IBM Tivoli Storage Manager for Mail
Version 7.1

Data Protection for IBM Domino for Windows

Installation and User's Guide

IBM Tivoli Storage Manager for Mail
Version 7.1

Data Protection for IBM Domino for Windows

Installation and User's Guide

Note:
Before using this information and the product it supports, read the information in "Notices" on page 201.

First edition (December 2013)

This edition applies to version 7, release 1, modification 0 of IBM Tivoli Storage Manager for Mail Data Protection for IBM Domino (product number 5608-E06), and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Tables

About this publication

IBM® Tivoli® Storage Manager for Mail: Data Protection for IBM Domino® is a storage management software product that provides storage management services in a multiplatform environment. This publication explains how to install, configure, and administer Data Protection for IBM Domino.

Who should read this guide

The target audience for this publication includes system installers, system users, administrators, database administrators, Domino administrators, and system administrators. It explains the procedures that are needed to install and customize Data Protection for IBM Domino.

In this publication, it is assumed that you have an understanding of the following applications:

- IBM DB2® UDB for Windows
- Lotus® Domino Server
- Tivoli Storage Manager server
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager Application Program Interface

It is assumed that you have an understanding of the Windows Server operating system.

Publications

Publications for the Tivoli Storage Manager family of products are available online. The Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM Tivoli.

To search across all publications or to download PDF versions of individual publications, go to the Tivoli Storage Manager information center at http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1.

You also can find the Tivoli Storage Manager product family information centers and other information centers that contain official product documentation for current and previous versions of Tivoli products at Tivoli Documentation Central. Tivoli Documentation Central is available at http://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli Documentation Central.

Reading syntax diagrams

This section describes how to read the syntax diagrams used in this book. To read a syntax diagram, follow the path of the line. Read from left to right, and top to bottom.

- The ►►— symbol indicates the beginning of a syntax diagram.
- The —► symbol at the end of a line indicates the syntax diagram continues on the next line.
- The ►— symbol at the beginning of a line indicates a syntax diagram continues from the previous line.
- The —►◄ symbol indicates the end of a syntax diagram.

Syntax items, such as a keyword or variable, can be:

- On the line (required element)
- Above the line (default element)
- Below the line (optional element).

Syntax Diagram Description	Example
Abbreviations: Uppercase letters denote the shortest acceptable truncation. If an item appears entirely in uppercase letters, it cannot be truncated. You can type the item in any combination of uppercase or lowercase letters. In this example, you can enter KEYWO, KEYWORD, or KEYWOrd.	►►—KEYWOrd—►◄
Symbols: Enter these symbols exactly as they appear in the syntax diagram.	* Asterisk { } Braces : Colon , Comma = Equal Sign - Hyphen () Parentheses . Period Space
Variables: Italicized lowercase items (*var_name*) denote variables. In this example, you can specify a *var_name* when you enter the KEYWORD command.	►►—KEYWOrd—*var_name*—►◄

Syntax Diagram Description	Example
Repetition: An arrow returning to the left means you can repeat the item. A character or space within the arrow means you must separate repeated items with that character or space. A footnote by the arrow references the number of times you can repeat the item.	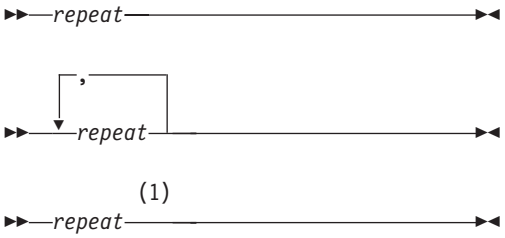 **Notes:** 1 Specify *repeat* as many as 5 times.
Required Choices: When two or more items are in a stack and one of them is on the line, you *must* specify one item. In this example, you *must* choose A, B, or C.	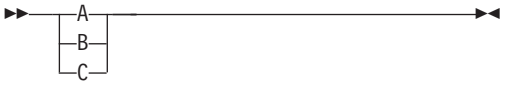
Optional Choice: When an item is below the line, that item is optional. In the first example, you can choose A or nothing at all. When two or more items are in a stack below the line, all of them are optional. In the second example, you can choose A, B, C, or nothing at all.	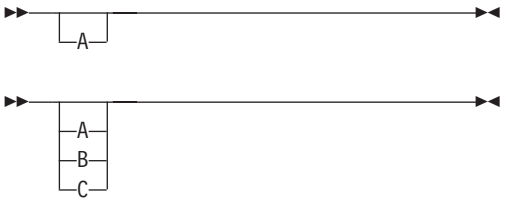
Defaults: Defaults are above the line. The default is selected unless you override it. You can override the default by including an option from the stack below the line. In this example, A is the default. You can override A by choosing B or C. You can also specify the default explicitly.	
Repeatable Choices: A stack of items followed by an arrow returning to the left means you can select more than one item or, in some cases, repeat a single item. In this example, you can choose any combination of A, B, or C.	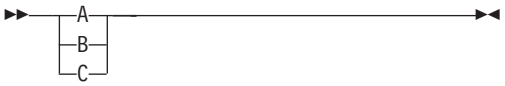

Syntax Diagram Description	Example
Syntax Fragments: Some diagrams, because of their length, must fragment the syntax. The fragment name appears between vertical bars in the diagram. The expanded fragment appears between vertical bars in the diagram after a heading with the same fragment name.	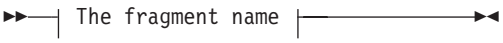 **The fragment name:** 

New for Data Protection for IBM Domino Version 7.1

Read about the new features and other changes in Data Protection for IBM Domino Version 7.1.

- Support for IBM Domino Servers running on 64-bit Linux x86 hardware
- Support for IBM Domino 9
- Support for Windows Server 2012
- Automated failover for data recovery

Chapter 1. Overview of Data Protection for IBM Domino

An overview of the features and capabilities of IBM Tivoli Storage Manager for Mail: Data Protection for IBM Domino is provided.

How to protect data with Data Protection for IBM Domino

The key features used to protect your data with Data Protection for IBM Domino are presented.

Data Protection for IBM Domino on Windows is used to back up and restore Domino databases and transaction logs. When archival logging is used on the Domino server, it archives transaction log files and retrieves them as required for a database recovery. Database backups and archived transaction log files are stored on theTivoli Storage Manager server.

Data Protection for IBM Domino communicates with a Tivoli Storage Manager server with the Tivoli Storage Manager application programming interface (API). Data Protection for IBM Domino communicates with a Domino server through the Domino API.

Note: Procedures that are run from the command line can also be run from the GUI and the Tivoli Storage Manager Web Client GUI.

Tasks

Protect and manage Domino server data with Data Protection for IBM Domino and these actions:

- Back up Domino NSF databases.
- Back up DB2 enabled Notes® databases when a DB2 enabled Domino server is available.
- Restore DB2 enabled Notes databases when a DB2 enabled Domino server is available.
- Maintain multiple backup versions of Domino databases.
- Archive Domino transaction log files when archival logging is in effect.
- Restore backup versions of a Domino database and apply changes since the last backup from the transaction log.
- Restore Domino databases to a specific point in time.
- Restore one or more archived transaction log files.
- Expire database backups that are automatically based on version limit and retention period.
- Expire archived transaction log files when no longer needed.
- Obtain online context-sensitive, task, and conceptual help.
- View online documentation for Data Protection for IBM Domino.
- Automate scheduled backups.
- Restore Domino databases to a different server or partition.
- Access Data Protection for IBM Domino remotely using the Tivoli Storage Manager web client.

- Access Data Protection for IBM Domino using the client GUI based on Oracle Java™.
- Access Data Protection for IBM Domino using the command-line interface.

Backup NSF databases

The different types of Domino NSF backups available with Data Protection for IBM Domino are described.

Domino NSF database backup and transaction log archive

Concepts that are associated with Data Protection for IBM Domino backups of Domino databases and transaction logs are provided.

The backup and recovery API in Domino provides the capability to do the following actions:

- Run full online backups of individual databases.
- Store archives of the transaction log when archival logging is in effect.

Domino server transaction log

Updates to a logged database are recorded in the Domino server transaction log so that full database backups are not required as frequently. Changes to a database since the last full backup can be applied from the transaction log after the backup is restored from the last full backup. Enabling transaction logs for all databases on a Domino server is not required, so the backup process must handle both logged and non-logged databases. Domino allows the active transaction log to be backed up also.

Transactions that are recorded in the transaction log are keyed by a Database Instance Identifier (DBIID), which is unique for each database on a Domino server. The DBIID must match that of a restored database for transactions in the log to be applied to the database. The most common reason for a DBIID to change is compaction of the database to reduce file size. When the DBIID changes, a full backup must be taken so that subsequent updates can be applied to a restored backup of that database. Transactions that are recorded since the DBIID change cannot be applied to prior backups of that database because the DBIID does not match. For more information about DBIID, see the Domino server documentation.

Types of NSF backup and archive logs

Data Protection for IBM Domino provides two types of database backups and an archive log function, incremental and selective.

Incremental Backup

An incremental backup provides a conditional backup function for a full online backup of Domino databases under the following conditions:

- The database is not excluded in the Tivoli Storage Manager include-exclude options file (standard include and exclude processing is supported).
- The database is not logged and was modified since the last active backup image for that database. Both data and non-data modification dates are checked. If either is different from the date of the active backup, the database is backed up.
- When circular logging is used on the Domino server, or when logging is disabled on the Domino server, transaction log files are not archived.

- The database is new or newly included in the backup and an active backup image does not exist on the Tivoli Storage Manager server.

The **incremental** command includes a function that determines when active backup database copies exist on the Tivoli Storage Manager server that are deleted from the Domino server or excluded from backup. If so, they are marked inactive so that automatic expiration of these backup copies can occur according to defined Tivoli Storage Manager management class parameters for backup files.

Selective Backup
A selective backup unconditionally runs a full online backup of the specified Domino databases. It does not run excluded backups that are specified through exclude statements in the Data Protection for IBM Domino `dsm.opt` file.

Archive Log
An archive log stores filled transaction log files on the Tivoli Storage Manager server so that space allocated to these files can be reused by the Domino logger. The **archivelog** command is available when transaction logging on the Domino server is enabled in archival mode. Filled transaction log files must be archived frequently enough to ensure the transaction log never fills completely and stops the Domino server.

Transaction log files that are stored on the Tivoli Storage Manager server are automatically restored as needed for a database recovery.

Archived transaction log files are retained on the Tivoli Storage Manager server when a database backup exists that needs these log files for a complete recovery. See "Expiration of NSF archived transaction log files" for further details.

Note: Transaction log files are not archived when circular or linear loop logging is used on the Domino server. If logging is disabled on the Domino Server, log files are not archived. For more information, see "NSF backup strategy" on page 4.

Expiration of NSF archived transaction log files

Concepts that are associated with expiring archived transaction log files are provided.

The **inactivatelogs** command expires transaction log files from backup storage. There is a single shared transaction log for all logged databases on a Domino server. Log files cannot be deactivated or allowed to expire until all databases that require that log file for recovery are inactive. This command queries the database backups on the Tivoli Storage Manager server to determine which log files are required by any active database backup. This command also deactivates log files that are no longer required when the database backups are deactivated. Deactivate the transaction logs for a deactivated database by running the **inactivatelogs** command after a full database backup completes.

NSF backup strategy

Implement NSF backup strategies that are based on your requirements such as network traffic, backup window, and acceptable restore times. Information about backup strategies, and sample strategies are provided.

Your choice of NSF backup strategy includes selecting the type of backup commands to use and the type of transaction logging to be done on the Domino server. Data Protection for IBM Domino can back up transaction logs from a Domino server that has archival logging enabled. Transaction logs cannot be backed up from a Domino server that has circular or linear loop logging in effect.

Archival logging allows transaction log data to be **archived** on the Tivoli Storage Manager server so that changes to logged databases can be stored on the Tivoli Storage Manager server without having to run a full backup. Less frequent full database backups result because changes to logged databases are available for restore in the archived transaction log files.

The **archivelog** command backs up Domino transaction log files when archival logging is in effect on the Domino server. The command queries the Domino server to determine whether any log extents are ready for archiving. If so, the log files are backed up to Tivoli Storage Manager server storage, and the Domino server is notified of their availability for reuse.

High and low threshold values can be specified as a percentage of the log capacity to control whether log files can be archived. These values allow the command to be scheduled regularly to protect against a log full condition. The log files are archived only if the log is getting close to being full.

Consider the following information when you are choosing a backup strategy:

- When you are using archival transaction logging, the frequency of **archivelog** command usage depends on the size of your log and the rate of change for logged databases. Run archival transaction logging several times a day if you generate a large volume of changes at a rapid rate.
- When a DBIID for a logged database changes, the database cannot be recovered until another backup of that database is run. The **incremental** command detects the changed DBIID. Any changes that are recorded in the log between the DBIID change and backup are not restored if the original database is lost. The Domino server sends a message to the server console when a DBIID change occurs. It is useful to monitor the server console and run a backup when the DBIID changes.
- When you are restoring a group of logged databases for which transactions must be applied, activate them together when possible. Activating them together avoids restoring the same transaction log files multiple times. Restored transaction log files are deleted during a database recovery by the Domino server. Activating and applying logs to the database separately requires retransmitting log files for each database.
- When querying the Tivoli Storage Manager server for databases to restore on the Data Protection for IBM Domino GUI, reduce query processing time by specifying a database name and typing letters with a wildcard character (*) in the **By Database Name** field.
- Data Protection for IBM Domino provides backup and restore functions for the Domino databases (including template files) and associated transaction logs. However, Data Protection for IBM Domino does not provide a complete disaster recovery solution for a Domino server by itself. There are executable and configuration files that are part of the Domino server installation, which are not

backed up. For example, database link files have an `nsf` extension but are not considered databases and are not backed up by Data Protection for IBM Domino. These files must be recovered in a disaster recovery situation. A comprehensive disaster recovery plan can be achieved using the Tivoli Storage Manager Backup Archive Client for your server together with Data Protection for IBM Domino.

- Personal copies (replicas) of Domino databases that are stored on Notes clients (not on the Domino server) are not protected by Data Protection for IBM Domino. You can use the Tivoli Storage Manager Backup-Archive Client on the Notes client to back up and restore these files or rely on Domino server replication if you must recover them.
- To restore an individual Notes document, you must restore the entire database to a different name. Choose a time when the document existed for both the **restore** `/pit` and **activate** `/applylogs` commands but before the document was deleted. Copy the required document on the Notes client.
- The Tivoli Storage Manager encryption, deduplication, and compression functions can be used with Data Protection for IBM Domino. For more information read the Using the application programming interface documentation on the Tivoli Storage Manager information center.

Sample strategies

Some strategies that you can employ are described here.

Full backups only

The following backup option can be implemented if your network capacity and backup window support regular full database backups:

- Perform regular Selective backups.
- Perform occasional Incremental backups to deactivate backup copies of databases that no longer exist on the Domino server.

Each backup takes longer to run, but the restore process is most efficient because only the most recent full backup is restored.

Note: You can apply updates to the restored database from the transaction log if the log is not wrapped since the backup was run. If the log is wrapped, the attempt to apply logs fails.

Full backup plus transaction log archives

It is often not practical to back up entire databases with each regular backup for large Domino installations. Archival logging captures changes to all logged databases in the archived transaction log files. Full database backups are run less frequently, reducing burdens on network and storage resources. To implement this strategy:

- Run regular log archives by running the **archivelog** command. The log file does not fill and it captures changes to logged databases.
- Run regular **Incremental** backups, which back up logged databases with unchanged DBIIDs.
- Run occasional **Selective** backups of all logged databases. The number of transaction log files to be processed during a restore is reduced.
- Issue the **inarchivelog** command (following Selective backups) to allow nonessential transaction log files to expire.

The **archivelog** command captures changes to all logged databases in between full backups of selected databases. To restore a database to its most recent state, restore the most recent database backup and specify

/applylogs. All the necessary archived transaction log files are restored so that updates for the database can be applied.

Backup DB2 enabled Notes databases

The types of DB2 enabled Notes database backups available with Data Protection for IBM Domino are described.

DB2 enabled Notes database backup

Concepts that are associated with Data Protection for IBM Domino backups of Domino DB2 enabled Notes databases are presented.

The following list provides a brief overview of key DB2 enabled Notes database backup features:

- The entire Domino DB2 database or separate DB2 Groups can be backed up.
- The backup can be restored to an alternative database.
- In a disaster recovery situation, the backup can be restored to the original Domino DB2 database.
- Individual DB2 enabled Notes databases are copied from the alternative DB2 database to the Domino DB2 database.

DB2 enabled Notes databases are stored in a DB2 database and managed by a DB2 server. Using Data Protection for IBM Domino you can back up the Domino 8 DB2 database and DB2 Groups (table space). Online and DB2 Group backups are only available when the Domino DB2 database is enabled for rollforward recovery. When rollforward recovery is not enabled, the backup is run offline.

Note, a DB2 enabled Notes backup is different from an NSF backup. An NSF database is backed up directly. A DB2 enabled Notes database is backed up indirectly as a DB2 Group. A DB2 Group (or DB2 table space) is really a collection of one or more DB2 enabled Notes databases.

The table space is the smallest entity that can be backed up in DB2 applications. Since it is possible for a table space (DB2 Group) to contain more than one DB2 enabled Notes database, it is not possible to back up and restore a single DB2 enabled Notes database unless there is only one DB2 enabled Notes database in a table space (DB2 Group). A single DB2 enabled Notes database can be copied from a restored table space to the original table space. Alternatively, it can be restored to a new table space with the Domino FastCopy feature. FastCopy can be used to restore only a single DB2 enabled Notes database when a DB2 redirected restore is used. Note, a DB2 redirected restore, and Domino FastCopy are only possible when DB2 is configured with federation enabled.

DB2 enabled Notes databases, NSF databases that are created and stored in DB2 databases, are represented in the Domino Data directory as regular files and are similar to Domino database link and directory link files. They contain the {DB2} text string to identify that the database is stored in a DB2 database. The DB2 enabled Notes link files are not backed up by Data Protection for IBM Domino. Only the full DB2 database or the DB2 Groups (that contain DB2 enabled Notes databases) are backed up by Data Protection for IBM Domino.

DB2 Tivoli Storage Manager Agent

DB2 provides a Tivoli Storage Manager Agent and utility program, `db2adutl`, that interfaces with the DB2 Recovery API. This interface helps you to manage Tivoli Storage Manager objects that are created on the DB2 server. Data Protection for IBM Domino uses the DB2 Tivoli Storage Manager Agent through the DB2 Recovery API to back up and restore the Domino DB2 database and DB2 Groups (table space). These Tivoli Storage Manager objects that are associated with DB2 backups are unique and there is only one Tivoli Storage Manager object that is created for each backup operation per session. The `db2adutl` program, for example, can be used to expire these objects.

Data Protection for IBM Domino and the DB2 API

Data Protection for IBM Domino uses the DB2 Recovery API to communicate with the DB2 Tivoli Storage Manager Agent to back up DB2 data to the Tivoli Storage Manager server. Configure the DB2 Tivoli Storage Manager Agent to use the same Tivoli Storage Manager `nodename` and to access the same Tivoli Storage Manager server as Data Protection for IBM Domino. The Tivoli Storage Manager objects that are created by the DB2 Recovery API belong to the same Tivoli Storage Manager node as the objects created by Data Protection for IBM Domino NSF databases. Specify the options file with the `adsmoptfile` parameter.

Types of DB2 backups

Data Protection for IBM Domino provides three types of database backups:

DB2 database Backup
Data Protection for IBM Domino DB2 database backups create a selective backup image that can be used for disaster recovery of the Domino 8 DB2 database or for restoring individual DB2 enabled Notes databases. Only selective backup `db2selective` is provided for DB2 enabled Notes databases.

DB2 Group (table space) Backup
Data Protection for IBM Domino DB2 Group backups create a selective table space backup image. This type of backup can be run only after the DB2 database is enabled for rollforward recovery.

Full DB2 database and NSF database Backup
Data Protection for IBM Domino can run a selective NSF database backup and a full Domino DB2 database backup in a single operation.

Expiration of DB2 backups and transaction log objects

Concepts that are associated with expiring DB2 backup objects and DB2 transaction log files are described.

Data Protection for IBM Domino uses the DB2 Recovery API to access the DB2 Tivoli Storage Manager Agent. When Data Protection for IBM Domino processes a backup, it informs the DB2 Tivoli Storage Manager Agent to back up the DB2 data to the Tivoli Storage Manager server. During backup processing, Data Protection for IBM Domino creates a group of Tivoli Storage Manager objects. These objects describe the contents of each Tivoli Storage Manager object that is created by the DB2 Tivoli Storage Manager Agent. Each object describes the type of backup and the name of the DB2 enabled Notes databases that are contained in the backup. The Tivoli Storage Manager group object has a reference to the object created by the DB2 Tivoli Storage Manager Agent. Policy settings are applied to the Tivoli

Storage Manager group object. As a result, when a backup version is no longer needed, the objects that are referenced by the Tivoli Storage Manager group object must also be deactivated.

These Tivoli Storage Manager group objects can be deactivated by using the `db2inactivateobjs` command. This command displays how to issue the DB2 Tivoli Storage Manager Agent `db2adutl` utility to deactivate these objects. The `db2adutl` utility ensures that information about the DB2 server remains consistent after objects are deactivated.

The Domino DB2 database transaction logs are archived automatically by the DB2 server (with the DB2 Tivoli Storage Manager Agent) to the Tivoli Storage Manager server. The `db2archivelog` command forces a backup of the Domino DB2 database transaction log file. This command can be used to ensure that the latest updates are available during an alternate DB2 database rollforward to the current time operation.

Note: Because transaction log file names are unique, they do not expire because of version limit.

Archived transaction log files are retained on the Tivoli Storage Manager server when a database backup exists that needs these log files for a complete recovery.

DB2 enabled Notes database backup strategy considerations

Factors to consider when you are planning your DB2 enabled Notes database backup strategy, and sample strategies are presented.

You can choose different backup strategies that are dependent on your specific requirements. Requirements to be considered include network traffic, backup window, and acceptable restore times. Your choice of strategy includes selecting the type of DB2 backup commands to use.

Note: The DB2 commands do not return information about when a backup to a Tivoli Storage Manager server was compressed, encrypted, sent LAN-free or de-duplicated.

Sample strategies

Some strategies that you can employ are described.

Full DB2 database backups only

This backup strategy can be followed when the Domino DB2 database is enabled for rollforward recovery:

- Perform full DB2 database backups regularly.
- Routinely deactivate and delete DB2 objects from the Tivoli Storage Manager server that are no longer needed.

A full DB2 database backup completes quicker and requires less storage space than DB2 Group backups. However, DB2 enabled Notes databases cannot be restored to a specific point-in-time since the database is not enabled for rollforward recovery and requires less storage space than backing up all the DB2 Groups individually.

Full DB2 database backups plus DB2 Group backups

This backup strategy can be followed when the Domino DB2 database is enabled for rollforward recovery:

- Perform full DB2 database backups regularly.
- Perform DB2 Group backups regularly in between full DB2 database backups. Note only those DB2 Groups with the strictest restore time requirements must be backed up.
- Maintain a complete set of transaction log files to a specified point-in-time. DB2 automatically archives the transaction logs when the DB2 database is enabled for rollforward recovery.
- Routinely deactivate and delete DB2 objects from the Tivoli Storage Manager server that are no longer needed.

To restore a DB2 enabled Notes database to its most recent time, select the most recent backup from a DB2 Group backup or from a full DB2 database backup that contains the DB2 enabled Notes database. If the most recent DB2 Group backup is not available, restore the DB2 Group from the most recent full DB2 database backup. Note, this type of restore is to an alternative DB2 database. Rollforward the DB2 Group or full DB2 database backup. Activate and copy the DB2 enabled Notes database to the alternate Domino DB2 database.

Environments that contain both NSF and DB2 enabled Notes databases

Domino 8 environments that contain both NSF and DB2 enabled Notes databases can implement the following backup strategy:

- Perform full DB2 database backups and NSF selective backups regularly.
- Perform routine incremental backups of NSF databases to deactivate backup copies that were deleted from the Domino server.
- Perform regular DB2 Group backups if the DB2 database is enabled for rollforward recovery.
- Perform routine archiving of the transaction log files if archival transaction logging is enabled on the Domino server.
- Routinely deactivate the Domino server log file and routinely deactivate and delete DB2 objects from the Tivoli Storage Manager server.

Restore (NSF databases)

Concepts that are associated with restoring and activating Domino databases and archived transaction logs are described.

A Domino database recovery can involve restoring several transaction log files in addition to the database backup file from the Tivoli Storage Manager server, depending on the backup strategy you choose. The function to restore database files is separate from the function that applies updates from the transaction log. You can restore database files separately while transaction logs are processed for all restored databases. This avoids restoring the same transaction log files multiple times. Restoring and updating a database with current changes from the transaction log is a two-step process that is implemented by the `restore` and `activatedbs` commands.

For more information about backup and restore strategies, see “NSF backup strategy” on page 4.

Domino database restore and activation

Concepts that are associated with restoring a Domino database and activating the archived transaction logs are described.

Restore

Restore is the first step of a two-stage recovery process. This function restores a single database or group of databases from Tivoli Storage Manager storage to the Domino server. You can restore the database to a different database file name or to a different Domino server. You can also restore a group of databases to a different directory and preserve existing file names. In addition, if you specify a point in time on the `restore` command, the most recent backup version before that time is restored. To restore a database without applying updates from the transaction log, the two steps can be combined into one step by specifying `/activate=yes` during the `restore` command.

Activation

Activation is the second step of the two stage recovery process. This function brings restored databases online for use by the Domino server. You can optionally apply transactions from the transaction log to update the database. Transactions can be applied up to a specific point in time or up through the most recent changes that are recorded in the transaction log. If archival logging is in effect, Data Protection for IBM Domino automatically restores archived transaction log files as needed.

The Domino server provides an alternative restore path feature used to specify the directory where transaction logs are restored. You can use this feature with the `activatedbs` command. See "**Domdsmc activatedbs**" on page 50 for details about this procedure.

The `query pendingdbs` command retrieves a list of restored databases not yet activated. Databases awaiting activation are assigned a temporary file name to avoid recognition as database files on the Domino Server.

Restore of archived transaction logs

Concepts that are associated with restoring archived transaction logs are described.

This function allows a single, archived transaction log file to be restored independently of a routine database restore. Restoring a single, archived transaction log file assists with disaster recovery operations. By retrieving the most recent archived log file, it is possible to rebuild the Domino transaction log control file. Even after a loss of the active transaction log, archived transaction log files can be used to recover restored database backups to a more current state. More than one archived transaction log file can be restored at a time.

Run the following command to restore an archived log:

```
Domdsmc restorelogarchive log_name
```

For more information about using archived transaction logs in disaster recovery procedures, see "Recovering from loss of Domino transaction logs for NSF databases" on page 166.

Restore at document level

Concepts that are associated with restoring a Domino database at the document level are described.

Data Protection for IBM Domino restores Domino databases at the database level. To restore a document in a database, the entire database must first be restored and the document copied.

A database can be restored to the production server under a temporary name, and the documents can be copied to the appropriate database. If for performance reasons, the production server cannot be used in the restore process, the database can be restored to an alternative server and copied to the production server. Running restore operations for the alternative server to reduce demands on the production Domino server is a good idea. Server restores can be run to an alternative partition or to a separate Domino server. For instructions for restoring to an alternate server or partition, see "NSF databases restore to alternate server and alternate partition" on page 167.

Restore (DB2 enabled Notes databases)

Concepts that are associated with restoring and activating Domino DB2 enabled Notes databases are described.

For more information about backup and restore strategies, see "DB2 enabled Notes database backup strategy considerations" on page 8.

Domino DB2 enabled Notes database restore, rollforward, and activation

Concepts that are associated with Domino DB2 enabled Notes database restore, rollforward, and activation are provided.

Restore

Using Data Protection for IBM Domino you can restore a single DB2 enabled Notes database or a group of DB2 enabled Notes databases. A Domino 8 DB2 Group can be restored from either a full DB2 database backup image or a DB2 table space backup image. Only one DB2 Group can be restored at a time if the DB2 Group is being restored from a DB2 Group backup. The DB2 Group is restored to an alternative DB2 database within the same DB2 instance. If more than one DB2 Group is restored, each DB2 Group must be restored to a different DB2 database. Otherwise, restoring more than one DB2 Group to the same alternate DB2 database overwrites the previously restored DB2 Group. If the DB2 Group is being restored from a full DB2 backup image, then more than one DB2 Group can be restored to the same alternate DB2 database.

A Domino 8 DB2 database can be restored from a full DB2 database backup image to an alternate DB2 database. Using the alternative database, frees the individual DB2 enabled Notes databases for restore. The DB2 database can also be restored directly to the Domino DB2 database. This type of restore operation is useful for disaster recovery purposes.

Rollforward

Rollforward is an intermediate step that is required when the Domino DB2 database is enabled for rollforward recovery. This task rolls the Domino DB2

database forward to the specified point in time and marks the rollforward as complete. The DB2 database can be an alternate DB2 database or the Domino DB2 database.

The **"Domdsmc query DB2rollforward" on page 153** command displays a list of DB2 databases available to rollforward.

Activation

Activation is the last step of the three stage recovery process. This function brings DB2 enabled Notes databases online for use by the Domino server. DB2 enabled Notes databases that are restored from a DB2 table space backup image can be activated after first rolling the alternate DB2 database forward to the wanted point-in-time. The DB2 enabled Notes database can be restored to a time later than the backup time by applying necessary transaction log files by specifying the `/applylogs` parameter during the rollforward operation. The logs are then applied to the alternate DB2 database or to the Domino DB2 database if it is an existing restore. The DB2 application automatically archives transaction log files when they become full. The active transaction log files must be archived before you start the rollforward operation to ensure that the latest transactions are available. The necessary logs that are archived are automatically restored during the rollforward operation. The DB2 enabled Notes databases are then copied into the Domino 8 DB2 database to their original filename location or to a new filename location.

DB2 enabled Notes databases that are restored from a full DB2 database backup image are activated in the same manner as described for activating DB2 enabled Notes databases that are restored from a DB2 table space backup image. However, DB2 enabled Notes databases on different table spaces can be rolled forward simultaneously if more than one table space is restored from the full backup image.

The **"Domdsmc query DB2pendingdbs" on page 151** command displays a list of restored DB2 enabled Notes databases that are available for activation.

Security

Concepts that are associated with security issues and Data Protection for IBM Domino are described.

Data Protection for IBM Domino must be registered to the Tivoli Storage Manager server and use the appropriate node name and password when you are connecting to the Tivoli Storage Manager server.

Data Protection for IBM Domino must run from the same system user ID the Domino server is running under.

The Tivoli Storage Manager API `enableclientencryptkey` option provides 128-bit transparent encryption of Domino databases during Data Protection for IBM Domino backup and restore processing. Transparent encryption is only available on Tivoli Storage Manager server Version 5.3 (or later). See "More configuration options" on page 33 for details.

Note: You can see whether an NSF backup is encrypted, by issuing the **`query DBBackup`** command or by using the web or Java GUI.

Performance

Many factors can affect the performance of Data Protection for IBM Domino. Performance can be improved by implementing some changes.

Many factors can affect the backup and restore performance of your Domino Server databases. Factors such as hardware configuration, network type, and capacity are beyond the control of Data Protection for IBM Domino. However, some parameters that are related to Data Protection for IBM Domino can be tuned for optimum performance.

Data Protection for IBM Domino uses multiple data buffers when it is transferring data between the Domino and Tivoli Storage Manager servers. The number and size of the buffers can be specified with the `/buffers` parameter. The number and size of buffers that are allocated by default can be configured through the `set` command or by selecting the `Preferences` item from the `Edit` menu on the Data Protection for IBM Domino GUI. The default number of buffers is 3 and the default buffer size is 1024 KB.

To reduce query processing time when querying the Tivoli Storage Manager server for databases to restore with the Data Protection for IBM Domino Windows GUI, specify a database name with letters and a wildcard character (*) in the `By Database Name` field. For example, specifying `a*` displays all databases that begin with the letter `a` regardless of the folder name. Specifying `folder\a*` selects all databases that begin with `a` in the specified folder and its sub folders. Make sure to click `Update` after you enter the query.

To improve throughput for backup and restore operations, run multiple sessions in parallel. This is most effective when work is partitioned by physical volume. For example, one Data Protection for IBM Domino session backs up all databases on one physical volume while a second Data Protection for IBM Domino session backs up all databases on another volume.

To improve throughput for backup operations, run multiple sessions in parallel.

On Windows systems, there are two ways to accomplish this.

- When the databases are cleanly partitioned by a physical volume, you can start one Data Protection for IBM Domino instance to back up the databases on one physical volume. Then, start a second Data Protection for IBM Domino instance to back up the databases on another volume.
- If the databases are not cleanly partitioned, you can start one Data Protection for IBM Domino instance to back up all databases and use the sessions parameter to create multiple independent threads and sessions with the Tivoli Storage Manager server. This is equivalent to starting multiple independent Data Protection for IBM Domino instances. The difference is that independent threads are used instead of independent instances. When you are using independent threads, you do not have explicit control of which databases are backed by the individual threads.

You can also specify `tcpnodelay yes` in the Data Protection for IBM Domino options file (dsm.opt) to improve backup and restore performance. Instead of buffering the data, this option sends the data as successive small packets across the network without delay.

The statistics option

The statistics option logs performance information about an individual database at the backup or restore level. Data Protection for IBM Domino processing is operated under two threads: a producer process (which reads the data) and a consumer process (which sends the data). During a backup, the producer reads the database and the consumer sends this data to the Tivoli Storage Manager server. During a restore, the producer receives the data from the Tivoli Storage Manager server and the consumer writes the restored database. The statistics option logs this information to help tuning Data Protection for IBM Domino optimal performance.

Example of the statistics option

In this example output, the consumer send rate is greater than the producer file read rate. Because the consumer completes sending the data before the producer completes filling the next buffer, it waits an average of 25 milliseconds for each read buffer that is filled by the producer. The best method for improving throughput would be to modify the input/output subsystem. If the send/receive rate was lower than the read/write rate, the best method for improving throughput would be to modify the TCP/IP subsystem. If both the producer and the consumer have significant average wait times and the send/receive and read/write rates are similar, then the best method for improving throughput would be to modify the processor. The standard Long Wait value is 0. A Long Wait value other than 0 is most likely caused by tape mounts being loaded during processing. As a result, the consumer send/receive time is artificially increased and not representative of the standard data transfer time.

```
================================================================================
Request : SELECTIVE
Database Input List : Sample.db1.nsf
Number of Buffers : 2
Buffer Size : 1024
Logged Databases Only? : No
Wait for Tape Mounts? : No
Process Subdirectories? : No
TSM Options File : c:\Program Files\Tivoli\TSM\domino\dsm.opt
TSM Nodename Override :
--------------------------------------------------------------------------------
Performance statistics for database Sample.db1.nsf
Section Total Time Wait Time Average Time Long Waits
(msec) (msec) (msec)
--------------------------------------------------------------------------
Producer 231 1 0 0
Consumer 393 75 25 0
--------------------------------------------------------------------------
Sub Section Total Time Bytes Transferred Transfer Rate
(msec) (bytes) (Kb/sec)
--------------------------------------------------------------------------
ReadWrite 160 458752 2867
SendRecv 70 458752 6553
Domino Server 1 0 0
--------------------------------------------------------------------------
Total Elapsed Time Total Bytes Transferred Rate
(msec) (bytes) (Kb/sec)
--------------------------------------------------------------------------
1493 458752 307
Total Domino databases inspected: 1
Total Domino databases backed up: 1
Total Domino databases excluded: 0
Total Domino databases deduplicated: 0
```

```
Throughput rate: 300.07 Kb/Sec
Total byes inspected: 458,752
Total bytes transferred: 458,752
Total LanFree bytes transferred: 0
Total bytes before deduplication: 0
Total bytes after deduplication: 0
Data compressed by: 0.00%
Deduplication reduction: 0.00%
Total data reduction ratio: 0.00%
Elapsed processing time: 1.49 Secs
```

You can find more information about the `statistics` option in the description of "Domdsmc set" on page 115.

The `sessions` option

The `sessions` option allows a specified number of TCP/IP sessions to be made available for communication with the Tivoli Storage Manager server when you are backing up Domino NSF databases. Since more than one TCP/IP session is made available for backup processing, improvements in performance are possible. For example, the `sessions` option must be specified when simultaneously backing up NSF data to multiple tape drives. You can specify from 1 to 64 sessions. The default value is 1. However, be aware that since network and hardware capabilities of the production environment can also impact the overall performance enhancements of the `sessions` option, environment conditions must be considered when you are using the `sessions` option.

In addition, be aware that each session requests a mount point from the Tivoli Storage Manager server when backup processing begins. If a mount point is in use, then it is not released for use by a new session until the backup is complete. Because of this behavior, it is possible that a session (waiting for an available mount point) might timeout, causing the backup attempt to fail. This situation can occur when the number of specified backup sessions exceeds the number of available mount points. To avoid this situation, make sure that the number of available mount points (from the Tivoli Storage Manager server) is equal to the number of sessions that are specified with the `sessions` option. It is the responsibility of the user to determine the number of available mount points as Data Protection for IBM Domino does not determine this information. Also, the Tivoli Storage Manager Administrator must set the `maxnummp` option to specify the maximum number of mount points to use (for the Domino Server ID) on the Tivoli Storage Manager server.

The `sessions` option is available with the following commands:

- "**`Domdsmc incremental`** " on page 73
- "Domdsmc selective" on page 110
- "Domdsmc fullselective" on page 129
- "Domdsmc set" on page 115

The `DOMTXNBYTELIMIT` option

The `DOMTXNBYTELIMIT=number` option specifies the number of bytes sent between Data Protection for IBM Domino and the Tivoli Storage Manager server in a single transaction. The default value is 0, which indicates no limit, and the maximum value is `2097152`. This number is multiplied by 1024 to calculate the limit in bytes.

This parameter is useful when you are backing up NSF databases to tape storage for these reasons:

- Processing for each transaction causes the tape to stop and start. Considerable time can be lost during the stop and start when you are using high speed tapes. This is true in a LAN-free environment.
- Errors that occur during backup processing are automatically tried again when `domtxnbytelimit` is set.

When a failure occurs during a backup, all of the backups in the transaction are tried again, not just the NSF database in error. Each backup is tried again in a separate transaction. After all backups are tried again, the `domtxnbytelimit` parameter is used to control the number of bytes per transaction.

The `DOMTXNGROUPmax` option

The `DOMTXNGROUPmax=number` option specifies the number of individual objects sent to the Tivoli Storage Manager server in a single transaction.

Two objects are sent to the Tivoli Storage Manager server for each database backup so the default value of this option is 2. The maximum value is 65000.

The `DOMTXNGROUPmax` option can be overridden by the Tivoli Storage Manager server `TXNGRPMAX` option. However, when `domtxngroupmax` is set, the minimum of the two values is used.

This parameter is useful when you are backing up NSF databases to tape storage for these reasons:

- Processing for each transaction causes the tape to stop and start. Considerable time can be lost during the stop and start when you are using high speed tapes. This is true in a LAN-free environment.
- Errors that occur during backup processing are automatically tried again when `domtxngroupmax` is set.

When a failure occurs during a backup, all of the backups in the transaction are tried again, not just the NSF database in error. Each backup is tried again in a separate transaction. After all backups are tried again, the `domtxngroupmax` parameter is used to control the number of individual objects per transaction. Consider running the `domtxngroupmax` parameter when you are backing up small NSF databases.

Automated failover for data recovery

When there is an outage on the Tivoli Storage Manager server, Data Protection for IBM Domino can fail over to a secondary server for data recovery operations.

The Tivoli Storage Manager server that Data Protection for IBM Domino connects to for backup operations is called the *primary server*. When the primary server, Data Protection for IBM Domino node, and the Tivoli Storage Manager backup-archive client node are set up for node replication on the primary server, the nodes can be replicated to another Tivoli Storage Manager server, called the *secondary server*.

During normal operations, connection information for the secondary server is automatically sent to Data Protection for IBM Domino from the primary server. The secondary server information is saved to the client options files on the Data Protection for IBM Domino node and the backup-archive client node. No manual intervention is required by you to add the information for the secondary server.

Each time the backup-archive client logs on to the server for backup services, it attempts to contact the primary server. If the primary server is unavailable, the backup-archive client automatically fails over to the secondary server. In failover mode, you can restore data that was replicated to the secondary server. When the primary server is online again, the backup-archive client automatically fails back to the primary server the next time the backup-archive client connects to the server.

