**Note:** If this option if set off, the fault monitor will not be actively monitoring, which means if the service goes down it will not try to bring it back.

- T *T1/T2*  
Overwrites the start and stop time-out.

For example:

- -T 15/10 updates the two time-outs respectively
- -T 15 updates the start time-out to 15 secs
- -T /10 updates the stop time-out to 10 secs

- I *I1/I2*  
Sets the status interval and time-out respectively.

- R *R1/R2*  
Sets the number of retries for the status method and action before giving up.

- n *email*  
Sets the email address for notification of events.

- h Prints usage.

- ? Prints usage.

---

## Chapter 7. High availability disaster recovery (HADR)

---

### High availability disaster recovery overview

DB2<sup>®</sup> Universal Databases (DB2 UDB) high availability disaster recovery (HADR) is a database replication feature that provides a high availability solution for both partial and complete site failures. HADR protects against data loss by replicating data changes from a source database, called the primary, to a target database, called the standby. A database that does not use HADR is referred to as a standard database.

HADR might be your best option if most or all of your database requires protection, or if you perform DDL operations that must be automatically replicated on the standby database.

Applications can only access the current primary database. Updates to the standby database occur by rolling forward log data that is generated on the primary database and shipped to the standby database.

A partial site failure can be caused by a hardware, network, or software (DB2 or operating system) failure. Without HADR, a partial site failure requires the database management system (DBMS) server or the machine where the database resides to be rebooted. The length of time it takes to restart the database and the machine where it resides is unpredictable. It can take several minutes before the database is brought back to a consistent state and made available. With HADR, the standby database can take over in seconds. Further, you can redirect the clients that were using the original primary database to the standby database (new primary database) by using automatic client reroute or retry logic in the application.

A complete site failure can occur when a disaster, such as a fire, causes the entire site to be destroyed. Because HADR uses TCP/IP for communication between the primary and standby databases, they can be situated in different locations. For example, your primary database might be located at your head office in one city, while your standby database is located at your sales office in another city. If a disaster occurs at the primary site, data availability is maintained by having the remote standby database take over as the primary database with full DB2 functionality. After a takeover operation occurs, you can bring the original primary database back up and return it to its primary database status; this is known as failback.

With HADR, you can choose the level of protection you want from potential loss of data by specifying one of three synchronization modes: synchronous, near synchronous, or asynchronous.

After the failed original primary server is repaired, it can rejoin the HADR pair as a standby database if the two copies of the database can be made consistent. After the original primary database is reintegrated into the HADR pair as the standby database, you can switch the roles of the databases to enable the original primary database to once again be the primary database.

HADR is only one of several replication solutions offered in the DB2 product family. Version 8.2 of DB2 Information Integrator and DB2 UDB include SQL replication and Q replication solutions that can also be used, in some configurations, to provide high availability. These functions maintain logically consistent copies of database tables at multiple locations. In addition, they provide flexibility and complex functionality such as support for column and row filtering, data transformation, updates to any copy of a table, and they can be used in partitioned database environments.

**Related concepts:**

- “High availability through log shipping” on page 163
- “Automatic client reroute and high availability disaster recovery” on page 184
- “Synchronization modes for high availability disaster recovery” on page 181
- “System requirements for high availability disaster recovery” on page 174
- “Restrictions for high availability disaster recovery” on page 175
- “Database configuration for high availability disaster recovery” on page 176
- “Comparison of Q replication to high availability disaster recovery (HADR)” in the *IBM DB2 Information Integrator Introduction to Replication and Event Publishing*

---

## System requirements for high availability disaster recovery

To achieve optimal performance with high availability disaster recovery (HADR), ensure that your system meets the following requirements for hardware, operating systems, and DB2<sup>®</sup> Universal Database (DB2 UDB).

**Recommendation:** For better performance, use the same hardware and software for the system where the primary database resides and for the system where the standby database resides. If the system where the standby database resides has fewer resources than the system where the primary database resides, it is possible that the standby database will be unable to keep up with the transaction load generated by the primary database. This can cause the standby database to fall behind or the performance of the primary database to degrade. In a failover situation, the new primary database should have the resources to service the client applications adequately.

**Hardware and operating system requirements:**

**Recommendation:** Use identical host computers for the HADR primary and standby databases. That is, they should be from the same vendor and have the same architecture.

The operating system on the primary and standby databases should be the same version, including patches. You can violate this rule for a short time during a rolling upgrade, but take extreme caution.

A TCP/IP interface must be available between the HADR host machines, and a high-speed, high-capacity network is recommended.

**DB2 UDB requirements:**

The database versions used by the primary and standby databases must be identical. During rolling upgrades, the database version of the standby database may be later than the primary database for a short time. The DB2 UDB version of the primary database can never be later than the standby database version.



The DB2 UDB software for both the primary and standby databases must have the same bit size (32 or 64 bit). Table spaces and their containers must be identical on the primary and standby databases. Properties that must be identical include the table space type (DMS or SMS), table space size, container path, container size, and container file type (raw device or file system). The amount of space allocated for log files should also be the same on both the primary and standby databases.

When you issue a table space statement on the primary database, such as CREATE TABLESPACE, ALTER TABLESPACE, or DROP TABLESPACE, it is replayed on the standby database. You must ensure that the devices involved are set up on both of the databases before you issue the table space statement on the primary database.

If the table space setup is not identical on the primary and standby databases, log replay on the standby database might encounter errors such as OUT OF SPACE or TABLE SPACE CONTAINER NOT FOUND. If this occurs, the affected table space is put in rollforward pending state and is ignored in subsequent log replay. If a takeover operation occurs, the table space will not be available to applications. If the necessary backup images and log file archives are available, you might be able to recover the table space by first issuing the SET TABLESPACE CONTAINERS statement with the IGNORE ROLLFORWARD CONTAINER OPERATIONS option and then issuing the ROLLFORWARD command.

The primary and standby databases do not require the same database path. If relative container paths are used, the same relative path may map to different absolute container paths on the primary and standby databases.

#### **Buffer pool requirements:**

Since buffer pool operations are also replayed on the standby database, it is important that the primary and standby databases have the same amount of memory.

#### **Related concepts:**

- “Restrictions for high availability disaster recovery” on page 175

#### **Related tasks:**

- “Performing a rolling upgrade in a high availability disaster recovery environment” on page 203

---

## **Restrictions for high availability disaster recovery**

The following list is a summary of high availability disaster recovery (HADR) restrictions:

- HADR is supported on DB2<sup>®</sup> UDB Enterprise Server Edition (ESE) only. However, it is not supported when you have multiple database partitions on ESE.
- The primary and standby databases must have the same operating system version and the same version of DB2 UDB, except for a short time during a rolling upgrade.
- The DB2 UDB release on the primary and standby databases must be the same bit size (32 or 64 bit).
- Reads on the standby database are not supported. Clients cannot connect to the standby database.
- Log archiving can only be performed by the current primary database.

- Backup operations are not supported on the standby database.
- Non-logged operations, such as changes to database configuration parameters and to the recovery history file, are not replicated to the standby database.
- Load operations with the COPY NO option specified are not supported.
- Use of Data Links is not supported.

**Related concepts:**

- “High availability disaster recovery overview” on page 173
- “System requirements for high availability disaster recovery” on page 174
- “Replicated operations for high availability disaster recovery” on page 186
- “Non-replicated operations for high availability disaster recovery” on page 187

---

## Database configuration for high availability disaster recovery

To achieve optimal performance with high availability disaster recovery (HADR), ensure that your database configuration meets the following requirements.

**Recommendation:** To the extent possible, the database configuration parameters and database manager configuration parameters should be identical on the systems where the primary and standby databases reside. If the configuration parameters are not properly set on the standby database the following problems might occur:

- Error messages might be returned on the standby database while replaying the log files that were shipped from the primary database.
- After a takeover operation, the new primary database will not be able to handle the workload, resulting in performance problems or in applications receiving error messages they were not receiving when they were connected to the original primary database.

Changes to the configuration parameters on the primary database are not automatically propagated to the standby database and must be made manually on the standby database. For dynamic configuration parameters, changes will take effect without shutting down and restarting the database management system (DBMS) or the database. For non-dynamic configuration parameters, changes will take effect after the standby database is restarted.

**Note:** An exception is the LOGFILSIZ database configuration parameter. Although this parameter is not replicated to the standby database, to guarantee identical log files on both databases, the standby database ignores the local LOGFILSIZ configuration and creates local log files that match the size of the log files on the primary database.

**Log receive buffer size on the standby database:**

By default, the log receive buffer size on the standby database will be two times the value specified for the LOGBUFSZ configuration parameter on the primary database. There might be times when this size is not sufficient. For example, when the HADR synchronization mode is asynchronous and the primary and standby databases are in peer state, if the primary database is experiencing a high transaction load, the log receive buffer on the standby database might fill to capacity and the log shipping operation from the primary database might stall. To manage these temporary peaks, you can increase the size of the log receive buffer on the standby database by modifying the DB2\_HADR\_BUF\_SIZE registry variable.

## Load operations and HADR:

If a load operation is executed on the primary database with the COPY YES option, the command will execute on the primary database and the data will be replicated to the standby database as long as the copy can be accessed through the path or device specified by the LOAD command. If the standby database cannot access the data, the table space in which the table is stored is marked bad on the standby database. The standby database will skip future log records that pertain to this table space.

If a load operation is executed on the primary database with the NONRECOVERABLE option, the command will execute on the primary database and the table on the standby database will be marked bad. The standby database will skip future log records that pertain to this table. You can choose to issue the LOAD command with the COPY YES and REPLACE options specified to bring the table back, or you can drop the table to recover the space.

Since executing a load operation with the COPY NO option is not supported with HADR, the command is automatically converted to a load operation with the NONRECOVERABLE option. To enable a load operation with the COPY NO option to be converted to a load operation with the COPY YES option, set the DB2\_LOAD\_COPY\_NO\_OVERRIDE registry variable on the primary database. This registry variable is ignored by the standby database. Ensure that the device or directory specified on the primary database can be accessed by the standby database using the same path, device, or load library.

If you are using Tivoli® Storage Manager (TSM) to perform a load operation with the COPY YES option, you might need to set the VENDOROPT configuration parameter on the primary and standby databases. Depending on how TSM is configured, the values on the primary and standby databases might not be the same. Also, when using TSM to perform a load operation with the COPY YES option, you must issue the **db2adutl** command with the GRANT option to give the standby database read access for the files that are loaded.

If table data is replicated by a load operation with the COPY YES option specified, the indexes will be replicated as follows:

- If the indexing mode is set to REBUILD and the table attribute is set to LOG INDEX BUILD, or the table attribute is set to DEFAULT and the LOGINDEXBUILD database configuration parameter is set to ON, the primary database will include the rebuilt index object in the copy file to enable the standby database to replicate the index object. If the index object on the standby database is marked bad before the load operation, it will become usable again after the load operation as a result of the index rebuild.
- If the indexing mode is set to INCREMENTAL and the table attribute is set to LOG INDEX BUILD, or the table attribute is set to NULL and LOGINDEXBUILD database configuration parameter on the primary database is set to ON, the index object on the standby database is updated only if it is not marked bad before the load operation. Otherwise, the index is marked bad on the standby database.

## HADR configuration parameters:

Several new database configuration parameters are available to support HADR. Setting these parameters does not change the role of a database. You must issue the START HADR or STOP HADR commands to change the role of a database.

Any changes made to an HADR configuration parameter are not effective until the database has been shut down and restarted. In a partitioned database environment, the HADR configuration parameters are visible and can be changed, but they are ignored.

The local host name of the primary database must be the same as the remote host name of the standby database, and the local host name of the standby database must be the same as the remote host name of the primary database. Use the HADR\_LOCAL\_HOST and HADR\_REMOTE\_HOST configuration parameters to set the local and remote hosts for each database. Configuration consistency for the local and remote host names is checked when a connection is established to ensure that the remote host specified is the expected node.

The synchronization mode (HADR\_SYNCMODE) and time-out period (HADR\_TIMEOUT) must be identical on both the primary and standby databases. The consistency of these configuration parameters is checked when an HADR pair establishes a connection.

TCP connections are used for communication between the primary and standby databases. A primary database that is not connected to a standby database, either because it is starting up or because the connection is lost, will listen on its local port for new connections. A standby database that is not connected to a primary database will continue issuing connection requests to its remote host.

Although the local host and local service parameters (HADR\_LOCAL\_HOST, HADR\_LOCAL\_SVC) are only used on the primary database, you should still set them on the standby database to ensure that they are ready if the standby database has to take over as the primary database.

The following sample configuration is for the primary and standby databases.

On the primary:

```
HADR_LOCAL_HOST host1.ibm.com
HADR_LOCAL_SVC  hadr_service
HADR_REMOTE_HOST host2.ibm.com
HADR_REMOTE_SVC hadr_service
HADR_REMOTE_INST dbinst2
HADR_TIMEOUT    120
HADR_SYNCMODE   NEARSYNC
```

On the standby:

```
HADR_LOCAL_HOST host2.ibm.com
HADR_LOCAL_SVC  hadr_service
HADR_REMOTE_HOST host1.ibm.com
HADR_REMOTE_SVC hadr_service
HADR_REMOTE_INST dbinst1
HADR_TIMEOUT    120
HADR_SYNCMODE   NEARSYNC
```

#### **Related concepts:**

- “High availability disaster recovery overview” on page 173

#### **Related reference:**

- “hadr\_db\_role - HADR database role configuration parameter” in the *Administration Guide: Performance*
- “hadr\_local\_host - HADR local host name configuration parameter” in the *Administration Guide: Performance*

- “hadr\_local\_svc - HADR local service name configuration parameter” in the *Administration Guide: Performance*
- “hadr\_remote\_host - HADR remote host name configuration parameter” in the *Administration Guide: Performance*
- “hadr\_remote\_inst - HADR instance name of the remote server configuration parameter” in the *Administration Guide: Performance*
- “hadr\_remote\_svc - HADR remote service name configuration parameter” in the *Administration Guide: Performance*
- “hadr\_syncmode - HADR synchronization mode for log write in peer state configuration parameter” in the *Administration Guide: Performance*
- “hadr\_timeout - HADR timeout value configuration parameter” in the *Administration Guide: Performance*
- “vendoropt - Vendor options configuration parameter” in the *Administration Guide: Performance*

---

## Standby database states in high availability disaster recovery

With the high availability disaster recovery (HADR) feature, when the standby database is started, it enters local catchup state and attempts to read the log files in its local log path. If it does not find a log file in the local log path and a log archiving method has been specified, the log file is retrieved using the specified method. After the log files are read, they are replayed on the standby database. During this time, a connection to the primary database is not required; however, if a connection does not exist, the standby database tries to connect to the primary database. When the end of local log files is reached, the standby database enters remote catchup pending state.

The standby database remains in remote catchup pending state until a connection to the primary database is established, at which time the standby database enters remote catchup state. During this time, the primary database reads log data from its log path or by way of a log archiving method and sends the log files to the standby database. The standby database receives and replays the log data. When all of the log files on disk have been replayed by the standby database, the primary and standby systems enter peer state.

When in peer state, log pages are shipped to the standby database whenever the primary database flushes a log page to disk. The log pages are written to the local log files on the standby database to ensure that the primary and standby databases have identical log file sequences. The log pages can then be replayed on the standby database.

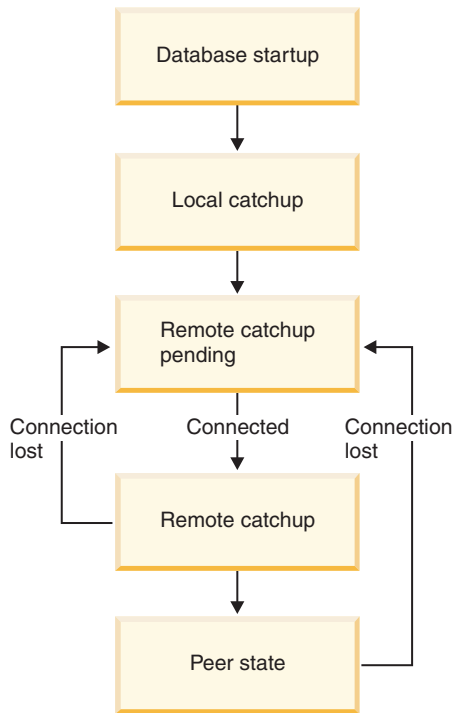


Figure 16. States of the standby database

You can see the state of the standby database by issuing the GET SNAPSHOT command with the DATABASE ON option. For example, if you have standby database MUSIC, you can issue the following command to see its state:

```
get snapshot for database on music
```

If the command is issued on the standby database, one of the standby database states is returned in the State field. If the query is issued on a primary database that is connected to a standby database, the state of the standby database is returned. If the primary database is not connected to a standby database, disconnected is returned.

The following output shows the HADR status section returned by the GET SNAPSHOT command:

```
HADR status

Role                = Primary
State               = Peer
Synchronization mode = Sync
Connection status   = Connected, 11-03-2002 12:23:09.35092
Heartbeat missed    = 0
Local host          = host1.ibm.com
Local service       = hadr_service
Remote host         = host2.ibm.com
Remote service      = hadr_service
Remote instance     = dbinst2
timeout(seconds)    = 120
Primary log position(file, page, LSN) = S0001234.LOG, 12, 0000000000BB800C
Standby log position(file, page, LSN) = S0001234.LOG, 12, 0000000000BB800C
Log gap running average(bytes) = 8723
```

#### Notes:

1. If the connection between the primary and standby databases is lost during remote catchup or peer state, the standby database will enter remote catchup pending state.
2. Because the standby database writes the log files it receives to its local log path, you must not use a shared network or local file system as the log path for both the primary and standby databases. You will receive an error message if DB2® detects a shared log path.
3. To speed up the catchup process, you can use a shared log archive device. However, if the shared device is a serial device such as a tape drive, you might experience performance degradation on both the primary and standby databases because of mixed read and write operations.
4. You can manually copy the primary database log files into the standby database log path to be used for local catchup. This must be done before you start the standby database, because when the end of the local log files is reached, the standby database will enter remote catchup pending state and will not try to access the log files again. Also, if the standby database enters remote catchup state, copying log files into its log path will interfere with the writing of local log files by the standby database. If more local log files become available after the standby database enters remote catchup pending state, you can shut down the standby database and restart it to ensure that it re-enters local catchup state.

#### Related concepts:

- “Managing log files” on page 42
- “Managing log files through log archiving” on page 45

#### Related reference:

- Appendix G, “User exit for database recovery,” on page 329
- “GET SNAPSHOT Command” in the *Command Reference*
- “Configuration parameters for database logging” on page 35
- “logarchmeth1 - Primary log archive method configuration parameter” in the *Administration Guide: Performance*
- “logarchmeth2 - Secondary log archive method configuration parameter” in the *Administration Guide: Performance*

---

## Synchronization modes for high availability disaster recovery

With high availability disaster recovery (HADR), you can specify one of three synchronization modes to choose your preferred level of protection from potential loss of data. The synchronization mode indicates how log writing is managed between the primary and standby databases. These modes apply only when the primary and standby databases are in peer state.

Use the HADR\_SYNCMODE configuration parameter to set the synchronization mode. Valid values are:

#### SYNC (synchronous)

This mode provides the greatest protection against transaction loss, and using it results in the longest transaction response time among the three modes.

In this mode, log writes are considered successful only when logs have been written to log files on the primary database and when the primary

database has received acknowledgement from the standby database that the logs have also been written to log files on the standby database. The log data is guaranteed to be stored at both sites.

If the standby database crashes before it can replay the log records, the next time it starts it can retrieve and replay them from its local log files. If the primary database fails, a failover to the standby database guarantees that any transaction that has been committed on the primary database has also been committed on the standby database. After the failover operation, when the client reconnects to the new primary database, there may be transactions committed on the new primary database that were never reported as committed to the original primary. This occurs when the primary database fails before it processes an acknowledgement message from the standby database. Client applications should consider querying the database to determine whether any such transactions exist.

If the primary database loses its connection to the standby database, the databases are no longer considered to be in peer state and transactions will not be held back waiting for acknowledgement from the standby database. If the failover operation is performed when the databases are disconnected, there is no guarantee that all of the transactions committed on the primary database will appear on the standby database.

If the primary database fails when the databases are in peer state, it can rejoin the HADR pair as a standby database after a failover operation. Because a transaction is not considered to be committed until the primary database receives acknowledgement from the standby database that the logs have also been written to log files on the standby database, the log sequence on the primary will be the same as the log sequence on the standby database. The original primary database (now a standby database) just needs to catch up by replaying the new log records generated on the new primary database since the failover operation.

If the primary database is not in peer state when it fails, its log sequence might be different from the log sequence on the standby database. If a failover operation has to be performed, the log sequence on the primary and standby databases might be different because the standby database starts its own log sequence after the failover. Because some operations cannot be undone (for example, dropping a table), it is not possible to revert the primary database to the point in time when the new log sequence was created. If the log sequences are different, an error message will be returned when you issue the START HADR command with the AS STANDBY option on the original primary. If the original primary database successfully rejoins the HADR pair, you can achieve failback of the database by issuing the TAKEOVER HADR command without specifying the BY FORCE option. If the original primary database cannot rejoin the HADR pair, you can reinitialize it as a standby database by restoring a backup image of the new primary database.

#### **NEARSYNC (near synchronous)**

While this mode has a shorter transaction response time than synchronous mode, it also provides slightly less protection against transaction loss.

In this mode, log writes are considered successful only when the log records have been written to the log files on the primary database and when the primary database has received acknowledgement from the standby system that the logs have also been written to main memory on the standby system. Loss of data occurs only if both sites fail



simultaneously and if the target site has not transferred to nonvolatile storage all of the log data that it has received.

If the standby database crashes before it can copy the log records from memory to disk, the log records will be lost on the standby database. Usually, the standby database can get the missing log records from the primary database when the standby database restarts. However, if a failure on the primary database or the network makes retrieval impossible and a failover is required, the log records will never appear on the standby database, and transactions associated with these log records will never appear on the standby database.

If transactions are lost, the new primary database is not identical to the original primary database after a failover operation. Client applications should consider resubmitting these transactions to bring the application state up to date.

If the primary database fails when the primary and standby databases are in peer state, it is possible that the original primary database cannot to rejoin the HADR pair as a standby database without being reinitialized using a full restore operation. If the failover involves lost log records (because both the primary and standby databases have failed), the log sequences on the primary and standby databases will be different and attempts to restart the original primary database as a standby database without first performing a restore operation will fail. If the original primary database successfully rejoins the HADR pair, you can achieve failback of the database by issuing the TAKEOVER HADR command without specifying the BY FORCE option. If the original primary database cannot rejoin the HADR pair, you can reinitialize it as a standby database by restoring a backup image of the new primary database.

#### **ASYNC (asynchronous)**

This mode has the highest chance of transaction loss if the primary system fails. It also has the shortest transaction response time among the three modes.

In this mode, log writes are considered successful only when the log records have been written to the log files on the primary database and have been delivered to the TCP layer of the primary system's host machine. Because the primary system does not wait for acknowledgement from the standby system, transactions might be considered committed when they are still on their way to the standby.

A failure on the primary database host machine, on the network, or on the standby database can cause log files in transit to be lost. If the primary database is available, the missing log files can be resent to the standby database when the pair reestablishes a connection. However, if a failover operation is required while there are missing log files, both the log files and the associated transactions will never reach the standby database. Permanent loss of transactions is caused by lost log files and a failure on the primary database.

If transactions are lost, the new primary database is not exactly the same as the original primary database after a failover operation. Client applications should consider resubmitting these transactions to bring the application state up to date.

If the primary database fails when the primary and standby databases are in peer state, it is possible that the original primary database will not be able to rejoin the HADR pair as a standby database without being

reinitialized using a full restore operation. If the failover involves lost log records, the log sequences on the primary and standby databases will be different, and attempts to restart the original primary database as a standby database will fail. Because there is a greater possibility of log records being lost if a failover occurs in asynchronous mode, there is also a greater possibility that the primary database will not be able to rejoin the HADR pair. If the original primary database successfully rejoins the HADR pair, you can achieve failback of the database by issuing the TAKEOVER HADR command without specifying the BY FORCE option. If the original primary database cannot rejoin the HADR pair, you can reinitialize it as a standby database by restoring a backup image of the new primary database.

**Related concepts:**

- “Standby database states in high availability disaster recovery” on page 179

**Related tasks:**

- “Reintegrating a database after a takeover operation” on page 202

**Related reference:**

- “`hadr_syncmode` - HADR synchronization mode for log write in peer state configuration parameter” in the *Administration Guide: Performance*

---

## Automatic client reroute and high availability disaster recovery

The automatic client reroute feature can be used with high availability disaster recovery to allow client applications to recover from a loss of communication with the server and to continue working with minimal interruption. Rerouting is only possible when an alternate database location has been specified at the server. Automatic client reroute is only supported with TCP/IP protocol.

You can use automatic client reroute with HADR to make client applications connect to the new primary database after a takeover operation. If automatic client reroute is not enabled, client applications will receive error message SQL30081, and no further attempts will be made to establish a connection with the server. The following example explains how to use the UPDATE ALTERNATE SERVER FOR DATABASE command to set up automatic client reroute with HADR.

**Example**

Your system is set up as follows:

- You have a client where database MUSIC is catalogued as being located at host HORNET.
- Database MUSIC is the primary database and its corresponding standby database, also MUSIC, resides on host MONTERO with port number 456.

To enable automatic client reroute, update the alternate server for database MUSIC on host HORNET:

```
db2 update alternate server for database music using hostname montero port 456
```

After this command is issued, the client must successfully connect to host HORNET to obtain the alternate server information. Then, if a communication error occurs between the client and database MUSIC at host HORNET, the client will

first attempt to reconnect to database MUSIC at host HORNET. If this fails, the client will then attempt to establish a connection with the standby database MUSIC on host MONTERO.

**Notes:**

1. The alternate host location is stored in the system database directory file at the server.
2. To enable the automatic client reroute feature, you must use the UPDATE ALTERNATE SERVER FOR DATABASE command. Automatic client reroute does not use the HADR\_REMOTE\_HOST and HADR\_REMOTE\_SVC database configuration parameters.

**Related concepts:**

- “High availability disaster recovery overview” on page 173
- “Automatic client reroute implementation” in the *Administration Guide: Implementation*

**Related reference:**

- “UPDATE ALTERNATE SERVER FOR DATABASE Command” in the *Command Reference*
- “db2UpdateAlternateServerForDB - Update Alternate Server for Database” in the *Administrative API Reference*

---

## Index logging and high availability disaster recovery

Consider the following recommendations when setting configuration parameters for high availability disaster recovery (HADR) databases.

### Using the LOGINDEXBUILD database configuration parameter

**Recommendation:** For HADR databases, set the LOGINDEXBUILD database configuration parameter to ON to ensure that complete information is logged for index creation, recreation, and reorganization. Although this means that index builds might take longer on the primary system and that more log space is required, the indexes will be rebuilt on the standby system during HADR log replay and will be available when a failover takes place. If index builds on the primary system are not logged and a failover occurs, any invalid indexes that remain after the failover is complete will have to be rebuilt before they can be accessed. While the indexes are being recreated, they cannot be accessed by any applications.

**Note:** If the LOG INDEX BUILD table attribute is set to its default value of NULL, DB2® will use the value specified for the LOGINDEXBUILD database configuration parameter. If the LOG INDEX BUILD table attribute is set to ON or OFF, the value specified for the LOGINDEXBUILD database configuration parameter will be ignored.

You might choose to set the LOG INDEX BUILD table attribute to OFF on one or more tables for either of the following reasons:

- You do not have enough active log space to support logging of the index builds.
- The index data is very large and the table is not accessed often; therefore, it is acceptable for the indexes to be recreated at the end of the takeover operation. In this case, set the INDEXREC configuration parameter to RESTART. Because the table is not frequently accessed, this setting will cause the system to recreate

the indexes at the end of the takeover operation instead of waiting for the first time the table is accessed after the takeover operation.

If the LOG INDEX BUILD table attribute is set to OFF on one or more tables, any index build operation on those tables might cause the indexes to be recreated any time a takeover operation occurs. Similarly, if the LOG INDEX BUILD table attribute is set to its default value of NULL, and the LOGINDEXBUILD database configuration parameter is set to OFF, any index build operation on a table might cause the indexes on that table to be recreated any time a takeover operation occurs. You can prevent the indexes from being recreated by taking one of the following actions:

- After all invalid indexes are recreated on the new primary database, take a backup of the database and apply it to the standby database. As a result of doing this, the standby database does not have to apply the logs used for recreating invalid indexes on the primary database, which would mark those indexes as rebuild required on the standby database.
- Set the LOG INDEX BUILD table attribute to ON, or set the LOG INDEX BUILD table attribute to NULL and the LOGINDEXBUILD configuration parameter to ON on the standby database to ensure that the index recreation will be logged.

#### Using the INDEXREC database configuration parameter

**Recommendation:** Set the INDEXREC database configuration parameter to RESTART (the default) on both the primary and standby databases. This will cause invalid indexes to be rebuilt after a takeover operation is complete. If any index builds have not been logged, this setting allows DB2 to check for invalid indexes and to rebuild them. This process takes place in the background, and the database will be accessible after the takeover operation has completed successfully.

If a transaction accesses a table that has invalid indexes before the indexes have been rebuilt by the background recreate index process, the invalid indexes will be rebuilt by the first transaction that accesses it.

#### Related reference:

- “indexrec - Index re-creation time configuration parameter” in the *Administration Guide: Performance*
- “ALTER TABLE statement” in the *SQL Reference, Volume 2*
- “logindexbuild - Log index pages created configuration parameter” in the *Administration Guide: Performance*

---

## Replicated operations for high availability disaster recovery

In high availability disaster recovery (HADR), the following operations are replicated from the primary to the standby database:

- Data definition language (DDL)
- Data manipulation language (DML)
- Buffer pool operations
- Table space operations
- Online reorganization
- Offline reorganization
- Meta-data for stored procedures and user defined functions (UDF) (but not the related object or library files)

During an online reorganization, all operations are logged in detail. As a result, HADR can replicate the operation without the standby database falling further behind than it would for more typical database updates. However, this behavior can potentially have a large impact on the system because of the large number of log records generated.

While offline reorganizations are not logged as extensively as online reorganizations, operations are typically logged per hundreds or thousands of affected rows. This means that the standby database could fall behind because it waits for each log record and then replays many updates at once. If the offline reorganization is non-clustered, a single log record is generated after the entire reorganization operation. This mode has the greatest impact on the ability of the standby database to keep up with the primary database. The standby database will perform the entire reorganization after it receives the log record from the primary database.

HADR does not replicate stored procedure and UDF object and library files. You must create the files on identical paths on both the primary and standby databases. If the standby database cannot find the referenced object or library file, the stored procedure or UDF invocation will fail on the standby database.

**Related concepts:**

- “Table reorganization” in the *Administration Guide: Performance*
- “Index reorganization” in the *Administration Guide: Performance*
- “Database configuration for high availability disaster recovery” on page 176
- “Non-replicated operations for high availability disaster recovery” on page 187

---

## Non-replicated operations for high availability disaster recovery

High availability disaster recovery (HADR) uses database logs to replicate data to the standby database. Non-logged operations are allowed on the primary database, but not replicated to the standby database. Non-replicated operations include, but are not limited to, the following operations:

- Tables created with the NOT LOGGED INITIALLY option specified.
- BLOBs and CLOBs are not replicated; however, the space for them will be allocated on the standby database.
- Datalinks are not supported in HADR databases. If the START HADR or the ACTIVATE DATABASE command is issued, or if a first client connection brings up a database that is in an HADR role (primary or standby), and the DATALINKS database manager configuration parameter is set to YES, the operation will fail. To use the HADR database, set the DATALINKS configuration parameter to NO and shutdown and restart the instance.

**Note:** Existing datalink columns are not affected when a database is converted from a standard database to a primary or standby database. Even if the DATALINKS configuration parameter is set to NO, you can still create new datalink columns in an HADR database. However, you cannot insert, update or select datalink columns.

- Updates to database configuration using the UPDATE DATABASE CONFIGURATION and UPDATE DATABASE MANAGER CONFIGURATION commands are not replicated.
- The recovery history file and changes to it are not automatically shipped from the primary database to the standby database.

You can place an initial copy of the history file (obtained from the backup image of the primary) on the standby database by issuing the RESTORE DATABASE command with the REPLACE HISTORY FILE option:

```
RESTORE DB KELLY REPLACE HISTORY FILE
```

After HADR is initialized and backup activities take place on the primary database, the history file on the standby database will become out of date. You can issue the following command to update the history file on the standby when backup operations occur on the primary database:

```
RESTORE DB KELLY HISTORY FILE
```

If a takeover operation occurs and the standby database has an up-to-date history file, backup and restore operations on the new primary will generate new records in the history file and blend seamlessly with the records generated on the original primary. If the history file is out of date or has missing entries, an automatic incremental restore might not be possible; instead, a manual incremental restore operation will be required.

**Related concepts:**

- “Table reorganization” in the *Administration Guide: Performance*
- “Index reorganization” in the *Administration Guide: Performance*
- “Understanding the recovery history file” on page 51
- “Database configuration for high availability disaster recovery” on page 176
- “Replicated operations for high availability disaster recovery” on page 186

---

## Cluster managers and high availability disaster recovery

You can use high availability disaster recovery (HADR) with a cluster manager to enhance the availability of the DBMS. There are two ways you can configure HADR to do this:

- Set up an HADR pair where the primary and standby databases are serviced by the same cluster manager.

This configuration is best suited to environments where the primary and standby databases are located at the same site and where the fastest possible failover is required. These environments would benefit from using HADR to maintain DBMS availability, rather using crash recovery or another recovery method.

You can use the cluster manager to quickly detect a problem and to initiate a takeover operation. Because HADR requires separate storage for the DBMS, the cluster manager should be configured with separate volume control. This configuration prevents the cluster manager from waiting for failover to occur on the volume before using the DBMS on the standby system. You can use the automatic client reroute feature to redirect client applications to the new primary database.

- Set up an HADR pair where the primary and standby databases are not serviced by the same cluster manager.

This configuration is best suited to environments where the primary and standby databases are located at different sites and where high availability is required for disaster recovery in the event of a complete site failure. There are several ways you can implement this configuration. When an HADR primary or standby database is part of a cluster, there are two possible failover scenarios.

- If a partial site failure occurs and a node to which the DBMS can fail over remains available, you can choose to perform a cluster failover. In this case, the IP address and volume failover is performed using the cluster manager; HADR is not affected.
- If a complete site failure occurs where the primary database is located, you can use HADR to maintain DBMS availability by initiating a takeover operation. If a complete site failure occurs where the standby database is located, you can repair the site or move the standby database to another site.

**Related concepts:**

- “High availability disaster recovery overview” on page 173
- “Automatic client reroute and high availability disaster recovery” on page 184

**Related tasks:**

- “Switching database roles in high availability disaster recovery” on page 198

## Initializing high availability disaster recovery

Use the following procedure to set up and initialize the primary and standby databases for high availability disaster recovery (HADR).

**Procedure:**

HADR can be initialized through the command line processor (CLP), the Set Up High Availability Disaster Recovery (HADR) wizard in the Control Center, or by invoking the corresponding application programming interfaces (APIs).

To use the CLP to initialize HADR on your system for the first time:

1. Determine the host name, host IP address, and the service name or port number for each of the HADR databases.

If a host has multiple network interfaces, ensure that the HADR host name or IP address maps to the intended one. The host name can only map to one IP address.

**Note:** The instance names for the primary and standby databases do not have to be the same.

2. Create the standby database by restoring a backup image or by initializing a split mirror, based on the existing database that is to be the primary.

In the following example, the BACKUP DATABASE and RESTORE DATABASE commands are used to initialize database SOCKS as a standby database. In this case, an NFS mounted file system is accessible at both sites.

Issue the following command at the primary database:

```
backup db socks to /nfs1/backups/db2/socks
```

Issue the following command at the standby database:

```
restore db socks from /nfs1/backups/db2/socks replace history file
```

The following example illustrates how to use the **db2inidb** utility to initialize the standby database using a split mirror of the primary database. This procedure is an alternative to the backup and restore procedure illustrated above.

Issue the following command at the standby database:

```
db2inidb socks as standby
```

**Notes:**

- a. The database names for the primary and standby databases must be the same.
  - b. It is recommended that you do not issue the `ROLLFORWARD DATABASE` command on the standby database after the restore operation or split mirror initialization. The results of using a rollforward operation may differ slightly from replaying the logs using HADR on the standby database. If the databases are not identical, issuing the `START HADR` command with the `AS STANDBY` option will fail.
  - c. When using the `RESTORE DATABASE` command, it is recommended that the `REPLACE HISTORY FILE` option is used.
  - d. The following `RESTORE DATABASE` command options should be avoided when setting up the standby database: `TABLESPACE`, `INTO`, `REDIRECT`, and `WITHOUT ROLLING FORWARD`.
  - e. When setting up the standby database using the `db2inidb` utility, do not use the `SNAPSHOT` or `MIRROR` options. You can specify the `RELOCATE USING` option to change one or more of the following configuration attributes: instance name, log path, and database path. However, you must not change the database name or the table space container paths.
3. Set the HADR configuration parameters on the primary and standby databases.

**Note:** It is very important that you set the following configuration parameters after the standby databases has been created:

- `HADR_LOCAL_HOST`
- `HADR_LOCAL_SVC`
- `HADR_REMOTE_HOST`
- `HADR_REMOTE_SVC`
- `HADR_REMOTE_INST`

If they are set prior to creating the standby database, the settings on the standby database will reflect what is set on the primary database.

4. Connect to the standby instance and start HADR on the standby database, as in the following example:

```
START HADR ON DB SOCKS AS STANDBY
```

**Note:** Usually, the standby database is started first. If you start the primary database first, this startup procedure will fail if the standby database is not started within the time period specified by the `HADR_TIMEOUT` database configuration parameter.

5. Connect to the primary instance and start HADR on the primary database, as in the following example:

```
START HADR ON DB SOCKS AS PRIMARY
```

6. HADR is now started on the primary and standby databases.

To open the Set Up High Availability Disaster Recovery (HADR) Databases wizard:

1. From the Control Center expand the object tree until you find the database for which you want to configure HADR.
2. Right-click the database and click High Availability Disaster Recovery → Set Up in the pop-up menu. The Set Up High Availability Disaster Recovery Databases wizard opens.



Detailed information is provided through the contextual help facility within the Control Center.

**Related concepts:**

- “Database configuration for high availability disaster recovery” on page 176

**Related tasks:**

- “Using restore” on page 82
- “Using a split mirror as a backup image” on page 168

**Related reference:**

- “START HADR” on page 191
- “STOP HADR” on page 196

---

## START HADR

Starts HADR operations for a database.

**Authorization:**


One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

**Required connection:**

Instance. The command establishes a database connection if one does not exist, and closes the database connection when the command completes.

**Command syntax:**



**Command parameters:**

**DATABASE** *database-alias*

Identifies the database on which HADR operations are to start.

**USER** *user-name*

Identifies the user name under which the HADR operations are to be started.

**USING** *password*

The password used to authenticate *user-name*.

**AS PRIMARY**

Specifies that HADR primary operations are to be started on the database.

**BY FORCE**

Specifies that the HADR primary database will not wait for the standby database to connect to it. After a start BY FORCE, the

## START HADR

primary database will still accept valid connections from the standby database whenever the standby later becomes available.

**Caution:** Use the START HADR command with the AS PRIMARY BY FORCE option with caution. If the standby database has been changed to a primary and the original primary database is restarted by issuing the START HADR command with the AS PRIMARY BY FORCE option, both copies of your database will be operating independently as primaries. (This is sometimes referred to as *split brain* or *dual primary*.) In this case, each primary database can accept connections and perform transactions, and neither receives and replays the updates made by the other. As a result, the two copies of the database will become inconsistent with each other.

### AS STANDBY

Specifies that HADR standby operations are to be started on the database. The standby database will attempt to connect to the HADR primary database until a connection is successfully established, or until the connection attempt is explicitly rejected by the primary. (The connection might be rejected by the primary database if an HADR configuration parameter is set incorrectly or if the database copies are inconsistent, both conditions for which continuing to retry the connection is not appropriate.)

### Usage notes:

The following table shows database behavior in various conditions:

Database status	Behavior upon START HADR command with the AS PRIMARY option	Behavior upon START HADR command with the AS STANDBY option
Inactive standard database	Activated as HADR primary database.	Database starts as an standby database if it is in rollforward-pending mode (which can be the result of a restore or a split mirror) or in rollforward in-progress mode. Otherwise, an error is returned.
Active standard database	Database enters HADR primary role.	Error message returned.
Inactive primary database	Activated as HADR primary database.	After a failover, this reintegrates the failed primary into the HADR pair as the new standby database. Some restrictions apply.
Active primary database	Warning message issued.	Error message returned.
Inactive standby database	Error message returned.	Starts the database as the standby database.
Active standby database	Error message returned.	Warning message issued.

---

## db2HADRStart - Start HADR

Starts HADR operations on a database.

### Authorization:

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

### Required connection:

Instance. The API establishes a database connection if one does not exist, and closes the database connection when the API completes.

### API include file:

*db2ApiDf.h*

### C API syntax:

```

/* File: db2ApiDf.h */
/* API: db2HADRStart */
/* ... */
SQL_API_RC SQL_API_FN
db2HADRStart (
    db2UInt32 versionNumber,
    void * pParmStruct,
    struct sqlca * pSqlca);

typedef SQL_STRUCTURE db2HADRStartStruct
{
    char                *piDbAlias;
    char                *piUserName;
    char                *piPassword;
    db2UInt32           iDbRole;
    db2UInt16           iByForce;
} db2HADRStartStruct;

```

### Generic API syntax:

```

/* File: db2ApiDf.h */
/* API: db2gHADRStart */
/* ... */
SQL_API_RC SQL_API_FN
db2gHADRStart (
    db2UInt32 versionNumber,
    void * pParmStruct,
    struct sqlca * pSqlca);

typedef SQL_STRUCTURE db2gHADRStartStruct
{
    char                *piDbAlias;
    db2UInt32           iAliasLen;
    char                *piUserName;
    db2UInt32           iUserNameLen;
    char                *piPassword;
    db2UInt32           iPasswordLen;
    db2UInt32           iDbRole;
    db2UInt16           iByForce;
} db2gHADRStartStruct;

```

## db2HADRStart - Start HADR

### API parameters:

#### versionNumber

Input. Specifies the version and release level of the structure passed as the second parameter *pParmStruct*.

#### pParmStruct

Input. A pointer to the *db2HADRStartStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### piDbAlias

Input. A pointer to the database alias.

#### iAliasLen

Input. Specifies the length in bytes of the database alias.

#### piUserName

Input. A pointer to the user name under which the command will be executed.

#### iUserNameLen

Input. Specifies the length in bytes of the user name.

#### piPassword

Input. A pointer to a string containing the password.

#### iPasswordLen

Input. Specifies the length in bytes of the password.

#### iDbRole

Input. Specifies whether the database is to be started as the HADR primary database or as the HADR standby database. Valid values are:

DB2HADR\_DB\_ROLE\_PRIMARY

DB2HADR\_DB\_ROLE\_STANDBY

#### iByForce

Input. This argument is ignored if *iDbRole* is set to DB2HADR\_DB\_ROLE\_STANDBY. Valid values are:

##### DB2HADR\_NO\_FORCE

Specifies that HADR is started on the primary database only if a standby database connects to it within a prescribed time limit.

##### DB2HADR\_FORCE

Specifies that HADR is to be started by force, without waiting for the standby database to connect to the primary database.

### Related reference:

- “db2HADRStop - Stop HADR” on page 197
- “db2HADRTakeover - Take Over as Primary Database” on page 206
- “START HADR” on page 191

---

## Stopping high availability disaster recovery

Use the STOP HADR command to stop high availability disaster recovery (HADR) operations on the primary or standby database. You can choose to stop HADR on one or both of the databases. If you are performing maintenance on the standby system, you only need to stop HADR on the standby database. If you want to stop using HADR completely, you can stop HADR on both databases.

**Restrictions:**

You can issue the STOP HADR command against a primary or a standby database only. If you issue this command against a standard database an error will be returned.

**Procedure:**

You can stop HADR by using the command line processor (CLP), the Manage High Availability Disaster Recovery (HADR) window in the Control Center, or the **db2HADRStop** application programming interface (API).

To use the CLP to stop HADR operations on the primary or standby database, issue the STOP HADR command on the database where you want to stop HADR operations.

In the following example, HADR operations are stopped on database SOCKS:

```
STOP HADR ON DATABASE SOCKS
```

If you issue this command against an inactive primary database, the database switches to a standard database and remains offline.

If you issue this command against an inactive standby database the database switches to a standard database, is placed in rollforward pending state, and remains offline.

If you issue this command on an active primary database, logs stop being shipped to the standby database and all HADR engine dispatchable units (EDUs) are shut down on the primary database. The database switches to a standard database and remains online. Transaction processing can continue. You can issue the START HADR AS PRIMARY command to switch the role of the database back to primary database.

If you issue this command on an active standby database, an error message is returned, indicating that you must deactivate the standby database before attempting to convert it to a standard database.

To open the Manage HADR window:

1. From the Control Center, expand the object tree until you find the database for which you want to manage HADR.
2. Right-click the database and click High Availability Disaster Recovery→Manage in the pop-up menu. The Manage High Availability Disaster Recovery (HADR) window opens.

Detailed information is provided through the contextual help facility within the Control Center.

**Related concepts:**

- “High availability disaster recovery overview” on page 173

**Related reference:**

- “START HADR” on page 191
- “STOP HADR” on page 196

## STOP HADR

Stops HADR operations for a database.

### Authorization:

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

### Required connection:

Instance. The command establishes a database connection if one does not exist, and closes the database connection when the command completes.

### Command syntax:

```

>> STOP HADR ON DATABASE database-alias
                |
                | DB
                |
                |----->
    |
    |----->>>
    |
    | USER user-name
    |         |
    |         | USING password
    |         |
    |         |----->
  
```

### Command parameters:

**DATABASE** *database-alias*

Identifies the database on which HADR operations are to stop.

**USER** *user-name*

Identifies the user name under which the HADR operations are to be stopped.

**USING** *password*

The password used to authenticate *user-name*.

### Usage notes:

The following table shows database behavior in various conditions:

Database status	Behavior upon STOP HADR command
Inactive standard database	Error message returned.
Active standard database	Error message returned.
Inactive primary database	Database role changes to standard. Database configuration parameter <i>hadr_db_role</i> is updated to STANDARD. Database remains offline. At the next restart, enters standard role.

Database status	Behavior upon STOP HADR command
Active primary database	Stops shipping logs to the HADR standby database and shuts down all HADR EDUs on the HADR primary database. Database role changes to standard and database remains online. Database remains in standard role until an explicit START HADR command with the AS PRIMARY option is issued. Open sessions and transactions are not affected by the STOP HADR command. You may repeatedly issue STOP HADR and START HADR commands while the database remains online. These commands take effect dynamically.
Inactive standby database	Database role changes to standard. Database configuration parameter <i>hadr_db_role</i> is updated to STANDARD. Database remains offline. Database is put into rollforward pending mode.
Active standby database	Error message returned: Deactivate the standby database before attempting to convert it to a standard database.

---

## db2HADRStop - Stop HADR

Stops HADR operations on a database.

### Authorization:

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

### Required connection:

Instance. The API establishes a database connection if one does not exist, and closes the database connection when the API completes.

### API include file:

*db2ApiDf.h*

### C API syntax:

```

/* File: db2ApiDf.h */
/* API: db2HADRStop */
/* ... */
SQL_API_RC SQL_API_FN
db2HADRStop (
    db2UInt32 versionNumber,
    void * pParmStruct,
    struct sqlca * pSqlca);

typedef SQL_STRUCTURE db2HADRStopStruct
{
    char                *piDbAlias;
    char                *piUserName;
    char                *piPassword;
} db2HADRStopStruct;

```

### Generic API syntax:

## db2HADRStop - Stop HADR

```
/* File: db2ApiDf.h */
/* API: db2gHADRStop */
/* ... */
SQL_API_RC SQL_API_FN
db2gHADRStop (
    db2UInt32 versionNumber,
    void * pParmStruct,
    struct sqlca * pSqlca);

typedef SQL_STRUCTURE db2gHADRStopStruct
{
    char                *piDbAlias;
    db2UInt32           iAliasLen;
    char                *piUserName;
    db2UInt32           iUserNameLen;
    char                *piPassword;
    db2UInt32           iPasswordLen;
} db2gHADRStopStruct;
```

### API parameters:

#### versionNumber

Input. Specifies the version and release level of the structure passed as the second parameter *pParmStruct*.

#### pParmStruct

Input. A pointer to the *db2HADRStopStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### piDbAlias

Input. A pointer to the database alias.

#### iAliasLen

Input. Specifies the length in bytes of the database alias.

#### piUserName

Input. A pointer to the user name under which the command will be executed.

#### iUserNameLen

Input. Specifies the length in bytes of the user name.

#### piPassword

Input. A pointer to a string containing the password.

#### iPasswordLen

Input. Specifies the length in bytes of the password.

### Related reference:

- “db2HADRStart - Start HADR” on page 193
- “db2HADRTakeover - Take Over as Primary Database” on page 206
- “STOP HADR” on page 196

---

## Switching database roles in high availability disaster recovery

During high availability disaster recovery (HADR), use the TAKEOVER HADR command to switch the roles of the primary and standby databases.

**Warning:** Before performing this procedure, ensure that the primary database is no longer processing database transactions.





FORCE option) could result in two primary databases. When there are two primary databases, each database will have different data, and the two databases can no longer be automatically synchronized.

- Deactivate the primary database or stop its instance, if possible. (This might not be possible if the primary system has hung, crashed, or is otherwise inaccessible.) After a takeover operation is performed, if the failed database is later restarted, it will not automatically assume the role of primary database.
- The likelihood and extent of transaction loss depends on your specific configuration and circumstances:
  - If the primary database fails while in peer state and the synchronization mode is synchronous (SYNC), the standby database will not lose transactions that were reported committed to an application before the primary database failed.
  - If the primary database fails while in peer state and the synchronization mode is near synchronous (NEARSYNC), the standby database can only lose transactions committed by the primary database if both the primary and the standby databases fail at the same time.
  - If the primary database fails while in peer state and the synchronization mode is asynchronous (ASYNC), the standby database can lose transactions committed by the primary database if the standby database did not receive all of the log records for the transactions before the takeover operation was performed. The standby database can also lose transactions committed by the primary database if both the primary and the standby databases fail at the same time.
  - If the primary database fails while in remote catchup pending state, transactions that have not been received and processed by the standby database will be lost.

**Note:** Any log gap shown in the database snapshot will represent the gap at the last time the primary and standby databases were communicating with each other; the primary database might have processed a very large number of transactions since that time.

**Restrictions:**

- The TAKEOVER HADR command can only be issued on the standby database.
- HADR does not interface with the DB2<sup>®</sup> fault monitor (db2fm) which can be used to automatically restart a failed database. If the fault monitor is enabled, you should be aware of possible fault monitor action on a presumably failed primary database.
- A takeover operation can only take place if the primary and standby databases are in peer state or the standby database is in remote catchup pending state. If the standby database is in any other state, an error will be returned.

**Note:** You can make a standby database that is in local catchup state available for normal use by converting it to a standard database. To do this, shut the database down by issuing the DEACTIVATE DATABASE command, and then issue the STOP HADR command. Once HADR has been stopped, you must complete a rollforward operation on the former standby database before it can be used. A database cannot rejoin an

HADR pair after it has been converted from a standby database to a standard database. To restart HADR on the two servers, follow the procedure for initializing HADR.

**Procedure:**

In a failover scenario, a takeover operation can be performed through the command line processor (CLP), the Manage High Availability Disaster Recovery window in the Control Center, or the **db2HADRTakeover** application programming interface (API).