You can confirm that Data Protection for IBM Domino has failed over by looking for entries about the secondary server in the `dsierror.log` file.

Requirements: To ensure that automated client failover can occur, Data Protection for IBM Domino must meet the following requirements:

- Data Protection for IBM Domino must be at the V7.1 level.
- The primary server, secondary server, and backup-archive client must be at the V7.1 level.
- The primary and secondary servers must be set up for node replication.
- The Data Protection and backup-archive client nodes must be configured for replication with the `replstate=enabled` option in each node definition on the server.
- Before the connection information for the secondary server can be sent to Tivoli Storage FlashCopy Manager, the following processes must occur:
  - You must back up data at least one time to the primary server.
  - The Data Protection for IBM Domino node on the primary server must be replicated at least one time to the secondary server.

Restriction: The following restrictions apply to Data Protection for IBM Domino during failover:

- Any operation that requires data to be stored on the Tivoli Storage Manager server, such as backup operations, are not available. You can use only data recovery functions, such as restore or query operations.
- Schedules are not replicated to the secondary server. Therefore, schedules are not run while the primary server is unavailable.
- If the primary server goes down before or during node replication, the most recent backup data is not successfully replicated to the secondary server. The replication status of the file space is not current. If you attempt to restore data in failover mode and the replication status is not current, the recovered data might not be usable. You must wait until the primary server comes back online before you can restore the data.
- For more information about the failover capabilities of Tivoli Storage Manager components, see http://www.ibm.com/support/docview.wss?uid=swg21649484.

For more information about automated client failover with the Tivoli Storage Manager backup-archive client, see *Automated client failover configuration and use* in the Tivoli Storage Manager information center.

Chapter 2. Installation of Data Protection for IBM Domino on a Windows system

Prerequisites and procedures for installing Data Protection for IBM Domino are provided.

Installation prerequisites

Before you install Data Protection for IBM Domino in a Windows environment, ensure that your system meets the hardware, software, and operating system requirements.

The minimum hardware and software requirements for the Data Protection for IBM Domino release are available in the hardware and software requirements technote for each particular release. For current requirements, review the *Hardware and Software Requirements* technote for your version of Data Protection for IBM Domino. This technote is available in the *TSM for Mail - All Requirements Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21219345. From the page, follow the link to the requirements technote for your specific release or update level.

Minimum hardware requirements

Your system must meet the minimum hardware requirements for operating Data Protection for IBM Domino in a Windows environment.

The minimum hardware requirements for the Data Protection for IBM Domino release are available in the hardware and software requirements technote for each particular release. For current requirements, review the *Hardware and Software Requirements* technote for your version of Data Protection for IBM Domino. This technote is available in the *TSM for Mail - All Requirements Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21219345. From the page, follow the link to the requirements technote for your specific release or update level.

Minimum software and operating system requirements

Your system must meet the minimum software requirements for operating Data Protection for IBM Domino in a Windows environment.

The minimum software and operating system requirements for the Data Protection for IBM Domino release are available in the hardware and software requirements technote for each particular release. For current requirements, review the *Hardware and Software Requirements* technote for your version of Data Protection for IBM Domino. This technote is available in the *TSM for Mail - All Requirements Documents* website at http://www.ibm.com/support/docview.wss?uid=swg21219345. From the page, follow the link to the requirements technote for your specific release or update level.

Virtualization support

Information about the virtualization environments that can be used with Data Protection for IBM Domino is available in the *IBM Tivoli Storage Manager guest support for virtual machines and virtualization* website at: http://www.ibm.com/

support/docview.wss?uid=swg21239546.

Installing in a Windows environment

These instructions guide you through the installation of Data Protection for IBM Domino on a Windows system.

Note: The Tivoli Storage Manager Backup-Archive client must be installed before you install Data Protection for IBM Domino. Installing in this order ensures that the Data Protection for IBM Domino plug-in feature that enables the Web Client GUI is in place.Data Protection for IBM Domino must be installed from an account with administrator privileges to the local system.

1. With the product DVD inserted, run one of the following commands, where `x` is the DVD drive letter:
   - `32-bit Domino server: x:\domino\windows\x32\client\setup`
   - `64-bit Domino server: x:\domino\windows\x64\client\setup`

   Select **OK**.
2. To install Data Protection for IBM Domino from a downloaded installation image run the executable installation file.
3. Select the language for the installation procedure. Select one of the following languages:

   Chinese (Simplified)
   Chinese (Traditional)
   Czech
   English (United States)
   French (France)
   German (Germany)
   Hungarian
   Italian (Italy)
   Japanese
   Korean
   Polish
   Portuguese (Brazil)
   Russian
   Spanish (Traditional Sort)
4. Click **OK** to start the installation program.
5. Read the License Agreement. You must accept the License Agreement in order for Data Protection for IBM Domino to install successfully.
6. Click **Next** and choose the folder where you would like to install Data Protection for IBM Domino.
7. Select one of the following setup types:

   Typical

   This choice installs the following program features:
   - Data Protection for IBM Domino base product.
   - Online help for the Data Protection for IBM Domino GUI.
   - License.

   You can also change the folder where you would like to install Data Protection for IBM Domino.

Custom
Use custom setup to select which program features from the `Typical` installation you would like to install. You can also change the folder where you would like to install Data Protection for IBM Domino.

To install the Data Protection for IBM Domino plug-in feature that enables the web client GUI, the `Custom` setup type must be selected.

8. Click **Finish** to complete the installation.
9. If Data Protection for IBM Domino Version 7.x is installed on your system and you want to reinstall it, select one of these setup options:

Modify
You can select which program features to install.

Repair
Repair any missing or corrupted files, shortcuts, and registry entries.

Remove
Remove Data Protection for IBM Domino from your system.

Installing language packs

Information about language packs and how to install them is provided.

If you want to view the Data Protection for IBM Domino GUI, command-line output, and messages in another language, install a different Language Pack from the product DVD.

Note: When you install the language pack, Data Protection for IBM Domino operates in the language that you installed. However, in rare cases it is possible that databases can contain characters that are not supported in that language and therefore are not processed.

The language packs are executable files that are in the following directories on the product DVD, where `x` is the DVD drive:

- 32-bit: `w:\domino\windows\x32\languages\`*xxx*
- 64-bit: `w:\domino\windows\x64\languages\`*xxx*

The *xxx* directory represents the three-letter country code that is associated with that language. After you install the Language Pack, activate the language by updating the Data Protection for IBM Domino configuration file, `domdsm.cfg` with either of these methods:

- Use the **`set`** command with the **`language`** parameter to specify the language. For example:

```
domdsmc set LANGuage=<language>
LANGuage=
   ENU  (English, United States)
      PTB  (Brazilian Portuguese)
      CHS  (Chinese, Simplified)
      CHT  (Chinese, Traditional)
      FRA  (Standard French)
      DEU  (Standard German)
      ITA  (Standard Italian)
      JPN  (Japanese)
      KOR  (Korean)
      ESP  (Standard Spanish)
      CSY  (Czech)
      HUN  (Hungarian)
      PLK  (Polish)
      RUS  (Russian)
```

See the description of the **language** parameter for a list of available languages and their three-letter country codes.

- Use the Configuration Editor in the Data Protection for IBM Domino GUI by selecting **Edit–>Preferences–>Regional–>Language**. The GUI Configuration Editor shows the languages in their long form. For example:

  ```
  English (United States)
  ```

Configuring Data Protection for IBM Domino in quick mode (Windows)

Instructions on how to configure Data Protection for IBM Domino on a Windows system in quick mode are provided.

Before you begin

This procedure uses default settings and requires minimal configuration tasks. It minimizes setup time so that you can proceed quickly to backing up your Domino databases. Detailed instructions on how to customize Data Protection for IBM Domino for your environment and processing needs are available in the configuration section.

Consider the following extra quick installation steps that are required for DB2 enabled Notes databases:

- Make sure the Tivoli Storage Manager API settings are defined for the DB2 environment.
- Set the user access to the DB2 environment by issuing this command:

  ```
  domdsmc set db2user=<DB2 user name>
  ```

Ensure that you note the installation path for your operating system before you begin the configuration. The following table shows the installation directories for Linux and AIX®:

Table 1. The installation paths for Linux and AIX:

Platform	Installation path
Linux on System z®	/opt/tivoli/tsm/domino/bin64
Linux_x86_64	/opt/tivoli/tsm/domino/bin
AIX	/opt/tivoli/tsm/client/domino/bin64

Procedure

1. Install Data Protection for IBM Domino. Detailed installation instructions are available in the installation section.
2. Change to the C:\Program Files\Tivoli\TSM\domino directory. Copy the dsm.smp file to dsm.opt.. Edit the dsm.opt file to include these options: TCPServeraddress x.x.x.x Replace x.x.x.x NODename YOURNODENAME with the IP address or host name of the Tivoli Storage Manager server to which Data Protection for IBM Domino backs up data. Replace YOURNODENAME with the node name by which Data Protection for IBM Domino is known to the Tivoli Storage Manager server. For more information about these options and the dsm.opt file, see “Options and preferences” on page 31.
3. Register the node (specified in step 2) to the Tivoli Storage Manager server with the following command: REG NODE <dpdomnode> password, where <dpdomnode> is the name of the system where Data Protection for IBM Domino

is installed, and `password` is the password for this node. All other options use default settings. If your Tivoli Storage Manager policy settings for Data Protection for IBM Domino backups are different from the default settings, ensure that you register the node to the DOMAIN containing your Data Protection for IBM Domino information.

4. If a password file for the Domino user ID does not exist, create one with the following command: `domdsmc query adsm -adsmpwd=password`. You need a password file to access the Domino Server and partitions with the Web client GUI.
5. Verify that you can communicate with the Domino Server by running the commands:

```
su -notes
$ . ./notes.profile
$ domdsmc query domino
```

6. Verify that you can communicate with the Tivoli Storage Manager server by running the **domdsmc query adsm** command as a Notes user.

Results

Data Protection for IBM Domino is now ready for backup and restore processing. For example, to run an incremental backup of your databases, enter the following command: `domdsmc incr *`

Silent installation of Data Protection for IBM Domino

You can install Data Protection for IBM Domino using a silent installation.

Silent installation runs on its own without any intervention, so that you do not need to monitor the installation or provide input. This method is especially useful when Data Protection for IBM Domino must be installed on a number of different computers with identical hardware. For example, a company might have 25 Domino Servers spread out across 25 different sites. To ensure a consistent configuration and to avoid having 25 different people enter Data Protection for IBM Domino parameters, an administrator can choose to produce an unattended installation. The installation can be made available to different sites by cutting and sending out DVDs or by placing the unattended installation package on a file server.

You can install silently with one of the following methods:

Setup Program
Use the **setup** command with the silent installation options.

Microsoft Installer (MSI)
Use **msiexec.exe** to install the MSI package.

The following options can be used with both silent installation methods:

Table 2. Silent installation options

Option	Description
/i	Specifies the program is to install the product.
/l*v	Specifies verbose logging.
/qn	Runs the installation without running the external user interface sequence.
/s	Specifies silent mode.

Table 2. Silent installation options (continued)

Option	Description
/v	Specifies the Setup Program to pass the parameter string to the call it makes to the MSI executable file, `msiexec.exe`. Note the following syntax requirements when you are using the **/v** option: • A backslash (\) must be placed in front of any quotation marks (" ") that are within existing quotation marks. • Do not include a space between the **/v** command-line option and its arguments. • Multiple parameters that are entered with the **/v** command-line option must be separated with a space. • You can create a log file by specifying the directory and file name at the end of the command. The directory must exist when the silent installation is run.
/x	Specifies the program is to uninstall the product.
addlocal	Specifies features to install.
allusers	Specifies which users can use the installation package.
installdir	Specifies the directory where Data Protection for IBM Domino is to be installed.
reboot	Specifies whether to prompt the user to restart the system after silent installation. • `Force` prompts the user to restart after silent installation. • `Suppress` prompts a restart after silent installation. • `ReallySuppress` suppresses all restarts, and prompts the user to restart after silent installation.
rebootyesno	Specifies whether to restart the system after silent installation. Specify **Yes** to restart the system after silent installation. Specify **No** not to restart the system after silent installation.
transforms	Specifies language to install.

The following features are used in this procedure and are case-sensitive:

Table 3. Silent installation features (base client only)

Feature	Description
Client	Data Protection for IBM Domino code
License_Paid	License file (Used when PAID versions of Data Protection for IBM Domino are installed)
Plug-in	Data Protection for IBM Domino plug-in (enables the web client GUI)

Table 4. Silent installation features (Language Packages only)

Feature	Description
LanguageFiles	Language-specific files

The following transforms are used in this procedure:

Table 5. Silent installation transforms

Transform	Language
1028.mst	CHT Chinese (Traditional)
1029.mst	CSY Czech
1031.mst	DEU German
1033.mst	ENG English
1034.mst	ESP Spanish
1036.mst	FRA French
1038.mst	HUN Hungarian
1040.mst	ITA Italian
1041.mst	JPN Japanese
1042.mst	KOR Korean
1045.mst	PLK Polish
1046.mst	PTB Portuguese
1049.mst	RUS Russian
2052.mst	CHS Chinese (Simplified)

Installing with the Setup executable file

Instructions on how to install Data Protection for IBM Domino using the **setup.exe**.

Before you begin

Data Protection for IBM Domino must be installed from an account that is a member of the local Administrators group for the processor on which the Domino server is running.

About this task

Note: During the Data Protection for IBM Domino silent installation, you must substitute the appropriate feature when you are installing a language other than English. See Table 4 on page 24.

Run the following command to silently install Data Protection for IBM Domino to the default installation directory:

```
setup /s /v/qn
```

This example silently installs Data Protection for IBM Domino to a directory other than the default installation directory and includes custom features:

```
setup /s /v"INSTALLDIR=\"c:\program files\tivoli\tsm\"
ADDLOCAL=\"Client,License_Paid\"
TRANSFORM=1033.mst /qn /l*v \"c:\temp\log.txt\"2
```

Note:

1. You must place a backslash \ before each quotation mark that is within an outer set of quotation marks ".
2. You must place quotation marks " around the following information:

- A directory path that contains spaces.
- An argument that specifies multiple features. Although quotation marks are needed around the complete argument, you must still place a backslash before each internal quotation mark.

3. All features that are listed in a custom installation must be listed after the `addlocal` option.

Creating batch files

Create a `setup.bat` file to install silently. An example is outlined.

A batch file can be created to begin the installation in silent mode, with the parameters that you choose. The following example shows a sample script `c:\setup.bat` to demonstrate unattended installation:

```
@echo off
 rem ==================================
 rem sample silent install script
 rem
 setup /s /v"INSTALLDIR=\"X:\Install Path\" /qn"
 rem ==================================
 rem code could be added after the
 rem installation completes to
 rem customize the dsm.opt files
 rem if desired
 rem ==================================
```

Installing Data Protection for IBM Domino with the MSI executable file

Instructions on how to install Data Protection for IBM Domino with the `msiexec.exe` file are presented.

Before you begin

Data Protection for IBM Domino must be installed from an account that is a member of the local Administrators group for the system on which the Domino server is running. Data Protection for IBM Domino installs the Microsoft Visual C++ 2012 Redistributable Package as a setup prerequisite. If you are installing with `msiexec.exe`, you must separately install the Microsoft Visual C++ 2012 Redistributable Package. The files are included in the installable packages:

Windows 32 bit, install `\ISSetupPrerequisites\Microsoft Visual C++ 2012 Service Update 1 Runtime Libraries (x86)\vcredist_x86.exe`

Windows 64 bit, install `\ISSetupPrerequisites\Microsoft Visual C++ 2012 Service Update 1 Runtime Libraries (x64)\vcredist_x64.exe`

About this task

Note: In the Data Protection for IBM Domino silent installation, you must substitute the appropriate `.msi` package file name and Language Package feature when you are installing a language other than English. See Table 4 on page 24.

This example silently installs Data Protection for IBM Domino to a directory other than the default installation directory and includes custom features:

```
msiexec /i <path to msi file>\"IBM Tivoli Storage Manager for Mail
 - Lotus Domino.msi
"RebootYesNo="No" Reboot="Suppress" ALLUSERS=1
```

```
INSTALLDIR="c:\program files\tivoli\tsm"
ADDLOCAL="Client,License_Paid"
TRANSFORM=1033.mst /norestart /qn /l*v "c:\temp\log.txt"
```

Note:

- You must place quotation marks " around the following items:
  - A directory path that contains spaces.
  - An argument that specifies multiple features. Although quotation marks are needed around the complete argument, you must still place a backslash before each internal quotation mark.
- All features that are listed in a custom installation must be specified after the `addlocal` option.

What to do when an installation fails

When an installation fails you need to gather some information including a log for the installation failure. Information about what you should do is listed.

If an installation fails, record the symptoms and environment information for the failing installation and contact customer support with that information. Collect the information that is listed:

- Operating system level.
- Service pack name.
- Hardware description.
- Installation package DVD or electronic download name and level.
- Any Windows event log that is relevant to the failed installation.
- Other Windows services active at the time of the installation, for example antivirus software.

Before you contact the support desk, check for the following items:

- You must be logged on to the local console not through a terminal server.
- You must be logged on as a local administrator, not a domain administrator. Cross-domain installations are not supported by Tivoli.

Gather a detailed log of the failing installation, and save it into a file called `setup.log`. Run the setup program as follows:

```
setup /v"/l*v setup.log"
```

Creating the package on a DVD or a file server

The package can be made available on a DVD or as a stored package in a shared directory on a file server. The package can contain the Data Protection for IBM Domino code distribution files and a batch file for silent installation.

Silent installation package for Data Protection for IBM Domino

Information about silently installing the Data Protection for IBM Domino package through a staging directory is listed. Example commands are provided.

First, you must choose a location for the package to be stored. If you are using a DVD, it is convenient to use a staging directory. If you are placing the package on a file server, you can use a staging directory or build the package directly on the file server. The following example uses `c:\tdpdpkg` as a staging directory. You need a minimum of 14 MB of free space in the staging directory. The following commands can be run to create the package.

Command	Description
`mkdir c:\tdpdpkg`	– Create a staging directory for the silent installation package
`cd /d c:\tdpdpkg`	– Go to the staging directory
`xcopy g:\*.* . /s`	– Copy the Data Protection for IBM Domino DVD distribution files to the staging directory
`copy c:\setup.bat`	– Replace the existing `setup.bat` with the one created in the previous step

Test the silent installation before you go any further. When the testing is complete, the package can be placed on a DVD or it can be made available from a shared directory for distribution.

Setup error messages

The `setup.exe` executable file might produce error messages if it cannot start properly.

In most cases, you will encounter these messages when a severe error occurs. Your users will rarely see these messages. Every error message has a unique number. There is no way to suppress error messages in your script.

If you encounter an error, you can go to the Knowledge Base, InstallShield support website at URL: `http://support.installshield.com/default.asp`, and search for information about the specific error.

Chapter 3. Configuring Data Protection for IBM Domino

How to configure Data Protection for IBM Domino to protect Domino databases is described.

Communication

The communication concepts between Data Protection for IBM Domino and the Tivoli Storage Manager server are described.

Data Protection for IBM Domino communicates with several product APIs to complete various functions. The Tivoli Storage Manager API is accessed in order for Data Protection for IBM Domino to communicate with the Tivoli Storage Manager server. The Domino API is accessed for communicating with the Domino server during database operations, and DB2 enabled Notes data is accessed by communicating with the DB2 Recovery API. The communication protocols and option parameters are specified in the `dsm.opt` options file. See *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide* for more information about specifying the communication method.

You can end a Data Protection for IBM Domino client session by running the **`Cancel Session`** command from a Tivoli Storage Manager admin client. Do not press `Ctrl-C` as it can lead to unexpected results.

Registering with the Tivoli Storage Manager server

The registration concepts between Data Protection for IBM Domino and the Tivoli Storage Manager server are presented.

Before you back up to and run a recover from a Tivoli Storage Manager server, you must have a Tivoli Storage Manager registered node name and password. The process of setting up a node name and password is called `registration`. When registered, you can begin to back up and restore Domino databases and transaction logs with Data Protection for IBM Domino.

If your system has a node name that is assigned to the Tivoli Storage Manager backup-archive client, you must have a different node name for Data Protection for IBM Domino.

For information about the registration process, see *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

Create policy

Information about Tivoli Storage Manager policy settings is provided.

Although Data Protection for IBM Domino operates in ways similar to other Tivoli Storage Manager clients, it is unlike regular Tivoli Storage Manager clients in that it does not always store complete replacements for objects on the Tivoli Storage Manager server. When a database file is backed up, it is a complete backup of the entire database and becomes a new backup version of that database. When archival logging is being used on the Domino server, then each archived transaction log file contains changes to one or more logged databases over time.

Each of these transaction log files has a unique name so there are never multiple versions of the same transaction log file. Because of this difference, Data Protection for IBM Domino requires special Tivoli Storage Manager policy settings.

How Data Protection for IBM Domino has an impact on policy (NSF databases)

Data Protection for IBM Domino has an impact on Tivoli Storage Manager backup policy for NSF databases in these ways:

- Regular use of the **Domdsmc Inactivatelogs** command deactivates the archived transaction log files when all NSF databases that would require that file for a complete recovery are inactive. Ensure to set the retention period for inactive transaction log files to be equal to, or greater than the database backup objects. Files are available when any inactive database file that might need them is available. A point in time recovery of an inactive database backup version can be accomplished. The same management class for the transaction log files for the database files is used.
- It is possible to have multiple versions of the same transaction log file under certain circumstances. Data Protection for IBM Domino and the Domino Server provide the capability to archive the currently filling transaction log. Thus, the same transaction log file can be backed up multiple times (while it is filling and again when it is full). If the Domino server is stopped abnormally after transaction log files are archived but not yet reused by the Domino server, those transaction log files can be archived. These files might be unchanged. As a result, version limit parameters for the management class that is used for transaction log files must be set to ensure that extra versions of a transaction log file are purged from the backup storage pools.
- To optimize the recovery process, use collocation for the file space that contains the transaction log files if they are stored on sequential media on the Tivoli Storage Manager server. The transaction log files are stored in a separate file space from the database files on the Tivoli Storage Manager server.

How Data Protection for IBM Domino has an impact on policy (DB2 enabled Notes databases)

Data Protection for IBM Domino has an impact on Tivoli Storage Manager backup policy for DB2 enabled Notes databases in these ways:

- Use the **Domdsmc DB2INActivateobjs** command regularly to deactivate the archived transaction log files. It backs up objects when all DB2 enabled Notes databases that would require those files for a complete recovery are inactive. Therefore, be sure to set the retention period for inactive transaction log files to be equal to or greater than the database backup objects. The retention period ensures that the files are available when any inactive database file that might need them is available. A point in time recovery of an inactive database backup can be accomplished with the same management class for the transaction log files. The same is true when the same backup objects for the database files are used.
- To optimize the recovery process, use collocation for the file space that contains the transaction log files if they are stored on sequential media on the Tivoli Storage Manager server. The transaction log files are stored in a separate file space from the database files on the Tivoli Storage Manager server.

Policy settings

Use default values for the following Backup Copy Group parameters because they are not applicable to Data Protection for IBM Domino:

- `frequency`
- `mode`
- `serialization`

Define a separate policy domain where the default management class has the settings that are required for your Domino backup data. Register all Domino Tivoli Storage Manager nodes to that domain. When you are defining a new management class within an existing policy domain that is not the default management class for that domain, you must add an `include` statement. The `include` statement in the Data Protection for IBM Domino options file binds all objects to that management class. For example: `include * mcname`.

See your Tivoli Storage Manager administrator or the appropriate Tivoli Storage Manager Administrator Guide for your server, for more information about defining or updating policy domains, and copy groups.

Data Protection for IBM Domino stores all objects as backup objects on Tivoli Storage Manager so an Archive Copy Group is not required, although it can exist.

All database backup objects are complete file backups so normal version controls available through Tivoli Storage Manager server policies apply. Set the `verdeleted`, `verexists`, `retonly`, and `retextra` parameters of the Backup Copy Group according to your needs for the number of backup versions to be kept and the retention period of these backup versions.

Options and preferences

Data Protection for IBM Domino uses options files to store configuration information. The files that have options and preferences that must be set are described.

If you have a Tivoli Storage Manager backup-archive client, on the same system with Data Protection for IBM Domino, you must use different node names for the two clients.

domdsm.cfg

The `domdsm.cfg` preferences file contains options specific to Data Protection for IBM Domino. Use the `set` command to set values for these options. Do not use a text editor to edit this file. You can display the current values in `domdsm.cfg` by issuing the `query preferences` command. You can also use the Preferences dialog from the **Edit** menu in the GUI. If the preferences file is corrupted and contains invalid values, the default values for the preferences are used. See "Domdsmc set" on page 115 for parameters that are stored in this file.

dsm.opt

The options file, `dsm.opt`, identifies the Tivoli Storage Manager server to contact and the node name by which Data Protection for IBM Domino is known to the Tivoli Storage Manager server. This file also contains some options that are related to back up and restore processing. Use the sample options file, `dsm.opt.smp`, to

create the dsm.opt file. The dsm.opt.smp file is in the Data Protection for IBM Domino installation directory.

Required options

Required options for Data Protection for IBM Domino are described.

After Data Protection for IBM Domino is registered to a Tivoli Storage Manager server, the following Tivoli Storage Manager options must be specified in the options file in the Data Protection for IBM Domino installation directory to communicate with the Tivoli Storage Manager server:

- nodename
- password
- tcpserveraddress

The default options file name is dsm.opt. The Tivoli Storage Manager administrator can provide you with the node name, password, and the communications method for connecting to the Tivoli Storage Manager server.

Preferred options

Customize Data Protection for IBM Domino by specifying options when configuring Data Protection for IBM Domino. Set options such as passwordaccess in the system options file in the Tivoli Storage Manager API installation directory.

passwordaccess

When passwordaccess is set to prompt, you are prompted for your password. When passwordaccess is set to generate, the Tivoli Storage Manager API saves the current password (encrypted) in the Windows registry. A new password is automatically generated when the current one expires. This method of password management is useful when you are running scheduled, unattended backups. This method also ensures that a backup never fails because of an expired password.

Specify this option in the client options file located in the Data Protection for IBM Domino installation directory.

compression

Specifying compression yes causes Data Protection for IBM Domino to compress data before it is sent to the Tivoli Storage Manager server. If you enable compression, it affects performance in two ways:

- Processor usage has a higher value on the server on which Data Protection for IBM Domino is running
- Network bandwidth usage has a lower value because fewer bytes are transmitted.

If the computer that is running Data Protection for IBM Domino has a processor overload, specify compression no because more usage can impact other applications such as the Domino server. It is better to specify compression yes when any of the following conditions exist:

- The network adapter has a data overload.
- Communications between Data Protection for IBM Domino and the Tivoli Storage Manager server are over a low-bandwidth connection.
- There is heavy network traffic.

Specifying `compression yes` results in reduced storage usage on the Tivoli Storage Manager server.

The Tivoli Storage Manager administrator can restrict use of the compression option by specifying, on the Tivoli Storage Manager server side, that a particular node:

- Always uses compression.
- Never uses compression.
- Leaves the decision up to the node to decide.

The value of the compression option for Data Protection for IBM Domino is recognized if the Tivoli Storage Manager administrator leaves the compression decision to the node. The default is to leave the decision to the node.

Specify this option in the client options file in the Data Protection for IBM Domino installation directory.

Exclude databases that increase in size during compression `compression yes` by using the client option, `exclude.compression`. See *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide* for information about this option. See "Include and exclude processing" on page 171 for examples of `include/exclude` statements.

The compression status of an NSF backup can be seen by issuing a query command or in the GUI.

tapeprompt

The `tapeprompt` option controls whether Data Protection for IBM Domino waits for tape mount requests to be resolved on the Tivoli Storage Manager server or terminates the current operation when the Tivoli Storage Manager server indicates that it is waiting for a tape mount. During a backup operation, Tivoli Storage Manager might issue a prompt to place a tape volume in a drive. Also, during a restore operation, the data you want to recover might be on a tape that is not currently mounted by the server. In either case, a Tivoli Storage Manager operator or autochanger must take time to mount the particular tape. During that time, Data Protection for IBM Domino continues to show activity and wait for a Tivoli Storage Manager server operation to complete. If this option is selected `tapepromptyes`, Data Protection for IBM Domino waits for a tape to be mounted before it continues. If this option is not selected `tapepromptno`, the operation ends.

Specify this option in the client options file located in the Data Protection for IBM Domino installation directory.

More configuration options

Customize your Data Protection for IBM Domino environment by configuring extra options for server communication, compression, encryption, and deduplication.

COMMRESTARTDURATION

Use `COMMRESTARTDURATION` to specify the total number of minutes that the server will attempt to restart a session after a communication failure. The range of values is 1 - 9999 and the default is `60`. This option must be set high on a network that is unreliable.

You can specify this option in the Data Protection for IBM Domino preferences file.

COMMRESTARTINTERVAL

Use COMMRESTARTINTERVAL to specify how many seconds the server will wait before it attempts to restart a session after a communication failure. The range of values is one through 9999 and the default is 15. The restart interval setting means that the network is not overloaded with restart requests. The COMMRESTARTINTERVAL time must always be less than or equal to the COMMRESTARTDURATION time.

You can specify this option in the Data Protection for IBM Domino preferences file.

deduplication

Use deduplication to specify whether data deduplication is used. The option can be set to deduplication yes or deduplication no depending on your requirements. You can specify this option in the Data Protection for IBM Domino preferences file.

domnode

Use domnode to use the web client GUI to back up and restore Domino server data. It provides the web client GUI with the Tivoli Storage Manager node name and respective Domino server to access for processing. It also provides important configuration information for the specified node. Specify the full path and name of the Data Protection for IBM Domino preferences file. For example: domnode c:\program files\notes\domdsm.cfg

Consider the following items when you are specifying the domnode option:

- It must be specified to access the web client GUI.
- It can be specified multiple times for as many Domino servers or Domino Partitioned Servers as are available.
- It can be used in short form domno and is not case-sensitive.
- Specify this option in the client **dsm.opt** file that is used by the backup-archive Client Acceptor Daemon (CAD). The options **dsm.opt** file is specified during setup of your web client by the Tivoli Storage Manager web client Configuration wizard.

domnode example

In this example, the backup-archive client can access Domino server A, Domino server B, and Domino server C:

- Contents of the **dsm.opt** file that is used by the backup-archive client CAD:

```
DOMNODE c:\Program Files\notesA\serverA.cfg
DOMNODE c:\Program Files\notesB\serverB.cfg
DOMNODE c:\Program Files\notesC\serverC.cfg
```

- Contents of dsm1.opt: NODENAME domservA TCPSERVERADDRESS tsmserv1.company.xyz.com PASSWORDACCESS generate
- Contents of dsm2.opt: NODENAME domservB TCPSERVERADDRESS tsmserv2.company.xyz.com PASSWORDACCESS generate
- Contents of dsm3.opt: NODENAME domservC TCPSERVERADDRESS tsmserv3.company.xyz.com PASSWORDACCESS generate
- Contents of serverA.cfg: NotesInipath C:\Lotus\DominoServer65\ tdpdominoserver1

- Contents of `serverB.cfg`: `NotesInipath D:\Lotus\DominoServer65\tdpdominoserver2`
- Contents of `serverC.cfg`: `NotesInipath E:\Lotus\DominoServer6\server`

enableclientencryptkey

When **enableclientencryptkey** is set to yes, Data Protection for IBM Domino provides 128-bit `Transparent encryption` of Domino databases during backup and restore processing. One random encryption key is generated per session and is stored on the Tivoli Storage Manager server with the object in the server database. Although Tivoli Storage Manager manages the key, a valid database must be available to restore an encrypted object. You can specify the databases that you want encrypted by adding an include statement with the **include.encrypt** option in the **dsm.opt** file, which is located in the Data Protection for IBM Domino installation directory. See "Include and exclude processing" on page 171 for an example of an `include.encrypt` statement. `Transparent encryption` is only available on Tivoli Storage Manager server Version 5.3 (or later). For more information, see *IBM Tivoli Storage Manager Using the Application Program Interface*.

The encryption status of a backup can be seen during the backup or restore process by using the **/DEtail** command.

asnodename

When `asnodename` is specified, Data Protection for IBM Domino backs up or restores databases on multiple clients under the single Tivoli Storage Manager node name that is specified by this option. Unlike the `nodename` option that requires you to enter the password for the node name you specify, the `asnodename` option requires that you enter the password for your client node to access data that you own. Specify this option in the Data Protection for IBM Domino options file **dsm.opt**, located by default in the Data Protection for IBM Domino installation directory. See *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide* for more information about the `asnodename` option.

Option and preference priority

Some options are specified in more than one options or preferences file. The sequence in which option values are prioritized is explained.

The source of options and preferences with the highest priority is seen in the following list:

1. Data Protection for IBM Domino command-line option has the highest precedence.
2. Data Protection for IBM Domino preferences file, `domdsm.cfg`.
3. Data Protection for IBM Domino options file, `dsm.opt`.

GUIs used to manage Domino databases

Requirements and procedures on how to access, start, and use the various GUIs to back up and restore your Domino databases and transaction log files.

The Tivoli Storage Manager web client GUI and the Java client GUI are available only when the Tivoli Storage Manager backup-archive client is installed. From the Tivoli Storage Manager web client GUI you can back up and restore Domino server data from a remote server through a web browser. It is useful for monitoring multiple servers.The Java client GUI is available locally on the desktop and has a similar interface to the web client GUI. Use Java client GUI to view servers that you have physical access to. The main difference is that the web client GUI uses a web browser and provides remote access.

For information about these tasks with the command-line interface, see "Command-line interface" on page 50.

Data Protection for IBM Domino GUI

Requirements and procedures about how to access, start, and use the Data Protection for IBM Domino GUI to back up and restore your Domino NSF databases and transaction log files are presented. You can use the Tivoli Storage Manager web client GUI or the Java client GUI to back up and restore DB2 enabled Notes databases.

Use the Data Protection for IBM Domino GUI to run the following tasks:

- Locate more information.
- Modify Data Protection for IBM Domino configuration
- Back up Domino NSF databases.
- Restore Domino NSF databases.
- Activate Domino NSF databases.
- Archive Domino NSF transaction logs.
- View and restore archived NSF transaction logs.
- Deactivate archived NSF log files.

Using the GUI

Some useful information for getting started with the Data Protection for IBM Domino GUI is presented.

Overview

The Data Protection for IBM Domino GUI is available on supported Windows operating systems. In addition, the Data Protection for IBM Domino GUI supports Domino NSF databases only. To back up and restore Domino DB2 enabled Notes databases, one of the following methods must be used:

- Tivoli Storage Manager Web client GUI or Java Client GUI
- "DB2 Commands" on page 122

Modifying preferences

Information about how to modify Data Protection for IBM Domino preferences from the Data Protection for IBM Domino GUI is presented.

Before you begin

Several configuration options can be adjusted for Data Protection for IBM Domino. The available parameters are organized into four groups. Each group can be accessed by selecting the tab for that group. The groups and related parameters are listed:

- General tab
  - **Wait for Tape Mounts**
  - **Replace Existing Files on Restore**
  - **Include Subdirectories**
  - **Notes .INI Path**
- Logging
  - **Log File Name**
  - **Prune Old Entries**
  - **Number of days to keep (Days)**
  - **Prune Now**
- Performance
  - **Number of Buffers to Use (2-8)**
  - **Buffer Size**
- Regional
  - **Language**
  - **Date Format**
  - **Time Format**
  - **Number Format**

About this task

Follow these steps to access options through the GUI.

Procedure

1. Click **Edit** to display the menu.
2. Select **Preferences** from the **Edit** menu.

What to do next

Refer to the online help for the Data Protection for IBM Domino settings for an explanation of each of these parameters.

Running an incremental backup

Information about how to use the Data Protection for IBM Domino GUI to run incremental backups is provided.

About this task

Incremental backup provides conditional backup function. To use incremental backup, follow the steps 1-7.

Procedure

1. Start the Data Protection for IBM Domino GUI.
2. Open the **Backup** tab, and select the **Incremental** option.
3. Expand the Domino server.
4. Expand the Data Directory to show its folders.
5. Select a folder to display Domino databases and subfolders within that folder.
6. Click the gray selection box to select the directory you want to back up. After your selection, the selection box changes color and contains a check mark. For incremental backups, only one directory can be selected at a time.
7. Click **Backup**. The progress box shows that the object that is being backed up, its status, and how many kb were transferred. Click **OK**.

Running a selective backup

Information about how to use the Data Protection for IBM Domino GUI to run selective backups is provided.

About this task

Use selective backup to unconditionally back up specified databases.

Follow these steps to use selective backup.

Procedure

1. Start the Data Protection for IBM Domino GUI.
2. Open the **Backup** tab, and click the **Selective** option.
3. Expand the Domino server.
4. Expand the Data Directory to show its folders.
5. Select a folder to display Domino databases and other subfolders within that folder.
6. Click the gray selection box to select the directory or individual database that you want to back up. After your selection, the selection box changes color and contains a check mark.
7. Click **Backup**. The progress box shows that the object that is being backed up, its status, and how many kb were transferred. Click **OK**.

Restoring a database

Information about how to use the Data Protection for IBM Domino GUI to restore Domino databases is provided.

Before you begin

Restore is the first step in the two-step recovery process, and activation is the second step. Restore your Domino databases after a device failure or after a database has become corrupted. Domino databases are restored by reloading a database backup, and optionally applying updates from the transaction logs that occurred after the backup was taken. Database backups can be restored over corrupted databases or to a different database file. When you restore a database to an existing file, the data in the database is overwritten and replaced by the data in the backup version.

About this task

Perform these steps to restore a database.

Procedure

1. Start the Data Protection for IBM Domino GUI.
2. Open the **Restore** tab.
3. Expand the Tivoli Storage Manager node in the list view. The tree expands showing backups from different servers.
4. Expand the tree to display the Domino Data Directory.
5. Under the Domino Data Directory, select a folder to display Domino databases.
6. Select the databases that you want to restore.
7. Select one of the following options:
   - **Restore Options**
     – If you want to restore to a different destination, enter the file path in the **Restore Into** field.
       - Example 1: Restore all selected databases to physical path `e:\temp`:
         ```
         e:\temp\=
         ```
       - Example 2: Restore all selected databases to the temporary directory within the Notes data directory:
         ```
         temp\=
         ```

       See “Domdsmc restore” on page 99 for more information about the **/into** parameter and the **restore** command.
     – If you want to activate and bring the database online, select the **Activate** check box.
     – Select **Replace Existing Files** to replace an existing database on the target machine with a restored database with the same path and name. This selection overrides settings in the Preferences file.

     Filtering Options

     – To reduce query processing time, specify a database name with letters and a wildcard character `*` in the **By Database Name** field. For example, if you specify `a*` all databases that begin with the letter `a` regardless of the folder name are displayed. Specifying `folder\a*` selects all databases that begin with `a` in the specified folder and its subfolders. Ensure to click **Update** after you type in the database query.

– If you want to restore information from a particular point in time, click **Set Point in Time**. The Point in Time Restore dialog opens with the following options:
  - Check the **Use Point in Time date when getting the list of databases** check box.
  - Specify the date and time, or click **Set Current Time**. Click **OK** to return to the Restore window.
  - Click **Update**. Note, changes specified in the Point in Time dialog cause the directory tree to refresh and lose currently selected backups. You are prompted to continue.

8. Click **Restore**. The progress box shows that the Data Protection for IBM Domino is restoring Domino databases. The **Status** shows the number of bytes restored.
9. Click **OK** when processing completes.

Activating a database

Information about how to use the Data Protection for IBM Domino GUI to activate Domino databases is presented.

Before you begin

You must restore the database before you activate it.

About this task

The second and final step to the recovery process is to activate the databases. Activation brings the restored databases online for use by the Domino server. Follow steps 1-10 to activate a database.

Procedure

1. Start the Data Protection for IBM Domino GUI.
2. Open the **Activate** tab.
3. Expand the **Activation to servername** in the tree list. **Pending dbs from local Domino server**, and **Pending dbs from other Domino servers** are listed.
4. Expand the server type and navigate to the database you want to activate.
5. Select a folder to display Domino databases.
6. Click the gray check box of the databases you want to activate. The box changes color and contains a check mark.
7. Check **Apply Logs** if you want to collect the transaction logs.
8. Click **Point in Time** if you want to activate the data from a specific point in time. Put in the date and time, or select **Set Current Time**.
9. Click **Activate**. The progress box shows that the Data Protection for IBM Domino is activating Domino databases.
10. Click **OK**.

Archiving transaction logs

Information about how to use the Data Protection for IBM Domino GUI to back up and archive Domino transaction log files is provided.

Before you begin

Ensure that archival logging is in effect on the Domino server.

About this task

With the Archive log function you can back up the Domino transaction log files when archival logging is in effect on the Domino server. Follow steps 1-3 to use archive logs.

Procedure

1. Start the Data Protection for IBM Domino GUI.
2. Click **Utilities** and select **Archive logs**.
3. The **Archive Log Progress** box shows the status of the archive operation. Click **OK**.

Viewing and restoring archived transaction logs

Details of how to use the Data Protection for IBM Domino GUI to view and restore archived transaction log files is presented.

Before you begin

You can back up (archive) Domino transaction log files when archival logging is in effect on the Domino server. This option stores filled transaction log files on the Tivoli Storage Manager server so that space allocated to these files can be reused by the Domino logger.

About this task

Perform these steps to view and restore archived transaction log files. To view archived transaction log files that are stored on the Tivoli Storage Manager server:

Procedure

1. Click **Utilities**, and select **View/Restore Log Archive**. A list of active transaction log file archives for the local Domino server is displayed.
2. Click **Active and Inactive** to view active and inactive log file archives.
3. To view log files that are archived from a different Domino server, select the appropriate Domino server from the drop-down list.
4. From the log file list, select the log file to be restored by clicking the box beside the **Log Archive Date** name.
5. Click **Restore**, the progress bar indicates that the log file is being restored.

Inactivating transaction logs

Information about how to use the Data Protection for IBM Domino GUI to inactivate your transaction log files is provided.

About this task

Inactivating Domino logs allows for expiration of the transaction log files from the Tivoli Storage Manager server. Follow steps 1-3 to inactivate the logs.

Procedure

1. Start the Data Protection for IBM Domino GUI.
2. Click **Utilities** and select **Inactivate Log Archives**.
3. The progress box shows that the Domino Application Client is inactivating the logs.

Tivoli Storage Manager Web Client GUI and Java client GUI

Information about the requirements and procedures on how to use the Tivoli Storage Manager Web Client GUI and Java client GUI. Use both GUIs to back up and restore your Domino NSF and DB2 enabled Notes databases and transaction log files.

Use the Tivoli Storage Manager Web Client GUI to back up and restore Domino server data from a remote system through a web browser. The Web Client GUI differs from the Data Protection for IBM Domino GUI. The GUI directory tree displays server nodes to other types of data and not only Domino data, depending on your environment setup. You can locate your Data Protection for IBM Domino node by opening the backup or restore window and expanding the directory tree next to the object, *Data Protection for IBM Domino*. The Java client GUI is also available locally on your desktop. Use the Java GUI in the same manner as the Web Client GUI. The main difference is that the Web Client GUI uses a web browser and provides remote access.

Getting started with the Web Client GUI

How to prepare your system for using the Tivoli Storage Manager Web Client GUI is described.

Make sure that the following requirements are satisfied before you use the Web Client GUI.

Software requirements

Details of the hardware and software requirements for IBM Tivoli Storage Manager for Mail: Data Protection for IBM Domino might change with maintenance updates, and other software currency support. For the most up-to-date requirements, visit the Hardware and Software Requirements technote that is associated with the release you are running, http://www.ibm.com/support/docview.wss?uid=swg21219345.

When the page opens, follow the link to the requirements technote for your specific release or update level.

Web browser

A web browser is required for using the Tivoli Storage Manager Web Client GUI with Data Protection for IBM Domino. The web browser must be installed on the same system as the Backup-Archive Client.

Tivoli Storage Manager Backup-Archive Client

The Tivoli Storage Manager Backup-Archive Client is required to use the Web Client GUI with Data Protection for IBM Domino operations. The Backup-Archive Client must be installed on the same system as Data Protection for IBM Domino.

Data Protection for IBM Domino plug-in

The Data Protection for IBM Domino plug-in is required to use the Web Client GUI. It must be installed on the same system as the Tivoli Storage Manager Backup-Archive Client.

You can verify that the Data Protection for IBM Domino plug-in is installed by entering the **dsmc show plugins** command. When the `Tivoli Storage Manager Domino Utility` opens, the plug-in is installed. For example:

```
<<< Installed plug-ins: >>>
******************************************************
Tivoli Storage Manager Domino Utility
******************************************************
plug-in name : PIDOM
library name : pidom.dll
library path : .\plugins\pidom.dll
function map : 0x00000001
plug-in type : Domino
plug-in ver. : 7.1.0.0
plug-in info.: NONE
plug-in lic. : C:\Program Files\Tivoli\TSM\baclient\plugins\pidomclient.lic


<<< Plug-in table information >>>
Plug-in directory search path       : .\plugins
Plug-in name criteria               : p
Plug-in load member name            : n/a
Return code from piTable creation   : 00000000
```

Environment requirements

The following environment settings must exist before you attempt to use the Tivoli Storage Manager Web Client GUI.

The Tivoli Storage Manager Web Client is installed and configured.

Use the GUI Setup wizard to install and configure the Web Client for your environment:

1. Start the Backup-Archive client GUI.
2. From the **Utilities** menu, select **Setup Wizard**.
3. Check the box next to **Help me configure the TSM Web Client**.
4. Click **Next** and follow the instructions.

Data Protection for IBM Domino is installed and configured.

You must specify the `domnode` option. For more information about these options, see “More configuration options” on page 33. In order for the web client to be able to access its partitions, specify `passwordaccess=generate` in the `dsm.opt` options file that is used by Data Protection for IBM Domino. This `dsm.opt` file is in the Data Protection for IBM Domino installation directory by default, or as specified by the **/adsmoptfile** parameter or preferences option. Ensure that you followed the instructions in the configuration section so that your system is ready to back up and restore Domino data.

Starting the Web Client GUI

1. Make sure that the software and environment requirements are met.
2. Specify the URL of the client workstation you are running the web client on, in the web browser. Specify the `httpport` number that is defined on the client workstation. The default value is 1581. For example, `http://myhost.mycompany.com:1581`. Tivoli Storage Manager logs information such as the `httpport` number, CAD activity, and errors to the `dsmwebcl.log` file. By default, this file is in the directory where the Tivoli Storage Manager backup-archive client is installed.
3. To start the Java client GUI, run the **dsm** command. The Java client GUI is only available for Linux for System z.

For more information

See the *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide* for detailed instructions about how to configure the Web Client GUI.

Backing up Domino NSF databases

How to configure Data Protection for Domino to back up Domino NSF databases is described.

Back up Domino NSF databases by following these steps:

1. Click the Backup in the Tivoli Storage Manager web client window. If you are not logged in to the server, the Tivoli Storage Manager Login window opens. Log in to the server and the Backup window opens.
2. In the `View` list at the top of the directory tree, select `Domino NSF`.
3. Expand the directory tree under Data Protection for Domino to reveal your Data Protection for Domino node. Note, this is the node name that is specified in the Tivoli Storage Manager options file.
4. Expand the directory tree next to the Domino Server to view the Domino Data Directory.
5. Expand the directory tree next to the Domino Data Directory.
6. Click the selection box next to the objects that you want to back up.
7. In the list in the Backup window, click the type of backup you want to run:

   Always backup
   This option runs an unconditional (selective) backup of the database.

   Incremental (complete)
   This option runs a conditional full backup of the database.

   Backup Domino
   This option runs a selective backup of NSF databases and a backup of the Domino DB2 database. This selection is only valid when the Domino Server node is selected.
8. Click **Backup**. After the backup completes, the Domino Backup Report window displays processing details, including the size of the backup.

To display information about a Domino object from the Backup window, select the Domino object, click the **View** menu and then **File Details**. You can then view detailed information about the backup, including whether it is compressed, encrypted, or deduplicated.

In a multiple Domino server partition environment, only one partition can be backed up at a time.

Backing up Domino DB2 enabled Notes databases and DB2 logs

To back up DB2 enabled Notes databases, the DB2 Group that contains the DB2 enabled Notes databases or the entire Domino DB2 database must be backed up.

For more information about backup operations, see "DB2 enabled Notes database backup" on page 6.

Back up Domino DB2 Groups, DB2 databases, and DB2 database logs by following these steps:

1. Click Backup in the Tivoli Storage Manager web client window. If you are not logged in to the server, the Tivoli Storage Manager login window opens. Log in to the Tivoli Storage Manager to open the Backup window.
2. From the View list in the directory tree, select Domino DB2.
3. Expand the directory tree under Data Protection for Domino to show your Data Protection for Domino node. Note, the node name that is listed is the one that is specified in the Tivoli Storage Manager options file.
4. Expand the directory tree next to the Domino Server and run one of the following backup operations:

   To run a full DB2 database backup and back up all the DB2 enabled Notes databases:

   a. Click the Domino Server.
   b. Select DB2 database from the list in the Backup window.
   c. Click `Backup`.

   To back up and archive a DB2 database log file:

   a. Click the Domino Server.
   b. Click `Archive Log`.

   To back up all the DB2 enabled Notes databases in a DB2 Group:

   a. Expand the directory tree for the DB2 Groups to view the Class Names.
   b. Expand the directory tree for the Class Names to view the individual DB2 Groups.
   c. Click the DB2 Groups you want to back up.
   d. Select `DB2 Group` in the Backup window.
   e. Click `Backup`.

   To back up the Domino DB2 enabled Notes database and all the NSF databases on the Domino Server:

   a. Click the Domino Server.
   b. Select `Backup Domino` in the Backup window.
   c. Click `Backup`.

After the backup completes, the Domino Backup Report window displays processing details, including the size of the backup.

To display information about a Domino object, select the Domino object, click the `View` menu and then `File Details`.

In a multiple Domino server partition environment, only one partition can be backed up at a time.

Restoring Domino NSF databases

The first step in the recovery process is restoring the data, activation is the second step. Domino databases must be restored after a device failure or after a database is corrupted. Domino databases are restored by reloading a database backup and optionally applying updates from the transaction logs that were made after the backup was taken. Database backups can be restored over corrupted databases or to a different database file. When you restore a database to an existing file, that existing database is overwritten with the information from the restored version.

Follow these tasks to restore a Domino NSF database:

1. Click **Restore** in the Data Protection for IBM Domino web client window. If you are not logged in to the server, the Data Protection for IBM Domino login window opens. Log in to open the Restore window.
2. In the `View` list in the directory tree, select **Domino NSF**.
3. Expand the directory tree for Data Protection for IBM Domino to show your Data Protection for IBM Domino node. This is the node name that is specified in the Data Protection for IBM Domino options file.
4. Expand the directory tree next to your Data Protection for IBM Domino node to show available Domino servers.
5. Expand the directory tree next to the Domino server that contains the databases to restore.
6. Expand the directory tree next to **Databases to Restore**.
7. Select the databases that you want to restore.
8. (Optional) If you want to run a point in time restore, click **Point In Time** and specify the date and time.
9. (Optional) If you want to activate the databases, click **Options** and check **Activate Databases on a Restore**.
10. Click **Restore**. In a multiple Domino server partition environment, only one Domino server can be restored at a time.

    Note: To display information about a Domino object from the Restore window, select the Domino object, click the `View` menu and then `File Details`. Information about whether the restore is compressed, deduplicated, or encrypted is available.

Restoring, rollforward, and activating Domino DB2 enabled Notes databases

How to restore, rollforward, and activate DominoDB2 enabled Notes databases is described.

Restore your Domino databases after a device failure or when databases are corrupted. Remember that to restore DB2 enabled Notes databases, the DB2 Group backup (that contains the DB2 enabled Notes databases) or the entire Domino DB2 database backup must be restored. Domino databases are restored by reloading a database backup and optionally applying updates from the transaction logs that occurred after the backup was taken. Database backups can be restored over corrupted databases or to a different database file. When you restore a database to an existing file, data which exists in the database is overwritten and replaced by the data in the backup version.

Follow these tasks to restore a DB2 enabled Notes database:

1. Click **Restore** in the Tivoli Storage Manager web client window or Java client window. If you are not logged in to the server, the Tivoli Storage Manager login window opens. Log in to the server, the Restore window opens.
2. From the **View** list, select Domino DB2.
3. Expand the directory tree under Data Protection for Domino to reveal your Data Protection for Domino node. Note, the node name is the one specified in the Tivoli Storage Manager options file.
4. Expand the directory tree next to your Data Protection for Domino node to reveal available Domino Servers.
5. Expand the directory tree next to the Domino Server that contains the data to restore.
   - To restore a DB2 Group that contains the DB2 enabled Notes database, expand the directory tree next to **Groups to Restore**. Find the DB2 enabled Notes database, and click the box next to the DB2 Group where it exists.
   - To restore a full DB2 database backup, expand the directory tree next to **Databases to Restore**. Click the box next to the database that you want to restore.
   - To restore a set of DB2 Groups, expand the directory tree next to **Databases to Restore**. Click the box next to the DB2 Groups that you want to restore.
   - To rollforward a DB2 database that was previously restored, expand the directory tree next to **Databases to Rollforward**. Click the box next to the database to rollforward. Use the **Rollforward options** to roll the database forward to a specific point in time. The transaction logs are applied to the DB2 database where the DB2 enabled Notes database exists.
   - To activate a DB2 database, expand the directory tree next to **Databases to Activate**. Click the box next to the database you want to activate.
6. Perform one of these actions:
   - Click **Restore** to restore a full DB2 database or a DB2 Group.
   - Click **Rollforward** to rollforward a DB2 database.
   - Click **Activate** to activate DB2 NSF databases.

Activating NSF databases

Information about how to use the Tivoli Storage Manager web client GUI or Java client GUI to activate Domino databases and apply the archived transaction logs is provided.

Before you begin

Activation brings the restored databases online for use by the Domino server. In a multiple Domino server partition environment, databases from only one Domino partition can be activated at a time.

About this task

Follow these steps to activate a Domino database.

Procedure

1. Click Restore in the Tivoli Storage Manager web client window or the Java client window. If you are not logged in to the server, the Tivoli Storage Manager login window opens. Log in to the server the Restore window opens.

2. Expand the directory tree under Data Protection for IBM Domino to show your Data Protection for IBM Domino node. Note, the node name is the one server that is specified by the **domnode** option.
3. Expand the directory tree next to the Domino server that contains the databases to activate.
4. Expand the directory tree next to **Databases to Activate**.
   - The databases under **Domino Data Directory** are databases that are restored to their original location.
   - The databases under **Other Databases** are databases that are restored to an alternative location.
5. Click the selection box next to the databases that you want to activate.
6. (Optional) If you want to activate databases to a time other than the current time, click **Point In Time**, and then use the **Apply Logs** option.
7. (Optional) If you want to activate to the current time, click **Options**, and check the box next to **Apply Logs**.
8. Click **Activate**.

Backing up and archiving NSF transaction logs

The Tivoli Storage Manager web client GUI or Java client GUI can be used to back up and archive Domino transaction log files.

Before you begin

You can back up (archive) Domino transaction log files when archival logging is in effect on the Domino server. This option stores filled transaction log files on the Tivoli Storage Manager server so that space allocated to these files can be reused by the Domino logger.

About this task

Perform these steps to back up Domino transaction log files.

Procedure

1. Click **Backup** in the Tivoli Storage Manager web client or the Java client window. If you are not logged in to the server, the Tivoli Storage Manager Login window opens. Enter the login details and click **Login**. The Backup window opens.
2. Expand the directory tree under Data Protection for IBM Domino to show your Data Protection for IBM Domino node. The node name is the one that is specified in the Tivoli Storage Manager options file.
3. Select the Domino server that contains the transaction log files to back up. In a multiple Domino server partition environment, only one Domino server can be archived at a time.
4. Click **Archive Log**. After processing completes, the Domino Backup Report window opens with processing details, including the size of the backup.

   Note: Information about the status of backups is also shown in this window. This indicates whether the backup is compressed or deduplicated, along with information about how many LAN-free bytes were transferred.

Restoring NSF transaction logs

Information about how to use the Tivoli Storage Manager web client GUI and Java client GUI to restore Domino transaction log files is provided.

Before you begin

Necessary Domino transaction log files are restored automatically during database restore processing. Restore these log files manually only in special circumstances. Note, in a multiple Domino server partition environment, only one transaction log file can be restored at a time.

About this task

Perform these steps to restore a Domino transaction log file from the Tivoli Storage Manager server.

Procedure

1. Click Restore in the Tivoli Storage Manager web client window. If you are not logged in to the server, the Tivoli Storage Manager login window opens. Log in to the server, the Restore window opens.
2. Expand the directory tree under Data Protection for IBM Domino to reveal your Data Protection for IBM Domino node. Note, the node name that is specified in the Tivoli Storage Manager options file.
3. Expand the Domino server that contains the transaction log files to restore.
4. If you want to restore all available transaction log files, click the selection box next to **Restore Log Archive** and then click **Restore**.
5. If you want to restore individual transaction log files, expand the directory tree next to **Restore Log Archive** and then click the box next to the transaction log files you want to restore. Click **Restore**.

   Note: The restore progress report shows details of how many LAN-free bytes were transferred. Click **Details** for detailed information about the backup. Details about encryption, deduplication and compression are shown. To display information about a Domino object from the Restore window, select the Domino object, click **File Details** from the **View** menu.

Inactivating archived NSF transaction logs

How to use the Tivoli Storage Manager web client GUI to inactivate archived Domino transaction log files is outlined.

About this task

Perform these steps to inactivate archived Domino transaction log files.

Procedure

1. Click **Restore** in the Tivoli Storage Manager web client window. If you are not logged in to the server, the Tivoli Storage Manager login window opens. Log in to the server to open the Restore window.
2. Expand the directory tree under Data Protection for IBM Domino to show your Data Protection for IBM Domino node. Note, this is the node name that is specified in the Tivoli Storage Manager options file.
3. Select the **Restore Log Archive** option.

4. Select **Inactivate Log Archives**. Note, in a multiple Domino server partition environment, only one Domino Server at a time is available for **Inactivate Log Archives**.

Command-line interface

How to use the Data Protection for IBM Domino command-line interface for tasks with Domino NSF and DB2 databases is provided.

NSF Commands

How to use the Data Protection for IBM Domino command-line interface with Domino NSF databases is provided.

Domdsmc activatedbs

The **Domdsmc activatedbs** command brings restored database backups online.

Purpose

Use **Domdsmc activatedbs** command to bring restored database backups online. If the database is logged, you can apply all applicable transactions from the transaction logs. Alternatively you can apply transactions up to a specific point in time to update the database.

You can put databases in a corrupted state if you press **CTRL-C** or cancel the job in the middle of an activate action. Canceling the action prevents the databases from being activated. Also, any databases that are activated before the **CTRL-C** cancel was run, can be corrupted, and must be restored and activated again.

This command acts on restored database backups that are pending activation, and that have been restored with the **/activate=no** parameter. The Tivoli Storage Manager server is not contacted unless archived transaction logs are needed for the **/applylogs** parameter.

If you receive the Domino message `Recovery Manager: Database is not latest copy.` when you issue the **domdsmc activatedbs** command with the **/applylogs** parameter, a problem might exist with your Domino Logger ID. Consult your Domino documentation to determine why you received this error message. You can also run the **domdsmc query logarchive** command to view archived transaction log extents for more information. If two `Logger IDs` display, as in the following example:

```
Domino Server: restroan01
--------------

Logger Id: OF88256BC1:005F8602-ON00000365:0136DCCF
----------
                          Transaction

Log Archive Date       Log Filename    A/I      Size
--------------------   ------------    ---     ------
05/22/2004 10:27:26      S0000000.TXN     A     64.00MB


Domino Server: restroan01
--------------

Logger Id: OF88256BC1:005EDBA5-ON00000365:0136DCCF
----------
            Transaction

Log Archive Date       Log Filename    A/I      Size
--------------------   ------------    ---     ------
05/22/2004 10:20:23      S0000000.TXN     A     64.00MB
```

Data Protection for IBM Domino uses an alternative restore path for the transaction logs on a Domino 6 or later environment when the *TRANSLOG_RECOVER_PATH* variable is specified in the NOTES.INI file. If the alternative log restore path specified in the NOTES.INI file is not a fully qualified path, Data Protection for IBM Domino does not use the alternative restore path.

You can use any of the displayed Logger IDs to restore logged databases. To use any of the Logger IDs other than the current one, you must use an alternative server to restore logged databases. See "NSF databases restore to alternate server and alternate partition" on page 167 for detailed instructions on how to run this procedure.

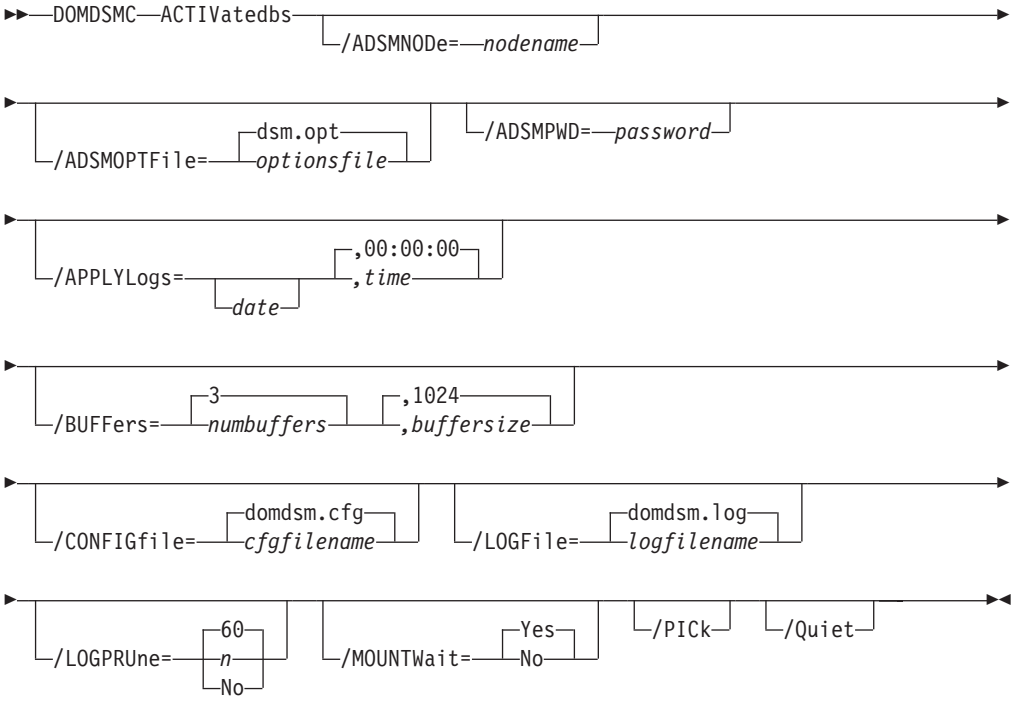

Parameters

/ADSMNODe=*nodename*
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile= *optionsfile*
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD= *password*
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify the command **passwordaccessgenerate** in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If **passwordaccess** is set to `generate` and you specify a password, the value is ignored unless a password for this node is stored. In this case, the specified password is stored and used for the current command execution.

If **passwordaccess** is set to `prompt` and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the requirement to prompt.

If **passwordaccess** is set to `prompt` and you do not specify a password on the command line, then you are prompted for a password.

/APPLYLogs =*date*,*time*
Specifies that transaction log recovery for the restored databases is run if they are logged. The **date** and **time** values must be specified in the same date and time format that is defined in the Data Protection for IBM Domino preferences file. The transaction logs are applied to a specified point in time or to the current date and time if no **date** and **time** values are specified.

date Specify a date string in the active date format. Transactions that are completed and committed before the specified date are applied to the restored database. The date that is specified must be after the backup date of the backup image that is being restored. The **/pit** option can be used with the **restore** command to automatically restore the most recent full backup image that is run before the wanted point-in-time.

Because there is one transaction log for all logged databases, all the databases must be activated together in one command. This situation applies when you are restoring multiple databases that requirement to have transactions that are applied from the log. This prevents the fetching of the same transaction logs multiple times from the Tivoli Storage Manager server. The databases can be restored separately if necessary, with the `/activate=no` parameter and then activated together with a single **activatedbs** command.

If you are restoring a database that is backed up from a different Domino server, logged transactions cannot be applied. In this case, you can activate only a full backup image. You must also use the Notes **fixup** utility to reset the internal sequence numbers of the restored and activated database.

Note: If circular logging is in effect, it might not be possible to properly apply transactions if the log is wrapped. If an attempt to apply transaction logs fails, the database or databases that are processed are marked as corrupted. The database or databases must be restored again.

The date must be specified with the same date format that is defined in the Data Protection for IBM Domino preferences file. See "Domdsmc set" on page 115 for a list of available date formats.

time Specify a time string in the active time format. If you specify a date without the time, 00:00:00 on a 24-hour clock is used.

The time must be specified with the same time format that is defined in the Data Protection for IBM Domino preferences file. See "Domdsmc set" on page 115 for a list of available time formats.

/BUFFers= *numbuffers,buffersize*
Specifies the number and size of data buffers that transfer data between the Domino server and the Tivoli Storage Manager API. Increasing the number or size (or both) of the data buffers can improve throughput.

You can specify 2 - 8 buffers. The default value is 3. The size of the buffers can be 64 - 8192 kb. The default value is 1024.

If the `/buffers` parameter is not specified on the command line or defined in the preferences file, Data Protection for IBM Domino uses the default values.

/CONFIGfile= *cfgfilename*
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file is in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile= *logfilename*
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. Logging for each instance is directed to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne= *60|n|No*
Specifies whether to prune log entries. By default, log pruning is enabled to run daily. You can use the **set** command to:

- Change the defaults so that log pruning is disabled.
- Change the number of days log entries are saved.

You can use the /logprune option to override these defaults for one command run. Note, when the value of /logprune is a number, the prune is done even if one is already completed for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the Data Protection for IBM Domino log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter is changed, run one of the following to prevent undesired pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or logfile setting.

/MOUNTWait= *Yes|No*

If the Tivoli Storage Manager server is configured to store transaction log backup data on removable media, then the Tivoli Storage Manager server can indicate to Data Protection for IBM Domino that it is waiting for a required storage volume to be mounted. With this option, you can specify whether Data Protection for IBM Domino waits for the media mount or stops the current operation. Removable media is media such as tapes or DVDs.

You can specify:

Yes Wait for tape mounts.

No Do not wait for tape mounts.

/PICk

Displays a list of the restored databases that are waiting for activation. The databases that are listed are the ones that match the dbname pattern specified. Databases to be activated can be selected from the list.

The pick list is presented as a scrollable list with the same manipulation functions as in the base Tivoli Storage Manager client PICK function.

/Quiet

Specifies that status information does not display. However, the information is written to the activity log.

Examples

Example 1: The following example brings all the restored database backups online and applies transactions from the transaction log to update the database to the date specified:

```
domdsmc activatedbs /applylogs=02/23/2007

Output example:
```

```
Starting Domino database activation...

Initializing Domino connection...
Logging on to the Tivoli Storage Manager server, please wait...

Starting archivelog recovery...

Media Recovery Replay: 100%
02/22/07 04:32:25 PM  Recovery Manager: Media Recovery complete for
i:\Lotus\Domino\Data\datadir3\yyyy.nsf.dad,last update applied 02/22/07
03:51:12 PM.

Archivelog recovery completed successfully.

Activating database datadir3\yyyy.nsf, 1 of 1,
Activate of datadir3\yyyy.nsf completed successfully.


Total pending databases inspected:                1
Total pending databases requested for activation:  1
Total pending databases activated:                1

Throughput rate:                                  0.00 Kb/Sec
Total bytes transferred:                          0
Elapsed processing time:                          0.00 Secs
```

Example 2: The following example brings all the restored database backups online:

```
domdsmc activatedbs
```

Output example:

```
Starting Domino database activation...

Initializing Domino connection...
Logging on to the Tivoli Storage Manager server, please wait...

Activating database testdb2.nsf, 1 of 1,
Activate of testdb2.nsf completed successfully.


Total pending databases inspected: 1
Total pending databases requested for activation: 1
Total pending databases activated: 1
```

Domdsmc archivelog

Domdsmc archivelog backs up Domino transaction log files when archival logging is in effect on the Domino server.

Purpose of Domdsmc archivelog

The **Domdsmc archivelog** command queries the Domino server to determine whether any log extents are ready for archiving. If so, the log files are backed up to Tivoli Storage Manager server storage, and the Domino server is notified of their availability for reuse. High and low threshold values can be specified as a percentage of the log capacity to control whether log files are archived when the command is run. This allows the command to be scheduled regularly to protect against a log full condition but to actually do the archive only if the log is getting close to being full. If enough log space is allocated to contain an average day of updates, it is possible to establish a strategy where log files are normally archived daily during low usage time. A daily schedule without threshold values but unusually high volumes of change can also be handled on an exception basis. For

example, an hourly schedule of the **archivelog** command with appropriate specified threshold values archives only if necessary.

Run this command frequently to ensure that allocated transaction log space is freed.

The active transaction log is also backed up with Domino Server 8.5.x or a more recent version.

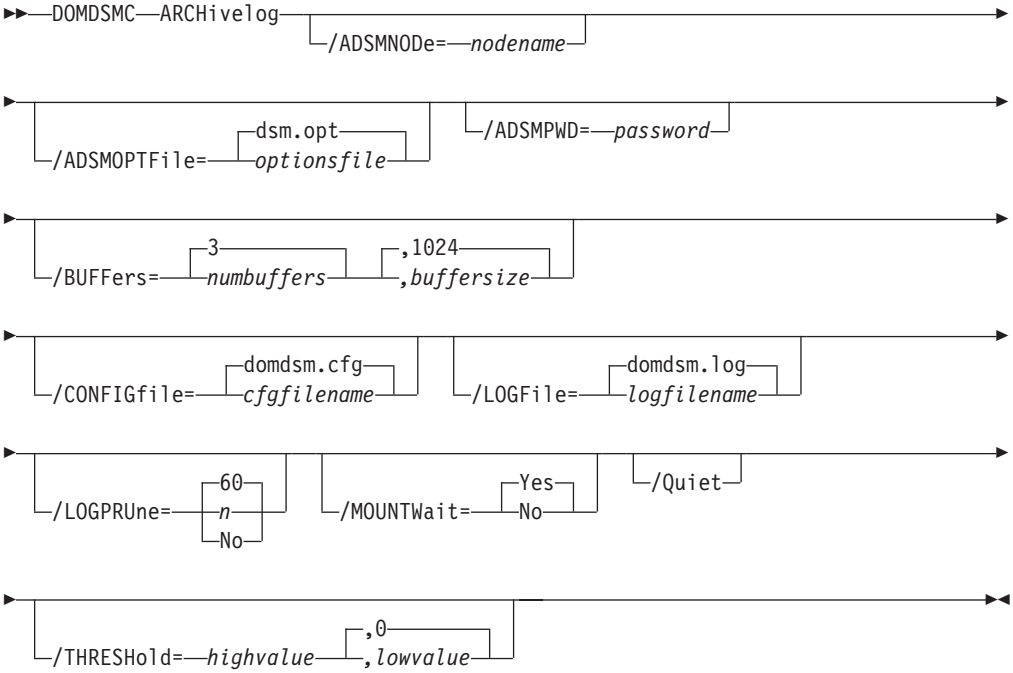

Parameters

/ADSMNODe=nodename
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify **passwordaccess generate** in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If **passwordaccess** is set to **generate** and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If **passwordaccess** is set to **prompt** and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the need to prompt.

If **passwordaccess** is set to **prompt** and you do not specify a password on the command line, then you are prompted for a password.

/BUFFers=numbuffers, buffersize
Use this parameter to specify the number of data buffers and the size of data buffers that transfer data between the Domino server and the Tivoli Storage Manager API. Increasing the number and size of the data buffers can improve throughput.

You can specify from 2 to 8 buffers, the default value is 3. The size of the buffers can be from 64 to 8192 kb, the default value is 1024.

If the `/buffers` parameter is not specified on the command line or defined in the preferences file, Data Protection for IBM Domino uses the default values.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the `set` command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and performed once per day. You can use the `set` command to:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is completed for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. This is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the `timeformat` or `dateformat` parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the `timeformat` or `dateformat` parameter has changed, do one of the following to prevent unwanted pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/MOUNTWait=Yes|No

If the Tivoli Storage Manager server is configured to store transaction log backup data on removable media, then the Tivoli Storage Manager server can indicate to Data Protection for IBM Domino that it is waiting for a required storage volume to be mounted. The mount command is used to specify whether Data Protection for IBM Domino waits for the media mount or stops the current operation. Removable media is media such as tapes or DVDs.

You can specify:

Yes Wait for tape mounts. This value is the default.

No Do not wait for tape mounts.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/THRESHold=highvaluelowvalue,

Use this option to specify when the `archivelog` command starts and stops archiving eligible transaction log files. The `highvalue`, specified as a percentage of the transaction log capacity, identifies the point at which log archiving begins. If the current occupancy of the transaction log equals or exceeds the value for this parameter, eligible log files are archived until the occupancy falls to or below the `lowvalue` which is also specified as a percentage of the log capacity.

The *highvalue* variable is an integer in the range from 1 to 99.

The *lowvalue* variable is an integer in the range from 0 to 98 but it must be less than the high value. The default is 0 which means that all log files eligible for archive are archived. Specify a low threshold value, greater than 0 to prevent the active transaction log from being backed up.

If the `/threshold` option is not specified, then all eligible transaction log files are archived. The active transaction log is an eligible transaction log file.

For example, specifying the following command will archive transaction log files after the log is at or more than 90% full. The archive process stops when sufficient space is reclaimed to make the log less than or equal to 50% full:

```
/threshold=90,50
```

It is important to note that the **/threshold** option is impacted by the total size of the transaction log. The total size of the transaction log is determined on the Domino server by the value of the TRANSLOG_UseAll and TRANSLOG_MaxSize options in the notes.ini file.

Example

The following example backs up the current archive log:

```
domdsmc archivelog
```

Output example:

```
Starting Domino transaction log archive...
Initializing Domino connection...
Logging on to the Tivoli Storage Manager server, please wait...

Archiving transaction log file d:\domino\logging\S0000050.TXN
Full: 0  Read: 67,109,888  Written: 169,119  Rate: 41.67 Kb/Sec
Archive of d:\domino\logging\S0000050.TXN completed successfully.


Total Domino transaction log files ready for archive:  1
Total Domino transaction log files archived:   1
Total Domino transaction log files deduplicated:  1

Throughput rate:     41.64 Kb/Sec
Total bytes inspected:     67,109,888
Total bytes transferred:    169,119
Total LanFree bytes transferred:   0
Total bytes before deduplication:  67,109,888
Total bytes after deduplication:   400,042
Data compressed by:   58%
Deduplication reduction:  99.41%
Total data reduction ratio:  99.75%
Elapsed processing time:  3.97 Secs
```

Domdsmc changeadsmpwd

How to use the **domdsmc changeadsmpwd** command is described.

Purpose

Domdsmc changeadsmpwd changes the Tivoli Storage Manager password that is used by Data Protection for IBM Domino. If you do not enter the old and new passwords on the command, you are prompted for them. When Data Protection for IBM Domino prompts you for the passwords, the password is not displayed on the screen.

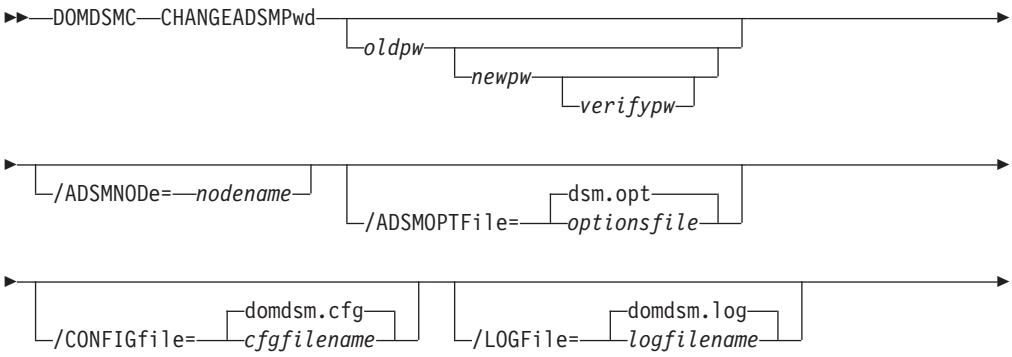

Parameters

oldpw The current password to change. You are prompted for this value if omitted.

newpw The new password. You are prompted for this value if omitted. When you choose a new password, you can use from 1 to 64 characters.

Valid password characters are as follows:

A-Z Any letter, A through Z, uppercase, or lowercase.

0–9 Any number, 0 through 9.

+ Plus.

. Period.

_ Underscore.

- Hyphen.

& Ampersand.

A password is not case-sensitive.

verifypw
The verify password is used to validate the password that is entered for `newpw`. You are prompted for this value if omitted.

/ADSMNODe=nodename
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.

The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

Example

The following example changes the Tivoli Storage Manager password to `secret`:

```
domdsmc changeadsmpwd oldpassword secret secret
```

Output example:

```
ACD0260I Password successfully changed.
```

Domdsmc help

How to use the **domdsmc help** command is described.

Purpose

The **domdsmc help** command provides online help for **domdsmc**. Running **domdsmc help** returns a list of commands and their parameters.

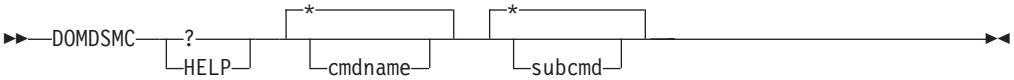

Parameters

***|cmdname**

Identifies the specific Data Protection for IBM Domino command that is to be displayed. If the wildcard character * is used, help for all Data Protection for IBM Domino commands is displayed. If you do not specify a command name, asterisk (*) is the default.

The valid command names are shown here:

```
ACTIVatedbs
ARCHivelog
CHANGEADSMPwd
DB2ACTivatedbs
DB2ARCHivelog
DB2DELetealternate
DB2INActivateobjs
DB2RESTore
DB2ROLLforward
DB2Selective
FULLSelective
HELP
INACTivatelogs
Incremental
Query
RESETdatabase
RESTore
RESTORELOGArchive
SELective
SET
UPDATEDB2Pwd
```

***|subcmd**

Help can be displayed for commands that have several subcommands, for example, the **query** commands. If you do not specify a subcommand or the wildcard character asterisk *, then help for all Data Protection for IBM Domino **query** commands is displayed.

The valid subcommand names for the **query** commands are shown in the following list:

```
Adsmserver
DBBackup
DB2Backup
DB2Pendingdbs
DB2ROLLforward
DOMino
LOGArchive
PENDingdbs
PREFerences
```

Examples

Example 1: The command **domdsmc** or **domdsmc help *** provides information about the syntax of all the commands.