The following procedure shows you how to initiate a failover on the primary or standby database using the CLP:

1. Completely disable the failed primary database. When a database encounters internal errors, normal shutdown commands may not completely shut it down. You may need to use operating system commands to remove resources such as processes, shared memory, or network connections.
2. Issue the TAKEOVER HADR command with the BY FORCE option on the standby database. In the following example the failover takes place on database LEAFS:

```
TAKEOVER HADR ON DB LEAFS BY FORCE
```

The BY FORCE option is required because the primary is expected to be offline. If the primary database is not completely disabled, the standby database will still have a connection to the primary and will send a message to the primary database asking it to shutdown. The standby database will still switch to the role of primary database whether or not it receives confirmation from that the primary database has been shutdown.

To open the Manage HADR window:

1. From the Control Center, expand the object tree until you find the database for which you want to manage HADR.
2. Right-click the database and click High Availability Disaster Recovery→Manage in the pop-up menu. The Manage High Availability Disaster Recovery window opens.

Detailed information is provided through the online help facility within the Control Center.

**Related concepts:**

- “High availability disaster recovery overview” on page 173
- “Synchronization modes for high availability disaster recovery” on page 181
- “Standby database states in high availability disaster recovery” on page 179

**Related tasks:**

- “Initializing high availability disaster recovery” on page 189
- “Switching database roles in high availability disaster recovery” on page 198

**Related reference:**

- “ROLLFORWARD DATABASE” on page 123
- “DEACTIVATE DATABASE Command” in the *Command Reference*
- “STOP HADR” on page 196

- “TAKEOVER HADR” on page 204

---

## Reintegrating a database after a takeover operation

If you executed a takeover operation in a high availability disaster recovery (HADR) environment because the primary database failed, you can bring the failed database back online and use it as a standby database or return it to its status as primary database.

To reintegrate the failed primary database into the HADR pair as the new standby database:

1. Repair the system where the original primary database resided. This could involve repairing failed hardware or rebooting the crashed operating system.
2. Restart the failed primary database as a standby database. In the following example, database LEAFS is started as a standby database:

```
START HADR ON DB LEAFS AS STANDBY
```

**Note:** This command will fail if the two copies of the database have incompatible log streams. In particular, HADR requires that the original primary database did not apply any logged operation that was never reflected on the original standby database before it took over as the new primary database. If this did occur, you can restart the original primary database as a standby database by restoring a backup image of the new primary database or by initializing a split mirror.

After the original primary database has rejoined the HADR pair as the standby database, you can choose to perform a failback operation to switch the roles of the databases to enable the original primary database to be once again the primary database. To perform this failback operation, issue the following command on the standby database:

```
TAKEOVER HADR ON DB LEAFS
```

**Notes:**

1. If the HADR databases are not in peer state or the pair is not connected, this command will fail.
2. Open sessions on the primary database are forced closed and inflight transactions are rolled back.
3. When switching the roles of the primary and standby databases, the BY FORCE option of the TAKEOVER HADR command cannot be specified.

**Related concepts:**

- “Synchronization modes for high availability disaster recovery” on page 181
- “Standby database states in high availability disaster recovery” on page 179

**Related tasks:**

- “Switching database roles in high availability disaster recovery” on page 198

**Related reference:**

- “TAKEOVER HADR” on page 204

---

## Performing a rolling upgrade in a high availability disaster recovery environment

Use this procedure in a high availability disaster recovery (HADR) environment when you upgrade software (operating system or DB2<sup>®</sup> UDB) or hardware, or when you make changes to database configuration parameters. This procedure keeps database service available throughout the upgrade process, with only a momentary service interruption when processing is switched from one database to the other. Because HADR performs optimally when both the primary and standby databases are on identical systems, you should apply changes to both systems as quickly as possible.

**Note:** All DB2 UDB FixPaks and upgrades should be implemented in a test environment before being applied to your production system.

### Prerequisites:

The HADR pair should be in peer state before starting the rolling upgrade.

### Restrictions:

This procedure will not work to upgrade from one major version of DB2 UDB to a later version.

### Procedure:

To perform a rolling upgrade in an HADR environment:

1. Upgrade the system where the standby database resides:
  - a. Use the DEACTIVATE DATABASE command to shut down the standby database.
  - b. If necessary, shut down the instance on the standby database.
  - c. Make changes to one or more of the following: the software, hardware, or the DB2 configuration parameters.
  - d. If necessary, restart the instance on the standby database.
  - e. Use the ACTIVATE DATABASE command to restart the standby database.
  - f. Ensure that the standby database enters peer state. Use the GET SNAPSHOT command to check this.
2. Switch the roles of the primary and standby databases:
  - a. Issue the TAKEOVER HADR command on the standby database.
  - b. Direct clients to the new primary database. This can be done using automatic client reroute.

**Note:** Because the standby database takes over as the primary database, the new primary database is now upgraded. If you are applying a DB2 UDB FixPak, the TAKEOVER HADR command will change the role of the original primary database to standby database. However, the command will not let the new standby database connect to the newly upgraded primary database. Because the new standby database uses an older version of DB2 UDB, it might not understand the new log records generated by the upgraded primary database, and it will be shut down. To connect to the primary database, the standby database must also be upgraded.

3. Upgrade original primary database (which is now the standby database) using the same procedure as in Step 1 above. When you have done this, both databases are upgraded and connected to each other in HADR peer state. The HADR system provides full database service and full high availability protection.
4. Optional. To return to your original configuration, switch the roles of the primary and standby database as in step 2.

**Related concepts:**

- “Automatic client reroute and high availability disaster recovery” on page 184
- “Standby database states in high availability disaster recovery” on page 179

**Related reference:**

- “GET SNAPSHOT Command” in the *Command Reference*
- “ACTIVATE DATABASE Command” in the *Command Reference*
- “DEACTIVATE DATABASE Command” in the *Command Reference*
- “TAKEOVER HADR” on page 204

## TAKEOVER HADR

Instructs an HADR standby database to take over as the new HADR primary database for the HADR pair.

**Authorization:**


One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

**Required connection:**

Instance. The command establishes a database connection if one does not exist, and closes the database connection when the command completes.

**Command syntax:**



**Command parameters:**

**DATABASE** *database-alias*  
 Identifies the current HADR standby database that should take over as the HADR primary database.

**USER** *user-name*  
 Identifies the user name under which the takeover operation is to be started.

**USING** *password*

The password used to authenticate *user-name*.

**BY FORCE**

Specifies that the database will not wait for confirmation that the original HADR primary database has been shut down. This option is required if the HADR pair is not in peer state.

**Usage notes:**

The following table shows the behavior of the TAKEOVER HADR command when issued on an active standby for each possible state and option combination. An error message is returned if this command is issued on an inactive standby database.

Standby state	BY FORCE option used	Takeover behavior
Local catchup or remote catchup	No	Error message returned
Local catchup or remote catchup	Yes	Error message returned
Peer	No	<p>Primary database and standby database switch roles.</p> <p>If no failure is encountered during takeover, there will be no data loss. However, if failures are encountered during takeover, data loss might occur and the roles of the primary and standby might or might not have been changed. The following is a guideline for handling failures during a takeover in which the primary and standby switch roles:</p> <ol style="list-style-type: none"> <li>1. If a failure occurs during a takeover operation, the roles of the HADR systems might or might not have been changed. If possible, make sure both databases are online. Check the HADR role of the available database or databases using the Snapshot Monitor, or by checking the value of the database configuration parameter <i>hadr_db_role</i>.</li> <li>2. If the intended new primary is still in standby role, and takeover is still desired, re-issue the TAKEOVER HADR command (see the next guideline regarding the BY FORCE option).</li> <li>3. It is possible to end up with both databases in standby role. In that case, the BY FORCE option can be issued at whichever node should now become the primary. The BY FORCE option is required in this case because the two standbys cannot establish the usual HADR primary-standby connection.</li> </ol>
Peer	Yes	The standby notifies the primary to shut itself (the primary) down. The standby stops receiving logs from the primary, finishes replaying the logs it has already received, and then becomes a primary. The standby does <i>not</i> wait for any acknowledgement from the primary to confirm that it has received the takeover notification or that it has shut down.

## TAKEOVER HADR

Standby state	BY FORCE option used	Takeover behavior
Remote catchup pending	No	Error message returned.
Remote catchup pending	Yes	The standby becomes a primary.

---

### db2HADRTakeover - Take Over as Primary Database

Instructs a standby database to take over as the primary database.

#### Authorization:

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

#### Required connection:

Instance. The API establishes a database connection if one does not exist, and closes the database connection when the API completes.

#### API include file:

*db2ApiDf.h*

#### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2HADRTakeover */
/* ... */
SQL_API_RC SQL_API_FN
db2HADRTakeover (
    db2UInt32 versionNumber,
    void * pParmStruct,
    struct sqlca * pSqlca);

typedef SQL_STRUCTURE db2HADRTakeoverStruct
{
    char                *piDbAlias;
    char                *piUserName;
    char                *piPassword;
    db2UInt16           iByForce;
} db2HADRTakeoverStruct;
```

#### Generic API syntax:

```
/* File: db2ApiDf.h */
/* API: db2gHADRTakeover */
/* ... */
SQL_API_RC SQL_API_FN
db2gHADRTakeover (
    db2UInt32 versionNumber,
    void * pParmStruct,
    struct sqlca * pSqlca);
```



## db2HADRTakeover - Take Over as Primary Database

```
typedef SQL_STRUCTURE db2gHADRTakeoverStruct
{
    char                *piDbAlias;
    db2UInt32           iAliasLen;
    char                *piUserName;
    db2UInt32           iUserNameLen;
    char                *piPassword;
    db2UInt32           iPasswordLen;
    db2UInt16           iByForce;
} db2gHADRTakeoverStruct;
```

### API parameters:

#### versionNumber

Input. Specifies the version and release level of the structure passed as the second parameter *pParmStruct*.

#### pParmStruct

Input. A pointer to the *db2HADRTakeoverStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### piDbAlias

Input. A pointer to the database alias.

#### iAliasLen

Input. Specifies the length in bytes of the database alias.

#### piUserName

Input. A pointer to the user name under which the command will be executed.

#### iUserNameLen

Input. Specifies the length in bytes of the user name.

#### piPassword

Input. A pointer to a string containing the password.

#### iPasswordLen

Input. Specifies the length in bytes of the password.

#### iByForce

Input. Valid values are:

##### DB2HADR\_NO\_FORCE

Specifies that a takeover occurs only if the two systems are in peer state with communication established; this results in a role reversal between the HADR primary and HADR standby databases.

##### DB2HADR\_FORCE

Specifies that the standby database takes over as the primary database without waiting for confirmation that the original primary database has been shut down.

### Related reference:

- “db2HADRTakeover - Start HADR” on page 193
- “db2HADRTakeover - Stop HADR” on page 197
- “TAKEOVER HADR” on page 204

## db2HADRTakeover - Take Over as Primary Database

---

## Chapter 8. Cluster support on AIX

---

### High Availability Cluster Multi-Processing support

Enhanced Scalability (ES) is a feature of High Availability Cluster Multi-Processing (HACMP) for AIX®. This feature provides the same failover recovery and has the same event structure as HACMP. Enhanced scalability also provides:

- Larger HACMP clusters.
- Additional error coverage through *user-defined events*. Monitored areas can trigger user-defined events, which can be as diverse as the death of a process, or the fact that paging space is nearing capacity. Such events include pre- and post-events that can be added to the failover recovery process, if needed. Extra functions that are specific to the different implementations can be placed within the HACMP pre- and post-event streams.

A *rules file* (`/usr/sbin/cluster/events/rules.hacmprd`) contains the HACMP events. User-defined events are added to this file. The script files that are to be run when events occur are part of this definition.

- HACMP client utilities for monitoring and detecting status changes (in one or more clusters) from AIX physical nodes outside of the HACMP cluster.

The nodes in HACMP ES clusters exchange messages called *heartbeats*, or *keepalive packets*, by which each node informs the other nodes about its availability. A node that has stopped responding causes the remaining nodes in the cluster to invoke recovery. The recovery process is called a *node\_down event* and may also be referred to as *failover*. The completion of the recovery process is followed by the re-integration of the node into the cluster. This is called a *node\_up event*.

There are two types of events: standard events that are anticipated within the operations of HACMP ES, and user-defined events that are associated with the monitoring of parameters in hardware and software components.

One of the standard events is the *node\_down* event. When planning what should be done as part of the recovery process, HACMP allows two failover options: hot (or idle) standby, and mutual takeover.

**Note:** When using HACMP, ensure that DB2® instances are not started at boot time by using the **db2iauto** utility as follows:

```
db2iauto -off InstName  
where
```

InstName is the login name of the instance.

#### Cluster Configuration

In a *hot standby* configuration, the AIX processor node that is the takeover node is *not* running any other workload. In a *mutual takeover* configuration, the AIX processor node that is the takeover node is running other workloads.

Generally, in a partitioned database environment, DB2 Universal Database runs in mutual takeover mode with partitions on each node. One exception is a scenario in which the catalog node is part of a hot standby configuration.

When planning a large DB2 installation on an RS/6000® SP™ using HACMP ES, you need to consider how to divide the nodes of the cluster within or between the RS/6000 SP frames. Having a node and its backup in different SP frames allows takeover in the event one frame goes down (that is, the frame power/switch board fails). However, such failures are expected to be exceedingly rare, because there are  $N+1$  power supplies in each SP frame, and each SP switch has redundant paths, along with  $N+1$  fans and power. In the case of a frame failure, manual intervention may be required to recover the remaining frames. This recovery procedure is documented in the SP Administration Guide. HACMP ES provides for recovery of SP node failures; recovery of frame failures is dependent on the proper layout of clusters within one or more SP frames.

Another planning consideration is how to manage big clusters. It is easier to manage a small cluster than a big one; however, it is also easier to manage one big cluster than many smaller ones. When planning, consider how your applications will be used in your cluster environment. If there is a single, large, homogeneous application running, for example, on 16 nodes, it is probably easier to manage the configuration as a single cluster rather than as eight two-node clusters. If the same 16 nodes contain many different applications with different networks, disks, and node relationships, it is probably better to group the nodes into smaller clusters. Keep in mind that nodes integrate into an HACMP cluster one at a time; it will be faster to start a configuration of multiple clusters rather than one large cluster. HACMP ES supports both single and multiple clusters, as long as a node and its backup are in the same cluster.

HACMP ES failover recovery allows pre-defined (also known as *cascading*) assignment of a resource group to a physical node. The failover recovery procedure also allows floating (or *rotating*) assignment of a resource group to a physical node. IP addresses, and external disk volume groups, or file systems, or NFS file systems, and application servers within each resource group specify either an application or an application component, which can be manipulated by HACMP ES between physical nodes by failover and reintegration. Failover and reintegration behavior is specified by the type of resource group created, and by the number of nodes placed in the resource group.

For example, consider a DB2 database partition (logical node). If its log and table space containers were placed on external disks, and other nodes were linked to those disks, it would be possible for those other nodes to access these disks and to restart the database partition (on a takeover node). It is this type of operation that is automated by HACMP. HACMP ES can also be used to recover NFS file systems used by DB2 instance main user directories.

Read the HACMP ES documentation thoroughly as part of your planning for recovery with DB2 UDB in a partitioned database environment. You should read the Concepts, Planning, Installation, and Administration guides, then build the recovery architecture for your environment. For each subsystem that you have identified for recovery, based on known points of failure, identify the HACMP clusters that you need, as well as the recovery nodes (either hot standby or mutual takeover).

It is strongly recommended that both disks and adapters be mirrored in your external disk configuration. For DB2 physical nodes that are configured for HACMP, care is required to ensure that nodes on the volume group can vary from the shared external disks. In a mutual takeover configuration, this arrangement requires some additional planning, so that the paired nodes can access each other's

volume groups without conflicts. In a partitioned database environment, this means that all container names must be unique across all databases.

One way to achieve uniqueness is to include the partition number as part of the name. You can specify a node expression for container string syntax when creating either SMS or DMS containers. When you specify the expression, the node number can be part of the container name or, if you specify additional arguments, the results of those arguments can be part of the container name. Use the argument " \$N" ([blank]\$N) to indicate the node expression. The argument must occur at the end of the container string, and can only be used in one of the following forms:

*Table 1. Arguments for Creating Containers.* The node number is assumed to be five.

Syntax	Example	Value
[blank]\$N	" \$N"	5
[blank]\$N+[number]	" \$N+1011"	1016
[blank]\$N%[number]	" \$N%3"	2
[blank]\$N+[number]%[number]	" \$N+12%13"	4
[blank]\$N%[number]+[number]	" \$N%3+20"	22
<b>Notes:</b>		
1. % is modulus.		
2. In all cases, the operators are evaluated from left to right.		

Following are some examples of how to create containers using this special argument:

- Creating containers for use on a two-node system.

```
CREATE TABLESPACE TS1 MANAGED BY DATABASE USING
(device '/dev/rcont $N' 20000)
```

The following containers would be used:

```
/dev/rcont0 - on Node 0
/dev/rcont1 - on Node 1
```

- Creating containers for use on a four-node system.

```
CREATE TABLESPACE TS2 MANAGED BY DATABASE USING
(file '/DB2/containers/TS2/container $N+100' 10000)
```

The following containers would be used:

```
/DB2/containers/TS2/container100 - on Node 0
/DB2/containers/TS2/container101 - on Node 1
/DB2/containers/TS2/container102 - on Node 2
/DB2/containers/TS2/container103 - on Node 3
```

- Creating containers for use on a two-node system.

```
CREATE TABLESPACE TS3 MANAGED BY SYSTEM USING
('/TS3/cont $N%2, '/TS3/cont $N%2+2')
```

The following containers would be used:

```
/TS3/cont0 - on Node 0
/TS3/cont2 - on Node 0
/TS3/cont1 - on Node 1
/TS3/cont3 - on Node 1
```

## Configuring DB2 Database Partitions for HACMP ES

Once configured, each database partition in an instance is started by HACMP ES, one physical node at a time. Multiple clusters are recommended for starting parallel DB2 configurations that are larger than four nodes. Note that in a 64-node parallel DB2 configuration, it is faster to start 32 two-node HACMP clusters in parallel, than four 16-node clusters.

A script file, `rc.db2pe`, is packaged with DB2 UDB Enterprise Server Edition (and installed on each node in `/usr/bin`) to assist in configuring for HACMP ES failover or recovery in either hot standby or mutual takeover nodes. In addition, DB2 buffer pool sizes can be customized during failover in mutual takeover configurations from within `rc.db2pe`. (Buffer pool sizes can be configured to ensure proper resource allocation when two database partitions run on one physical node.)

### HACMP ES Event Monitoring and User-defined Events

Initiating a failover operation if a process dies on a given node, is an example of a user-defined event. Examples that illustrate user-defined events, such as shutting down a database partition and forcing a transaction abort to free paging space, can be found in the `samples/hacmp/es` subdirectory.

A rules file, `/user/sbin/cluster/events/rules.hacmprd`, contains HACMP events. Each event description in this file has the following nine components:

- Event name, which must be unique.
- State, or qualifier for the event. The event name and state are the rule triggers. HACMP ES Cluster Manager initiates recovery only if it finds a rule with a trigger corresponding to the event name and state.
- Resource program path, a full-path specification of the `xxx.rp` file containing the recovery program.
- Recovery type. This is reserved for future use.
- Recovery level. This is reserved for future use.
- Resource variable name, which is used for Event Manager events.
- Instance vector, which is used for Event Manager events. This is a set of elements of the form "name=value". The values uniquely identify the copy of the resource in the system and, by extension, the copy of the resource variable.
- Predicate, which is used for Event Manager events. This is a relational expression between a resource variable and other elements. When this expression is true, the Event Management subsystem generates an event to notify the Cluster Manager and the appropriate application.
- Rearm predicate, which is used for Event Manager events. This is a predicate used to generate an event that alters the status of the primary predicate. This predicate is typically the inverse of the primary predicate. It can also be used with the event predicate to establish an upper and a lower boundary for a condition of interest.

Each object requires one line in the event definition, even if the line is not used. If these lines are removed, HACMP ES Cluster Manager cannot parse the event definition properly, and this may cause the system to hang. Any line beginning with `"#"` is treated as a comment line.

**Note:** The rules file requires exactly nine lines for each event definition, not counting any comment lines. When adding a user-defined event at the bottom of the rules file, it is important to remove the unnecessary empty line at the end of the file, or the node will hang.

HACMP ES uses PSSP event detection to treat user-defined events. The PSSP Event Management subsystem provides comprehensive event detection by monitoring various hardware and software resources.

The process can be summarized as follows:

1. Either Group Services/ES (for predefined events) or Event Management (for user-defined events) notifies HACMP ES Cluster Manager of the event.
2. Cluster Manager reads the `rules.hacmprd` file, and determines the recovery program that is mapped to the event.
3. Cluster Manager runs the recovery program, which consists of a sequence of recovery commands.
4. The recovery program executes the recovery commands, which may be shell scripts or binary commands. (In HACMP for AIX, the recovery commands are the same as the HACMP event scripts.)
5. Cluster Manager receives the return status from the recovery commands. An unexpected status "hangs" the cluster until manual intervention (using `smitty cm_rec_aids` or the `/usr/sbin/cluster/utilities/clruncmd` command) is carried out.

For detailed information on the implementation and design of highly available IBM® DB2 Universal Database™ environments on AIX see the following white papers which are available from the "DB2 UDB and DB2 Connect™ Support" web site (<http://www.ibm.com/software/data/pubs/papers/>):

- "IBM DB2 Universal Database Enterprise Edition for AIX and HACMP/ES"
- "IBM DB2 Universal Database Enterprise - Extended Edition for AIX and HACMP/ES"

**Related reference:**

- "db2start - Start DB2 Command" in the *Command Reference*





---

## Chapter 9. Cluster support on the Windows operating system

---

### Microsoft Cluster Server support

#### Introduction

Microsoft Cluster Server (MSCS) is a feature of Windows<sup>®</sup> NT Server, Windows 2000 Server, and Windows Server 2003 operating systems. It is the software that supports the connection of two servers (up to four servers in DataCenter Server) into a cluster for high availability and easier management of data and applications. MSCS can also automatically detect and recover from server or application failures. It can be used to move server workloads to balance machine utilization and to provide for planned maintenance without downtime.

The following DB2<sup>®</sup> products have support for MSCS:

- DB2 Universal Database<sup>™</sup> Workgroup Server Edition
- DB2 Universal Database Enterprise Server Edition (DB2 ESE)
- DB2 Universal Database Connect Enterprise Edition (DB2 CEE)

#### DB2 MSCS Components

A cluster is a configuration of two or more nodes, each of which is an independent computer system. The cluster appears to network clients as a single server.




Figure 17. Example MSCS Configuration

The nodes in an MSCS cluster are connected using one or more shared storage buses and one or more physically independent networks. The network that connects only the servers but does not connect the clients to the cluster is referred to as a *private* network. The network that supports client connections is referred to as the *public* network. There are one or more local disks on each node. Each shared storage bus attaches to one or more disks. Each disk on the shared bus is owned by only one node of the cluster at a time. The DB2 software resides on the local

disk. DB2 database files (tables, indexes, log files, etc.) reside on the shared disks. Because MSCS does not support the use of raw partitions in a cluster, it is not possible to configure DB2 to use raw devices in an MSCS environment.

### **The DB2 Resource**

In an MSCS environment, a resource is an entity that is managed by the clustering software. For example, a disk, an IP address, or a generic service can be managed as a resource. DB2 integrates with MSCS by creating its own resource type called DB2. Each DB2 resource manages a DB2 instance, and when running in a partitioned database environment, each DB2 resource manages a database partition. The name of the DB2 resource is the instance name, although in the case of a partitioned database environment, the name of the DB2 resource consists of both the instance name and the partition (or node) number.

### **Pre-online and Post-online Script**

You can run scripts both before and after a DB2 resource is brought online. These scripts are referred to as pre-online and post-online scripts respectively. Pre-online and post-online scripts are .BAT files that can run DB2 and system commands.

In a situation when multiple instances of DB2 may be running on the same machine, you can use the pre-online and post-online scripts to adjust the configuration so that both instances can be started successfully. In the event of a failover, you can use the post-online script to perform manual database recovery. Post-online script can also be used to start any applications or services that depend on DB2.

### **The DB2 Group**

Related or dependent resources are organized into resource groups. All resources in a group move between cluster nodes as a unit. For example, in a typical DB2 single partition cluster environment, there will be a DB2 group that contains the following resources:

1. DB2 resource. The DB2 resource manages the DB2 instance (or node).
2. IP Address resource. The IP Address resource allows client applications to connect to the DB2 server.
3. Network Name resource. The Network Name resource allows client applications to connect to the DB2 server by using a name rather than using an IP address. The Network Name resource has a dependency on the IP Address resource. The Network Name resource is optional. (Configuring a Network Name resource may affect the failover performance.)
4. One or more Physical Disk resources. Each Physical Disk resource manages a shared disk in the cluster.

**Note:** The DB2 resource is configured to depend on all other resources in the same group so the DB2 server can only be started after all other resources are online.

### **Failover Configurations**

Two types of configuration are available:

- Hot standby
- Mutual takeover

In a partitioned database environment, the clusters do not all have to have the same type of configuration. You can have some clusters that are set up to use hot standby, and others that are set up for mutual takeover. For example, if your DB2 instance consists of five workstations, you can have two machines set up to use a mutual takeover configuration, two to use a hot standby configuration, and one machine not configured for failover support.

### Hot Standby Configuration

In a hot standby configuration, one machine in the MSCS cluster provides dedicated failover support, and the other machine participates in the database system. If the machine participating in the database system fails, the database server on it will be started on the failover machine. If, in a partitioned database system, you are running multiple logical nodes on a machine and it fails, the logical nodes will be started on the failover machine. Figure 18 shows an example of a hot standby configuration.




Figure 18. Hot Standby Configuration

### Mutual Takeover Configuration

In a mutual takeover configuration, both workstations participate in the database system (that is, each machine has at least one database server running on it). If one of the workstations in the MSCS cluster fails, the database server on the failing machine will be started to run on the other machine. In a mutual takeover configuration, a database server on one machine can fail independently of the database server on another machine. Any database server can be active on any machine at any given point in time. Figure 19 on page 218 shows an example of a mutual takeover configuration.




Figure 19. Mutual Takeover Configuration

For detailed information on the implementation and design of highly available IBM® DB2 Universal Database environments on the Windows Operating System see the following white papers which are available from the "DB2 UDB and DB2 Connect™ Support" web site

(<http://www.ibm.com/software/data/pubs/papers/>):

- "Implementing IBM DB2 Universal Database Enterprise - Extended Edition with Microsoft® Cluster Server"
- "Implementing IBM DB2 Universal Database Enterprise Edition with Microsoft Cluster Server"
- "DB2 Universal Database for Windows: High Availability Support Using Microsoft Cluster Server - Overview"

---

## Chapter 10. Cluster support for the Solaris Operating Environment

---

### Cluster support for the Solaris Operating Environment

High availability in the Solaris™ Operating Environment can be achieved through DB2® working with Sun™ Cluster, or Veritas Cluster Server (VCS). For information about Sun Cluster, see the white paper entitled “DB2 Universal Database and High Availability on Sun Cluster 3.X”, which is available from the “DB2 UDB and DB2 Connect Online Support” web site (<http://www.ibm.com/software/data/pubs/papers/>). For information about VERITAS Cluster Server, see the white paper entitled “DB2 and High Availability on VERITAS Cluster Server”, which is available from the “IBM® Support and downloads” Web site (<http://www.ibm.com/support/docview.wss?uid=swg21045033>).

**Note:** When using Sun Cluster 3.0 or Veritas Cluster Server, ensure that DB2 instances are not started at boot time by using the **db2iauto** utility as follows:

```
db2iauto -off InstName
where
```

InstName is the login name of the instance.

#### High Availability

The computer systems that host data services contain many distinct components, and each component has a “mean time before failure” (MTBF) associated with it. The MTBF is the average time that a component will remain usable. The MTBF for a quality hard drive is in the order of one million hours (approximately 114 years). While this seems like a long time, one out of 200 disks is likely to fail within a 6-month period.

Although there are a number of methods to increase availability for a data service, the most common is an HA cluster. A cluster, when used for high availability, consists of two or more machines, a set of private network interfaces, one or more public network interfaces, and some shared disks. This special configuration allows a data service to be moved from one machine to another. By moving the data service to another machine in the cluster, it should be able to continue providing access to its data. Moving a data service from one machine to another is called a *failover*, as illustrated in Figure 20 on page 220.

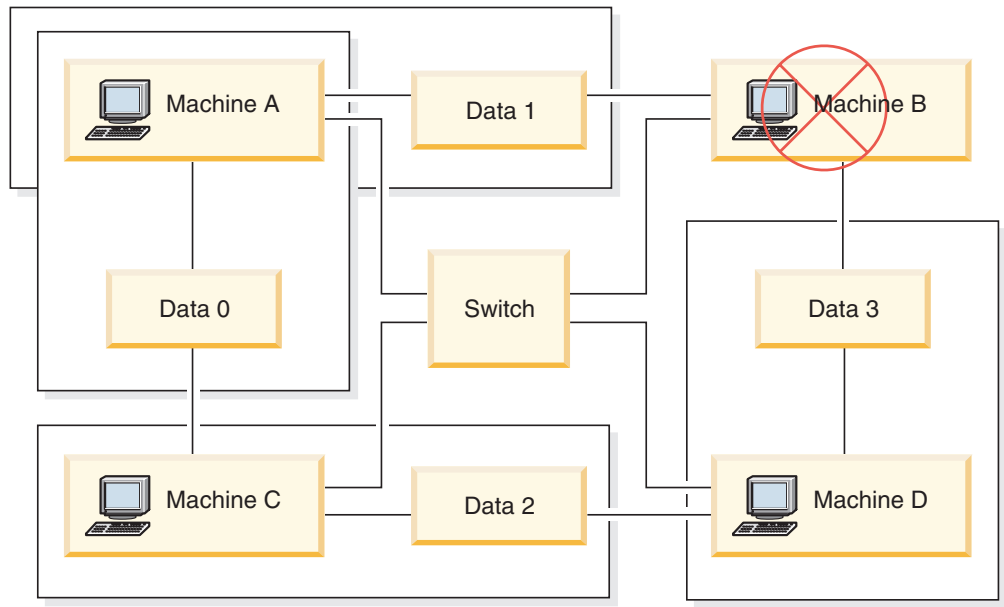


Figure 20. Failover. When Machine B fails its data service is moved to another machine in the cluster so that the data can still be accessed.

The private network interfaces are used to send *heartbeat* messages, as well as control messages, among the machines in the cluster. The public network interfaces are used to communicate directly with clients of the HA cluster. The disks in an HA cluster are connected to two or more machines in the cluster, so that if one machine fails, another machine has access to them.

A data service running on an HA cluster has one or more logical public network interfaces and a set of disks associated with it. The clients of an HA data service connect via TCP/IP to the logical network interfaces of the data service only. If a failover occurs, the data service, along with its logical network interfaces and set of disks, are moved to another machine.

One of the benefits of an HA cluster is that a data service can recover without the aid of support staff, and it can do so at any time. Another benefit is redundancy. All of the parts in the cluster should be redundant, including the machines themselves. The cluster should be able to survive any single point of failure.

Even though highly available data services can be very different in nature, they have some common requirements. Clients of a highly available data service expect the network address and host name of the data service to remain the same, and expect to be able to make requests in the same way, regardless of which machine the data service is on.

Consider a web browser that is accessing a highly available web server. The request is issued with a URL (Uniform Resource Locator), which contains both a host name, and the path to a file on the web server. The browser expects both the host name and the path to remain the same after failover of the web server. If the browser is downloading a file from the web server, and the server is failed over, the browser will need to reissue the request.

Availability of a data service is measured by the amount of time the data service is available to its users. The most common unit of measurement for availability is the percentage of "up time"; this is often referred to as the number of "nines":

99.99% => service is down for (at most) 52.6 minutes / yr  
99.999% => service is down for (at most) 5.26 minutes / yr  
99.9999% => service is down for (at most) 31.5 seconds / yr

When designing and testing an HA cluster:

1. Ensure that the administrator of the cluster is familiar with the system and what should happen when a failover occurs.
2. Ensure that each part of the cluster is truly redundant and can be replaced quickly if it fails.
3. Force a test system to fail in a controlled environment, and make sure that it fails over correctly each time.
4. Keep track of the reasons for each failover. Although this should not happen often, it is important to address any issues that make the cluster unstable. For example, if one piece of the cluster caused a failover five times in one month, find out why and fix it.
5. Ensure that the support staff for the cluster is notified when a failover occurs.
6. Do not overload the cluster. Ensure that the remaining systems can still handle the workload at an acceptable level after a failover.
7. Check failure-prone components (such as disks) often, so that they can be replaced before problems occur.

### **Fault Tolerance**

Another way to increase the availability of a data service is fault tolerance. A *fault tolerant* machine has all of its redundancy built in, and should be able to withstand a single failure of any part, including CPU and memory. Fault tolerant machines are most often used in niche markets, and are usually expensive to implement. An HA cluster with machines in different geographical locations has the added advantage of being able to recover from a disaster affecting only a subset of those locations.

An HA cluster is the most common solution to increase availability because it is scalable, easy to use, and relatively inexpensive to implement.

### **Related concepts:**

- “Sun Cluster 3.0 support” on page 221
- “VERITAS Cluster Server support” on page 224

---

## **Sun Cluster 3.0 support**

This section provides an overview of how DB2<sup>®</sup> works with Sun<sup>™</sup> Cluster 3.0 to achieve high availability, and includes a description of the high availability agent, which acts as a mediator between the two software products (see Figure 21 on page 222).



Figure 21. DB2, Sun Cluster 3.0, and High Availability. The relationship between DB2, Sun Cluster 3.0 and the high availability agent.

### Failover

Sun Cluster 3.0 provides high availability by enabling application failover. Each node is periodically monitored and the cluster software automatically relocates a cluster-aware application from a failed primary node to a designated secondary node. When a failover occurs, clients may experience a brief interruption in service and may have to reconnect to the server. However, they will not be aware of the physical server from which they are accessing the application and the data. By allowing other nodes in a cluster to automatically host workloads when the primary node fails, Sun Cluster 3.0 can significantly reduce downtime and increase productivity.

### Multihost Disks

Sun Cluster 3.0 requires multihost disk storage. This means that disks can be connected to more than one node at a time. In the Sun Cluster 3.0 environment, multihost storage allows disk devices to become highly available. Disk devices that reside on multihost storage can tolerate single node failures since there is still a physical path to the data through the alternate server node. Multihost disks can be accessed globally through a primary node. If client requests are accessing the data through one node and that node fails, the requests are switched over to another node that has a direct connection to the same disks. A volume manager provides for mirrored or RAID 5 configurations for data redundancy of the multihost disks. Currently, Sun Cluster 3.0 supports Solstice DiskSuite and VERITAS Volume Manager as volume managers. Combining multihost disks with disk mirroring and striping protects against both node failure and individual disk failure.

### Global Devices

Global devices are used to provide cluster-wide, highly available access to any device in a cluster, from any node, regardless of the device's physical location. All disks are included in the global namespace with an assigned device ID (DID) and are configured as global devices. Therefore, the disks themselves are visible from all cluster nodes.

### File systems/Global File Systems

A cluster or global file system is a proxy between the kernel (on one node) and the underlying file system volume manager (on a node that has a physical connection to one or more disks). Cluster file systems are dependent on global devices with physical connections to one or more nodes. They are independent of the underlying file system and volume manager. Currently, cluster file systems can be



built on UFS using either Solstice DiskSuite or VERITAS Volume Manager. The data only becomes available to all nodes if the file systems on the disks are mounted globally as a cluster file system.

### **Device Group**

All multihost disks must be controlled by the Sun Cluster framework. Disk groups, managed by either Solstice DiskSuite or VERITAS Volume Manager, are first created on the multihost disk. Then, they are registered as Sun Cluster disk device groups. A disk device group is a type of global device. Multihost device groups are highly available. Disks are accessible through an alternate path if the node currently mastering the device group fails. The failure of the node mastering the device group does not affect access to the device group except for the time required to perform the recovery and consistency checks. During this time, all requests are blocked (transparently to the application) until the system makes the device group available.

### **Resource Group Manager (RGM)**

The RGM, provides the mechanism for high availability and runs as a daemon on each cluster node. It automatically starts and stops resources on selected nodes according to pre-configured policies. The RGM allows a resource to be highly available in the event of a node failure or to reboot by stopping the resource on the affected node and starting it on another. The RGM also automatically starts and stops resource-specific monitors that can detect resource failures and relocate failing resources onto another node.

### **Data Services**

The term data service is used to describe a third-party application that has been configured to run on a cluster rather than on a single server. A data service includes the application software and Sun Cluster 3.0 software that starts, stops and monitors the application. Sun Cluster 3.0 supplies data service methods that are used to control and monitor the application within the cluster. These methods run under the control of the Resource Group Manager (RGM), which uses them to start, stop, and monitor the application on the cluster nodes. These methods, along with the cluster framework software and multihost disks, enable applications to become highly available data services. As highly available data services, they can prevent significant application interruptions after any single failure within the cluster, regardless of whether the failure is on a node, on an interface component or in the application itself. The RGM also manages resources in the cluster, including network resources (logical host names and shared addresses) and application instances.

### **Resource Type, Resource and Resource Group**

A resource type is made up of the following:

1. A software application to be run on the cluster.
2. Control programs used as callback methods by the RGM to manage the application as a cluster resource.
3. A set of properties that form part of the static configuration of a cluster.

The RGM uses resource type properties to manage resources of a particular type.

A resource inherits the properties and values of its resource type. It is an instance of the underlying application running on the cluster. Each instance requires a

unique name within the cluster. Each resource must be configured in a resource group. The RGM brings all resources in a group online and offline together on the same node. When the RGM brings a resource group online or offline, it invokes callback methods on the individual resources in the group.

The nodes on which a resource group is currently online are called its primary nodes, or its primaries. A resource group is mastered by each of its primaries. Each resource group has an associated Nodelist property, set by the cluster administrator, to identify all potential primaries or masters of the resource group.

For detailed information on the implementation and design of highly available IBM® DB2 Universal Database environments on the Sun Cluster 3.0 platform see the white paper entitled "DB2 and High Availability on Sun Cluster 3.0" which is available from the "DB2 UDB and DB2 Connect™ Support" web site (<http://www.ibm.com/software/data/pubs/papers/>).

**Related concepts:**

- "Cluster support for the Solaris Operating Environment" on page 219
- "VERITAS Cluster Server support" on page 224

---

## VERITAS Cluster Server support

VERITAS Cluster Server can be used to eliminate both planned and unplanned downtime. It can facilitate server consolidation and effectively manage a wide range of applications in heterogeneous environments. VERITAS Cluster Server supports up to 32 node clusters in both storage area network (SAN) and traditional client/server environments, VERITAS Cluster Server can protect everything from a single critical database instance, to very large multi-application clusters in networked storage environments. This section provides a brief summary of the features of VERITAS Cluster Server.

### Hardware Requirements

Following is a list of hardware currently supported by VERITAS Cluster Server:

- For server nodes:
  - Any SPARC/Solaris server from Sun™ Microsystems running Solaris™ 2.6 or later with a minimum of 128MB RAM.
- For disk storage:
  - EMC Symmetrix, IBM® Enterprise Storage Server®, HDS 7700 and 9xxx, Sun T3, Sun A5000, Sun A1000, Sun D1000 and any other disk storage supported by VCS 2.0 or later; your VERITAS representative can confirm which disk subsystems are supported or you can refer to VCS documentation.
  - Typical environments will require mirrored private disks (in each cluster node) for the DB2® UDB binaries and shared disks between nodes for the DB2 UDB data.
- For network interconnects:
  - For the public network connections, any network connection supporting IP-based addressing.
  - For the heartbeat connections (internal to the cluster), redundant heartbeat connections are required; this requirement can be met through the use of two additional Ethernet controllers per server or one additional Ethernet controller per server and the use of one shared GABdisk per cluster

## Software Requirements

The following VERITAS software components are qualified configurations:

- VERITAS Volume Manager 3.2 or later, VERITAS File System 3.4 or later, VERITAS Cluster Server 2.0 or later.
- DB Edition for DB2 for Solaris 1.0 or later.

While VERITAS Cluster Server does not require a volume manager, the use of VERITAS Volume Manager is strongly recommended for ease of installation, configuration and management.

## Failover

VERITAS Cluster Server is an availability clustering solution that manages the availability of application services, such as DB2 UDB, by enabling application failover. The state of each individual cluster node and its associated software services are regularly monitored. When a failure occurs that disrupts the application service (in this case, the DB2 UDB service), VERITAS Cluster Server and/or the VCS HA-DB2 Agent detect the failure and automatically take steps to restore the service. This can include restarting DB2 UDB on the same node or moving DB2 UDB to another node in the cluster and restarting it on that node. If an application needs to be migrated to a new node, VERITAS Cluster Server moves everything associated with the application (i.e., network IP addresses, ownership of underlying storage) to the new node so that users will not be aware that the service is actually running on another node. They will still access the service using the same IP addresses, but those addresses will now point to a different cluster node.

When a failover occurs with VERITAS Cluster Server, users may or may not see a disruption in service. This will be based on the type of connection (stateful or stateless) that the client has with the application service. In application environments with stateful connections (like DB2 UDB), users may see a brief interruption in service and may need to reconnect after the failover has completed. In application environments with stateless connections (like NFS), users may see a brief delay in service but generally will not see a disruption and will not need to log back on.

By supporting an application as a service that can be automatically migrated between cluster nodes, VERITAS Cluster Server can not only reduce unplanned downtime, but can also shorten the duration of outages associated with planned downtime (i.e., for maintenance and upgrades). Failovers can also be initiated manually. If a hardware or operating system upgrade must be performed on a particular node, DB2 UDB can be migrated to another node in the cluster, the upgrade can be performed, and then DB2 UDB can be migrated back to the original node.

Applications recommended for use in these types of clustering environments should be crash tolerant. A crash tolerant application can recover from an unexpected crash while still maintaining the integrity of committed data. Crash tolerant applications are sometimes referred to as *cluster friendly* applications. DB2 UDB is a crash tolerant application.

For information on how to decrease the amount of time it takes to perform a failover using a VERITAS CFS, CVM, and VCS solution, see the white paper entitled “DB2 UDB Version 8 and VERITAS Database Edition: Accelerating Failover

Times in DB2 UDB Database Environments”, which is available from the “DB2 UDB and DB2 Connect™ Online Support” web site (<http://www.ibm.com/software/data/pubs/papers/>).

### **Shared Storage**

When used with the VCS HA-DB2 Agent, Veritas Cluster Server requires shared storage. Shared storage is storage that has a physical connection to multiple nodes in the cluster. Disk devices resident on shared storage can tolerate node failures since a physical path to the disk devices still exists through one or more alternate cluster nodes.

Through the control of VERITAS Cluster Server, cluster nodes can access shared storage through a logical construct called “disk groups”. Disk groups represent a collection of logically defined storage devices whose ownership can be atomically migrated between nodes in a cluster. A disk group can only be imported to a single node at any given time. For example, if Disk Group A is imported to Node 1 and Node 1 fails, Disk Group A can be exported from the failed node and imported to a new node in the cluster. VERITAS Cluster Server can simultaneously control multiple disk groups within a single cluster.

In addition to allowing disk group definition, a volume manager can provide for redundant data configurations, using mirroring or RAID 5, on shared storage. VERITAS Cluster Server supports VERITAS Volume Manager and Solstice DiskSuite as logical volume managers. Combining shared storage with disk mirroring and striping can protect against both node failure and individual disk or controller failure.

### **VERITAS Cluster Server Global Atomic Broadcast(GAB) and Low Latency Transport (LLT)**

An internode communication mechanism is required in cluster configurations so that nodes can exchange information concerning hardware and software status, keep track of cluster membership, and keep this information synchronized across all cluster nodes. The Global Atomic Broadcast (GAB) facility, running across a low latency transport (LLT), provides the high speed, low latency mechanism used by VERITAS Cluster Server to do this. GAB is loaded as a kernel module on each cluster node and provides an atomic broadcast mechanism that ensures that all nodes get status update information at the same time.

By leveraging kernel-to-kernel communication capabilities, LLT provides high speed, low latency transport for all information that needs to be exchanged and synchronized between cluster nodes. GAB runs on top of LLT. VERITAS Cluster Server does not use IP as a heartbeat mechanism, but offers two other more reliable options. GAB with LLT, can be configured to act as a heartbeat mechanism, or a GABdisk can be configured as a disk-based heartbeat. The heartbeat must run over redundant connections. These connections can either be two private Ethernet connections between cluster nodes, or one private Ethernet connection and one GABdisk connection. The use of two GABdisks is not a supported configuration since the exchange of cluster status between nodes requires a private Ethernet connection.

For more information about GAB or LLT, or how to configure them in VERITAS Cluster Server configurations, please consult the VERITAS Cluster Server 2.0 User’s Guide for Solaris.

## Bundled and Enterprise Agents

An agent is a program that is designed to manage the availability of a particular resource or application. When an agent is started, it obtains the necessary configuration information from VCS and then periodically monitors the resource or application and updates VCS with the status. In general, agents are used to bring resources online, take resources offline, or monitor resources and provide four types of services: start, stop, monitor and clean. Start and stop are used to bring resources online or offline, monitor is used to test a particular resource or application for its status, and clean is used in the recovery process.

A variety of bundled agents are included as part of VERITAS Cluster Server and are installed when VERITAS Cluster Server is installed. The bundled agents are VCS processes that manage predefined resource types commonly found in cluster configurations (i.e., IP, mount, process and share), and they help to simplify cluster installation and configuration considerably. There are over 20 bundled agents with VERITAS Cluster Server.

Enterprise agents tend to focus on specific applications such as DB2 UDB. The VCS HA-DB2 Agent can be considered an Enterprise Agent, and it interfaces with VCS through the VCS Agent framework.

## VCS Resources, Resource Types and Resource Groups

A resource type is an object definition used to define resources within a VCS cluster that will be monitored. A resource type includes the resource type name and a set of properties associated with that resource that are salient from a high availability point of view. A resource inherits the properties and values of its resource type, and resource names must be unique on a cluster-wide basis.

There are two types of resources: persistent and standard (non-persistent). Persistent resources are resources such as network interface controllers (NICs) that are monitored but are not brought online or taken offline by VCS. Standard resources are those whose online and offline status is controlled by VCS.

The lowest level object that is monitored is a resource, and there are various resource types (i.e., share, mount). Each resource must be configured into a resource group, and VCS will bring all resources in a particular resource group online and offline together. To bring a resource group online or offline, VCS will invoke the start or stop methods for each of the resources in the group. There are two types of resource groups: failover and parallel. A highly available DB2 UDB configuration, regardless of whether it is partitioned or not, will use failover resource groups.

A "primary" or "master" node is a node that can potentially host a resource. A resource group attribute called `systemlist` is used to specify which nodes within a cluster can be primaries for a particular resource group. In a two node cluster, usually both nodes are included in the `systemlist`, but in larger, multi-node clusters that may be hosting several highly available applications there may be a requirement to ensure that certain application services (defined by their resources at the lowest level) can never fail over to certain nodes.

Dependencies can be defined between resource groups, and VERITAS Cluster Server depends on this resource group dependency hierarchy in assessing the impact of various resource failures and in managing recovery. For example, if the resource group `ClientApp1` can not be brought online unless the resource group

DB2 has already been successfully started, resource group ClientApp1 is considered dependent on resource group DB2.

For detailed information on the implementation and design of highly available IBM DB2 Universal Database environments with the VERITAS Cluster Server see the technote entitled "DB2 UDB and High Availability with VERITAS Cluster Server" which you can view by going to the following web site: <http://www.ibm.com/support>, and searching for the keyword "1045033".

**Related concepts:**

- "Cluster support for the Solaris Operating Environment" on page 219
- "Sun Cluster 3.0 support" on page 221

---

## Part 3. Appendixes





---

## Appendix A. How to read the syntax diagrams

A syntax diagram shows how a command should be specified so that the operating system can correctly interpret what is typed.

Read a syntax diagram from left to right, and from top to bottom, following the horizontal line (the main path). If the line ends with an arrowhead, the command syntax is continued, and the next line starts with an arrowhead. A vertical bar marks the end of the command syntax.

When typing information from a syntax diagram, be sure to include punctuation, such as quotation marks and equal signs.

Parameters are classified as keywords or variables:

- Keywords represent constants, and are shown in uppercase letters; at the command prompt, however, keywords can be entered in upper, lower, or mixed case. A command name is an example of a keyword.
- Variables represent names or values that are supplied by the user, and are shown in lowercase letters; at the command prompt, however, variables can be entered in upper, lower, or mixed case, unless case restrictions are explicitly stated. A file name is an example of a variable.

A parameter can be a combination of a keyword and a variable.

Required parameters are displayed on the main path:


►►—COMMAND—*required parameter*—►►

Optional parameters are displayed below the main path:


►►—COMMAND—  
└—*optional parameter*—┘—►►

## How to read the syntax diagrams

A parameter's default value is displayed above the path:



A stack of parameters, with the first parameter displayed on the main path, indicates that one of the parameters must be selected:



A stack of parameters, with the first parameter displayed below the main path, indicates that one of the parameters can be selected:



An arrow returning to the left, above the path, indicates that items can be repeated in accordance with the following conventions:

- If the arrow is uninterrupted, the item can be repeated in a list with the items separated by blank spaces:




- If the arrow contains a comma, the item can be repeated in a list with the items separated by commas:




Items from parameter stacks can be repeated in accordance with the stack conventions for required and optional parameters discussed previously.

Some syntax diagrams contain parameter stacks within other parameter stacks. Items from stacks can only be repeated in accordance with the conventions discussed previously. That is, if an inner stack does not have a repeat arrow above it, but an outer stack does, only one parameter from the inner stack can be chosen and combined with any parameter from the outer stack, and that combination can be repeated. For example, the following diagram shows that one could combine parameter *choice2a* with parameter *choice2*, and then repeat that combination again (*choice2* plus *choice2a*):




Some commands are preceded by an optional path parameter:



If this parameter is not supplied, the system searches the current directory for the command. If it cannot find the command, the system continues searching for the command in all the directories on the paths listed in the `.profile`.

Some commands have syntactical variants that are functionally equivalent:



## How to read the syntax diagrams

---

## Appendix B. Warning, error and completion messages

Messages generated by the various utilities are included among the SQL messages. These messages are generated by the database manager when a warning or error condition has been detected. Each message has a message identifier that consists of a prefix (SQL) and a four- or five-digit message number. There are three message types: notification, warning, and critical. Message identifiers ending with an N are error messages. Those ending with a W indicate warning or informational messages. Message identifiers ending with a C indicate critical system errors.

The message number is also referred to as the *SQLCODE*. The *SQLCODE* is passed to the application as a positive or negative number, depending on its message type (N, W, or C). N and C yield negative values, whereas W yields a positive value. DB2 returns the *SQLCODE* to the application, and the application can get the message associated with the *SQLCODE*. DB2 also returns an *SQLSTATE* value for conditions that could be the result of an SQL statement. Some *SQLCODE* values have associated *SQLSTATE* values.

You can use the information contained in this book to identify an error or problem, and to resolve the problem by using the appropriate recovery action. This information can also be used to understand where messages are generated and logged.

SQL messages, and the message text associated with *SQLSTATE* values, are also accessible from the operating system command line. To access help for these error messages, enter the following at the operating system command prompt:

```
db2 ? SQLnnnn
```

where *nnnnn* represents the message number. On UNIX based systems, the use of double quotation mark delimiters is recommended; this will avoid problems if there are single character file names in the directory:

```
db2 "? SQLnnnn"
```

The message identifier accepted as a parameter for the **db2** command is not case sensitive, and the terminating letter is not required. Therefore, the following commands will produce the same result:

```
db2 ? SQL0000N
db2 ? sql0000
db2 ? SQL0000n
```

If the message text is too long for your screen, use the following command (on UNIX based operating systems and others that support the "more" pipe):

```
db2 ? SQLnnnn | more
```

You can also redirect the output to a file which can then be browsed.

Help can also be invoked from interactive input mode. To access this mode, enter the following at the operating system command prompt:

```
db2
```

To get DB2 message help in this mode, type the following at the command prompt (db2 =>):

? SQLnnnnn

The message text associated with SQLSTATEs can be retrieved by issuing:

db2 ? nnnnn

or

db2 ? nn

where *nnnnn* is a five-character SQLSTATE value (alphanumeric), and *nn* is a two-digit SQLSTATE class code (the first two digits of the SQLSTATE value).

# Appendix C. Additional DB2 commands

This appendix describes recovery-related system and CLP commands that are not discussed in detail in this manual.

## System commands

### db2adutl - Managing DB2 objects within TSM

Allows users to query, extract, verify, and delete backup images, logs, and load copy images saved using Tivoli Storage Manager. Also allows users to grant and revoke access to objects on a TSM server.

On UNIX-based operating systems, this utility is located in the `sql1lib/adsm` directory. On Windows it is located in `sql1lib\bin`.

**Authorization:**

None

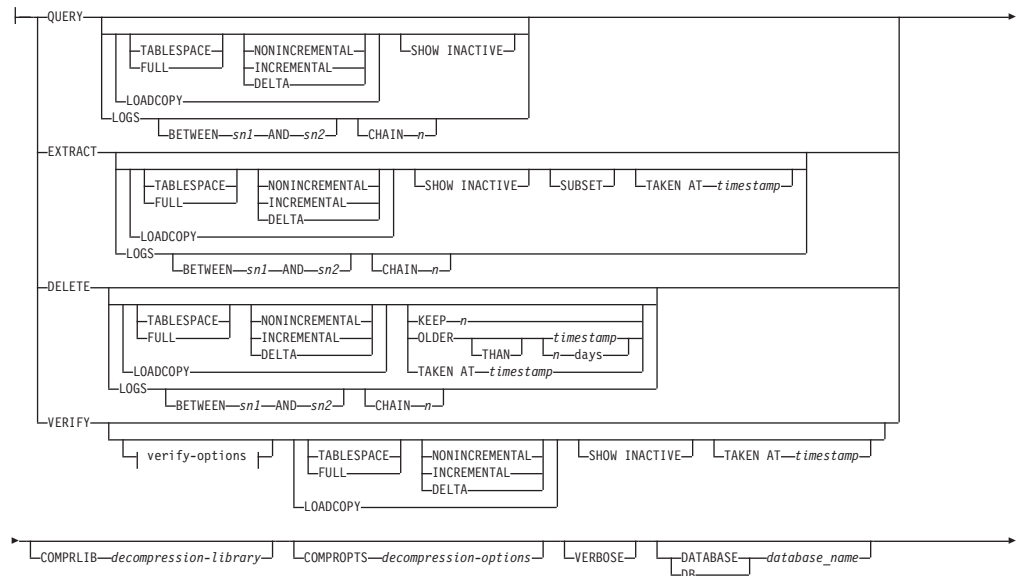
**Required connection:**

None

**Command syntax:**



**db2-object-options:**



## db2adutl - Managing DB2 objects within TSM

```
┌DBPARTITIONNUM—db-partition-number┐ ┌PASSWORD—password┐ ┌NODENAME—node_name┐ ┌OWNER—owner┐ ┌WITHOUT PROMPTING┐
```

### verify-options:

```
┌ALL┐
├CHECK┐
├DMS┐
├HEADER┐
├LFH┐
├TABLESPACES┐
├HEADERONLY┐
├TABLESPACESONLY┐
├OBJECT┐
└PAGECOUNT┐
```

### access-control-options:

```
┌GRANT┐ ┌ALL┐ ┌ON┐ ┌ALL┐ ┌FOR┐ ┌DATABASE┐ ┌database_name┐
├USER—user_name┐ ├NODENAME—node_name┐ ├DB┐
└REVOKE┐ ┌ALL┐ ┌ON┐ ┌ALL┐ ┌FOR┐ ┌ALL┐ ┌DATABASE┐ ┌database_name┐
├USER—user_name┐ ├NODENAME—node_name┐ ├DB┐
└QUERYACCESS┐ ┌FOR┐ ┌ALL┐ ┌DATABASE┐ ┌database_name┐
├DB┐
└PASSWORD—password┐
```

### Command parameters:

#### QUERY

Queries the TSM server for DB2 objects.

#### EXTRACT

Copies DB2 objects from the TSM server to the current directory on the local machine.

#### DELETE

Either deactivates backup objects or deletes log archives on the TSM server.

#### VERIFY

Performs consistency checking on the backup copy that is on the server.

**Note:** This parameter causes the entire backup image to be transferred over the network.

**ALL** Displays all available information.

#### CHECK

Displays results of checkbits and checksums.

#### DMS

Displays information from headers of DMS table space data pages.

#### HEADER

Displays the media header information.

#### HEADERONLY

Displays the same information as HEADER but only reads the 4 K media header information from the beginning of the image. It does not validate the image.

#### LFH

Displays the log file header (LFH) data.



**OBJECT**

Displays detailed information from the object headers.

**PAGECOUNT**

Displays the number of pages of each object type found in the image.

**TABLESPACES**

Displays the table space details, including container information, for the table spaces in the image.

**TABLESPACESONLY**

Displays the same information as TABLESPACES but does not validate the image.

**TABLESPACE**

Includes only table space backup images.

**FULL** Includes only full database backup images.

**NONINCREMENTAL**

Includes only non-incremental backup images.

**INCREMENTAL**

Includes only incremental backup images.

**DELTA**

Includes only incremental delta backup images.

**LOADCOPY**

Includes only load copy images.

**LOGS** Includes only log archive images

**BETWEEN *sn1* AND *sn2***

Specifies that the logs between log sequence number 1 and log sequence number 2 are to be used.

**CHAIN *n***

Specifies the chain ID of the logs to be used.

**SHOW INACTIVE**

Includes backup objects that have been deactivated.

**SUBSET**

Extracts pages from an image to a file. To extract pages, you will need an input and an output file. The default input file is called extractPage.in. You can override the default input file name by setting the DB2LISTFILE environment variable to a full path. The format of the input file is as follows:

For SMS table spaces:

S <tblspID> <objID> <objType> <startPage> <numPages>

For DMS table spaces:

D <tblspID> <objType> <startPage> <numPages>

**Note:** <objType> is only needed if verifying DMS load copy images

For log files:

L <log num> <startPos> <numPages>

For other data (for example, initial data):

O <objType> <startPos> <numBytes>

## db2adutl - Managing DB2 objects within TSM

| The default output file is extractPage.out. You can override the default  
| output file name by setting the DB2EXTRACTFILE environment variable to  
| a full path.

**TAKEN AT** *timestamp*

Specifies a backup image by its time stamp.

**KEEP** *n*

Deactivates all objects of the specified type except for the most recent *n* by time stamp.

**OLDER THAN** *timestamp* or *n days*

Specifies that objects with a time stamp earlier than *timestamp* or *n days* will be deactivated.

**COMPRLIB** *decompression-library*

| Indicates the name of the library to be used to perform the decompression.  
| The name must be a fully qualified path referring to a file on the server. If  
| this parameter is not specified, DB2 will attempt to use the library stored  
| in the image. If the backup was not compressed, the value of this  
| parameter will be ignored. If the specified library cannot be loaded, the  
| operation will fail.

**COMPROPTS** *decompression-options*

| Describes a block of binary data that will be passed to the initialization  
| routine in the decompression library. DB2 will pass this string directly from  
| the client to the server, so any issues of byte reversal or code page  
| conversion will have to be handled by the decompression library. If the  
| first character of the data block is '@', the remainder of the data will be  
| interpreted by DB2 as the name of a file residing on the server. DB2 will  
| then replace the contents of the data block with the contents of this file  
| and will pass this new value to the initialization routine instead. The  
| maximum length for this string is 1024 bytes.

**DATABASE** *database\_name*

Considers only those objects associated with the specified database name.

**DBPARTITIONNUM** *db-partition-number*

Considers only those objects created by the specified database partition number.

**PASSWORD** *password*

Specifies the TSM client password for this node, if required. If a database is specified and the password is not provided, the value specified for the *tsm\_password* database configuration parameter is passed to TSM; otherwise, no password is used.

**NODENAME** *node\_name*

Considers only those images associated with a specific TSM node name.

**OWNER** *owner*

Considers only those objects created by the specified owner.

**WITHOUT PROMPTING**

The user is not prompted for verification before objects are deleted.

**VERBOSE**

Displays additional file information.

**GRANT ALL / USER** *user\_name*

Adds access rights to the TSM files on the current TSM node to all users or

to the users specified. Granting access to users gives them access for all current and future files related to the database specified.

### **REVOKE ALL / USER** *user\_name*

Removes access rights to the TSM files on the current TSM node from all users or to the users specified.

### **QUERYACCESS**

Retrieves the current access list. A list of users and TSM nodes is displayed.

### **ON ALL / NODENAME** *node\_name*

Specifies the TSM node for which access rights will be changed.

### **FOR ALL / DATABASE** *database\_name*

Specifies the database to be considered.

### **Examples:**

1. The following is sample output from the command `db2 backup database rawsampl use tsm`

```
Backup successful. The timestamp for this backup is : 20031209184503
```

The following is sample output from the command `db2adutl query` issued following the backup operation:

```
Query for database RAWSAMPL
```

```
Retrieving FULL DATABASE BACKUP information.
```

```
1 Time: 20031209184403, Oldest log: S0000050.LOG, Sessions: 1
```

```
Retrieving INCREMENTAL DATABASE BACKUP information.
```

```
No INCREMENTAL DATABASE BACKUP images found for RAWSAMPL
```

```
Retrieving DELTA DATABASE BACKUP information.
```

```
No DELTA DATABASE BACKUP images found for RAWSAMPL
```

```
Retrieving TABLESPACE BACKUP information.
```

```
No TABLESPACE BACKUP images found for RAWSAMPL
```

```
Retrieving INCREMENTAL TABLESPACE BACKUP information.
```

```
No INCREMENTAL TABLESPACE BACKUP images found for RAWSAMPL
```

```
Retrieving DELTA TABLESPACE BACKUP information.
```

```
No DELTA TABLESPACE BACKUP images found for RAWSAMPL
```

```
Retrieving LOCAL COPY information.
```

```
No LOCAL COPY images found for RAWSAMPL
```

```
Retrieving log archive information.
```

```
Log file: S0000050.LOG, Chain Num: 0, DB Partition Number: 0,
```

```
  Taken at 2003-12-09-18.46.13
```

```
Log file: S0000051.LOG, Chain Num: 0, DB Partition Number: 0,
```

```
  Taken at 2003-12-09-18.46.43
```

```
Log file: S0000052.LOG, Chain Num: 0, DB Partition Number: 0,
```

```
  Taken at 2003-12-09-18.47.12
```

```
Log file: S0000053.LOG, Chain Num: 0, DB Partition Number: 0,
```

```
  Taken at 2003-12-09-18.50.14
```

```
Log file: S0000054.LOG, Chain Num: 0, DB Partition Number: 0,
```

```
  Taken at 2003-12-09-18.50.56
```

```
Log file: S0000055.LOG, Chain Num: 0, DB Partition Number: 0,
```

```
  Taken at 2003-12-09-18.52.39
```

2. The following is sample output from the command `db2adutl delete full taken at 20031209184503 db rawsampl`

## db2adutl - Managing DB2 objects within TSM

Query for database RAWSAMPL

Retrieving FULL DATABASE BACKUP information.

Taken at: 20031209184503 DB Partition Number: 0 Sessions: 1

Do you want to delete this file (Y/N)? y

Are you sure (Y/N)? y

Retrieving INCREMENTAL DATABASE BACKUP information.

No INCREMENTAL DATABASE BACKUP images found for RAWSAMPL

Retrieving DELTA DATABASE BACKUP information.

No DELTA DATABASE BACKUP images found for RAWSAMPL

The following is sample output from the command db2adutl query issued following the operation that deleted the full backup image. Note the timestamp for the backup image.

Query for database RAWSAMPL

Retrieving FULL DATABASE BACKUP information.

1 Time: 20031209184403, Oldest log: S0000050.LOG, Sessions: 1

Retrieving INCREMENTAL DATABASE BACKUP information.

No INCREMENTAL DATABASE BACKUP images found for RAWSAMPL

Retrieving DELTA DATABASE BACKUP information.

No DELTA DATABASE BACKUP images found for RAWSAMPL

Retrieving TABLESPACE BACKUP information.

No TABLESPACE BACKUP images found for RAWSAMPL

Retrieving INCREMENTAL TABLESPACE BACKUP information.

No INCREMENTAL TABLESPACE BACKUP images found for RAWSAMPL

Retrieving DELTA TABLESPACE BACKUP information.

No DELTA TABLESPACE BACKUP images found for RAWSAMPL

Retrieving LOCAL COPY information.

No LOCAL COPY images found for RAWSAMPL

Retrieving log archive information.

Log file: S0000050.LOG, Chain Num: 0, DB Partition Number: 0,

Taken at 2003-12-09-18.46.13

Log file: S0000051.LOG, Chain Num: 0, DB Partition Number: 0,

Taken at 2003-12-09-18.46.43

Log file: S0000052.LOG, Chain Num: 0, DB Partition Number: 0,

Taken at 2003-12-09-18.47.12

Log file: S0000053.LOG, Chain Num: 0, DB Partition Number: 0,

Taken at 2003-12-09-18.50.14

Log file: S0000054.LOG, Chain Num: 0, DB Partition Number: 0,

Taken at 2003-12-09-18.50.56

Log file: S0000055.LOG, Chain Num: 0, DB Partition Number: 0,

Taken at 2003-12-09-18.52.39

3. The following is sample output from the command db2adutl queryaccess for all

Node	User	Database Name	type
bar2	jchisan	sample	B
<all>	<all>	test	B

Access Types: B – Backup images L – Logs A – both

### Usage Notes:

One parameter from each group below can be used to restrict what backup images types are included in the operation:

#### Granularity:

- FULL - include only database backup images.
- TABLESPACE - include only table space backup images.

#### Cumulativeness:

- NONINCREMENTAL - include only non-incremental backup images.
- INCREMENTAL - include only incremental backup images.
- DELTA - include only incremental delta backup images.

#### Compatibilities:

For compatibility with versions earlier than Version 8:

- The keyword NODE can be substituted for DBPARTITIONNUM.

## db2ckbkp - Check Backup

This utility can be used to test the integrity of a backup image and to determine whether or not the image can be restored. It can also be used to display the metadata stored in the backup header.

#### Authorization:

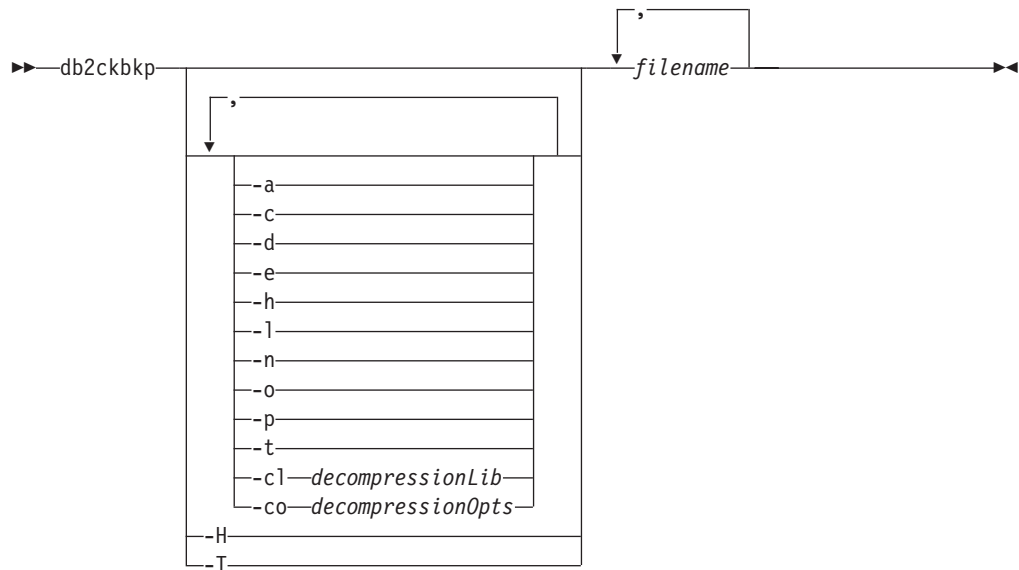
Anyone can access the utility, but users must have read permissions on image backups in order to execute this utility against them.

#### Required connection:

None

#### Command syntax:

## db2ckbkp - Check Backup



### Command parameters:

- a Displays all available information.
- c Displays results of checkbits and checksums.
- cl *decompressionLib*  
Indicates the name of the library to be used to perform the decompression. The name must be a fully qualified path referring to a file on the server. If this parameter is not specified, DB2 will attempt to use the library stored in the image. If the backup was not compressed, the value of this parameter will be ignored. If the specified library cannot be loaded, the operation will fail.
- co *decompressionOpts*  
Describes a block of binary data that will be passed to the initialization routine in the decompression library. DB2 will pass this string directly from the client to the server, so any issues of byte reversal or code page conversion will have to be handled by the decompression library. If the first character of the data block is '@', the remainder of the data will be interpreted by DB2 as the name of a file residing on the server. DB2 will then replace the contents of *string* with the contents of this file and will pass this new value to the initialization routine instead. The maximum length for string is 1024 bytes.
- d Displays information from the headers of DMS table space data pages.
- e Extracts pages from an image to a file. To extract pages, you will need an input and an output file. The default input file is called extractPage.in. You can override the default input file name by setting the DB2LISTFILE environment variable to a full path. The format of the input file is as follows:  
  
For SMS table spaces:  
S <tbodyID> <objID> <objType> <startPage> <numPages>  
  
For DMS table spaces:  
D <tbodyID> <objType> <startPage> <numPages>

**Note:** <objType> is only needed if verifying DMS load copy images  
For log files:

L <log num> <startPos> <numPages>

For other data (for example, initial data):

0 <objType> <startPos> <numBytes>

The default output file is extractPage.out. You can override the default output file name by setting the DB2EXTRACTFILE environment variable to a full path.

**-h** Displays media header information including the name and path of the image expected by the restore utility.

**-H** Displays the same information as -h but only reads the 4K media header information from the beginning of the image. It does not validate the image.

**Note:** This option cannot be used in combination with any other options.

**-l** Displays log file header (LFH) and mirror log file header (MFH) data.

**-n** Prompt for tape mount. Assume one tape per device.

**-o** Displays detailed information from the object headers.

**-p** Displays the number of pages of each object type.

**-t** Displays table space details, including container information, for the table spaces in the image.

**-T** Displays the same information as -t but does not validate the image.

**Note:** This option cannot be used in combination with any other options.

#### **filename**

The name of the backup image file. One or more files can be checked at a time.

#### **Notes:**

1. If the complete backup consists of multiple objects, the validation will only succeed if **db2ckbkp** is used to validate all of the objects at the same time.
2. When checking multiple parts of an image, the first backup image object (.001) must be specified first.

#### **Examples:**

##### Example 1 (on UNIX platforms)

```
db2ckbkp SAMPLE.0.krodger.NODE0000.CATN0000.19990817150714.001
SAMPLE.0.krodger.NODE0000.CATN0000.19990817150714.002
SAMPLE.0.krodger.NODE0000.CATN0000.19990817150714.003
```

```
[1] Buffers processed: ##
[2] Buffers processed: ##
[3] Buffers processed: ##
Image Verification Complete - successful.
```

##### Example 2 (on Windows platforms)

## db2ckbkp - Check Backup

```
db2ckbkp SAMPLE.0\krodger\NODE0000\CATN0000\19990817\150714.001
SAMPLE.0\krodger\NODE0000\CATN0000\19990817\150714.002
SAMPLE.0\krodger\NODE0000\CATN0000\19990817\150714.003
```

```
[1] Buffers processed: ##
[2] Buffers processed: ##
[3] Buffers processed: ##
Image Verification Complete - successful.
```

### Example 3

```
db2ckbkp -h SAMPLE2.0.krodger.NODE0000.CATN0000.19990818122909.001
```

```
=====
MEDIA HEADER REACHED:
=====
Server Database Name          -- SAMPLE2
Server Database Alias        -- SAMPLE2
Client Database Alias        -- SAMPLE2
Timestamp                     -- 19990818122909
Database Partition Number    -- 0
Instance                      -- krodger
Sequence Number              -- 1
Release ID                   -- 900
Database Seed                 -- 65E0B395
DB Comment's Codepage (Volume) -- 0
DB Comment (Volume)          --
DB Comment's Codepage (System) -- 0
DB Comment (System)          --
Authentication Value         -- 255
Backup Mode                   -- 0
Include Logs                  -- 0
Compression                   -- 0
Backup Type                   -- 0
Backup Gran.                  -- 0
Status Flags                  -- 11
System Cats inc               -- 1
Catalog Database Partition No. -- 0
DB Codeset                    -- ISO8859-1
DB Territory                  --
LogID                         -- 1074717952
LogPath                       -- /home/krodger/krodger/NODE0000/
                               SQL00001/SQLOGDIR
Backup Buffer Size             -- 4194304
Number of Sessions            -- 1
Platform                      -- 0
```

```
The proper image file name would be:
SAMPLE2.0.krodger.NODE0000.CATN0000.19990818122909.001
```

```
[1] Buffers processed: ###
Image Verification Complete - successful.
```

### Usage notes:

1. If a backup image was created using multiple sessions, **db2ckbkp** can examine all of the files at the same time. Users are responsible for ensuring that the session with sequence number 001 is the first file specified.
2. This utility can also verify backup images that are stored on tape (except images that were created with a variable block size). This is done by preparing the tape as for a restore operation, and then invoking the utility, specifying the tape device name. For example, on UNIX based systems:

```
db2ckbkp -h /dev/rmt0
```

and on Windows:



```
db2ckbkp -d \\.\tape1
```

3. If the image is on a tape device, specify the tape device path. You will be prompted to ensure it is mounted, unless option '-n' is given. If there are multiple tapes, the first tape must be mounted on the first device path given. (That is the tape with sequence 001 in the header).

The default when a tape device is detected is to prompt the user to mount the tape. The user has the choice on the prompt. Here is the prompt and options: (where the device I specified is on device path /dev/rmt0)

```
Please mount the source media on device /dev/rmt0.
Continue(c), terminate only this device(d), or abort this tool(t)?
(c/d/t)
```

The user will be prompted for each device specified, and when the device reaches the end of tape.

**Related reference:**

- “db2adutl - Managing DB2 objects within TSM” on page 237

## db2ckrst - Check Incremental Restore Image Sequence

Queries the database history and generates a list of timestamps for the backup images that are required for an incremental restore. A simplified restore syntax for a manual incremental restore is also generated.


**Authorization:**

None

**Required connection:**

None

**Command syntax:**



**Command parameters:**

- d *database name*  
Specifies the alias name for the database that will be restored.
- t *timestamp*  
Specifies the timestamp for a backup image that will be incrementally restored.
- r  
Specifies the type of restore that will be executed. The default is database.

**Note:** If tablespace is chosen and no table space names are given, the utility looks into the history entry of the specified image and uses the table space names listed to do the restore.

## db2ckrst - Check Incremental Restore Image Sequence

**-n** *tablespace name*

Specifies the name of one or more table spaces that will be restored.

**Note:** If a database restore type is selected and a list of table space names is specified, the utility will continue as a tablespace restore using the table space names given.

**-h/-u/-?**

Displays help information. When this option is specified, all other options are ignored, and only the help information is displayed.

### Examples:

```
db2ckrst -d mr -t 20001015193455 -r database
db2ckrst -d mr -t 20001015193455 -r tablespace
db2ckrst -d mr -t 20001015193455 -r tablespace -n tbsp1 tbsp2
```

```
> db2 backup db mr
```

```
Backup successful. The timestamp for this backup image is : 20001016001426
```

```
> db2 backup db mr incremental
```

```
Backup successful. The timestamp for this backup image is : 20001016001445
```

```
> db2ckrst -d mr -t 20001016001445
```

```
Suggested restore order of images using timestamp 20001016001445 for
database mr.
```

```
=====
db2 restore db mr incremental taken at 20001016001445
db2 restore db mr incremental taken at 20001016001426
db2 restore db mr incremental taken at 20001016001445
=====
```

```
> db2ckrst -d mr -t 20001016001445 -r tablespace -n userspace1
Suggested restore order of images using timestamp 20001016001445 for
database mr.
```

```
=====
db2 restore db mr tablespace ( USERSPACE1 ) incremental taken at
20001016001445
db2 restore db mr tablespace ( USERSPACE1 ) incremental taken at
20001016001426
db2 restore db mr tablespace ( USERSPACE1 ) incremental taken at
20001016001445
=====
```

### Usage notes:

The database history must exist in order for this utility to be used. If the database history does not exist, specify the HISTORY FILE option in the RESTORE command before using this utility.

If the FORCE option of the PRUNE HISTORY command is used, you may delete entries that are required for automatic incremental restoration of databases. Manual restores will still work correctly. Use of this command can also prevent the dbckrst utility from being able to correctly analyse the complete chain of required backup images. The default operation of the PRUNE HISTORY command prevents required entries from being deleted. It is recommended that you do not use the FORCE option of the PRUNE HISTORY command.

This utility should not be used as a replacement for keeping records of your backups.

## db2ckrst - Check Incremental Restore Image Sequence

### Related tasks:

- “Restoring from incremental backup images” on page 26

### Related reference:

- “RESTORE DATABASE” on page 87
- “PRUNE HISTORY/LOGFILE” on page 260

## db2flsn - Find Log Sequence Number

Returns the name of the file that contains the log record identified by a specified log sequence number (LSN).

### Authorization:

None

### Command syntax:

```
▶▶ db2flsn  input_LSN ▶▶
```

### Command parameters:

**-q** Specifies that only the log file name be printed. No error or warning messages will be printed, and status can only be determined through the return code. Valid error codes are:

- -100 Invalid input
- -101 Cannot open LFH file
- -102 Failed to read LFH file
- -103 Invalid LFH
- -104 Database is not recoverable
- -105 LSN too big
- -500 Logical error.

Other valid return codes are:

- 0 Successful execution
- 99 Warning: the result is based on the last known log file size.

### input\_LSN

A 12-character string that represents the internal (6-byte) hexadecimal value with leading zeros.

### Examples:

```
db2flsn 000000BF0030
Given LSN is contained in log file S0000002.LOG
```

```
db2flsn -q 000000BF0030
S0000002.LOG
```

```
db2flsn 000000BE0030
Warning: the result is based on the last known log file size.
The last known log file size is 23 4K pages starting from log extent 2.
```

## db2flsn - Find Log Sequence Number

Given LSN is contained in log file S0000001.LOG

```
db2flsn -q 000000BE0030
S0000001.LOG
```

### Usage notes:

The log header control file SQLLOGCTL.LFH must reside in the current directory. Since this file is located in the database directory, the tool can be run from the database directory, or the control file can be copied to the directory from which the tool will be run.

The tool uses the *logfilsiz* database configuration parameter. DB2 records the three most recent values for this parameter, and the first log file that is created with each *logfilsiz* value; this enables the tool to work correctly when *logfilsiz* changes. If the specified LSN predates the earliest recorded value of *logfilsiz*, the tool uses this value, and returns a warning. The tool can be used with database managers prior to UDB Version 5.2; in this case, the warning is returned even with a correct result (obtained if the value of *logfilsiz* remains unchanged).

This tool can only be used with recoverable databases. A database is recoverable if it is configured with the *logarchmeth1* or *logarchmeth2* configuration parameters set to a value other than OFF.

## db2inidb - Initialize a Mirrored Database

Initializes a mirrored database in a split mirror environment. The mirrored database can be initialized as a clone of the primary database, placed in roll forward pending state, or used as a backup image to restore the primary database. This command can only be run against a split mirror database, and it must be run before the split mirror can be used.

### Authorization:

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*

### Required connection:

None

### Command syntax:

```
db2inidb database_alias AS 

|          |
|----------|
| SNAPSHOT |
| STANDBY  |
| MIRROR   |

 RELOCATE USING configFile
```

### Command parameters:

#### database\_alias

Specifies the alias of the database to be initialized.

#### SNAPSHOT

Specifies that the mirrored database will be initialized as a clone of the primary database.

### STANDBY

Specifies that the database will be placed in roll forward pending state.

**Note:** New logs from the primary database can be fetched and applied to the standby database. The standby database can then be used in place of the primary database if it goes down.

### MIRROR

Specifies that the mirrored database is to be used as a backup image which can be used to restore the primary database.

### RELOCATE USING *configFile*

Specifies that the database files are to be relocated based on the information listed in the specified *configFile* prior to initializing the database as a snapshot, standby, or mirror. The format of *configFile* is described in `db2relocatedb - Relocate Database Command`.

### Usage notes:

In a partitioned database environment, **db2inidb** must be run on every partition before the split mirror from any of the partitions can be used. **db2inidb** can be run on all partitions simultaneously using the **db2\_all** command.

If; however, you are using the RELOCATE USING option, you cannot use the **db2\_all** command to run **db2inidb** on all of the partitions simultaneously. A separate configuration file must be supplied for each partition, that includes the NODENUM value of the partition being changed. For example, if the name of a database is being changed, every partition will be affected and the **db2relocatedb** command must be run with a separate configuration file on each partition. If containers belonging to a single database partition are being moved, the **db2relocatedb** command only needs to be run once on that partition.

If the RELOCATE USING *configFile* parameter is specified and the database is relocated successfully, the specified *configFile* will be copied into the database directory and renamed to `db2path.cfg`. During a subsequent crash recovery or rollforward recovery, this file will be used to rename container paths as log files are being processed.

If a clone database is being initialized, the specified *configFile* will be automatically removed from the database directory after a crash recovery is completed.

If a standby database or mirrored database is being initialized, the specified *configFile* will be automatically removed from the database directory after a rollforward recovery is completed or canceled. New container paths can be added to the `db2path.cfg` file after **db2inidb** has been run. This would be necessary when CREATE or ALTER TABLESPACE operations are done on the original database and different paths must be used on the standby database.

### Related tasks:

- “Using a split mirror to clone a database” on page 165
- “Using a split mirror as a standby database” on page 167
- “Using a split mirror as a backup image” on page 168

### Related reference:

- “`db2relocatedb - Relocate Database Command`” in the *Command Reference*

## db2inidb - Initialize a Mirrored Database

- “rah and db2\_all command descriptions” in the *Administration Guide: Implementation*

## db2mscs - Set up Windows Failover Utility

Creates the infrastructure for DB2 failover support on Windows using Microsoft Cluster Server (MSCS). This utility can be used to enable failover in both single-partition and partitioned database environments.

### Authorization:

The user must be logged on to a domain user account that belongs to the Administrators group of each machine in the MSCS cluster.

### Command syntax:

```
db2mscs -f:input_file -u:instance_name
```

### Command parameters:

#### -f:input\_file

Specifies the DB2MSCS.CFG input file to be used by the MSCS utility. If this parameter is not specified, the DB2MSCS utility reads the DB2MSCS.CFG file that is in the current directory.

#### -u:instance\_name

This option allows you to undo the db2mscs operation and revert the instance back to the non-MSCS instance specified by instance\_name.

### Usage notes:

The DB2MSCS utility is a standalone command line utility used to transform a non-MSCS instance into an MSCS instance. The utility will create all MSCS groups, resources, and resource dependencies. It will also copy all DB2 information stored in the Windows registry to the cluster portion of the registry as well as moving the instance directory to a shared cluster disk. The DB2MSCS utility takes as input a configuration file provided by the user specifying how the cluster should be set up. The DB2MSCS.CFG file is an ASCII text file that contains parameters that are read by the DB2MSCS utility. You specify each input parameter on a separate line using the following format: PARAMETER\_KEYWORD=parameter\_value. For example:

```
CLUSTER_NAME=FINANCE
GROUP_NAME=DB2 Group
IP_ADDRESS=9.21.22.89
```

Two example configuration files can be found in the CFG subdirectory under the DB2 install directory. The first, DB2MSCS.EE, is an example for single-partition database environments. The second, DB2MSCS.EEE, is an example for partitioned database environments.

The parameters for the DB2MSCS.CFG file are as follows:

#### DB2\_INSTANCE

The name of the DB2 instance. This parameter has a global scope and should be specified only once in the DB2MSCS.CFG file.

### **DAS\_INSTANCE**

The name of the DB2 Admin Server instance. Specify this parameter to migrate the DB2 Admin Server to run in the MSCS environment. This parameter has a global scope and should be specified only once in the DB2MSCS.CFG file.

### **CLUSTER\_NAME**

The name of the MSCS cluster. All the resources specified following this line are created in this cluster until another CLUSTER\_NAME parameter is specified.

### **DB2\_LOGON\_USERNAME**

The user name of the domain account for the DB2 service (specified as *domain\user*). This parameter has a global scope and should be specified only once in the DB2MSCS.CFG file.

### **DB2\_LOGON\_PASSWORD**

The password of the domain account for the DB2 service. This parameter has a global scope and should be specified only once in the DB2MSCS.CFG file.

### **GROUP\_NAME**

The name of the MSCS group. If this parameter is specified, a new MSCS group is created if it does not exist. If the group already exists, it is used as the target group. Any MSCS resource specified after this parameter is created in this group or moved into this group until another GROUP\_NAME parameter is specified. Specify this parameter once for each group.

### **DB2\_NODE**

The partition number of the database partition server (or database partition) to be included in the current MSCS group. If multiple logical database partitions exist on the same machine, each database partition requires a separate DB2\_NODE parameter. Specify this parameter after the GROUP\_NAME parameter so that the DB2 resources are created in the correct MSCS group. This parameter is required for a multi-partitioned database system.

### **IP\_NAME**

The name of the IP Address resource. The value for the IP\_NAME is arbitrary, but it must be unique in the cluster. When this parameter is specified, an MSCS resource of type IP Address is created. This parameter is required for remote TCP/IP connections. This parameter is optional in a single partition environment. A recommended name is the hostname that corresponds to the IP address.

### **IP\_ADDRESS**

The TCP/IP address for the IP resource specified by the preceding IP\_NAME parameter. This parameter is required if the IP\_NAME parameter is specified. This is a new IP address that is not used by any machine in the network.

### **IP\_SUBNET**

The TCP/IP subnet mask for the IP resource specified by the preceding IP\_NAME parameter. This parameter is required if the IP\_NAME parameter is specified.

### **IP\_NETWORK**

The name of the MSCS network to which the preceding IP Address resource belongs. This parameter is optional. If it is not specified, the first

## db2mscs - Set up Windows Failover Utility

MSCS network detected by the system is used. The name of the MSCS network must be entered exactly as seen under the Networks branch in Cluster Administrator.

**Note:** The previous four IP keywords are used to create an IP Address resource.

### NETNAME\_NAME

The name of the Network Name resource. Specify this parameter to create the Network Name resource. This parameter is optional for single partition database environment. You must specify this parameter for the instance owning machine in a partitioned database environment.

### NETNAME\_VALUE

The value for the Network Name resource. This parameter must be specified if the NETNAME\_NAME parameter is specified.

### NETNAME\_DEPENDENCY

The name for the IP resource that the Network Name resource depends on. Each Network Name resource must have a dependency on an IP Address resource. This parameter is optional. If it is not specified, the Network Name resource has a dependency on the first IP resource in the group.

### SERVICE\_DISPLAY\_NAME

The display name of the Generic Service resource. Specify this parameter if you want to create a Generic Service resource.

### SERVICE\_NAME

The service name of the Generic Service resource. This parameter must be specified if the SERVICE\_DISPLAY\_NAME parameter is specified.

### SERVICE\_STARTUP

Optional startup parameter for the Generic Resource service.

### DISK\_NAME

The name of the physical disk resource to be moved to the current group. Specify as many disk resources as you need. The disk resources must already exist. When the DB2MSCS utility configures the DB2 instance for failover support, the instance directory is copied to the first MSCS disk in the group. To specify a different MSCS disk for the instance directory, use the INSTPROF\_DISK parameter. The disk name used should be entered exactly as seen in Cluster Administrator.

### INSTPROF\_DISK

An optional parameter to specify an MSCS disk to contain the DB2 instance directory. If this parameter is not specified the DB2MSCS utility uses the first disk that belongs to the same group.

### INSTPROF\_PATH

An optional parameter to specify the exact path where the instance directory will be copied. This parameter *must* be specified when using IPSHADisks, a ServerRAID Netfinity disk resource (for example, INSTPROF\_PATH=p:\db2profs). INSTPROF\_PATH will take precedence over INSTPROF\_DISK if both are specified.

### TARGET\_DRVMAP\_DISK

An optional parameter to specify the target MSCS disk for database drive mapping for a the multi-partitioned database system. This parameter will specify the disk the database will be created on by mapping it from the



drive the create database command specifies. If this parameter is not specified, the database drive mapping must be manually registered using the DB2DRVMP utility.

### DB2\_FALLBACK

An optional parameter to control whether or not the applications should be forced off when the DB2 resource is brought offline. If not specified, then the setting for DB2\_FALLBACK will be YES. If you do not want the applications to be forced off, then set DB2\_FALLBACK to NO.

## db2rfpen - Reset rollforward pending state

Puts a database in rollforward pending state. If you are using high availability disaster recovery (HADR), the database is reset to a standard database.

### Authorization:

None

### Required connection:

None

### Command syntax:

```
▶▶ db2rfpen -ON database_alias -log logfile_path ▶▶
```

### Command parameters:

#### database\_alias

Specifies the name of the database to be placed in rollforward pending state. If you are using high availability disaster recovery (HADR), the database is reset to a standard database.

#### -log logfile\_path

Specifies the log file path.

### Related concepts:

- “High availability disaster recovery overview” on page 173

---

## CLP commands

### ARCHIVE LOG

Closes and truncates the active log file for a recoverable database.

### Authorization:

One of the following:

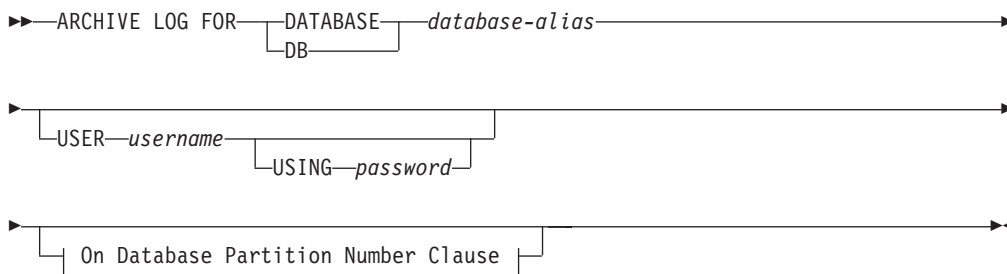
- *sysadm*
- *sysctrl*
- *sysmaint*
- *dbadm*

### Required connection:


## ARCHIVE LOG

None. This command establishes a database connection for the duration of the command.


### Command syntax:



### On Database Partition Number Clause:



### Database Partition Number List Clause:



### Command parameters:

#### DATABASE *database-alias*

Specifies the alias of the database whose active log is to be archived.

#### USER *username*

Identifies the user name under which a connection will be attempted.

#### USING *password*

Specifies the password to authenticate the user name.

#### ON ALL DBPARTITIONNUMS

Specifies that the command should be issued on all database partitions in the `db2nodes.cfg` file. This is the default if a database partition number clause is not specified.

#### EXCEPT

Specifies that the command should be issued on all database partitions in the `db2nodes.cfg` file, except those specified in the database partition number list.

#### ON DBPARTITIONNUM/ON DBPARTITIONNUMS

Specifies that the logs should be archived for the specified database on a set of database partitions.

#### *db-partition-number*

Specifies a database partition number in the database partition number list.

**TO db-partition-number**

Used when specifying a range of database partitions for which the logs should be archived. All database partitions from the first database partition number specified up to and including the second database partition number specified are included in the database partition number list.

**Usage notes:**

This command can be used to collect a complete set of log files up to a known point. The log files can then be used to update a standby database.

This command can only be executed when the invoking application or shell does not have a database connection to the specified database. This prevents a user from executing the command with uncommitted transactions. As such, the ARCHIVE LOG command will not forcibly commit the user's incomplete transactions. If the invoking application or shell already has a database connection to the specified database, the command will terminate and return an error. If another application has transactions in progress with the specified database when this command is executed, there will be a slight performance degradation since the command flushes the log buffer to disk. Any other transactions attempting to write log records to the buffer will have to wait until the flush is complete.

If used in a partitioned database environment, a subset of database partitions may be specified by using a database partition number clause. If the database partition number clause is not specified, the default behavior for this command is to close and archive the active log on all database partitions.

Using this command will use up a portion of the active log space due to the truncation of the active log file. The active log space will resume its previous size when the truncated log becomes inactive. Frequent use of this command may drastically reduce the amount of the active log space available for transactions.

**Compatibilities:**

For compatibility with versions earlier than Version 8:

- The keyword NODE can be substituted for DBPARTITIONNUM.
- The keyword NODES can be substituted for DBPARTITIONNUMS.

**INITIALIZE TAPE**

When running on Windows NT-based operating systems, DB2 supports backup and restore operations to streaming tape devices. Use this command for tape initialization.

**Authorization:**

None

**Required connection:**

None

**Command syntax:**

## INITIALIZE TAPE



### Command parameters:

#### ON device

Specifies a valid tape device name. The default value is `\\.\TAPE0`.

#### USING blksize

Specifies the block size for the device, in bytes. The device is initialized to use the block size specified, if the value is within the supported range of block sizes for the device.

**Note:** The buffer size specified for the BACKUP DATABASE command and for RESTORE DATABASE must be divisible by the block size specified here.

If a value for this parameter is not specified, the device is initialized to use its default block size. If a value of zero is specified, the device is initialized to use a variable length block size; if the device does not support variable length block mode, an error is returned.

### Related reference:

- “BACKUP DATABASE” on page 66
- “RESTORE DATABASE” on page 87
- “REWIND TAPE” on page 262
- “SET TAPE POSITION” on page 262

## LIST HISTORY

Lists entries in the history file. The history file contains a record of recovery and administrative events. Recovery events include full database and table space level backup, incremental backup, restore, and rollforward operations. Additional logged events include create, alter, drop, or rename table space, reorganize table, drop table, and load.


### Authorization:


None

### Required connection:

Instance. You must attach to any remote database in order to run this command against it. For a local database, an explicit attachment is not required.

### Command syntax:



**Command parameters:****HISTORY**

Lists all events that are currently logged in the history file.

**BACKUP**

Lists backup and restore operations.

**ROLLFORWARD**

Lists rollforward operations.

**DROPPED TABLE**

Lists dropped table records. A dropped table record is created only when the table is dropped and the table space containing it has the DROPPED TABLE RECOVERY option enabled.

**LOAD**

Lists load operations.

**CREATE TABLESPACE**

Lists table space create and drop operations.

**RENAME TABLESPACE**

Lists table space renaming operations.

**REORG**

Lists reorganization operations.

**ALTER TABLESPACE**

Lists alter table space operations.

**ARCHIVE LOG**

Lists archive log operations and the archived logs.

**ALL** Lists all entries of the specified type in the history file.

**SINCE timestamp**

A complete time stamp (format *yyyymmddhhmmss*), or an initial prefix (minimum *yyyy*) can be specified. All entries with time stamps equal to or greater than the time stamp provided are listed.

**CONTAINING schema.object\_name**

This qualified name uniquely identifies a table.

**CONTAINING object\_name**

This unqualified name uniquely identifies a table space.

**FOR DATABASE database-alias**

Used to identify the database whose recovery history file is to be listed.

**Examples:**

```

db2 list history since 19980201 for sample
db2 list history backup containing userspace1 for sample
db2 list history dropped table all for db sample
  
```

**Usage notes:**

The report generated by this command contains the following symbols:

## LIST HISTORY

### Operation

- A - Create table space
- B - Backup
- C - Load copy
- D - Dropped table
- F - Roll forward
- G - Reorganize table
- L - Load
- N - Rename table space
- O - Drop table space
- Q - Quiesce
- R - Restore
- T - Alter table space
- U - Unload
- X - Archive log

### Type

#### Archive Log types:

- P - Primary log path
- M - Secondary (mirror) log path
- F - Failover archive path
- 1 - Primary log archive method
- 2 - Secondary log archive method

#### Backup types:

- F - Offline
- N - Online
- I - Incremental offline
- O - Incremental online
- D - Delta offline
- E - Delta online

#### Rollforward types:

- E - End of logs
- P - Point in time

#### Load types:

- I - Insert
- R - Replace

#### Alter table space types:

- C - Add containers
- R - Rebalance

#### Quiesce types:

- S - Quiesce share
- U - Quiesce update
- X - Quiesce exclusive
- Z - Quiesce reset

## PRUNE HISTORY/LOGFILE

Used to delete entries from the recovery history file, or to delete log files from the active log file path. Deleting entries from the recovery history file may be necessary if the file becomes excessively large and the retention period is high.

### Authorization:


One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*
- *dbadm*

**Required connection:**

Database

**Command syntax:**



**Command parameters:**

**HISTORY timestamp**

Identifies a range of entries in the recovery history file that will be deleted. A complete time stamp (in the form *yyyymmddhhmmss*), or an initial prefix (minimum *yyyy*) can be specified. All entries with time stamps equal to or less than the time stamp provided are deleted from the recovery history file.

**WITH FORCE OPTION**

Specifies that the entries will be pruned according to the time stamp specified, even if some entries from the most recent restore set are deleted from the file. A restore set is the most recent full database backup including any restores of that backup image. If this parameter is not specified, all entries from the backup image forward will be maintained in the history.

**AND DELETE**

Specifies that the associated log archives will be physically deleted (based on the location information) when the history file entry is removed. This option is especially useful for ensuring that archive storage space is recovered when log archives are no longer needed.

**Note:** If you are archiving logs via a user exit program, the logs cannot be deleted using this option.

**LOGFILE PRIOR TO logfile-name**

Specifies a string for a log file name, for example *S0000100.LOG*. All log files prior to (but not including) the specified log file will be deleted. The LOGRETAIN database configuration parameter must be set to RECOVERY or CAPTURE.

**Examples:**

To remove the entries for all restores, loads, table space backups, and full database backups taken before and including December 1, 1994 from the recovery history file, enter:

```
db2 prune history 199412
```

**Note:** 199412 is interpreted as 19941201000000.

## PRUNE HISTORY/LOGFILE

### Usage notes:

If the **FORCE** option is used, you might delete entries that are required for automatic incremental restoration of databases. Manual restores will still work correctly. Use of this command can also prevent the **dbckrst** utility from being able to correctly analyze the complete chain of required backup images. Using the **PRUNE HISTORY** command without the **FORCE** option prevents required entries from being deleted.

Pruning backup entries from the history file causes related file backups on DB2 Data Links Manager servers to be deleted.

## REWIND TAPE

When running on Windows NT-based operating systems, DB2 supports backup and restore operations to streaming tape devices. Use this command for tape rewinding.

### Authorization:

None

### Required connection:

None

### Command syntax:

```
➤➤ REWIND TAPE [ON device] ➤➤
```

### Command parameters:

#### ON device

Specifies a valid tape device name. The default value is `\\.\TAPE0`.

### Related reference:

- “INITIALIZE TAPE” on page 257
- “SET TAPE POSITION” on page 262

## SET TAPE POSITION

When running on Windows NT-based operating systems, DB2 supports backup and restore operations to streaming tape devices. Use this command for tape positioning.

### Authorization:


None

### Required connection:

None

### Command syntax:





**Command parameters:**

**ON device**

Specifies a valid tape device name. The default value is \\.\TAPE0.

**TO position**

Specifies the mark at which the tape is to be positioned. DB2 for Windows writes a tape mark after every backup image. A value of 1 specifies the first position, 2 specifies the second position, and so on. If the tape is positioned at tape mark 1, for example, archive 2 is positioned to be restored.

**Related reference:**

- "INITIALIZE TAPE" on page 257
- "REWIND TAPE" on page 262

## UPDATE HISTORY FILE

Updates the location, device type, or comment in a history file entry.

**Authorization:**


One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*
- *dbadm*

**Required connection:**

Database

**Command syntax:**



**Command parameters:**

**EID *eid***

Specifies the history entry ID.

**FOR *object-part***

Specifies the identifier for the backup or copy image. It is a time stamp with an optional sequence number from 001 to 999.

**LOCATION *new-location***

Specifies the new physical location of a backup image. The interpretation of this parameter depends on the device type.

## UPDATE HISTORY FILE

### DEVICE TYPE *new-device-type*

Specifies a new device type for storing the backup image. Valid device types are:

D	Disk
K	Diskette
T	Tape
A	TSM
U	User exit
P	Pipe
N	Null device
X	XBSA
Q	SQL statement
O	Other

### COMMENT *new-comment*

Specifies a new comment to describe the entry.

### STATUS *new-status*

Specifies a new status for an entry. Valid values are:

A	Marks the entry as active.
I	Marks the entry as inactive.

### Examples:

To update the history file entry for a full database backup taken on April 13, 1997 at 10:00 a.m., enter:

```
db2 update history for 19970413100000001 with
location /backup/dbbackup.1 device type d
```

### Usage notes:

The primary purpose of the database history file is to record information, but the data contained in the history is used directly by automatic restore operations. During any restore where the AUTOMATIC option is specified, the history of backup images and their locations will be referenced and used by the restore utility to fulfill the automatic restore request. If the automatic restore function is to be used and backup images have been relocated since they were created, it is recommended that the database history record for those images be updated to reflect the current location. If the backup image location in the database history is not updated, automatic restore will not be able to locate the backup images, but manual restore commands can still be used successfully.

### Related reference:

- “PRUNE HISTORY/LOGFILE” on page 260

---

## Appendix D. Additional APIs and associated data structures

This appendix describes recovery-related APIs and their data structures that are not discussed in detail in this manual.

---

### db2ArchiveLog - Archive Active Log

Closes and truncates the active log file for a recoverable database. If user exit is enabled, issues an archive request.

#### Authorization:

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*
- *dbadm*

#### Required connection:

This API automatically establishes a connection to the specified database. If a connection to the specified database already exists, the API will return an error.

#### API include file:

*db2ApiDf.h*

#### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2ArchiveLog */
/* ... */
SQL_API_RC SQL_API_FN
db2ArchiveLog (
    db2UInt32 version,
    void *pDB2ArchiveLogStruct,
    struct sqlca *pSqlca);

typedef struct
{
    char          *piDatabaseAlias;
    char          *piUserName;
    char          *piPassword;
    db2UInt16     iAllNodeFlag;
    db2UInt16     iNumNodes;
    SQL_PDB_NODE_TYPE *piNodeList;
    db2UInt32     iOptions;
} db2ArchiveLogStruct
/* ... */
```

#### Generic API syntax:

```
/* File: db2ApiDf.h */
/* API: db2gArchiveLog */
/* ... */
SQL_API_RC SQL_API_FN
db2gArchiveLog (
```

## db2ArchiveLog - Archive Active Log

```
    db2UInt32 version,
    void *pDB2ArchiveLogStruct,
    struct sqlca *pSqlca);

typedef struct
{
    db2UInt32          iAliasLen;
    db2UInt32          iUserNameLen;
    db2UInt32          iPasswordLen;
    char               *piDatabaseAlias;
    char               *piUserName;
    char               *piPassword;
    db2UInt16          iAllNodeFlag;
    db2UInt16          iNumNodes;
    SQL_PDB_NODE_TYPE *piNodeList;
    db2UInt32          iOptions;
} db2ArchiveLogStruct
/* ... */
```

### API parameters:

#### version

Input. Specifies the version and release level of the variable passed in as the second parameter, *pDB2ArchiveLogStruct*.

#### pDB2ArchiveLogStruct

Input. A pointer to the *db2ArchiveLogStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### iAliasLen

Input. A 4-byte unsigned integer representing the length in bytes of the database alias.

#### iUserNameLen

Input. A 4-byte unsigned integer representing the length in bytes of the user name. Set to zero if no user name is used.

#### iPasswordLen

Input. A 4-byte unsigned integer representing the length in bytes of the password. Set to zero if no password is used.

#### piDatabaseAlias

Input. A string containing the database alias (as cataloged in the system database directory) of the database for which the active log is to be archived.