Output example:

```
IBM Tivoli Storage Manager for Mail:
Data Protection for Lotus Domino
Version 7, Release 1, Level 0.0
(C) Copyright IBM Corporation 1999, 2011, 2013. All rights reserved.


 Choose from the following commands:

DOMDSMC ACTIVatedbs
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/APPLYLogs=date[,time]] (default: currentdate,currenttime)
  [/BUFFers=numbuffers[,buffersize]] (default: 3,1024)
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/MOUNTWait=Yes|No] (default: Yes)
  [/PICk]
  [/Quiet]

DOMDSMC ARCHivelog
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/BUFFers=numbuffers[,buffersize]] (default: 3,1024)
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/MOUNTWait=Yes|No] (default: Yes)
  [/Quiet]
  [/THRESHold=highvalue[,lowvalue]]

DOMDSMC CHANGEADSMPwd [oldpw [newpw [verifypw]]]
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC DB2ACTivatedbs dbname[dbname,...]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/DB2ALtdbnames=db2database] (default: DOM_ALT)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/INTO=filename]
  [/ISOLATE]
  [/LOCKGroup]
  [/LOGPRUne=60|n|No] (default: 60)
  [/PICk=[SHOWAl]]
  [/REPlace=Yes|No] (default: Yes)
  [/Quiet]

DOMDSMC DB2ARCHivelog
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/Quiet]

DOMDSMC DB2DELetealternate db2database
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
```

```
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC DB2INActivateobjs
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/Quiet]
  [/SERVer=currentserver|servername]

DOMDSMC DB2RESTore db2group[,db2group,...]
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/DB2ALtdbname=db2datbase] (default: DOM_ALT)
  [/DB2DATAbase=db2database]
  [/DB2CONTainerpath=path]
  [/DB2LOGPath=path]
  [/DB2REPlace=Yes|No] (default;Yes)
  [/DB2RESTIntopath=path]
  [/DB2SESSIONS=numsessions] (default: 1)
  [/FULL]
  [/INPlace]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/PICk=SHOWACtive|SHOWALl] (default: SHOWACtive)
  [/PIT=date[,time]] (default: currentdate,currenttime)
  [/Quiet]
  [/SERVer=currentserver|servername]

DOMDSMC DB2ROLLforward db2database
  [/APPLYLogs=date[,time]]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/PICk=[SHOWALl]]
  [/Quiet]

DOMDSMC DB2Selective db2group[,db2group,...]
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/DB2SESSIONS=numsessions] (default: 1)
  [/FULL]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/Quiet]

DOMDSMC FULLSelective
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/BUFFers=numbuffers[,buffersize]] (default: 3,1024)
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/DB2SESSions=numsessions (default: 1)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/MOUNTWait=Yes|No] (default: Yes)
  [/Quiet]
  [/SESSions=numsessions (default: 1)

DOMDSMC HELP|? [*|command] [*|subcmd]
```

```
Valid command names :        Valid subcmds :
   ACTIVatedbs                  Adsmserver
   ARCHivelog                   DBBackup
   CHANGEADSMPwd                DB2Backup
   DB2ACTivatedbs               DB2Pendingdbs
   DB2ARCHivelog                DB2ROLLforward
   DB2DELetealternate           DOMino
   DB2INActivateobjs            LOGArchive
   DB2RESTore                   PENDingdbs
   DB2ROLLforward               PREFerences
   DB2Selective
   FULLSelective
   HELP
   INACTivatelogs
   Incremental
   Query
   RESETdatabase
   RESTore
   RESTORELOGArchive
   Selective
   SET

DOMDSMC INACTivatelogs
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/Quiet]
  [/SERVer=currentserver|servername]

DOMDSMC Incremental dbname[,dbname,...]
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/BUFFers=numbuffers[,buffersize]] (default: 3,1024)
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/MOUNTWait=Yes|No] (default: Yes)
  [/Quiet]
  [/SUBDir=No|Yes] (default: No)
  [/SESSions=numsessions (default: 1)

DOMDSMC Query Adsmserver
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC Query DBBackup *|dbname
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/DEtail]
  [/INACTive]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/SERVer=currentserver|servername]
  [/SUBDir=No|Yes]  (default: No)

DOMDSMC Query DB2Backup *|db2group
  [/ADSMNODe=nodename]
```

```
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/DB2DATAbase=db2database]
  [/DEtail]
  [/FULL]
  [/INACTive]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/SERVer=currentserver|servername]

DOMDSMC Query DB2Pendingdbs
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC Query DB2ROLLforward
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC Query DOMino [*|dbname]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/SUBDir=No|Yes] (default: No)
  [/TYpe=All|Nsf|Db2] (default: All)

DOMDSMC Query LOGArchive
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/DEtail]
  [/INACTive]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/SERVer=currentserver|servername]
  [/FROMDate=date]
  [/TODate=date]

DOMDSMC Query PENDingdbs
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC Query PREFerences
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC RESETdatabase [dbname|dbname,...]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC RESTore dbname[,dbname,...]
  [/ACTIVate=No|Yes] (default: No)
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/BUFFers=numbuffers[,buffersize]] (default: 3,1024)
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/INTO=filename]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/MOUNTWait=Yes|No] (default: Yes)
```

```
  [/PICk=SHOWACtive|SHOWALl] (default: SHOWACtive)
  [/PIT=date[,time]] (default: currentdate,currenttime)
  [/Quiet]
  [/REPlace=Yes|No]  (default: Yes)
  [/SERVer=currentserver|servername]
  [/SUBDir=No|Yes] (default: No)

DOMDSMC RESTORELOGArchive [logname[,logname,...]] (default: 'last')
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/BUFFers=numbuffers[,buffersize]] (default: 3,1024)
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/INTOPath=pathname]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/MOUNTWait=Yes|No] (default: Yes)
  [/PICk=SHOWACtive|SHOWALl] (default: SHOWACtive)
  [/Quiet]
  [/REPlace=Yes|No]  (default: Yes)
  [/SERVer=currentserver|servername]
  [/FROMDate=date]
  [/TODate=date]

DOMDSMC Selective dbname[,dbname,...]
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/BUFFers=numbuffers[,buffersize]] (default: 3,1024)
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGGedonly]
  [/LOGPRUne=60|n|No] (default: 60)
  [/MOUNTWait=Yes|No] (default: Yes)
  [/Quiet]
  [/SESSions=numsessions (default: 1)
  [/SUBDir=No|Yes] (default:No)

DOMDSMC SET PARMname=value
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)

  where PARMname and default values are:
     ADSMLOGDir=
     ADSMOPTFile=
     BUFFers=3  (2..8)
     BUFFERSIze=1024  (64..8192)
     DATEformat=
       1    MM/DD/YYYY
       2    DD-MM-YYYY
       3    YYYY-MM-DD
       4    DD.MM.YYYY
       5    YYYY.MM.DD
     DB2ALtdbname=DOM_ALT
     DB2CONTainerpath=
     DB2LOGPath=
     DB2LOGTarget=
     DB2REPlace=Yes (Yes|No)
     DB2RESTIntopath=
     DB2SESSions=1 (1..64)
     DB2USER=
     LANGuage=
       ENU  (English, United States)
       PTB  (Brazilian Portuguese)
       CHS  (Chinese, Simplified)
       CHT  (Chinese, Traditional)
       FRA  (Standard French)
       DEU  (Standard German)
```

```
        ITA  (Standard Italian)
        JPN  (Japanese)
        KOR  (Korean)
        ESP  (Standard Spanish)
        CSY  (Czech)
        HUN  (Hungarian)
        PLK  (Polish)
        RUS  (Russian)
      LOGFile=domdsm.log
      LOGPRUne=60  (0..9999 | No)
      MOUNTWait=Yes  (Yes|No)
      NOTESInipath=
      NUMberformat=
        1    n,nnn.dd
        2    n,nnn,dd
        3    n nnn,dd
        4    n nnn.dd
        5    n.nnn,dd
        6    n'nnn,dd
      REPlace=Yes  (Yes|No)
      SESSions=1 (1..64)
      STATistics=No (No|Yes)
      SUBDir=No  (No|Yes)
      TIMEformat=
        1    HH:MM:SS
        2    HH,MM,SS
        3    HH.MM.SS
        4    HH:MM:SSA/P
    DOMTXNBYTElimit=0 (0..2097152)
    DOMTXNGROUPmax=2 (2,65000)
    COMMRESTARTDuration=60(1,9999)
    COMMRESTARTInterval=15(1,9999)

DOMDSMC UPDATEDB2Pwd [oldpw [newpw [verifypw]]]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

EXAMPLES:

  DOMDSMC Selective adb.nsf
  DOMDSMC Query Domino
```

Example 2: To display help for all the **query** commands, enter the following command:

```
domdsmc help query *
```

Output example:

```
DOMDSMC Query Adsmserver
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC Query DBBackup *|dbname
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/INACTive]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/SERVer=currentserver|servername]
  [/SUBDir=No|Yes]  (default: No)

DOMDSMC Query DB2Backup *|db2group
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/DB2DATAbase=db2database]
  [/DEtail]
  [/FULL]
  [/INACTive]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/SERVer=currentserver|servername]
```

```
DOMDSMC Query DB2Pendingdbs
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC Query DB2ROLLforward
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/DELetedb=db2database]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC Query DOMino [*|dbname]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/SUBDir=No|Yes] (default: No)
  [/TYpe=All|Nsf|Db2] (default: All)

DOMDSMC Query LOGArchive
  [/ADSMNODe=nodename]
  [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
  [/ADSMPWD=password]
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/INACTive]
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
  [/SERVer=currentserver|servername]

DOMDSMC Query PENDingdbs
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)

DOMDSMC Query PREFerences
  [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
  [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
  [/LOGPRUne=60|n|No] (default: 60)
```

Example 3: To display help for the **query domino** command, enter the following command:

```
domdsmc help query domino
```

Output example:

```
DOMDSMC Query DOMino [*|dbname]
[/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
[/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
[/LOGPRUne=60|n|No] (default: 60)
[/SUBDir=No|Yes] (default: No)
[/TYpe=All|Nsf|Db2] (default: All)
```

Example 4: To display help for the **db2selective** command, enter the following command:

```
domdsmc help db2selective
```

Output example:

```
DOMDSMC DB2Selective db2group[,db2group,...]
 [/ADSMNODe=nodename]
 [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
 [/ADSMPWD=password]
 [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
 [/DB2SESSIONS=numsessions] (default: 1)
 [/FULL]
 [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
 [/LOGPRUne=60|n|No] (default: 60)
 [/Quiet]
```

Example 5: To display help for the **query db2backup** command, enter the following command:

```
domdsmc help query db2backup
```

Output example:

```
DOMDSMC Query DB2Backup *|db2group
 [/ADSMNODe=nodename]
 [/ADSMOPTFile=dsm.opt|filename] (default: dsm.opt)
 [/ADSMPWD=password]
 [/CONFIGfile=domdsm.cfg|filename] (default: domdsm.cfg)
 [/DB2DATAbase=db2database]
 [/DEtail]
 [/FULL]
 [/INACTive]
 [/LOGFile=domdsm.log|logfilename] (default: domdsm.log)
 [/LOGPRUne=60|n|No] (default: 60)
 [/SERVer=currentserver|servername]
```

Domdsmc inactivatelogs

How to use the **Domdsmc inactivatelogs** command is described.

Purpose

Domdsmc inactivatelogs expires transaction log files from backup storage. Because there is a single shared transaction log for all logged databases on a Domino server, log files cannot be deactivated. This can happen only when all databases that require that log file for recovery are inactive. This command queries the database backups on the Tivoli Storage Manager server to determine which log files are required by any active database backup. This command also deactivates log files that are no longer required. Run this command after full database backups are completed to deactivate the transaction logs at the same time the database backups that require them are deactivated.

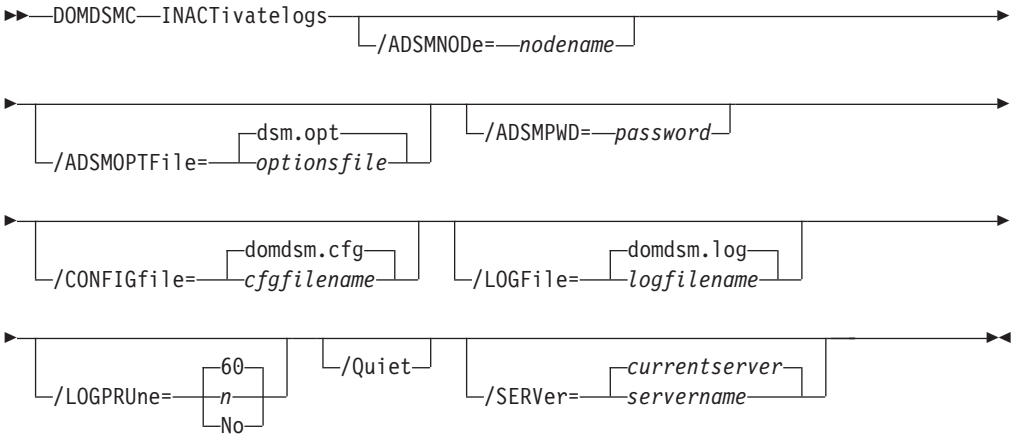

Parameters

/ADSMNODe=nodename

Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to logon to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile

Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password

Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify `passwordaccessgenerate` in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If `passwordaccess` is set to `generate` and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If passwordaccess is set to prompt and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the prompt.

If passwordaccess is set to prompt and you do not specify a password on the command line, then you are prompted for a password.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is domdsm.cfg.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is domdsm.log.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the /logfile parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the /logprune option to override these defaults for one command run. Note, when the value of /logprune is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that

prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/SERVer=currentserver|servername
Specifies the Domino server name. If not specified, Data Protection for IBM Domino uses the current Domino server.

Example

This example causes the archive log extents that are no longer needed to expire:

```
domdsmc inactivatelogs
```

Output example:

```
Number of Logs Inactivated: 1
```

Domdsmc incremental

How to use the **domdsmc incremental** command is described.

Purpose

Running **domdsmc incremental** runs the following functions:

- Backs up new databases since the last backup, or newly included ones.
- Backs up any non-logged databases that changed since the last backup (based on modification dates of both data and metadata).
- Backs up any logged databases with a changed DBIID when archival logging is in effect.
- Inactivates any active database backups on the Tivoli Storage Manager server that are excluded from backup or no longer exist on the Domino server.

A query of the current backup objects from the Tivoli Storage Manager server is required before any actions take place.

This command backs up a database by matching the `dbname` pattern with the following conditions:

- The database is not excluded in the Tivoli Storage Manager options file. Standard `include` and `exclude` processing is supported.
- The database is not logged, and was modified since the last active backup image for that database. Both data and non-data modification dates are checked. If either date is different from the active backup, the database is backed up.
- Archival logging is in effect, and the DBIID of a logged database is changed. If the DBIID is not changed, then logged databases are not backed up. The changes are captured in the transaction log backups. In this case, periodic selective backups of all logged databases must be done to refresh the active backup images. This reduces the number of transaction logs to be applied during a recovery.

Note: When circular logging is used on the Domino server (or when logging is disabled on the Domino server), transaction log files are not archived. For more information about logging, see "NSF backup strategy" on page 4.

- The database is new or newly included in the backup (an active backup image does not exist on the Tivoli Storage Manager server).

The incremental command can also deactivate active backup images for databases that are deleted from the Domino server or excluded from backup. Backups can automatically be expired according to the retention parameters that are defined in the Tivoli Storage Manager management class.

Use the incremental command to back up a single directory or all databases within the Notes data path by specifying an appropriate `dbname` pattern.

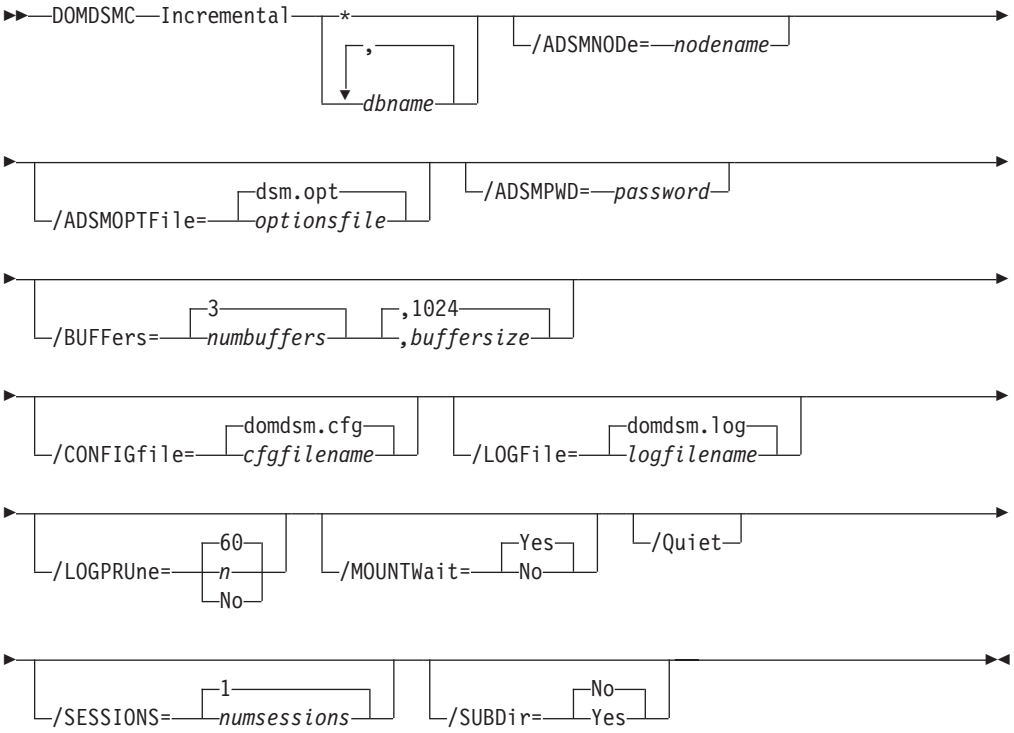

Parameters

***|dbname,dbname,...,**

Specifies the file path of a database or file path pattern for a group of databases. The file path pattern can represent a group of databases to be conditionally backed up. The wildcard character asterisk (*) is used to specify a group of databases when used in the `dbname`. Multiple `dbnames` can be specified separated with commas.

The file path must be relative to the Notes data directory.

Symbolic links are referred to by their symbolic names. To reference a database in a directory outside of the Notes data directory and any subsequent directories pointed to by a directory link in the data path, use the directory link name as the directory name. For example, if database `xyz.nsf` is in a directory, and pointed to by the link `vol1.dir`, refer to it as

vol1\xyz.nsf. If a symbolic directory link is created with the same name as a physical directory in the Notes data path, only the physical directory is searched.

The wildcard character (*) is used to represent any number of any characters when used in the file name portion of the file path. The wildcard character is not supported within directory names. The following example backs up all databases within the dir_A directory beginning with the characters **ter**:

```
domdsmc incremental dir_A\ter*
```

The following example backs up all databases on the server that meet the criteria of the incremental backup:

```
domdsmc incremental * /subdir=yes
```

The following example backs up all databases whose file name ends in **acct**:

```
domdsmc incremental *acct.n* /subdir=yes
```

Note: Standard include and exclude processing applies to Domino database names. Wildcards can be used on the backup command, and specific databases can be excluded from the backup with the include-exclude list in the Tivoli Storage Manager options file. For example, to exclude all databases on a volume pointed to by the symbolic directory link temp.dir, use the following statement:

```
exclude \temp\*
```

The exclude statement refers to the relative file name that includes symbols, and not the physical file path. For more information about include and exclude options, see "Include and exclude processing" on page 171 and *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

/ADSMNODe=nodename

Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile

Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is dsm.opt.

/ADSMPWD=password

Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify **passwordaccess generate** in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If **passwordaccess** is set to **generate** and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If **passwordaccess** is set to **prompt** and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the requirement to prompt.

If **passwordaccess** is set to **prompt** and you do not specify a password on the command line, then you are prompted for a password.

/BUFFers=numbuffers, buffersize
Specifies the number and size of data buffers that transfer data between the Domino server and the Tivoli Storage Manager API. Increasing the number or size (or both) of the data buffers can improve throughput.

You can specify 2 - 8 buffers. The default value is 3. The size of the buffers can be 64 - 8192 kb. The default value is 1024.

If the `/buffers` parameter is not specified on the command line or defined in the preferences file, Data Protection for IBM Domino uses the default values.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino, use the `/logfile` parameter to specify a different log file for each instance. Logging is collected for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/MOUNTWait=Yes|No
If the Tivoli Storage Manager server is configured to store transaction log backup data on removable media, then the Tivoli Storage Manager server can indicate to Data Protection for IBM Domino that it is waiting for a required storage volume to be mounted. In this case, you can specify whether Data Protection for IBM Domino waits for the media mount or stops the current operation. Removable media is media such as tapes.

You can specify:

Yes Wait for tape mounts. This value is the default.

No Do not wait for tape mounts.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/SESSions=numsessions|1
Specifies the number of Tivoli Storage Manager server sessions to be used by Data Protection for IBM Domino. You can specify 1 - 64 sessions. The default value is 1.

/SUBDir=No|Yes
Specifies whether subdirectories within the specified file path are searched for databases that match the file pattern. If this option is not specified, Data Protection for IBM Domino uses the value of the **/subdir** parameter in the Data Protection for IBM Domino preferences file.

You can specify:

No Do not search the subdirectories within the specified file path for databases that match the file pattern. This value is the default unless reset in the Data Protection for IBM Domino preferences file.

Yes Search the subdirectories within the specified file path for databases that match the file pattern.

Example

This example shows a backup of all databases that are changed. Changes to any database IDs for logged databases, and data changes for non-logged databases are backed up. This example shows deactivated database backups that refer to databases that no longer exist on the Domino server, or ones that are excluded.

```
domdsmc incremental * /subdir=yes

Output example:
```

```
Starting Domino database backup...
Initializing Domino connection...
Querying Domino for a list of databases, please wait...
Querying Domino for a list of databases, please wait...

Backing up database mail\hvargas.nsf, 1 of 1.
Full: 0   Read: 25,690,112   Written: 8,379,837   Rate: 2,707.06 Kb/Sec
Backup of mail\hvargas.nsf   completed successfully.


Total Domino databases inspected:    1
Total Domino databases backed up:    1
Total Domino databases excluded:   0
Total Domino backup objects expired:    0
Total Domino databases deduplicated:    1

Throughput rate:   2,678.70 Kb/Sec
Total bytes inspected:   25,690,112
Total bytes transferred:  8,379,837
Total LanFree bytes transferred:  0
Total bytes before deduplication: 25,690,112
Total bytes after deduplication:   16,359,696
Data compressed by:    49%
Deduplication reduction:  36.32%
Total data reduction ratio:  67.39%
Elapsed processing time:  3.06 Secs
```

Domdsmc query adsmserver

How to use the **Domdsmc query adsmserver** command is described.

Purpose

Use this command to provide the following information about the Tivoli Storage Manager server:

- Tivoli Storage Manager server name.
- Tivoli Storage Manager server level.
- Tivoli Storage Manager server platform.
- Tivoli Storage Manager nodename of the server
- NetWork host name of the server
- Options in effect at the server that affects this node (for example, management class information)

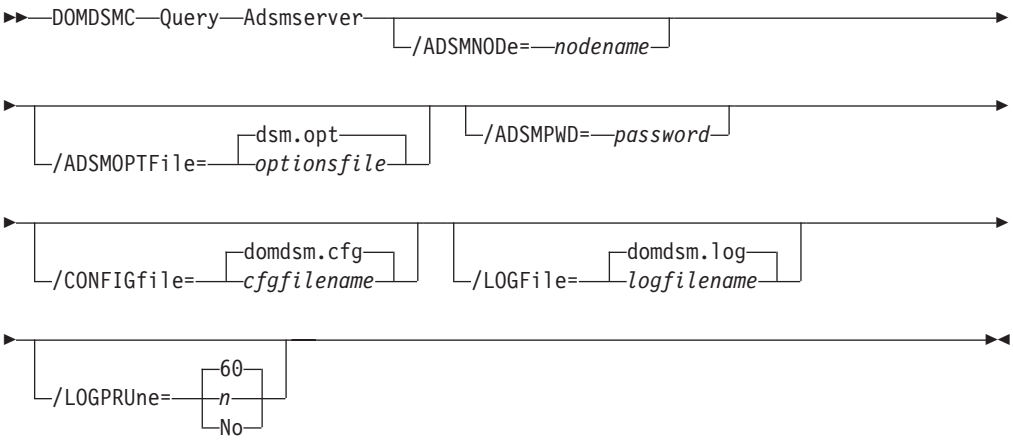

Parameters

/ADSMNODe=nodename
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify **passwordaccess generate** in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If **passwordaccess** is set to **generate** and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If **passwordaccess** is set to **prompt** and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the requirement to prompt.

If **passwordaccess** is set to **prompt** and you do not specify a password on the command line, then you are prompted for a password.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

Example

The following example queries your Tivoli Storage Manager server:

```
domdsmc query adsmserver
```

Output example:

```
Tivoli Storage Manager Server Connection Information
---------------------------------------------------

Nodename ............................... NODE1
NetWork Host Name of Server ............ 172.16.15.46
TSM API Version ........................ Version 7, Release 1, Level 0.54

Server Name ............................ TEST_TSM
Server Type ............................ Windows
Server Version ......................... Version 7, Release 1, Level 0.0
Compression Mode ....................... Client Determined
Domain Name ............................ STANDARD
Active Policy Set ...................... STANDARD
Default Management Class ............... STANDARD
```

Domdsmc query dbbackup

How to use the **Domdsmc query dbbackup** command is described.

Purpose

This command displays a list of database backups that are stored on the Tivoli Storage Manager server that match the dbname pattern. Active and inactive objects can be displayed. By default, only the active backup objects are displayed. To include inactive backup versions in the list, use the **/inactive** parameter.

The following information is provided:

- Database title.
- Database relative path name.
- Database size.
- Database backup date and time.
- Domino server name.
- Whether the backup is active or inactive.
- Whether the database is logged or not.
- Whether the database is encrypted.
- Whether the database is compressed.
- Whether the database is deduplicated.

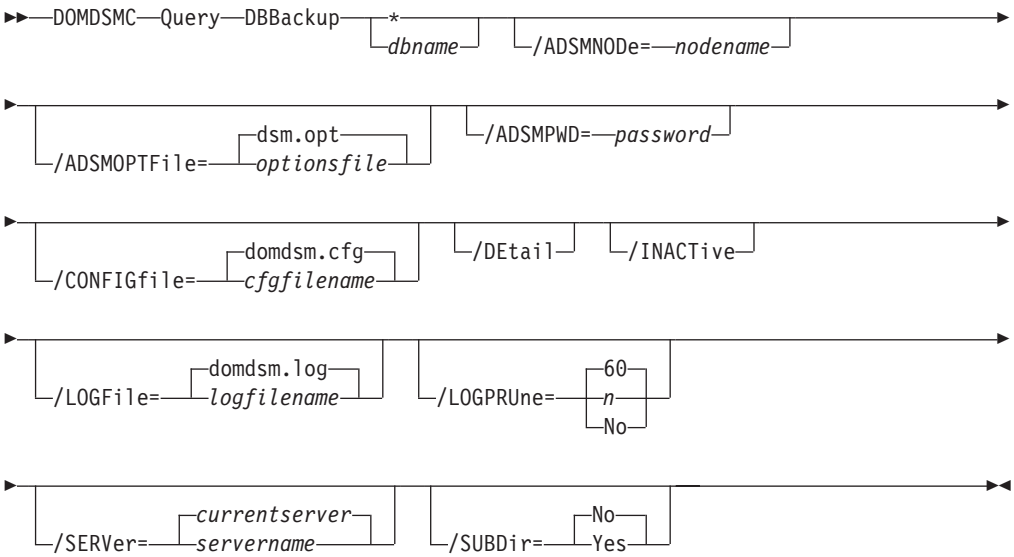

Parameters

*** | dbname**

Specifies the file path of a database or file path pattern. The file path pattern can represent a group of databases. You can also specify a group of databases by using the wildcard character asterisk (*).

The file path must be relative to the Notes data directory.

Symbolic links are referred to by their symbolic names. To reference a database in a directory pointed to by a directory link in the data path, use the directory link name as the directory name. For example, if database xyz.nsf is in a directory, pointed to by the link vol1.dir, refer to it as

vol1\xyz.nsf. If a symbolic directory link is created with the same name as a physical directory in the Notes data path, only the physical directory is searched.

The wildcard character (*) is used to represent any number of any characters. For example:

```
domdsmc query dbbackup abc*
```

This example lists all databases that begin with the characters abc. When used with the **/subdir** parameter, all databases within all subdirectories are listed.

/ADSMNODe=nodename
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is dsm.opt.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify **passwordaccess generate** in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If **passwordaccess** is set to **generate** and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If **passwordaccess** is set to **prompt** and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the need to prompt.

If **passwordaccess** is set to **prompt** and you do not specify a password on the command line, then you are prompted for a password.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is domdsm.cfg.

/DEtail
Displays information about the backup, such as whether it is encrypted, compressed, or deduplicated.

/INACTive
Specifies that both active and inactive backup objects are displayed. The default is to display only the active backup objects.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file.

You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/SERVer=currentserver|servername

Specifies the Domino server name. If not specified, Data Protection for IBM Domino uses the current Domino server.

/SUBDir=No|Yes

Specifies whether subdirectories within the specified file path are searched for databases that match the file pattern. If this option is not specified, Data Protection for IBM Domino uses the value of the **/subdir** parameter in the Data Protection for IBM Domino preferences file.

You can specify:

No Do not search the subdirectories within the specified file path for databases that match the file pattern. `No` is the default unless it is reset in the Data Protection for IBM Domino preferences file.

Yes Search the subdirectories within the specified file path for databases that match the file pattern.

Examples

Example 1: This example displays information for all the active database backups that are stored on the local Tivoli Storage Manager server.

```
domdsmc query dbbackup *
```

Output example (only a portion of the actual output is shown):

```
                          Database Backup List
                          --------------------

    Domino Server: chilly
    --------------------

   DB Backup Date    Size     A/I Logged Database Title  Database File
------------------- ---------  --- ------ -------------- -------------
01/22/2008 14:29:42  819.00KB   A   Yes   Administration  admin4.nsf
01/22/2008 14:29:44  501.50KB   A   No    Administration  admin4.ntf
01/22/2008 14:29:46  384.00KB   A   Yes   Ja AgeRun       AgeRun.nsf
01/22/2008 14:29:47  153.50KB   A   No    Agent Log       alog4.ntf
01/22/2008 14:29:48  246.50KB   A   No    Archive Log     archlg80.ntf
01/22/2008 14:29:49  226.00KB   A   No    Billing         billing.ntf
01/22/2008 14:29:54  1226.50KB  A   No    Bookmarks       bookmark.ntf
01/22/2008 14:29:56  320.00KB   A   Yes   Local free time busytime.nsf
01/22/2008 14:29:58  143.00KB   A   No    Local free time busytime.ntf
01/22/2008 14:29:58  140.50KB   A   No    Local Document  cache.ntf
01/22/2008 14:29:59  1170.00KB  A   Yes   Catalog         catalog.nsf
01/22/2008 14:30:02  799.00KB   A   No    Catalog         catalog.ntf
01/22/2008 14:30:04  1896.00KB  A   No    Domino    Certi cca80.ntf
01/22/2008 14:30:08  159.00KB   A   No    Certification L certlog.ntf
```

Example 2: This example displays information for all the database backups that are stored on the local Tivoli Storage Manager server. The information includes inactive backup objects and subdirectories within the file path.

```
domdsmc query dbbackup * /inactive /subdir=yes
```

Output example (only a portion of the actual output is shown):

```
                          Database Backup List
                          --------------------

   Domino Server: chilly
   --------------------

  DB Backup Date     Size     A/I Logged Database Title Database File
------------------- --------- --- ------ -------------- -------------
01/22/2008 14:29:42 819.00KB  A   Yes    Administration admin4.nsf
01/07/2008 12:11:07 729.00KB  I   Yes    Administration admin4.nsf
01/07/2008 12:29:04 729.00KB  I   Yes    Administration admin4.nsf
01/07/2008 12:46:21 729.00KB  I   Yes    Administration admin4.nsf
01/19/2008 13:50:27 819.00KB  I   Yes    Administration admin4.nsf
01/22/2008 14:29:44 501.50KB  A   No     Administration admin4.ntf
01/07/2008 12:29:06 389.50KB  I   No     Administration admin4.ntf
01/07/2008 12:46:23 389.50KB  I   No     Administration admin4.ntf
01/19/2008 13:50:33 501.50KB  I   No     Administration admin4.ntf
01/19/2008 13:56:48 501.50KB  I   No     Administration admin4.ntf
01/22/2008 14:30:24 300.75MB  A   Yes    A new database data2\db1.nsf
01/07/2008 11:51:39 300.75MB  I   Yes    A new database data2\db1.nsf
01/07/2008 12:11:42 300.75MB  I   Yes    A new database data2\db1.nsf
01/07/2008 12:29:42 300.75MB  I   Yes    A new database data2\db1.nsf
01/07/2008 12:47:04 300.75MB  I   Yes    A new database data2\db1.nsf
```

Example 3: The following example queried the Tivoli Storage Manager server and included the **/adsmpwd** parameter:

```
domdsmc q dbb * /adsmpwd=neil
```

Output example:

```
                  Database Backup List
                  --------------------

   Domino Server: Server1
   ----------------------

  DB Backup Date     Size     A/I  Logged Database Title Database File
------------------- ---------  --- ------ -------------- -------------
01/16/2008 11:14:19 1019.50KB  A   Yes    db1            db1.nsf
01/16/2008 10:56:28 1019.50KB  A   Yes    db2            db2.nsf
01/16/2008 10:56:29 1019.50KB  A   Yes    db3            db3.nsf
01/16/2008 10:56:30 1170.00KB  A   Yes    newdb          dblink.nsf
01/16/2008 10:56:31 1019.50KB  A   Yes    SERVER1 Mailbox mail.box
```

Example 4: This example displays detailed information for a specific database that is stored on the local Tivoli Storage Manager server.

```
domdsmc query dbbackup mynotes1.nsf /detail
```

Output example:

```
Database Backup List
--------------------


Backup Object Information
-------------------------

Domino Server Name ..................... DOMINOTESTSERVER
Database Title ......................... mynotes1
Database File .......................... mail\mynotes1.nsf
Database Backup Date ................... 07/22/2011 12:56:24
Database Size .......................... 24.50MB
Database Backup State .................. Active
Database Logged ........................ Yes
Database Compressed .................... Yes
Database Encryption Type ............... None
Database Client-deduplicated ........... Yes
```

Domdsmc query domino

How to use the **domdsmc query domino** command is described.

Purpose

This command displays general information and an optional list of databases on the local Domino server. If you do not specify a dbname pattern, then only general server information is displayed. If you specify a dbname pattern, then a list of databases on the Domino server that match the dbname pattern is displayed.

The information that is provided is listed:

- Domino server name.
- Domino server level.
- Domino server build.
- Logging type in effect.

- When the Domino server is enabled for DB2, the DB2 enabled status and the name of the Domino DB2 database are provided.
- Optionally lists current databases with their specific details, database title, and relative path name.

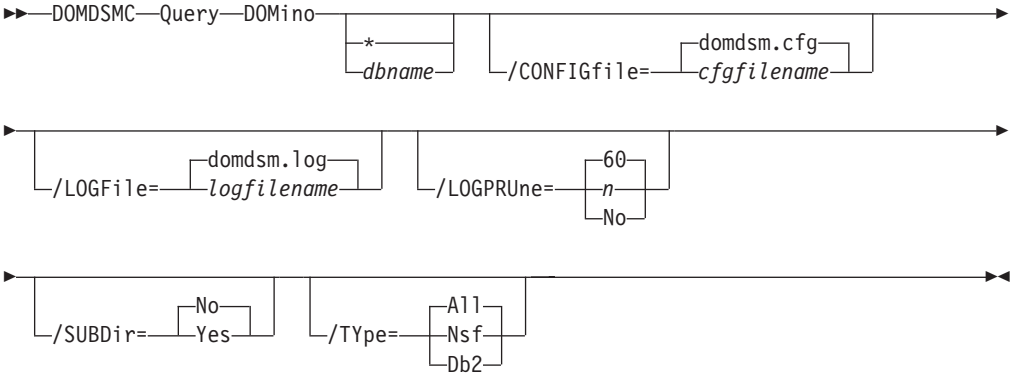

Parameters

*** | dbname**

Specifies the file path of a database or file path pattern. The file path pattern can represent a group of databases.

The file path must be relative to the Notes data directory.

Symbolic links are referred to by their symbolic names. To reference a database in a directory pointed to by a directory link in the data path, use the directory link name as the directory name. For example, if database `xyz.nsf` is in a directory, pointed to by the link `vol1.dir`, refer to it as `vol1\xyz.nsf`. If a symbolic directory link is created with the same name as a physical directory in the Notes data path, only the physical directory is searched.

The wildcard character (*) is used to represent any number of any characters. For example, the following command reveals domino server information and lists all databases:

```
domdsmc query domino abc*
```

This command lists all databases that begin with the characters `abc` in the Notes data directory. When used with the **query dbbackup** command, a list of all database backups that are stored on the Tivoli Storage Manager server is provided. When used with the **/subdir** parameter, all databases within all subdirectories are listed.

Note: "All" databases on a Domino server are defined to mean all databases within the Notes data directory or symbolically linked to the Notes data directory. This means that databases with nonstandard file extensions are not included. Databases with a file extension `.nsf` are standard file extensions. Templates have a standard file extension of `.ntf` and are included.

/CONFIGfile=cfgfilename

Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/SUBDir=No|Yes

Specifies whether subdirectories within the specified file path are searched for databases that match the file pattern. If this option is not specified, Data Protection for IBM Domino uses the value of the **/subdir** parameter in the Data Protection for IBM Domino preferences file.

You can specify:

No Do not search the subdirectories within the specified file path for

databases that match the file pattern. No is the default unless it is reset in the Data Protection for IBM Domino preferences file.

Yes Search the subdirectories within the specified file path for databases that match the file pattern.

/TYpe=All|Nsf|Db2
Specifies the type of database to display.

You can specify:

All Specifies that both Domino NSF and Domino DB2 enabled Notes databases are displayed. This value is the default.

Nsf Specifies that only Domino NSF databases are displayed.

Db2 Specifies that only Domino DB2 enabled Notes databases are displayed.

Examples

Example 1: This command example shows how to list all the databases in the Notes data directory on the Domino server that match the dbname pattern xyz.

```
domdsmc query domino xyz*
```

Output example (only a portion of the actual output is shown):

```
Domino Server Information
-------------------------

  Domino Server Name:  Windows2003
  Domino Server Level: 8.5.3.0
  Domino Server Build: 390
  Logging:             Archival

Domino NSF Database Information
-------------------------------

  Last Modified Date         Size    Logged   Database Title        Database
  ------------------         ----    ------   --------------        --------
  08/29/2011 05:05:08       648.00KB   Yes    Domino xyz License   userlicenses.ntf
  08/29/2011 05:05:11       448.00KB   Yes    xyz Registration Qu  userreg.ntf
```

Example 2: The following example queried the Tivoli Storage Manager server and listed databases on the Domino server with the wildcard character *:

```
domdsmc q dom *
```

Output example (only a portion of the actual output is shown):

```
Domino Server Information
-------------------------

  Domino Server Name:  abraxis
  Domino Server Level: 7.0
  Domino Server Build: 166
  Logging:             No Logging

Domino Database Information
---------------------------

  Last Modified Date         Size    Logged   Database Title        Database
  ------------------         ----    ------   --------------       --------
  09/09/2007 12:18:12       320.00KB   No     Java AgentRunner      AgentRunner.nsf
  10/13/2007 13:34:38        75.00KB   No     Agent Log             alog4.ntf
  04/19/2007 07:08:51       144.00KB   No     Archive Log (R5.0)    archlg50.ntf
  06/19/2007 00:00:18     1,876.00KB   No                           bookmark.nsf
  04/04/2007 05:37:10     1,250.00KB   No     Bookmarks             bookmark.ntf
  10/13/2007 13:35:06        64.00KB   No     Local free time info  busytime.ntf
  05/21/2007 07:06:16        61.00KB   No     Local Document Cache  cache.ntf
  06/07/2007 10:00:55     1,750.00KB   No     Domino R5 Certificat  cca50.ntf
  05/21/2007 06:58:22        91.00KB   No     Cluster Analysis      clusta4.ntf
  05/31/2007 11:19:28       821.00KB   No     Server Certificate A  csrv50.ntf
  03/06/2007 07:52:11       111.00KB   No     Database Analysis     dba4.ntf
  10/13/2007 13:34:49        95.00KB   No     Database Library      dblib4.ntf
  03/19/2007 05:23:05       130.00KB   No     Decommission Server   decomsrv.ntf
  01/17/2007 09:49:12       867.00KB   No     Discussion - Notes &  discsw50.ntf
  06/19/2007 11:43:08       455.00KB   No     Microsoft Office Lib  doclbm46.ntf
  12/29/2007 06:20:07       506.00KB   No     Microsoft Office Lib  doclbm50.ntf
  01/04/2007 08:40:47       503.00KB   No     Lotus SmartSuite Lib  doclbs5.ntf
  12/13/2007 07:51:50       550.00KB   No     Doc Library - Notes   doclbw50.ntf
  08/31/2007 14:06:25       400.00KB   No     Document Library (R4  doclib4.ntf
  08/08/2007 07:41:46       440.00KB   No     Lotus SmartSuite 96   doclibl4.ntf
  02/12/2007 08:43:13     2,816.00KB   No     Domino Administrator  domadmin.nsf
```

Example 3: The following example queries (and displays) the DB2 enabled Notes databases available on the Domino server with the wildcard character *:

```
domdsmc query domino * /type=db2
```

Output example:

```
Domino Server Information
-------------------------

  Domino Server Name:  domino7
  Domino Server Level: 7.0.0.0
  Domino Server Build: 259
  Logging:             Archival
  DB2 Enabled:         Yes
  DB2 Database Name:   DOMINO7


Domino NSF DB2 Database Information
-----------------------------------

  Class Name: class0
  ------------

  Last Modified Date         Size     Database Title Group    Database
  ------------------         ----     -------------- -------- --------
  01/20/2008 02:00:19       111.00KB  db2 nsf 1      GRP4     cb2nsf1.nsf
  01/21/2008 02:00:25       110.00KB  db2 nsf 1      GRP4     db2f.nsf
  01/20/2008 02:00:36       114.00KB  db2 nsf 2      GRP4     pb2nsf2.nsf

  Class Name: class1
  ------------

  Last Modified Date         Size     Database Title Group    Database
  ------------------         ----     -------------- -------- --------
  01/21/2008 02:00:24       107.00KB  db2 nsf 1      GRP1     db1.nsf
  01/21/2008 02:00:24       107.00KB  db2 nsf 2      GRP1     db2a.nsf
  01/21/2008 02:00:24       110.00KB  db2 nsf 1      GRP1     db2c.nsf
  01/21/2008 02:00:25       107.00KB  db2 nsf 1      GRP1     db2e.nsf
  01/20/2008 02:00:17       112.00KB  db2 nsf 1      GRP2     ab2nsf1.nsf
  01/21/2008 02:00:24       107.00KB  db2 nsf 2      GRP2     db2b.nsf
  01/21/2008 02:00:25       107.00KB  db2 nsf 1      GRP2     db2g.nsf
  01/20/2008 02:00:37       113.00KB  db2 nsf 1      GRP2     xb2nsf1.nsf

  Class Name: class2
  ------------

  Last Modified Date         Size     Database Title Group    Database
  ------------------         ----     -------------- -------- --------
  01/20/2008 02:00:18       112.00KB  db2 nsf 2      GRP3     bb2nsf2.nsf
  01/21/2008 02:00:25       107.00KB  db2 nsf 1      GRP3     db2d.nsf
  01/21/2008 02:00:26       107.00KB  db2 nsf 1      GRP3     dbi.nsf
  01/20/2008 02:00:36       114.00KB  db2 nsf 1      GRP3     ob2nsf1.nsf
```

Domdsmc query logarchive

How to use the **Domdsmc query logarchive** command is described.