#### piUserName

Input. A string containing the user name to be used when attempting a connection.

#### piPassword

Input. A string containing the password to be used when attempting a connection.

#### iAllNodeFlag

Partitioned database environment only. Input. Flag indicating whether the operation should apply to all nodes listed in the *db2nodes.cfg* file. Valid values are:

#### DB2ARCHIVELOG\_NODE\_LIST

Apply to nodes in a node list that is passed in *piNodeList*.

**DB2ARCHIVELOG\_ALL\_NODES**

Apply to all nodes. *piNodeList* should be NULL. This is the default value.

**DB2ARCHIVELOG\_ALL\_EXCEPT**

Apply to all nodes except those in the node list passed in *piNodeList*.

**iNumNodes**

Partitioned database environment only. Input. Specifies the number of nodes in the *piNodeList* array.

**piNodeList**

Partitioned database environment only. Input. A pointer to an array of node numbers against which to apply the archive log operation.

**iOptions**

Input. Reserved for future use.

**Related reference:**

- “ARCHIVE LOG” on page 255

---

## db2HistoryCloseScan - Close History File Scan

Ends a history file scan and frees DB2 resources required for the scan. This API must be preceded by a successful call to `db2HistoryOpenScan`.

**Authorization:**

None

**Required connection:**

Instance. It is not necessary to call `sqlcattin` before calling this API.

**API include file:**

*db2ApiDf.h*

**C API syntax:**

```
/* File: db2ApiDf.h */
/* API: db2HistoryCloseScan */
/* ... */
SQL_API_RC SQL_API_FN
db2HistoryCloseScan (
    db2UInt32 version,
    void *piHandle,
    struct sqlca *pSqlca);
/* ... */
```

**Generic API syntax:**

```
/* File: db2ApiDf.h */
/* API: db2GenHistoryCloseScan */
/* ... */
SQL_API_RC SQL_API_FN
db2GenHistoryCloseScan (
```

## db2HistoryCloseScan - Close History File Scan

```
db2Uint32 version,  
void *piHandle,  
struct sqlca *pSqlca);  
/* ... */
```

### API parameters:

#### version

Input. Specifies the version and release level of the second parameter, *piHandle*.

#### piHandle

Input. Specifies a pointer to the handle for scan access that was returned by `db2HistoryOpenScan`.

#### pSqlca

Output. A pointer to the *sqlca* structure.

### REXX API syntax:

```
CLOSE RECOVERY HISTORY FILE :scanid
```

### REXX API parameters:

**scanid** Host variable containing the scan identifier returned from `OPEN RECOVERY HISTORY FILE SCAN`.

### Usage notes:

For a detailed description of the use of the history file APIs, see `db2HistoryOpenScan`.

### Related reference:

- “`db2Prune - Prune History File`” on page 278
- “`db2HistoryUpdate - Update History File`” on page 275
- “`db2HistoryOpenScan - Open History File Scan`” on page 271
- “`db2HistoryGetEntry - Get Next History File Entry`” on page 268
- “SQLCA” in the *Administrative API Reference*

### Related samples:

- “`dbrecov.sqc -- How to recover a database (C)`”
- “`dbrecov.sqC -- How to recover a database (C++)`”

---

## db2HistoryGetEntry - Get Next History File Entry

Gets the next entry from the history file. This API must be preceded by a successful call to `db2HistoryOpenScan`.

### Authorization:

None

### Required connection:

Instance. It is not necessary to call `sqleatin` before calling this API.

### API include file:

## db2HistoryGetEntry - Get Next History File Entry

*db2ApiDf.h*

### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2HistoryGetEntry */
/* ... */
SQL_API_RC SQL_API_FN
db2HistoryGetEntry (
    db2UInt32 version,
    void *pDB2HistoryGetEntryStruct,
    struct sqlca *pSqlca);

typedef struct
{
    db2UInt16 iHandle,
    db2UInt16 iCallerAction,
    struct db2HistData *pioHistData
} db2HistoryGetEntryStruct;
/* ... */
```

### Generic API syntax:

```
/* File: db2ApiDf.h */
/* API: db2GenHistoryGetEntry */
/* ... */
SQL_API_RC SQL_API_FN
db2GenHistoryGetEntry (
    db2UInt32 version,
    void *pDB2GenHistoryGetEntryStruct,
    struct sqlca *pSqlca);

typedef struct
{
    db2UInt16 iHandle,
    db2UInt16 iCallerAction,
    struct db2HistData *pioHistData
} db2GenHistoryGetEntryStruct;
/* ... */
```

### API parameters:

#### version

Input. Specifies the version and release level of the structure passed in as the second parameter, *pDB2HistoryGetEntryStruct*.

#### pDB2HistoryGetEntryStruct

Input. A pointer to the *db2HistoryGetEntryStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### iHandle

Input. Contains the handle for scan access that was returned by *db2HistoryOpenScan*.

#### iCallerAction

Input. Specifies the type of action to be taken. Valid values (defined in *db2ApiDf*) are:

#### DB2HISTORY\_GET\_ENTRY

Get the next entry, but without any command data.

#### DB2HISTORY\_GET\_DDL

Get only the command data from the previous fetch.

## db2HistoryGetEntry - Get Next History File Entry

### DB2HISTORY\_GET\_ALL

Get the next entry, including all data.

#### pioHistData

Input. A pointer to the *db2HistData* structure.

#### REXX API syntax:

```
GET RECOVERY HISTORY FILE ENTRY :scanid [USING :value]
```

#### REXX API parameters:

**scanid** Host variable containing the scan identifier returned from OPEN RECOVERY HISTORY FILE SCAN.

**value** A compound REXX host variable into which the history file entry information is returned. In the following, XXX represents the host variable name:

XXX.0	Number of first level elements in the variable (always 15)
XXX.1	Number of table space elements
XXX.2	Number of used table space elements
XXX.3	OPERATION (type of operation performed)
XXX.4	OBJECT (granularity of the operation)
XXX.5	OBJECT_PART (time stamp and sequence number)
XXX.6	OPTYPE (qualifier of the operation)
XXX.7	DEVICE_TYPE (type of device used)
XXX.8	FIRST_LOG (earliest log ID)
XXX.9	LAST_LOG (current log ID)
XXX.10	BACKUP_ID (identifier for the backup)
XXX.11	SCHEMA (qualifier for the table name)
XXX.12	TABLE_NAME (name of the loaded table)
XXX.13.0	NUM_OF_TABLESPACES (number of table spaces involved in backup or restore)
XXX.13.1	Name of the first table space backed up/restored
XXX.13.2	Name of the second table space backed up/restored
XXX.13.3	and so on
XXX.14	LOCATION (where backup or copy is stored)
XXX.15	COMMENT (text to describe the entry).

#### Usage notes:

The records that are returned will have been selected using the values specified on the call to *db2HistoryOpenScan*.

For a detailed description of the use of the history file APIs, see *db2HistoryOpenScan*.

#### Related reference:

- “db2Prune - Prune History File” on page 278



## db2HistoryGetEntry - Get Next History File Entry

- “db2HistoryUpdate - Update History File” on page 275
- “db2HistoryOpenScan - Open History File Scan” on page 271
- “db2HistoryCloseScan - Close History File Scan” on page 267
- “SQLCA” in the *Administrative API Reference*
- “db2HistData” on page 289

### Related samples:

- “dbrecov.sqc -- How to recover a database (C)”
- “dbrecov.sqC -- How to recover a database (C++)”

---

## db2HistoryOpenScan - Open History File Scan

Starts a history file scan.

### Authorization:

None

### Required connection:

Instance. If the database is cataloged as remote, call `sqlcatin` before calling this API.

### API include file:

*db2ApiDf.h*

### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2HistoryOpenScan */
/* ... */
SQL_API_RC SQL_API_FN
db2HistoryOpenScan (
    db2UInt32 version,
    void *pDB2HistoryOpenStruct,
    struct sqlca *pSqlca);
```

```
typedef struct
{
    char *piDatabaseAlias,
    char *piTimestamp,
    char *piObjectName,
    db2UInt32 oNumRows,
    db2UInt32 oMaxTbspaces,
    db2UInt16 iCallerAction,
    db2UInt16 oHandle
} db2HistoryOpenStruct;
/* ... */
```

### Generic API syntax:

```
/* File: db2ApiDf.h */
/* API: db2GenHistoryOpenScan */
/* ... */
SQL_API_RC SQL_API_FN
db2GenHistoryOpenScan (
    db2UInt32 version,
    void *pDB2GenHistoryOpenStruct,
    struct sqlca *pSqlca);
```

## db2HistoryOpenScan - Open History File Scan

```
typedef struct
{
    char *piDatabaseAlias,
    char *piTimestamp,
    char *piObjectName,
    db2Uint32 oNumRows,
    db2Uint32 oMaxTbspaces,
    db2Uint16 iCallerAction,
    db2Uint16 oHandle
} db2GenHistoryOpenStruct;
/* ... */
```

### API parameters:

#### version

Input. Specifies the version and release level of the structure passed in as the second parameter, *pDB2HistoryOpenStruct*.

#### pDB2HistoryOpenStruct

Input. A pointer to the *db2HistoryOpenStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### piDatabaseAlias

Input. A pointer to a string containing the database alias.

#### piTimestamp

Input. A pointer to a string specifying the time stamp to be used for selecting records. Records whose time stamp is equal to or greater than this value are selected. Setting this parameter to NULL, or pointing to zero, prevents the filtering of entries using a time stamp.

#### piObjectName

Input. A pointer to a string specifying the object name to be used for selecting records. The object may be a table or a table space. If it is a table, the fully qualified table name must be provided. Setting this parameter to NULL, or pointing to zero, prevents the filtering of entries using the object name.

#### oNumRows

Output. Upon return from the API, this parameter contains the number of matching history file entries.

#### oMaxTbspaces

Output. The maximum number of table space names stored with any history entry.

#### iCallerAction

Input. Specifies the type of action to be taken. Valid values (defined in *db2ApiDf*) are:

##### DB2HISTORY\_LIST\_HISTORY

Lists all events that are currently logged in the history file.

##### DB2HISTORY\_LIST\_BACKUP

Lists backup and restore operations.

##### DB2HISTORY\_LIST\_ROLLFORWARD

Lists rollforward operations.

##### DB2HISTORY\_LIST\_DROPPED\_TABLE

Lists dropped table records. The DDL field associated with an entry is not returned. To retrieve the DDL information for an entry,

## db2HistoryOpenScan - Open History File Scan

db2HistoryGetEntry must be called with a caller action of DB2HISTORY\_GET\_DDL immediately after the entry is fetched.

### DB2HISTORY\_LIST\_LOAD

Lists load operations.

### DB2HISTORY\_LIST\_CRT\_TABLESPACE

Lists table space create and drop operations.

### DB2HISTORY\_LIST\_REN\_TABLESPACE

Lists table space renaming operations.

### DB2HISTORY\_LIST\_ALT\_TABLESPACE

Lists alter table space operations. The DDL field associated with an entry is not returned. To retrieve the DDL information for an entry, db2HistoryGetEntry must be called with a caller action of DB2HISTORY\_GET\_DDL immediately after the entry is fetched.

### DB2HISTORY\_LIST\_REORG

Lists REORGANIZE TABLE operations. This value is not currently supported.

### oHandle

Output. Upon return from the API, this parameter contains the handle for scan access. It is subsequently used in db2HistoryGetEntry, and db2HistoryCloseScan.

### REXX API syntax:

```
OPEN [BACKUP] RECOVERY HISTORY FILE FOR database_alias  
[OBJECT objname] [TIMESTAMP :timestamp]  
USING :value
```

### REXX API parameters:

#### database\_alias

The alias of the database whose history file is to be listed.

#### objname

Specifies the object name to be used for selecting records. The object may be a table or a table space. If it is a table, the fully qualified table name must be provided. Setting this parameter to NULL prevents the filtering of entries using *objname*.

#### timestamp

Specifies the time stamp to be used for selecting records. Records whose time stamp is equal to or greater than this value are selected. Setting this parameter to NULL prevents the filtering of entries using *timestamp*.

**value** A compound REXX host variable to which history file information is returned. In the following, XXX represents the host variable name.

XXX.0 Number of elements in the variable (always 2)

XXX.1 Identifier (handle) for future scan access

XXX.2 Number of matching history file entries.

### Usage notes:

The combination of time stamp, object name and caller action can be used to filter records. Only records that pass all specified filters are returned.

The filtering effect of the object name depends on the value specified:

## db2HistoryOpenScan - Open History File Scan

- Specifying a table will return records for load operations, because this is the only information for tables in the history file.
- Specifying a table space will return records for backup, restore, and load operations for the table space.

**Note:** To return records for tables, they must be specified as *schema.tablename*. Specifying *tablename* will only return records for table spaces.

A maximum of eight history file scans per process is permitted.

To list every entry in the history file, a typical application will perform the following steps:

1. Call `db2HistoryOpenScan`, which will return *oNumRows*.
2. Allocate an *db2HistData* structure with space for *n oTablespace* fields, where *n* is an arbitrary number.
3. Set the *iDB2NumTablespace* field of the *db2HistData* structure to *n*.
4. In a loop, perform the following:
  - Call `db2HistoryGetEntry` to fetch from the history file.
  - If `db2HistoryGetEntry` returns an SQLCODE of `SQL_RC_OK`, use the *sqlid* field of the *db2HistData* structure to determine the number of table space entries returned.
  - If `db2HistoryGetEntry` returns an SQLCODE of `SQLUH_SQLUHINFO_VARS_WARNING`, not enough space has been allocated for all of the table spaces that DB2 is trying to return; free and reallocate the *db2HistData* structure with enough space for *oDB2UsedTablespace* table space entries, and set *iDB2NumTablespace* to *oDB2UsedTablespace*.
  - If `db2HistoryGetEntry` returns an SQLCODE of `SQLC_RC_NOMORE`, all history file entries have been retrieved.
  - Any other SQLCODE indicates a problem.
5. When all of the information has been fetched, call `db2HistoryCloseScan` to free the resources allocated by the call to `db2HistoryOpenScan`.

The macro `SQLUHINFOSIZE(n)` (defined in `sqlutil`) is provided to help determine how much memory is required for an *db2HistData* structure with space for *n oTablespace* fields.

### Related reference:

- “db2Prune - Prune History File” on page 278
- “db2HistoryUpdate - Update History File” on page 275
- “db2HistoryGetEntry - Get Next History File Entry” on page 268
- “db2HistoryCloseScan - Close History File Scan” on page 267
- “SQLCA” in the *Administrative API Reference*

### Related samples:

- “dbrecov.sqc -- How to recover a database (C)”
- “dbrecov.sqC -- How to recover a database (C++)”

## db2HistoryUpdate - Update History File

Updates the location, device type, or comment in a history file entry.

### Authorization:

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*
- *dbadm*

### Required connection:

Database. To update entries in the history file for a database other than the default database, a connection to the database must be established before calling this API.

### API include file:

*db2ApiDf.h*

### C API syntax:

```

/* File: db2ApiDf.h */
/* API: db2HistoryUpdate */
/* ... */
SQL_API_RC SQL_API_FN
db2HistoryUpdate (
    db2UInt32 version,
    void *pDB2HistoryUpdateStruct,
    struct sqlca *pSqlca);

typedef SQL_STRUCTURE db2HistoryUpdateStruct
{
    char          *piNewLocation;
    char          *piNewDeviceType;
    char          *piNewComment;
    char          *piNewStatus;
    db2HistoryEID iEID;
} db2HistoryUpdateStruct;

/* Structure db2HistoryEID */
typedef SQL_STRUCTURE db2HistoryEID
{
    SQL_PDB_NODE_TYPE ioNode;
    db2UInt32         ioHID;
} db2HistoryEID;

/* ... */

```

### Generic API syntax:

```

/* File: db2ApiDf.h */
/* API: db2gHistoryUpdate */
/* ... */
SQL_API_RC SQL_API_FN
db2GenHistoryUpdate (
    db2UInt32 version,
    void *pDB2GenHistoryUpdateStruct,
    struct sqlca *pSqlca);

typedef SQL_STRUCTURE db2gHistoryUpdateStruct

```

## db2HistoryUpdate - Update History File

```
|
|         {
|             char                *piNewLocation;
|             char                *piNewDeviceType;
|             char                *piNewComment;
|             char                *piNewStatus;
|             db2Uint32           iNewLocationLen;
|             db2Uint32           iNewDeviceLen;
|             db2Uint32           iNewCommentLen;
|             db2Uint32           iNewStatusLen;
|             db2HistoryEID       iEID;
|         } db2gHistoryUpdateStruct;
|         /* ... */
```

### API parameters:

#### version

Input. Specifies the version and release level of the structure passed in as the second parameter, *pDB2HistoryUpdateStruct*.

#### pDB2HistoryUpdateStruct

Input. A pointer to the *db2HistoryUpdateStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### piNewLocation

Input. A pointer to a string specifying a new location for the backup, restore, or load copy image. Setting this parameter to NULL, or pointing to zero, leaves the value unchanged.

#### piNewDeviceType

Input. A pointer to a string specifying a new device type for storing the backup, restore, or load copy image. Setting this parameter to NULL, or pointing to zero, leaves the value unchanged.

#### piNewComment

Input. A pointer to a string specifying a new comment to describe the entry. Setting this parameter to NULL, or pointing to zero, leaves the comment unchanged.

#### piNewStatus

Input. A pointer to a string specifying a new status type for the entry. Setting this parameter to NULL, or pointing to zero, leaves the status unchanged.

#### iNewLocationLen

Input. Length of the *piNewLocationLen* field.

#### iNewDeviceLen

Input. Length of the *piNewDeviceLen* field.

#### iNewCommentLen

Input. Length of the *piNewCommentLen* field.

#### iNewStatusLen

Input. Length of the *piNewStatusLen* field.

**iEID** Input. A unique identifier that can be used to update a specific entry in the history file.

#### ioNode

This parameter can be used as either an input or output parameter.

Indicates the node number.

**ioHID** This parameter can be used as either an input or output parameter.  
Indicates the local history file entry ID.

### REXX API syntax:

```
UPDATE RECOVERY HISTORY USING :value
```

### REXX API parameters:

**value** A compound REXX host variable containing information pertaining to the new location of a history file entry. In the following, XXX represents the host variable name:

**XXX.0** Number of elements in the variable (must be between 1 and 4)

**XXX.1** OBJECT\_PART (time stamp with a sequence number from 001 to 999)

**XXX.2** New location for the backup or copy image (this parameter is optional)

**XXX.3** New device used to store the backup or copy image (this parameter is optional)

**XXX.4** New comment (this parameter is optional).

### Usage notes:

This is an update function, and all information prior to the change is replaced and cannot be recreated. These changes are not logged.

The primary purpose of the database history file is to record information, but the data contained in the history is used directly by automatic restore operations. During any restore where the AUTOMATIC option is specified, the history of backup images and their locations will be referenced and used by the restore utility to fulfill the automatic restore request. If the automatic restore function is to be used and backup images have been relocated since they were created, it is recommended that the database history record for those images be updated to reflect the current location. If the backup image location in the database history is not updated, automatic restore will not be able to locate the backup images, but manual restore commands can still be used successfully.

### Related reference:

- “db2Rollforward - Rollforward Database” on page 131
- “db2Prune - Prune History File” on page 278
- “db2HistoryOpenScan - Open History File Scan” on page 271
- “db2HistoryGetEntry - Get Next History File Entry” on page 268
- “db2HistoryCloseScan - Close History File Scan” on page 267
- “SQLCA” in the *Administrative API Reference*
- “UPDATE HISTORY FILE” on page 263
- “db2Backup - Backup database” on page 71

### Related samples:

- “dbrecov.sqc -- How to recover a database (C)”
- “dbrecov.sqC -- How to recover a database (C++)”

### db2Prune - Prune History File

Deletes entries from the history file or log files from the active log path.

#### Authorization:

One of the following:

- *sysadm*
- *sysctrl*
- *sysmaint*
- *dbadm*

#### Required connection:

Database. To delete entries from the history file for any database other than the default database, a connection to the database must be established before calling this API.

#### API include file:

*db2ApiDf.h*

#### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2Prune */
/* ... */
SQL_API_RC SQL_API_FN
db2Prune (
    db2UInt32 version,
    void *pDB2PruneStruct,
    struct sqlca *pSqlca);
```

```
typedef struct
{
    char *piString,
    db2UInt32 iEID,
    db2UInt32 iCallerAction,
    db2UInt32 iOptions
} db2PruneStruct;
/* ... */
```

#### Generic API syntax:

```
/* File: db2ApiDf.h */
/* API: db2GenPrune */
/* ... */
SQL_API_RC SQL_API_FN
db2GenPrune (
    db2UInt32 version,
    void *pDB2GenPruneStruct,
    struct sqlca *pSqlca);
```

```
typedef struct
{
    db2UInt32 iStringLen;
    char *piString,
    db2UInt32 iEID,
```



```

    db2Uint32 iCallerAction,
    db2Uint32 iOptions
} db2GenPruneStruct;
/* ... */

```

**API parameters:****version**

Input. Specifies the version and release level of the structure passed in as the second parameter, *pDB2PruneStruct*.

**pDB2PruneStruct**

Input. A pointer to the *db2PruneStruct* structure.

**pSqlca**

Output. A pointer to the *sqlca* structure.

**iStringLen**

Input. Specifies the length in bytes of *piString*.

**piString**

Input. A pointer to a string specifying a time stamp or a log sequence number (LSN). The time stamp or part of a time stamp (minimum *yyyy*, or year) is used to select records for deletion. All entries equal to or less than the time stamp will be deleted. A valid time stamp must be provided; there is no default behavior for a NULL parameter.

This parameter can also be used to pass an LSN, so that inactive logs can be pruned.

**iEID** Input. Specifies a unique identifier that can be used to prune a single entry from the history file.

**iCallerAction**

Input. Specifies the type of action to be taken. Valid values (defined in *db2ApiDf.h*) are:

**DB2PRUNE\_ACTION\_HISTORY**

Remove history file entries.

**DB2PRUNE\_ACTION\_LOG**

Remove log files from the active log path.

**iOptions**

Input. Valid values (defined in *db2ApiDf.h*) are:

**DB2PRUNE\_OPTION\_FORCE**

Force the removal of the last backup.

**DB2PRUNE\_OPTION\_DELETE**

Delete log files that are pruned from the history file.

**DB2PRUNE\_OPTION\_LSNSTRING**

Specify that the value of *piString* is an LSN, used when a caller action of *DB2PRUNE\_ACTION\_LOG* is specified.

**REXX API syntax:**

```
PRUNE RECOVERY HISTORY BEFORE :timestamp [WITH FORCE OPTION]
```

**REXX API parameters:**

## db2Prune - Prune History File

### timestamp

A host variable containing a time stamp. All entries with time stamps equal to or less than the time stamp provided are deleted from the history file.

### WITH FORCE OPTION

If specified, the history file will be pruned according to the time stamp specified, even if some entries from the most recent restore set are deleted from the file. If not specified, the most recent restore set will be kept, even if the time stamp is less than or equal to the time stamp specified as input.

### Usage notes:

Pruning the history file does not delete the actual backup or load files. The user must manually delete these files to free up the space they consume on storage media.

### CAUTION:

**If the latest full database backup is deleted from the media (in addition to being pruned from the history file), the user must ensure that all table spaces, including the catalog table space and the user table spaces, are backed up. Failure to do so may result in a database that cannot be recovered, or the loss of some portion of the user data in the database.**

### Related reference:

- “db2HistoryUpdate - Update History File” on page 275
- “db2HistoryOpenScan - Open History File Scan” on page 271
- “db2HistoryGetEntry - Get Next History File Entry” on page 268
- “db2HistoryCloseScan - Close History File Scan” on page 267
- “SQLCA” in the *Administrative API Reference*

### Related samples:

- “dbrecov.sqc -- How to recover a database (C)”
- “dbrecov.sqC -- How to recover a database (C++)”

---

## db2ReadLogNoConn - Read Log Without a Database Connection

Extract log records from the DB2 UDB database logs and query the Log Manager for current log state information. Prior to using this API, use db2ReadLogNoConnInit to allocate the memory that is passed as an input parameter to this API. After using this API, use db2ReadLogNoConnTerm to deallocate the memory.

### Authorization:

None

### Required connection:

None

### API include file:

*db2ApiDf.h*

## db2ReadLogNoConn - Read Log Without a Database Connection

### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2ReadLogNoConn */
/* ... */
SQL_API_RC SQL_API_FN
db2ReadLogNoConn (
    db2UInt32 versionNumber,
    void *pDB2ReadLogNoConnStruct,
    struct sqlca *pSqlca);

typedef SQL_STRUCTURE db2ReadLogNoConnStruct
{
    db2UInt32          iCallerAction;
    SQLU_LSN          *piStartLSN;
    SQLU_LSN          *piEndLSN;
    char              *poLogBuffer;
    db2UInt32          iLogBufferSize;
    char              *piReadLogMemPtr;
    db2ReadLogNoConnInfoStruct *poReadLogInfo;
} db2ReadLogNoConnStruct;

typedef SQL_STRUCTURE db2ReadLogNoConnInfoStruct
{
    SQLU_LSN          firstAvailableLSN;
    SQLU_LSN          firstReadLSN;
    SQLU_LSN          nextStartLSN;
    db2UInt32          logRecsWritten;
    db2UInt32          logBytesWritten;
    db2UInt32          lastLogFullyRead;
    db2TimeOfLog       currentTimeValue;
} db2ReadLogNoConnInfoStruct;

/* ... */
```

### API parameters:

#### versionNumber

Input. Specifies the version and release level of the structure passed as the second parameter *pDB2ReadLogNoConnStruct*.

#### pDB2ReadLogNoConnStruct

Input. A pointer to the *db2ReadLogNoConnStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### iCallerAction

Input. Specifies the action to be performed. Valid values are:

##### DB2READLOG\_READ

Read the database log from the starting log sequence to the ending log sequence number and return log records within this range.

##### DB2READLOG\_READ\_SINGLE

Read a single log record (propagatable or not) identified by the starting log sequence number.

##### DB2READLOG\_QUERY

Query the database log. Results of the query will be sent back via the *db2ReadLogNoConnInfoStruct* structure.

#### piStartLSN

Input. The starting log sequence number specifies the starting relative byte address for the reading of the log. This value must be the start of an actual log record.

## db2ReadLogNoConn - Read Log Without a Database Connection

### **piEndLSN**

Input. The ending log sequence number specifies the ending relative byte address for the reading of the log. This value must be greater than *piStartLsn*, and does not need to be the end of an actual log record.

### **poLogBuffer**

Output. The buffer where all the propagatable log records read within the specified range are stored sequentially. This buffer must be large enough to hold a single log record. As a guideline, this buffer should be a minimum of 32 bytes. Its maximum size is dependent on the size of the requested range. Each log record in the buffer is prefixed by a six byte log sequence number (LSN), representing the LSN of the following log record.

### **iLogBufferSize**

Input. Specifies the size, in bytes, of the log buffer.

### **piReadLogMemPtr**

Input. Block of memory of size *iReadLogMemoryLimit* that was allocated in the initialization call. This memory contains persistent data that the API requires at each invocation. This memory block must not be reallocated or altered in any way by the caller.

### **poReadLogInfo**

Output. A pointer to the *db2ReadLogNoConnInfoStruct* structure.

### **firstAvailableLSN**

First available LSN in available logs.

### **firstReadLSN**

First LSN read on this call.

### **nextStartLSN**

Next readable LSN.

### **logRecsWritten**

Number of log records written to the log buffer field, *poLogBuffer*.

### **logBytesWritten**

Number of bytes written to the log buffer field, *poLogBuffer*.

### **lastLogFullyRead**

Number indicating the last log file that was read to completion.

### **Usage notes:**

The *db2ReadLogNoConn* API requires a memory block that must be allocated using the *db2ReadLogNoConnInit* API. The memory block must be passed as an input parameter to all subsequent *db2ReadLogNoConn* API calls, and must not be altered.

When requesting a sequential read of log, the API requires a log sequence number (LSN) range and the allocated memory. The API will return a sequence of log records based on the filter option specified when initialized and the LSN range. When requesting a query, the read log information structure will contain a valid starting LSN, to be used on a read call. The value used as the ending LSN on a read can be one of the following:

- A value greater than the caller-specified startLSN.
- FFFF FFFF FFFF which is interpreted by the asynchronous log reader as the end of the available logs.

## db2ReadLogNoConn - Read Log Without a Database Connection

The propagatable log records read within the starting and ending LSN range are returned in the log buffer. A log record does not contain its LSN, it is contained in the buffer before the actual log record. Descriptions of the various DB2 UDB log records returned by db2ReadLogNoConn can be found in the DB2 UDB Log Records section.

After the initial read, in order to read the next sequential log record, use the nextStartLSN value returned in db2ReadLogNoConnInfoStruct. Resubmit the call, with this new starting LSN and a valid ending LSN and the next block of records is then read. An sqlca code of SQLU\_RLOG\_READ\_TO\_CURRENT means the log reader has read to the end of the available log files.

When the API will no longer be used, use db2ReadLogNoConnTerm to terminate the memory.

### Related reference:

- “SQLCA” in the *Administrative API Reference*
- “db2ReadLogNoConnInit - Initialize Read Log Without a Database Connection” on page 283
- “db2ReadLogNoConnTerm - Terminate Read Log Without a Database Connection” on page 285

---

## db2ReadLogNoConnInit - Initialize Read Log Without a Database Connection

Allocates the memory to be used by db2ReadLogNoConn in order to extract log records from the DB2 UDB database logs and query the Log Manager for current log state information.

### Authorization:

None

### Required connection:

None

### API include file:

*db2ApiDf.h*

### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2ReadLogNoConnInit */
/* ... */
SQL_API_RC SQL_API_FN
db2ReadLogNoConnInit (
    db2UInt32 versionNumber,
    void * pDB2ReadLogNoConnInitStruct,
    struct sqlca * pSqlca);

typedef SQL_STRUCTURE db2ReadLogNoConnInitStruct
{
    db2UInt32          iFilterOption;
    char               *piLogFilePath;
    char               *piOverflowLogPath;
```

## db2ReadLogNoConnInit - Initialize Read Log Without a Database Connection

```
    db2UInt32          iRetrieveLogs;
    char              *piDatabaseName;
    char              *piNodeName;
    db2UInt32          iReadLogMemoryLimit;
    char              **poReadLogMemPtr;
} db2ReadLogNoConnInitStruct;
/* ... */
```

### API parameters:

#### versionNumber

Input. Specifies the version and release level of the structure passed as the second parameter *pDB2ReadLogNoConnInitStruct*.

#### pDB2ReadLogNoConnInitStruct

Input. A pointer to the *db2ReadLogNoConnInitStruct* structure.

#### pSqlca

Output. A pointer to the *sqlca* structure.

#### iFilterOption

Input. Specifies the level of log record filtering to be used when reading the log records. Valid values are:

##### DB2READLOG\_FILTER\_OFF

Read all log records in the given LSN range.

##### DB2READLOG\_FILTER\_ON

Reads only log records in the given LSN range marked as propagatable. This is the traditional behavior of the asynchronous log read API.

#### piLogFilePath

Input. Path where the log files to be read are located.

#### piOverflowLogPath

Input. Alternate path where the log files to be read may be located.

#### iRetrieveLogs

Input. Option specifying if *userexit* should be invoked to retrieve log files that cannot be found in either the log file path or the overflow log path. Valid values are:

##### DB2READLOG\_RETRIEVE\_OFF

*Userexit* should not be invoked to retrieve missing log files.

##### DB2READLOG\_RETRIEVE\_LOGPATH

*Userexit* should be invoked to retrieve missing log files into the specified log file path.

##### DB2READLOG\_RETRIEVE\_OVERFLOW

*Userexit* should be invoked to retrieve missing log files into the specified overflow log path.

#### piDatabaseName

Input. Name of the database that owns the recovery logs being read. This is required if the retrieve option above is specified.

#### piNodeName

Input. Name of the node that owns the recovery logs being read. This is required if the retrieve option above is specified.

#### iReadLogMemoryLimit

Input. Maximum number of bytes that the API may allocate internally.

## db2ReadLogNoConnInit - Initialize Read Log Without a Database Connection

### poReadLogMemPtr

Output. API-allocated block of memory of size *iReadLogMemoryLimit*. This memory contains persistent data that the API requires at each invocation. This memory block must not be reallocated or altered in any way by the caller.

### Usage notes:

The memory initialized by `db2ReadLogNoConnInit` must not be altered.

When `db2ReadLogNoConn` will no longer be used, invoke `db2ReadLogNoConnTerm` to deallocate the memory initialized by `db2ReadLogNoConnInit`.

### Related reference:

- “SQLCA” in the *Administrative API Reference*
- “`db2ReadLogNoConn` - Read Log Without a Database Connection” on page 280
- “`db2ReadLogNoConnTerm` - Terminate Read Log Without a Database Connection” on page 285

---

## db2ReadLogNoConnTerm - Terminate Read Log Without a Database Connection

Deallocates the memory used by `db2ReadLogNoConn`, originally initialized by `db2ReadLogNoConnInit`.

### Authorization:

None

### Required connection:

None

### API include file:

*db2ApiDf.h*

### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2ReadLogNoConnTerm */
/* ... */
SQL_API_RC SQL_API_FN
db2ReadLogNoConnTerm (
    db2UInt32 versionNumber,
    void * pDB2ReadLogNoConnTermStruct,
    struct sqlca * pSqlca);

typedef SQL_STRUCTURE db2ReadLogNoConnTermStruct
{
    char **poReadLogMemPtr;
} db2ReadLogNoConnTermStruct;
/* ... */
```

### API parameters:

## db2ReadLogNoConnTerm - Terminate Read Log Without a Database Connection

### versionNumber

Input. Specifies the version and release level of the structure passed as the second parameter *pDB2ReadLogNoConnTermStruct*.

### pDB2ReadLogNoConnTermStruct

Input. A pointer to the *db2ReadLogNoConnTermStruct* structure.

### pSqlca

Output. A pointer to the *sqlca* structure.

### poReadLogMemPtr

Output. Pointer to the block of memory allocated in the initialization call. This pointer will be freed and set to NULL.

### Related reference:

- “SQLCA” in the *Administrative API Reference*
- “db2ReadLogNoConn - Read Log Without a Database Connection” on page 280
- “db2ReadLogNoConnInit - Initialize Read Log Without a Database Connection” on page 283

---

## db2ReadLog - Asynchronous Read Log

Extract log records from the DB2 UDB database logs and the Log Manager for current log state information. This API can only be used with recoverable databases. A database is recoverable if it is configured with *logretain* set to RECOVERY or *userexit* set to 0N.

### Authorization:

One of the following:

- *sysadm*
- *dbadm*

### Required connection:

Database

### API include file:

*db2ApiDf.h*

### C API syntax:

```
/* File: db2ApiDf.h */
/* API: db2ReadLog */
/* ... */
SQL_API_RC SQL_API_FN
db2ReadLog (
    db2UInt32 versionNumber,
    void *pDB2ReadLogStruct,
    struct sqlca *pSqlca);

typedef SQL_STRUCTURE db2ReadLogStruct
{
    db2UInt32          iCallerAction;
    SQLU_LSN          *piStartLSN;
    SQLU_LSN          *piEndLSN;
    char              *poLogBuffer;
    db2UInt32          iLogBufferSize;
```



```

    db2UInt32          iFilterOption;
    db2ReadLogInfoStruct *poReadLogInfo;

typedef SQL_STRUCTURE db2ReadLogInfoStruct
{
    SQLU_LSN          initialLSN;
    SQLU_LSN          firstReadLSN;
    SQLU_LSN          nextStartLSN;
    db2UInt32         logRecsWritten;
    db2UInt32         logBytesWritten;
    SQLU_LSN          firstReusedLSN;
    db2UInt32         timeOfLSNReuse;
    db2TimeOfLog      currentTimeValue;
} db2ReadLogInfoStruct;

typedef SQL_STRUCTURE db2TimeOfLog
{
    db2UInt32         seconds;
    db2UInt32         accuracy;
} db2TimeOfLog;
/* ... */

```

**API parameters:****versionNumber**

Input. Specifies the version and release level of the structure passed as the second parameter, *pDB2ReadLogStruct*.

**pDB2ReadLogStruct**

Input. A pointer to the *db2ReadLogStruct*.

**pSqlca**

Output. A pointer to the *sqlca* structure.

**iCallerAction**

Input. Specifies the action to be performed.

**DB2READLOG\_READ**

Read the database log from the starting log sequence to the ending log sequence number and return log records within this range.

**DB2READLOG\_READ\_SINGLE**

Read a single log record (propagatable or not) identified by the starting log sequence number.

**DB2READLOG\_QUERY**

Query the database log. Results of the query will be sent back via the *db2ReadLogInfoStruct* structure.

**piStartLsn**

Input. The starting log sequence number specifies the starting relative byte address for the reading of the log. This value must be the start of an actual log record.

**piEndLsn**

Input. The ending log sequence number specifies the ending relative byte address for the reading of the log. This value must be greater than *startLsn*, and does not need to be the end of an actual log record.

**poLogBuffer**

Output. The buffer where all the propagatable log records read within the specified range are stored sequentially. This buffer must be large enough to hold a single log record. As a guideline, this buffer should be a minimum of 32 bytes. Its maximum size is dependent on the size of the requested

## db2ReadLog - Asynchronous Read Log

range. Each log record in the buffer is prefixed by a six byte log sequence number (LSN), representing the LSN of the following log record.

### **iLogBufferSize**

Input. Specifies the size, in bytes, of the log buffer.

### **iFilterOption**

Input. Specifies the level of log record filtering to be used when reading the log records. Valid values are:

#### **DB2READLOG\_FILTER\_OFF**

Read all log records in the given LSN range.

#### **DB2READLOG\_FILTER\_ON**

Reads only log records in the given LSN range marked as propagatable. This is the traditional behaviors of the asynchronous log read API.

### **poReadLogInfo**

Output. A structure detailing information regarding the call and the database log.

### **Usage notes:**

If the requested action is to read the log, the caller will provide a log sequence number range and a buffer to hold the log records. This API reads the log sequentially, bounded by the requested LSN range, and returns log records associated with tables having the DATA CAPTURE option CHANGES, and a db2ReadLogInfoStruct structure with the current active log information. If the requested action is query, the API returns an db2ReadLogInfoStruct structure with the current active log information.

To use the Asynchronous Log Reader, first query the database log for a valid starting LSN. Following the query call, the read log information structure (db2ReadLogInfoStruct) will contain a valid starting LSN (in the initialLSN member), to be used on a read call. The value used as the ending LSN on a read can be one of the following:

- A value greater than initialLSN
- FFFF FFFF FFFF, which is interpreted by the asynchronous log reader as the end of the current log.

The propagatable log records read within the starting and ending LSN range are returned in the log buffer. A log record does not contain its LSN; it is contained in the buffer before the actual log record. Descriptions of the various DB2 log records returned by db2ReadLog the DB2 UDB Log Records section.

To read the next sequential log record after the initial read, use the nextStartLSN field returned in the db2ReadLogStruct structure. Resubmit the call, with this new starting LSN and a valid ending LSN. The next block of records is then read. An *sqlca* code of `SQLU_RLOG_READ_TO_CURRENT` means that the log reader has read to the end of the current active log.

### **Related reference:**

- “SQLCA” in the *Administrative API Reference*
- “db2Reorg - Reorganize” in the *Administrative API Reference*

### **Related samples:**

- “dbrecov.sqc -- How to recover a database (C)”
- “dbrecov.sqC -- How to recover a database (C++)”

---

## db2HistData

This structure is used to return information after a call to db2HistoryGetEntry.

Table 2. Fields in the db2HistData Structure

Field Name	Data Type	Description
ioHistDataID	char(8)	An 8-byte structure identifier and “eye-catcher” for storage dumps. The only valid value is “SQLUHINF”. No symbolic definition for this string exists.
oObjectPart	db2Char	The first 14 characters are a time stamp with format <i>yyyymmddhhmss</i> , indicating when the operation was begun. The next 3 characters are a sequence number. Each backup operation can result in multiple entries in this file when the backup image is saved in multiple files or on multiple tapes. The sequence number allows multiple locations to be specified. Restore and load operations have only a single entry in this file, which corresponds to sequence number ‘001’ of the corresponding backup. The time stamp, combined with the sequence number, must be unique.
oEndTime	db2Char	A time stamp with format <i>yyyymmddhhmss</i> , indicating when the operation was completed.
oFirstLog	db2Char	The earliest log file ID (ranging from S0000000 to S9999999): <ul style="list-style-type: none"> <li>• Required to apply rollforward recovery for an online backup</li> <li>• Required to apply rollforward recovery for an offline backup</li> <li>• Applied after restoring a full database or table space level backup that was current when the load started.</li> </ul>
oLastLog	db2Char	The latest log file ID (ranging from S0000000 to S9999999): <ul style="list-style-type: none"> <li>• Required to apply rollforward recovery for an online backup</li> <li>• Required to apply rollforward recovery to the current point in time for an offline backup</li> <li>• Applied after restoring a full database or table space level backup that was current when the load operation finished (will be the same as <i>oFirstLog</i> if roll forward recovery is not applied).</li> </ul>
oID	db2Char	Unique backup or table identifier.
oTableQualifier	db2Char	Table qualifier.
oTableName	db2Char	Table name.

Table 2. Fields in the db2HistData Structure (continued)

Field Name	Data Type	Description
oLocation	db2Char	For backups and load copies, this field indicates where the data has been saved. For operations that require multiple entries in the file, the sequence number defined by <i>oObjectPart</i> identifies which part of the backup is found in the specified location. For restore and load operations, the location always identifies where the first part of the data restored or loaded (corresponding to sequence '001' for multi-part backups) has been saved. The data in <i>oLocation</i> is interpreted differently, depending on <i>oDeviceType</i> : <ul style="list-style-type: none"> <li>• For disk or diskette (D or K), a fully qualified file name</li> <li>• For tape (T), a volume label</li> <li>• For TSM (A), the server name</li> <li>• For user exit or other (U or O), free form text.</li> </ul>
oComment	db2Char	Free form text comment.
oCommandText	db2Char	Command text, or DDL.
oLastLSN	SQLU_LSN	Last log sequence number.
oEID	Structure	Unique entry identifier.
poEventSQLCA	Structure	Result <i>sqlca</i> of the recorded event.
poTablespace	db2Char	A list of table space names.
ioNumTablespaces	db2UInt32	Number of entries in the <i>poTablespace</i> list. Each table space backup contains one or more table spaces. Each table space restore operation replaces one or more table spaces. If this field is not zero (indicating a table space level backup or restore), the next lines in this file contain the name of the table space backed up or restored, represented by an 18-character string. One table space name appears on each line.
oOperation	char	See Table 3.
oObject	char	Granularity of the operation: D for full database, P for table space, and T for table.
oOtype	char	See Table 4 on page 291.
oStatus	char	Entry status: A for action, D for deleted (future use), E for expired, I for inactive, N for not yet committed, Y for committed or active, a for active backup, but some datalink servers have not yet completed the backup, and i for inactive backup, but some datalink servers have not yet completed the backup.
oDeviceType	char	Device type. This field determines how the <i>oLocation</i> field is interpreted: A for TSM, C for client, D for disk, K for diskette, L for local, O for other (for other vendor device support), P for pipe, Q for cursor, S for server, T for tape, and U for user exit.

Table 3. Valid oOperation Values in the db2HistData Structure

Value	Description	C Definition	COBOL/FORTRAN Definition
A	add table space	DB2HISTORY_OP_ADD_TABLESPACE	DB2HIST_OP_ADD_TABLESPACE
B	backup	DB2HISTORY_OP_BACKUP	DB2HIST_OP_BACKUP
C	load copy	DB2HISTORY_OP_LOAD_COPY	DB2HIST_OP_LOAD_COPY

Table 3. Valid oOperation Values in the db2HistData Structure (continued)

Value	Description	C Definition	COBOL/FORTRAN Definition
D	dropped table	DB2HISTORY_OP_DROPPED_ TABLE	DB2HIST_OP_DROPPED_TABLE
F	rollforward	DB2HISTORY_OP_ROLLFWD	DB2HIST_OP_ROLLFWD
G	reorganize table	DB2HISTORY_OP_REORG	DB2HIST_OP_REORG
L	load	DB2HISTORY_OP_LOAD	DB2HIST_OP_LOAD
N	rename table space	DB2HISTORY_OP_REN_ TABLESPACE	DB2HIST_OP_REN_ TABLESPACE
O	drop table space	DB2HISTORY_OP_DROP_ TABLESPACE	DB2HIST_OP_DROP_ TABLESPACE
Q	quiesce	DB2HISTORY_OP_QUIESCE	DB2HIST_OP_QUIESCE
R	restore	DB2HISTORY_OP_RESTORE	DB2HIST_OP_RESTORE
T	alter table space	DB2HISTORY_OP_ALT_ TABLESPACE	DB2HIST_OP_ALT_TBS
U	unload	DB2HISTORY_OP_UNLOAD	DB2HIST_OP_UNLOAD

Table 4. Valid oOtype Values in the db2HistData Structure

oOperation	oOtype	Description	C/COBOL/FORTRAN Definition
B	F	offline	DB2HISTORY_OPTYPE_OFFLINE
	N	online	DB2HISTORY_OPTYPE_ONLINE
	I	incremental offline	DB2HISTORY_OPTYPE_INCR_ OFFLINE
	O	incremental online	DB2HISTORY_OPTYPE_INCR_ ONLINE
	D	delta offline	DB2HISTORY_OPTYPE_DELTA_ OFFLINE
	E	delta online	DB2HISTORY_OPTYPE_DELTA_ ONLINE
F	E	end of logs	DB2HISTORY_OPTYPE_EOL
	P	point in time	DB2HISTORY_OPTYPE_PIT
G	F	offline	DB2HISTORY_OPTYPE_OFFLINE
	N	online	DB2HISTORY_OPTYPE_ONLINE
L	I	insert	DB2HISTORY_OPTYPE_INSERT
	R	replace	DB2HISTORY_OPTYPE_REPLACE
Q	S	quiesce share	DB2HISTORY_OPTYPE_SHARE
	U	quiesce update	DB2HISTORY_OPTYPE_UPDATE
	X	quiesce exclusive	DB2HISTORY_OPTYPE_EXCL
	Z	quiesce reset	DB2HISTORY_OPTYPE_RESET
R	F	offline	DB2HISTORY_OPTYPE_OFFLINE
	N	online	DB2HISTORY_OPTYPE_ONLINE
	I	incremental offline	DB2HISTORY_OPTYPE_INCR_ OFFLINE
	O	incremental online	DB2HISTORY_OPTYPE_INCR_ ONLINE
T	C	add containers	DB2HISTORY_OPTYPE_ADD_CONT
	R	rebalance	DB2HISTORY_OPTYPE_REB

Table 5. Fields in the db2Char Structure

Field Name	Data Type	Description
pioData	char	A pointer to a character data buffer. If NULL, no data will be returned.
iLength	db2UInt32	Input. The size of the <i>pioData</i> buffer.
oLength	db2UInt32	Output. The number of valid characters of data in the <i>pioData</i> buffer.

Table 6. Fields in the db2HistoryEID Structure

Field Name	Data Type	Description
ioNode	SQL_PDB_NODE_TYPE	Node number.
ioHID	db2UInt32	Local history file entry ID.

**Language syntax:****C Structure**

```

/* File: db2ApiDf.h */
/* ... */
typedef SQL_STRUCTURE db2HistoryData
{
    char ioHistDataID[8];
    db2Char oObjectPart;
    db2Char oEndTime;
    db2Char oFirstLog;
    db2Char oLastLog;
    db2Char oID;
    db2Char oTableQualifier;
    db2Char oTableName;
    db2Char oLocation;
    db2Char oComment;
    db2Char oCommandText;
    SQLU_LSN oLastLSN;
    db2HistoryEID oEID;
    struct sqlca * poEventSQLCA;
    db2Char * poTablespace;
    db2UInt32 ioNumTablespaces;
    char oOperation;
    char oObject;
    char oOptype;
    char oStatus;
    char oDeviceType
} db2HistoryData;

typedef SQL_STRUCTURE db2Char
{
    char * pioData;
    db2UInt32 ioLength
} db2Char;

typedef SQL_STRUCTURE db2HistoryEID
{
    SQL_PDB_NODE_TYPE ioNode;
    db2UInt32 ioHID
} db2HistoryEID;
/* ... */

```

**Related reference:**

- “db2HistoryGetEntry - Get Next History File Entry” on page 268
- “SQLCA” in the *Administrative API Reference*

## SQLU-LSN

This union, used by the `db2ReadLog` API, contains the definition of the log sequence number. A log sequence number (LSN) represents a relative byte address within the database log. All log records are identified by this number. It represents the log record's byte offset from the beginning of the database log.

*Table 7. Fields in the SQLU-LSN Union*

Field Name	Data Type	Description
<code>lsnChar</code>	Array of UNSIGNED CHAR	Specifies the 6-member character array log sequence number.
<code>lsnWord</code>	Array of UNSIGNED SHORT	Specifies the 3-member short array log sequence number.