Purpose

This command displays a list of the archived transaction log extents that are stored on the Tivoli Storage Manager server. By default, only the active log extents are listed. To display inactive extents, use the **/inactive** parameter.

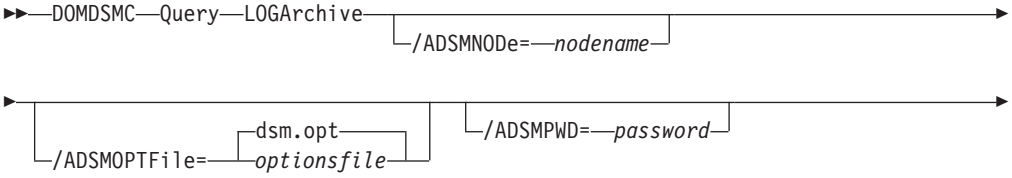

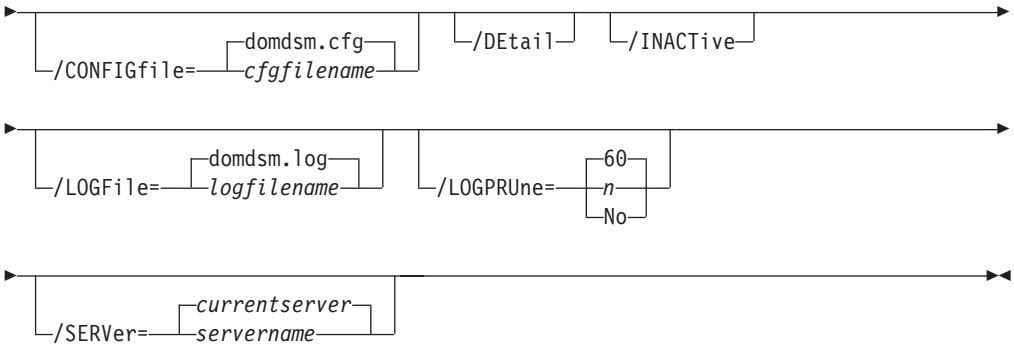

Parameters

/ADSMNODe=nodename

Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile

Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password

Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify **passwordaccess generate** in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If **passwordaccess** is set to **generate** and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If **passwordaccess** is set to **prompt** and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the requirement to prompt.

If **passwordaccess** is set to **prompt** and you do not specify a password on the command line, then you are prompted for a password.

/CONFIGfile=cfgfilename

Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/DEtail

Displays information about the backup, such as whether it is encrypted, compressed, or deduplicated.

/INACTive

Specifies that both active and inactive backup objects are displayed. The default is to display only the active backup objects.

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. Logging is directed for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/SERVer=currentserver|servername

Specifies the Domino server name. If not specified, Data Protection for IBM Domino uses the current Domino server.

Examples

`Example 1:` The following example displays the list of archived log extents that are stored on the Tivoli Storage Manager server.

```
domdsmc query logarchive
```

Output example:

```
   Domino Server: chilly
   --------------

   Logger Id: OF8525679F:004266F1-ON000003EC:BD9D27DB
   ----------

                        Transaction
  Log Archive Date      Log Filename    A/I     Size
--------------------    ------------    ---    ------
01/22/2013 16:52:15     S0000001.TXN     A       64MB
```

Example 2: The following example displays the list of archived log extents that are stored on the Tivoli Storage Manager server, including inactive backup objects. This example uses a Tivoli Storage Manager options file named `aserver.opt`.

```
domdsmc query logarchive /inactive /adsmoptfile=aserver.opt
```

Output example:

```
   Domino Server: chilly
   --------------

   Logger Id: OF8525679F:004266F1-ON000003EC:BD9D27DB
   ----------

                        Transaction
  Log Archive Date      Log Filename    A/I     Size
--------------------    ------------    ---    ------
01/23/2013 09:43:16     S0000003.TXN     A       64MB
01/23/2013 09:41:22     S0000002.TXN     A       64MB
01/22/2013 16:52:15     S0000001.TXN     I       64MB
```

Example 3: The following example queried the Tivoli Storage Manager server and included the **/adsmpwd** parameter:

```
domdsmc q loga /adsmpwd=neil
```

Output example:

```
ACD5819I There are no archived logs for the server named SnailTrail.
```

Example 4: The example displays detailed information for the list of archived log extents that are stored on the Tivoli Storage Manager server:

```
domdsmc query logarchive /detail
```

Output example:

```
Backup Object Information
-------------------------

Domino Server Name .................. DOMINOTESTSERVER
Logger Id ........................... OF09A666BF:AACF9914-ON00000302:A048A61F
Log Archive File .................... S0000035.TXN
Log Archive Date .................... 07/22/2013 12:52:37
Log Archive Size .................... 64.00MB
Log Archive State ................... Active
Log Archive Compressed .............. Yes
Log Archive Encryption Type ......... None
Log Archive Client-deduplicated ..... Yes
```

Domdsmc query pendingdbs

How to use the **domdsmc query pendingdbs** command is described.

Purpose

This command displays a list of all the databases that are restored, but not yet activated.

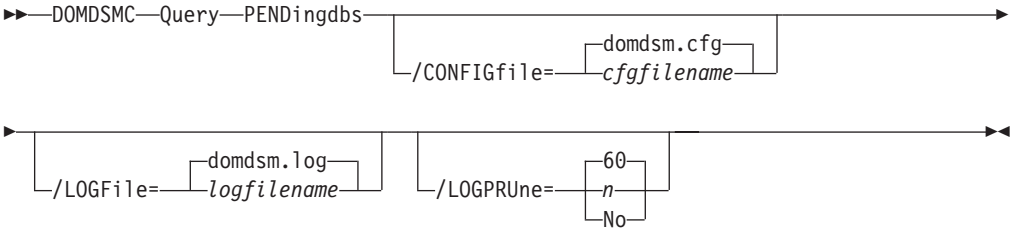

Parameters

/CONFIGfile=cfgfilename

Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file is located in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and run daily. You can use the **set** command to do the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. When the value of `/logprune` is a number, the prune is run even if one has already been run for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or **logfile** setting.

Example

The following example command shows how to list all the pending databases on the Domino server:

```
domdsmc query pendingdbs
```

```
                          Pending Database List
                          ---------------------

   Domino Server: chilly
   ---------------

 Backup Time Stamp  Size     A/I Logged Database Title Database File
------------------- -------- --- ------ -------- ----- -------------
01/22/2008 14:36:50 896.00KB A    No    Vacat    Plan  datar3\db8.nsf
```

Domdsmc query preferences

The **domdsmc query preferences** command displays a list of the current values that are set in the preferences file for Data Protection for IBM Domino.

Purpose

You can view a list of the current values that are set in the preferences file for Data Protection for IBM Domino by running the **domdsmc query preferences** command

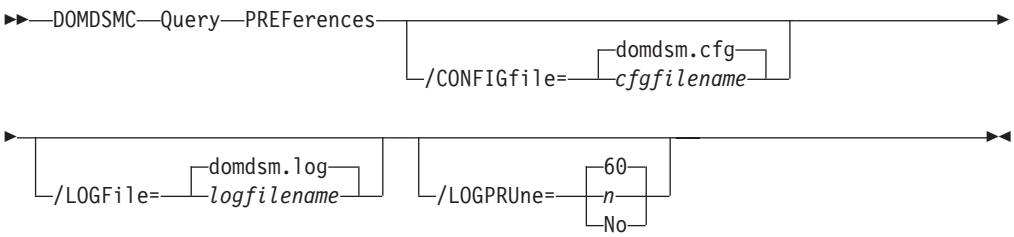

Parameters

/CONFIGfile=cfgfilename

Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file is stored in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino for operations, use the **/logfile** parameter to specify a different log file for each instance used. This parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and run daily. Use the **set** command to do the following actions:

- Change the defaults so that log pruning is disabled.
- Change the number of days log entries are saved.

You can use the `/logprune` option to override these defaults for one command run. When the value of **/logprune** is a number, the prune is run even if one has completed for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. This is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the *timeformat* or *dateformat* parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the *timeformat* or *dateformat* parameter changes, run one of the following actions to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the */logfile* parameter or *logfile* setting.

Example

This command shows the current values that are stored in the default preferences file for Data Protection for IBM Domino.

```
domdsmc query preferences

Output example:
```

```
Data Protection for Domino Preferences
------------------------------

ADSMLOGDir .....................
ADSMOPTFile ....................
BUFFers ........................ 3
BUFFERSIze ..................... 1024
DATEformat .....................
LANGUAGE .......................
LOGFile ........................ domdsm.log
LOGPRUne ....................... 60
MOUNTWait ...................... Yes
NOTESInipath ...................
NUMberformat ................... 1
REPlace ........................ Yes
SESSions ....................... 1
STATistics ..................... No
SUBDir ......................... No
TIMEformat ..................... 1
DOMTXNBYTElimit................. 0
DOMTXNGROUPmaX.................. 2
COMMRESTARTDuration ............ 60
COMMRESTARTInterval ............ 15

DB2ALTDBname ................... DOM_ALT
DB2CONTainerpath ...............
DB2LOGPath .....................
DB2LOGTarget ...................
DB2REPlace ..................... Yes
DB2RESTIntopath ................
DB2SESSions .................... 1
DB2USER ........................
```

Domdsmc resetdatabase

How to use the **domdsmc resetdatabase** command is described.

Purpose

Running the **domdsmc resetdatabase** command resets a Domino server database that is incomplete because of an unexpected termination during a Data Protection for IBM Domino backup. An unexpected termination during a Data Protection for IBM Domino backup can be caused by the following actions:

- A program check or segment violation.
- A command to end the process that is entered through the task manager window.
- Exiting a debugging program.

Subsequent attempts to back up such a database fails until the **resetdatabase** command is issued.

You can specify one or more Domino databases to be reset. You must specify the database name. The **resetdatabase** command does not accept wildcard characters.

For more information about disaster recovery procedures, see “Recovering from loss of Domino transaction logs for NSF databases” on page 166.

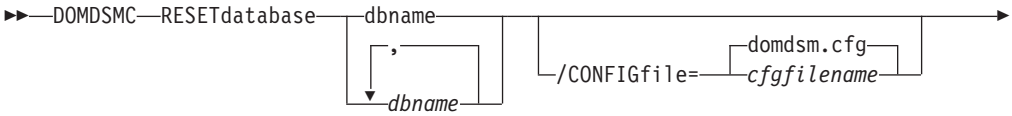

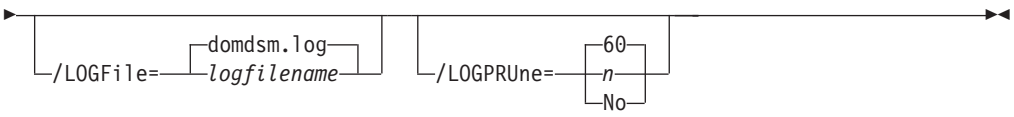

Parameters

dbname *dbname,dbname,...,*
Specifies the database to be reset. Multiple *dbnames* can be specified separated with commas.

/CONFIGfile=*cfgfilename*
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file is stored in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=*logfilename*
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the **/logfile** parameter to specify a different log file for each instance used. Logging is directed for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=*60*|*n*|*No*
Specifies whether to prune log entries. By default, log pruning is enabled and run daily. Use the **set** command to do the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the **/logprune** option to override these defaults for one command run. When the value of **/logprune** is a number, the prune is run even if one has already been run for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the *timeformat* or *dateformat* parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the *timeformat* or *dateformat* parameter changes, do one of the following to prevent undesired pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or **logfile** setting.

Example

The following example resets the database `testdata.nsf`:

```
domdsmc resetdatabase testdata.nsf
```

`Output example:`

```
Database testdata.nsf successfully reset.
```

Domdsmc restore

How to use the **domdsmc restore** command is described.

Purpose

Running the **domdsmc restore** command restores a single database or a group of databases from Tivoli Storage Manager storage to the Domino server. If you are planning to apply transaction logs to the restored databases to get a more current state, use the **/activate**=*no* parameter. You can apply transaction logs by running the **activatedbs** command.

Note: If you receive the error message ACD5223E, you must check the permissions of the directory where the `<name of the Tivoli Storage Manager server>.pdb` file is created.

The `.pdb` file is created in the directory indicated by the value of the following registry entry: `HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\domclient\Path`

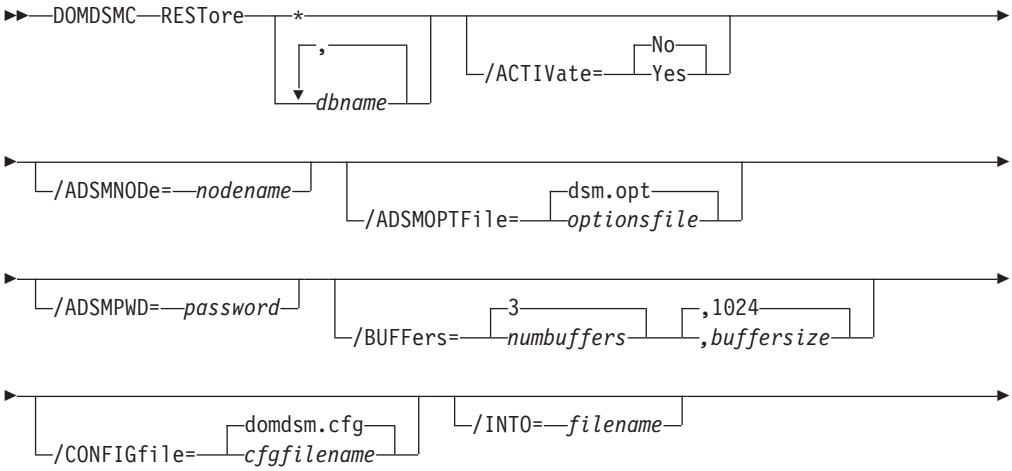

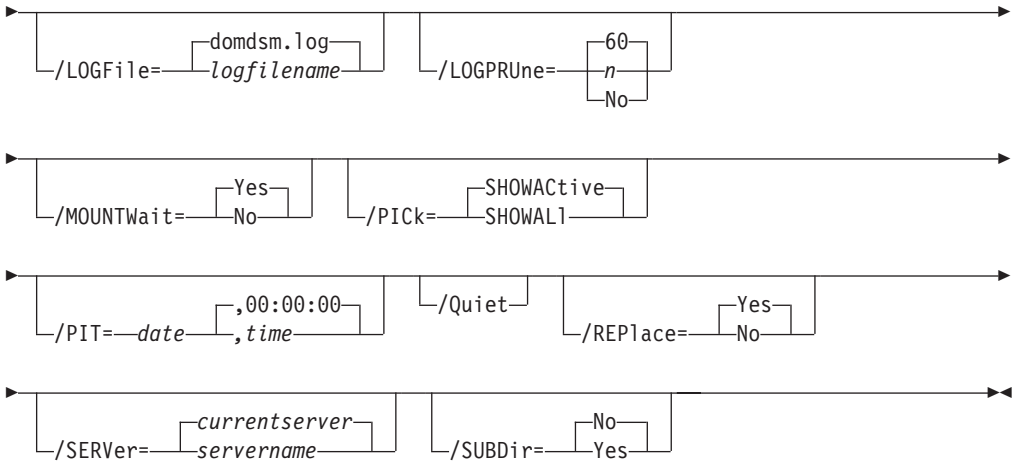

Parameters

*|*dbname,...,dbname*

Specifies the file path of a database or file path pattern for a group of databases. The file path pattern can represent a group of databases to be restored from the Tivoli Storage Manager server. The wildcard character asterisk (*) is used to specify a group of databases when used in the *dbname*. Multiple *dbnames* can be specified when they are separated with commas.

The file path must be relative to the Notes data directory.

Symbolic links are referred to by their symbolic names. For example, if `mydata.dir` is a directory link in the Notes data directory that points to `x:\data`, database `mydb.nsf` in the `x:\data` directory would be named `mydata\mydb.nsf`. The physical file path for the relative name is resolved according to the symbolic values at the time of the restore.

If a symbolic link used in the name of a database backup image does not exist, the restore must be done with the **/into** parameter. This parameter specifies where the database is placed.

The wildcard character (*) can be used in the file name portion of the file path. The wildcard character is not supported within directory names. The * is used to represent any number of any characters. For example, the following command restores the active backup of all databases that begin with the characters **ter**:

```
domdsmc restore ter*
```

For example, the following command lists all active database backups on the Tivoli Storage Manager server so that you can select which ones you want to restore:

```
domdsmc restore * /pick
```

Note: The value of the **/subdir** parameter determines whether only the specified directory or all subdirectories are searched for databases that match the file pattern.

There is no default for **dbname**.

/ACTIVate=*No*|*Yes*

Specifies whether the databases that are being restored are to be brought

online. If the restored database is to be rolled forward to a more current state by applying transaction logs, then **/activate**=*no* must be specified so that transaction logs can be applied with the **activatedbs** command.

Because there is a single transaction log for all logged databases, all databases must be activated together by running one command. This prevents the fetching of the same transaction logs multiple times from the Tivoli Storage Manager server. The databases can be restored separately by specifying **/activate**=*no*. The databases can then be activated together with a single **activatedbs** command.

If the **/activate** parameter is not specified, **/activate**=*no* is the default value.

No Do not activate the database. This value is the default.

Yes Activate the database.

/ADSMNODe=nodename
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify **passwordaccess generate** in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If **passwordaccess** is set to **generate** and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If **passwordaccess** is set to **prompt** and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the requirement to prompt.

If **passwordaccess** is set to **prompt** and you do not specify a password on the command line, then you are prompted for a password.

/BUFFers=numbuffers, buffersize
Specifies the number and size of data buffers that transfer data between the Domino server and the Tivoli Storage Manager API. Increasing the number or size (or both) of the data buffers can improve throughput.

You can specify 2 - 8 buffers. The default value is 3. The size of the buffers can be 64 - 8192 kb. The default value is 1024.

If the `/buffers` parameter is not specified on the command line or defined in the preferences file, Data Protection for IBM Domino uses the default values.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a

path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/INTO=*filepath*

Specifies the file path and file name to be used for the restored database. The file path that is specified must be relative to the Notes data directory or can be a fully qualified physical path. If a relative file path is specified, symbolic names can be included when the symbolic links exist to resolve the names. The specified path is considered a physical file path if it begins with a directory delimiter or a drive letter that is followed by a colon.

If multiple databases are being restored at one time, the file name must be specified as a pattern with a single equal sign, =, represents the entire file name and extension of the database backup. For example, the following command restores all backups from the `vola` directory into the `tempvol` directory with the same file names:

```
domdsmc restore vola\* /into=tempvol\=
```

For example, the following command restores all backups from the `vola` directory into the `vola` directory with the file names from the backup version that is prefixed with a `t`:

```
domdsmc restore vola\* /into=vola\t=
```

If you entered `domdsmc restore vola\* /into=vola\=xyz /activate=yes`, `xyz` is appended to the database suffix. For example, a database that is called `abc.nsf` is restored as `abc.nsfxyz`.

If you restore without running **activatedbs**, `.dad` is appended to the suffix of the database name. When you run **activatedbs** or select the ACTIVATE tab on the GUI, the `.dad` append is removed from the suffix of the database name.

If multiple databases in a subdirectory branch are being restored and you must preserve the directory structure, the file name must be specified as a pattern with two equal signs, ==, representing the filepath of the database backup. For example, the following command restores all backups from the `vola` directory and its subdirectories into the `tempvol` directory with the same file names and directory structure. The == is replaced by the full relative path for each database file that is restored, including the `vola` directory:

```
domdsmc restore vola\* /subdir=yes /into=tempvol\==
```

Note:

- When the **/into** parameter is used with the **restore** command, replication is disabled for the restored databases.
- If the **/into** parameter is not used, replication settings remain as they were in the backup version that is restored.

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.

The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/MOUNTWait=Yes|No

If the Tivoli Storage Manager server is configured to store transaction log backup data on removable media, then the Tivoli Storage Manager server can indicate to Data Protection for IBM Domino that it is waiting for a required storage volume to be mounted. In this case, you can specify whether Data Protection for IBM Domino waits for the media mount or stops the current operation. Removable media is media such as tapes.

You can specify:

Yes Wait for tape mounts. This value is the default.

No Do not wait for tape mounts.

/PICk=*SHOWACtive*|*SHOWALl*

Displays a list of database backups that match the `dbname` pattern that can be selected for restore. The pick list is displayed as a scrollable list from which you can select the database backups for restore.

You can specify:

SHOWACtive
Displays a list of active database backup versions.

SHOWALl
Displays a list of both active and inactive database backup versions. All the backup versions that match the **dbname** pattern are shown.

/PIT=*currentdate,currenttime*|*date,time*
Specifies a point in time when the specified databases are restored. The date and time values must be specified in the same date and time format that is defined in the Data Protection for IBM Domino preferences file. The most recent database backup images that are taken before the specified point in time are restored. Deleted backup images are not restored. Logged databases can then be rolled forward to that point by specifying the same date and time values on the */applylogs* option of the **activatedbs** command.

date Specify a date string in the active date format. If you do not specify a date, the specified databases are restored unless the **/pick** parameter was used to select inactive backup versions.

The date must be specified with the same date format that is defined in the Data Protection for IBM Domino preferences file. See "Domdsmc set" on page 115 for a list of available date formats.

time Specify a time string in the active time format. If you specify a date without the time, HH:MM:SS on a 24-hour clock is used.

The time must be specified with the same time format that is defined in the Data Protection for IBM Domino preferences file. See "Domdsmc set" on page 115 for a list of available time formats.

Note: If this parameter is used with the **/pick** parameter, the *showactive* and *showall* variables for the **/pick** parameter are ignored. The pick list contains the database backup images that meet the **/PIT** criteria.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/REPlace=*Yes*|*No*
Specifies whether to replace existing databases on the target system.

You can specify:

Yes Allows an existing database on the target system to be replaced during the restore process.

No Prevents an existing database on the target system from being overwritten during the restore process.

/SERVer=*currentserver*|*servername*
Specifies the Domino server name. If not specified, Data Protection for IBM Domino uses the current Domino server.

/SUBDir=No|Yes
Specifies whether subdirectories within the specified file path are searched for databases that match the file pattern. If this option is not specified, Data Protection for IBM Domino uses the value of the **/subdir** parameter in the Data Protection for IBM Domino preferences file.

You can specify:

No Do not search the subdirectories within the specified file path for

databases that match the file pattern. This value is the default unless reset in the Data Protection for IBM Domino preferences file.

Yes Search the subdirectories within the specified file path for databases that match the file pattern.

Examples

Example 1: This example restores all your databases and subdirectories.

```
domdsmc restore * /subdir=yes
```

Example 2: The following example restores a database to the specified date and time.

```
domdsmc restore datadir3\yyy.nsf /subdir=yes /pit=01/11/2004,10:00:00
```

Output example:

```
Starting Domino database restore...

Initializing Domino connection...
Logging on to the Tivoli Storage Manager server, please wait...
Querying Tivoli Storage Manager server for a list of database backups, please wait...


Restoring database datadir3\yyy.nsf, 1 of 1,
to /data/testdata1/notes1/notesdata/userlicenses.ntf.dad
Full: 0 Read: 663,552 Written: 663,552 Rate: 3,600.00 Kb/Sec
Restore of userlicenses.ntf completed successfully.


Total database backups inspected: 1
Total database backups requested for restore: 1
Total database backups restored: 1
Total database activated: 0

Throughput rate: 3,600.00 Kb/Sec
Total bytes transferred: 663,552
Total LanFree bytes transferred: 0
Elapsed processing time: 0.18 Secs
```

Example 3: The following example restores a database into the same directory but with a different name.

```
domdsmc restore a_dir\db1.nsf /into=a_dir\db8.nsf
```

Output example:

```
Starting Domino database restore...

Initializing Domino connection...
Logging on to the Tivoli Storage Manager server, please wait...
Querying Tivoli Storage Manager server for a list of database backups, please wa
it...

[0270:0002-0F1C] Clearing DBIID E03E8718 for DB
C:\Program Files\IBM\Lotus\Domino\data\a_dir\db8.nsf.dad
[0270:0002-0F1C] 09/29/2011 05:47:25 AM Recovery Manager: Assigning new DBIID for
C:\Program Files\IBM\Lotus\Domino\data\a_dir\db8.nsf.dad
(need new backup for media recovery).

Restoring database a_dir\db1.nsf, 1 of 1,
to C:\Program Files\IBM\Lotus\Domino\data\a_dir\db8.nsf.dad
Full: 0 Read: 327,680 Written: 327,680 Rate: 1,545.89 Kb/Sec
Restore of a_dir\db1.nsf completed successfully.


Total database backups inspected: 1
Total database backups requested for restore: 1
Total database backups restored: 1
Total database activated: 0

Throughput rate: 1,538.46 Kb/Sec
Total bytes transferred: 327,680
Total LanFree bytes transferred: 0
Elapsed processing time: 0.21 Secs
```

Domdsmc restorelogarchive

How to use the **domdsmc restorelogarchive** command is described.

Purpose

Domdsmc restorelogarchive restores archived transaction logs from Tivoli Storage Manager storage to the Domino server. This command assists with disaster recovery operations. By retrieving the most recent archived log file, it is possible to rebuild the Domino transaction log control file. In this way archived transaction log files can be used to recover restored database backups to a more current state, even after a loss of the active transaction log.

When you are restoring a transaction log file from an old Logger ID during an alternate server or alternate partition restore procedure, you must specify the **/pick** parameter with the **restorelogarchive** command to choose the log extent. For more information see "NSF databases restore to alternate server and alternate partition" on page 167.

For more information on disaster recovery procedures, see "Recovering from loss of Domino transaction logs for NSF databases" on page 166.

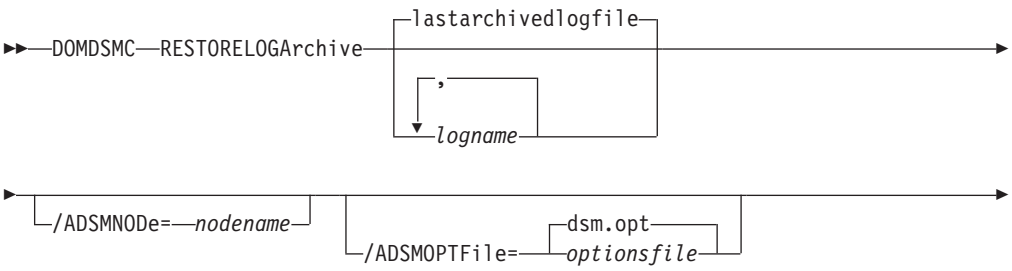

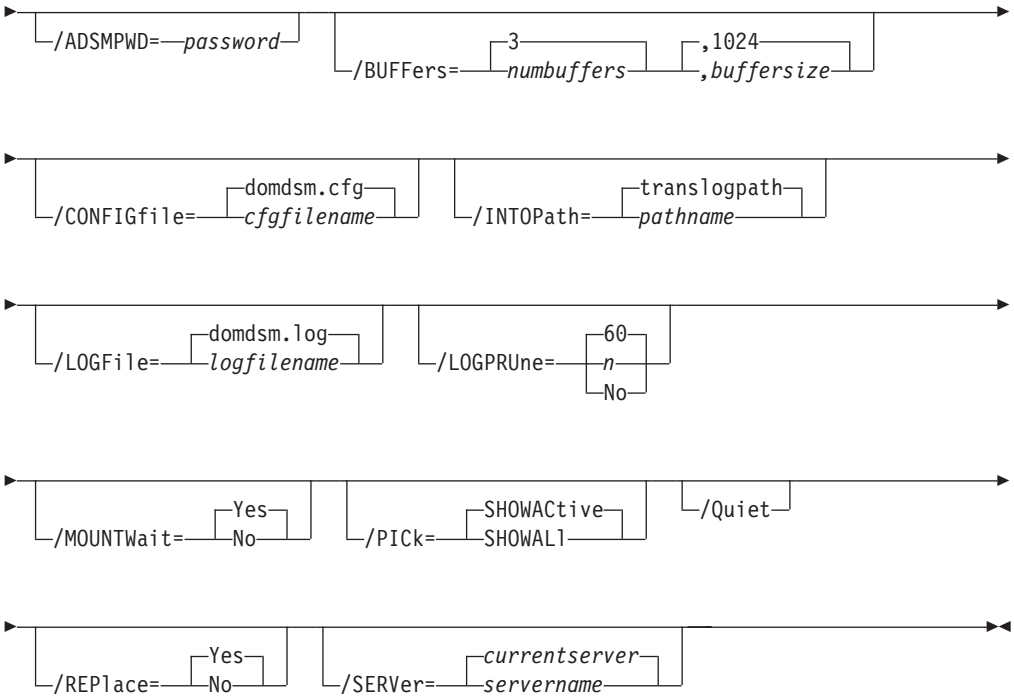

Parameters

logname,...,logname

The **logname** optional parameter specifies the logname of the archived transaction log to be restored. Multiple *lognames* can be specified as long as they are separated with commas. Use the wildcard character (*) to specify a group of files when used in *logname*.

When a logname is not specified with the **restorelogarchive** command, the last transaction log archived to the Tivoli Storage Manager server (that is still active on the Tivoli Storage Manager server) is restored. The *lastarchivedlogfile* variable shown in the syntax diagram represents the default behavior and is not a keyword that can be specified on the command line.

To restore an inactive transaction log file from the Tivoli Storage Manager server, use the **/pick**=*showall* parameter and select the desired file from the list.

/ADSMNODe=nodename

Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile

Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password

Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to logon to the Tivoli Storage Manager server. If you specify

passwordaccessgenerate in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If passwordaccess is set to generate and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If passwordaccess is set to prompt and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the prompt.

If passwordaccess is set to prompt and you do not specify a password on the command line, then you are prompted for a password.

/BUFFers=numbuffers, buffersize
Specifies the number and size of data buffers that transfer data between the Domino server and the Tivoli Storage Manager API. Increasing the number or size (or both) of the data buffers can improve throughput.

You can specify 2 - 8 buffers. The default value is 3. The size of the buffers can be 64 - 8192 kb. The default value is 1024.

If the /buffers parameter is not specified on the command line or defined in the preferences file, Data Protection for IBM Domino uses the default values.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is domdsm.cfg.

/INTOPath=*translogpath* | *pathname*
Specifies the file path that is used for the restored transaction logs. The file path must be a fully qualified physical path. The *translogpath* variable that is shown in the syntax diagram represents the default location of the Domino server transaction log files and is not a keyword that can be specified on the command line. The default location of the Domino server transaction log files is defined by the TRANSLOG_Path variable in the notes.ini file.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is domdsm.log.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the /logfile parameter to specify a different log file for each instance. Using this parameter directs logging for

each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/MOUNTWait=Yes|No
If the Tivoli Storage Manager server is configured to store transaction log backup data on removable media, then the Tivoli Storage Manager server can indicate to Data Protection for IBM Domino that it is waiting for a required storage volume to be mounted. In this case, you can specify whether Data Protection for IBM Domino waits for the media mount or stops the current operation. Removable media is media such as tapes.

You can specify:

Yes Wait for tape mounts. This value is the default.

No Do not wait for tape mounts.

/PICk=*SHOWACtive*|*SHOWALl*
Displays a list of database backups that match the `dbname` pattern that can be selected for restore. The pick list is displayed as a scrollable list from which you can select the database backups for restore.

You can specify:

SHOWACtive
Displays a list of active database backup versions.

SHOWALl
Displays a list of both active and inactive database backup versions. All the backup versions that match the **dbname** pattern are shown.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/REPlace=*Yes* | *No*
Specifies whether to replace existing databases on the target system.

You can specify:

Yes Allows an existing database on the target system to be replaced during the restore process.

No Prevents an existing database on the target system from being overwritten during the restore process.

/SERVer=*currentserver* | *servername*
Specifies the Domino server name. If not specified, Data Protection for IBM Domino uses the current Domino server.

Example

The following example restores the last transaction log archived to the Tivoli Storage Manager server.

```
domdsmc restorelogarchive /intopath=c:\restoredLogs
```

Output Example:

```
Starting transaction log file restore...

Initializing Domino connection...
Logging on to the Tivoli Storage Manager server, please wait...
Querying Tivoli Storage Manager server for a list of transaction log file archives,
please wait...


Restoring transaction log file S0000524.TXN
to c:\restoredLogs\S0000524.TXN
Full: 0 Read: 67,109,888 Written: 67,109,888 Rate: 2,326.56 Kb/Sec
Restore of S0000524.TXN completed successfully.


Total transaction log file archives inspected: 27
Total transaction log file archives requested for restore: 1
Total transaction log file archives restored: 1

Throughput rate: 2,326.40 Kb/Sec
Total bytes transferred: 67,109,888
Total LanFree bytes transferred: 0
Elapsed processing time: 28.17 Secs
```

Domdsmc selective

How to use the **domdsmc selective** command is described.

Purpose

The **domdsmc selective** parameter backs up the databases you specify. You can exclude databases from backup with the exclude statement in the options file. This command does not do comparisons of attributes with the active backup images as does the incremental command. It backs up all databases that match the dbname pattern and passes the include-exclude filter.

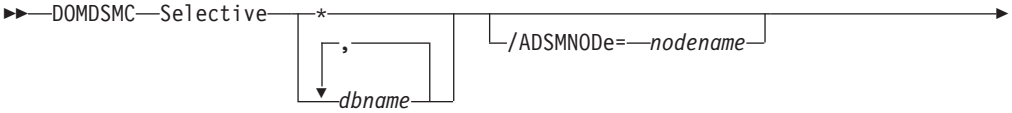

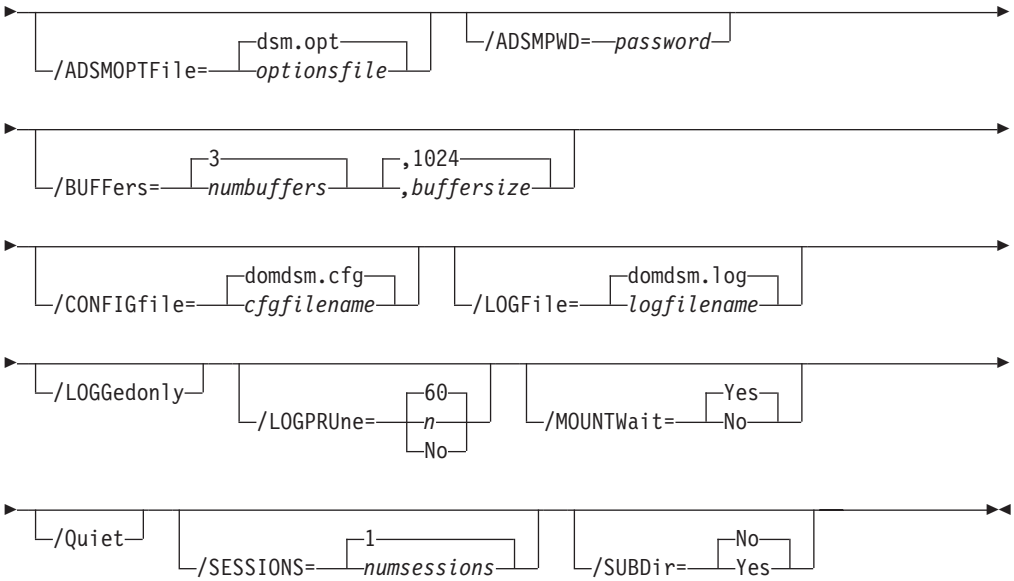

Parameters

*** | dbname,dbname,...**

Specifies the file path of a database or file path pattern for a group of databases. The file path pattern can represent a group of databases to be restored from the Tivoli Storage Manager server. The wildcard character asterisk (*) is used to specify a group of databases when used in the **dbname**. Multiple dbnames can be specified when they are separated with commas.

The file path must be relative to the Notes data directory.

Symbolic links are referred to by their symbolic names. To reference a database in a directory pointed to by a directory link in the data path, use the directory link name as the directory name. For example, if database xyz.nsf is in a directory, pointed to by the link vol1.dir, refer to it as vol1\xyz.nsf. If a symbolic directory link is created with the same name as a physical directory in the Notes data path, only the physical directory is searched.

The wildcard character (*) is used to represent any number of any characters when used in the file name portion of the file path. The wildcard character is not supported within directory names. The following example backs up all databases within the dir_A directory beginning with the characters *ter*:

```
domdsmc selective dir_A\ter*
```

The following example backs up all databases on the server:

```
domdsmc selective * /subdir=yes
```

The following example backs up all databases whose file name ends in acct:

```
domdsmc selective *acct.n* /subdir=yes
```

Note: Standard *include* and *exclude* processing applies to Domino database names. Specific databases can be excluded from the backup with the include-exclude list in the Tivoli Storage Manager options file. Wildcards

can be used on the backup command. For example, to exclude all databases on a volume pointed to by the symbolic directory link `temp.dir`, use the following statement:

```
exclude \temp\*
```

The exclude statement refers to the relative file name that includes symbolics and not the physical file path. For more information on include and exclude options, see "Include and exclude processing" on page 171 and *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide*.

/ADSMNODe=nodename
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify `passwordaccessgenerate` in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If `passwordaccess` is set to `generate` and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If `passwordaccess` is set to `prompt` and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the prompt.

If `passwordaccess` is set to `prompt` and you do not specify a password on the command line, then you are prompted for a password.

/BUFFers=numbuffers, buffersize
Specifies the number and size of data buffers that transfer data between the Domino server and the Tivoli Storage Manager API. Increasing the number or size (or both) of the data buffers can improve throughput.

You can specify from 2 to 8 buffers. The default value is 3. The size of the buffers can be 64 - 8192 kb. The default value is 1024.

If the `/buffers` parameter is not specified on the command line or defined in the preferences file, Data Protection for IBM Domino uses the default values.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGGedonly

Specifies that only logged databases that match the `dbname` pattern must be backed up. This option is used to force periodic refreshes of the backup for logged databases. Without a refresh, the databases are not backed up by the Incremental command on a Domino server when archival logging is in effect.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/MOUNTWait=Yes|No

If the Tivoli Storage Manager server is configured to store transaction log backup data on removable media, then the Tivoli Storage Manager server

can indicate to Data Protection for IBM Domino that it is waiting for a required storage volume to be mounted. In this case, you can specify whether Data Protection for IBM Domino waits for the media mount or stops the current operation. Removable media is media such as tapes.

You can specify:

Yes Wait for tape mounts. This value is the default.

No Do not wait for tape mounts.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/SESSions=numsessions|1
Specifies the number of Tivoli Storage Manager server sessions to be used by Data Protection for IBM Domino. You can specify 1 - 64 sessions. The default value is 1.

/SUBDir=No|Yes
Specifies whether subdirectories within the specified file path are searched for databases that match the file pattern. If this option is not specified, Data Protection for IBM Domino uses the value of the **/subdir** parameter in the Data Protection for IBM Domino preferences file.

You can specify:

No Do not search the subdirectories within the specified file path for databases that match the file pattern. This value is the default unless reset in the Data Protection for IBM Domino preferences file.

Yes Search the subdirectories within the specified file path for databases that match the file pattern.

Examples

Example 1:

The following example backs up all databases that are contained in the a_dir directory and its subdirectories:

```
domdsmc selective a_dir\* /subdir=yes
```

Output example:

```
Starting Domino database backup...
Initializing Domino connection...
Querying Domino for a list of databases, please wait...

Backing up database mail\ovargas.nsf, 1 of 1.
Full: 0  Read: 25,690,112  Written: 21,406  Rate: 11.90 Kb/Sec
Backup of mail\ovargas.nsf completed successfully.


Total Domino databases inspected:   1
Total Domino databases backed up:   1
Total Domino databases excluded:   0
Total Domino databases deduplicated:   1

Throughput rate:    11.73 Kb/Sec
Total bytes inspected:   25,690,112
Total bytes transferred:   21,406
Total LanFree bytes transferred:  0
Total bytes before deduplication:  25,690,112
Total bytes after deduplication:  50,002
Data compressed by:    58%
Deduplication reduction:  99.81%
Total data reduction ratio:  99.92%
Elapsed processing time:  1.78 Secs
```

Example 2: This example assumes that the Tivoli Storage Manager options file contains the following statements:

```
exclude c_dir\*
```

With this exclude statement, the following example backs up all databases, including databases contained in subdirectories, except for databases that are contained in the c_dir directory:

```
domdsmc selective * /subdir=yes
```

Domdsmc set

How to use the **domdsmc set** command is described.

Purpose

The **Domdsmc set** command sets the configuration options and values in the Data Protection for IBM Domino preferences file. The value that is saved in the preferences file is used as the default value when a parameter is not specified on a command invocation.

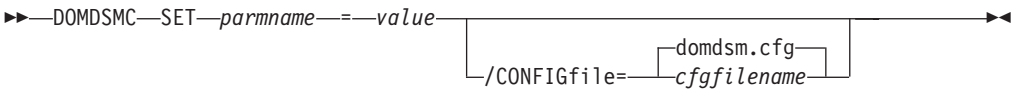

Parameters

parmname=*value*

Specifies the parameter and value to save in the preferences file. You can set only one value with the **domdsmc set** command run.

The **parmname**=*value* is one of the following strings:

ADSMLOGDir=*directory path*

Specify the full path name to where the Tivoli Storage Manager API error log file (dsierror.log) is stored. The default directory is the Data Protection for IBM Domino installation directory. Specify

adsmlogdir in the Data Protection for IBM Domino preferences file (domdsm.cfg) when you are using the web client GUI.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is dsm.opt.

/BUFFers=numbuffers, buffersize
Use this parameter to specify the number of data buffers and the size of data buffers that transfer data between the Domino server and the Tivoli Storage Manager API. Increasing the number and size of the data buffers can improve throughput.

You can specify 2 - 8 buffers, the default value is 3. The size of the buffers can be from 64 to 8192 kb. The default value is 1024.

If the /buffers parameter is not specified on the command line or defined in the preferences file, Data Protection for IBM Domino uses the default values.

BUFFERSIze=*size*
Specifies the size of the buffers. The size can be 64 - 8192 kilobytes. The default is 1024.

COMMRESTARTDURATION =*number*
Specifies the total number of minutes that the server continues trying to restart a session. The valid range is 1 - 9999 and the default is 60.

COMMRESTARTINTERVAL=*number*
Specifies the number of seconds the server waits between attempts to restart a session. The valid range is 1 - 9999 and the default is 15.

DATEformat=*formatnumber*
Specifies the format that you want to use to display dates.

The default value is defined in the message file. As a result, each language has a different default value.

The *formatnumber* variable displays the date in one of the formats listed, select the number that corresponds to the format you want to use:

1 The format is MM/DD/YYYY.

2 The format is DD–MM–YYYY.

3 The format is YYYY–MM–DD.

4 The format is DD.MM.YYYY.

5 The format is YYYY.MM.DD.

Changes to the value of the **dateformat** parameter can result in an undesired pruning of the Data Protection for IBM Domino log file (tdpdom.log by default). You can avoid losing existing log file data by doing one of the following actions:

- After changes are made to the value of the **dateformat** parameter, make a copy of the existing log file before you run Data Protection for IBM Domino.

- Specify a new log file with the **/logfile** parameter.

DB2ALtdbname=*name of alternate database*
Specifies the name of the different DB2 database. The default value is DOM_ALT.

DB2CONTainerpath=*directory path*
Specifies the default path for table space containers that are used on DB2 redirected restores. Redirected DB2 restores are selected automatically by Data Protection for IBM Domino when running an alternate DB2 database restore. If this option is not specified, the target path for the table space container is defined relative to the value of the **db2restoreintopath** option.

Note: **db2containerpath** is required when the DB2DIRECTORY option is specified in the Domino server `notes.ini` file. Otherwise, the restore fails. That is because DB2 attempts to place the alternate DB2 database data in the directory that is specified by the DB2DIRECTORY option, which is already used by the Domino DB2 database.

DB2LOGPath=*directory path*
Specifies the absolute path name of a directory that will be used for active log files after a restore operation.

DB2LOGTarget=*directory path*
Specifies the location for the logs from the backup image during a restore operation.

DB2REPlace=*Yes | No*
Specifies whether to replace the existing alternate DB2 database when you are running a restore operation. This parameter defines the default behavior if the **db2replace** parameter is not specified during the **db2restore** command. The default value is Yes.

DB2RESTIntopath*drive*
Specifies the target DB2 database directory drive letter when you are restoring to an alternate DB2 database. The specified drive must be local. The default is the value of the DB2 instance default database path configuration option.

DB2SESSions=*number*
Specifies the number of sessions to be created between DB2 and the Tivoli Storage Manager server. This parameter is used by the Tivoli Storage Manager DB2 agent to back up DB2 data. You can specify 1-64 sessions. The default is one session.

DB2USER=*user name*
Specifies the DB2 user name.

DOMTXNBYTElimit=*number*
Specifies the number of bytes sent between Data Protection for IBM Domino and the Tivoli Storage Manager server in a single transaction. The default value is 0 (which indicates no limit) and the maximum value is 2097152. This number is multiplied by 1024 to calculate the limit in bytes. This parameter is useful when you are backing up NSF databases to tape storage for these reasons:

- Processing for each transaction causes the tape to stop and start. Considerable time can be lost during the stop and start when you are using high speed tapes. This time loss is true in a LAN-free environment.
- Errors that are encountered during backup processing are automatically tried again if **domtxnbytelimit** is set.
- When a failure occurs during a backup, all of the backups in the transaction are tried again, not just the NSF database in error. Each backup is tried in a separate transaction, and the **domtxnbytelimit** parameter is used to control the number of bytes per transaction.

Stopping and starting tape storage for large NSF databases is small when compared to data transfer time. Use the **domtxnbytelimit** parameter to adjust the behavior for large database backups.

DOMTXNGROUPmax=*number*

Specifies the number of individual objects sent to the Tivoli Storage Manager server in a single transaction. Two objects are sent to the Tivoli Storage Manager server for each database backup. The default value of 2 specifies that there is one database per transaction and that each database is stored as two objects on the Tivoli Storage Manager server. The maximum value is 65000. This parameter can be overridden by the Tivoli Storage Manager server TXNGRPMAX option. However, when **domtxngroupmax** is set, the minimum of the two values is used. This parameter is useful when you are backing up NSF databases to tape storage for these reasons:

- Processing for each transaction causes the tape to stop and start. Considerable time can be lost during the stop and start when you are using high speed tapes. This is true in a LAN-free environment.
- Errors that occur during backup processing are automatically tried when **domtxngroupmax** is set.
- When a failure occurs during a backup, all of the backups in the transaction are tried again, not just the NSF database in error. Each backup is tried in a separate transaction. After all backups are tried, the **domtxngroupmax** parameter is used to control the number of individual objects per transaction.

The **domtxngroupmax** parameter must be used when you are backing up small NSF databases.

LANGuage=*language*

Specifies the language that you want to use to display messages.

The *language* variable displays messages in one of the languages that are listed. Select the entry that corresponds to the language you want to use.

CHS Simplified Chinese

CHT Traditional Chinese

CSY Czech

DEU Standard German

ENU English (United States). This language is the default .

ESP Standard Spanish

FRA Standard French

HUN Standard Spanish

ITA Standard Italian

JPN Japanese

KOR Korean

PLK Polish

PTB Brazilian Portuguese

RUS Russian

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the `set` command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. This parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and run daily. You can use the `set` command to do the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is completed for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the `timeformat` or `dateformat` parameter can result in an undesired pruning of the log file. If you are

running a command that prunes the log file and the value of the `timeformat` or `dateformat` parameter changes, do one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/MOUNTWait=Yes|No
If the Tivoli Storage Manager server is configured to store transaction log backup data on removable media, then the Tivoli Storage Manager server indicates to Data Protection for IBM Domino that it is waiting for a required storage volume to be mounted. The mount command is used to specify whether Data Protection for IBM Domino waits for the media mount or stops the current operation. Removable media is media such as tapes or DVDs.

You can specify:

Yes Wait for tape mounts. This value is the default.

No Do not wait for tape mounts.

NOTESInipath=*dirpath*
Specifies the directory path where the `notes.ini` file is stored for the target Domino server. In a multiple Domino server partition environment, the **notesinipath** parameter must be specified for each partition to identify the Domino server for Data Protection for IBM Domino.

NUMberformat=*fmtnum*
The **numberformat** parameter specifies the format that you want to use to display numbers.

The default value is defined in the message file. As a result, each language has a different default value.

The *fmtnum* variable displays numbers by using one of the formats that are listed. Select the format number that corresponds to the format you want to use.

1 The format is `n,nnn.dd`.

2 The format is `n,nnn,dd`.

3 The format is `n nnn,dd`.

4 The format is `n nnn.dd`.

5 The format is `n.nnn,dd`.

6 The format is `n'nnn,dd`.

/REPlace=*Yes* | *No*
Specifies whether to replace existing databases or log files on the target system when you are running a restore operation. This parameter defines the default behavior if the **/replace** parameter is not specified on the restore command.

If the target path name for a database or log file to be restored exists, you can specify:

Yes A Yes value activates the restore procedure and replaces the existing database or log file on the target system. This value is the default.

No A No value prevents the existing database or log file to be replaced, so the restore of that database or log file is skipped.

/SESSions=*number*
Specifies the number of sessions to open to the Tivoli Storage Manager server. This option applies to NSF database backups only. You can specify a 1 - 64 sessions. The default value is 1.

/STATistics=*No* | *Yes*
Specifies whether to log backup and restore performance statistics about an individual database at the backup or restore level. Statistics are logged to the Data Protection for IBM Domino log file (`domdsm.log` by default). The statistics contain information such as the database read/write time and transfer rate, the send/receive time and transfer rate, and the Domino server data transfer time and transfer rate. This information can help tuning Data Protection for IBM Domino for optimal performance.

You can specify:

No A No value results in no collection of log backup and restore performance statistics for an individual database. This value is the default.

Yes A Yes value results in the collection of log backup and restore performance statistics for an individual database.

SUBDir=*No* | *Yes*
Specifies whether subdirectories within the specified file path are searched for databases that match the file pattern.

You can specify:

No Do not search the subdirectories within the specified file path for databases that match the file pattern. This value is the default unless reset in the Data Protection for IBM Domino preferences file.

Yes Search the subdirectories within the specified file path for databases that match the file pattern.

TIMEformat=*formatnumber*
Specifies the format in which you want system time that is displayed.

The default value is defined in the message file. As a result, each language has a different default value.

The *formatnumber* variable displays time in one of the formats that are listed. Select the format number that corresponds to the format you want to use.

1 The format is HH:MM:SS.

2 The format is HH,MM,SS.

3 The format is HH.MM.SS.

4 The format is HH:MM:SSA/P.

Changes to the value of the **timeformat** parameter can result in an undesired pruning of the log file (`tdpdom.log` by default). You can avoid losing existing log file data by running one of the following actions:

- After you change the value of the **timeformat** parameter, make a copy of the existing log file before you run Data Protection for IBM Domino.
- Specify a new log file with the **/logfile** parameter.

/CONFIGfile=*cfgfilename*
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file that is stored in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

Example

Example: The following example sets the number of buffers to 8.

```
domdsmc set buffers=8
```

Output example:

```
ACD5217I The preference has been set successfully.
```

DB2 Commands

How to use the Data Protection for IBM Domino command-line interface with Domino DB2 enabled Notes databases is described.

Domdsmc DB2activatedbs

How to use the **domdsmc db2activatedbs** command is described.

Purpose

The **domdsmc db2activatedbs** parameter activates DB2 enabled Notes databases that are restored to an alternate database. This command copies the specified DB2 enabled Notes database into the Domino DB2 database and brings the database online. When a list of database is activated and the **/isolate** parameter is specified, each DB2 enabled Notes database is activated to a new DB2 Group.

The DB2 enabled Notes databases in the alternate DB2 database are available for restore as long the alternate DB2 database is available. The alternate DB2 database is considered available when it is not manually deleted through DB2 server interface, overwritten by another restore operation, or removed from the list of DB2 databases (that contain DB2 enabled Notes databases) available for activation. See the "Domdsmc DB2deletealternate" on page 127 command for information about how to remove a database from the activation list.

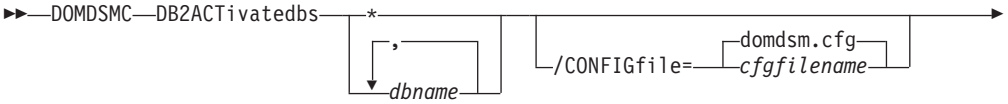

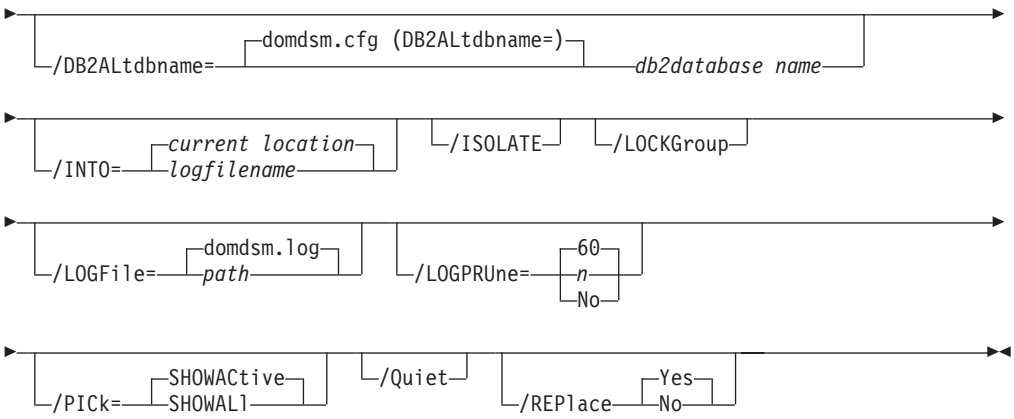

Parameters

***|dbname,dbname,...**
Specifies the DB2 enabled Notes databases to activate. The DB2 enabled Notes databases are activated into the current DB2 Group.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/DB2ALtdbname=*database name*
Specify the name of the alternate DB2 database to use for activation. If the **/DB2ALtdbname** parameter is not specified, the value of the **/DB2ALtdbname** configuration option (defined in the Data Protection for IBM Domino `domdsm.cfg` preferences file) is used. If **/DB2ALtdbname** is not defined in the preferences file, the alternate database name DOM_ALT is used.

/INTOPath=*translogpath|pathname*
Specifies the file path that is used for the restored transaction logs. The file path must be a fully qualified physical path. The *translogpath* variable that is shown in the syntax diagram represents the default location of the Domino server transaction log files and is not a keyword that can be specified on the command line. The default location of the Domino server transaction log files is defined by the TRANSLOG_Path variable in the `notes.ini` file.

/ISOLATE
Specify this parameter to activate the database into a new DB2 Group.

/LOCKGroup
Specify whether to lock the DB2 Group after the DB2 enabled Notes database is activated.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.

The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/PICk=*SHOWACtive*|*SHOWALl*
Displays a list of database backups that match the `dbname` pattern that can be selected for restore. The pick list is displayed as a scrollable list from which you can select the database backups for restore.

You can specify:

SHOWACtive
Displays a list of active database backup versions.

SHOWALl
Displays a list of both active and inactive database backup versions. All the backup versions that match the **dbname** pattern are shown.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/REPlace=*Yes*|*No*
Specifies whether to replace existing databases on the target system.

You can specify:

Yes Allows an existing database on the target system to be replaced during the restore process.

No Prevents an existing database on the target system from being overwritten during the restore process.

Example

The following example displays a list of DB2 enabled Notes databases that are ready to activate from specified alternate DB2 database:

```
domdsmc db2activatedbs * /pick=showall
```

Output example:

```
  #       Backup Time Stamp            Size   DB2 DB   GROUP  NSFDB2 Database
       ------------------------------------------------------------------------
     1. | 01/19/2012 13:55:33         0.00B  DOM_ALT  GRP1   db2nsf1.nsf
     2. | 01/19/2012 13:55:33         0.00B  DOM_ALT  GRP1   db2nsf2.nsf
     2. | 01/19/2012 13:55:33         0.00B  DOM_ALT  GRP1   db2nsf3.nsf
     3. | 01/20/2012 13:55:33         0.00B  DOMFULL1 GRP1   db2nsf1.nsf
     4. | 01/20/2012 13:55:33         0.00B  DOMFULL1 GRP1   db2nsf2.nsf
     4. | 01/20/2012 13:55:33         0.00B  DOMFULL1 GRP1   db2nsf3.nsf
     4. | 01/20/2012 13:55:33         0.00B  DOMFULL1 GRP2   db2nsf4.nsf
     4. | 01/20/2012 13:55:33         0.00B  DOMFULL1 GRP2   db2nsf5.nsf
     4. | 01/20/2012 13:55:33         0.00B  DOMFULL1 GRP2   db2nsf6.nsf
        |
        |
        |
        |
        |
        |
        |
        |
        0---------10--------20--------30--------40--------50--------60--------7
<U>=Up  <D>=Down  <T>=Top  <B>=Bottom  <R>=Right  <L>=Left <G#>=Goto Line #
<#>=Toggle Entry  <+>=Select All  <->=Deselect All <#:#+>=Select A Range
<#:#->=Deselect A Range  <O>=Ok  <C>=Cancel
```

Domdsmc DB2archivelog

How to use the **domdsmc db2archivelog** command is described.

Purpose

The **domdsmc db2archivelog** parameter archives the Domino DB2 database log files. Although DB2 automatically archives the log file, it is possible to force an archive of the log so that the latest transactions are available when the alternate database is rolled forward.

An archive copy group is required to archive the Domino DB2 database log file. If an archive copy group is not defined on the target management class, the **db2archivelog** command completes successfully but the log is not archived to the Tivoli Storage Manager server. The logs are archived to the path specified by the Domino DB2 database configuration option, FAILARCHPATH.

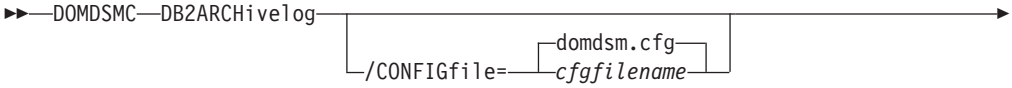

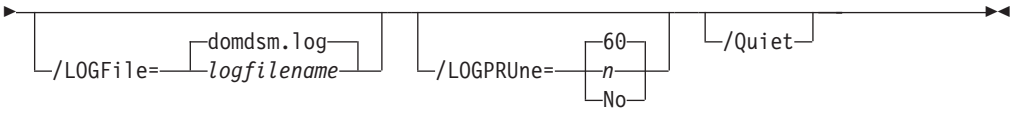

Parameters

/CONFIGfile=cfgfilename

Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that

prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

Example

The following example archives the DB2 transaction log file:

```
domdsmc db2archivelog
```

Output example:

```
Starting Domino DB2 database transaction log archive...
Initializing Domino connection...
Initializing DB2 connection...

Archiving Domino DB2 transaction logs

Domino DB2 transaction log archive completed successfully.
```

Domdsmc DB2deletealternate

How to use the **domdsmc db2deletealternate** command is described.

Purpose

The **domdsmc db2deletealternate** command deletes the specified alternate DB2 database from the pending DB2 file.

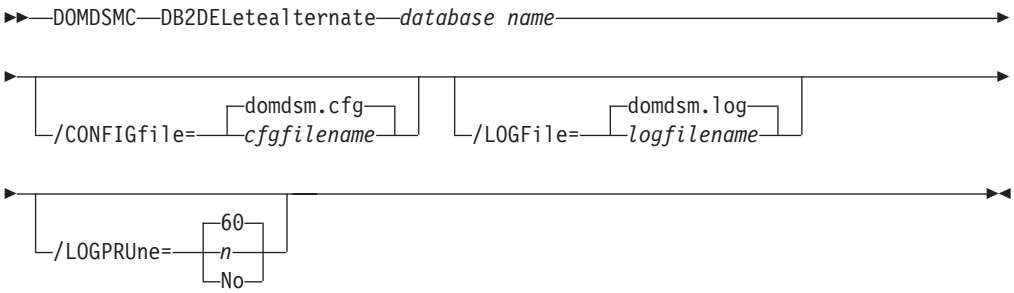

Parameters

dbname Specifies the alternate DB2 database to deleted. If not specified, the default alternate DB2 database (DB2ALTDBNAME) is used.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection

for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

Example

The following example deletes the alternate DB2 database:

```
domdsmc db2deletealternate
```

Output example:

```
Starting Domino DB2 database transaction log archive...
Initializing Domino connection...
Initializing DB2 connection...

Archiving Domino DB2 transaction logs

Domino DB2 transaction log archive completed successfully.
```

Domdsmc fullselective

The **domdsmc fullselective** command is a method of backing up both NSF and DB2 enabled Notes databases.

Purpose

The **domdsmc fullselective** command first backs up all of the Domino NSF databases and then runs a full backup of the Domino DB2 database to back up all of the DB2 enabled Notes databases. It differs from the **selective** command, which backs up NSF databases only. For non-DB2 enabled servers, use the **selective** command.

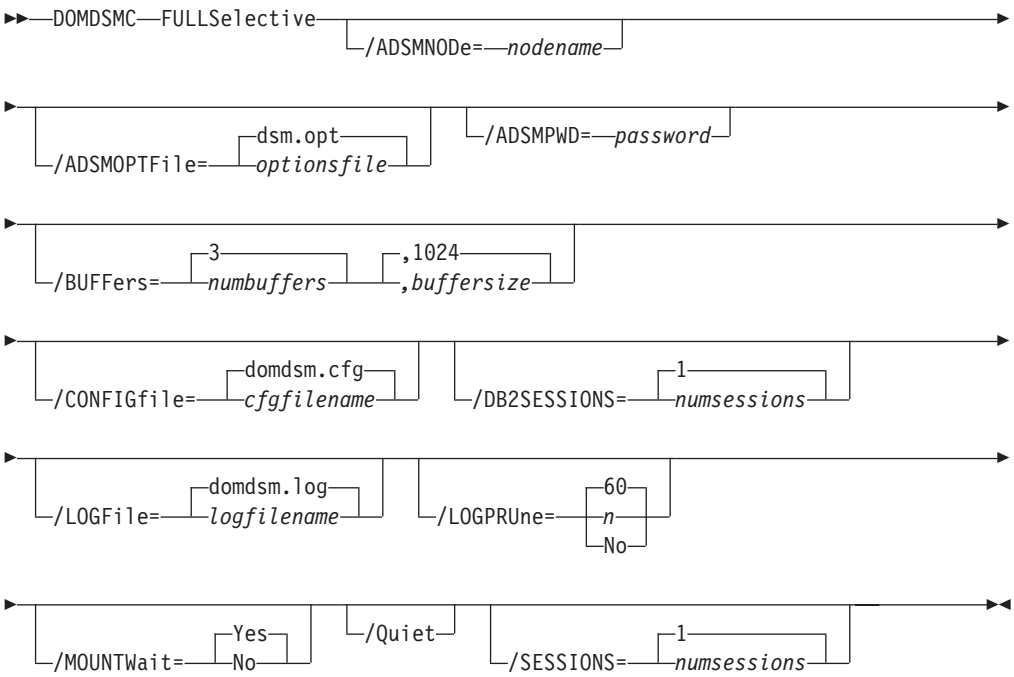

Parameters

/ADSMNODe=nodename

Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile

Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to logon to the Tivoli Storage Manager server. If you specify `passwordaccessgenerate` in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If `passwordaccess` is set to `generate` and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If `passwordaccess` is set to `prompt` and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the prompt.

If `passwordaccess` is set to `prompt` and you do not specify a password on the command line, then you are prompted for a password.

/BUFFers=numbuffers, buffersize
Specifies the number and size of data buffers that transfer data between the Domino server and the Tivoli Storage Manager API. Increasing the number or size (or both) of the data buffers can improve throughput.

You can specify 2 - 8 buffers. The default value is 3. The size of the buffers can be 64 - 8192 kb. The default value is 1024.

If the `/buffers` parameter is not specified on the command line or defined in the preferences file, Data Protection for IBM Domino uses the default values.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/DB2SESSIONS=*numsessions*
Specify the number of Tivoli Storage Manager sessions that the DB2 Tivoli Storage Manager agent uses. You can specify 1-64 sessions. The default value is one.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for

each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/MOUNTWait=Yes|No
If the Tivoli Storage Manager server is configured to store transaction log backup data on removable media, then the Tivoli Storage Manager server can indicate to Data Protection for IBM Domino that it is waiting for a required storage volume to be mounted. In this case, you can specify whether Data Protection for IBM Domino waits for the media mount or stops the current operation. Removable media is media such as tapes.

You can specify:

Yes Wait for tape mounts. This value is the default.

No Do not wait for tape mounts.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/SESSions=numsessions|1
Specifies the number of Tivoli Storage Manager server sessions to be used by Data Protection for IBM Domino. You can specify from 1 to 64 sessions. The default value is 1.

Example

The following example backs up both NSF and DB2 enabled Notes databases using two sessions for the DB2 Tivoli Storage Manager agent to access the Tivoli Storage Manager server: `domdsmc fullselective /db2session=2`

```
Starting Domino full backup...
Initializing Domino connection...
Querying Domino for a list of databases, please wait...

Restart Analysis (0 MB): 100%
09/28/2011 06:55:58 AM Recovery Manager: Restart Recovery complete.
(0/0 databases needed full/partial recovery)
09/28/2011 06:56:01 AM The map for DB2 errors was successfully created.


...

Backing up database statrep5.ntf, 114 of 122.
Full: 0 Read: 1,686,528 Written: 1,686,528 Rate: 4,844.12 Kb/Sec

Backup of statrep5.ntf completed successfully.

Backing up database teamrm7.ntf, 115 of 122.
Full: 0 Read: 2,883,584 Written: 2,883,584 Rate: 6,834.95 Kb/Sec

Backup of teamrm7.ntf completed successfully.

Backing up database toolbox.ntf, 116 of 122.
Full: 0 Read: 688,128 Written: 688,128 Rate: 5,209.30 Kb/Sec

Backup of toolbox.ntf completed successfully.

Backing up database updatesite.ntf, 117 of 122.
Full: 0 Read: 2,883,584 Written: 2,883,584 Rate: 8,233.92 Kb/Sec

Backup of updatesite.ntf completed successfully.

Backing up database userlicenses.ntf, 118 of 122.
Full: 0 Read: 663,552 Written: 663,552 Rate: 4,729.93 Kb/Sec

Backup of userlicenses.ntf completed successfully.

Backing up database userreg.ntf, 119 of 122.
Full: 0 Read: 458,752 Written: 458,752 Rate: 3,612.90 Kb/Sec

Backup of userreg.ntf completed successfully.

Backing up database webadmin.nsf, 120 of 122.
Full: 0 Read: 8,388,608 Written: 8,388,608 Rate: 11,457.34 Kb/Sec

Backup of webadmin.nsf completed successfully.

Backing up database webadmin.ntf, 121 of 122.
Full: 0 Read: 10,747,904 Written: 10,747,904 Rate: 11,371.61 Kb/Sec

Backup of webadmin.ntf completed successfully.

Backing up database xxx\busytime.ntf, 122 of 122.
Full: 0 Read: 248,832 Written: 248,832 Rate: 1,840.91 Kb/Sec

Backup of xxx\busytime.ntf completed successfully.

Backing up DB2 database DOMINO7.

Domino DB2 database backup completed successfully.


Total Domino NSF databases inspected: 122
Total Domino NSF backed up: 122
Total Domino NSF excluded: 0
Total Domino NSF deduplicated: 0
Total Domino NSF bytes inspected: 1,734,538,240
```

```
Total Domino NSF bytes transferred: 1,734,538,240
Total Domino NSF LanFree bytes transferred: 0
Total Domino NSF bytes before deduplication: 0
Total Domino NSF bytes after deduplication: 0
Total Domino NSF data compressed by: 0%
Total Domino NSF deduplication reduction: 0.00%
Total Domino NSF data reduction ratio: 0.00%

Domino DB2 database inspected: 1
Domino DB2 database backed up: 1

Throughput rate: 19,821.81 Kb/Sec
Total bytes transferred: 2,325,935,104
Elapsed processing time: 114.59 Secs
```

Domdsmc DB2INActivateobjs

How to use the **domdsmc db2inactivateobjs** command is described.

Purpose

The **domdsmc db2inactivateobjs** parameter displays the db2adutl utility commands that are required to inactivate Tivoli Storage Manager objects that have been created by the DB2 API and are no longer referenced by any Data Protection for IBM Domino Tivoli Storage Manager objects. The db2adutl utility is part of the DB2 Tivoli Storage Manager Agent and is used to manage Tivoli Storage Manager objects. The db2adutl commands (displayed by the **domdsmc db2inactivateobjs** command) must be issued from a DB2 command window and should be issued regularly after DB2 Group or full DB2 backups.

When Data Protection for IBM Domino backs up a DB2 Group or a DB2 database, the backup objects are created by the DB2 API. These objects have a unique name and must be inactivated when they are no longer referenced by Tivoli Storage Manager objects that are expired as a result of management policies. This command also inactivates table spaces, full DB2 database backups, and archived logs. This command must be run regularly after DB2 Group or full DB2 backups.

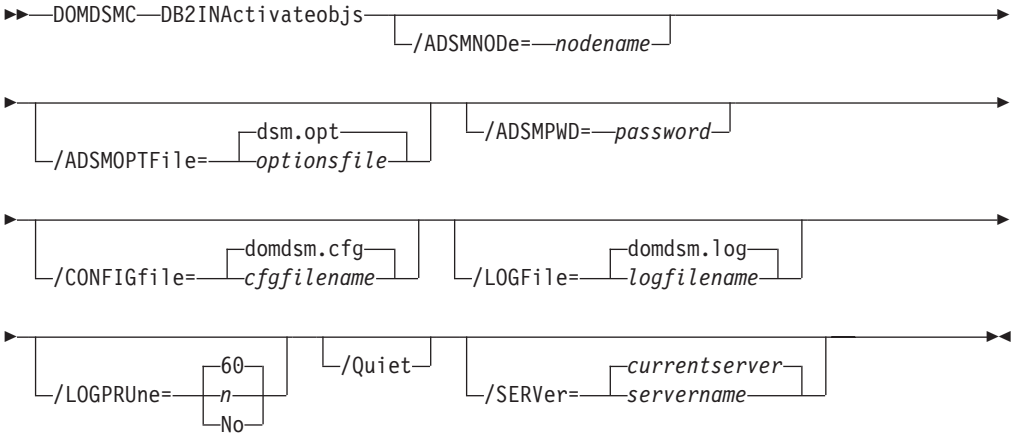

Parameters

/ADSMNODe=nodename

Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to logon to the Tivoli Storage Manager server. If you specify `passwordaccessgenerate` in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If `passwordaccess` is set to `generate` and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If `passwordaccess` is set to `prompt` and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the prompt.

If `passwordaccess` is set to `prompt` and you do not specify a password on the command line, then you are prompted for a password.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the /logprune option to override these defaults for one command run. Note, when the value of /logprune is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the /logfile parameter or logfile setting.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/SERVer=*currentserver*|*servername*
Specifies the Domino server name. If not specified, Data Protection for IBM Domino uses the current Domino server.

Example

The following example displays the db2adutl utility commands that are required to inactivate Tivoli Storage Manager objects that are created by the DB2 API and are no longer referenced by any Data Protection for IBM Domino Tivoli Storage Manager objects:

```
domdsmc db2inactivateobjs
```

Output example:

```
Issue the following DB2 command to delete unneeded log archives:
db2adutl DELETE FULL OLDER THAN 20070925091903 DATABASE DOMINO

Issue the following DB2 command to delete unneeded tablespace backups:
db2adutl DELETE TABLESPACE OLDER THAN 20070925082543 DATABASE DOMINO

Issue the following DB2 command to delete unneeded full database backups:
db2adutl DELETE LOGS BETWEEN S0000000 AND S0000016 DATABASE DOMINO
```

Domdsmc DB2restore

How to use the **domdsmc db2restore** command is described.

Purpose

The **domdsmc db2restore** command restores Domino DB2 enabled Notes databases from one of the following sources:

- a DB2 Group backup image.
- a set of DB2 Groups from a full DB2 database backup image
- a full DB2 database from a full DB2 database backup image

Note that all of the DB2 enabled Notes databases that reside in the DB2 Group, the set of DB2 Groups, or the full DB2 database backup image are restored.

The pending DB2 file is updated during a successful **db2restore** restore operation. Note that when performing an alternate database restore, the alternate database can exist as long as **/db2replace**=*yes* is specified and the log directory must not be in use by another DB2 database during the first alternate database restore operation.

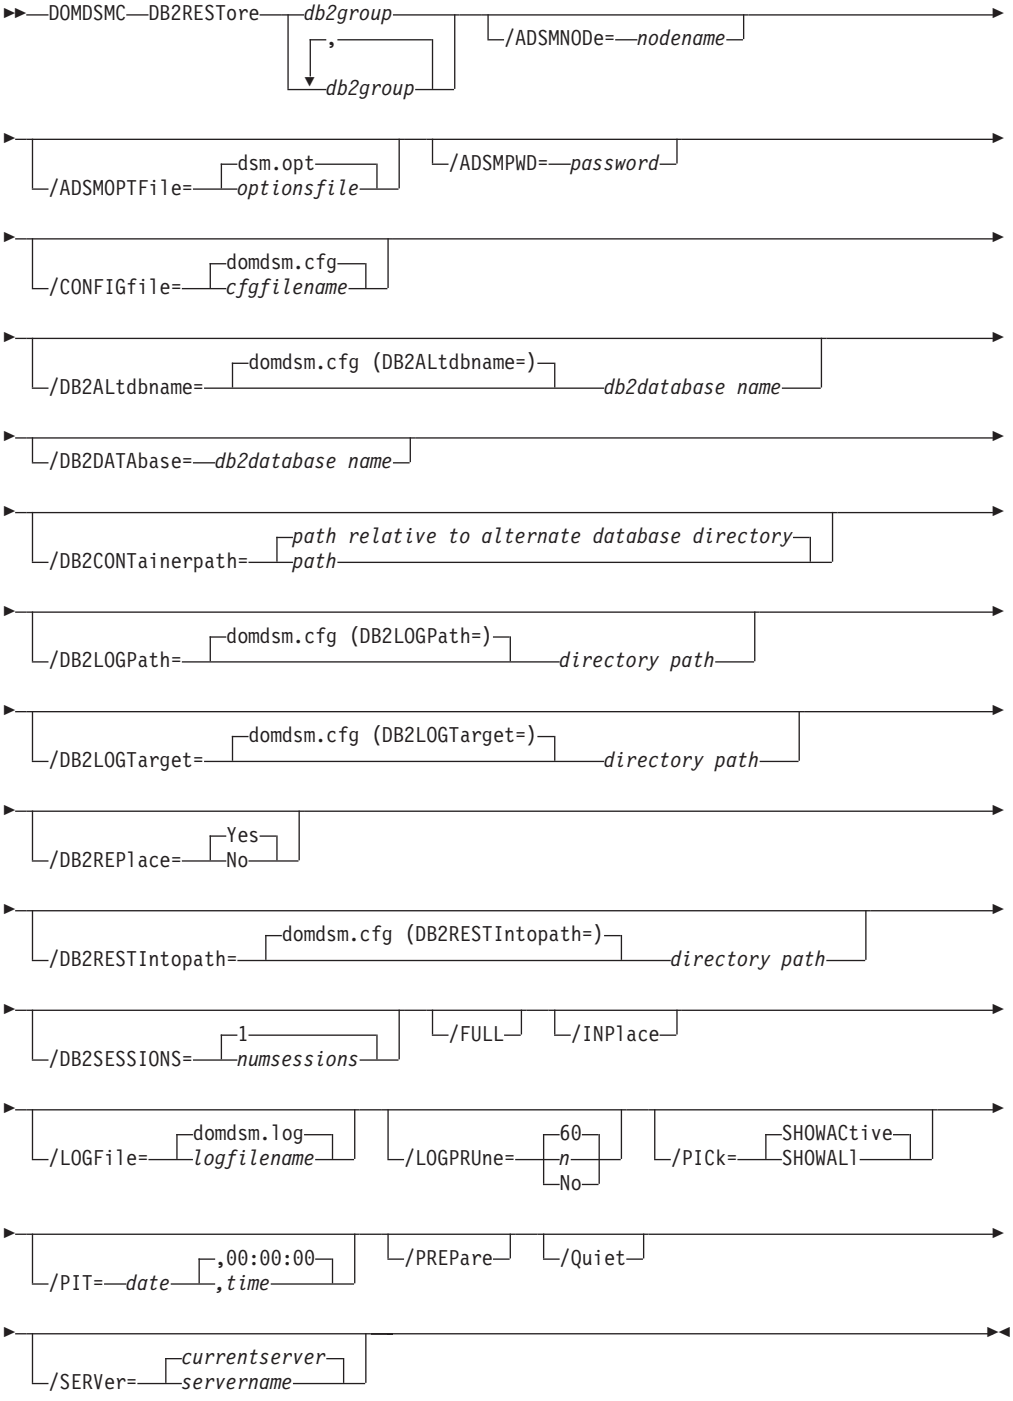

Parameters

db2group,...
Specifies the DB2 Group to restore from a table space backup image. Only one DB2 Group can be specified when you are restoring from a table space backup image. When you are restoring a full DB2 database backup image (**/full=yes**), you can specify multiple DB2 Groups by name or you can specify the wildcard character (*).

/ADSMNODe=nodename
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to logon to the Tivoli Storage Manager server. If you specify `passwordaccessgenerate` in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If `passwordaccess` is set to `generate` and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If `passwordaccess` is set to `prompt` and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the prompt.

If `passwordaccess` is set to `prompt` and you do not specify a password on the command line, then you are prompted for a password.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/DB2ALtdbname=*database name*
Specify the name of the alternate DB2 database to use for activation. If the **/db2altdbname** parameter is not specified, the value of the **db2altdbname** configuration option (defined in the Data Protection for IBM Domino `domdsm.cfg` preferences file) is used. If **db2altdbname** is not defined in the preferences file, the alternate database name DOM_ALT is used.

/DB2DATAbase=*database name*
Specify the name of the DB2 database to restore. The name that is specified overrides the default name (which is the name of the Domino DB2 database that is used by the Domino server). If the **/DB2DATAbase** parameter is not specified, the current Domino DB2 database name is used.