### Language syntax:

#### C Structure

```
typedef union SQLU_LSN
{
  unsigned char lsnChar [6] ;
  unsigned short lsnWord [3] ;
} SQLU_LSN;
```

### Related reference:

- “`db2ReadLog` - Asynchronous Read Log” on page 286





---

## Appendix E. Recovery sample programs

---

### Sample programs with embedded SQL

The following sample programs show how to use DB2 backup and restore APIs to:

- Read and update a database recovery file entry
- Read database log files with a database connection
- Read database log files with no database connection
- Restore a database from a backup image
- Perform a rollforward operation after a database restore operation

**Note:** These sample files can be found in `sqlib/samples/c` and `sqlib/samples/cpp` directory.

#### dbredirect sample program:

The `dbredirect` sample files show how to perform a redirected restore of database.

```

/*****
** Licensed Materials - Property of IBM
**
** Governed under the terms of the International
** License Agreement for Non-Warranted Sample Code.
**
** (C) COPYRIGHT International Business Machines Corp. 2003
** All Rights Reserved.
**
** US Government Users Restricted Rights - Use, duplication or
** disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
*****/
**
** SOURCE FILE NAME: dbredirect.sqc
**
** SAMPLE: How to perform Redirected Restore of a database
**
**      This program ends in ".sqc" even though it does not contain
**      embedded SQL statements. It links in the embedded SQL utility
**      file for database connection and disconnection, so it needs the
**      embedded SQL extension for the precompiler.
**
**      Note:
**      You must be disconnected from the sample database to run
**      this program. To ensure you are, enter 'db2 connect reset'
**      on the command line prior to running dbredirect. If the target
**      database for the redirected restore already exists, SQLCODE 2529
**      will be displayed.
**
** DB2 API USED:
**      db2CfgSet -- Set Configuration
**      db2Restore -- Restore Database
**
** OUTPUT FILE: dbredirect.out (available in the online documentation)
*****/
**
** For detailed information about database backup and database recovery, see
** the Data Recovery and High Availability Guide and Reference. This manual
** will help you to determine which database and table space recovery methods
** are best suited to your business environment.
```

## Sample Programs with embedded SQL

```
**
** For more information on the sample programs, see the README file.
**
** For information on developing C applications, see the Application
** Development Guide.
**
** For information on using SQL statements, see the SQL Reference.
**
** For information on DB2 APIs, see the Administrative API Reference.
**
** For the latest information on programming, building, and running DB2
** applications, visit the DB2 application development website:
**   http://www.software.ibm.com/data/db2/udb/ad
**/*****
#include "utilrecov.c"

/* local function prototypes */
int DbBackupAndRedirectedRestore(char *, char *, char *, char *, char *);

/* support function called by DbBackupAndRedirectedRestore() */
int InaccessibleContainersRedefine(char *);

int main(int argc, char *argv[])
{
    int rc = 0;
    char nodeName[SQL_INSTNAME_SZ + 1] = { 0 };
    char serverWorkingPath[SQL_PATH_SZ + 1] = { 0 };
    char redirectedRestoredDbAlias[SQL_ALIAS_SZ + 1] = { 0 };
    char dbAlias[SQL_ALIAS_SZ + 1] = { 0 };
    char user[USERID_SZ + 1] = { 0 };
    char pswd[PSWD_SZ + 1] = { 0 };

    /* check the command line arguments */
    rc = CmdLineArgsCheck3(argc, argv, dbAlias, nodeName, user, pswd);
    CHECKRC(rc, "CmdLineArgsCheck3");

    printf("\nTHIS SAMPLE SHOWS HOW TO PERFORM A REDIRECTED RESTORE\n");
    printf("FROM A DATABASE BACKUP.\n");

    strcpy(redirectedRestoredDbAlias, "RRDB");

    /* attach to a local or remote instance */
    rc = InstanceAttach(nodeName, user, pswd);
    CHECKRC(rc, "Instance Attach");

    /* get the server working path */
    rc = ServerWorkingPathGet(dbAlias, serverWorkingPath);
    CHECKRC(rc, "ServerWorkingPathGet");

    printf("\nNOTE: Backup images will be created on the server\n");
    printf("      in the directory %s,\n", serverWorkingPath);
    printf("      and will not be deleted by the program.\n");

    /* call the sample function */
    rc = DbRecoveryHistoryFilePrune(dbAlias, user, pswd);
    CHECKRC(rc, "DbRecoveryHistoryFilePrune");

    rc = DbBackupAndRedirectedRestore(dbAlias,
                                     redirectedRestoredDbAlias,
                                     user, pswd, serverWorkingPath);
    CHECKRC(rc, "DbBackupAndRedirectedRestore");

    /* Detach from the local or remote instance */
    rc = InstanceDetach(nodeName);
    CHECKRC(rc, "InstanceDetach");

    return 0;
}
```

```

} /* end main */

int DbBackupAndRedirectedRestore(char dbAlias[],
                                char restoredDbAlias[],
                                char user[],
                                char pswd[], char serverWorkingPath[])
{
    int rc = 0;
    struct sqlca sqlca;
    db2CfgParam cfgParameters[1];
    db2Cfg cfgStruct;
    unsigned short logretain = 0;

    char restoreTimestamp[SQLU_TIME_STAMP_LEN + 1] = { 0 };

    db2BackupStruct backupStruct;
    db2TablespaceStruct tablespaceStruct;
    db2MediaListStruct mediaListStruct;
    db2UInt32 backupImageSize = 0;
    db2RestoreStruct restoreStruct;
    db2TablespaceStruct rtablespaceStruct;
    db2MediaListStruct rmediaListStruct;

    printf("\n*****\n");
    printf("*** REDIRECTED RESTORE ***\n");
    printf("*****\n");
    printf("\nUSE THE DB2 APIs:\n");
    printf(" db2CfgSet -- Update Configuration\n");
    printf(" db2Backup -- Backup Database\n");
    printf(" sqlcrea -- Create Database\n");
    printf(" db2Restore -- Restore Database\n");
    printf(" sqlbmtsq -- Tablespace Query\n");
    printf(" sqlbtcq -- Tablespace Container Query\n");
    printf(" sqlbstsc -- Set Tablespace Containers\n");
    printf(" sqlfmem -- Free Memory\n");
    printf(" sqledrpd -- Drop Database\n");
    printf("TO BACK UP AND DO A REDIRECTED RESTORE OF A DATABASE.\n");

    printf("\n Update \'%s\' database configuration:\n", dbAlias);
    printf(" - Disable the database configuration parameter LOGRETAIN\n");
    printf(" i.e., set LOGRETAIN = OFF/NO\n");

    /* initialize cfgParameters */
    /* SQLF_DBTN_LOG_RETAIN is a token of the updatable database configuration
       parameter 'logretain'; it is used to update the database configuration
       file */
    cfgParameters[0].flags = 0;
    cfgParameters[0].token = SQLF_DBTN_LOG_RETAIN;
    cfgParameters[0].ptrvalue = (char *)&logretain;

    /* disable the database configuration parameter 'logretain' */
    logretain = SQLF_LOGRETAIN_DISABLE;

    /* initialize cfgStruct */
    cfgStruct.numItems = 1;
    cfgStruct.paramArray = cfgParameters;
    cfgStruct.flags = db2CfgDatabase | db2CfgDelayed;
    cfgStruct.dbname = dbAlias;

    /* get database configuration */
    db2CfgSet(db2Version810, (void *)&cfgStruct, &sqlca);
    DB2_API_CHECK("Db Log Retain -- Disable");

    /*****
    /* BACK UP THE DATABASE */
    *****/

```

## Sample Programs with embedded SQL

```
/* Calling up the routine for database backup */
rc = DbBackup(dbAlias, user, pswd, serverWorkingPath, &backupStruct);
CHECKRC(rc, "DbBackup");

/*****
/*   RESTORE THE DATABASE   */
*****/

strcpy(restoreTimestamp, backupStruct.oTimestamp);

rtablespaceStruct tablespaces = NULL;
rtablespaceStruct numTablespaces = 0;

rmediaListStruct.locations = &serverWorkingPath;
rmediaListStruct.numLocations = 1;
rmediaListStruct.locationType = SQLU_LOCAL_MEDIA;

restoreStruct.piSourceDBAlias = dbAlias;
restoreStruct.piTargetDBAlias = restoredDbAlias;
restoreStruct.piTimestamp = restoreTimestamp;
restoreStruct.piTargetDBPath = NULL;
restoreStruct.piReportFile = NULL;
restoreStruct.piTablespaceList = &rtablespaceStruct;
restoreStruct.piMediaList = &rmediaListStruct;
restoreStruct.piUsername = user;
restoreStruct.piPassword = pswd;
restoreStruct.piNewLogPath = NULL;
restoreStruct.piVendorOptions = NULL;
restoreStruct.iVendorOptionsSize = 0;
restoreStruct.iParallelism = 1;
restoreStruct.iBufferSize = 1024; /* 1024 x 4KB */
restoreStruct.iNumBuffers = 2;
restoreStruct.iOptions = DB2RESTORE_OFFLINE | DB2RESTORE_DB |
    DB2RESTORE_NODATALINK | DB2RESTORE_NOROLLFWD;

printf("\n Restoring a database ...\n");
printf(" - source image alias      : %s\n", dbAlias);
printf(" - source image time stamp: %s\n", restoreTimestamp);
printf(" - target database         : %s\n", restoredDbAlias);

restoreStruct.iCallerAction = DB2RESTORE_RESTORE_STORDEF;

/* The API db2Restore is used to restore a database that has been backed
   up using the API db2Backup. */
db2Restore(db2Version810, &restoreStruct, &sqlca);

/* If restoring to a different database and restoreDbAlias already exists,
   SQLCODE 2529 is expected. */
if (strcmp(dbAlias, restoredDbAlias))
{
    printf("\n SQLCODE 2529 is expected if target database '%s' already exists\n",
        restoredDbAlias);
}

EXPECTED_WARN_CHECK("database restore -- start");

while (sqlca.sqlcode != 0)
{
    /* continue the restore operation */
    printf("\n Continuing the restore operation...\n");

    /* depending on the sqlca.sqlcode value, user action may be
       required, such as mounting a new tape */

    if (sqlca.sqlcode == SQLUD_INACCESSABLE_CONTAINER)
    {
        /* redefine the table space container layout */
    }
}
```

```

        printf("\n Find and redefine inaccessible containers.\n");
        rc = InaccessibleContainersRedefine(serverWorkingPath);
        CHECKRC(rc, "InaccessibleContainersRedefine");
    }

    restoreStruct.iCallerAction = DB2RESTORE_CONTINUE;

    /* restore the database */
    db2Restore(db2Version810, &restoreStruct, &sqlca);
    DB2_API_CHECK("database restore -- continue");
}

printf("\n Restore finished.\n");

/* drop the restored database */
rc = DbDrop(restoredDbAlias);
CHECKRC(rc, "DbDrop");

return 0;
} /* DbBackupAndRedirectedRestore */

int InaccessibleContainersRedefine(char serverWorkingPath[])
{
    struct sqlca sqlca;
    sqluint32 numTablespaces = 0;
    struct SQLB_TBSPQRY_DATA **ppTablespaces = NULL;
    sqluint32 numContainers = 0;
    struct SQLB_TBSCONTQRY_DATA *pContainers = NULL;
    int tspNb = 0;
    int contNb = 0;
    char pathSep[2] = { 0 };

    /* The API sqlbmtsq provides a one-call interface to the table space query
       data. The query data for all table spaces in the database is returned
       in an array. */
    sqlbmtsq(&sqlca,
             &numTablespaces, &ppTablespaces, SQLB_RESERVED1, SQLB_RESERVED2);
    DB2_API_CHECK("tablespaces -- get");

    /* refeedine the inaccessible containers */
    for (tspNb = 0; tspNb < numTablespaces; tspNb++)
    {
        /* The API sqlbtcq provides a one-call interface to the table space
           container query data. The query data for all the containers in a table
           space, or for all containers in all table spaces, is returned in an
           array. */
        sqlbtcq(&sqlca, ppTablespaces[tspNb]->id, &numContainers, &pContainers);
        DB2_API_CHECK("tablespace containers -- get");

        for (contNb = 0; contNb < numContainers; contNb++)
        {
            if (!pContainers[contNb].ok)
            {
                /* redefine inaccessible container */
                printf("\n   Redefine inaccessible container:\n");
                printf("       - table space name: %s\n", ppTablespaces[tspNb]->name);
                printf("       - default container name: %s\n",
                    pContainers[contNb].name);
                if (strstr(pContainers[contNb].name, "/"))
                {
                    /* UNIX */
                    strcpy(pathSep, "/");
                }
                else
                {
                    /* Intel */
                    strcpy(pathSep, "\\");
                }
                switch (pContainers[contNb].contType)

```

## Sample Programs with embedded SQL

```
        {
        case SQLB_CONT_PATH:
            printf("        - container type: path\n");

            sprintf(pContainers[contNb].name, "%s%sSQLT%04d.%d",
                serverWorkingPath, pathSep,
                ppTablespaces[tspNb]->id, pContainers[contNb].id);
            printf("        - new container name: %s\n",
                pContainers[contNb].name);
            break;
        case SQLB_CONT_DISK:
        case SQLB_CONT_FILE:
        default:
            printf("        Unknown container type.\n");
            break;
        }
    }
}

/* The API sqlbstsc is used to set or redefine table space containers
   while performing a 'redirected' restore of the database. */
sqlbstsc(&sqlca,
        SQLB_SET_CONT_FINAL_STATE,
        ppTablespaces[tspNb]->id, numContainers, pContainers);
DB2_API_CHECK("tablespace containers -- redefine");

/* The API sqlfmem is used here to free memory allocated by DB2 for use
   with the API sqlbtcq (Tablespace Container Query). */
sqlfmem(&sqlca, pContainers);
DB2_API_CHECK("tablespace containers memory -- free");
}

/* The API sqlfmem is used here to free memory allocated by DB2 for
   use with the API sqlbmtsq (Tablespace Query). */
sqlfmem(&sqlca, ppTablespaces);
DB2_API_CHECK("tablespaces memory -- free");

return 0;
} /* InaccessableContainersRedefine */
```

### dbhistfile sample program:

The dbhistfile sample files show how to read and update a database recovery file entry.

```
/******
** Licensed Materials - Property of IBM
**
** Governed under the terms of the International
** License Agreement for Non-Warranted Sample Code.
**
** (C) COPYRIGHT International Business Machines Corp. 2003
** All Rights Reserved.
**
** US Government Users Restricted Rights - Use, duplication or
** disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
*****
** SOURCE FILE NAME: dbhistfile.sqc
**
** SAMPLE: How to read and update a database recovery history file entry.
**
** This program ends in ".sqc" even though it does not contain
** embedded SQL statements. It links in the embedded SQL utility
** file for database connection and disconnection, so it needs the
** embedded SQL extension for the precompiler.
**
```

```

** DB2 APIs USED:
**      db2HistoryCloseScan -- Close Recovery History File Scan
**      db2HistoryGetEntry  -- Get Next Recovery History File Entry
**      db2HistoryOpenScan  -- Open Recovery History File Scan
**      db2HistoryUpdate    -- Update Recovery History File
**
** OUTPUT FILE: dbhistfile.out (available in the online documentation)
*****
**
** For more information on the sample programs, see the README file.
**
** For information on developing C applications, see the Application
** Development Guide.
**
** For information on using SQL statements, see the SQL Reference.
**
** For information on DB2 APIs, see the Administrative API Reference.
**
** For the latest information on programming, building, and running DB2
** applications, visit the DB2 application development website:
**      http://www.software.ibm.com/data/db2/udb/ad
*****/
#include "utilrecov.c"

/* local function prototypes */
int DbRecoveryHistoryFileRead(char *);
int DbFirstRecoveryHistoryFileEntryUpdate(char *, char *, char *);
int HistoryEntryDataFieldsAlloc(struct db2HistoryData *);
int HistoryEntryDisplay(struct db2HistoryData );
int HistoryEntryDataFieldsFree(struct db2HistoryData *);

int main(int argc, char *argv[])
{
    int rc = 0;
    char nodeName[SQL_INSTNAME_SZ + 1] = { 0 };
    char serverWorkingPath[SQL_PATH_SZ + 1] = { 0 };
    sqluint16 savedLogRetainValue = 0;
    char dbAlias[SQL_ALIAS_SZ + 1] = { 0 };
    char user[USERID_SZ + 1] = { 0 };
    char pswd[PSWD_SZ + 1] = { 0 };

    /* check the command line arguments */
    rc = CmdLineArgsCheck3(argc, argv, dbAlias, nodeName, user, pswd);
    CHECKRC(rc, "CmdLineArgsCheck3");

    printf("\nTHIS SAMPLE SHOWS HOW TO READ A DATABASE RECOVERY HISTORY FILE \n");
    printf("AND UPDATE A RECOVERY HISTORY FILE ENTRY. \n");

    /* attach to a local or remote instance */
    rc = InstanceAttach(nodeName, user, pswd);
    CHECKRC(rc, "Instance Attach");

    /* get the server working path */
    rc = ServerWorkingPathGet(dbAlias, serverWorkingPath);
    CHECKRC(rc, "ServerWorkingPathGet");

    rc = DbRecoveryHistoryFileRead(dbAlias);
    CHECKRC(rc, "DbRecoveryHistoryFileRead");

    rc = DbFirstRecoveryHistoryFileEntryUpdate(dbAlias, user, pswd);
    CHECKRC(rc, "DbFirstRecoveryHistoryFileEntryUpdate");

    /* Detach from the local or remote instance */
    rc = InstanceDetach(nodeName);
    CHECKRC(rc, "InstanceDetach");

    return 0;
}

```

## Sample Programs with embedded SQL

```
    } /* end main */

int DbRecoveryHistoryFileRead(char dbAlias[])
{
    int rc = 0;
    struct sqlca sqlca;
    struct db2HistoryOpenStruct dbHistoryOpenParam;
    sqluint32 numEntries = 0;
    sqluint16 recoveryHistoryFileHandle = 0;
    sqluint32 entryNb = 0;
    struct db2HistoryGetEntryStruct dbHistoryEntryGetParam;
    struct db2HistoryData histEntryData;

    printf("\n*****\n");
    printf("*** READ A DATABASE RECOVERY HISTORY FILE ***\n");
    printf("*****\n");
    printf("\nUSE THE DB2 APIs:\n");
    printf(" db2HistoryOpenScan -- Open Recovery History File Scan\n");
    printf(" db2HistoryGetEntry -- Get Next Recovery History File Entry\n");
    printf(" db2HistoryCloseScan -- Close Recovery History File Scan\n");
    printf("TO READ A DATABASE RECOVERY HISTORY FILE.\n");

    /* initialize the data structures */
    dbHistoryOpenParam.piDatabaseAlias = dbAlias;
    dbHistoryOpenParam.piTimestamp = NULL;
    dbHistoryOpenParam.piObjectName = NULL;
    dbHistoryOpenParam.iCallerAction = DB2HISTORY_LIST_HISTORY;
    dbHistoryEntryGetParam.pioHistData = &histEntryData;
    dbHistoryEntryGetParam.iCallerAction = DB2HISTORY_GET_ALL;

    rc = HistoryEntryDataFieldsAlloc(&histEntryData);
    CHECKRC(rc, "HistoryEntryDataFieldsAlloc");

    /*****
    /* OPEN THE DATABASE RECOVERY HISTORY FILE */
    *****/
    printf("\n Open recovery history file for '%s' database.\n", dbAlias);

    /* open the recovery history file to scan */
    db2HistoryOpenScan(db2Version810, &dbHistoryOpenParam, &sqlca);
    DB2_API_CHECK("database recovery history file -- open");
    numEntries = dbHistoryOpenParam.oNumRows;

    /* dbHistoryOpenParam.oHandle returns the handle for scan access */
    recoveryHistoryFileHandle = dbHistoryOpenParam.oHandle;
    dbHistoryEntryGetParam.iHandle = recoveryHistoryFileHandle;

    /*****
    /* READ AN ENTRY IN THE RECOVERY HISTORY FILE */
    *****/
    for (entryNb = 0; entryNb < numEntries; entryNb++)
    {
        printf("\n Read entry number %u.\n", entryNb);

        /* get the next entry from the recovery history file */
        db2HistoryGetEntry(db2Version810, &dbHistoryEntryGetParam, &sqlca);
        DB2_API_CHECK("database recovery history file entry -- read")

        /* display the entries in the recovery history file */
        printf("\n Display entry number %u.\n", entryNb);
        rc = HistoryEntryDisplay(histEntryData);
        CHECKRC(rc, "HistoryEntryDisplay");
    }

    /*****
    /* CLOSE THE DATABASE RECOVERY HISTORY FILE */
    *****/
}
```



```

printf("\n Close recovery history file for '%s' database.\n", dbAlias);

/* The API db2HistoryCloseScan ends the recovery history file scan and
   frees DB2 resources required for the scan. */
db2HistoryCloseScan(db2Version810, &recoveryHistoryFileHandle, &sqlca);
DB2_API_CHECK("database recovery history file -- close");

/* free the allocated memory */
rc = HistoryEntryDataFieldsFree(&histEntryData);
CHECKRC(rc, "HistoryEntryDataFieldsFree");

return 0;
} /* DbRecoveryHistoryFileRead */

int DbFirstRecoveryHistoryFileEntryUpdate(char dbAlias[], char user[],
                                          char pswd[])
{
    int rc = 0;
    struct sqlca sqlca;
    struct db2HistoryOpenStruct dbHistoryOpenParam;
    sqluint16 recoveryHistoryFileHandle = 0;
    struct db2HistoryGetEntryStruct dbHistoryEntryGetParam;
    struct db2HistoryData histEntryData;
    char newLocation[DB2HISTORY_LOCATION_SZ + 1] = { 0 };
    char newComment[DB2HISTORY_COMMENT_SZ + 1] = { 0 };
    struct db2HistoryUpdateStruct dbHistoryUpdateParam;

    printf("\n*****\n");
    printf("*** UPDATE A DATABASE RECOVERY HISTORY FILE ENTRY ***\n");
    printf("*****\n");
    printf("\nUSE THE DB2 APIs:\n");
    printf(" db2HistoryOpenScan -- Open Recovery History File Scan\n");
    printf(" db2HistoryGetEntry -- Get Next Recovery History File Entry\n");
    printf(" db2HistoryUpdate -- Update Recovery History File\n");
    printf(" db2HistoryCloseScan -- Close Recovery History File Scan\n");
    printf("TO UPDATE A DATABASE RECOVERY HISTORY FILE ENTRY.\n");

    /* initialize data structures */
    dbHistoryOpenParam.piDatabaseAlias = dbAlias;
    dbHistoryOpenParam.piTimestamp = NULL;
    dbHistoryOpenParam.piObjectName = NULL;
    dbHistoryOpenParam.iCallerAction = DB2HISTORY_LIST_HISTORY;

    dbHistoryEntryGetParam.pioHistData = &histEntryData;
    dbHistoryEntryGetParam.iCallerAction = DB2HISTORY_GET_ALL;
    rc = HistoryEntryDataFieldsAlloc(&histEntryData);
    CHECKRC(rc, "HistoryEntryDataFieldsAlloc");

    /******
    /* OPEN THE DATABASE RECOVERY HISTORY FILE */
    /******
    printf("\n Open the recovery history file for '%s' database.\n",
           dbAlias);

    /* The API db2HistoryOpenScan starts a recovery history file scan */
    db2HistoryOpenScan(db2Version810, &dbHistoryOpenParam, &sqlca);
    DB2_API_CHECK("database recovery history file -- open");

    /* dbHistoryOpenParam.oHandle returns the handle for scan access */
    recoveryHistoryFileHandle = dbHistoryOpenParam.oHandle;
    dbHistoryEntryGetParam.iHandle = recoveryHistoryFileHandle;

    /******
    /* READ THE FIRST ENTRY IN THE RECOVERY HISTORY FILE */
    /******
    printf("\n Read the first entry in the recovery history file.\n");

```

## Sample Programs with embedded SQL

```
/* The API db2HistoryGetEntry gets the next entry from the recovery
history file. */
db2HistoryGetEntry(db2Version810, &dbHistoryEntryGetParam, &sqlca);
DB2_API_CHECK("first recovery history file entry -- read");
printf("\n Display the first entry.\n");

/* HistoryEntryDisplay is a support function used to display the entries
in the recovery history file. */
rc = HistoryEntryDisplay(histEntryData);
CHECKRC(rc, "HistoryEntryDisplay");

/* update the first history file entry */
rc = DbConn(dbAlias, user, pswd);
CHECKRC(rc, "DbConn");

strcpy(newLocation, "this is the NEW LOCATION");
strcpy(newComment, "this is the NEW COMMENT");
printf("\n Update the first entry in the history file:\n");
printf("    new location = '%s'\n", newLocation);
printf("    new comment = '%s'\n", newComment);

dbHistoryUpdateParam.piNewLocation = newLocation;
dbHistoryUpdateParam.piNewDeviceType = NULL;
dbHistoryUpdateParam.piNewComment = newComment;
dbHistoryUpdateParam.iEID.ioNode = histEntryData.oEID.ioNode;
dbHistoryUpdateParam.iEID.ioHID = histEntryData.oEID.ioHID;

/* The API db2HistoryUpdate can be used to update the location,
device type, or comment in a history file entry. */

/* Call this API to update the location and comment of the first
entry in the history file: */
db2HistoryUpdate(db2Version810, &dbHistoryUpdateParam, &sqlca);
DB2_API_CHECK("first history file entry -- update");

rc = DbDisconn(dbAlias);
CHECKRC(rc, "DbDisconn");

/*****
/* CLOSE THE DATABASE RECOVERY HISTORY FILE */
*****/
printf("\n Close recovery history file for '%s' database.\n", dbAlias);

/* The API db2HistoryCloseScan ends the recovery history file scan and
frees DB2 resources required for the scan. */
db2HistoryCloseScan(db2Version810, &recoveryHistoryFileHandle, &sqlca);
DB2_API_CHECK("database recovery history file -- close");

/*****
/* RE-OPEN THE DATABASE RECOVERY HISTORY FILE */
*****/
printf("\n Open the recovery history file for '%s' database.\n",
dbAlias);

/* starts a recovery history file scan */
db2HistoryOpenScan(db2Version810, &dbHistoryOpenParam, &sqlca);
DB2_API_CHECK("database recovery history file -- open");

recoveryHistoryFileHandle = dbHistoryOpenParam.oHandle;
dbHistoryEntryGetParam.iHandle = recoveryHistoryFileHandle;
printf("\n Read the first recovery history file entry.\n");

/*****
/* READ THE FIRST ENTRY IN THE RECOVERY HISTORY FILE AFTER MODIFICATION */
*****/
db2HistoryGetEntry(db2Version810, &dbHistoryEntryGetParam, &sqlca);
DB2_API_CHECK("first recovery history file entry -- read");
```

```

printf("\n Display the first entry.\n");
rc = HistoryEntryDisplay(histEntryData);
CHECKRC(rc, "HistoryEntryDisplay");

/*****
/* CLOSE THE DATABASE RECOVERY HISTORY FILE */
*****/
printf("\n Close the recovery history file for '%s' database.\n",
      dbAlias);

/* ends the recovery history file scan */
db2HistoryCloseScan(db2Version810, &recoveryHistoryFileHandle, &sqlca);
DB2_API_CHECK("database recovery history file -- close");

/* free the allocated memory */
rc = HistoryEntryDataFieldsFree(&histEntryData);
CHECKRC(rc, "HistoryEntryDataFieldsFree");

return 0;
} /* DbFirstRecoveryHistoryFileEntryUpdate */

/*****
/* HistoryEntryDataFieldsAlloc      */
/* Allocates memory for all the fields in a database recovery history      */
/* file entry                       */
*****/
int HistoryEntryDataFieldsAlloc(struct db2HistoryData *pHistEntryData)
{
    int rc = 0;
    sqluint32 tsNb = 0;

    strcpy(pHistEntryData->ioHistDataID, "SQLUHINF");

    pHistEntryData->oObjectPart.pioData = malloc(DB2HISTORY_OBJPART_SZ + 1);
    pHistEntryData->oObjectPart.iLength = DB2HISTORY_OBJPART_SZ + 1;

    pHistEntryData->oEndTime.pioData = malloc(DB2HISTORY_TIMESTAMP_SZ + 1);
    pHistEntryData->oEndTime.iLength = DB2HISTORY_TIMESTAMP_SZ + 1;

    pHistEntryData->oFirstLog.pioData = malloc(DB2HISTORY_LOGFILE_SZ + 1);
    pHistEntryData->oFirstLog.iLength = DB2HISTORY_LOGFILE_SZ + 1;

    pHistEntryData->oLastLog.pioData = malloc(DB2HISTORY_LOGFILE_SZ + 1);
    pHistEntryData->oLastLog.iLength = DB2HISTORY_LOGFILE_SZ + 1;

    pHistEntryData->oID.pioData = malloc(DB2HISTORY_ID_SZ + 1);
    pHistEntryData->oID.iLength = DB2HISTORY_ID_SZ + 1;

    pHistEntryData->oTableQualifier.pioData =
        malloc(DB2HISTORY_TABLE_QUAL_SZ + 1);
    pHistEntryData->oTableQualifier.iLength = DB2HISTORY_TABLE_QUAL_SZ + 1;

    pHistEntryData->oTableName.pioData = malloc(DB2HISTORY_TABLE_NAME_SZ + 1);
    pHistEntryData->oTableName.iLength = DB2HISTORY_TABLE_NAME_SZ + 1;

    pHistEntryData->oLocation.pioData = malloc(DB2HISTORY_LOCATION_SZ + 1);
    pHistEntryData->oLocation.iLength = DB2HISTORY_LOCATION_SZ + 1;

    pHistEntryData->oComment.pioData = malloc(DB2HISTORY_COMMENT_SZ + 1);
    pHistEntryData->oComment.iLength = DB2HISTORY_COMMENT_SZ + 1;

    pHistEntryData->oCommandText.pioData = malloc(DB2HISTORY_COMMAND_SZ + 1);
    pHistEntryData->oCommandText.iLength = DB2HISTORY_COMMAND_SZ + 1;

    pHistEntryData->poEventSQLCA =
        (struct sqlca *)malloc(sizeof(struct sqlca));

```

## Sample Programs with embedded SQL

```
pHistEntryData->poTablespace = (db2Char *) malloc(3 * sizeof(db2Char));
for (tsNb = 0; tsNb < 3; tsNb++)
{
    pHistEntryData->poTablespace[tsNb].pioData = malloc(18 + 1);
    pHistEntryData->poTablespace[tsNb].iLength = 18 + 1;
}

pHistEntryData->iNumTablespaces = 3;

return 0;
} /* HistoryEntryDataFieldsAlloc */

/*****
/* HistoryEntryDisplay          */
/* Displays the fields of an entry in the database recovery history file */
*****/
int HistoryEntryDisplay(struct db2HistoryData histEntryData)
{
    int rc = 0;
    int bufLen = 0;
    char *buf = NULL;
    sqluint32 tsNb = 0;

    bufLen =
        MIN(histEntryData.oObjectPart.oLength,
histEntryData.oObjectPart.iLength);
    buf = malloc(bufLen + 1);
    memcpy(buf, histEntryData.oObjectPart.pioData, bufLen);
    buf[bufLen] = '\0';
    printf("    object part: %s\n", buf);
    free(buf);

    bufLen =
        MIN(histEntryData.oEndTime.oLength, histEntryData.oEndTime.iLength);
    buf = malloc(bufLen + 1);
    memcpy(buf, histEntryData.oEndTime.pioData, bufLen);
    buf[bufLen] = '\0';
    printf("    end time: %s\n", buf);
    free(buf);

    bufLen =
        MIN(histEntryData.oFirstLog.oLength, histEntryData.oFirstLog.iLength);
    buf = malloc(bufLen + 1);
    memcpy(buf, histEntryData.oFirstLog.pioData, bufLen);
    buf[bufLen] = '\0';
    printf("    first log: %s\n", buf);
    free(buf);

    bufLen =
        MIN(histEntryData.oLastLog.oLength, histEntryData.oLastLog.iLength);
    buf = malloc(bufLen + 1);
    memcpy(buf, histEntryData.oLastLog.pioData, bufLen);
    buf[bufLen] = '\0';
    printf("    last log: %s\n", buf);
    free(buf);

    bufLen = MIN(histEntryData.oID.oLength, histEntryData.oID.iLength);
    buf = malloc(bufLen + 1);
    memcpy(buf, histEntryData.oID.pioData, bufLen);
    buf[bufLen] = '\0';
    printf("    ID: %s\n", buf);
    free(buf);

    bufLen =
        MIN(histEntryData.oTableQualifier.oLength,
histEntryData.oTableQualifier.iLength);
```

```

buf = malloc(bufLen + 1);
memcpy(buf, histEntryData.oTableQualifier.pioData, bufLen);
buf[bufLen] = '\0';
printf("    table qualifier: %s\n", buf);
free(buf);

bufLen =
    MIN(histEntryData.oTableName.oLength, histEntryData.oTableName.iLength);
buf = malloc(bufLen + 1);
memcpy(buf, histEntryData.oTableName.pioData, bufLen);
buf[bufLen] = '\0';
printf("    table name: %s\n", buf);
free(buf);

bufLen =
    MIN(histEntryData.oLocation.oLength, histEntryData.oLocation.iLength);
buf = malloc(bufLen + 1);
memcpy(buf, histEntryData.oLocation.pioData, bufLen);
buf[bufLen] = '\0';
printf("    location: %s\n", buf);
free(buf);

bufLen =
    MIN(histEntryData.oComment.oLength, histEntryData.oComment.iLength);
buf = malloc(bufLen + 1);
memcpy(buf, histEntryData.oComment.pioData, bufLen);
buf[bufLen] = '\0';
printf("    comment: %s\n", buf);
free(buf);

bufLen =
    MIN(histEntryData.oCommandText.oLength,
histEntryData.oCommandText.iLength);

buf = malloc(bufLen + 1);
memcpy(buf, histEntryData.oCommandText.pioData, bufLen);
buf[bufLen] = '\0';
printf("    command text: %s\n", buf);
printf("    history file entry ID: %u\n", histEntryData.oEID.ioHID);
printf("    table spaces:\n");
free(buf);

for (tsNb = 0; tsNb < histEntryData.oNumTablespaces; tsNb++)
{
    bufLen =
        MIN(histEntryData.poTablespace[tsNb].oLength,
histEntryData.poTablespace[tsNb].iLength);
    buf = malloc(bufLen + 1);
    memcpy(buf, histEntryData.poTablespace[tsNb].pioData, bufLen);
    buf[bufLen] = '\0';
    printf("        %s\n", buf);
    free(buf);
}

printf("    type of operation: %c\n", histEntryData.oOperation);
printf("    granularity of the operation: %c\n", histEntryData.oObject);
printf("    operation type: %c\n", histEntryData.oOptype);
printf("    entry status: %c\n", histEntryData.oStatus);
printf("    device type: %c\n", histEntryData.oDeviceType);
printf("    SQLCA:\n");
printf("        sqlcode: %ld\n", histEntryData.poEventSQLCA->sqlcode);

bufLen = SQLUDF_SQLSTATE_LEN;
buf = malloc(bufLen + 1);
memcpy(buf, histEntryData.poEventSQLCA->sqlstate, bufLen);
buf[bufLen] = '\0';
printf("        sqlstate: %s\n", buf);

```

## Sample Programs with embedded SQL

```
        free(buf);

        bufLen = histEntryData.poEventSQLCA->sqlerrml;
        buf = malloc(bufLen + 1);
        memcpy(buf, histEntryData.poEventSQLCA->sqlerrmc, bufLen);
        buf[bufLen] = '\0';
        printf("        message: %s\n", buf);
        free(buf);

        return 0;
} /* HistoryEntryDisplay */

/*****
/* HistoryEntryDataFieldsFree          */
/* Deallocates the memory for database recovery history file structures */
/*****
int HistoryEntryDataFieldsFree(struct db2HistoryData *pHistEntryData)
{
    int rc = 0;
    sqluint32 tsNb = 0;

    free(pHistEntryData->oObjectPart.pioData);
    free(pHistEntryData->oEndTime.pioData);
    free(pHistEntryData->oFirstLog.pioData);
    free(pHistEntryData->oLastLog.pioData);
    free(pHistEntryData->oID.pioData);
    free(pHistEntryData->oTableQualifier.pioData);
    free(pHistEntryData->oTableName.pioData);
    free(pHistEntryData->oLocation.pioData);
    free(pHistEntryData->oComment.pioData);
    free(pHistEntryData->oCommandText.pioData);
    free(pHistEntryData->poEventSQLCA);
    pHistEntryData->oObjectPart.pioData = NULL;
    pHistEntryData->oEndTime.pioData = NULL;
    pHistEntryData->oFirstLog.pioData = NULL;
    pHistEntryData->oLastLog.pioData = NULL;
    pHistEntryData->oID.pioData = NULL;
    pHistEntryData->oTableQualifier.pioData = NULL;
    pHistEntryData->oTableName.pioData = NULL;
    pHistEntryData->oLocation.pioData = NULL;
    pHistEntryData->oComment.pioData = NULL;
    pHistEntryData->oCommandText.pioData = NULL;
    pHistEntryData->poEventSQLCA = NULL;

    for (tsNb = 0; tsNb < 3; tsNb++)
    {
        free(pHistEntryData->poTablespace[tsNb].pioData);
        pHistEntryData->poTablespace[tsNb].pioData = NULL;
    }

    free(pHistEntryData->poTablespace);
    pHistEntryData->poTablespace = NULL;

    return 0;
} /* HistoryEntryDataFieldsFree */
```

### dblogconn sample program:

The dblogconn sample files show how to read database log files with a database connection.

```
/*****
** Licensed Materials - Property of IBM
**
** Governed under the terms of the International
** License Agreement for Non-Warranted Sample Code.
**
```

## Sample Programs with embedded SQL

```
** (C) COPYRIGHT International Business Machines Corp. 2003
** All Rights Reserved.
**
** US Government Users Restricted Rights - Use, duplication or
** disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
*****
**
** SOURCE FILE NAME: dblogconn.sqc
**
** SAMPLE: How to read database log files asynchronously with a database
**         connection
**
**         Note:
**         You must be initially disconnected from the sample database
**         to run this program. To ensure you are, enter 'db2 connect
**         reset' on the command line prior to running dblogconn.
**
** DB2 API USED:
**         db2CfgSet -- Set Configuration
**         db2ReadLog -- Asynchronous Read Log
**
** SQL STATEMENTS USED:
**         ALTER TABLE
**         COMMIT
**         DELETE
**         INSERT
**         ROLLBACK
**         CONNECT RESET
**
** OUTPUT FILE: dblogconn.out (available in the online documentation)
*****
**
** For detailed information about database backup and database recovery, see
** the Data Recovery and High Availability Guide and Reference. This manual
** will help you to determine which database and table space recovery methods
** are best suited to your business environment.
**
** For more information on the sample programs, see the README file.
**
** For information on developing C applications, see the Application
** Development Guide.
**
** For information on using SQL statements, see the SQL Reference.
**
** For information on DB2 APIs, see the Administrative API Reference.
**
** For the latest information on programming, building, and running DB2
** applications, visit the DB2 application development website:
**     http://www.software.ibm.com/data/db2/udb/ad
*****/
#include "utilrecov.c"
#include "utilemb.h"

/* local function prototypes */
int DbLogRecordsForCurrentConnectionRead(char *, char *, char *, char *);

int main(int argc, char *argv[])
{
    int rc = 0;
    char nodeName[SQL_INSTNAME_SZ + 1] = { 0 };
    char serverWorkingPath[SQL_PATH_SZ + 1] = { 0 };
    sqluint16 savedLogRetainValue = 0;
    char dbAlias[SQL_ALIAS_SZ + 1] = { 0 };
    char user[USERID_SZ + 1] = { 0 };
    char pswd[PSWD_SZ + 1] = { 0 };

    /* check the command line arguments */
```

## Sample Programs with embedded SQL

```
rc = CmdLineArgsCheck3(argc, argv, dbAlias, nodeName, user, pswd);
CHECKRC(rc, "CmdLineArgsCheck3");

printf("\nTHIS SAMPLE SHOWS HOW TO READ DATABASE LOGS ASYNCHRONOUSLY\n");
printf("WITH A DATABASE CONNECTION.\n");

/* attach to a local or remote instance */
rc = InstanceAttach(nodeName, user, pswd);
CHECKRC(rc, "Instance Attach");

/* get the server working path */
rc = ServerWorkingPathGet(dbAlias, serverWorkingPath);
CHECKRC(rc, "ServerWorkingPathGet");

/* call the function to do asynchronous log read */
rc = DbLogRecordsForCurrentConnectionRead(dbAlias,
                                           user, pswd, serverWorkingPath);
CHECKRC(rc, "DbLogRecordsForCurrentConnectionRead");

/* Detach from the local or remote instance */
rc = InstanceDetach(nodeName);
CHECKRC(rc, "InstanceDetach");

return 0;
} /* end main */

int DbLogRecordsForCurrentConnectionRead(char dbAlias[],
                                         char user[],
                                         char pswd[],
                                         char serverWorkingPath[])
{
    int rc = 0;
    struct sqlca sqlca;
    db2CfgParam cfgParameters[1];
    db2Cfg cfgStruct;
    unsigned short logretain = 0;

    db2BackupStruct backupStruct;
    db2TablespaceStruct tablespaceStruct;
    db2MediaListStruct mediaListStruct;
    db2UInt32 backupImageSize = 0;
    db2RestoreStruct restoreStruct;
    db2TablespaceStruct rtablespaceStruct;
    db2MediaListStruct rmediaListStruct;

    SQLU_LSN startLSN;
    SQLU_LSN endLSN;
    char *logBuffer = NULL;
    sqluint32 logBufferSize = 0;
    db2ReadLogInfoStruct readLogInfo;
    db2ReadLogStruct readLogInput;
    int i = 0;

    printf("\n*****\n");
    printf("*** ASYNCHRONOUS READ LOG ***\n");
    printf("*****\n");
    printf("\nUSE THE DB2 APIs:\n");
    printf(" db2CfgSet -- Set Configuration\n");
    printf(" db2Backup -- Backup Database\n");
    printf(" db2ReadLog -- Asynchronous Read Log\n");
    printf("AND THE SQL STATEMENTS:\n");
    printf(" CONNECT\n");
    printf(" ALTER TABLE\n");
    printf(" COMMIT\n");
    printf(" INSERT\n");
    printf(" DELETE\n");
    printf(" ROLLBACK\n");
}
```



```

printf(" CONNECT RESET\n");
printf("TO READ LOG RECORDS FOR THE CURRENT CONNECTION.\n");

printf("\n Update \'%s\' database configuration:\n", dbAlias);
printf(" - Enable the database configuration parameter LOGRETAIN \n");
printf(" i.e., set LOGRETAIN = RECOVERY/YES\n");

/* initialize cfgParameters */
cfgParameters[0].flags = 0;
cfgParameters[0].token = SQLF_DBTN_LOG_RETAIN;
cfgParameters[0].ptrvalue = (char *)&logretain;

/* enable LOGRETAIN */
logretain = SQLF_LOGRETAIN_RECOVERY;

/* initialize cfgStruct */
cfgStruct.numItems = 1;
cfgStruct.paramArray = cfgParameters;
cfgStruct.flags = db2CfgDatabase | db2CfgDelayed;
cfgStruct.dbname = dbAlias;

/* get database configuration */
db2CfgSet(db2Version810, (void *)&cfgStruct, &sqlca);
DB2_API_CHECK("Db Log Retain -- Enable");

tablespaceStruct tablespaces = NULL;
tablespaceStruct.numTablespaces = 0;

mediaListStruct.locations = &serverWorkingPath;
mediaListStruct.numLocations = 1;
mediaListStruct.locationType = SQLU_LOCAL_MEDIA;

/* Calling up the routine for database backup */
rc = DbBackup(dbAlias, user, pswd, serverWorkingPath, &backupStruct);
CHECKRC(rc, "DbBackup");

/* connect to the database */
rc = DbConn(dbAlias, user, pswd);
CHECKRC(rc, "DbConn");

/* invoke SQL statements to fill database log */
printf("\n Invoke the following SQL statements:\n"
" ALTER TABLE emp_resume DATA CAPTURE CHANGES;\n"
" COMMIT;\n"
" INSERT INTO emp_resume\n"
" VALUES('000777', 'ascii', 'This is a new resume.);\n"
" ('777777', 'ascii', 'This is another new resume');\n"
" COMMIT;\n"
" DELETE FROM emp_resume WHERE empno = '000777';\n"
" DELETE FROM emp_resume WHERE empno = '777777';\n"
" COMMIT;\n"
" DELETE FROM emp_resume WHERE empno = '000140';\n"
" ROLLBACK;\n"
" ALTER TABLE emp_resume DATA CAPTURE NONE;\n" " COMMIT;\n");

EXEC SQL ALTER TABLE emp_resume DATA CAPTURE CHANGES;
EMB_SQL_CHECK("SQL statement 1 -- invoke");

EXEC SQL COMMIT;
EMB_SQL_CHECK("SQL statement 2 -- invoke");

EXEC SQL INSERT INTO emp_resume
VALUES('000777', 'ascii', 'This is a new resume.'),
('777777', 'ascii', 'This is another new resume');
EMB_SQL_CHECK("SQL statement 3 -- invoke");

EXEC SQL COMMIT;

```

## Sample Programs with embedded SQL

```
EMB_SQL_CHECK("SQL statement 4 -- invoke");

EXEC SQL DELETE FROM emp_resume WHERE empno = '000777';
EMB_SQL_CHECK("SQL statement 5 -- invoke");

EXEC SQL DELETE FROM emp_resume WHERE empno = '777777';
EMB_SQL_CHECK("SQL statement 6 -- invoke");

EXEC SQL COMMIT;
EMB_SQL_CHECK("SQL statement 7 -- invoke");

EXEC SQL DELETE FROM emp_resume WHERE empno = '000140';
EMB_SQL_CHECK("SQL statement 8 -- invoke");

EXEC SQL ROLLBACK;
EMB_SQL_CHECK("SQL statement 9 -- invoke");

EXEC SQL ALTER TABLE emp_resume DATA CAPTURE NONE;
EMB_SQL_CHECK("SQL statement 10 -- invoke");

EXEC SQL COMMIT;
EMB_SQL_CHECK("SQL statement 11 -- invoke");

printf("\n Start reading database log.\n");

logBuffer = NULL;
logBufferSize = 0;

/*
 * The API db2ReadLog (Asynchronous Read Log) is used to extract
 * records from the database logs, and to query the log manager for
 * current log state information. This API can only be used on
 * recoverable databases.
 */

/* Query the log manager for current log state information. */
readLogInput.iCallerAction = DB2READLOG_QUERY;
readLogInput.piStartLSN = NULL;
readLogInput.piEndLSN = NULL;
readLogInput.poLogBuffer = NULL;
readLogInput.iLogBufferSize = 0;
readLogInput.iFilterOption = DB2READLOG_FILTER_ON;
readLogInput.poReadLogInfo = &readLogInfo;

db2ReadLog(db2Version810, &readLogInput, &sqlca);
DB2_API_CHECK("database log info -- get");

logBufferSize = 64 * 1024; /* Maximum size of a log buffer */
logBuffer = (char *)malloc(logBufferSize);

memcpy(&startLSN, &(readLogInfo.initialLSN), sizeof(startLSN));
memcpy(&endLSN, &(readLogInfo.nextStartLSN), sizeof(endLSN));

/*
 * Extract a log record from the database logs, and read the first
 * log sequence asynchronously.
 */
readLogInput.iCallerAction = DB2READLOG_READ;
readLogInput.piStartLSN = &startLSN;
readLogInput.piEndLSN = &endLSN;
readLogInput.poLogBuffer = logBuffer;
readLogInput.iLogBufferSize = logBufferSize;
readLogInput.iFilterOption = DB2READLOG_FILTER_ON;
readLogInput.poReadLogInfo = &readLogInfo;

db2ReadLog(db2Version810, &readLogInput, &sqlca);
if (sqlca.sqlcode != SQLU_RLOG_READ_TO_CURRENT)
```

```

    {
        DB2_API_CHECK("database logs -- read");
    }
else
    {
        if (readLogInfo.logRecsWritten == 0)
        {
            printf("\n Database log empty.\n");
        }
    }

/* display log buffer */
rc = LogBufferDisplay(logBuffer, readLogInfo.logRecsWritten);
CHECKRC(rc, "LogBufferDisplay");

while (sqlca.sqlcode != SQLU_RLOG_READ_TO_CURRENT)
{
    /* read the next log sequence */

    memcpy(&startLSN, &(readLogInfo.nextStartLSN), sizeof(startLSN));

    /*
     * Extract a log record from the database logs, and read the
     * next log sequence asynchronously.
     */
    db2ReadLog(db2Version810, &readLogInput, &sqlca);
    if (sqlca.sqlcode != SQLU_RLOG_READ_TO_CURRENT)
    {
        DB2_API_CHECK("database logs -- read");
    }
    /* display log buffer */
    rc = LogBufferDisplay(logBuffer, readLogInfo.logRecsWritten);
    CHECKRC(rc, "LogBufferDisplay");
}

/* free the log buffer */
free(logBuffer);
logBuffer = NULL;
logBufferSize = 0;

/* disconnect from the database */
rc = DbDisconn(dbAlias);
CHECKRC(rc, "DbDisconn");

return 0;
} /* DbLogRecordsForCurrentConnectionRead */

```

**dblognoconn sample program:**

The dblognoconn sample files show how to read database log files with no database connection.

```

/*****
** Licensed Materials - Property of IBM
**
** Governed under the terms of the International
** License Agreement for Non-Warranted Sample Code.
**
** (C) COPYRIGHT International Business Machines Corp. 2003
** All Rights Reserved.
**
** US Government Users Restricted Rights - Use, duplication or
** disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
*****/
**
** SOURCE FILE NAME: dblognoconn.sqc
**

```

## Sample Programs with embedded SQL

```
** SAMPLE: How to read database log files asynchronously
**         with no database connection
**
**         This program ends in ".sql" even though it does not contain
**         embedded SQL statements. It links in the embedded SQL utility
**         file for database connection and disconnection, so it needs the
**         embedded SQL extension for the precompiler.
**
**         Note:
**         You must be disconnected from the sample database to run
**         this program. To ensure you are, enter 'db2 connect reset'
**         on the command line prior to running dblognoconn.
**
** DB2 API USED:
**         db2CfgSet -- Set Configuration
**         db2ReadLogNoConnInit -- Read log without a db connection
**         db2ReadLog -- Asynchronous Read Log
**
** OUTPUT FILE: dblognoconn.out (available in the online documentation)
*****
** For detailed information about database backup and database recovery, see
** the Data Recovery and High Availability Guide and Reference. This manual
** will help you to determine which database and table space recovery methods
** are best suited to your business environment.
**
** For more information on the sample programs, see the README file.
**
** For information on developing C applications, see the Application
** Development Guide.
**
** For information on using SQL statements, see the SQL Reference.
**
** For information on DB2 APIs, see the Administrative API Reference.
**
** For the latest information on programming, building, and running DB2
** applications, visit the DB2 application development website:
**     http://www.software.ibm.com/data/db2/udb/ad
*****/
#include "utilrecov.c"
#include "utilemb.h"

/* local function prototypes */
int DbReadLogRecordsNoConn(char *);

int main(int argc, char *argv[])
{
    int rc = 0;
    char nodeName[SQL_INSTNAME_SZ + 1] = { 0 };
    char serverWorkingPath[SQL_PATH_SZ + 1] = { 0 };
    sqluint16 savedLogRetainValue = 0;
    char dbAlias[SQL_ALIAS_SZ + 1] = { 0 };
    char user[USERID_SZ + 1] = { 0 };
    char pswd[PSWD_SZ + 1] = { 0 };

    /* check the command line arguments */
    rc = CmdLineArgsCheck3(argc, argv, dbAlias, nodeName, user, pswd);
    CHECKRC(rc, "CmdLineArgsCheck3");

    printf("\nTHIS SAMPLE SHOWS HOW TO READ DATABASE LOGS ASYNCHRONOUSLY\n");
    printf("WITH NO DATABASE CONNECTION.\n");

    /* get the server working path */
    rc = ServerWorkingPathGet(dbAlias, serverWorkingPath);
    CHECKRC(rc, "ServerWorkingPathGet");

    rc = DbRecoveryHistoryFilePrune(dbAlias, user, pswd);
```

```

CHECKRC(rc, "DbRecoveryHistoryFilePrune");

/* Detach from the local or remote instance */
rc = InstanceDetach(nodeName);
CHECKRC(rc, "InstanceDetach");

rc = DbReadLogRecordsNoConn(dbAlias);
CHECKRC(rc, "DbReadLogRecordsNoConn");

return 0;
} /* end main */

int DbReadLogRecordsNoConn(char dbAlias[])
{
    int rc = 0;
    struct sqlca sqlca;
    char logPath[SQL_PATH_SZ + 1] = { 0 };
    db2CfgParam cfgParameters[1];
    db2Cfg cfgStruct;
    char nodeName[] = "NODE0000\0";
    db2Uint32 readLogMemSize = 0;
    char *readLogMemory = NULL;
    struct db2ReadLogNoConnInitStruct readLogInit;
    struct db2ReadLogNoConnInfoStruct readLogInfo;
    struct db2ReadLogNoConnStruct readLogInput;
    SQLU_LSN startLSN;
    SQLU_LSN endLSN;
    char *logBuffer = NULL;
    db2Uint32 logBufferSize = 0;
    struct db2ReadLogNoConnTermStruct readLogTerm;

    printf("\n*****\n");
    printf("*** NO DB CONNECTION READ LOG ***\n");
    printf("*****\n");
    printf("\nUSE THE DB2 APIs:\n");
    printf(" db2ReadLogNoConnInit -- Initialize No Db Connection Read Log\n");
    printf(" db2ReadLogNoConn -- No Db Connection Read Log\n");
    printf(" db2ReadLogNoConnTerm -- Terminate No Db Connection Read Log\n");
    printf("TO READ LOG RECORDS FROM A DATABASE LOG DIRECTORY.\n");

    /* Determine the logpath to read log files from */
    cfgParameters[0].flags = 0;
    cfgParameters[0].token = SQLF_DBTN_LOGPATH;
    cfgParameters[0].ptrvalue =
        (char *)malloc((SQL_PATH_SZ + 1) * sizeof(char));

    /* Initialize cfgStruct */
    cfgStruct.numItems = 1;
    cfgStruct.paramArray = cfgParameters;
    cfgStruct.flags = db2CfgDatabase;
    cfgStruct.dbname = dbAlias;

    db2CfgGet(db2Version810, (void *)&cfgStruct, &sqlca);
    DB2_API_CHECK("log path -- get");

    strcpy(logPath, cfgParameters[0].ptrvalue);
    free(cfgParameters[0].ptrvalue);
    cfgParameters[0].ptrvalue = NULL;

    /*
     * First we must allocate memory for the API's control blocks and log
     * buffer
     */
    readLogMemSize = 4 * 4096;
    readLogMemory = (char *)malloc(readLogMemSize);

    /* Invoke the initialization API to set up the control blocks */