/DB2CONTainerpath=*path*
Specify container path to be used for a redirected restore operation. A

redirected restore operation is run when you are restoring to an alternate database (*inplace=no*) or when you are restoring in place and redefining the table space containers. If **/DB2CONTainerpath** is not specified during an alternate restore operation, the table space containers are defined relative to the alternate database directory.

Attention: **/DB2CONTainerpath** is required when the DB2DIRECTORY option is specified in the Domino server `notes.ini` file. Otherwise, the restore fails because DB2 attempts to place the alternate DB2 database data in the directory that is specified by the DB2DIRECTORY option, which is already used by the Domino DB2 database.

/DB2LOGPath=*path*
Specify the base log directory for the alternate database. The directory must exist and must not contain any files before the **db2restore** command is run. When the **/DB2LOGPath** parameter is not specified, the configuration option **db2altdbname** is used. Since the log path cannot be shared by more than one DB2 database, this option must be specified if one alternate database exists.

/DB2LOGTarget=*path*
Specify the target directory for extracting log files from a backup image during an *inplace* restore. If the **/DB2LOGTarget** parameter is not specified, the value of the **/DB2LOGTarget** configuration option (defined in the Data Protection for IBM Domino `domdsm.cfg` preferences file) is used.

/DB2REPlace=*Yes | No*
Specify whether to replace an alternate database (if it exists). The default value is Yes.

/DB2RESTIntopath=*drive*
Specify the base target DB2 database directory drive letter for the alternate database when you are restoring to an alternate DB2 database. The specified drive must be local. If the **/DB2RESTIntopath** parameter is not specified, the configuration option **/DB2RESTIntopath** is used. If the **/DB2RESTIntopath** configuration option is not specified, the DB2 database default database directory configuration setting is used.

/DB2SESSIONS=*numsessions*
Specify the number of Tivoli Storage Manager sessions that the DB2 Tivoli Storage Manager agent uses. You can specify 1 - 64 sessions. The default value is one.

/FULL=*Yes | No*
Specify whether a full DB2 database is restored.

/INPlace
Specify an *inplace* restore. An *inplace* restore is allowed only when you are restoring a full DB2 database backup image.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.

The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/PICk=*SHOWACtive*|*SHOWALl*
Displays a list of database backups that match the `dbname` pattern that can be selected for restore. The pick list is displayed as a scrollable list from which you can select the database backups for restore.

You can specify:

SHOWACtive
Displays a list of active database backup versions.

SHOWALl
Displays a list of both active and inactive database backup versions. All the backup versions that match the **dbname** pattern are shown.

/pit=*currentdate,currenttime*|*date,time*
Specifies a point in time when the specified databases are restored. The date and time values must be specified in the same date and time format that is defined in the Data Protection for IBM Domino preferences file. The

most recent database backup images that are taken before the specified point in time are restored. Deleted backup images are not restored. Logged databases can then be rolled forward to that point by specifying the same date and time values on the **/applylogs** option of the **activatedbs** command.

date Specify a date string in the active date format. If you do not specify a date, the specified databases are restored unless the **/pick** parameter was used to select inactive backup versions.

The date must be specified with the same date format that is defined in the Data Protection for IBM Domino preferences file. See "Domdsmc set" on page 115 for a list of available date formats.

time Specify a time string in the active time format. If you specify a date without the time, HH:MM:SS on a 24-hour clock is used.

The time must be specified with the same time format that is defined in the Data Protection for IBM Domino preferences file. See "Domdsmc set" on page 115 for a list of available time formats.

Note: If this parameter is used with the **/pick** parameter, the *showactive* and *showall* variables for the **/pick** parameter are ignored. The pick list contains the database backup images that meet the **/pit** criteria.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

/SERVer=*currentserver*|*servername*
Specifies the Domino server name. If not specified, Data Protection for IBM Domino uses the current Domino server.

Examples

Example 1: The following example restores the full backup image of all DB2 Groups (created when originally backed up with the `db2selective /full=yes` command) into the existing database (*inplace* restore):

```
domdsmc db2restore /full /inplace
```

Output example:

```
Starting Domino DB2 database restore...
Restoring Domino DB2 Database DOMINO to DOMINO
Restore of DOMINO completed successfully.
```

Example 2: The following example restores a DB2 enabled Notes database from a full backup image of all DB2 Groups (created when originally backed up with the `db2selective /full=yes` command) into an alternate DB2 database:

```
domdsmc db2restore /full
```

Output example:

```
Starting Domino DB2 database restore...
Restoring Domino DB2 Database DOMINO to DOM_ALT
Restore of DOMINO completed successfully.
```

After the restore completes:

1. The DB2 Groups in the restored alternate DB2 database are updated with the changes in the DB2 transaction logs by issuing the "Domdsmc DB2rollforward" command.
2. The DB2 enabled Notes databases (in the restored DB2 Groups) must be copied from the alternate restored DB2 database to their final DB2 enabled Notes location by issuing the "Domdsmc DB2activatedbs" on page 122 command.

Example 3: The following example restores a DB2 enabled Notes database from a backup image of DB2 Group `GRP1` (created when originally backed up with the `db2selective GRP1` command) into an alternate DB2 database:

```
domdsmc db2restore GRP1
```

Output example:

```
Starting Domino DB2 database restore...
Restoring Domino DB2 group GRP1 to DB2 Database DOM_ALT
Restore of GRP1 completed successfully.
```

After the restore completes:

1. The restored DB2 Group is updated with the changes in the DB2 transaction logs by issuing the "Domdsmc DB2rollforward" command.
2. The DB2 enabled Notes databases (in the restored DB2 Group) must be copied from the temporary restored DB2 Group to their final DB2 enabled Notes location by issuing the "Domdsmc DB2activatedbs" on page 122 command.

Domdsmc DB2rollforward

How to use the **domdsmc db2rollforward** command is described.

Purpose

The **domdsmc db2rollforward** command rolls a DB2 database forward to the specified point in time and marks the rollforward as complete. The DB2 database can be an alternate DB2 database or the Domino DB2 database. When the Domino DB2 database is enabled for rollforward recovery, the rollforward command must be executed after the restore. To recover a database to a time greater than the backup time, use the **/applylogs** parameter. The list of available DB2 databases to rollforward is obtained from the pending DB2 database file. To view the pending DB2 list, use the **domdsmc query db2rollforward** command. The **db2rollforward** command is only valid when the Domino DB2 database has been enabled for rollforward recovery.

When the **/applylogs** parameter is specified and the database is being rolled forward after a restore, it is not necessary to manually extract the logs.

If the DB2 database is being rolled forward after an inplace restore or an alternate database restore, the archived logs (required to roll forward the database) are automatically restored.

DB2 automatically archives the transaction log files when they become full. However, the user can also initiate an archive of the log to archive active log files and have them available for alternate database rollforward command.

Transaction log files that are stored on the Tivoli Storage Manager server are automatically restored as needed for a database recovery.

When a DB2 database is enabled for rollforward recovery and the database is used for an inplace restore, the Domino server cannot connect to the DB2 database until after the rollforward operation completes. As a result, the command output displays this message text:

```
Starting Domino DB2 database rollforward...
Initializing Domino connection...
Restart Analysis (0 MB): 100%
04/21/2012 12:02:57 AM  A RM error occurred.: An error occurred accessing the db
2 datasource.

DB2 CONNECTION ERROR:  Domino unable to connect to DB2 database 'DOMDB2' as user
 'db2admin'...
[IBM][CLI Driver] SQL1117N  A connection to or activation of database "DOMDB2" c
annot be made because of ROLL-FORWARD PENDING.  SQLSTATE=57019

DB2 CONNECTION ERROR:  set DEBUG_DB2CONNECT=0 to suppress this message.
04/21/2012 12:02:57 AM  Unable to initialize DB2 services.  DB2-based nsfs will
be unusable.: An error occurred accessing the db2 datasource.
```

There is no DB2 connection error and therefore, this message text can be ignored.

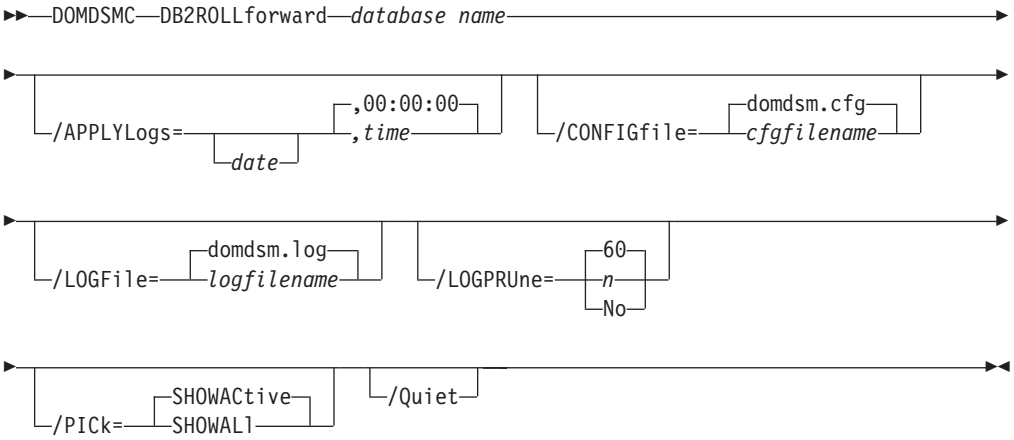

Parameters

dbname Specifies the DB2 database to rollforward. If not specified, the default alternate DB2 database (DB2ALTDBNAME) is used.

/APPLYLogs=*date,time*

Specifies that transaction log recovery for the restored databases is run if they are logged. The date and time values must be specified in the same date and time format defined in the Data Protection for IBM Domino preferences file. The transaction logs are applied to a specified point in time or to the current date and time if no date and time values are specified.

date Specify a date string in the active date format. When specified, transactions that are completed and committed before the specified date is applied to the restored database. The date that is specified must be after the backup date of the backup image that is being restored. The */pit* option can be used with the **restore** command to automatically restore the most recent full backup image that is run before the specified point in time.

The date must be specified with the date format that is defined in the Data Protection for IBM Domino preferences file. See "Domdsmc set" on page 115 for a list of available date formats.

time Specify a time string in the active time format. If you specify a date without the time, 00:00:00 on a 24-hour clock is used.

The time must be specified with the same time format that is defined in the Data Protection for IBM Domino preferences file. See "Domdsmc set" on page 115 for a list of available time formats.

/CONFIGfile=cfgfilename

Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/PICk=*SHOWACtive* | *SHOWALl*
Displays a list of database backups that match the `dbname` pattern that can be selected for restore. The pick list is displayed as a scrollable list from which you can select the database backups for restore.

You can specify:

SHOWACtive
Displays a list of active database backup versions.

SHOWALl
Displays a list of both active and inactive database backup versions. All the backup versions that match the **dbname** pattern are shown.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

Example

This command sequence shows rollforward processing for a full inplace DB2 database restore:

Command 1: `domdsmc query db2rollforward`

Output:

```
Backup Date        Size      Group     DB2 Database State
-----------       --------- ---------- -------
01/26/12   05:34:22   57.00MB   DOMINO     Pending
```

`Command 2: domdsmc db2rollforward DOMINO`

Output:

```
Starting Domino DB2 database rollforward...
Rollforward DB2 database DOMINO.
Rollforward of DOMINO completed successfully.
```

Domdsmc DB2selective

How to use the **domdsmc db2selective** command is described.

Purpose

The **domdsmc db2selective** command backs up DB2 Groups and the Domino DB2 database.

- If the Domino DB2 database is enabled for rollforward recovery, an online backup is run.
- If the Domino DB2 database is not enabled for rollforward recovery, Data Protection for IBM Domino cannot back it up.

- DB2 Group backups are only available when the Domino DB2 database is enabled for rollforward recovery.

Tip: When you are backing up multiple DB2 groups, increase the value of the Tivoli Storage Manager server COMMTIMEOUT option to avoid a backup failure because of a session timeout.

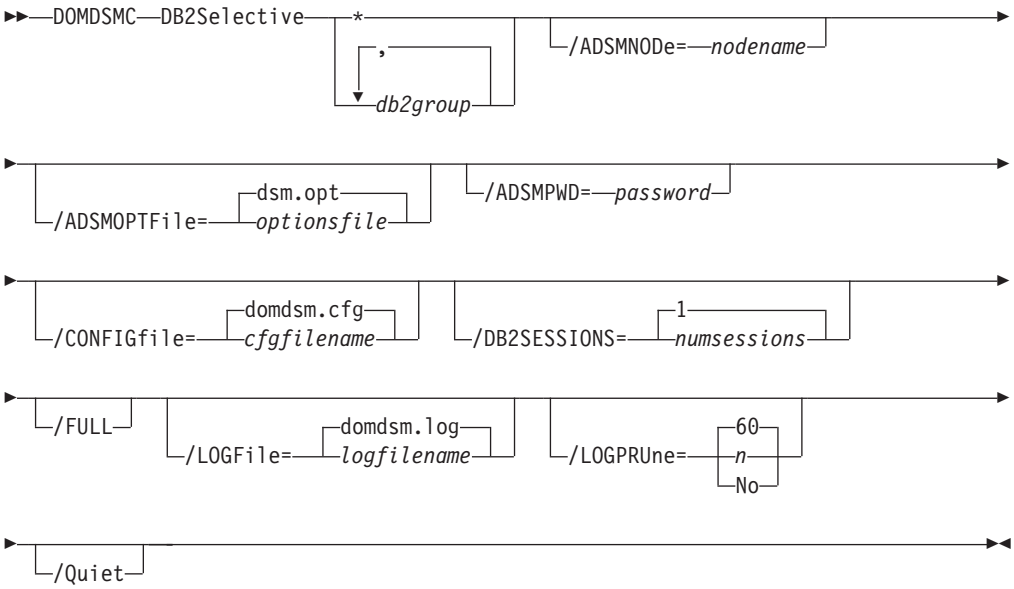

Parameters

***|db2group,db2group,...**
Specifies the DB2 Groups to back up. When a DB2 Group is not specified and the **/full** parameter is specified, a full DB2 database backup is run. Otherwise, a table space backup is run. The wildcard character asterisk (*) is used to specify a group of databases when used in the db2group. Multiple db2groups can be specified separated with commas.

/ADSMNODe=nodename
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify `passwordaccessgenerate` in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If `passwordaccess` is set to `generate` and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If passwordaccess is set to prompt and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the prompt.

If passwordaccess is set to prompt and you do not specify a password on the command line, then you are prompted for a password.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is domdsm.cfg.

/DB2SESSIONS=*numsessions*
Specify the number of Tivoli Storage Manager sessions that the DB2 Tivoli Storage Manager agent uses. You can specify 1-64 sessions. The default value is 1.

/FULL Specify whether a full DB2 database is backed up.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is domdsm.log.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the /logfile parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the /logprune option to override these defaults for one command run. Note, when the value of /logprune is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/Quiet Specifies that status information does not display. However, the information is written to the activity log.

Examples

Example 1: The following example backs up the Domino DB2 database:

```
domdsmc db2selective /full
```

Output example:

```
Starting Domino DB2 database backup...
Backing up DB2 database DOMINO.
Domino DB2 database backup completed successfully.
```

Example 2: The following example backs up the DB2 Groups `GRP1` and `GRP2`:

```
domdsmc db2selective GRP1,GRP2
```

Output example:

```
Starting Domino DB2 group backup...
Backing up DB2 group Default/GRP1, 1 of 2.
Backup of GRP1 completed successfully.

Backing up DB2 group Default/GRP2, 2 of 2.
Backup of GRP1 completed successfully
```

Domdsmc query DB2backup

How to use the **domdsmc query db2backup** command is described.

Purpose

domdsmc query db2backup lists DB2 backup objects.

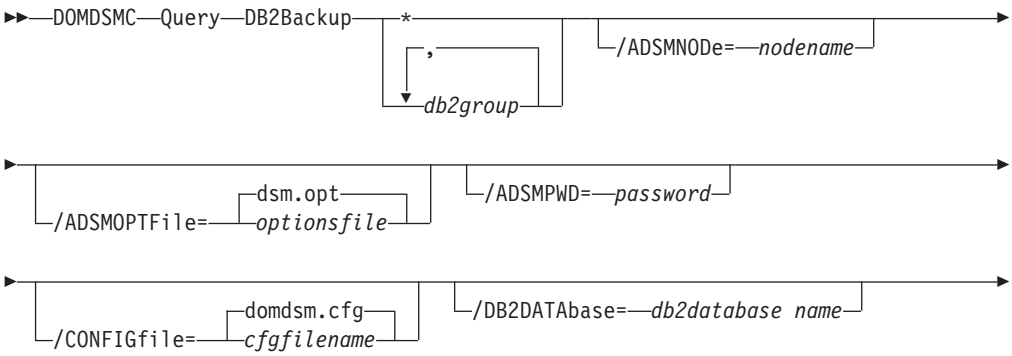

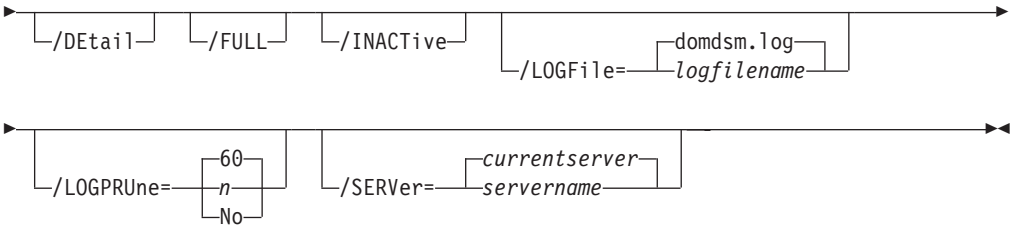

Parameters

***|db2group,db2group,...**
Specifies the DB2 Group to query.

/ADSMNODe=nodename
Specifies the Tivoli Storage Manager node name Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. The command-line value overrides the value in the Tivoli Storage Manager options file.

/ADSMOPTFile=optionsfile
Specifies the Tivoli Storage Manager options file name. The file name can include a fully qualified path name. If you do not specify a path, the installation directory and then the current directory are searched for the specified file. The default is `dsm.opt`.

/ADSMPWD=password
Specifies the Tivoli Storage Manager password Data Protection for IBM Domino uses to log on to the Tivoli Storage Manager server. If you specify **passwordaccess generate** in the Tivoli Storage Manager options file, then the password is not required. In this case, Data Protection for IBM Domino uses the password that is stored by the Tivoli Storage Manager API.

If **passwordaccess** is set to **generate** and you specify a password, the value is ignored unless a password for this node is not stored. In this case, the specified password is stored and used for the current command execution.

If **passwordaccess** is set to **prompt** and you specify a password on the command line, you are not prompted for a password. The command-line value overrides the requirement to prompt.

If **passwordaccess** is set to **prompt** and you do not specify a password on the command line, then you are prompted for a password.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/DB2DATAbase=*database name*
Specify the name of the alternate DB2 database to use for restore. If the **db2altdbname** parameter is not specified, the configuration option **db2altdbname** is used.

/DEtail
Specify whether to display a detailed output of the DB2 Groups and databases that are contained in the backup images.

/FULL Specify whether a full DB2 database backup image is queried.

/INACTive
Specify that both active and inactive backup objects are displayed. The default value is to display only the active backup objects.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. 60 days is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

/SERVer=*currentserver|servername*
Specifies the Domino server name. If not specified, Data Protection for IBM Domino uses the current Domino server.

Examples

Example 1:

The following example queries a list of active and inactive DB2 Group backups:

```
domdsmc query db2backup * /inactive
```

Output example:

```
   Domino Server: domino7     DB2 Database Name: DOMINO
   --------------             ------------------

Group Backup Date          Size      A/I  Type   Class         Group(TID)
-----------------          ----      ---  ----   ---------     ----------
01.08.2008 14:46:40       162.00B     A    T     DEFAULT       GRP1(6)
01.08.2008 14:44:37       166.00B     A    T     DEFAULT       GRP2(8)
```

Example 2: The following example queries (and displays) a list of active DB2 Group backups and the DB2 enabled Notes databases that are contained within the DB2 Group backup:

```
domdsmc query db2backup * /detail
```

Output example:

```
   Domino Server: domino7     DB2 Database Name: DOMINO
   --------------             ------------------

Group Backup Date          Size      A/I  Type   DB2 Group     Group Name
-----------------          ----      ---  ----   ---------     ----------
01.08.2008 11:46:52       160.00B     A    T      DEFAULT       GRP2

                           Size      Database Title   Database File
                           ----      --------------   -------------
                          160.00B    db2 nsf 1        ab2nsf1.nsf
                          160.00B    db2 nsf 2        db2b.nsf
                          160.00B    db2 nsf 1        db2g.nsf
                          160.00B    db2 nsf 1        xb2nsf1.nsf
```

Example 3:

The following example queries a full DB2 database backup and the DB2 Groups and DB2 enabled Notes databases that are contained within the full DB2 database backup:

```
domdsmc query db2backup * /detail /full
```

Output example:

```
  Domino Server: domino7     DB2 Database Name: DOMINO
  --------------             ------------------

Group Backup Date          Size      A/I  Type   Class         Group(TID)
-----------------          ----      ---  ----   ---------     ----------
01.08.2008 14:46:40       162.00B     A    T     DEFAULT       GRP1(6)

                           Size      Database Title   Database File
                           ----      --------------   -------------
                          162.00B    db2 nsf 1        db1.nsf
                          162.00B    db2 nsf 2        db2a.nsf
                          162.00B    db2 nsf 1        db2c.nsf
                          162.00B    db2 nsf 1        db2e.nsf

Group Backup Date          Size      A/I  Type   Class         Group(TID)
-----------------          ----      ---  ----   ---------     ----------
01.08.2008 14:44:37       166.00B     A    T     DEFAULT       GRP2(8)

                           Size      Database Title   Database File
                           ----      --------------   -------------
                          166.00B    db2 nsf 1        ab2nsf1.nsf
                          166.00B    db2 nsf 2        db2b.nsf
                          166.00B    db2 nsf 1        db2g.nsf
                          166.00B    db2 nsf 1        xb2nsf1.nsf
```

Domdsmc query DB2pendingdbs

How to use the **domdsmc query db2pendingdbs** command is described.

Purpose

The **domdsmc query db2pendingdbs** parameter lists the DB2 enabled Notes databases that are pending activation. These databases reside in an alternate database and the activate process (**"Domdsmc DB2activatedbs" on page 122**) copies them to the Domino DB2 database. The alternate DB2 database is considered available when it is not manually deleted through DB2 server interface, overwritten by another restore operation, or removed from the list of DB2 databases (that contain DB2 enabled Notes databases) available for activation.

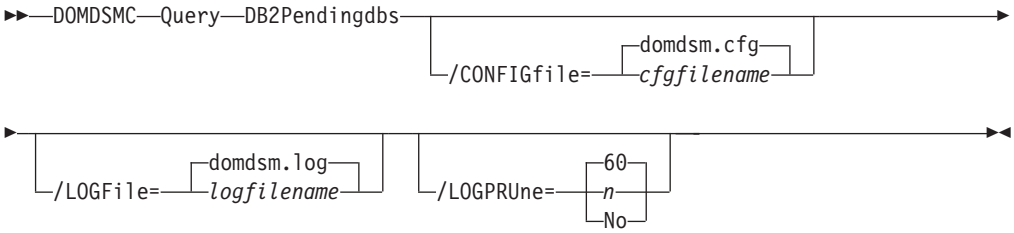

Parameters

/CONFIGfile=cfgfilename

Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file.

You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No

Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

Examples

Example 1: The following example queries (and displays) a list of DB2 enabled Notes databases that are pending activation:

```
domdsmc query db2pendingdbs
```

Output example:

```
   Domino Server: Server01
   --------------

   Backup Date           Size      Group    DB2 Database  Database
--------------------  -----------  -----    ------------  --------
01/21/2012 11:53:30        64.00B   GRP1     DOMALT1       db2nsf1.nsf
01/21/2012 11:53:30        64.00B   GRP1     DOMALT1       db2nsf2.nsf
01/21/2012 11:53:30        64.00B   GRP1     DOMALT1       db2nsf3.nsf
01/21/2012 11:53:30        64.00B   GRP1     DOMFULL1      db2nsf1.nsf
01/21/2012 11:53:30        64.00B   GRP1     DOMFULL1      db2nsf2.nsf
01/21/2012 11:53:30        64.00B   GRP1     DOMFULL1      db2nsf3.nsf
01/21/2012 11:53:30        64.00B   GRP2     DOMFULL1      db2nsf4.nsf
01/21/2012 11:53:30        64.00B   GRP2     DOMFULL1      db2nsf5.nsf
01/21/2012 11:53:30        64.00B   GRP2     DOMFULL1      db2nsf6.nsf
```

Example 2: The following example queries a list of DB2 enabled Notes databases that are pending activation. However, there are no databases waiting for activation:

```
domdsmc query db2pendingdbs
```

Output example:

```
ACD5418I There are no databases pending activation.
```

Domdsmc query DB2rollforward

How to use the **domdsmc query db2rollforward** command is described.

Purpose

Domdsmc query db2rollforward lists the DB2 database rollforward status.

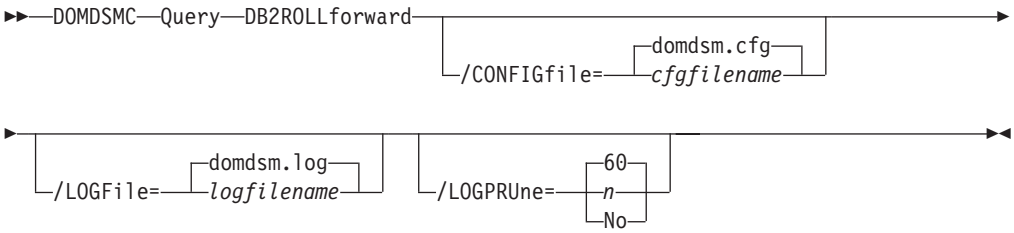

Parameters

/CONFIGfile =*cfgfilename*

Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file is in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=*logfilename*

Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file.

The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the **/logfile** parameter to specify a different log file for each instance used. This directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=*60 | n |No*

Specifies whether to prune log entries. By default, log pruning is enabled and run daily. You can use the **set** command to do the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the **/logprune** option to override these defaults for one command run. When the value of **/logprune** is a number, the prune is run even if one was run for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning. This is the default.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the *timeformat* or *dateformat* parameters can result in an undesired pruning of the Data Protection for IBM Domino log file. If you are running a command that prunes the log file and the value of the *timeformat* or *dateformat* parameter changes, do one of the following actions to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the **/logfile** parameter or *logfile* setting.

Examples

Example 1: The following example queries a list of DB2 databases that are available for rollforward processing after a restore operation:

```
domdsmc query db2rollforward
```

Output example:

```
Backup Date         Size      Group   DB2 Database State
-----------  --------- ---------- -------  ---------
01/26/08 05:34:22   57.00MB   DOMINO       Pending
```

Example 2: The following example queries (and displays) a list of DB2 databases that are available for rollforward processing after an alternate DB2 database restore:

```
domdsmc query db2rollforward
```

Output example:

```
                  DB2 Database Rollforward Status
                  -------------------------------

   Domino Server: polar1
   --------------

 Backup Date  Size  Group  DB2 Database State
 ----------  ---------- ----------------- -
 01/19/08   13:11:44     1,078.00MB    GRP8       DOM_ALT         Pending
 01/19/08   14:12:01      659.00MB    GRP5       DOM_ALT1        Pending
 01/18/08   13:36:46     1,031.00MB     -         DOM_FULL        Pending
```

Domdsmc updatedb2pwd

How to use the **domdsmc updatedb2pwd** command is described.

Purpose

The **domdsmc updatedb2pwd** parameter updates the DB2 user password. The password is required to access the DB2 instance where the Domino DB2 database resides. Data Protection for IBM Domino prompts the user for the password the first time and saves the password encrypted in a file. The password is read from this file when access to DB2 is required. The command allows the user to change the password in the file in case the DB2 user password is changed. If you do not enter the old and new passwords on the command, you are prompted for them. When Data Protection for IBM Domino prompts you for the passwords, the password is not displayed on the screen.

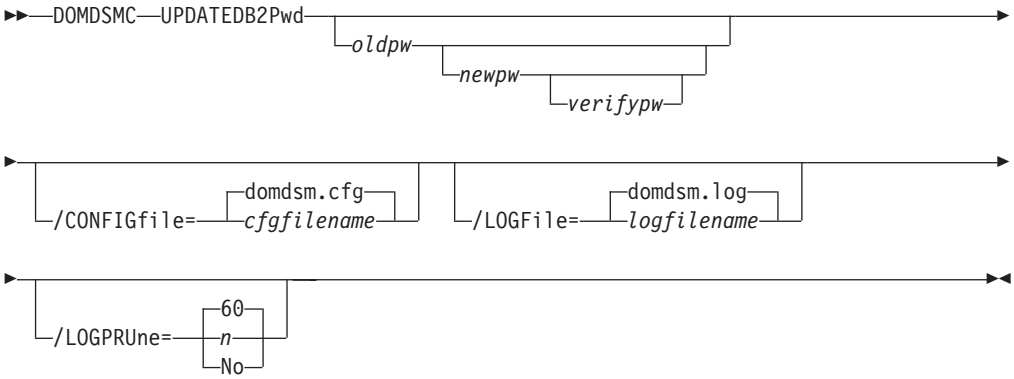

Parameters

oldpw The current password to change. You are prompted for this value if omitted.

newpw The new password. You are prompted for this value if omitted. When you choose a new password, you can use 1 - 64 characters.

Valid password characters are as follows:

A-Z Any letter, A through Z, upper-case, or lower-case

0–9 Any number, 0 - 9

+ Plus

. Period

_ Underscore

- Hyphen

& Ampersand

A password is not case-sensitive.

verifypw
The verify password is used to validate the password that is entered for **newpw**. You are prompted for this value if omitted.

/CONFIGfile=cfgfilename
Specifies the name of the Data Protection for IBM Domino preferences file. The file name can include a fully qualified path. If you do not specify a path, it is assumed the preferences file in the directory where Data Protection for IBM Domino is installed.

The default preferences file is `domdsm.cfg`.

/LOGFile=logfilename
Specifies the name of the activity log that is generated by Data Protection for IBM Domino. The log file name is used for the current command and does not update the default log file that is stored in the preferences file. You can use the **set** command to change the default log file name that is stored in the preferences file. The command-line parameter can be used to override the default for one command run. If the specified file does not exist, it is created. If it does exist, new log entries are appended to the file. The file name can include a fully qualified path. If you do not specify a path, the file is written to the directory where Data Protection for IBM Domino is installed.

The default log file is `domdsm.log`.

When you are using multiple simultaneous instances of Data Protection for IBM Domino to run operations, use the `/logfile` parameter to specify a different log file for each instance. Using this parameter directs logging for each instance to a different log file and prevents interspersed log file records. Failure to specify a different log file for each instance can result in unreadable log files.

/LOGPRUne=60|n|No
Specifies whether to prune log entries. By default, log pruning is enabled and done daily. You can use the **set** command to do one of the following actions:

- Change the defaults so that log pruning is disabled
- Change the number of days log entries are saved

You can use the `/logprune` option to override these defaults for one command run. Note, when the value of `/logprune` is a number, the prune is done even if one is complete for the day.

You can specify:

60 Specifies that log entries are saved for 60 days before pruning.

n Specifies the number of days to save log entries. The range of values is 0 - 9999. A value of 0 deletes all entries in the log except for the current command run entries.

No Do not prune the log.

Changes to the value of the **timeformat** or **dateformat** parameter can result in an undesired pruning of the log file. If you are running a command that

prunes the log file and the value of the **timeformat** or **dateformat** parameter changes, run one of the following to prevent pruning of the log file:

- Make a copy of the existing log file.
- Specify a new log file with the `/logfile` parameter or `logfile` setting.

Example

The following example changes the DB2 user password to *secret*:

```
domdsmc updatedb2pwd oldpassword secret secret
```

Output example:

```
ACD0260I Password successfully changed.
```

Chapter 4. Protecting IBM Domino Server data

Details of how to use Data Protection for IBM Domino to protect Domino servers are provided.

Automating backups

Use the Tivoli Storage Manager scheduler with Data Protection for IBM Domino to automate online backups of Domino server databases.

Install the latest Tivoli Storage Manager backup-archive client. The backup-archive client must be on the same system as Data Protection for IBM Domino to use the scheduler service.

After Data Protection for IBM Domino is registered to a Tivoli Storage Manager server and installed and configured on the Domino server, run the following steps:

1. **On the Tivoli Storage Manager server:**
   a. Define a schedule in the policy domain to which Data Protection for IBM Domino is registered.
   b. Associate the Data Protection for IBM Domino node to the defined schedule.
2. **On the Domino server where Data Protection for IBM Domino is installed:**
   a. Install the Tivoli Storage Manager scheduler client as a Windows service. If a scheduler exists for the regular Tivoli Storage Manager backup-archive client, install and configure another scheduler for Data Protection for IBM Domino. The Tivoli Storage Manager scheduler must have a different node name from the regular Tivoli Storage Manager Backup-Archive client.
   b. Define a command file that contains the Data Protection for IBM Domino commands to run a backup.
   c. If you are running in a cluster server environment, install the Tivoli Storage Manager scheduler as a Windows service `on both cluster nodes`.
   d. If you are running in a cluster server environment, create a cluster resource that represents the Tivoli Storage Manager scheduler. Verify that the cluster resource is started.
   e. Start the scheduler that is installed for Data Protection for IBM Domino.

Setting up a scheduler scenario

An example scenario of the Tivoli Storage Manager scheduler is provided. Scheduler considerations and tasks are presented.

The example that is outlined assumes the following details:

- Data Protection for IBM Domino is registered to a Tivoli Storage Manager server:
  - The node name is `mynode`.
  - The password is `mypassword`.
  - The policy domain is `mydomain`.
- For Domino NSF databases, there are several events that can be scheduled. For this example, assume that the Domino server is running archival logging and using the backup strategy of full backups plus transaction log archives. For this backup strategy, it is suggested that you do periodic archival of the transaction

logs. You must also do incremental backups, selective backups of the logged databases, and inactivation of transaction logs. Each of these tasks must have its own schedule as they must be done at various times.

- For DB2 enabled Notes databases, a scheduled backup strategy can consist of a weekly full DB2 database backup and a daily DB2 backup for the most critical DB2 Group.
- This example shows how to schedule hourly archiving of the transaction logs. From this example and the sample files that are installed, you can schedule the remaining tasks that must be done.

This method is flexible because you can define a command file with any set of commands you choose. You can also use the same method to schedule other back ups on Windows.

Defining a Tivoli Storage Manager server schedule

How to define a schedule for archiving transaction logs is provided.

1. Enter the following command to define the schedule to do an hourly archival of the transaction logs.

   Note: You can enter this command on the server console or from an administrative client. The administrative client does not have to be running on the same system as the Tivoli Storage Manager server.

   ```
   def sched domagents dom_hourly_archive desc="Domino Hourly Archive"
   action=command objects="c:\domarc.cmd" priority=2 starttime=5:00
   duration=15 duru=minutes period=1 perunits=hours dayofweek=any
   ```

   Tivoli Storage Manager displays this message:

   ```
   ANR2500I Schedule DOM_HOURLY_ARCHIVE
   defined in policy domain DOMAGENTS.
   ```

2. To associate Data Protection for IBM Domino to this schedule, issue the following command:

   ```
   define association domagents dom_hourly_archive mars
   ```

   Tivoli Storage Manager displays this message:

   ```
   ANR2510I Node MARS associated with schedule
   DOM_HOURLY_ARCHIVE in policy domain DOMAGENTS.
   ```

   At this point, a schedule is defined on the Tivoli Storage Manager server that runs a command file called `c:\domarc.cmd`. The schedule starts around 5:00 am. The schedule is re-executed every hour and can start on any day of the week.

   Note: If you want to confirm that the schedule and association is set correctly, you can use the Tivoli Storage Manager administrative commands **query schedule** and **query association**. See the appropriate Tivoli Storage Manager Administrator's Guide for your server platform for more information.

Running Tivoli Storage Manager client tasks

Set up a Tivoli Storage Manager schedule for Domino with the backup archive-client setup wizard.

This example assumes that the Tivoli Storage Manager backup-archive client is installed on the Domino server in the default installation directory (`c:\program files\tivoli\tsm\baclient`) directory. It also assumes that the Data Protection for IBM Domino for the Domino server is stored in the default installation directory (`c:\program files\tivoli\tsm\domino`) directory. The options files in each of the installation directories must be updated so that the communication parameters point to the Tivoli Storage Manager server.

1. Start the backup archive client GUI.
2. In the menu, go to **Utilities –> Setup Wizard**.
3. When the TSM Client Configuration Wizard opens, select **Help me configure the TSM Client Scheduler** and click Next.
4. Follow the prompts to install a new scheduler service. You need the following information:
   - Schedule name.
   - Tivoli Storage Manager options file.
   - Tivoli Storage Manager node name.
   - Tivoli Storage Manager node name password.
   - Schedule log file.
   - Schedule error log file.
5. The options file that is defined by Data Protection for IBM Domino is used by the scheduler when it is validating the node and password. The options file is also used when it contacts the Tivoli Storage Manager server for schedule information. The example assumes that the `dsm.opt` file is updated so that the communication parameters point to the Tivoli Storage Manager server to which the Domino databases are to be backed up.

   If the following message is shown:

   ```
   A communications error occurred connecting to the
   Tivoli Storage Manager server
   ```

   Ensure that the options file contains entries that point to the correct Tivoli Storage Manager server. Also, ensure that the Tivoli Storage Manager server is running.
6. Create a batch file that is called `c:\domarc.cmd`. In the directory where Data Protection for IBM Domino was installed, there is a sample command file, `domarc.smp`. You can use this sample as a starting point to coding this command file.

   When you are using the Tivoli Storage Manager scheduler to run the commands in a command file, you must use the complete path names for all file names and non-system commands. The scheduler runs from the Windows 2003 system directory where the scheduler service looks for input and produces its output by default.
7. The scheduler is installed and configured, but is not started.

   To start the service, issue the following command in the Windows console window:

   ```
   net start "Tivoli Storage Manager Data Protection for Domino Archive Schedule"
   ```

   The following output is displayed:

   ```
   The Tivoli Storage Manager Data Protection for Domino Archive
   Schedule service is starting.

   The Tivoli Storage Manager Data Protection for Domino Archive
   Schedule service was started successfully.
   ```

   Note, because the **/autostart:yes** option is used, the Tivoli Storage Manager scheduler is automatically started each time the Windows system is rebooted.

   Your system is now ready to run automatic hourly archival of the transactions logs.

The Tivoli Storage Manager client scheduler and associated log files

Guidance about editing your Data Protection for IBM Domino options files, and collecting and viewing log files for scheduled backups with the Tivoli Storage Manager server is provided.

Tivoli Storage Manager server prompted scheduling

Editing the options files, viewing the domsched.log file and other logs for status, specifying the Tivoli Storage Manager password, and creating nodes for backups are detailed.

To use the Tivoli Storage Manager server prompted scheduling mode, ensure the dsm.opt options file has the tcpclientaddress and tcpclientport options specified. If you want to run more than one scheduler service, use the same tcpclientaddress. However, you must use different values for tcpclientport (in addition to the different node names). An example of running more than one scheduler service is when you are scheduling Data Protection for IBM Domino at the same time as the regular backup-archive client. Server-prompted scheduling is supported only when TCP/IP communication is being used. By default, Tivoli Storage Manager uses the client polling schedule mode.

If changes that affect the scheduler are made to the Data Protection for IBM Domino options file, the scheduler must be restarted to pick up the changes. Such changes might include the Tivoli Storage Manager server address, the schedule mode, or the client TCP address or port. Restarting is done by issuing the following commands:

```
net stop "Tivoli Storage Manager Data Protection for Domino Archive Schedule"
net start "Tivoli Storage Manager Data Protection for Domino Archive Schedule"
```

.

The domsched.log

The domsched.log file contains status information for the Tivoli Storage Manager scheduler. In this example, the file is at this path:

```
c:\program files\tivoli\tsm\baclient\domsched.log
```

You can override this file name by specifying the schedlogname option in the Tivoli Storage Manager options file.