```

## Sample Programs with embedded SQL

```
readLogInit.iFilterOption = DB2READLOG_FILTER_ON;
readLogInit.piLogFilePath = logPath;
readLogInit.piOverflowLogPath = NULL;
readLogInit.iRetrieveLogs = DB2READLOG_RETRIEVE_OFF;
readLogInit.piDatabaseName = dbAlias;
readLogInit.piNodeName = nodeName;
readLogInit.iReadLogMemoryLimit = readLogMemSize;
readLogInit.poReadLogMemPtr = &readLogMemory;

db2ReadLogNoConnInit(db2Version810, &readLogInit, &sqlca);
if (sqlca.sqlcode != SQLU_RLOG_LSNS_REUSED)
{
    DB2_API_CHECK("database logs no db conn -- initialization");
}
/* Query for the current log information */
readLogInput.iCallerAction = DB2READLOG_QUERY;
readLogInput.piStartLSN = NULL;
readLogInput.piEndLSN = NULL;
readLogInput.poLogBuffer = NULL;
readLogInput.iLogBufferSize = 0;
readLogInput.piReadLogMemPtr = readLogMemory;
readLogInput.poReadLogInfo = &readLogInfo;

db2ReadLogNoConn(db2Version810, &readLogInput, &sqlca);
if (sqlca.sqlcode != 0)
{
    DB2_API_CHECK("database logs no db conn -- query");
}
/* Read some log records */
logBufferSize = 64 * 1024; /* Maximum size of a log buffer */
logBuffer = (char *)malloc(logBufferSize);

memcpy(&startLSN, &(readLogInfo.nextStartLSN), sizeof(startLSN));
endLSN.lsnWord[0] = 0xffff;
endLSN.lsnWord[1] = 0xffff;
endLSN.lsnWord[2] = 0xffff;

readLogInput.iCallerAction = DB2READLOG_READ;
readLogInput.piStartLSN = &startLSN;
readLogInput.piEndLSN = &endLSN;
readLogInput.poLogBuffer = logBuffer;
readLogInput.iLogBufferSize = logBufferSize;
readLogInput.piReadLogMemPtr = readLogMemory;
readLogInput.poReadLogInfo = &readLogInfo;

db2ReadLogNoConn(db2Version810, &readLogInput, &sqlca);
if (sqlca.sqlcode != SQLU_RLOG_READ_TO_CURRENT)
{
    DB2_API_CHECK("database logs no db conn -- read");
}
else
{
    if (readLogInfo.logRecsWritten == 0)
    {
        printf("\n Database log empty.\n");
    }
}

/* Display the log records read */
rc = LogBufferDisplay(logBuffer, readLogInfo.logRecsWritten);
CHECKRC(rc, "LogBufferDisplay");

while (sqlca.sqlcode != SQLU_RLOG_READ_TO_CURRENT)
{
    /* read the next log sequence */
    memcpy(&startLSN, &(readLogInfo.nextStartLSN), sizeof(startLSN));
```

```

/*
 * Extract a log record from the database logs, and read the
 * next log sequence asynchronously.
 */
db2ReadLogNoConn(db2Version810, &readLogInput, &sqlca);
if (sqlca.sqlcode != SQLU_RLOG_READ_TO_CURRENT)
{
    DB2_API_CHECK("database logs no db conn -- read");
}
/* display log buffer */
rc = LogBufferDisplay(logBuffer, readLogInfo.logRecsWritten);
CHECKRC(rc, "LogBufferDisplay");
}

printf("\nRead to end of logs.\n\n");
free(logBuffer);
logBuffer = NULL;
logBufferSize = 0;

readLogTerm.poReadLogMemPtr = &readLogMemory;

db2ReadLogNoConnTerm(db2Version810, &readLogTerm, &sqlca);
DB2_API_CHECK("database logs no db conn -- terminate");

return 0;
} /* DbReadLogRecordsNoConn */

```

#### dbrestore sample program:

The dbrestore sample files show how to restore a database from a backup image.

```

/*****
** Licensed Materials - Property of IBM
**
** Governed under the terms of the International
** License Agreement for Non-Warranted Sample Code.
**
** (C) COPYRIGHT International Business Machines Corp. 2003
** All Rights Reserved.
**
** US Government Users Restricted Rights - Use, duplication or
** disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
*****/
**
** SOURCE FILE NAME: dbrestore.sqc
**
** SAMPLE: How to restore a database from a backup
**
**      This program ends in ".sqc" even though it does not contain
**      embedded SQL statements. It links in the embedded SQL utility
**      file for database connection and disconnection, so it needs the
**      embedded SQL extension for the precompiler.
**
**      Note:
**      You must be disconnected from the sample database to run
**      this program. To ensure you are, enter 'db2 connect reset'
**      on the command line prior to running dbrestore.
**
** DB2 API USED:
**      db2CfgSet -- Set Configuration
**      db2Restore -- Restore Database
**
** OUTPUT FILE: dbrestore.out (available in the online documentation)
*****/
**
** For detailed information about database backup and database recovery, see
** the Data Recovery and High Availability Guide and Reference. This manual

```

## Sample Programs with embedded SQL

```
** will help you to determine which database and table space recovery methods
** are best suited to your business environment.
**
** For more information on the sample programs, see the README file.
**
** For information on developing C applications, see the Application
** Development Guide.
**
** For information on using SQL statements, see the SQL Reference.
**
** For information on DB2 APIs, see the Administrative API Reference.
**
** For the latest information on programming, building, and running DB2
** applications, visit the DB2 application development website:
**   http://www.software.ibm.com/data/db2/udb/ad
**
*****/
#include "utilrecov.c"

/* local function prototypes */
int DbBackupAndRestore(char *, char *, char *, char *, char *);

int main(int argc, char *argv[])
{
    int rc = 0;
    char nodeName[SQL_INSTNAME_SZ + 1] = { 0 };
    char serverWorkingPath[SQL_PATH_SZ + 1] = { 0 };
    char restoredDbAlias[SQL_ALIAS_SZ + 1] = { 0 };
    char dbAlias[SQL_ALIAS_SZ + 1] = { 0 };
    char user[USERID_SZ + 1] = { 0 };
    char pswd[PSWD_SZ + 1] = { 0 };

    /* check the command line arguments */
    rc = CmdLineArgsCheck3(argc, argv, dbAlias, nodeName, user, pswd);
    CHECKRC(rc, "CmdLineArgsCheck3");

    printf("\nTHIS SAMPLE SHOWS HOW TO RESTORE A DATABASE FROM A\n");
    printf("BACKUP.\n");

    strcpy(restoredDbAlias, dbAlias);

    /* attach to a local or remote instance */
    rc = InstanceAttach(nodeName, user, pswd);
    CHECKRC(rc, "Instance Attach");

    /* get the server working path */
    rc = ServerWorkingPathGet(dbAlias, serverWorkingPath);
    CHECKRC(rc, "ServerWorkingPathGet");

    printf("\nNOTE: Backup images will be created on the server\n");
    printf("        in the directory %s,\n", serverWorkingPath);
    printf("        and will not be deleted by the program.\n");

    /* prune the recovery history file */
    rc = DbRecoveryHistoryFilePrune(dbAlias, user, pswd);
    CHECKRC(rc, "DbRecoveryHistoryFilePrune");

    rc = DbBackupAndRestore(dbAlias,
                           restoredDbAlias, user, pswd, serverWorkingPath);
    CHECKRC(rc, "DbBackupAndRestore");

    /* Detach from the local or remote instance */
    rc = InstanceDetach(nodeName);
    CHECKRC(rc, "InstanceDetach");

    return 0;
} /* end main */
```



```

int DbBackupAndRestore(char dbAlias[],
                      char restoredDbAlias[], char user[],
                      char pswd[], char serverWorkingPath[])
{
    int rc = 0;
    struct sqlca sqlca;
    db2CfgParam cfgParameters[1];
    db2Cfg cfgStruct;
    unsigned short logretain = 0;
    char restoreTimestamp[SQLU_TIME_STAMP_LEN + 1] = { 0 };

    db2BackupStruct backupStruct;
    db2TablespaceStruct tablespaceStruct;
    db2MediaListStruct mediaListStruct;
    db2Uint32 backupImageSize = 0;
    db2RestoreStruct restoreStruct;
    db2TablespaceStruct rtablespaceStruct;
    db2MediaListStruct rmediaListStruct;

    printf("\n*****\n");
    printf("*** BACK UP AND RESTORE A DATABASE ***\n");
    printf("*****\n");
    printf("\nUSE THE DB2 APIs:\n");
    printf(" db2CfgSet -- Set Configuration\n");
    printf(" db2Backup -- Backup Database\n");
    printf(" db2Restore -- Restore Database\n");
    printf("TO BACK UP AND RESTORE A DATABASE.\n");

    printf("\n Update \'%s\' database configuration:\n", dbAlias);
    printf(" - Disable the database configuration parameter LOGRETAIN\n");
    printf(" i.e., set LOGRETAIN = OFF/NO\n");

    /* initialize cfgParameters */
    /* SQLF_DBTN_LOG_RETAIN is a token of the updatable database configuration
       parameter 'logretain'; it is used to update the database configuration
       file */
    cfgParameters[0].flags = 0;
    cfgParameters[0].token = SQLF_DBTN_LOG_RETAIN;
    cfgParameters[0].ptrvalue = (char *)&logretain;

    /* disable the database configuration parameter 'logretain' */
    logretain = SQLF_LOGRETAIN_DISABLE;

    /* initialize cfgStruct */
    cfgStruct.numItems = 1;
    cfgStruct.paramArray = cfgParameters;
    cfgStruct.flags = db2CfgDatabase | db2CfgDelayed;
    cfgStruct.dbname = dbAlias;

    /* set database configuration */
    db2CfgSet(db2Version810, (void *)&cfgStruct, &sqlca);
    DB2_API_CHECK("Db Log Retain -- Disable");

    /******
    /* BACKUP THE DATABASE */
    /******

    /* Calling up the routine for database backup */
    rc = DbBackup(dbAlias, user, pswd, serverWorkingPath, &backupStruct);
    CHECKRC(rc, "DbBackup");

    /******
    /* RESTORE THE DATABASE */
    /******

    strcpy(restoreTimestamp, backupStruct.oTimestamp);

```

## Sample Programs with embedded SQL

```
printf("\n Restoring a database ...\n");
printf(" - source image alias      : %s\n", dbAlias);
printf(" - source image time stamp: %s\n", restoreTimestamp);
printf(" - target database         : %s\n", restoredDbAlias);

rtablespaceStruct tablespaces = NULL;
rtablespaceStruct numTablespaces = 0;

rmediaListStruct locations = &serverWorkingPath;
rmediaListStruct numLocations = 1;
rmediaListStruct locationType = SQLU_LOCAL_MEDIA;

restoreStruct.piSourceDBAlias = dbAlias;
restoreStruct.piTargetDBAlias = restoredDbAlias;

restoreStruct.piTimestamp = restoreTimestamp;
restoreStruct.piTargetDBPath = NULL;
restoreStruct.piReportFile = NULL;
restoreStruct.piTablespaceList = &rtablespaceStruct;
restoreStruct.piMediaList = &rmediaListStruct;
restoreStruct.piUsername = user;
restoreStruct.piPassword = pswd;
restoreStruct.piNewLogPath = NULL;
restoreStruct.piVendorOptions = NULL;
restoreStruct.iVendorOptionsSize = 0;
restoreStruct.iParallelism = 1;
restoreStruct.iBufferSize = 1024; /* 1024 x 4KB */
restoreStruct.iNumBuffers = 2;
restoreStruct.iCallerAction = DB2RESTORE_RESTORE;
restoreStruct.iOptions =
    DB2RESTORE_OFFLINE | DB2RESTORE_DB | DB2RESTORE_NODATALINK |
    DB2RESTORE_NOROLLFWD;

/* The API db2Restore is used to restore a database that has been backed
   up using the API db2Backup. */
db2Restore(db2Version810, &restoreStruct, &sqlca);
EXPECTED_WARN_CHECK("database restore -- start");

while (sqlca.sqlcode != 0)
{
    /* continue the restore operation */
    printf("\n Continuing the restore operation...\n");

    /* depending on the sqlca.sqlcode value, user action may be
       required, such as mounting a new tape */

    restoreStruct.iCallerAction = DB2RESTORE_CONTINUE;

    /* restore the database */
    db2Restore(db2Version810, &restoreStruct, &sqlca);
    DB2_API_CHECK("database restore -- continue");
}

printf("\n Restore finished.\n");

return 0;
} /* DbBackupAndRestore */
```

### dbrollfwd sample program:

The dbrollfwd sample files show how to perform rollforward operation after a database restore operation.

```
/******
** Licensed Materials - Property of IBM
**
** Governed under the terms of the International
```

## Sample Programs with embedded SQL

```
** License Agreement for Non-Warranted Sample Code.
**
** (C) COPYRIGHT International Business Machines Corp. 2003
** All Rights Reserved.
**
** US Government Users Restricted Rights - Use, duplication or
** disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
*****
**
** SOURCE FILE NAME: dbrollfwd.sqc
**
** SAMPLE: How to perform rollforward after restore of a database
**
** This program ends in ".sqc" even though it does not contain
** embedded SQL statements. It links in the embedded SQL utility
** file for database connection and disconnection, so it needs the
** embedded SQL extension for the precompiler.
**
** Note:
** You must be disconnected from the sample database to run
** this program. To ensure you are, enter 'db2 connect reset'
** on the command line prior to running dbrollfwd.
**
** DB2 APIs USED:
** db2CfgSet -- Set Configuration
** db2Restore -- Restore Database
** db2Rollforward -- Rollforward Database
**
** OUTPUT FILE: dbrollfwd.out (available in the online documentation)
*****
** For detailed information about database backup and database recovery, see
** the Data Recovery and High Availability Guide and Reference. This manual
** will help you to determine which database and table space recovery methods
** are best suited to your business environment.
**
** For more information on the sample programs, see the README file.
**
** For information on developing C applications, see the Application
** Development Guide.
**
** For information on using SQL statements, see the SQL Reference.
**
** For information on DB2 APIs, see the Administrative API Reference.
**
** For the latest information on programming, building, and running DB2
** applications, visit the DB2 application development website:
** http://www.software.ibm.com/data/db2/udb/ad
*****/
#include "utilrecov.c"

/* local function prototypes */
int DbBackupRestoreAndRollforward(char *, char *, char *, char *, char *);

int main(int argc, char *argv[])
{
    int rc = 0;
    char nodeName[SQL_INSTNAME_SZ + 1] = { 0 };
    char serverWorkingPath[SQL_PATH_SZ + 1] = { 0 };
    char rolledForwardDbAlias[SQL_ALIAS_SZ + 1] = { 0 };
    char dbAlias[SQL_ALIAS_SZ + 1] = { 0 };
    char user[USERID_SZ + 1] = { 0 };
    char pswd[PSWD_SZ + 1] = { 0 };

    /* check the command line arguments */
    rc = CmdLineArgsCheck3(argc, argv, dbAlias, nodeName, user, pswd);
    CHECKRC(rc, "CmdLineArgsCheck3");
}
```

## Sample Programs with embedded SQL

```
printf("\nTHIS SAMPLE SHOWS HOW TO PERFORM ROLLFORWARD AFTER\n");
printf("RESTORE OF A DATABASE.\n");
strcpy(rolledForwardDbAlias, "RFDB");

/* attach to a local or remote instance */
rc = InstanceAttach(nodeName, user, pswd);
CHECKRC(rc, "Instance Attach");

/* get the server working path */
rc = ServerWorkingPathGet(dbAlias, serverWorkingPath);
CHECKRC(rc, "ServerWorkingPathGet");

printf("\nNOTE: Backup images will be created on the server\n");
printf("      in the directory %s,\n", serverWorkingPath);
printf("      and will not be deleted by the program.\n");

rc = DbBackupRestoreAndRollforward(dbAlias, rolledForwardDbAlias, user,
                                   pswd, serverWorkingPath);
CHECKRC(rc, "DbBackupRestoreAndRollforward");

/* Detach from the local or remote instance */
rc = InstanceDetach(nodeName);
CHECKRC(rc, "InstanceDetach");

return 0;
} /* end main */

int DbBackupRestoreAndRollforward(char dbAlias[],
                                  char rolledForwardDbAlias[],
                                  char user[], char pswd[],
                                  char serverWorkingPath[])
{
    int rc = 0;
    struct sqlca sqlca;
    db2CfgParam cfgParameters[1];
    db2Cfg cfgStruct;
    unsigned short logretain = 0;
    char restoreTimestamp[SQLU_TIME_STAMP_LEN + 1] = { 0 };
    db2BackupStruct backupStruct;
    db2TablespaceStruct tablespaceStruct;
    db2MediaListStruct mediaListStruct;
    db2Uint32 backupImageSize;
    db2RestoreStruct restoreStruct;
    db2TablespaceStruct rtablespaceStruct;
    db2MediaListStruct rmediaListStruct;
    db2RfwdInputStruct rfwdInput;
    db2RfwdOutputStruct rfwdOutput;
    db2RollforwardStruct rfwdStruct;
    char rollforwardAppId[SQLU_APPLID_LEN + 1] = { 0 };
    sqlint32 numReplies = 0;
    struct sqlurf_info nodeInfo;

    printf("\n*****\n");
    printf("*** ROLLFORWARD RECOVERY ***\n");
    printf("*****\n");
    printf("\nUSE THE DB2 APIs:\n");
    printf("  db2CfgSet -- Set Configuration\n");
    printf("  db2Backup -- Backup Database\n");
    printf("  sqlecrea -- Create Database\n");
    printf("  db2Restore -- Restore Database\n");
    printf("  db2Rollforward -- Rollforward Database\n");
    printf("  sqledrpd -- Drop Database\n");
    printf("TO BACK UP, RESTORE, AND ROLLFORWARD A DATABASE.\n");
    printf("\n  Update \'%s\' database configuration:\n", dbAlias);
    printf("    - Enable the configuration parameter LOGRETAIN\n");
    printf("      i.e., set LOGRETAIN = RECOVERY/YES\n");
```

```

/* initialize cfgParameters */
cfgParameters[0].flags = 0;
cfgParameters[0].token = SQLF_DBTN_LOG_RETAIN;
cfgParameters[0].ptrvalue = (char *)&logretain;

/* enable the configuration parameter 'logretain' */
logretain = SQLF_LOGRETAIN_RECOVERY;

/* initialize cfgStruct */
cfgStruct.numItems = 1;
cfgStruct.paramArray = cfgParameters;
cfgStruct.flags = db2CfgDatabase | db2CfgDelayed;
cfgStruct.dbname = dbAlias;

/* get database configuration */
db2CfgSet(db2Version810, (void *)&cfgStruct, &sqlca);
DB2_API_CHECK("Db Log Retain -- Enable");

/*****
/*   BACKUP THE DATABASE   */
*****/

/* Calling the routine for database backup */
rc = DbBackup(dbAlias, user, pswd, serverWorkingPath, &backupStruct);
CHECKRC(rc, "DbBackup");

/* To restore a remote database, you will first need to create an empty database
   if the client's code page is different from the server's code page.
   If this is the case, uncomment the call to DbCreate(). It will create
   an empty database on the server with the server's code page. */

/*
rc = DbCreate(dbAlias, rolledForwardDbAlias);
CHECKRC(rc, "DbCreate");
*/

/*****
/*   RESTORE THE DATABASE   */
*****/
strcpy(restoreTimestamp, backupStruct.oTimestamp);
rtablespaceStruct tablespaces = NULL;
rtablespaceStruct.numTablespaces = 0;
rmediaListStruct.locations = &serverWorkingPath;
rmediaListStruct.numLocations = 1;
rmediaListStruct.locationType = SQLU_LOCAL_MEDIA;
restoreStruct.piSourceDBAlias = dbAlias;
restoreStruct.piTargetDBAlias = rolledForwardDbAlias;
restoreStruct.piTimestamp = restoreTimestamp;
restoreStruct.piTargetDBPath = NULL;
restoreStruct.piReportFile = NULL;
restoreStruct.piTablespaceList = &rtablespaceStruct;
restoreStruct.piMediaList = &rmediaListStruct;
restoreStruct.piUsername = user;
restoreStruct.piPassword = pswd;
restoreStruct.piNewLogPath = NULL;
restoreStruct.piVendorOptions = NULL;
restoreStruct.iVendorOptionsSize = 0;
restoreStruct.iParallelism = 1;
restoreStruct.iBufferSize = 1024; /* 1024 x 4KB */
restoreStruct.iNumBuffers = 2;
restoreStruct.iCallerAction = DB2RESTORE_RESTORE;
restoreStruct.iOptions =
DB2RESTORE_OFFLINE | DB2RESTORE_DB | DB2RESTORE_NODATALINK |
DB2RESTORE_ROLLFWD;

printf("\n Restoring a database ...\n");

```

## Sample Programs with embedded SQL

```
printf("    - source image alias      : %s\n", dbAlias);
printf("    - source image time stamp: %s\n", restoreTimestamp);
printf("    - target database          : %s\n", rolledForwardDbAlias);

/* The API db2Restore is used to restore a database that has been backed
   up using the API db2Backup. */
db2Restore(db2Version810, &restoreStruct, &sqlca);
DB2_API_CHECK("database restore -- start");
while (sqlca.sqlcode != 0)

{

    /* continue the restore operation */
    printf("\n Continuing the restore operation...\n");

    /* Depending on the sqlca.sqlcode value, user action may be
       required, such as mounting a new tape. */
    restoreStruct.iCallerAction = DB2RESTORE_CONTINUE;

    /* restore the database */
    db2Restore(db2Version810, &restoreStruct, &sqlca);
    DB2_API_CHECK("database restore -- continue");
}
printf("\n Restore finished.\n");

/*****
/*    ROLLFORWARD RECOVERY    */
*****/
printf("\n Rolling forward database '%s'...\n", rolledForwardDbAlias);
rfdInput.iVersion = SQLUM_RFWD_VERSION;
rfdInput.piDbAlias = rolledForwardDbAlias;
rfdInput.iCallerAction = DB2ROLLFORWARD_RFWD_STOP;
rfdInput.piStopTime = SQLUM_INFINITY_TIMESTAMP;
rfdInput.piUserName = user;
rfdInput.piPassword = pswd;
rfdInput.piOverflowLogPath = serverWorkingPath;
rfdInput.iNumChngLgOvrflw = 0;
rfdInput.piChngLogOvrflw = NULL;
rfdInput.iConnectMode = DB2ROLLFORWARD_OFFLINE;
rfdInput.piTablespaceList = NULL;
rfdInput.iAllNodeFlag = DB2_ALL_NODES;
rfdInput.iNumNodes = 0;
rfdInput.piNodeList = NULL;
rfdInput.piDroppedTblID = NULL;
rfdInput.piExportDir = NULL;
rfdInput.iNumNodeInfo = 1;
rfdInput.iRollforwardFlags = DB2ROLLFORWARD_EMPTY_FLAG;
rfdOutput.poApplicationId = rollforwardAppId;
rfdOutput.poNumReplies = &numReplies;
rfdOutput.poNodeInfo = &nodeInfo;
rfdStruct.piRfdInput = &rfdInput;
rfdStruct.poRfdOutput = &rfdOutput;

/* rollforward database */
/* The API db2Rollforward rollforward recovers a database by
   applying transactions recorded in the database log files. */
db2Rollforward(db2Version810, &rfdStruct, &sqlca);
DB2_API_CHECK("rollforward -- start");
printf(" Rollforward finished.\n");

/* drop the restored database */
rc = DbDrop(rolledForwardDbAlias);
CHECKRC(rc, "DbDrop");
return 0;
} /* DbBackupRestoreAndRollforward */
```

---

## Appendix F. Tivoli Storage Manager

When calling the `BACKUP DATABASE` or `RESTORE DATABASE` commands, you can specify that you want to use the Tivoli Storage Manager (TSM) product to manage database or table space backup or restore operation. The minimum required level of TSM client API is Version 4.2.0, except on the following:

- 64-bit Solaris systems which require TSM client API Version 4.2.1.
- 64-bit Windows NT operating systems which require TSM client API Version 5.1.
- SuSE Linux Enterprise Server 8 for IA64 and LinuxPPC which require a minimum TSM client API Version 5.1.5.
- 64-bit Linux on AMD Opteron systems which require a minimum TSM client API Version 5.2.0.
- 64-bit Linux for S/390 which requires a minimum TSM client API Version 5.2.2.

---

### Configuring a Tivoli Storage Manager client

Before the database manager can use the TSM option, the following steps may be required to configure the TSM environment:

1. A functioning TSM client and server must be installed and configured. In addition, the TSM client API must be installed on each DB2 server.
2. Set the environment variables used by the TSM client API:

**DSMI\_DIR** Identifies the user-defined directory path where the API trusted agent file (`dsmtca`) is located.

**DSMI\_CONFIG** Identifies the user-defined directory path to the `dsm.opt` file, which contains the TSM user options. Unlike the other two variables, this variable should contain a fully qualified path and file name.

**DSMI\_LOG** Identifies the user-defined directory path where the error log (`dsierror.log`) will be created.

**Note:** In a multi-partition database environment these settings must be specified in the `sqllib/userprofile` directory.

3. If any changes are made to these environment variables and the database manager is running, you should:
  - Stop the database manager using the **db2stop** command.
  - Start the database manager using the **db2start** command.
4. Depending on the server's configuration, a Tivoli client may require a password to interface with a TSM server. If the TSM environment is configured to use `PASSWORDACCESS=generate`, the Tivoli client needs to have its password established.

The executable file `dsmapipw` is installed in the `sqllib/adsm` directory of the instance owner. This executable allows you to establish and reset the TSM password.

To execute the `dsmapipw` command, you must be logged in as the local administrator or "root" user. When this command is executed, you will be prompted for the following information:

- *Old password*, which is the current password for the TSM node, as recognized by the TSM server. The first time you execute this command, this password will be the one provided by the TSM administrator at the time your node was registered on the TSM server.
- *New password*, which is the new password for the TSM node, stored at the TSM server. (You will be prompted twice for the new password, to check for input errors.)

**Note:** Users who invoke the BACKUP DATABASE or RESTORE DATABASE commands do not need to know this password. You only need to run the dsmapipw command to establish a password for the initial connection, and after the password has been reset on the TSM server.

## Considerations for using Tivoli Storage Manager

To use specific features within TSM, you may be required to give the fully qualified path name of the object using the feature. (Remember that on Windows operating systems, the \ will be used instead of /.) The fully qualified path name of:

- A full database backup object is:  
/<database>/NODEnnnn/FULL\_BACKUP.timestamp.seq\_no
- An incremental database backup object is:  
/<database>/NODEnnnn/DB\_INCR\_BACKUP.timestamp.seq\_no
- An incremental delta database backup object is\_  
/<database>/NODEnnnn/DB\_DELTA\_BACKUP.timestamp.seq\_no
- A full table space backup object is:  
/<database>/NODEnnnn/TSP\_BACKUP.timestamp.seq\_no
- An incremental table space backup object is:  
/<database>/NODEnnnn/TSP\_INCR\_BACKUP.timestamp.seq\_no
- An incremental delta table space backup object is:  
/<database>/NODEnnnn/TSP\_DELTA\_BACKUP.timestamp.seq\_no

where <database> is the database alias name, and NODEnnnn is the node number. The names shown in uppercase characters must be entered as shown.

- In the case where you have multiple backup images using the same database alias name, the time stamp and sequence number become the distinguishing part of a fully qualified name. You will need to query TSM to determine which backup version to use.
- Individual backup images are pooled into file spaces that TSM manages. Individual backup images can only be manipulated through the TSM APIs, or through **db2adutl** which uses these APIs.
- The TSM server will time out a session if the Tivoli client does not respond within the period of time specified by the COMMTIMEOUT parameter in the server's configuration file. Three factors can contribute to a timeout problem:
  - The COMMTIMEOUT parameter may be set too low at the TSM server. For example, during a restore operation, a timeout can occur if large DMS table spaces are being created. The recommended value for this parameter is 6000 seconds.
  - The DB2 backup or restore buffer may be too large.
  - Database activity during an online backup operation may be too high.
- Use multiple sessions to increase throughput (only if sufficient hardware is available on the TSM server).



**Related concepts:**

- “Managing log files” on page 42
- “Tivoli Space Manager Hierarchical Storage Manager (AIX)” in the *Quick Beginnings for Data Links Manager*

**Related reference:**

- “db2adutl - Managing DB2 objects within TSM” on page 237



---

## Appendix G. User exit for database recovery

You can develop a *user exit program* to automate log file archiving and retrieval. Before invoking a user exit program for log file archiving or retrieval, ensure that the *logarchmeth1* database configuration parameter has been set to USEREXIT. This also enables your database for rollforward recovery.

When a user exit program is invoked, the database manager passes control to the executable file, *db2uext2*. The database manager passes parameters to *db2uext2* and, on completion, the program passes a return code back to the database manager. Because the database manager handles a limited set of return conditions, the user exit program should be able to handle error conditions (see “Error handling” on page 331). And because only one user exit program can be invoked within a database manager instance, it must have a section for each of the operations it may be asked to perform.

The following topics are covered:

- “Sample user exit programs”
- “Calling format” on page 330
- “Error handling” on page 331

---

### Sample user exit programs

Sample user exit programs are provided for all supported platforms. You can modify these programs to suit your particular requirements. The sample programs are well commented with information that will help you to use them most effectively.

You should be aware that user exit programs must *copy* log files from the active log path to the archive log path. Do not remove log files from the active log path. (This could cause problems during database recovery.) DB2<sup>®</sup> removes archived log files from the active log path when these log files are no longer needed for recovery.

Following is a description of the sample user exit programs that are shipped with DB2.

- **UNIX<sup>®</sup> based systems**

The user exit sample programs for DB2 for UNIX based systems are found in the *sql11ib/samples/c* subdirectory. Although the samples provided are coded in C, your user exit program can be written in a different programming language.

Your user exit program must be an executable file whose name is *db2uext2*.

There are four sample user exit programs for UNIX based systems:

- *db2uext2.ctsm*

This sample uses Tivoli<sup>®</sup> Storage Manager to archive and retrieve database log files.

- *db2uext2.ctape*

This sample uses tape media to archive and retrieve database log files .

- *db2uext2.cdisk*

This sample uses the operating system COPY command and disk media to archive and retrieve database log files.

- db2uxt2.cxbsa

This sample works with the XBSA Draft 0.8 published by the X/Open group. It can be used to archive and retrieve database log files. This sample is only supported on AIX.

- **Windows® operating systems**

The user exit sample programs for DB2 for Windows operating systems are found in the `sql11ib\samples\c` subdirectory. Although the samples provided are coded in C, your user exit program can be written in a different programming language.

Your user exit program must be an executable file whose name is `db2uxt2`.

There are two sample user exit programs for Windows operating systems:

- db2uxt2.ctsm

This sample uses Tivoli Storage Manager to archive and retrieve database log files.

- db2uxt2.cdisk

This sample uses the operating system COPY command and disk media to archive and retrieve database log files.

---

## Calling format

When the database manager calls a user exit program, it passes a set of parameters (of data type CHAR) to the program. The calling format is dependent on your operating system:

```
db2uxt2 -OS<os> -RL<db2rel> -RQ<request> -DB<dbname>
-NN<nodenum> -LP<logpath> -LN<logname> -AP<tsmpasswd>
-SP<startpage> -LS<logsize>
```

<b>os</b>	Specifies the platform on which the instance is running. Valid values are: AIX®, Solaris, HP-UX, SCO, Linux, and NT.
<b>db2rel</b>	Specifies the DB2 release level. For example, SQL07020.
<b>request</b>	Specifies a request type. Valid values are: ARCHIVE and RETRIEVE.
<b>dbname</b>	Specifies a database name.
<b>nodenum</b>	Specifies the local node number, such as 5, for example.
<b>logpath</b>	Specifies the fully qualified path to the log files. The path must contain the trailing path separator. For example, <code>/u/database/log/path/</code> , or <code>d:\logpath\</code> .
<b>logname</b>	Specifies the name of the log file that is to be archived or retrieved, such as <code>S0000123.LOG</code> , for example.
<b>tsmpasswd</b>	Specifies the TSM password. (If a value for the database configuration parameter <code>tsm_password</code> has previously been specified, that value is passed to the user exit program.)
<b>startpage</b>	Specifies the number of 4-KB offset pages of the device at which the log extent starts.
<b>logsize</b>	Specifies the size of the log extent, in 4-KB pages. This parameter is only valid if a raw device is used for logging.

---

## Error handling

Your user exit program should be designed to provide specific and meaningful return codes, so that the database manager can interpret them correctly. Because the user exit program is called by the underlying operating system command processor, the operating system itself could return error codes. And because these error codes are not remapped, use the operating system message help utility to obtain information about them.

Table 8 shows the codes that can be returned by a user exit program, and describes how these codes are interpreted by the database manager. If a return code is not listed in the table, it is treated as if its value were 32.

*Table 8. User Exit Program Return Codes.* Applies to archiving and retrieval operations only.

Return Code	Explanation
0	Successful.
4	Temporary resource error encountered. <sup>a</sup>
8	Operator intervention is required. <sup>a</sup>
12	Hardware error. <sup>b</sup>
16	Error with the user exit program or a software function used by the program. <sup>b</sup>
20	Error with one or more of the parameters passed to the user exit program. Verify that the user exit program is correctly processing the specified parameters. <sup>b</sup>
24	The user exit program was not found. <sup>b</sup>
28	Error caused by an input/output (I/O) failure, or by the operating system. <sup>b</sup>
32	The user exit program was terminated by the user. <sup>b</sup>
255	Error caused by the user exit program not being able to load the library file for the executable. <sup>c</sup>

Table 8. User Exit Program Return Codes (continued). Applies to archiving and retrieval operations only.

Return Code	Explanation
	<p><sup>a</sup> For archiving or retrieval requests, a return code of 4 or 8 causes a retry in five minutes. If the user exit program continues to return 4 or 8 on retrieve requests for the same log file, DB2 will continue to retry until successful. (This applies to rollforward operations, or calls to the <b>db2ReadLog</b> API, which is used by the replication utility.)</p>
	<p><sup>b</sup> User exit requests are suspended for five minutes. During this time, all requests are ignored, including the request that caused the error condition. Following this five-minute suspension, the next request is processed. If this request is processed without error, processing of new user exit requests continues, and DB2 reissues the archive request that failed or was suspended previously. If a return code greater than 8 is generated during the retry, requests are suspended for an additional five minutes. The five-minute suspensions continue until the problem is corrected, or the database is stopped and restarted. Once all applications have disconnected from the database, DB2 issues an archive request for any log file that may not have been successfully archived previously. If the user exit program fails to archive log files, your disk may become filled with log files, and performance may be degraded. Once the disk becomes full, the database manager will not accept further application requests for database updates. If the user exit program was called to retrieve log files, rollforward recovery is suspended, but not stopped, unless the ROLLFORWARD STOP option was specified. If the STOP option was not specified, you can correct the problem and resume recovery.</p>
	<p><sup>c</sup> If the user exit program returns error code 255, it is likely that the program cannot load the library file for the executable. To verify this, manually invoke the user exit program. More information is displayed.</p>
	<p><b>Note:</b> During archiving and retrieval operations, an alert message is issued for all return codes except 0, and 4. The alert message contains the return code from the user exit program, and a copy of the input parameters that were provided to the user exit program.</p>

---

## Appendix H. Backup and restore APIs for vendor products

---

### APIs for backup and restore to storage managers

DB2 provides interfaces that can be used by third-party media management products to store and retrieve data for backup and restore operations and log files. This function is designed to augment the backup, restore, and log archiving data targets of diskette, disk, tape, and Tivoli Storage Manager, that are supported as a standard part of DB2.

These third-party media management products will be referred to as vendor products in the remainder of this appendix.

DB2 defines a set of function prototypes that provide a general purpose data interface to backup, restore, and log archiving that can be used by many vendors. These functions are to be provided by the vendor in a shared library on UNIX based systems, or DLL on the Windows operating system. When the functions are invoked by DB2, the shared library or DLL specified by the calling backup, restore, or log archiving routine is loaded and the functions provided by the vendor are called to perform the required tasks.

Sample files demonstrating the DB2 vendor functionality are located on UNIX platforms in the `sql1lib/samples/BARVendor` directory, and on Windows in the `sql1lib\samples\BARVendor` directory.

### Operational overview

Seven functions are defined to interface DB2 and the vendor product:

- `sqluvint` - Initialize and Link to Device
- `sqluvget` - Reading Data from Device
- `sqluvput` - Writing Data to Device
- `sqluvend` - Unlink the Device
- `sqluvdel` - Delete Committed Session
- `db2VendorQueryApiVersion` - Query Device Supported API Level
- `db2VendorGetNextObj` - Get Next Object on Device

DB2 will call these functions, and they should be provided by the vendor product in a shared library on UNIX based systems, or in a DLL on the Windows operating system.

**Note:** The shared library or DLL code will be run as part of the database engine code. Therefore, it must be reentrant and thoroughly debugged. An errant function may compromise data integrity of the database.

The sequence of functions that DB2 will call during a specific backup or restore operation depends on:

- The number of sessions that will be utilized.
- Whether it is a backup, a restore, a log archive, or a log retrieve operation.
- The PROMPTING mode that is specified on the backup or restore operation.
- The characteristics of the device on which the data is stored.

## APIs for backup and restore to storage managers

- The errors that may be encountered during the operation.

### Number of sessions

DB2 supports the backup and restore of database objects using one or more data streams or sessions. A backup or restore using three sessions would require three physical or logical devices to be available. When vendor device support is being used, it is the vendor's functions that are responsible for managing the interface to each physical or logical device. DB2 simply sends or receives data buffers to or from the vendor provided functions.

The number of sessions to be used is specified as a parameter by the application that calls the backup or restore database function. This value is provided in the INIT-INPUT structure used by sqluvint.

DB2 will continue to initialize sessions until the specified number is reached, or it receives an SQLUV\_MAX\_LINK\_GRANT warning return code from an sqluvint call. In order to warn DB2 that it has reached the maximum number of sessions that it can support, the vendor product will require code to track the number of active sessions. Failure to warn DB2 could lead to a DB2 initialize session request that fails, resulting in a termination of all sessions and the failure of the entire backup or restore operation.

When the operation is backup, DB2 writes a media header record at the beginning of each session. The record contains information that DB2 uses to identify the session during a restore operation. DB2 uniquely identifies each session by appending a sequence number to the name of the backup image. The number starts at one for the first session, and is incremented by one each time another session is initiated with an sqluvint call for a backup or a restore operation.

When the backup operation completes successfully, DB2 writes a media trailer to the last session it closes. This trailer includes information that tells DB2 how many sessions were used to perform the backup operation. During a restore operation, this information is used to ensure all the sessions, or data streams, have been restored.

### Operation with no errors, warnings, or prompting

For backup, the following sequence of calls is issued by DB2 for *each* session.

```
sqluvint, action = SQLUV_WRITE
```

followed by 1 to n

```
sqluvput
```

followed by 1

```
sqluvend, action = SQLUV_COMMIT
```

When DB2 issues an sqluvend call (action SQLUV\_COMMIT), it expects the vendor product to appropriately save the output data. A return code of SQLUV\_OK to DB2 indicates success.

The DB2-INFO structure, used on the sqluvint call, contains the information required to identify the backup. A sequence number is supplied. The vendor product may choose to save this information. DB2 will use it during restore to identify the backup that will be restored.

For restore, the sequence of calls for each session is:



## APIs for backup and restore to storage managers

```
sqluvint, action = SQLUV_READ
```

followed by 1 to n

```
sqluvget
```

followed by 1

```
sqluvend, action = SQLUV_COMMIT
```

The information in the DB2-INFO structure used on the sqluvint call will contain the information required to identify the backup. A sequence number is not supplied. DB2 expects that all backup objects (session outputs committed during a backup) will be returned. The first backup object returned is the object generated with sequence number 1, and all other objects are restored in no specific order. DB2 checks the media tail to ensure that all objects have been processed.

**Note:** Not all vendor products will keep a record of the names of the backup objects. This is most likely when the backups are being done to tapes, or other media of limited capacity. During the initialization of restore sessions, the identification information can be utilized to stage the necessary backup objects so that they are available when required; this may be most useful when juke boxes or robotic systems are used to store the backups. DB2 will always check the media header (first record in each session's output) to ensure that the correct data is being restored.

### Prompting mode

When a backup or a restore operation is initiated, two prompting modes are possible:

- WITHOUT PROMPTING or NOINTERRUPT, where there is no opportunity for the vendor product to write messages to the user, or for the user to respond to them.
- PROMPTING or INTERRUPT, where the user can receive and respond to messages from the vendor product.

For PROMPTING mode, backup and restore define three possible user responses:

- Continue  
The operation of reading or writing data to the device will resume.
- Device terminate  
The device will receive no additional data, and the session is terminated.
- Terminate  
The entire backup or restore operation is terminated.

The use of the PROMPTING and WITHOUT PROMPTING modes is discussed in the sections that follow.

### Device characteristics

For purposes of the vendor device support APIs, two general types of devices are defined:

- Limited capacity devices requiring user action to change the media; for example, a tape drive, diskette, or CDROM drive.
- Very large capacity devices, where normal operations do not require the user to handle media; for example, a juke box, or an intelligent robotic media handling device.

## APIs for backup and restore to storage managers

A limited capacity device may require that the user be prompted to load additional media during the backup or restore operation. Generally DB2 is not sensitive to the order in which the media is loaded for either backup or restore operations. It also provides facilities to pass vendor media handling messages to the user. This prompting requires that the backup or restore operation be initiated with PROMPTING on. The media handling message text is specified in the description field of the return code structure.

If PROMPTING is on, and DB2 receives an SQLUV\_ENDOFMEDIA or an SQLUV\_ENDOFMEDIA\_NO\_DATA return code from a sqluvput (write) or a sqluvget (read) call, DB2:

- Marks the last buffer sent to the session to be resent, if the call was sqluvput. It will be put to a session later.
- Calls the session with sqluwend (action = SQLUV\_COMMIT). If successful (SQLUV\_OK return code), DB2:
  - Sends a vendor media handling message to the user from the return code structure that signaled the end-of-media condition.
  - Prompts the user for a continue, device terminate, or terminate response.
- If the response is *continue*, DB2 initializes another session using the sqluvint call, and if successful, begins writing data to or reading data from the session. To uniquely identify the session when writing, DB2 increments the sequence number. The sequence number is available in the DB2-INFO structure used with sqluvint, and is in the media header record, which is the first data record sent to the session.

DB2 will not start more sessions than requested when a backup or a restore operation is started, or indicated by the vendor product with a SQLUV\_MAX\_LINK\_GRANT warning on an sqluvint call.

- If the response is *device terminate*, DB2 does not attempt to initialize another session, and the number of active sessions is reduced by one. DB2 does not allow all sessions to be terminated by device terminate responses; at least one session must be kept active until the backup or the restore operation completes.
- If the response is *terminate*, DB2 terminates the backup or the restore operation. For more information on exactly what DB2 does to terminate the sessions, see “If error conditions are returned to DB2” on page 337.

Because backup or restore performance is often dependent on the number of devices being used, it is important that parallelism be maintained. For backup operations, users are encouraged to respond with a *continue*, unless they know that the remaining active sessions will hold the data that is still to be written out. For restore operations, users are also encouraged to respond with a *continue* until all media have been processed.

If the backup or the restore mode is WITHOUT PROMPTING, and DB2 receives an SQLUV\_ENDOFMEDIA or an SQLUV\_ENDOFMEDIA\_NO\_DATA return code from a session, it will terminate the session and not attempt to open another session. If all sessions return end-of-media to DB2 before the backup or the restore operation is complete, the operation will fail. Because of this, WITHOUT PROMPTING should be used carefully with limited capacity devices; it does, however, make sense to operate in this mode with very large capacity devices.

It is possible for the vendor product to hide media mounting and switching actions from DB2, so that the device appears to have infinite capacity. Some very large capacity devices operate in this mode. In these cases, it is critical that all the data that was backed up be returned to DB2 in the same order when a restore operation

is in progress. Failure to do so could result in missing data, but DB2 assumes a successful restore operation, because it has no way of detecting the missing data.

DB2 writes data to the vendor product with the assumption that each buffer will be contained on one and only one media (for example, a tape). It is possible for the vendor product to split these buffers across multiple media without DB2's knowledge. In this case, the order in which the media is processed during a restore operation is critical, because the vendor product will be responsible for returning reconstructed buffers from the multiple media to DB2. Failure to do so will result in a failed restore operation.

### If error conditions are returned to DB2

When performing a backup or a restore operation, DB2 expects that all sessions will complete successfully; otherwise, the entire backup or restore operation fails. A session signals successful completion to DB2 with an `SQLUV_OK` return code on the `sqluvend` call, `action = SQLUV_COMMIT`.

If unrecoverable errors are encountered, the session is terminated by DB2. These can be DB2 errors, or errors returned to DB2 from the vendor product. Because all sessions must commit successfully to have a complete backup or restore operation, the failure of one causes DB2 to terminate the other sessions associated with the operation.

If the vendor product responds to a call from DB2 with an unrecoverable return code, the vendor product can optionally provide additional information, using message text placed in the description field of the `RETURN-CODE` structure. This message text is presented to the user, along with the DB2 information, so that corrective action can be taken.

There will be backup scenarios in which a session has committed successfully, and another session associated with the backup operation experiences an unrecoverable error. Because all sessions must complete successfully before a backup operation is considered successful, DB2 must delete the output data in the committed sessions: DB2 issues a `sqlvdel` call to request deletion of the object. This call is not considered an I/O session, and is responsible for initializing and terminating any connection that may be necessary to delete the backup object.

The `DB2-INFO` structure will not contain a sequence number; `sqlvdel` will delete all backup objects that match the remaining parameters in the `DB2-INFO` structure.

### Warning conditions

It is possible for DB2 to receive warning return codes from the vendor product; for example, if a device is not ready, or some other correctable condition has occurred. This is true for both read and write operations.

On `sqlvput` and `sqlvget` calls, the vendor can set the return code to `SQLUV_WARNING`, and optionally provide additional information, using message text placed in the description field of the `RETURN-CODE` structure. This message text is presented to the user so that corrective action can be taken. The user can respond in one of three ways: *continue*, *device terminate*, or *terminate*:

- If the response is *continue*, DB2 attempts to rewrite the buffer using `sqlvput` during a backup operation. During a restore operation, DB2 issues an `sqlvget` call to read the next buffer.

## APIs for backup and restore to storage managers

- If the response is *device terminate* or *terminate*, DB2 terminates the entire backup or restore operation in the same way that it would respond after an unrecoverable error (for example, it will terminate active sessions and delete committed sessions).

## Operational hints and tips

This section provides some hints and tips for building vendor products.

### History file

The history file can be used as an aid in database recovery operations. It is associated with each database, and is automatically updated with each backup or restore operation. Information in the file can be viewed, updated, or pruned through the following facilities:

- Control Center
- Command line processor (CLP)
  - LIST HISTORY command
  - UPDATE HISTORY FILE command
  - PRUNE HISTORY command
- APIs
  - db2HistoryOpenScan
  - db2HistoryGetEntry
  - db2HistoryCloseScan
  - db2HistoryUpdate
  - db2Prune

For information about the layout of the file, see db2HistData.

When a backup operation completes, one or more records is written to the file. If the output of the backup operation was directed to vendor devices and the LOAD keyword was used, the DEVICE field in the history record contains an 0. If the backup operation was directed to TSM, the DEVICE field contains an A. The LOCATION field contains either:

- The vendor file name specified when the backup operation was invoked.
- The name of the shared library, if no vendor file name was specified.

For more information about specifying this option, see “Invoking a backup or a restore operation using vendor products.”

The LOCATION field can be updated using the Control Center, the CLP, or an API. The location of backup information can be updated if limited capacity devices (for example, removable media) have been used to hold the backup image, and the media is physically moved to a different (perhaps off-site) storage location. If this is the case, the history file can be used to help locate a backup image if a recovery operation becomes necessary.

## Invoking a backup or a restore operation using vendor products

Vendor products can be specified when invoking the DB2 backup or the DB2 restore utility from:

- The Control Center
- The command line processor (CLP)
- An application programming interface (API).

### The Control Center

The Control Center is the graphical user interface for database administration that is shipped with DB2.

To specify	The Control Center input variable for backup or restore operations
Use of vendor device and library name	Is <i>Use Library</i> . Specify the library name (on UNIX based systems) or the DLL name (on the Windows operating system).
Number of sessions	Is <i>Sessions</i> .
Vendor options	Is not supported.
Vendor file name	Is not supported.
Transfer buffer size	Is (for backup) <i>Size of each Buffer</i> , and (for restore) not applicable.

### The command line processor (CLP)

The command line processor (CLP) can be used to invoke the DB2 BACKUP DATABASE or the RESTORE DATABASE command.

To specify	The command line processor parameter	
	for backup is	for restore is
Use of vendor device and library name	<i>library-name</i>	<i>shared-library</i>
Number of sessions	<i>num-sessions</i>	<i>num-sessions</i>
Vendor options	not supported	not supported
Vendor file name	not supported	not supported
Transfer buffer size	<i>buffer-size</i>	<i>buffer-size</i>

### Backup and restore API function calls

Two API function calls support backup and restore operations: db2Backup for backup and db2Restore for restore.

To specify	The API parameter (for both db2Backup and db2Restore) is
Use of vendor device and library name	as follows: In structure <i>sqlu_media_list</i> , specify a media type of <code>SQLU_OTHER_MEDIA</code> , and then in structure <i>sqlu_vendor</i> , specify a shared library or DLL in <i>shr_lib</i> .
Number of sessions	as follows: In structure <i>sqlu_media_list</i> , specify <i>sessions</i> .
Vendor options	<i>PVendorOptions</i>
Vendor file name	as follows: In structure <i>sqlu_media_list</i> , specify a media type of <code>SQLU_OTHER_MEDIA</code> , and then in structure <i>sqlu_vendor</i> , specify a file name in <i>filename</i> .
Transfer buffer size	<i>BufferSize</i>

**Related reference:**

## APIs for backup and restore to storage managers

- “sqluvint - Initialize and Link to Device” on page 340
- “sqluvget - Reading Data from Device” on page 343
- “sqluvput - Writing Data to Device” on page 344
- “sqluvend - Unlink the Device and Release its Resources” on page 346
- “sqluvdel - Delete Committed Session” on page 348
- “DB2-INFO” on page 351
- “VENDOR-INFO” on page 353
- “INIT-INPUT” on page 354
- “INIT-OUTPUT” on page 354
- “DATA” on page 355
- “RETURN-CODE” on page 355
- “db2VendorQueryApiVersion - Query Device Supported API Level” on page 349
- “db2VendorGetNextObj - Get Next Object on Device” on page 349

---

## sqluvint - Initialize and Link to Device

This function is called to provide information for initialization and establishment of a logical link between DB2 and the vendor device.

### Authorization:

One of the following:

- *sysadm*
- *dbadm*

### Required connection:

Database

### API include file:

*sql.h*

### C API syntax:

```
/* File: sqluvend.h */
/* API: Initialize and Link to Device */
/* ... */
int sqluvint (
    struct Init_input  *,
    struct Init_output *,
    struct Return_code *);
/* ... */
```

### API parameters:

#### Init\_input

Input. Structure that contains information provided by DB2 to establish a logical link with the vendor device.

#### Init\_output

Output. Structure that contains the output returned by the vendor device.

**Return\_code**

Output. Structure that contains the return code to be passed to DB2, and a brief text explanation.

**Usage notes:**

For each media I/O session, DB2 will call this function to obtain a device handle. If for any reason, the vendor function encounters an error during initialization, it will indicate it via a return code. If the return code indicates an error, DB2 may choose to terminate the operation by calling the **sqluvend** function. Details on possible return codes, and the DB2 reaction to each of these, is contained in the return codes table (see Table 9).

The INIT-INPUT structure contains elements that can be used by the vendor product to determine if the backup or restore can proceed:

- size\_HI\_order and size\_LOW\_order  
This is the estimated size of the backup. They can be used to determine if the vendor devices can handle the size of the backup image. They can be used to estimate the quantity of removable media that will be required to hold the backup. It might be beneficial to fail at the first **sqluvint** call if problems are anticipated.
- req\_sessions  
The number of user requested sessions can be used in conjunction with the estimated size and the prompting level to determine if the backup or restore operation is possible.
- prompt\_lvl  
The prompting level indicates to the vendor if it is possible to prompt for actions such as changing removable media (for example, put another tape in the tape drive). This might suggest that the operation cannot proceed since there will be no way to prompt the user.  
If the prompting level is WITHOUT PROMPTING and the quantity of removable media is greater than the number of sessions requested, DB2 will not be able to complete the operation successfully.

DB2 names the backup being written or the restore to be read via fields in the DB2-INFO structure. In the case of an action = SQLUV\_READ, the vendor product must check for the existence of the named object. If it cannot be found, the return code should be set to SQLUV\_OBJ\_NOT\_FOUND so that DB2 will take the appropriate action.

After initialization is completed successfully, DB2 will continue by issuing other data transfer functions, but may terminate the session at any time with an **sqluvend** call.

**Return codes:**

*Table 9. Valid Return Codes for sqluvint and Resulting DB2 Action*

Literal in Header File	Description	Probable Next Call	Other Comments
SQLUV_OK	Operation successful.	sqluvput, sqluvget (see comments)	If action = SQLUV_WRITE, the next call will be sqluvput (to BACKUP data). If action = SQLUV_READ, verify the existence of the named object prior to returning SQLUV_OK; the next call will be sqluvget to RESTORE data.
SQLUV_LINK_EXIST	Session activated previously.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.

## sqluvint - Initialize and Link to Device

Table 9. Valid Return Codes for sqluvint and Resulting DB2 Action (continued)

Literal in Header File	Description	Probable Next Call	Other Comments
SQLUV_COMM_ERROR	Communication error with device.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_INV_VERSION	The DB2 and vendor products are incompatible.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_INV_ACTION	Invalid action is requested. This could also be used to indicate that the combination of parameters results in an operation which is not possible.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_NO_DEV_AVAIL	No device is available for use at the moment.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_OBJ_NOT_FOUND	Object specified cannot be found. This should be used when the action on the sqluvint call is 'R' (read) and the requested object cannot be found based on the criteria specified in the DB2-INFO structure.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_OBJS_FOUND	More than 1 object matches the specified criteria. This will result when the action on the sqluvint call is 'R' (read) and more than one object matches the criteria in the DB2-INFO structure.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_INV_USERID	Invalid userid specified.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_INV_PASSWORD	Invalid password provided.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_INV_OPTIONS	Invalid options encountered in the vendor options field.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_INIT_FAILED	Initialization failed and the session is to be terminated.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_DEV_ERROR	Device error.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_MAX_LINK_GRANT	Max number of links established.	sqluvput, sqluvget (see comments)	This is treated as a warning by DB2. The warning tells DB2 not to open additional sessions with the vendor product, because the maximum number of sessions it can support has been reached (note: this could be due to device availability). If action = SQLUV_WRITE (BACKUP), the next call will be sqluvput. If action = SQLUV_READ, verify the existence of the named object prior to returning SQLUV_MAX_LINK_GRANT; the next call will be sqluvget to RESTORE data.
SQLUV_IO_ERROR	I/O error.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.
SQLUV_NOT_ENOUGH_SPACE	There is not enough space to store the entire backup image; the size estimate is provided as a 64-bit value in bytes.	no further calls	Session initialization fails. Free up memory allocated for this session and terminate. A sqluvend call will not be received, since the session was never established.



## sqluvget - Reading Data from Device

After initialization, this function can be called to read data from the device.

### Authorization:

One of the following:

- *sysadm*
- *dbadm*

### Required connection:

Database

### API include file:

*sqluvend.h*

### C API syntax:

```
/* File: sqluvend.h */
/* API: Reading Data from Device */
/* ... */
int sqluvget (
    void * pVendorCB,
    struct Data      *,
    struct Return_code *);
/* ... */
```

```
typedef struct Data
}
    sqlint32  obj_num;
    sqlint32  buff_size;
    sqlint32  actual_buff_size;
    void      *dataptr;
    void      *reserve;
{ Data;
```

### API parameters:

#### **pVendorCB**

Input. Pointer to space allocated for the DATA structure (including the data buffer) and Return\_code.

**Data** Input/output. A pointer to the *data* structure.

#### **Return\_code**

Output. The return code from the API call.

#### **obj\_num**

Specifies which backup object should be retrieved.

#### **buff\_size**

Specifies the buffer size to be used.

#### **actual\_buff\_size**

Specifies the actual bytes read or written. This value should be set to output to indicate how many bytes of data were actually read.

#### **dataptr**

A pointer to the data buffer.

## sqluvget - Reading Data from Device

### reserve

Reserved for future use.

### Usage notes:

This function is used by the restore utility.

### Return codes:

Table 10. Valid Return Codes for sqluvget and Resulting DB2 Action

Literal in Header File	Description	Probable Next Call	Other Comments
SQLUV_OK	Operation successful.	sqluvget	DB2 processes the data
SQLUV_COMM_ERROR	Communication error with device.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_INV_ACTION	Invalid action is requested.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_INV_DEV_HANDLE	Invalid device handle.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_INV_BUFF_SIZE	Invalid buffer size specified.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_DEV_ERROR	Device error.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_WARNING	Warning. This should not be used to indicate end-of-media to DB2; use SQLUV_ENDOFMEDIA or SQLUV_ENDOFMEDIA_NO_DATA for this purpose. However, device not ready conditions can be indicated using this return code.	sqluvget, or sqluvend, action = SQLU_ABORT	
SQLUV_LINK_NOT_EXIST	No link currently exists.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_MORE_DATA	Operation successful; more data available.	sqluvget	
SQLUV_ENDOFMEDIA_NO_DATA	End of media and 0 bytes read (for example, end of tape).	sqluvend	
SQLUV_ENDOFMEDIA	End of media and > 0 bytes read, (for example, end of tape).	sqluvend	DB2 processes the data, and then handles the end-of-media condition.
SQLUV_IO_ERROR	I/O error.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
<b>Next call:</b>			
<sup>a</sup> If the next call is an sqluvend, action = SQLU_ABORT, this session and all other active sessions will be terminated.			

## sqluvput - Writing Data to Device

After initialization, this function can be used to write data to the device.

### Authorization:

One of the following:

- *sysadm*
- *dbadm*

### Required connection:

Database

### API include file:

*sqluvend.h*

### C API syntax:

```

/* File: sqluvend.h */
/* API: Writing Data to Device */
/* ... */
int sqluvput (
    void * pVendorCB,
    struct Data *,
    struct Return_code *);
/* ... */

typedef struct Data
}
    sqlint32  obj_num;
    sqlint32  buff_size;
    sqlint32  actual_buff_size;
    void      *dataptr;
    void      *reserve;
{ Data;

```

**API parameters:**

**pVendorCB**

Input. Pointer to space allocated for the DATA structure (including the data buffer) and Return\_code.

**Data** Output. Data buffer filled with data to be written out.

**Return\_code**

Output. The return code from the API call.

**obj\_num**

Specifies which backup object should be retrieved.

**buff\_size**

Specifies the buffer size to be used.

**actual\_buff\_size**

Specifies the actual bytes read or written. This value should be set to indicate how many bytes of data were actually read.

**dataptr**

A pointer to the data buffer.

**reserve**

Reserved for future use.

**Usage notes:**

This function is used by the backup utility.

**Return codes:**

Table 11. Valid Return Codes for sqluvput and Resulting DB2 Action

Literal in Header File	Description	Probable Next Call	Other Comments
SQLUV_OK	Operation successful.	sqluvput or sqluvend, if complete (for example, DB2 has no more data)	Inform other processes of successful operation.
SQLUV_COMM_ERROR	Communication error with device.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_INV_ACTION	Invalid action is requested.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_INV_DEV_HANDLE	Invalid device handle.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_INV_BUFF_SIZE	Invalid buffer size specified.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_ENDOFMEDIA	End of media reached, for example, end of tape.	sqluvend	

## sqluvput - Writing Data to Device

Table 11. Valid Return Codes for sqluvput and Resulting DB2 Action (continued)

Literal in Header File	Description	Probable Next Call	Other Comments
SQLUV_DATA_RESEND	Device requested to have buffer sent again.	sqluvput	DB2 will retransmit the last buffer. This will only be done once.
SQLUV_DEV_ERROR	Device error.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_WARNING	Warning. This should not be used to indicate end-of-media to DB2; use SQLUV_ENDOFMEDIA for this purpose. However, device not ready conditions can be indicated using this return code.	sqluvput	
SQLUV_LINK_NOT_EXIST	No link currently exists.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
SQLUV_IO_ERROR	I/O error.	sqluvend, action = SQLU_ABORT <sup>a</sup>	The session will be terminated.
<b>Next call:</b>			
<sup>a</sup> If the next call is an sqluvend, action = SQLU_ABORT, this session and all other active sessions will be terminated. Committed sessions are deleted with an sqluvint, sqluvdel, and sqluvend sequence of calls.			

## sqluvend - Unlink the Device and Release its Resources

Ends or unlinks the device, and frees all of its related resources. The vendor must free or release unused resources (for example, allocated space and file handles) before returning to DB2.

### Authorization:

One of the following:

- *sysadm*
- *dbadm*

### Required connection:

Database

### API include file:

*sql.h*

### C API syntax:

```
/* File: sqluvend.h */
/* API: Unlink the Device and Release its Resources */
/* ... */
int sqluvend (
    sqlint32 action,
    void * pVendorCB,
    struct Init_output *,
    struct Return_code *);
/* ... */
```

### API parameters:

**action** Input. Used to commit or abort the session:

- SQLUV\_COMMIT ( 0 = to commit )
- SQLUV\_ABORT ( 1 = to abort )

### pVendorCB

Input. Pointer to the Init\_output structure.

## sqluvend - Unlink Device and Release Resources

### Init\_output

Output. Space for Init\_output de-allocated. The data has been committed to stable storage for a backup if action is to commit. The data is purged for a backup if the action is to abort.

### Return code

Output. The return code from the API call.

### Usage notes:

This function is called for each session that has been opened. There are two possible action codes:

- Commit

Output of data to this session, or the reading of data from the session, is complete.

For a write (backup) session, if the vendor returns to DB2 with a return code of SQLUV\_OK, DB2 assumes that the output data has been appropriately saved by the vendor product, and can be accessed if referenced in a later **sqluvint** call.

For a read (restore) session, if the vendor returns to DB2 with a return code of SQLUV\_OK, the data should not be deleted, because it may be needed again.

If the vendor returns SQLUV\_COMMIT\_FAILED, DB2 assumes that there are problems with the entire backup or restore operation. All active sessions are terminated by **sqluvend** calls with action = SQLUV\_ABORT. For a backup operation, committed sessions receive a **sqluvint**, **sqluvdel**, and **sqluvend** sequence of calls.

- Abort

A problem has been encountered by DB2, and there will be no more reading or writing of data to the session.

For a write (backup) session, the vendor should delete the partial output dataset, and use a SQLUV\_OK return code if the partial output is deleted. DB2 assumes that there are problems with the entire backup. All active sessions are terminated by **sqluvend** calls with action = SQLUV\_ABORT, and committed sessions receive a **sqluvint**, **sqluvdel**, and **sqluvend** sequence of calls.

For a read (restore) session, the vendor should not delete the data (because it may be needed again), but should clean up and return to DB2 with a SQLUV\_OK return code. DB2 terminates all the restore sessions by **sqluvend** calls with action = SQLUV\_ABORT. If the vendor returns SQLUV\_ABORT\_FAILED to DB2, the caller is not notified of this error, because DB2 returns the first fatal failure and ignores subsequent failures. In this case, for DB2 to have called **sqluvend** with action = SQLUV\_ABORT, an initial fatal error must have occurred.

### Return codes:

Table 12. Valid Return Codes for sqluvend and Resulting DB2 Action

Literal in Header File	Description	Probable Next Call	Other Comments
SQLUV_OK	Operation successful.	no further calls	Free all memory allocated for this session and terminate.
SQLUV_COMMIT_FAILED	Commit request failed.	no further calls	Free all memory allocated for this session and terminate.
SQLUV_ABORT_FAILED	Abort request failed.	no further calls	

## sqluvdel - Delete Committed Session

Deletes committed sessions.

### Authorization:

One of the following:

- *sysadm*
- *dbadm*

### Required connection:

Database

### API include file:

*sqluvend.h*

### C API syntax:

```
/* File: sqluvend.h */
/* API: Delete Committed Session */
/* ... */
int sqluvdel (
    struct Init_input *,
    struct Init_output *,
    struct Return_code *);
/* ... */
```

### API parameters:

#### Init\_input

Input. Space allocated for Init\_input and Return\_code.

#### Return\_code

Output. Return code from the API call. The object pointed to by the Init\_input structure is deleted.

### Usage notes:

If multiple sessions are opened, and some sessions are committed, but one of them fails, this function is called to delete the committed sessions. No sequence number is specified; **sqluvdel** is responsible for finding all of the objects that were created during a particular backup operation, and deleting them. Information in the INIT-INPUT structure is used to identify the output data to be deleted. The call to **sqluvdel** is responsible for establishing any connection or session that is required to delete a backup object from the vendor device. If the return code from this call is SQLUV\_DELETE\_FAILED, DB2 does not notify the caller, because DB2 returns the first fatal failure and ignores subsequent failures. In this case, for DB2 to have called **sqluvdel**, an initial fatal error must have occurred.

### Return codes:

Table 13. Valid Return Codes for sqluvdel and Resulting DB2 Action

Literal in Header File	Description	Probable Next Call	Other Comments
SQLUV_OK	Operation successful.	no further calls	
SQLUV_DELETE_FAILED	Delete request failed.	no further calls	

### db2VendorQueryApiVersion - Query Device Supported API Level

This function is called to determine which level of the vendor API is supported by the vendor library. If the vendor library is not compatible with DB2, then that vendor library will not be used.

If a vendor library does not have this API implemented for logs, the vendor library cannot be used and DB2 will report an error. This will not affect images that currently work with existing vendor libraries.

**Authorization:**

One of the following:

- *sysadm*
- *dbadm*

**Required connection:**

Database.

**API include file:**

*db2VendorApi.h*

**C API syntax:**

```
void db2VendorQueryApiVersion(db2UInt32 *supportedVersion);
```

**API parameters:**

**supportedVersion**

Output. Returns the version of the vendor API the vendor library supports.

**Usage notes:**

This function will be called before any other vendor APIs are invoked.

---

### db2VendorGetNextObj - Get Next Object on Device

This function is called after a query has been set up (using *sqluvint*) to get the next object that matches the search criteria. Only one search for either images or log files can be set up at one time.

**Authorization:**

One of the following:

- *sysadm*
- *dbadm*

**Required connection:**

Database.

**API include file:**

## db2VendorGetNextObj - Get Next Object on Device

*db2VendorApi.h*

### C API syntax:

```
int db2VendorGetNextObj(void *vendorCB,  
    struct db2VendorQueryInfo *queryInfo,  
    struct Return_code *returnCode);  
  
typedef struct db2VendorQueryInfo  
{  
    char          db2Instance[SQL_INSTNAME_SZ + 1];  
    char          dbname[SQL_DBNAME_SZ + 1];  
    char          dbalias[SQL_ALIAS_SZ + 1];  
    char          timestamp[SQLU_TIME_STAMP_LEN + 1];  
    char          filename[DB2VENDOR_MAX_FILENAME_SZ + 1];  
    char          owner[DB2VENDOR_MAX_OWNER_SZ + 1];  
    char          mgmtClass[DB2VENDOR_MAX_MGMTCLASS_SZ + 1];  
    char          oldestLogfile[DB2_LOGFILE_NAME_LEN + 1];  
    db2UInt16     sequenceNum  
    SQL_PDB_NODE_TYPE dbPartitionNum;  
    db2UInt32     type;  
    db2UInt64     sizeEstimate;  
} db2VendorQueryInfo;
```

### API parameters:

#### vendorCB

Input. Pointer to space allocated by the vendor library.

#### queryInfo

Output. Pointer to a *db2VendorQueryInfo* structure to be filled in by the vendor library.

#### returnCode

Output. The return code from the API call.

#### db2Instance

Specifies the name of the instance that the object belongs to.

#### dbname

Specifies the name of the database that the object belongs to.

#### dbalias

Specifies the alias of the database that the object belongs to.

#### timestamp

Specifies the time stamp used to identify the backup image. Valid only if the object is a backup image.

#### filename

Specifies the name of the object if the object is a load copy image or an archived log file.

**owner** Specifies the owner of the object.

#### mgmtClass

Specifies the management class the object was stored under (used by TSM).

#### oldestLogfile

Specifies the oldest log file stored with a backup image.

#### sequenceNum

Specifies the file extension for the backup image. Valid only if the object is a backup.



## db2VendorGetNextObj - Get Next Object on Device

### dbPartitionNum

Specifies the number of the database partition that the object belongs to.

**type** Specifies the image type if the object is a backup image.

### sizeEstimate

Specifies the estimated size of the object.

### Usage notes:

Not all fields will pertain to each object or each vendor. The mandatory fields that need to be filled out are db2Instance, dbname, dbalias, timestamp (for images), filename (for logs and load copy images), owner, sequenceNum (for images) and dbPartitionNum. The remaining fields will be left for the specific vendors to define. If a field does not pertain, then it should be initialized to "" for strings and 0 for numeric types.

---

## DB2-INFO

This structure contains information identifying DB2 to the vendor device.

Table 14. Fields in the DB2-INFO Structure. All fields are NULL-terminated strings.

Field Name	Data Type	Description
DB2_id	char	An identifier for the DB2 product. Maximum length of the string it points to is 8 characters.
version	char	The current version of the DB2 product. Maximum length of the string it points to is 8 characters.
release	char	The current release of the DB2 product. Set to NULL if it is insignificant. Maximum length of the string it points to is 8 characters.
level	char	The current level of the DB2 product. Set to NULL if it is insignificant. Maximum length of the string it points to is 8 characters.
action	char	Specifies the action to be taken. Maximum length of the string it points to is 1 character.
filename	char	The file name used to identify the backup image. If it is NULL, the <i>server_id</i> , <i>db2instance</i> , <i>dbname</i> , and <i>timestamp</i> will uniquely identify the backup image. Maximum length of the string it points to is 255 characters.
server_id	char	A unique name identifying the server where the database resides. Maximum length of the string it points to is 8 characters.
db2instance	char	The db2instance ID. This is the user ID invoking the command. Maximum length of the string it points to is 8 characters.
type	char	Specifies the type of backup being taken or the type of restore being performed. The following are possible values:  When action is SQLUV_WRITE: 0 - full database backup 3 - table space level backup  When action is SQLUV_READ: 0 - full restore 3 - online table space restore 4 - table space restore 5 - history file restore

Table 14. Fields in the DB2-INFO Structure (continued). All fields are NULL-terminated strings.

Field Name	Data Type	Description
dbname	char	The name of the database to be backed up or restored. Maximum length of the string it points to is 8 characters.
alias	char	The alias of the database to be backed up or restored. Maximum length of the string it points to is 8 characters.
timestamp	char	The time stamp used to identify the backup image. Maximum length of the string it points to is 26 characters.
sequence	char	Specifies the file extension for the backup image. For write operations, the value for the first session is 1 and each time another session is initiated with an sqluvint call, the value is incremented by 1. For read operations, the value is always zero. Maximum length of the string it points to is 3 characters.
obj_list	struct sqlu_gen_list	Reserved for future use.
max_bytes_per_txn	sqlint32	Specifies to the vendor in bytes, the transfer buffer size specified by the user.
image_filename	char	Reserved for future use.
reserve	void	Reserved for future use.
nodename	char	Name of the node at which the backup was generated.
password	char	Password for the node at which the backup was generated.
owner	char	ID of the backup originator.
mcNameP	char	Management class.
nodeNum	SQL_PDB_NODE_TYPE	Node number. Numbers greater than 255 are supported by the vendor interface.

The *filename*, or *server\_id*, *db2instance*, *type*, *dbname* and *timestamp* uniquely identifies the backup image. The sequence number, specified by *sequence*, identifies the file extension. When a backup image is to be restored, the same values must be specified to retrieve the backup image. Depending on the vendor product, if *filename* is used, the other parameters may be set to NULL, and vice versa.

#### Language syntax:

##### C Structure

```

/* File: sqluvend.h */
/* ... */
typedef struct DB2_info
{
    char          *DB2_id;
    char          *version;
    char          *release;
    char          *level;
    char          *action;
    char          *filename;
    char          *server_id;
    char          *db2instance;
    char          *type;
    char          *dbname;
    char          *alias;
    char          *timestamp;
}

```

```

char          *sequence;
struct sqlu_gen_list *obj_list;
long         max_bytes_per_txn;
char        *image_filename;
void        *reserve;
char        *nodename;
char        *password;
char        *owner;
char        *mcNameP;
char        *nodeNum;
SQL_PDB_NODE_TYPE
} DB2_info;
/* ... */

```

---

## VENDOR-INFO

This structure contains information identifying the vendor and version of the device.

*Table 15. Fields in the VENDOR-INFO Structure. All fields are NULL-terminated strings.*

Field Name	Data Type	Description
vendor_id	char	An identifier for the vendor. Maximum length of the string it points to is 64 characters.
version	char	The current version of the vendor product. Maximum length of the string it points to is 8 characters.
release	char	The current release of the vendor product. Set to NULL if it is insignificant. Maximum length of the string it points to is 8 characters.
level	char	The current level of the vendor product. Set to NULL if it is insignificant. Maximum length of the string it points to is 8 characters.
server_id	char	A unique name identifying the server where the database resides. Maximum length of the string it points to is 8 characters.
max_bytes_per_txn	sqlint32	The maximum supported transfer buffer size. Specified by the vendor, in bytes. This is used only if the return code from the vendor initialize function is SQLUV_BUFF_SIZE, indicating that an invalid buffer size was specified.
num_objects_in_backup	sqlint32	The number of sessions that were used to make a complete backup. This is used to determine when all backup images have been processed during a restore operation.
reserve	void	Reserved for future use.

### Language syntax:

#### C Structure

```

typedef struct Vendor_info
{
    char        *vendor_id;
    char        *version;
    char        *release;
    char        *level;
    char        *server_id;
    sqlint32    max_bytes_per_txn;
    sqlint32    num_objects_in_backup;
    void        *reserve;
} Vendor_info;

```

---

**INIT-INPUT**

This structure contains information provided by DB2 to set up and to establish a logical link with the vendor device.

*Table 16. Fields in the INIT-INPUT Structure. All fields are NULL-terminated strings.*

Field Name	Data Type	Description
DB2_session	struct DB2_info	A description of the session from the perspective of DB2.
size_options	unsigned short	The length of the options field. When using the DB2 backup or restore function, the data in this field is passed directly from the <i>VendorOptionsSize</i> parameter.
size_HI_order	sqluint32	High order 32 bits of DB size estimate in bytes; total size is 64 bits.
size_LOW_order	sqluint32	Low order 32 bits of DB size estimate in bytes; total size is 64 bits.
options	void	This information is passed from the application when the backup or the restore function is invoked. This data structure must be flat; in other words, no level of indirection is supported. Byte-reversal is not done, and the code page for this data is not checked. When using the DB2 backup or restore function, the data in this field is passed directly from the <i>pVendorOptions</i> parameter.
reserve	void	Reserved for future use.
prompt_lvl	char	Prompting level requested by the user when a backup or a restore operation was invoked. Maximum length of the string it points to is 1 character.
num_sessions	unsigned short	Number of sessions requested by the user when a backup or a restore operation was invoked.

### Language syntax:

#### C Structure

```
typedef struct Init_input
{
    struct DB2_info *DB2_session;
    unsigned short size_options;
    sqluint32      size_HI_order;
    sqluint32      size_LOW_order;
    void           *options;
    void           *reserve;
    char           *prompt_lvl;
    unsigned short num_sessions;
} Init_input;
```

---

**INIT-OUTPUT**

This structure contains the output returned by the vendor device.

*Table 17. Fields in the INIT-OUTPUT Structure*

Field Name	Data Type	Description
vendor_session	struct Vendor_info	Contains information to identify the vendor to DB2.

Table 17. Fields in the INIT-OUTPUT Structure (continued)

Field Name	Data Type	Description
pVendorCB	void	Vendor control block.
reserve	void	Reserved for future use.

**Language syntax:****C Structure**

```
typedef struct Init_output
{
    struct Vendor_info *vendor_session;
    void *pVendorCB;
    void *reserve;
} Init_output;
```

---

**DATA**

This structure contains data transferred between DB2 and the vendor device.

Table 18. Fields in the DATA Structure

Field Name	Data Type	Description
obj_num	sqlint32	The sequence number assigned by DB2 during a backup operation.
buff_size	sqlint32	The size of the buffer.
actual_buf_size	sqlint32	The actual number of bytes sent or received. This must not exceed <i>buff_size</i> .
dataptr	void	Pointer to the data buffer. DB2 allocates space for the buffer.
reserve	void	Reserved for future use.

**Language syntax:****C Structure**

```
typedef struct Data
{
    sqlint32 obj_num;
    sqlint32 buff_size;
    sqlint32 actual_buff_size;
    void *dataptr;
    void *reserve;
} Data;
```

---

**RETURN-CODE**

This structure contains the return code and a short explanation of the error being returned to DB2.

Table 19. Fields in the RETURN-CODE Structure

Field Name	Data Type	Description
return_code <sup>a</sup>	sqlint32	Return code from the vendor function.
description	char	A short description of the return code.
reserve	void	Reserved for future use.

<sup>a</sup> This is a vendor-specific return code that is not the same as the value returned by various DB2 APIs. See the individual API descriptions for the return codes that are accepted from vendor products.

**Language syntax:****C Structure**

```
typedef struct Return_code
{
    sqlint32  return_code,
    char      description[30],
    void      *reserve,
} Return_code;
```

---

**APIs for compressed backups****Compression plug-in interface**

DB2 will provide the definition for the `COMPR_DB2INFO` structure; the vendor will provide definitions for each of the others of the following structures and APIs. The following structures, prototypes, and constants are defined in the file `sqlucompr.h`, which is shipped with DB2.

**Description of the DB2 environment - `COMPR_DB2INFO`:**

```
struct COMPR_DB2INFO {
    char      tag[16];
    db2Uint32 version;
    db2Uint32 size;
    char      dbalias[SQLU_ALIAS_SZ+1];
    char      instance[SQL_INSTNAME_SZ+1];
    SQL_PDB_NODE_TYPE node;
    SQL_PDB_NODE_TYPE catnode;
    char      timestamp[SQLU_TIME_STAMP_LEN+1];
    db2Uint32 bufferSize;
    db2Uint32 options;
    db2Uint32 bkOptions;

    db2Uint32 db2Version;
    db2Uint32 platform;
    db2int32  comprOptionsByteOrder;
    db2Uint32 comprOptionsSize;
    void      *comprOptions;
    db2Uint32 savedBlockSize;
    void      *savedBlock;
};
```

**`COMPR_DB2INFO`**

DB2 will allocate and define this structure and will pass it in as a parameter on the `InitCompression` and `InitDecompression` APIs. This structure describes the database being backed up or restored and gives details about DB2 environment where the operation is occurring. The fields in the structure are:

**tag[16]**

Used as an eye catcher for the structure. This is always set to the string `"COMPR_DB2INFO \0"`.

**version**

Indicates which version of the structure is being used so APIs can indicate the presence of additional fields. Currently, the version is 1. In the future there may be more fields added to this structure.

**size** Specifies the size of the `COMPR_DB2INFO` structure in bytes.

**dbalias**[SQLU\_ALIAS\_SZ+1]

**instance**[SQL\_INSTNAME\_SZ+1]

**node**

**catnode**

**timestamp**[SQLU\_TIME\_STAMP\_LEN+1]

Describes the database being backed up or restored. These are the fields that are used to name the backup image. For restore operations, dbalias refers to the alias of the source database.

**bufferSize**

Specifies the size of a transfer buffer (in 4 K pages).

**options**

The iOptions field specified on the db2Backup API or the db2Restore API.

**bkOptions**

For restore operations, specifies the iOptions field that was used on the db2Backup API when the backup was created. For backup operations, it is set to zero.

**db2Version**

Specifies the version of the DB2 engine.

**platform**

Specifies the platform on which the DB2 engine is running. The value will be one of the ones listed in <sqlmon.h>.

**comprOptionsByteOrder**

Specifies the byte-order used on the client where the API is being run. DB2 will do no interpretation or conversion of the data passed through as *comprOptions*, so this field should be used to determine whether the data needs to be byte reversed before being used. Any conversion must be done by the plug-in library itself.

**comprOptionsSize**

Specifies the value of the *piComprOptionsSize* parameter on the db2Backup and db2Restore APIs.

**\*comprOptions**

Specifies the value of the *piComprOptions* field on the db2Backup and db2Restore APIs.

**savedBlockSize**

**\*savedBlock**

DB2 allows the plug-in library to save an arbitrary block of data in the backup image. If such a block of data was saved with a particular backup, it will be returned on these fields on the restore operation. For backup operations, these fields are set to zero.

### Description of the plug-in - COMPR\_PIINFO:

```
struct COMPR_PIINFO {
    char    tag[16];
    db2UInt32  version;
    db2UInt32  size;
    db2UInt32  useCRC;
```

## Compression plug-in interface

```
db2UInt32 useGran;  
db2UInt32 useAllBlocks;  
db2UInt32 savedBlockSize;  
};
```

### COMPR\_PIINFO

This structure is used by the plug-in library to describe itself to DB2. This structure is allocated and initialized by DB2, and the key fields are filled in by the plug-in library on the InitCompression call.

#### tag[16]

Used as an eye catcher for the structure. (It is set by DB2.) This is always set to the string "COMPR\_PIINFO \0".

#### version

Indicates which version of the structure is being used so APIs can indicate the presence of additional fields. Currently, the version is 1. (It is set by DB2.) In the future there may be more fields added to this structure.

**size** Indicates the size of the COMPR\_PIINFO structure (in bytes). (It is set by DB2.)

#### useCRC

DB2 allows compression plug-ins to use a 32-bit CRC or checksum value to validate the integrity of the data being compressed and decompressed. If the library uses such a check, it will set this field to 1. Otherwise, it will set the field to 0.

#### useGran

If the compression routine is able to compress data in arbitrarily-sized increments, the library will set this field to 1. If the compression routine compresses data only in byte-sized increments, the library will set this field to 0. See the description of the useGran parameter of Compress for details of the implications of setting this indicator. For restore operations, this field is ignored.

#### useAllBlocks

Specifies whether DB2 should back up a compressed block of data that is larger than the original uncompressed block. By default, DB2 will store data uncompressed if the compressed version is larger, but under some circumstances the plug-in library will wish to have the compressed data backed up anyway. If DB2 is to save the compressed version of the data for all blocks, the library will set this value to 1. If DB2 is to save the compressed version of the data only when it is smaller than the original data, the library will set this value to 0. For restore operations, this field is ignored.

#### savedBlockSize

DB2 allows the plug-in library to save an arbitrary block of data in the backup image. If such a block of data is to be saved with a particular backup, the library will set this field to the size of the block to be allocated for this data. (The actual data will be passed to DB2 on a subsequent API call.) If no data is to be saved, the plug-in library will set this field to zero. For restore operations, this field is ignored.

### Description of the control block - COMPR\_CB:

```
struct COMPR_CB;  
  
extern "C" {
```



```

int InitCompression(
    const COMPR_DB2INFO *db2Info,
    COMPR_PIINFO *piInfo,
    COMPR_CB **pCB);

int GetSavedBlock(
    COMPR_CB *pCB,
    db2Uint32 blockSize,
    void *data);

int Compress(
    COMPR_CB *pCB,
    const char *src,
    db2int32 srcLen,
    db2Uint32 srcGran,
    char *tgt,
    db2int32 tgtSize,
    db2int32 *srcAct,
    db2int32 *tgtAct,
    db2Uint32 *tgtCRC);

int GetMaxCompressedSize(
    COMPR_CB *pCB,
    db2Uint32 srcLen);

int TermCompression(
    COMPR_CB *pCB);

int InitDecompression(
    const COMPR_DB2INFO *db2Info,
    COMPR_CB **pCB);

int Decompress(
    COMPR_CB *pCB,
    const char *src,
    db2int32 srcLen,
    char *tgt,
    db2int32 tgtSize,
    db2int32 *tgtAct,
    db2Uint32 *tgtCRC);

int TermDecompression(
    COMPR_CB *pCB);
}

```

### COMPR\_CB

This is a structure that will be used internally by the plug-in library. It contains data used internally by compression and decompression routines. DB2 passes the structure to each call it makes to the plug-in library, but all aspects of the structure are left up to the library, including the definition of the structure's fields and memory management of the structure.

```

int InitCompression(
    const COMPR_DB2INFO *db2Info,
    COMPR_PIINFO *piInfo,
    COMPR_CB **pCB);

```

Initializes the compression library. DB2 will pass the db2Info and piInfo structures. The library will fill in the appropriate fields of piInfo, and will allocate pCB and return a pointer to the allocated memory.

```

int GetSavedBlock(
    COMPR_CB *pCB,
    db2Uint32 blockSize,
    void *data);

```

## Compression plug-in interface

Gets the vendor-specific block of data to be saved with the backup image. If the library returned a non-zero value for `piInfo->savedBlockSize`, DB2 will call `GetSavedBlock` using that value as *blockSize*. The plug-in library writes data of the given size to the memory referenced by *data*. This function will be called during initial data processing in BM1 for backup only. Even if `parallelism > 1` is specified on the `db2Backup` API, this function will be called only once per backup.

```
int Compress(  
    COMPR_CB          *pCB,  
    const char        *src,  
    db2int32          srcLen,  
    db2int32          srcGran,  
    char              *tgt,  
    db2int32          tgtSize,  
    db2int32          *srcAct,  
    db2int32          *tgtAct,  
    db2uint32         *tgtCRC);
```

Compress a block of data. *src* points to a block of data that is *srcLen* bytes in size. *tgt* points to a buffer that is *tgtSize* bytes in size. The plug-in library compresses the data at address *src* and writes the compressed data to the buffer at address *tgt*. The actual amount of uncompressed data that was compressed is stored in *srcAct*. The actual size of the compressed data is returned as *tgtAct*.

If the library returned a value of 1 for `piInfo->useCRC`, the CRC value of the uncompressed block is returned as *tgtCRC*. If the library returned a value of 0 for `piInfo->useCRC`, *tgtCRC* will be a null pointer.

If the library returned a value of 1 for `piInfo->useGran`, *srcGran* specifies the log2 of the page size of the data. (For example, if the page size of the data is 4096 bytes, *srcGran* is 12.) The library ensures that the amount of data actually compressed (*srcAct*) is an exact multiple of this page size. If the library sets the `useGran` flag, DB2 is able to use a more efficient algorithm for fitting the compressed data into the backup image. This means that both the performance of the plug-in will be better and that the compressed backup image will be smaller. If the library returned a value of 0 for `piInfo->srcGran`, the granularity is 1 byte.

```
int GetMaxCompressedSize(  
    COMPR_CB          *pCB,  
    db2uint32         srcLen,  
    db2uint32         *tgtLen);
```

Estimates the size of the largest possible buffer required to compress a block of data. *srcLen* indicates the size of a block of data about to be compressed. The library returns the theoretical maximum size of the buffer after compression as *tgtLen*.

DB2 will use the value returned as *tgtLen* to optimize its use of memory internally. The penalty for not calculating a value or for calculating an incorrect value is that DB2 will have to call the `Compress` API more than once for a single block of data, or that it may waste memory from the utility heap. The backup will still be created correctly, regardless of the values returned.

```
int TermCompression(  
    COMPR_CB          *pCB);
```

Terminates the compression library. The library will free the memory used for `pCB`.

## Compression plug-in interface

```
int InitDecompression(
    const COMPR_DB2INFO *db2Info,
    COMPR_CB **pCB);
```

Initializes the decompression library. DB2 will pass the `db2Info` structure. The library will allocate `pCB` and return a pointer to the allocated memory.

```
int Decompress(
    COMPR_CB *pCB,
    const char *src,
    db2int32 srcLen,
    char *tgt,
    db2int32 tgtSize,
    db2int32 *tgtLen,
    db2uint32 *tgtCRC);
```

Decompresses a block of data. `src` points to a block of data that is `srcLen` bytes in size. `tgt` points to a buffer that is `tgtSize` bytes in size. The plug-in library decompresses the data at address `src` and writes the uncompressed data to the buffer at address `tgt`. The actual size of the uncompressed data is returned as `tgtLen`. If the library returned a value of 1 for `piInfo->useCRC`, the CRC of the uncompressed block is returned as `tgtCRC`. If the library returned a value of 0 for `piInfo->useCRC`, `tgtLen` will be a null pointer.

```
int TermDecompression(
    COMPR_CB *pCB);
```

Terminates the decompression library. The library will free the memory used for `pCB`. All the memory used internally by these APIs will be managed by the vendor. The plug-in library will manage memory used by the `COMPR_CB` structure. DB2 will manage the memory used for the data buffers (the `src` and `tgt` parameters on the APIs).

### Plug-in interface return codes:

These are the return codes that the APIs may return. Except where specified, DB2 will terminate the backup or restore when any non-zero return code is returned.

SQLUV_OK	0	Operation succeeded
SQLUV_BUFFER_TOO_SMALL	100	Target buffer is too small. When indicated on backup, the <code>tgtAct</code> field shall indicate the estimated size required to compress the object. DB2 will retry the operation with a buffer at least as large as specified. When indicated on restore, the operation will fail.
SQLUV_PARTIAL_BUFFER	101	A buffer was partially compressed. When indicated on backup, the <code>srcAct</code> field shall indicate the actual amount of data actually compressed and the <code>tgtAct</code> field shall indicate the actual size of the compressed data. When indicated on restore, the operation will fail.
SQLUV_NO_MEMORY	102	Out of memory
SQLUV_EXCEPTION	103	A signal or exception was raised in the code.
SQLUV_INTERNAL_ERROR	104	An internal error was detected.

## Compression plug-in interface

|                   The difference between SQLUV\_BUFFER\_TOO\_SMALL and  
|                   SQLUV\_PARTIAL\_BUFFER is that when SQLUV\_PARTIAL\_BUFFER is returned,  
|                   DB2 will consider the data in the output buffer to be valid.

|                   **Related reference:**

- |                   • “db2Backup - Backup database” on page 71
- |                   • “db2Restore - Restore database” on page 96

---

## Appendix I. DB2 Universal Database technical information

---

### DB2 documentation and help

DB2<sup>®</sup> technical information is available through the following tools and methods:

- DB2 Information Center
  - Topics
  - Help for DB2 tools
  - Sample programs
  - Tutorials
- Downloadable PDF files, PDF files on CD, and printed books
  - Guides
  - Reference manuals
- Command line help
  - Command help
  - Message help
  - SQL state help
- Installed source code
  - Sample programs

You can access additional DB2 Universal Database<sup>™</sup> technical information such as technotes, white papers, and Redbooks<sup>™</sup> online at [ibm.com](http://ibm.com)<sup>®</sup>. Access the DB2 Information Management software library site at [www.ibm.com/software/data/pubs/](http://www.ibm.com/software/data/pubs/).

### DB2 documentation updates

IBM<sup>®</sup> may periodically make documentation FixPaks and other documentation updates to the DB2 Information Center available. If you access the DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/>, you will always be viewing the most up-to-date information. If you have installed the DB2 Information Center locally, then you need to install any updates manually before you can view them. Documentation updates allow you to update the information that you installed from the *DB2 Information Center CD* when new information becomes available.

The Information Center is updated more frequently than either the PDF or the hardcopy books. To get the most current DB2 technical information, install the documentation updates as they become available or go to the DB2 Information Center at the [www.ibm.com](http://www.ibm.com) site.

#### Related concepts:

- “CLI sample programs” in the *CLI Guide and Reference, Volume 1*
- “Java sample programs” in the *Application Development Guide: Building and Running Applications*
- “DB2 Information Center” on page 364

#### Related tasks:

- “Invoking contextual help from a DB2 tool” on page 381

- “Updating the DB2 Information Center installed on your computer or intranet server” on page 373
- “Invoking message help from the command line processor” on page 382
- “Invoking command help from the command line processor” on page 382
- “Invoking SQL state help from the command line processor” on page 383

**Related reference:**

- “DB2 PDF and printed documentation” on page 375

## DB2 Information Center

The DB2<sup>®</sup> Information Center gives you access to all of the information you need to take full advantage of DB2 family products, including DB2 Universal Database<sup>™</sup>, DB2 Connect<sup>™</sup>, DB2 Information Integrator and DB2 Query Patroller<sup>™</sup>. The DB2 Information Center also contains information for major DB2 features and components including replication, data warehousing, and the DB2 extenders.

The DB2 Information Center has the following features if you view it in Mozilla 1.0 or later or Microsoft<sup>®</sup> Internet Explorer 5.5 or later. Some features require you to enable support for JavaScript<sup>™</sup>:

### Flexible installation options

You can choose to view the DB2 documentation using the option that best meets your needs:

- To effortlessly ensure that your documentation is always up to date, you can access all of your documentation directly from the DB2 Information Center hosted on the IBM<sup>®</sup> Web site at <http://publib.boulder.ibm.com/infocenter/db2help/>
- To minimize your update efforts and keep your network traffic within your intranet, you can install the DB2 documentation on a single server on your intranet
- To maximize your flexibility and reduce your dependence on network connections, you can install the DB2 documentation on your own computer

### Search

You can search all of the topics in the DB2 Information Center by entering a search term in the **Search** text field. You can retrieve exact matches by enclosing terms in quotation marks, and you can refine your search with wildcard operators (\*, ?) and Boolean operators (AND, NOT, OR).

### Task-oriented table of contents

You can locate topics in the DB2 documentation from a single table of contents. The table of contents is organized primarily by the kind of tasks you may want to perform, but also includes entries for product overviews, goals, reference information, an index, and a glossary.

- Product overviews describe the relationship between the available products in the DB2 family, the features offered by each of those products, and up to date release information for each of these products.
- Goal categories such as installing, administering, and developing include topics that enable you to quickly complete tasks and develop a deeper understanding of the background information for completing those tasks.

- Reference topics provide detailed information about a subject, including statement and command syntax, message help, and configuration parameters.

#### **Show current topic in table of contents**

You can show where the current topic fits into the table of contents by clicking the **Refresh / Show Current Topic** button in the table of contents frame or by clicking the **Show in Table of Contents** button in the content frame. This feature is helpful if you have followed several links to related topics in several files or arrived at a topic from search results.

**Index** You can access all of the documentation from the index. The index is organized in alphabetical order by index term.

#### **Glossary**

You can use the glossary to look up definitions of terms used in the DB2 documentation. The glossary is organized in alphabetical order by glossary term.

#### **Integrated localized information**

The DB2 Information Center displays information in the preferred language set in your browser preferences. If a topic is not available in your preferred language, the DB2 Information Center displays the English version of that topic.

For iSeries™ technical information, refer to the IBM eServer™ iSeries information center at [www.ibm.com/eserver/series/infocenter/](http://www.ibm.com/eserver/series/infocenter/).

#### **Related concepts:**

- “DB2 Information Center installation scenarios” on page 365

#### **Related tasks:**

- “Updating the DB2 Information Center installed on your computer or intranet server” on page 373
- “Displaying topics in your preferred language in the DB2 Information Center” on page 374
- “Invoking the DB2 Information Center” on page 372
- “Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)” on page 368
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” on page 370

---

## **DB2 Information Center installation scenarios**

Different working environments can pose different requirements for how to access DB2® information. The DB2 Information Center can be accessed on the IBM® Web site, on a server on your organization’s network, or on a version installed on your computer. In all three cases, the documentation is contained in the DB2 Information Center, which is an architected web of topic-based information that you view with a browser. By default, DB2 products access the DB2 Information Center on the IBM Web site. However, if you want to access the DB2 Information Center on an intranet server or on your own computer, you must install the DB2 Information Center using the DB2 Information Center CD found in your product Media Pack. Refer to the summary of options for accessing DB2 documentation which follows, along with the three installation scenarios, to help determine which

method of accessing the DB2 Information Center works best for you and your work environment, and what installation issues you might need to consider.

### Summary of options for accessing DB2 documentation:

The following table provides recommendations on which options are possible in your work environment for accessing the DB2 product documentation in the DB2 Information Center.

Internet access	Intranet access	Recommendation
Yes	Yes	Access the DB2 Information Center on the IBM Web site, or access the DB2 Information Center installed on an intranet server.
Yes	No	Access the DB2 Information Center on the IBM Web site.
No	Yes	Access the DB2 Information Center installed on an intranet server.
No	No	Access the DB2 Information Center on a local computer.

### Scenario: Accessing the DB2 Information Center on your computer:

Tsu-Chen owns a factory in a small town that does not have a local ISP to provide him with Internet access. He purchased DB2 Universal Database™ to manage his inventory, his product orders, his banking account information, and his business expenses. Never having used a DB2 product before, Tsu-Chen needs to learn how to do so from the DB2 product documentation.

After installing DB2 Universal Database on his computer using the typical installation option, Tsu-Chen tries to access the DB2 documentation. However, his browser gives him an error message that the page he tried to open cannot be found. Tsu-Chen checks the installation manual for his DB2 product and discovers that he has to install the DB2 Information Center if he wants to access DB2 documentation on his computer. He finds the *DB2 Information Center CD* in the media pack and installs it.

From the application launcher for his operating system, Tsu-Chen now has access to the DB2 Information Center and can learn how to use his DB2 product to increase the success of his business.

### Scenario: Accessing the DB2 Information Center on the IBM Web site:

Colin is an information technology consultant with a training firm. He specializes in database technology and SQL and gives seminars on these subjects to businesses all over North America using DB2 Universal Database. Part of Colin's seminars includes using DB2 documentation as a teaching tool. For example, while teaching courses on SQL, Colin uses the DB2 documentation on SQL as a way to teach basic and advanced syntax for database queries.

Most of the businesses at which Colin teaches have Internet access. This situation influenced Colin's decision to configure his mobile computer to access the DB2 Information Center on the IBM Web site when he installed the latest version of DB2 Universal Database. This configuration allows Colin to have online access to the latest DB2 documentation during his seminars.



However, sometimes while travelling Colin does not have Internet access. This posed a problem for him, especially when he needed to access to DB2 documentation to prepare for seminars. To avoid situations like this, Colin installed a copy of the DB2 Information Center on his mobile computer.

Colin enjoys the flexibility of always having a copy of DB2 documentation at his disposal. Using the **db2set** command, he can easily configure the registry variables on his mobile computer to access the DB2 Information Center on either the IBM Web site, or his mobile computer, depending on his situation.

**Scenario: Accessing the DB2 Information Center on an intranet server:**

Eva works as a senior database administrator for a life insurance company. Her administration responsibilities include installing and configuring the latest version of DB2 Universal Database on the company's UNIX<sup>®</sup> database servers. Her company recently informed its employees that, for security reasons, it would not provide them with Internet access at work. Because her company has a networked environment, Eva decides to install a copy of the DB2 Information Center on an intranet server so that all employees in the company who use the company's data warehouse on a regular basis (sales representatives, sales managers, and business analysts) have access to DB2 documentation.

Eva instructs her database team to install the latest version of DB2 Universal Database on all of the employee's computers using a response file, to ensure that each computer is configured to access the DB2 Information Center using the host name and the port number of the intranet server.

However, through a misunderstanding Migual, a junior database administrator on Eva's team, installs a copy of the DB2 Information Center on several of the employee computers, rather than configuring DB2 Universal Database to access the DB2 Information Center on the intranet server. To correct this situation Eva tells Migual to use the **db2set** command to change the DB2 Information Center registry variables (DB2\_DOCHOST for the host name, and DB2\_DOCPORT for the port number) on each of these computers. Now all of the appropriate computers on the network have access to the DB2 Information Center, and employees can find answers to their DB2 questions in the DB2 documentation.

**Related concepts:**

- "DB2 Information Center" on page 364

**Related tasks:**

- "Updating the DB2 Information Center installed on your computer or intranet server" on page 373
- "Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)" on page 368
- "Installing the DB2 Information Center using the DB2 Setup wizard (Windows)" on page 370

**Related reference:**

- "db2set - DB2 Profile Registry Command" in the *Command Reference*

---

## Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)

DB2 product documentation can be accessed in three ways: on the IBM Web site, on an intranet server, or on a version installed on your computer. By default, DB2 products access DB2 documentation on the IBM Web site. If you want to access the DB2 documentation on an intranet server or on your own computer, you must install the documentation from the *DB2 Information Center CD*. Using the DB2 Setup wizard, you can define your installation preferences and install the DB2 Information Center on a computer that uses a UNIX operating system.

### Prerequisites:

This section lists the hardware, operating system, software, and communication requirements for installing the DB2 Information Center on UNIX computers.

- **Hardware requirements**

You require one of the following processors:

- PowerPC (AIX)
- HP 9000 (HP-UX)
- Intel 32-bit (Linux)
- Solaris UltraSPARC computers (Solaris Operating Environment)

- **Operating system requirements**

You require one of the following operating systems:

- IBM AIX 5.1 (on PowerPC)
- HP-UX 11i (on HP 9000)
- Red Hat Linux 8.0 (on Intel 32-bit)
- SuSE Linux 8.1 (on Intel 32-bit)
- Sun Solaris Version 8 (on Solaris Operating Environment UltraSPARC computers)

**Note:** The DB2 Information Center runs on a subset of the UNIX operating systems on which DB2 clients are supported. It is therefore recommended that you either access the DB2 Information Center from the IBM Web site, or that you install and access the DB2 Information Center on an intranet server.

- **Software requirements**

- The following browser is supported:
  - Mozilla Version 1.0 or greater

- The DB2 Setup wizard is a graphical installer. You must have an implementation of the X Window System software capable of rendering a graphical user interface for the DB2 Setup wizard to run on your computer. Before you can run the DB2 Setup wizard you must ensure that you have properly exported your display. For example, enter the following command at the command prompt:

```
export DISPLAY=9.26.163.144:0.
```

- **Communication requirements**

- TCP/IP

### Procedure:

To install the DB2 Information Center using the DB2 Setup wizard:

1. Log on to the system.
2. Insert and mount the DB2 Information Center product CD on your system.
3. Change to the directory where the CD is mounted by entering the following command:

```
cd /cd
```

where */cd* represents the mount point of the CD.

4. Enter the **./db2setup** command to start the DB2 Setup wizard.
5. The IBM DB2 Setup Launchpad opens. To proceed directly to the installation of the DB2 Information Center, click **Install Product**. Online help is available to guide you through the remaining steps. To invoke the online help, click **Help**. You can click **Cancel** at any time to end the installation.
6. On the **Select the product you would like to install** page, click **Next**.
7. Click **Next** on the **Welcome to the DB2 Setup wizard** page. The DB2 Setup wizard will guide you through the program setup process.
8. To proceed with the installation, you must accept the license agreement. On the **License Agreement** page, select **I accept the terms in the license agreement** and click **Next**.
9. Select **Install DB2 Information Center on this computer** on the **Select the installation action** page. If you want to use a response file to install the DB2 Information Center on this or other computers at a later time, select **Save your settings in a response file**. Click **Next**.
10. Select the languages in which the DB2 Information Center will be installed on **Select the languages to install** page. Click **Next**.
11. Configure the DB2 Information Center for incoming communication on the **Specify the DB2 Information Center port** page. Click **Next** to continue the installation.
12. Review the installation choices you have made in the **Start copying files** page. To change any settings, click **Back**. Click **Install** to copy the DB2 Information Center files onto your computer.

You can also install the DB2 Information Center using a response file.

The installation logs `db2setup.his`, `db2setup.log`, and `db2setup.err` are located, by default, in the `/tmp` directory.

The `db2setup.log` file captures all DB2 product installation information, including errors. The `db2setup.his` file records all DB2 product installations on your computer. DB2 appends the `db2setup.log` file to the `db2setup.his` file. The `db2setup.err` file captures any error output that is returned by Java, for example, exceptions and trap information.

When the installation is complete, the DB2 Information Center will be installed in one of the following directories, depending upon your UNIX operating system:

- AIX: `/usr/opt/db2_08_01`
- HP-UX: `/opt/IBM/db2/V8.1`
- Linux: `/opt/IBM/db2/V8.1`
- Solaris Operating Environment: `/opt/IBM/db2/V8.1`

**Related concepts:**

- “DB2 Information Center” on page 364
- “DB2 Information Center installation scenarios” on page 365

#### Related tasks:

- “Installing DB2 using a response file (UNIX)” in the *Installation and Configuration Supplement*
- “Updating the DB2 Information Center installed on your computer or intranet server” on page 373
- “Displaying topics in your preferred language in the DB2 Information Center” on page 374
- “Invoking the DB2 Information Center” on page 372
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” on page 370

---

## Installing the DB2 Information Center using the DB2 Setup wizard (Windows)

DB2 product documentation can be accessed in three ways: on the IBM Web site, on an intranet server, or on a version installed on your computer. By default, DB2 products access DB2 documentation on the IBM Web site. If you want to access the DB2 documentation on an intranet server or on your own computer, you must install the DB2 documentation from the *DB2 Information Center CD*. Using the DB2 Setup wizard, you can define your installation preferences and install the DB2 Information Center on a computer that uses a Windows operating system.

#### Prerequisites:

This section lists the hardware, operating system, software, and communication requirements for installing the DB2 Information Center on Windows.

- **Hardware requirements**

You require one of the following processors:

- 32-bit computers: a Pentium or Pentium compatible CPU

- **Operating system requirements**

You require one of the following operating systems:

- Windows 2000
- Windows XP

**Note:** The DB2 Information Center runs on a subset of the Windows operating systems on which DB2 clients are supported. It is therefore recommended that you either access the DB2 Information Center on the IBM Web site, or that you install and access the DB2 Information Center on an intranet server.

- **Software requirements**

– The following browsers are supported:

- Mozilla 1.0 or greater
- Internet Explorer Version 5.5 or 6.0 (Version 6.0 for Windows XP)

- **Communication requirements**

- TCP/IP

#### Restrictions:

- You require an account with administrative privileges to install the DB2 Information Center.

## Procedure:

To install the DB2 Information Center using the DB2 Setup wizard:

1. Log on to the system with the account that you have defined for the DB2 Information Center installation.
2. Insert the CD into the drive. If enabled, the auto-run feature starts the IBM DB2 Setup Launchpad.
3. The DB2 Setup wizard determines the system language and launches the setup program for that language. If you want to run the setup program in a language other than English, or the setup program fails to auto-start, you can start the DB2 Setup wizard manually.

To start the DB2 Setup wizard manually:

- a. Click **Start** and select **Run**.
- b. In the **Open** field, type the following command:

```
x:\setup.exe /i 2-letter language identifier
```

where *x*: represents your CD drive, and *2-letter language identifier* represents the language in which the setup program will be run.

- c. Click **OK**.
4. The IBM DB2 Setup Launchpad opens. To proceed directly to the installation of the DB2 Information Center, click **Install Product**. Online help is available to guide you through the remaining steps. To invoke the online help, click **Help**. You can click **Cancel** at any time to end the installation.
  5. On the **Select the product you would like to install** page, click **Next**.
  6. Click **Next** on the **Welcome to the DB2 Setup wizard** page. The DB2 Setup wizard will guide you through the program setup process.
  7. To proceed with the installation, you must accept the license agreement. On the **License Agreement** page, select **I accept the terms in the license agreement** and click **Next**.
  8. Select **Install DB2 Information Center on this computer** on the **Select the installation action** page. If you want to use a response file to install the DB2 Information Center on this or other computers at a later time, select **Save your settings in a response file**. Click **Next**.
  9. Select the languages in which the DB2 Information Center will be installed on **Select the languages to install** page. Click **Next**.
  10. Configure the DB2 Information Center for incoming communication on the **Specify the DB2 Information Center port** page. Click **Next** to continue the installation.
  11. Review the installation choices you have made in the **Start copying files** page. To change any settings, click **Back**. Click **Install** to copy the DB2 Information Center files onto your computer.

You can install the DB2 Information Center using a response file. You can also use the **db2rspgn** command to generate a response file based on an existing installation.

For information on errors encountered during installation, see the `db2.log` and `db2wi.log` files located in the 'My Documents'\DB2LOG\ directory. The location of the 'My Documents' directory will depend on the settings on your computer.

The `db2wi.log` file captures the most recent DB2 installation information. The `db2.log` captures the history of DB2 product installations.

| **Related concepts:**

- “DB2 Information Center” on page 364
- “DB2 Information Center installation scenarios” on page 365

| **Related tasks:**

- “Installing a DB2 product using a response file (Windows)” in the *Installation and Configuration Supplement*
- “Updating the DB2 Information Center installed on your computer or intranet server” on page 373
- “Displaying topics in your preferred language in the DB2 Information Center” on page 374
- “Invoking the DB2 Information Center” on page 372
- “Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)” on page 368

| **Related reference:**

- “db2rspgn - Response File Generator Command (Windows)” in the *Command Reference*

---

## Invoking the DB2 Information Center

| The DB2 Information Center gives you access to all of the information that you  
| need to use DB2 products for Linux, UNIX, and Windows operating systems such  
| as DB2 Universal Database, DB2 Connect, DB2 Information Integrator, and DB2  
| Query Patroller.

You can invoke the DB2 Information Center from one of the following places:

- Computers on which a DB2 UDB client or server is installed
- An intranet server or local computer on which the DB2 Information Center installed
- The IBM Web site

| **Prerequisites:**

Before you invoke the DB2 Information Center:

- *Optional:* Configure your browser to display topics in your preferred language
- *Optional:* Configure your DB2 client to use the DB2 Information Center installed on your computer or intranet server

| **Procedure:**

To invoke the DB2 Information Center on a computer on which a DB2 UDB client or server is installed:

- From the Start Menu (Windows operating system): Click **Start** → **Programs** → **IBM DB2** → **Information** → **Information Center**.
- From the command line prompt:
  - For Linux and UNIX operating systems, issue the **db2icdocs** command.
  - For the Windows operating system, issue the **db2icdocs.exe** command.

To open the DB2 Information Center installed on an intranet server or local computer in a Web browser:

- Open the Web page at <http://<host-name>:<port-number>/>, where <host-name> represents the host name and <port-number> represents the port number on which the DB2 Information Center is available.

To open the DB2 Information Center on the IBM Web site in a Web browser:

- Open the Web page at [publib.boulder.ibm.com/infocenter/db2help/](http://publib.boulder.ibm.com/infocenter/db2help/).

**Related concepts:**

- “DB2 Information Center” on page 364

**Related tasks:**

- “Displaying topics in your preferred language in the DB2 Information Center” on page 374
- “Invoking contextual help from a DB2 tool” on page 381
- “Updating the DB2 Information Center installed on your computer or intranet server” on page 373
- “Invoking message help from the command line processor” on page 382
- “Invoking command help from the command line processor” on page 382
- “Invoking SQL state help from the command line processor” on page 383

---

## Updating the DB2 Information Center installed on your computer or intranet server

The DB2 Information Center available from <http://publib.boulder.ibm.com/infocenter/db2help/> will be periodically updated with new or changed documentation. IBM may also make DB2 Information Center updates available to download and install on your computer or intranet server. Updating the DB2 Information Center does not update DB2 client or server products.

**Prerequisites:**

You must have access to a computer that is connected to the Internet.

**Procedure:**

To update the DB2 Information Center installed on your computer or intranet server:

1. Open the DB2 Information Center hosted on the IBM Web site at: <http://publib.boulder.ibm.com/infocenter/db2help/>
2. In the Downloads section of the welcome page under the Service and Support heading, click the **DB2 Universal Database documentation** link.
3. Determine if the version of your DB2 Information Center is out of date by comparing the latest refreshed documentation image level to the documentation level you have installed. The documentation level you have installed is listed on the DB2 Information Center welcome page.
4. If a more recent version of the DB2 Information Center is available, download the latest refreshed *DB2 Information Center* image applicable to your operating system.
5. To install the refreshed *DB2 Information Center* image, follow the instructions provided on the Web page.

**Related concepts:**

- “DB2 Information Center installation scenarios” on page 365

**Related tasks:**

- “Invoking the DB2 Information Center” on page 372
- “Installing the DB2 Information Center using the DB2 Setup wizard (UNIX)” on page 368
- “Installing the DB2 Information Center using the DB2 Setup wizard (Windows)” on page 370

---

## Displaying topics in your preferred language in the DB2 Information Center

The DB2 Information Center attempts to display topics in the language specified in your browser preferences. If a topic has not been translated into your preferred language, the DB2 Information Center displays the topic in English.

**Procedure:**

To display topics in your preferred language in the Internet Explorer browser:

1. In Internet Explorer, click the **Tools** —> **Internet Options** —> **Languages...** button. The Language Preferences window opens.
2. Ensure your preferred language is specified as the first entry in the list of languages.
  - To add a new language to the list, click the **Add...** button.

**Note:** Adding a language does not guarantee that the computer has the fonts required to display the topics in the preferred language.

- To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
3. Refresh the page to display the DB2 Information Center in your preferred language.

To display topics in your preferred language in the Mozilla browser:

1. In Mozilla, select the **Edit** —> **Preferences** —> **Languages** button. The Languages panel is displayed in the Preferences window.
2. Ensure your preferred language is specified as the first entry in the list of languages.
  - To add a new language to the list, click the **Add...** button to select a language from the Add Languages window.
  - To move a language to the top of the list, select the language and click the **Move Up** button until the language is first in the list of languages.
3. Refresh the page to display the DB2 Information Center in your preferred language.

**Related concepts:**

- “DB2 Information Center” on page 364



---

## DB2 PDF and printed documentation

The following tables provide official book names, form numbers, and PDF file names. To order hardcopy books, you must know the official book name. To print a PDF file, you must know the PDF file name.

The DB2 documentation is categorized by the following headings:

- Core DB2 information
- Administration information
- Application development information
- Business intelligence information
- DB2 Connect information
- Getting started information
- Tutorial information
- Optional component information
- Release notes

The following tables describe, for each book in the DB2 library, the information needed to order the hard copy, or to print or view the PDF for that book. A full description of each of the books in the DB2 library is available from the IBM Publications Center at [www.ibm.com/shop/publications/order](http://www.ibm.com/shop/publications/order)

### Core DB2 information

The information in these books is fundamental to all DB2 users; you will find this information useful whether you are a programmer, a database administrator, or someone who works with DB2 Connect, DB2 Warehouse Manager, or other DB2 products.

*Table 20. Core DB2 information*

Name	Form Number	PDF File Name
<i>IBM DB2 Universal Database Command Reference</i>	SC09-4828	db2n0x81
<i>IBM DB2 Universal Database Glossary</i>	No form number	db2t0x81
<i>IBM DB2 Universal Database Message Reference, Volume 1</i>	GC09-4840, not available in hardcopy	db2m1x81
<i>IBM DB2 Universal Database Message Reference, Volume 2</i>	GC09-4841, not available in hardcopy	db2m2x81
<i>IBM DB2 Universal Database What's New</i>	SC09-4848	db2q0x81

### Administration information

The information in these books covers those topics required to effectively design, implement, and maintain DB2 databases, data warehouses, and federated systems.

*Table 21. Administration information*

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Administration Guide: Planning</i>	SC09-4822	db2d1x81

*Table 21. Administration information (continued)*

<b>Name</b>	<b>Form number</b>	<b>PDF file name</b>
<i>IBM DB2 Universal Database Administration Guide: Implementation</i>	SC09-4820	db2d2x81
<i>IBM DB2 Universal Database Administration Guide: Performance</i>	SC09-4821	db2d3x81
<i>IBM DB2 Universal Database Administrative API Reference</i>	SC09-4824	db2b0x81
<i>IBM DB2 Universal Database Data Movement Utilities Guide and Reference</i>	SC09-4830	db2dmx81
<i>IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference</i>	SC09-4831	db2hax81
<i>IBM DB2 Universal Database Data Warehouse Center Administration Guide</i>	SC27-1123	db2ddx81
<i>IBM DB2 Universal Database SQL Reference, Volume 1</i>	SC09-4844	db2s1x81
<i>IBM DB2 Universal Database SQL Reference, Volume 2</i>	SC09-4845	db2s2x81
<i>IBM DB2 Universal Database System Monitor Guide and Reference</i>	SC09-4847	db2f0x81

## Application development information

The information in these books is of special interest to application developers or programmers working with DB2 Universal Database (DB2 UDB). You will find information about supported languages and compilers, as well as the documentation required to access DB2 UDB using the various supported programming interfaces, such as embedded SQL, ODBC, JDBC, SQLJ, and CLI. If you are using the DB2 Information Center, you can also access HTML versions of the source code for the sample programs.

*Table 22. Application development information*

<b>Name</b>	<b>Form number</b>	<b>PDF file name</b>
<i>IBM DB2 Universal Database Application Development Guide: Building and Running Applications</i>	SC09-4825	db2axx81
<i>IBM DB2 Universal Database Application Development Guide: Programming Client Applications</i>	SC09-4826	db2a1x81
<i>IBM DB2 Universal Database Application Development Guide: Programming Server Applications</i>	SC09-4827	db2a2x81
<i>IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 1</i>	SC09-4849	db2l1x81

Table 22. Application development information (continued)

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Call Level Interface Guide and Reference, Volume 2</i>	SC09-4850	db2l2x81
<i>IBM DB2 Universal Database Data Warehouse Center Application Integration Guide</i>	SC27-1124	db2adx81
<i>IBM DB2 XML Extender Administration and Programming</i>	SC27-1234	db2sxx81

## Business intelligence information

The information in these books describes how to use components that enhance the data warehousing and analytical capabilities of DB2 Universal Database.

Table 23. Business intelligence information

Name	Form number	PDF file name
<i>IBM DB2 Warehouse Manager Standard Edition Information Catalog Center Administration Guide</i>	SC27-1125	db2dix81
<i>IBM DB2 Warehouse Manager Standard Edition Installation Guide</i>	GC27-1122	db2idx81
<i>IBM DB2 Warehouse Manager Standard Edition Managing ETI Solution Conversion Programs with DB2 Warehouse Manager</i>	SC18-7727	iwhe1mstx80

## DB2 Connect information

The information in this category describes how to access data on mainframe and midrange servers using DB2 Connect Enterprise Edition or DB2 Connect Personal Edition.

Table 24. DB2 Connect information

Name	Form number	PDF file name
<i>IBM Connectivity Supplement</i>	No form number	db2h1x81
<i>IBM DB2 Connect Quick Beginnings for DB2 Connect Enterprise Edition</i>	GC09-4833	db2c6x81
<i>IBM DB2 Connect Quick Beginnings for DB2 Connect Personal Edition</i>	GC09-4834	db2c1x81
<i>IBM DB2 Connect User's Guide</i>	SC09-4835	db2c0x81

## Getting started information

The information in this category is useful when you are installing and configuring servers, clients, and other DB2 products.

Table 25. Getting started information

Name	Form number	PDF file name
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Clients</i>	GC09-4832, not available in hardcopy	db2itx81
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Servers</i>	GC09-4836	db2isx81
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Personal Edition</i>	GC09-4838	db2i1x81
<i>IBM DB2 Universal Database Installation and Configuration Supplement</i>	GC09-4837, not available in hardcopy	db2iyx81
<i>IBM DB2 Universal Database Quick Beginnings for DB2 Data Links Manager</i>	GC09-4829	db2z6x81

## Tutorial information

Tutorial information introduces DB2 features and teaches how to perform various tasks.

Table 26. Tutorial information

Name	Form number	PDF file name
<i>Business Intelligence Tutorial: Introduction to the Data Warehouse</i>	No form number	db2tux81
<i>Business Intelligence Tutorial: Extended Lessons in Data Warehousing</i>	No form number	db2tax81
<i>Information Catalog Center Tutorial</i>	No form number	db2aix81
<i>Video Central for e-business Tutorial</i>	No form number	db2twx81
<i>Visual Explain Tutorial</i>	No form number	db2tvx81

## Optional component information

The information in this category describes how to work with optional DB2 components.

Table 27. Optional component information

Name	Form number	PDF file name
<i>IBM DB2 Cube Views Guide and Reference</i>	SC18-7298	db2aax81
<i>IBM DB2 Query Patroller Guide: Installation, Administration and Usage Guide</i>	GC09-7658	db2dwx81
<i>IBM DB2 Spatial Extender and Geodetic Extender User's Guide and Reference</i>	SC27-1226	db2sbx81

Table 27. Optional component information (continued)

Name	Form number	PDF file name
IBM DB2 Universal Database Data Links Manager Administration Guide and Reference	SC27-1221	db2z0x82
DB2 Net Search Extender Administration and User's Guide	SH12-6740	N/A

**Note:** HTML for this document is *not* installed from the HTML documentation CD.

## Release notes

The release notes provide additional information specific to your product's release and FixPak level. The release notes also provide summaries of the documentation updates incorporated in each release, update, and FixPak.

Table 28. Release notes

Name	Form number	PDF file name
DB2 Release Notes	See note.	See note.
DB2 Installation Notes	Available on product CD-ROM only.	Not available.

**Note:** The Release Notes are available in:

- XHTML and Text format, on the product CDs
- PDF format, on the PDF Documentation CD

In addition the portions of the Release Notes that discuss *Known Problems and Workarounds* and *Incompatibilities Between Releases* also appear in the DB2 Information Center.

To view the Release Notes in text format on UNIX-based platforms, see the Release.Notes file. This file is located in the DB2DIR/Readme/%L directory, where %L represents the locale name and DB2DIR represents:

- For AIX operating systems: /usr/opt/db2\_08\_01
- For all other UNIX-based operating systems: /opt/IBM/db2/V8.1

### Related concepts:

- "DB2 documentation and help" on page 363

### Related tasks:

- "Printing DB2 books from PDF files" on page 380
- "Ordering printed DB2 books" on page 380
- "Invoking contextual help from a DB2 tool" on page 381

---

## Printing DB2 books from PDF files

You can print DB2 books from the PDF files on the *DB2 PDF Documentation* CD. Using Adobe Acrobat Reader, you can print either the entire book or a specific range of pages.

### Prerequisites:

Ensure that you have Adobe Acrobat Reader installed. If you need to install Adobe Acrobat Reader, it is available from the Adobe Web site at [www.adobe.com](http://www.adobe.com)

### Procedure:

To print a DB2 book from a PDF file:

1. Insert the *DB2 PDF Documentation* CD. On UNIX operating systems, mount the DB2 PDF Documentation CD. Refer to your *Quick Beginnings* book for details on how to mount a CD on UNIX operating systems.
2. Open `index.htm`. The file opens in a browser window.
3. Click on the title of the PDF you want to see. The PDF will open in Acrobat Reader.
4. Select **File** → **Print** to print any portions of the book that you want.

### Related concepts:

- “DB2 Information Center” on page 364

### Related tasks:

- “Mounting the CD-ROM (AIX)” in the *Quick Beginnings for DB2 Servers*
- “Mounting the CD-ROM (HP-UX)” in the *Quick Beginnings for DB2 Servers*
- “Mounting the CD-ROM (Linux)” in the *Quick Beginnings for DB2 Servers*
- “Ordering printed DB2 books” on page 380
- “Mounting the CD-ROM (Solaris Operating Environment)” in the *Quick Beginnings for DB2 Servers*

### Related reference:

- “DB2 PDF and printed documentation” on page 375

---

## Ordering printed DB2 books

If you prefer to use hardcopy books, you can order them in one of three ways.

### Procedure:

Printed books can be ordered in some countries or regions. Check the IBM Publications website for your country or region to see if this service is available in your country or region. When the publications are available for ordering, you can:

- Contact your IBM authorized dealer or marketing representative. To find a local IBM representative, check the IBM Worldwide Directory of Contacts at [www.ibm.com/planetwide](http://www.ibm.com/planetwide)
- Phone 1-800-879-2755 in the United States or 1-800-IBM-4YOU in Canada.

- Visit the IBM Publications Center at <http://www.ibm.com/shop/publications/order>. The ability to order books from the IBM Publications Center may not be available in all countries.

At the time the DB2 product becomes available, the printed books are the same as those that are available in PDF format on the *DB2 PDF Documentation CD*. Content in the printed books that appears in the *DB2 Information Center CD* is also the same. However, there is some additional content available in DB2 Information Center CD that does not appear anywhere in the PDF books (for example, SQL Administration routines and HTML samples). Not all books available on the DB2 PDF Documentation CD are available for ordering in hardcopy.

**Note:** The DB2 Information Center is updated more frequently than either the PDF or the hardcopy books; install documentation updates as they become available or refer to the DB2 Information Center at <http://publib.boulder.ibm.com/infocenter/db2help/> to get the most current information.

**Related tasks:**

- “Printing DB2 books from PDF files” on page 380

**Related reference:**

- “DB2 PDF and printed documentation” on page 375

---

## Invoking contextual help from a DB2 tool

Contextual help provides information about the tasks or controls that are associated with a particular window, notebook, wizard, or advisor. Contextual help is available from DB2 administration and development tools that have graphical user interfaces. There are two types of contextual help:

- Help accessed through the **Help** button that is located on each window or notebook
- Infopops, which are pop-up information windows displayed when the mouse cursor is placed over a field or control, or when a field or control is selected in a window, notebook, wizard, or advisor and F1 is pressed.

The **Help** button gives you access to overview, prerequisite, and task information. The infopops describe the individual fields and controls.

**Procedure:**

To invoke contextual help:

- For window and notebook help, start one of the DB2 tools, then open any window or notebook. Click the **Help** button at the bottom right corner of the window or notebook to invoke the contextual help.

You can also access the contextual help from the **Help** menu item at the top of each of the DB2 tools centers.

Within wizards and advisors, click on the Task Overview link on the first page to view contextual help.

- For infopop help about individual controls on a window or notebook, click the control, then click **F1**. Pop-up information containing details about the control is displayed in a yellow window.

**Note:** To display infopops simply by holding the mouse cursor over a field or control, select the **Automatically display infopops** check box on the **Documentation** page of the Tool Settings notebook.

Similar to infopops, diagnosis pop-up information is another form of context-sensitive help; they contain data entry rules. Diagnosis pop-up information is displayed in a purple window that appears when data that is not valid or that is insufficient is entered. Diagnosis pop-up information can appear for:

- Compulsory fields.
- Fields whose data follows a precise format, such as a date field.

**Related tasks:**

- “Invoking the DB2 Information Center” on page 372
- “Invoking message help from the command line processor” on page 382
- “Invoking command help from the command line processor” on page 382
- “Invoking SQL state help from the command line processor” on page 383
- “How to use the DB2 UDB help: Common GUI help”

---

## Invoking message help from the command line processor

Message help describes the cause of a message and describes any action you should take in response to the error.

**Procedure:**

To invoke message help, open the command line processor and enter:

```
? XXXnnnnn
```

where *XXXnnnnn* represents a valid message identifier.

For example, ? SQL30081 displays help about the SQL30081 message.

**Related concepts:**

- “Introduction to messages” in the *Message Reference Volume 1*

**Related reference:**

- “db2 - Command Line Processor Invocation Command” in the *Command Reference*

---

## Invoking command help from the command line processor

Command help explains the syntax of commands in the command line processor.

**Procedure:**

To invoke command help, open the command line processor and enter:

```
? command
```

where *command* represents a keyword or the entire command.

For example, ? catalog displays help for all of the CATALOG commands, while ? catalog database displays help only for the CATALOG DATABASE command.



| **Related tasks:**

- | • “Invoking contextual help from a DB2 tool” on page 381
- | • “Invoking the DB2 Information Center” on page 372
- | • “Invoking message help from the command line processor” on page 382
- | • “Invoking SQL state help from the command line processor” on page 383

| **Related reference:**

- | • “db2 - Command Line Processor Invocation Command” in the *Command Reference*

---

## | Invoking SQL state help from the command line processor

| DB2 Universal Database returns an SQLSTATE value for conditions that could be  
| the result of an SQL statement. SQLSTATE help explains the meanings of SQL  
| states and SQL state class codes.

| **Procedure:**

| To invoke SQL state help, open the command line processor and enter:

|     ? *sqlstate* or ? *class code*

| where *sqlstate* represents a valid five-digit SQL state and *class code* represents the  
| first two digits of the SQL state.

| For example, ? 08003 displays help for the 08003 SQL state, and ? 08 displays help  
| for the 08 class code.

| **Related tasks:**

- | • “Invoking the DB2 Information Center” on page 372
- | • “Invoking message help from the command line processor” on page 382
- | • “Invoking command help from the command line processor” on page 382

---

## | DB2 tutorials

| The DB2® tutorials help you learn about various aspects of DB2 Universal  
| Database. The tutorials provide lessons with step-by-step instructions in the areas  
| of developing applications, tuning SQL query performance, working with data  
| warehouses, managing metadata, and developing Web services using DB2.

| **Before you begin:**

| You can view the XHTML versions of the tutorials from the Information Center at  
| <http://publib.boulder.ibm.com/infocenter/db2help/>.

| Some tutorial lessons use sample data or code. See each tutorial for a description  
| of any prerequisites for its specific tasks.

| **DB2 Universal Database tutorials:**

| Click on a tutorial title in the following list to view that tutorial.

*Business Intelligence Tutorial: Introduction to the Data Warehouse Center*

Perform introductory data warehousing tasks using the Data Warehouse Center.

*Business Intelligence Tutorial: Extended Lessons in Data Warehousing*

Perform advanced data warehousing tasks using the Data Warehouse Center.

*Information Catalog Center Tutorial*

Create and manage an information catalog to locate and use metadata using the Information Catalog Center.

*Visual Explain Tutorial*

Analyze, optimize, and tune SQL statements for better performance using Visual Explain.

---

## DB2 troubleshooting information

A wide variety of troubleshooting and problem determination information is available to assist you in using DB2® products.

### DB2 documentation

Troubleshooting information can be found throughout the DB2 Information Center, as well as throughout the PDF books that make up the DB2 library. You can refer to the "Support and troubleshooting" branch of the DB2 Information Center navigation tree (in the left pane of your browser window) to see a complete listing of the DB2 troubleshooting documentation.

### DB2 Technical Support Web site

Refer to the DB2 Technical Support Web site if you are experiencing problems and want help finding possible causes and solutions. The Technical Support site has links to the latest DB2 publications, TechNotes, Authorized Program Analysis Reports (APARs), FixPaks and the latest listing of internal DB2 error codes, and other resources. You can search through this knowledge base to find possible solutions to your problems.

Access the DB2 Technical Support Web site at

<http://www.ibm.com/software/data/db2/udb/winos2unix/support>

### DB2 Problem Determination Tutorial Series

Refer to the DB2 Problem Determination Tutorial Series Web site to find information on how to quickly identify and resolve problems you might encounter while working with DB2 products. One tutorial introduces you to the DB2 problem determination facilities and tools available, and helps you decide when to use them. Other tutorials deal with related topics, such as "Database Engine Problem Determination", "Performance Problem Determination", and "Application Problem Determination".

See the full set of DB2 problem determination tutorials on the DB2 Technical Support site at

<http://www.ibm.com/software/data/support/pdm/db2tutorials.html>

### Related concepts:

- "DB2 Information Center" on page 364
- "Introduction to problem determination - DB2 Technical Support tutorial" in the *Troubleshooting Guide*

---

## Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. The following list specifies the major accessibility features in DB2<sup>®</sup> Version 8 products:

- All DB2 functionality is available using the keyboard for navigation instead of the mouse. For more information, see “Keyboard input and navigation.”
- You can customize the size and color of the fonts on DB2 user interfaces. For more information, see “Accessible display.”
- DB2 products support accessibility applications that use the Java<sup>™</sup> Accessibility API. For more information, see “Compatibility with assistive technologies” on page 386.
- DB2 documentation is provided in an accessible format. For more information, see “Accessible documentation” on page 386.

### Keyboard input and navigation

#### Keyboard input

You can operate the DB2 tools using only the keyboard. You can use keys or key combinations to perform operations that can also be done using a mouse. Standard operating system keystrokes are used for standard operating system operations.

For more information about using keys or key combinations to perform operations, see Keyboard shortcuts and accelerators: Common GUI help.

#### Keyboard navigation

You can navigate the DB2 tools user interface using keys or key combinations.

For more information about using keys or key combinations to navigate the DB2 Tools, see Keyboard shortcuts and accelerators: Common GUI help.

#### Keyboard focus

In UNIX<sup>®</sup> operating systems, the area of the active window where your keystrokes will have an effect is highlighted.

### Accessible display

The DB2 tools have features that improve accessibility for users with low vision or other visual impairments. These accessibility enhancements include support for customizable font properties.

#### Font settings

You can select the color, size, and font for the text in menus and dialog windows, using the Tools Settings notebook.

For more information about specifying font settings, see Changing the fonts for menus and text: Common GUI help.

#### Non-dependence on color

You do not need to distinguish between colors in order to use any of the functions in this product.

## Compatibility with assistive technologies

The DB2 tools interfaces support the Java Accessibility API, which enables you to use screen readers and other assistive technologies with DB2 products.

## Accessible documentation

Documentation for DB2 is provided in XHTML 1.0 format, which is viewable in most Web browsers. XHTML allows you to view documentation according to the display preferences set in your browser. It also allows you to use screen readers and other assistive technologies.

Syntax diagrams are provided in dotted decimal format. This format is available only if you are accessing the online documentation using a screen-reader.

### Related concepts:

- “Dotted decimal syntax diagrams” on page 386

---

## Dotted decimal syntax diagrams

Syntax diagrams are provided in dotted decimal format for users accessing the Information Center using a screen reader.

In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), they can appear on the same line, because they can be considered as a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that your screen reader is set to read out punctuation. All the syntax elements that have the same dotted decimal number (for example, all the syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, you know that your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, it is preceded by the backslash (\) character. The \* symbol can be used next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element \*FILE with dotted decimal number 3 is given the format 3 \\* FILE. Format 3\* FILE indicates that syntax element FILE repeats. Format 3\* \\* FILE indicates that syntax element \* FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol giving information about the syntax elements. For example, the lines 5.1\*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the

LASTRUN and DELETE syntax elements, the elements must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, this indicates a reference that is defined elsewhere. The string following the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you should refer to separate syntax fragment OP1.

The following words and symbols are used next to the dotted decimal numbers:

- ? means an optional syntax element. A dotted decimal number followed by the ? symbol indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element, (for example 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that syntax elements NOTIFY and UPDATE are optional; that is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.
- ! means a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicates that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the same dotted decimal number can specify a ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In this example, if you include the FILE keyword but do not specify an option, default option KEEP will be applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP only applies to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.
- \* means a syntax element that can be repeated 0 or more times. A dotted decimal number followed by the \* symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1\* data area, you know that you can include one data area, more than one data area, or no data area. If you hear the lines 3\*, 3 HOST, and 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

**Notes:**

1. If a dotted decimal number has an asterisk (\*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
  2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you could write HOST STATE, but you could not write HOST HOST.
  3. The \* symbol is equivalent to a loop-back line in a railroad syntax diagram.
- + means a syntax element that must be included one or more times. A dotted decimal number followed by the + symbol indicates that this syntax element must be included one or more times; that is, it must be included at least once

| and can be repeated. For example, if you hear the line 6.1+ data area, you must  
| include at least one data area. If you hear the lines 2+, 2 HOST, and 2 STATE,  
| you know that you must include HOST, STATE, or both. Similar to the \* symbol,  
| the + symbol can only repeat a particular item if it is the only item with that  
| dotted decimal number. The + symbol, like the \* symbol, is equivalent to a  
| loop-back line in a railroad syntax diagram.

| **Related concepts:**

- | • "Accessibility" on page 385

| **Related tasks:**

- | • "Contents : Common help"

| **Related reference:**

- | • "How to read the syntax diagrams" in the *SQL Reference, Volume 2*

---

## | **Common Criteria certification of DB2 Universal Database products**

| DB2 Universal Database is being evaluated for certification under the Common  
| Criteria at evaluation assurance level 4 (EAL4). For more information about  
| Common Criteria, see the Common Criteria web site at: [http://niap.nist.gov/cc-  
| scheme/](http://niap.nist.gov/cc-scheme/).

---

## Appendix J. Notices

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

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

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

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

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

**The following paragraph does not apply to the United Kingdom or any other country/region 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 that has been exchanged, should contact:

IBM Canada Limited  
Office of the Lab Director  
8200 Warden Avenue  
Markham, Ontario  
L6G 1C7  
CANADA

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 may contain 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 may contain 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:



© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. *\_enter the year or years\_*. All rights reserved.

---

## Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both, and have been used in at least one of the documents in the DB2 UDB documentation library.

ACF/VTAM	iSeries
AISPO	LAN Distance
AIX	MVS
AIXwindows	MVS/ESA
AnyNet	MVS/XA
APPN	Net.Data
AS/400	NetView
BookManager	OS/390
C Set++	OS/400
C/370	PowerPC
CICS	pSeries
Database 2	QBIC
DataHub	QMF
DataJoiner	RACF
DataPropagator	RISC System/6000
DataRefresher	RS/6000
DB2	S/370
DB2 Connect	SP
DB2 Extenders	SQL/400
DB2 OLAP Server	SQL/DS
DB2 Information Integrator	System/370
DB2 Query Patroller	System/390
DB2 Universal Database	SystemView
Distributed Relational Database Architecture	Tivoli
DRDA	VisualAge
eServer	VM/ESA
Extended Services	VSE/ESA
FFST	VTAM
First Failure Support Technology	WebExplorer
IBM	WebSphere
IMS	WIN-OS/2
IMS/ESA	z/OS
	zSeries

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the documents in the DB2 UDB documentation library:

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

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

---

# Index

## A

- accessibility
  - dotted decimal syntax diagrams 386
  - features 385
- active logs 31
- administration notification log 9
- APIs
  - db2ArchiveLog 265
  - db2Backup 71
  - db2HADRStart 193
  - db2HADRStop 197
  - db2HADRTakeover 206
  - db2HistoryCloseScan 267
  - db2HistoryGetEntry 268
  - db2HistoryOpenScan 271
  - db2HistoryUpdate 275
  - db2Prune 278
  - db2ReadLog 286
  - db2ReadLogNoConn 280
  - db2ReadLogNoConnInit 283
  - db2ReadLogNoConnTerm 285
  - db2Recover 152
  - db2Restore 96
  - db2Rollforward 131
  - db2VendorGetNextObj 349
  - db2VendorQueryApiVersion 349
  - sqluvdel 348
  - sqluvend 346
  - sqluvget 343
  - sqluvint 340
  - sqlvput 344
- Archive Active Log API 265
- ARCHIVE LOG command 255
- archive logging 31
- archived logs
  - offline 31
  - online 31
- archiving logs on demand 47
- ASYNCH
  - synchronization mode 181
- Asynchronous Read Log API 286
- automatic client reroute
  - high availability disaster recovery (HADR) 184
- automatic incremental restore, limitations 28
- automatic maintenance
  - backup 3
- automatic restart 9

## B

- backup and restore vendor products 333
- Backup database API 71
- BACKUP DATABASE command 66
- backup image
  - including log files with 49
- Backup Services APIs (XBSA) 66

- backup utility
  - authorities and privileges required to use 62
  - displaying information 59
  - examples 78
  - overview 59
  - performance 79
  - restrictions 62
  - troubleshooting 59
- backups
  - active 53
  - automated 3
  - container names 59
  - expired 53
  - frequency 6
  - images 59
  - inactive 53
  - incremental 25
  - log chain 53
  - log sequence 53
  - offline 6
  - online 6
  - operating system restrictions 9
  - storage considerations 8
  - to named pipes 65
  - to tape 64
  - user exit program 8
- blklogdskful database configuration parameter 35

## C

- cascading assignment 209
- check backup command 243
- check incremental restore image sequence
  - command 247
- circular logging 31, 44
- client reroute
  - high availability disaster recovery (HADR) 184
- clone database, creating 165
- close history file scan API 267
- cluster managers
  - high availability disaster recovery (HADR) 188
- clusters, HACMP 209
- command help
  - invoking 382
- command syntax
  - interpreting 231
- commands
  - ARCHIVE LOG 255
  - BACKUP DATABASE 66
  - db2adutl 237
  - db2ckbkp 243
  - db2ckrst 247
  - db2flsn 249
  - db2fm 171
  - db2inidb 250
  - db2mscs 252
  - INITIALIZE TAPE 257

- commands (*continued*)
  - LIST HISTORY 258
  - PRUNE HISTORY//LOGFILE 260
  - RESTORE DATABASE 87
  - REWIND TAPE 262
  - ROLLFORWARD DATABASE 123
  - SET TAPE POSITION 262
  - UPDATE HISTORY FILE 263
- completion messages 235
- Compression plug-in interface 356
- configuration parameters
  - database logging 35
- configuring
  - high availability disaster recovery 176
- containers
  - names 59
- continuous availability 219
- crash recovery 9
- creating
  - clone database 165

## D

- damaged table space 11
  - non-recoverable 12
  - recoverable 12
- data and parity striping by sectors (RAID level 5) 13
- DATA structure 355
- data structures
  - DB2-INFO 351
  - db2HistData 289
  - INIT-OUTPUT 354
  - RETURN-CODE 355
  - SQLU-LSN 293
  - used by vendor APIs 333
  - VENDOR-INFO 353
- database configuration parameters
  - autorestart 9
- database logs 31
  - configuration parameters 35
- database objects
  - recovery history file 3
  - recovery log file 3
  - table space change history file 3
- database partitions
  - synchronization 121
- databases
  - backup history file 260
  - nonrecoverable 3
  - recoverable 3
  - recovering 123
  - restoring (rebuilding) 87
  - rollforward recovery 22, 123
- DB2 books
  - printing PDF files 380
- DB2 Data Links Manager
  - garbage collection 53
- DB2 Fault Monitor command 171
- DB2 Information Center 364

- DB2 Information Center (*continued*)
  - invoking 372
- DB2 sync point manager (SPM)
  - recovery of indoubt transactions 19
- DB2 tutorials 383
- DB2-INFO structure 351
- db2adutl command 237
- db2ArchiveLog API 265
- db2Backup API 71
- db2ckbkp command 243
- db2ckrst command 247
- db2flsn command 249
- db2fm command 171
- db2HADRStart API 193
- db2HADRStop API 197
- db2HADRTakeover API 206
- db2HistData structure 289
- db2HistoryCloseScan API 267
- db2HistoryGetEntry API 268
- db2HistoryOpenScan API 271
- db2HistoryUpdate API 275
- db2inidb command 250
- db2inidb tool 164
- DB2LOADREC registry variable 120
- db2mscs command 252
- db2Prune API 278
- db2ReadLog API 286
- db2ReadLogNoConn API 280
- db2ReadLogNoConnInit API 283
- db2ReadLogNoConnTerm API 285
- db2Recover API 152
- db2Restore API 96
- db2rfpen command 255
- db2Rollforward API 131
- db2VendorGetNextObj API 349
- db2VendorQueryApiVersion API 349
- Delete Committed Session API 348
- disability 385
- disaster recovery 21
  - high availability (HADR)
    - overview 173
- disk arrays
  - hardware 13
  - reducing failure 13
  - software 13
- disk failure protection 13
- disk mirroring (RAID level 1) 13
- disks
  - RAID (redundant array of independent disks) 13
  - striping 13
- displaying information
  - backup utility 59
- documentation
  - displaying 372
- dotted decimal syntax diagrams 386
- dropped table recovery 118
- DSMICONFIG 325
- DSMIDIR 325
- DSMILOG 325
- dual logging 33
- duplexing (RAID level 1) 13

## E

- enhanced scalability (ES) 209

- error handling
  - log full 35
- error messages
  - during rollforward 131
  - overview 235
- ES (enhanced scalability) 209
- event monitors
  - high-availability on AIX 209

## F

- failback operation 202
- failed database partition server 15
- failover support 161
  - AIX 209
  - idle standby 161
  - mutual takeover 161
  - overview 219
  - Solaris Operating Environment 219
  - Sun Cluster 3.0 221
  - Windows 215
- failure transaction 9
- Fault Monitor Facility 169
- fault tolerance 219
- file systems
  - journaled 161
- Find Log Sequence Number
  - command 249
- flushing logs 31

## G

- garbage collection 53
- Get Next History File Entry API 268

## H

- HACMP (high availability cluster multi-processing) 209
- HADR
  - cluster managers 188
  - configuring 176
  - load operations 176
  - performing a rolling upgrade 203
  - replicated operations 186, 187
  - restrictions 175
  - sample configuration 176
  - synchronization modes 181
  - system requirements 174
- hardware disk arrays 13
- heartbeat 209, 219
- help
  - displaying 372, 374
  - for commands
    - invoking 382
  - for messages
    - invoking 382
  - for SQL statements
    - invoking 383
- high availability 161, 215, 219
- high availability cluster multi-processing (HACMP) 209
- high availability disaster recovery (HADR)
  - system requirements 174

- high availability disaster recovery (HADR)
  - automatic client reroute 184
  - cluster managers 188
  - configuring 176
  - failback 202
  - load operations 176
  - overview 173
  - performing a rolling upgrade 203
  - primary reintegration 202
  - replicated operations 186, 187
  - restrictions 175
  - sample configuration 176
  - switching database roles 202
  - synchronization modes 181
- hot standby configuration 209
- HP-UX
  - backup and restore support 9
- HTML documentation
  - updating 373

## I

- images
  - backup 59
- incremental backup and recovery 25
- incremental restore 26, 83
- index logging
  - high availability disaster recovery (HADR) 185
- indoubt transactions
  - recovering
    - on the host 19
  - recovery
    - with DB2 syncpoint manager 19
    - without DB2 syncpoint manager 20
- Information Center
  - installing 365, 368, 370
- INIT-INPUT structure 354
- INIT-OUTPUT structure 354
- Initialize a Mirrored Database
  - command 250
- Initialize and Link to Device API 340
- Initialize Read Log Without a Database Connection API 283
- INITIALIZE TAPE command 257
- installing
  - Information Center 365, 368, 370
- invoking
  - command help 382
  - message help 382
  - SQL statement help 383

## J

- JFS (Journaled File System)
  - AIX considerations 161
- Journaled File System (JFS)
  - AIX considerations 161

## K

- keepalive packets 209
- keyboard shortcuts
  - support for 385

keywords  
syntax 231

## L

LIST HISTORY command 258  
load copy location file, using  
rollforward 120  
local catchup state 179  
log chain 53  
log file management  
ACTIVATE DATABASE command 42  
log files  
including in backup image 49  
log sequence 53  
log shipping 163  
LOGBUFFSZ configuration parameter 35  
LOGFILSIZ configuration parameter 35  
logging  
archive 31  
circular 31  
indexes  
high availability disaster recovery  
(HADR) 185  
raw devices 48  
LOGPRIMARY configuration  
parameter 35  
logretain configuration parameter 35  
logs  
active 31  
allocation 44  
archiving on demand 47  
circular logging 44  
database 31  
directory, full 47  
flushing 31  
listing during roll forward 123  
managing 42  
mirroring 33  
offline archived 31  
online archived 31  
preventing loss 50  
removal 44  
storage required 8  
user exit program 8  
LOGSECOND configuration parameter  
description 35

## M

media failure  
catalog node considerations 13  
logs 8  
reducing the impact of 13  
message help  
invoking 382  
messages  
overview 235  
Microsoft Cluster Server (MSCS) 215  
mincommit database configuration  
parameter 35  
mirroring  
logs 33  
MIRRORLOGPATH configuration  
parameter 33

mirrorlogpath database configuration  
parameter 35  
monitoring  
progress  
backup 30  
crash recovery 30  
restore 30  
rollforward 30  
MSCS (Microsoft Cluster Server) 215  
multiple instances  
use with Tivoli Storage Manager 325  
mutual takeover configuration 209

## N

Named Pipes  
backing up to 65  
NEARSYNC  
synchronization mode 181  
newlogpath database configuration  
parameter 35  
node synchronization 121  
nodedown event 209  
nodeup event 209  
non-recoverable database  
backup and recovery 3

## O

offline archived logs 31  
on demand log archiving 47  
online  
archived logs 31  
help, accessing 381  
Open History File Scan API 271  
optimizing  
backup performance 79  
restore performance 109  
ordering DB2 books 380  
overflowlogpath database configuration  
parameter 35

## P

parallel recovery 56  
parameters  
syntax 231  
partitioned database environments  
transaction failure recovery in 15  
peer state 179  
pending states 55  
performance  
recovery 56  
point of consistency, database 9  
primary reintegration  
high availability disaster recovery  
(HADR) 202  
printed books, ordering 380  
printing  
PDF files 380  
privileges  
backup 62  
restore utility 82  
roll-forward utility 113  
problem determination  
online information 384

problem determination (*continued*)  
tutorials 384  
protecting against disk failure 13  
Prune History File API 278  
PRUNE HISTORY/LOGFILE  
command 260

## R

RAID (Redundant Array of Independent  
Disks) devices  
description 13  
level 1 (disk mirroring or  
duplexing) 13  
level 5 (data and parity striping by  
sectors) 13  
raw logs 48  
Read Log Without a Database Connection  
API 280  
Reading Data from Device API 343  
Recover Database API 152  
recoverable databases 3  
recovering databases  
overview 145  
recovery  
crash 9  
damaged table spaces 11, 12  
database 87  
dropped table 118  
dropped tables, roll-forward  
utility 118  
history file 3, 51  
incremental 25  
log file 3  
objects 3  
operating system restrictions 9  
overview 3  
parallel 56  
performance 56  
point-in-time 22  
reducing logging 34  
roll-forward 22  
storage considerations 8  
table space change history file 3  
time required 6  
to end of logs 22  
two-phase commit protocol 15  
user exit 329  
version 22  
with rolling forward 123  
without roll forward 87  
redefining table space containers, restore  
utility 85  
redirected restore 85  
reducing impacts  
media failure 13  
transaction failure 15  
reducing logging  
Declared temporary tables 34  
NOT LOGGED INITIALLY  
parameter 34  
Redundant Array of Independent Disks  
(RAID)  
reducing the impact of media  
failure 13  
registry variables  
DB2LOADREC 120

- relationships
  - between tables 9
- remote catchup pending state 179
- remote catchup state 179
- replicated operations
  - high availability disaster recovery (HADR) 186, 187
- Reset rollforward pending state
  - command 255
- RESTART DATABASE command 9
- Restore database API 96
- RESTORE DATABASE command 87
- restore utility
  - authorities and privileges required to use 82
  - examples 107
  - overview 81
  - performance 81, 109
  - redefining table space containers 85
  - restoring to a new database 87
  - restoring to an existing database 86
  - restrictions 82
- restoring
  - automatic incremental, limitations 28
  - data to a new database 87
  - data to an existing database 86
  - databases
    - incremental 25
    - rollforward recovery 22
  - earlier versions of DB2 databases 87
  - incremental 26, 83
- restrictions
  - high availability disaster recovery (HADR) 175
- RETURN-CODE structure 355
- REWIND TAPE command 262
- roll-forward recovery
  - configuration file parameters
    - supporting 35
  - database 22
  - log management considerations 42
  - log sequence 42
  - table space 22, 114
- Rollforward Database API 131
- ROLLFORWARD DATABASE
  - command 123
- rollforward utility
  - authorities and privileges required to use 113
  - examples 140
  - load copy location file, using 120
  - overview 111
  - recovering a dropped table 118
  - restrictions 113
- rolling upgrade
  - performing 203
- rotating assignment 209
- rules file 209

## S

- scalability 209
- seed database 86, 87
- SET TAPE POSITION command 262
- Set Up Windows Failover utility
  - command 252

- site failure
  - high availability disaster recovery (HADR) 173
- software disk arrays 13
- Solaris Operating Environment
  - backup and restore support 9
- SP frame 209
- split mirror
  - as backup image 168
  - as standby database 167
  - handling 164
- SQL messages 235
- SQL statement help
  - invoking 383
- SQLCODE
  - overview 235
- SQLSTATE
  - overview 235
- SQLU-LSN structure 293
- sqluvdel API 348
- sqluvend API 346
- sqluvget API 343
- sqluvint API 340
- sqluvput API 344
- standby database
  - states 179
- Start HADR API 193
- states
  - pending 55
  - standby database 179
- Stop HADR API 197
- storage
  - media failure 8
  - requirements
    - backup and recovery 8
- Sun Cluster 3.0, high availability 221
- suspended I/O to support continuous availability 164
- switching database roles
  - high availability disaster recovery (HADR) 202
- SYNC
  - synchronization mode 181
- synchronization
  - database partition 121
  - node 121
  - recovery considerations 121
- synchronization modes
  - high availability disaster recovery (HADR) 181
- syntax diagrams
  - reading 231
- system requirements
  - high availability disaster recovery (HADR) 174

## T

- table spaces
  - recovery 11, 12
  - restoring 22
  - roll-forward recovery 22
- tables
  - relationships 9
- Take Over as Primary Database API 206
- tape backup 64, 66

- Terminate Read Log Without a Database Connection API 285
- time
  - database recovery time 6
- timestamps
  - conversion, client/server environment 122, 146
- Tivoli Storage Manager (TSM)
  - backup restrictions 325
  - client setup 325
  - timeout problem resolution 325
  - using 325
  - with BACKUP DATABASE
    - command 325
  - with RESTORE DATABASE
    - command 325
- transactions
  - blocking when log directory is full 47
  - failure recovery
    - crashes 15
    - n the failed database partition server 15
    - on active database partition server 15
    - reducing the impact of failure 9
- troubleshooting
  - online information 384
  - tutorials 384
- TSM archived images 237
- tutorials 383
  - troubleshooting and problem determination 384
- two-phase commit
  - protocol 15

## U

- Unlink the Device and Release its Resources API 346
- Update History File API 275
- UPDATE HISTORY FILE command 263
- Updating
  - HTML documentation 373
- user exit programs
  - archive and retrieve
    - considerations 45
  - backup 8
  - calling format 329
  - error handling 329
  - for database recovery 329
  - logs 8
  - sample programs 329
- user-defined events 209
- userexit database configuration
  - parameter 35

## V

- variables
  - syntax 231
- vendor products
  - backup and restore 333
  - DATA structure 355
  - description 333
  - INIT-INPUT structure 354

- vendor products (*continued*)
  - operation 333
- VENDOR-INFO structure 353
- VERITAS Cluster Server 224
  - high availability 224
- version levels
  - version recovery of the database 22

## **W**

- warning messages
  - overview 235
- Windows
  - failover 215
- Work with TSM Archived Images
  - command 237
- Writing Data to Device API 344

## **X**

- XBSA (Backup Services APIs) 66





---

## Contacting IBM

In the United States, call one of the following numbers to contact IBM:

- 1-800-IBM-SERV (1-800-426-7378) for customer service
- 1-888-426-4343 to learn about available service options
- 1-800-IBM-4YOU (426-4968) for DB2 marketing and sales

In Canada, call one of the following numbers to contact IBM:

- 1-800-IBM-SERV (1-800-426-7378) for customer service
- 1-800-465-9600 to learn about available service options
- 1-800-IBM-4YOU (1-800-426-4968) for DB2 marketing and sales

To locate an IBM office in your country or region, check IBM's Directory of Worldwide Contacts on the web at <http://www.ibm.com/planetwide>

---

## Product information

Information regarding DB2 Universal Database products is available by telephone or by the World Wide Web at <http://www.ibm.com/software/data/db2/udb>

This site contains the latest information on the technical library, ordering books, product downloads, newsgroups, FixPaks, news, and links to web resources.

If you live in the U.S.A., then you can call one of the following numbers:

- 1-800-IBM-CALL (1-800-426-2255) to order products or to obtain general information.
- 1-800-879-2755 to order publications.

For information on how to contact IBM outside of the United States, go to the IBM Worldwide page at [www.ibm.com/planetwide](http://www.ibm.com/planetwide)







Printed in USA

SC09-4831-01



Spine information:



IBM® DB2 Universal Database™

Data Recovery and High Availability Guide and Reference

Version 8.2