Output from scheduled commands is sent to the log file. After scheduled work completes, check the log to ensure that it completed successfully.

```
Scheduled event eventname completed successfully
```

This log entry indicates that Tivoli Storage Manager successfully issued the scheduled command that is associated with the eventname. No attempt is made to determine the success or failure of the command. In the dsmsched.log file, there is a log entry with the following text, Finished command. Return code is: 0. This log entry indicates that the command file started successfully. The return code does not indicate the outcome of the scheduled command. To determine the success or failure of the scheduled command view the Data Protection for IBM Domino log file.

If any scheduled backups fail, the scheduler script exits with the same error code as the failed backup command. An error code without a zero means that the backup failed. Follow the instructions in the troubleshooting section when the command fails.

Specified log files, domsarc.log

Data Protection for IBM Domino creates its own log file with statistics about the archived transaction log objects when the /logfile parameter is specified during the domdsmc command. In the domarc.smp file, the log file is domsarc.log. This file is different from the Tivoli Storage Manager scheduler log file and must also be different from the file to which the domdsmc command output is redirected. In the example, domarc.smp, this file is domasch.log.

Specifying the Tivoli Storage Manager password

When Data Protection for IBM Domino is not configured to automatically generate the Tivoli Storage Manager password, it expires. Use the domdsmc command to specify the Tivoli Storage Manager password. Use the /adsmpwd option in the command file that is being run by the scheduler, domarc.cmd .

Use a node and a scheduler service for backups

The Tivoli Storage Manager client scheduler allows only one scheduler process at a time. Other schedules must wait for the first scheduled backup to complete before they can begin. When a scheduled backup is processing at the same time a backup is scheduled to begin archivelog, it might cause an issue.

For example, assume that there are two scheduled Data Protection for IBM Domino backups that run under the same node and scheduler service. The database backup runs at 06:00 and the archivelog backup runs every hour. If the database backup takes longer than 1 hour, the archivelog backup fails. You can avoid this issue by creating a node and a scheduler service for the archivelog backup. This new node and service starts the batch .cmd file for the archivelog backup only. Ensure that the Data Protection for IBM Domino dsm.opt file contains the name of the regular Data Protection for IBM Domino node. For example:

- Data Protection for IBM Domino backups use node name DOMClient as specified in the Data Protection for IBM Domino dsm.opt file.
- The new node name for the archivelog backup is DOMSched.
- The scheduler is associated with the DOMSchedt node. The batch .cmd file that runs the domdsmc archivelog command, points to the Data Protection for IBM Domino dsm.opt file that contains the node name DOMClient.

As a result, backup objects are associated correctly to the Data Protection for IBM Domino node, not the node for the scheduler.

Setting up other schedules

Important information to consider when you are setting up other schedules.

Run several other schedules for Data Protection for IBM Domino as part of a complete backup strategy for Lotus Domino databases and transaction logs.

Use the **SESSIONS** command to improve performance when scheduling tasks in Data Protection for IBM Domino. More information is available in the Performance section.

Incremental backup of all databases

- Frequency is daily.
- Sample command file, `dominc.smp`, exists in the Data Protection for IBM Domino installation directory
- `Domdsmc` log file is created, `dominc.log`.
- Output is redirected, `domisch.log`

Selective backup of all logged databases

- Frequency is weekly, maybe on a Saturday.
- Sample command file, `domsel.smp`, exists in the Data Protection for IBM Domino installation directory
- `Domdsmc` log file is created, `domsel.log`.
- Output is redirected, `domssch.log`.

Inactive logs

- Frequency is weekly, maybe on a Sunday, to ensure that the selective backup completes.
- Sample command file, `domina.smp`, exists in the Data Protection for IBM Domino installation directory
- `Domdsmc` log file is created, `domina.log`.
- Output is redirected, `domiasch.log`.

Back up all DB2 Groups (DB2 enabled Notes databases)

- Frequency is daily.
- Sample command file, `domdb2grp.smp`, exists in the Data Protection for IBM Domino installation directory
- `Domdsmc` log file is created, `domdb2grp.log`.
- Output is redirected, `domdb2grp.out`.

Full DB2 database backup (DB2 enabled Notes databases)

- The frequency is weekly.
- Sample command file, `domdb2db.smp`, exists in the Data Protection for IBM Domino installation directory
- `Domdsmc` log file is created, `domdb2db.log`.
- Output is redirected to `domdb2db.out`.

Sample command file

Sample command files are provided.

Sample

This is an example of the domarc.smp file.

```
@ECHO OFF
rem  ====================================================================
rem   Sample Command File - domarc.smp
rem
rem   Sample command file containing commands to do a scheduled archive
rem   of transaction logs to an IBM Tivoli Storage Manager server.
rem
rem   This file is meant to be executed by the IBM Tivoli Storage Manager
rem   central scheduler in response to a defined schedule on the
rem   IBM Tivoli Storage Manager server.
rem
rem   Complete paths must be given for all file names and non-system
rem   commands.
rem
rem   Copy this file to domarc.cmd and edit it to match your
rem   local environment.
rem
rem  ====================================================================

rem  ====================================================================
rem   Replace "X:" with the drive where the Data Protection for IBM
rem   Domino is installed.
rem  ====================================================================

set dom_dir="X:\Program Files\Tivoli\TSM\domino"

cd /d %dom_dir%

rem  ====================================================================
rem   The 2 lines below put a date and time stamp in a log file for
rem   you.
rem
rem   Note: You can change "domarc.log" to whatever you prefer.
rem  ====================================================================
echo Current date is: >> domarc.log
date /t < NUL >> domarc.log
echo Current time is: >> domarc.log
time /t < NUL >> domarc.log

rem  ====================================================================
rem   Now call the command line to do the archive of the logs:
rem
rem   Note: You can change "domasch.log" to whatever you prefer.
rem  ====================================================================
start /B domdsmc archivelog /adsmoptfile=dsm.opt
/logfile=domasch.log >> domarc.log
```

Recovering from loss of Domino transaction logs for NSF databases

How to recover from a loss of the Domino server, and transaction log failure for NSF databases is outlined. How to use the archived transaction log files for database recovery is provided.

Before you begin

Recovery of a database to the most current available backup requires restoring the last full backup plus applying updates to that backup from the archived transaction log files. For archived transaction log files to be used for database recovery, the current transaction log ID must match that of the archived log files.

Note: If the current transaction log is lost, creating one results in a new log ID and thus the archived log files would not be usable for database recovery.

About this task

When you are using archival transaction logging, archived transaction log files contain updates to logged databases that might not yet be captured in a full database backup.

Procedure

1. Recover the non-database Domino server files. If necessary, reinstall but do not configure the server. Restore the non-database Domino files that include `notes.ini`, `cert.id`, and `server.id` with your file backup solution such as the Backup-Archive Client. Make sure that the new installation is configured in the same way as the damaged one. It must have the same directory structure, directory location, and `logdir` path. Do not start the new server.
2. Using a text editor, modify the `notes.ini` file for the Domino server with this setting: `TRANSLOG_Status=0`
3. Using Data Protection for IBM Domino, restore the transaction log file to be used in the log recovery procedure. This is the last transaction log file that is archived before the loss of the active transaction log.

   Note: Use the `-replace=no` option as an added safety measure.
4. Close the Data Protection for IBM Domino GUI (if in use).
5. Delete the contents of the Domino transaction log directory except for the log file that is restored in Step 3.
6. Modify the `notes.ini` file for the Domino server with these settings: `TRANSLOG_Recreate_Logctrl=1TRANSLOG_Status=1`
7. Restore but do not activate the databases that you want to recover to the latest state within the archived log extents with Data Protection for IBM Domino.
8. Use Data Protection for IBM Domino to activate the databases you are recovering and apply transaction logs. The **`TRANSLOG_Recreate_Logctrl`** parameter in the `notes.ini` file is automatically reset to 0.
9. Start the Domino server. With the disaster recovery complete, it is now safe to start the Domino server and run server tasks and functions.
10. Use the Selective backup function in Data Protection for IBM Domino to run full backups of all databases. Having full backups ensures recoverability with subsequent transaction log files.

11. Use Data Protection for IBM Domino to archive the transaction log. The transaction log file that is used in the recovery procedure is modified and available for archiving. This transaction log will also have the ID of the current logger.

NSF databases restore to alternate server and alternate partition

To reduce demands on the Domino production server, restore data to an alternate server.

A restore operation involves two steps. First, the backup copies of the databases are retrieved for the Tivoli Storage Manager server. Second, the recorded transactions in the log files are applied to the databases. If the transaction log files required to recover the databases are archived, they can be retrieved from the Tivoli Storage Manager server. These steps can have an impact on performance in the processor, and in disk input and output.

The transaction log directory must be on a dedicated physical disk drive for optimal performance. When a dedicated physical disk drive is used, the Domino server can write transactions sequentially to the log. Using this disk is faster than writing transactions to random nonsequential parts of a disk. If the restore operation is run on a Domino production server, the restore and application of the transactions interferes with sequential writing of transactions to the log. Performance of the Domino server is affected, and the time that is required to run the restore operation increases. The application of the transaction logs competes for processor cycles with the Domino server.

Restore operations must be run on an alternate server or on an alternate partition for these reasons.

An alternate server restore is the preferred method since the restore operation has no impact on the performance of the production Domino server. However, the production Domino server and the server on the alternate partition can use separate disk drives for their transaction log directory. If separate disk drives are used, the production Domino server access to the transaction log is not affected by the restore operation on the alternate partition.

Note: With Domino 6 you can specify an alternate path to restore the archived transaction logs. When a separate disk drive is used, the alternate path feature helps to minimize the cost of a restore operation.

Restoring NSF databases to an alternate server

Instructions for how to restore NSF databases to an alternate server are described.

About this task

This procedure describes how to use an alternate server to restore logged databases.

Production Server Domino Environment

- Installation directory: `D:\Lotus\Domino`
- Data Directory: `D:\Lotus\Domino\Data`
- Database to be restored: `restoredb.nsf`

Alternate Server Domino Environment

- Installation directory: E:\Lotus\Domino
- Data Directory: E:\Lotus\Domino\Data

Procedure

1. Install Domino server on a separate system.
   a. The level of Domino server that is installed must be the same as on the production server. Do not configure this Domino server.
   b. If you are using an existing Domino server, make sure that the server is stopped.
2. Install Data Protection for IBM Domino on the same system. Run the following tasks:
   a. Update the dsm.opt file so it contains the same settings as the one on the production server.
   b. Verify that you can successfully run the **domdsmc q adsm** command.
3. Create the following directories on the alternate server:
   a. A directory to contain the restored databases. If you are using an existing directory, make sure that the directory is empty. For example:E:\Lotus\Domino\Data\restoredb
   b. A directory to contain the restored log files. If you are using an existing directory, make sure that the directory is empty. For example: F:\alternatelog
4. Create a notes.ini file on the alternate server with the following values: [Notes] Directory=<directory for restored databases> KeyFilename=<directory for restored databases>\server.id TRANSLOG_Status=0 TRANSLOG_Style=1 TRANSLOG_Path=<directory for restored transaction logs> TRANSLOG_MEDIAONLY=1. This notes.ini file can be stored in any directory of your choice.
   a. If you place the notes.ini file in the alternate server data directory, save a copy of the existing notes.ini file. For example:rename notes.ini notes.save
   b. If you place the notes.ini file in a directory other than the alternate server data directory, update the Data Protection for IBM Domino preferences file, domdsm.cfg, to point to the location of the notes.ini file: DOMDSMC SET Notesinipath=<directory for notes.ini>
   c. This notes.ini file is used only during this alternate server restore process. Note, transaction logging is disabled. For example, E:\Lotus\Domino\notes.ini[Notes] Directory=E:\Lotus\Domino\Data\restoredb KeyFilename=E:\Lotus\Domino\Data\restoredb\server.id TRANSLOG_Status=0 TRANSLOG_Style=1 TRANSLOG_Path=F:\alternatelog TRANSLOG_MEDIAONLY=1
5. Place a copy of the server.id file, from the production Domino server, on the alternate server in the directory that is created for restored databases.
6. Run an archive of the transaction log on the production server. For example: domdsmc archivelog To use the Data Protection for IBM Domino GUI, run the **domdsm** command, archive the transaction log, and then close the GUI.
7. Restore one of the following on the alternate server:
   a. The last archived transaction log file is the transaction log file to be used in the log recovery procedure. For example: domdsmc restorelogarchive.To use the Data Protection for IBM Domino GUI, run the **domdsm** command, restore the transaction log, then close the GUI.

b. A transaction log file to be restored from an old Logger ID. This might be necessary if you are trying to restore and apply transactions for a logged database that used an old Logger ID. Run the **restorelogarchive** command with the `pick` option, and choose the wanted log extent. For example: `domdsmc restorelogarchive logname /pick=showall`

8. On the alternate server, modify the `notes.ini` file to enable transaction logging: `TRANSLOG_Status=1`. This is the `notes.ini` file that is created in Step 3 for the alternate server restore process only.

   Note: The IBM Domino server must be restarted and then stopped so that it can recognize the changes that are made to the `notes.ini` file.

9. On the alternate server, restore but do not activate the databases that you want to recover to their latest state. Activation at this step triggers the Domino transaction log recovery process, which requires considerable processing time. For example,`domdsmc restore restoredb.nsf`. To use the Data Protection for IBM Domino GUI, run the **domdsm** command, restore the `restoredb.nsf` database, then close the GUI.
10. On the alternate server, activate the databases you are recovering and apply transaction logs. For example, `domdsmc activate /applylogs`. To use the Data Protection for IBM Domino GUI, run the `domdsm` command, activate, and apply logs to the `restoredb.nsf` database, then close the GUI.
11. Now you can do the following tasks:
    - Copy the recovered databases to the production Domino server, or
    - Access the recovered databases through a remote Notes client to copy individual documents.

    Do not attempt to open or access the restored databases with the alternate Domino server if the databases are to be copied to the production Domino server. If you access them with the alternate Domino server, they require corrections to resolve inconsistencies on the production Domino server.
12. If the Domino server used for the recovery is a configured server and you saved the `notes.ini` file, copy that `notes.save` file back to `notes.ini` to be able to start the server.

Restoring NSF databases to an alternate partition

How to restore NSF databases to alternate partitions is described.

Before you begin

This procedure assumes the following environment:

Domino Environment

- Installation directory: `D:\Lotus\Domino`
- Production Partition Data Directory: `D:\production`
- Alternate Partition Data Directory: `E:\alternate`
- Database to be restored: `restoredb.nsf`

Data Protection for IBM Domino

- Production Server Preferences File: `production.cfg`.

  The `notesinipath` option in the `production.cfg` file specifies `D:\production`.
- Alternate Partition Preferences File: `alternate.cfg`.

The notesinipath option in the alternate.cfg file specifies E:\alternate.

About this task

Procedure

1. Install an alternate partition if one is not available. See your Domino Server documentation for information about how to install an alternate partition.
   a. Do not configure this alternate partition.
   b. If you are using an existing alternate partition, make sure the server on that partition is stopped.
2. Create the following directories :
   a. A directory to contain the restored databases. If you are using an existing directory, make sure that the directory is empty. For example:E:\ alternate\restoredb
   b. A directory to contain the restored log files. If you using an existing directory, make sure that the directory is empty. For example: F:\alternatelog
3. Create a notes.ini file with the following values: [Notes] Directory=<directory for restored databases> KeyFilename=<directory for restored databases>\server.id TRANSLOG_Status=0 TRANSLOG_Style=1 TRANSLOG_Path=<directory for restored transaction logs> TRANSLOG_MEDIAONLY=1 This notes.ini file can be in any directory of your choice.
   a. If you place the notes.ini file in the alternate partition data directory, save a copy of the existing notes.ini file. For example: rename notes.ini notes.save
   b. If you place the notes.ini file in a directory other than the alternate partition data directory, update the Data Protection for IBM Domino preferences file, domdsm.cfg by default, to point to the location of this notes.ini file: domdsmc set notesinipath=<directory for notes.ini> /config=alternate.cfg
   c. This notes.ini file is used only during this alternate partition restore process. Note, transaction logging is disabled. For example, in the file C:\alternate\notes.ini[Notes] Directory=E:\alternate\restoredb KeyFilename=E:\alternate\restoredb\server.id TRANSLOG_Status=0 TRANSLOG_Style=1 TRANSLOG_Path=F:\alternatelog TRANSLOG_MEDIAONLY=1
4. Place a copy of the server.id file from the production Domino server, on the alternate server in the directory that is created for restored databases.
5. Perform an archive of the transaction log on the production server. You can apply the latest updates from the transaction log to the restored database. For example: domdsmc archivelog /config=production.cfg To use the Data Protection for IBM Domino GUI, run the **domdsm** command with the**/config=production.cfg** parameter, archive the transaction logs, then close the GUI.
6. Restore one of the following on the alternate partition:
   a. The last archived transaction log file which is the log file to be used in the log recovery procedure. For example: domdsmc restorelogarchive /config=alternate.cfg To use the Data Protection for IBM Domino GUI, run the **domdsm** command with the**/config=alternate.cfg** parameter, restore the transaction logs, then close the GUI.

b. A transaction log file to be restored from an old Logger ID. Use this log file if you are trying to restore and apply transactions for a logged database that used an old Logger ID. Run the **restorelogarchive** command with the `pick` option and choose the log extent. For example:`domdsmc restorelogarchive logname /config=alternate.cfg /pick=showall`.

See "**Domdsmc activatedbs**" on page 50 for a description of when this type of restore might be necessary.

7. On the alternate partition, modify the `notes.ini` file to enable transaction logging: `TRANSLOG_Status=1`. This is the `notes.ini` file that is created in Step 3 for the alternate partition restore process only.
8. On the alternate partition, restore (but do not activate) the databases you want to recover to their latest state.

   Attention: Warning! Activation at this step triggers the Domino transaction log recovery process which requires considerable processing time. For example: `domdsmc restore restoredb.nsf /config=alternate.cfg` To use the Data Protection for IBM Domino GUI, run the **domdsm** command with the**/config=alternate.cfg** parameter, restore the `restoredb.nsf` database, then close the GUI.
9. On the alternate partition, activate the databases that you are recovering and apply transaction logs. For example: `domdsmc activate /applylogs /config=alternate.cfg`To use the Data Protection for IBM Domino GUI, run the `domdsm` command with the**/config=alternate.cfg** parameter, activate, and apply logs to the `restoredb.nsf` database, then close the GUI.
10. You can now do the following tasks:
    - Copy the recovered databases to the production Domino server, or
    - Access the recovered databases through a remote Notes client to copy individual documents.
    - Do not attempt to open or access the restored databases with the alternate Domino server if the databases are to be copied to the production Domino server. If you access them with the alternate Domino server, they require corrections to resolve inconsistencies on the production Domino server.
11. If the alternate partition is configured and you saved the `notes.ini` file in Step 3, copy that `notes.save` file back to`notes.ini` to be able to start the server.

Include and exclude processing

Information about file inclusions and file exclusion is presented.

Considerations

Data Protection for IBM Domino deals only with Domino databases and transaction log files if archival logging is in enabled on the Domino server. Other files that might exist on the server are not backed up by Data Protection for IBM Domino so they do not have to be excluded. If you want to limit the backups to a subset of the databases on your Domino server, the standard include and exclude syntax can be used.

Read the documentation about include and exclude processing for the base Tivoli Storage Manager backup-archive client as a thorough introduction to processing concepts. See the Examples in this section regarding Data Protection for IBM Domino.

Examples

Domino databases are stored by their relative names on the Tivoli Storage Manager server. As result, relative names must be used in include and exclude statements. The notes data directory should not be specified, and databases that are linked to the notes data directory by database or directory links must be referenced by the symbolic name. Do not use fully qualified physical file names.

A single database backup is stored as two objects on the Tivoli Storage Manager server. The objects that are created are the relative database name and the relative database name with a `.DATA` extension. For example, a backup of database `mail6\user1.nsf` would result in the following two objects:

1. The relative name of the database:

   ```
   mail6\user1.nsf
   ```

2. The relative name of the database with `.DATA`:

   ```
   mail6\user1.nsf.DATA
   ```

When you exclude a group of databases and then include a specific subset of that group, you must be sure to include both objects. For example, to exclude all databases in directory `mail6` except for database `user1.nsf`, use the following statements:

```
EXCLUDE mail6\*
INCLUDE mail6\user1.nsf
INCLUDE mail6\user1.nsf.DATA
```

Note: When you exclude a specific database, the `.DATA` object does not have to be excluded because the `.DATA` object is not created unless the database is included. When you assign a group of databases to a management class, you must assign both objects. For example, to assign all databases that match `*.nsf` in the `mail6` subdirectory to the `DOMINO` management class, code the following statement:

```
INCLUDE mail6\*.nsf* DOMINO
```

If archival logging is in effect on the domino server, you must be sure not to exclude the transaction log files from backup. The transaction logs have a base object name of `S######.TXN`, the "#" character represents a number. If you use a broad exclude statement, make sure to include the transaction log files by coding a statement as follows:

```
INCLUDE S*.TXN
```

Exclude databases that increase in size during compression `compressionyes` by using the client option, `exclude.compression`. You must specify the `.DATA` object to exclude a database from compression. For example, to exclude the database `mail6\user1.nsf` from compression, enter:

```
EXCLUDE.COMPRESSION mail6\user1.nsf.*
```

See *IBM Tivoli Storage Manager for Windows Backup-Archive Client Installation and User's Guide* for more information about the `exclude.compression` option.

You can encrypt Domino databases during backup and restore processing by specifying `enableclientencryptkey=yes` in the Data Protection for IBM Domino options file `dsm.opt`, located by default in the Data Protection for IBM Domino installation directory.. In the same file, specify the databases that you want to encrypt by adding an include statement with the `include.encrypt` option. For example, to encrypt all data, use:

```
include.encrypt *\...\*
```

To encrypt only the `Mydb.nsf` database in the default directory, use:

```
include.encrypt Mydb.nsf
include.encrypt Mydb.nsf.DATA
```

or

```
include.encrypt Mydb.nsf*
```

To encrypt all databases in the `mail65` directory, use:

```
include.encrypt mail65\...\*
```

Transparent encryption is available on Tivoli Storage Manager server Version 5.3 (or later).

You can choose to include backup or archive files for data deduplication. To refine the list of files to be included, the `include.dedup` option can be used in combination with the `exclude.dedup` option. By default, all eligible objects are included for data deduplication. The following examples show how to use the include and exclude options:

```
exclude.dedup E:\myfiles\...\*

Include.dedup E:\myfiles\archive\*
```

Exclude all databases named `db1.nsf` regardless of where they display:

```
EXCLUDE db1.nsf
```

Exclude all databases that match `help5_*` in the `help` subdirectory:

```
EXCLUDE help\help5_*
```

Include all databases in the `mail6` directory:

```
INCLUDE mail6\...\*
```

Assign all databases that match `*.nsf` in the `mail` subdirectory to the `MAILDB` management class:

```
INCLUDE mail\*.nsf* MAILDB
```

Exclude all databases in the `mail6` subdirectory from compression:

```
EXCLUDE.COMPRESSION mail6\...\*
```

The default `INCLUDE/EXCLUDE` lists.

```
EXCLUDE mail.box
EXCLUDE log.nsf
```

Note: You can back up the `log.nsf` database, but you can only restore it to an alternate name.

Include all transaction logs:

```
INCLUDE S*.TXN
```

Domino DB2

Include and exclude statements can be specified for table space backups and for full DB2 database backups. The administrator can use `include/exclude` statements to manage the meta objects (created by Data Protection for IBM Domino) and the

data objects (created by DB2 API). The management class assigned to meta objects is forced on the data objects created by the DB2 API. The include/exclude statements specification for meta objects is based on the naming convention for the meta object group leaders. To assign management classes to DB2NSF databases, the user must use the Tivoli Storage Manager object name of the full DB2 group leader object, or the Tivoli Storage Manager object name of the table space group leader object.

Domino DB2 example (full DB2 backup)

The following statement includes all DB2 databases that are assigned to management class MGMTC1, on partition NODE000, on the Domino 7 Server during a full DB2 backup:

```
INCLUDE \domino7.DOMDBS\NODE0000\FULL\DOMINO  MGMTC1
```

The following statement includes all DB2 databases that are assigned to management class MGMTC1, on all partitions that reside on all available Domino servers during a full DB2 backup:

```
INCLUDE \...\FULL\*    MGMTC1
```

Domino DB2 example (table space backup)

This statement includes DB2 Group GRP1:

```
INCLUDE GRP1
```

This statement assigns DB2 Group GRP2 to management class DB2GROUPS:

```
INCLUDE GRP2 DB2GROUPS
```

This statement excludes all DB2 Groups in CLASS1:

```
EXCLUDE CLASS1\*
```

This statement includes all DB2 Groups in CLASS2:

```
INCLUDE CLASS2\*
```

This statement excludes DB2 Group GRP1 in CLASS3:

```
EXCLUDE CLASS3\GRP1
```

Using multiple Domino server partitions

How to set up Data Protection for IBM Domino in multiple Domino server partitions is provided.

About this task

To use Data Protection for IBM Domino with multiple Domino server partitions on a single system, you must specify which partition you want to work with by identifying the location of the `notes.ini` file for that partition. In addition, when you are working with multiple Domino partitions, you must have separate Data Protection for IBM Domino log files for each server instance. The log file to be used is also specified in the Data Protection for IBM Domino preferences by the `logfile` option. To support multiple Domino partitions, create multiple preference files as follows:

Procedure

1. Use the **set** command with the `configfile` option to define a preferences file for each Domino partition to be supported. Make sure to specify the full path to the preferences file and be sure to set the `logfile` value appropriately. For example:

   ```
   domdsmc set notesinipath=c:\notes\data1\  /configfile=domino1.cfg
   domdsmc set logfile=domdsm1.log  /configfile=domino1.cfg
   domdsmc set notesinipath=c:\notes\data2\  /configfile=domino2.cfg
   domdsmc set logfile=domdsm2.log  /configfile=domino2.cfg
   ```

   Other Data Protection for IBM Domino preferences can be set for each partition. However, to identify the Domino server, the **notesinipath** parameter must be specified for each partition.
2. Use the `configfile` option when you are running the Data Protection for IBM Domino executable files to identify which preferences file is used for the command execution and thus which Domino partition to access. For example:

   ```
   domdsmc selective *  /configfile=domino1.cfg
   domdsm /configfile=domino1.cfg
   domdsm  /configfile=domino2.cfg
   ```

What to do next

For use of the GUI, create separate shortcuts with the appropriate `configfile` value to make it easy to start the Data Protection for IBM Domino GUI from an icon or from the Windows start menu.

Multiple Tivoli Storage Manager servers

How to use Data Protection for IBM Domino with multiple Tivoli Storage Manager servers is presented.

To use Data Protection for IBM Domino with multiple Tivoli Storage Manager servers, create multiple Data Protection for IBM Domino option files (one for each Tivoli Storage Manager server) and then use the `/adsmoptfile` parameter with the Data Protection for IBM Domino executable files to identify the server.

For example, assuming that you created `dsmserv1.opt` and `dsmserv2.opt` to identify the address and communication parameters necessary to access two Tivoli Storage Manager servers. Access the two servers using:

```
domdsmc selective *  /adsmoptfile=dsmserv1.opt
```

```
domdsmc selective *  /adsmoptfile=dsmserv2.opt
domdsm  /adsmoptfile=dsmserv2.opt
domdsm /adsmoptfile=dsmserv1.opt
```

For use of the GUI, separate shortcuts must be created with the appropriate `/adsmoptfile` value to make it easy to start the Data Protection for IBM Domino GUI from an icon or the windows start menu.

Problem determination

When an error occurs during a Data Protection for IBM Domino event, there is information that you can view to determine what the problem might be. Information about how to locate information to resolve problems is provided.

Data Protection for IBM Domino logs information, by default, to the `domdsm.log` file in the directory where Data Protection for IBM Domino is installed. This file indicates the date and time of a backup, data backed up, and any error messages or completion codes. This file is important and should be monitored daily.

The Tivoli Storage Manager API logs API error information, by default, to the `dsierror.log` file in the directory where Data Protection for IBM Domino is installed. This file does not contain backup statistics.

The Domino server logs information to the Windows Domino Event Log. Domino server error information can be obtained by viewing the Windows Domino Event Log.

The Tivoli Storage Manager scheduler logs information to both the `dsmsched.log` and the `dsmerror.log` files. By default, these files are located in the directory where the Tivoli Storage Manager backup-archive client is installed.

Note: When a scheduled command is processed, the schedule log might contain the following entry: `Scheduled event eventname completed successfully`. This indicates that Tivoli Storage Manager successfully issued the scheduled command that is associated with the `eventname`. No attempt is made to determine the success or failure of the command. Assess the success or failure of the command by evaluating the return code from the scheduled command in the schedule log. The schedule log entry for the command return code is prefaced with the following text: `Finished command. Return code is:`

The `statistics` option provides performance information at the individual database backup or restore level. Statistics are logged to the Data Protection for IBM Domino log file `domdsm.log` by default. Make sure that this option is specified in the Data Protection for IBM Domino preferences file `domdsm.cfg` by default, during backup and restore processing.

If the sources of information that are listed, do not provide an answer to your problem, contact your IBM service representative. The IBM service representative can provide more ways to gather diagnostic information.

You might be asked to provide these files as part of the troubleshooting process:

- `dsm.opt`
- `dsmerror.log`
- `dsmsched.log`
- `domdsm.log`

- domdsm.cfg
- dsierror.log
- tdpdommustgather.out

You might be asked to run these commands:

```
echo "------domdsmc query adsm------" >> tdpdommustgather.out
domdsmc query adsm >> tdpdommustgather.out
echo "------domdsmc query domino------" >> tdpdommustgather.out
domdsmc query domino >> tdpdommustgather.out
echo "------domdsmc query preferences------" >> tdpdommustgather.out
domdsmc query preferences >> tdpdommustgather.out
echo "------set------" >> tdpdommustgather.out
set >> tdpdommustgather.out
reg query HKLM\software\ibm\adsm\currentversion /s >> tdpdommustgather.out
```

Information about how to configure Tivoli Storage Manager compression, encryption, and deduplication is available in the Using the Application Programming Interface guide in the Tivoli Storage Manager Information Center.

Migration

Different migration scenarios for Data Protection for IBM Domino are presented.

Backups of older IBM Domino Server levels can be restored with Data Protection for IBM Domino on a newer Domino Server. However, backups of newer IBM Domino Server can be restored only to an IBM Domino Server of the same level.

For example, backups of Domino 8.5 Servers can be restored to IBM Domino 9 Servers. Backups of Domino 9 Servers can be restored only to IBM Domino 9 Servers.

The following sections provide two possible migration scenarios. The scenario that you choose depends on whether your database environment is using replicated servers. Replicated servers must be used for migration. This environment allows for a smooth transition from a Domino 8 Server to Domino 9 Server. All existing backup data is kept and available until you determine that it is no longer required.

Migrating in a replicated server environment

The steps that are required to migrate in a replicated server environment are presented.

About this task

In a replicated server environment, there are two or more servers with replicated databases. Follow the steps to migrate in a replicated server environment:

Procedure

1. Install the Domino 9 Server on one of the replicated Notes servers.
2. On the same server, install Data Protection for IBM Domino 7.1.
3. Stop taking backups on older Domino servers and begin full backups with Data Protection for IBM Domino 7.1. Alternatively, continue backing up in parallel on both servers until the new production environment is in place. Because Domino supports replication between servers, restores of backup data that is taken with the current Notes Tivoli Storage Manager can be done using the current Notes Tivoli Storage Manager on the Domino server. Replication action

propagates the restore operation to the new Domino server. When a restore operation is run to the Domino server with Data Protection for IBM Domino 7.1, the restored database can be replicated to the Domino server.

4. When the Domino server and backup scheme are stable, the other replicated servers can be upgraded with the new Domino Server and Data Protection for IBM Domino 7.1.

Migrating in a non-replicated server environment

Tasks that are required to migrate in a non-replicated server environment are provided.

About this task

When a second server is not available for replication, follow these steps:

Procedure

1. Before you upgrade the server to Domino 8.x or 9 take a full offline backup of the databases with theTivoli Storage Manager client.
2. Install Domino 8.5.2 on the server.
3. Install Data Protection for IBM Domino 7.1.
4. Begin with full backups with Data Protection for IBM Domino.

What to do next

If necessary, the Domino 6.5.x or 7.x version of the databases can be recovered with the Tivoli Storage Manager Backup Archive client.

Backing up and restoring Domino databases with DAOS

Data Protection for Domino can back up and restore online and offline Domino 8.5 (or later) NSF and NTF databases and transaction logs. This document provides an overview of the Data Protection for IBM Domino product, and describes the additional steps necessary to back up and restore Domino 8.5 (or later) NLO files.

Domino DAOS information

Some detailed information about using Data Protection for IBM Domino and Domino DAOS is provided.

Domino 8.5 and Data Protection for IBM Domino

IBM Domino server version 8.5 (and later) employs the Domino attachment and object service to save significant space at the file level by sharing data that is identified as identical between databases on the same server in NLO files. In databases that use DAOS, the Domino server no longer saves a separate and complete copy of every document attachment. The server saves a reference to each attached file in an NLO file. It then refers to the same NLO file from multiple documents in one or more databases on the same server.

NLO files can be present on a Domino 8.5 server if DAOS is enabled on the server, and DAOS participation is chosen for some of the NSF databases. NLO files are not supported by Data Protection for IBM Domino. They must be backed up and restored with the Tivoli Storage Manager Backup Archive Client. If DAOS is not enabled for an individual NSF, or on the server, the current Data Protection for IBM Domino backup and restore procedures are followed.

Domino server backup and restore strategy

Choose different backup strategies that are suitable for your specific requirements for network traffic, backup window, and acceptable restore times. Your choice of strategy includes you selecting the type of backup commands to use, and the type of transaction logging to be done on the Domino server.

Backing up NSF and NTF databases with Data Protection for IBM Domino

Data Protection for IBM Domino provides backup and restore functions for the Domino NSF and NTF databases (including template files) and associated transaction logs. However, Data Protection for IBM Domino does not provide a complete disaster recovery solution for a Domino server by itself. You must use the Tivoli Storage Manager backup-archive client with Data Protection for IBM Domino for a complete disaster recovery solution.

Backing up NLO files with the Tivoli Storage Manager backup-archive client

- Data Protection for IBM Domino does not process NLO files. The Tivoli Storage Manager backup-archive client must be used to back up and restore NLO files.
- NLO files are not modified after they are created. Only new NLO files must be processed during an incremental backup. The backup of the NLO files is done after the backup of the NSF and NTF databases. Thus ensuring that the NLO files in the backup are a superset of what is referenced by the NSF databases.
- The DAOS deferred delete interval is set to be longer than the interval between backups. If the backups are done weekly, the shortest DAOS deferred delete interval is eight days. Setting the delete interval longer than the backup interval ensures that all NLO files are backed up.
- If there is a retention limit for backups, the DAOS delete interval must be set to longer than that retention limit. This measure ensures that all referenced NLO files exist if an NSF database is restored from the oldest backup.

Backing up execution and configuration files with the Tivoli Storage Manager backup-archive client

Some files that are part of the Domino server installation, such as execution and configuration files are not supported by Data Protection for IBM Domino. For example, the DAOS configuration file, `daos.cfg`, is not supported by Data Protection for IBM Domino, and can be part of your recovery strategy.

Another example is database link files. These files have an `.nsf` extension but are not considered databases, and are not backed up by Data Protection for IBM Domino. These files must be recovered in a disaster recovery situation. A comprehensive disaster recovery plan can be achieved with the Tivoli Storage Manager backup-archive client with Data Protection for IBM Domino.

Personal copies (replicas) of Domino databases that are stored on Notes clients (not on the Domino server) are not protected by Data Protection for IBM Domino. You can use the Tivoli Storage Manager backup-archive client on the Notes client system to back up and restore these files. Alternatively, the Domino server replication recovers them.

Restoring Domino documents

To restore an individual Notes document, you must restore the entire database with another name. Choose a time when the document existed for both the restore **/pit** and activate **/applylogs** commands, but before the document was deleted. Then, copy the document from the Notes client.

Data Protection for IBM Domino backs up transaction logs from a Domino server that has archival logging in effect only. Transaction logs cannot be backed up from a Domino server that has circular or linear logging in effect.

When you are using archival transaction logging, the frequency of the **archivelog** command use depends on the size of your log and the rate of change for logged databases. Run archival transaction logging several times per day if you generate a large volume of changes at a rapid rate.

When you are restoring a group of logged databases for which transactions must be applied, activate them together when possible. This action avoids restoring the same transaction log files multiple times. Restored transaction log files are deleted during a database recovery by the Domino server. Activating and applying logs to the database separately requires retransmitting log files for each database.

Backing up a Lotus Domino database with DAOS

Data Protection for IBM Domino provides two types of database backups, incremental and selective. Descriptions of how to run backups are presented.

Data Protection for IBM Domino can run full and incremental online backups of individual NSF and NTF databases when archival logging is in effect. If archival logging is not in effect, only full offline backups of NSF databases can be run. NSF and NTF databases must be backed up using Data Protection for IBM Domino to ensure that an internally consistent image of the NSF is saved, and that the Domino transaction logs are archived as part of the backup process. Backing up an NSF or NTF database without using Data Protection for IBM Domino while the Domino server is running, and while modifying the file, might result in an unusable image being saved.

Running an incremental backup

An incremental backup runs a full online backup of Domino databases under the following conditions:

1. The database is within the Domino data path or is symbolically linked to the Domino data path by directory or database links.
2. The database is not excluded from backup by exclude statements within the Tivoli Storage Manager include-exclude options file.
3. If the database is logged, the DBIID changes.
4. If the database is not logged, it is modified since the last backup occurred. Data and non-data modification dates are checked.
5. The database is new or newly included in the backup.

The **incremental** command includes a function that determines when active backup database copies exist on the Tivoli Storage Manager server that were deleted from the Domino server or excluded from backup. If so, they are marked inactive so automatic expiration of these backup copies can occur according to defined management class parameters for backup files.

The **incremental** command normally specifies a wildcard qualified name. The databases that match the wildcard qualification and meet the selection criteria for an incremental backup are backed up. Use the Tivoli Storage Manager backup-archive client **incremental** command to back up NLO files because when they are written they are never changed.

If selective backups of the NLO files are made, each additional backup of the same file results in the saving of an identical backup. If incremental backups of the NLO files are made, only a single backup file is created. When you are backing up NSF and NLO files, first backup the NSF databases and then backup the NLO files.

The Data Protection for IBM Domino command to incrementally back up NSF databases is:

```
domdsmc incr database_selection_criteria
```

The Tivoli Storage Manager backup-archive client command to incrementally back up all NLO files is:

```
dsmc incr /local.notesdata/daos/* -su=yes
```

The Tivoli Storage Manager backup-archive client incremental backup command of the NLO files must specify the fully qualified path name of the NLO files to be backed up.

Running a selective backup

A selective backup unconditionally backs up the specified databases, unless they are excluded from backup through exclude statements within the Tivoli Storage Manager `include /exclude` options file.

Do not use the Tivoli Storage Manager backup-archive client selective command to back up NLO files because when they are written they are never changed. If selective backups of the NLO files are made, each additional backup of the same file results in an additional identical backup. When you back up NSF, NTF, and NLO files, first backup the NSF and NTF databases, and then incrementally backup the NLO files.

The Data Protection for IBM Domino command to selectively back up NSF databases is:

```
domdsmc sel database_selection_criteria
```

The Tivoli Storage Manager backup-archive client command to incrementally back up all NLO files is:

```
dsmc incr /local.notesdata/daos/* -su=yes
```

Domino Transaction Log Archive

Data Protection for IBM Domino provides the capability to create archives of transaction logs when archival logging is in effect. There are no changes required in archiving Domino transaction logs when DAOS is enabled. A transaction log captures database changes so full database backups are not required as frequently. Updates to a logged database are recorded in the Domino server transaction log. Changes to a database since the last full backup can be applied from the transaction log after the backup is restored from the last full backup. Enabling transaction logging for all databases on a Domino server is not required, so the backup process must handle databases that are logged and not logged.

Domino allows the active transaction log to be backed up as well. The Data Protection for IBM Domino archive log capability stores filled transaction log files on the Tivoli Storage Manager server so that space allocated for these files can be reused by the Domino logger.

The **archivelog** command is available when transaction logging on the Domino server is enabled in archival mode. Filled transaction log files must be archived frequently enough to ensure the transaction log never fills completely and stops the Domino server. Transaction log files that are stored on the Tivoli Storage Manager server are automatically restored as needed for a database recovery. Archived transaction log files are retained on the Tivoli Storage Manager server as long as a database backup exists that needs these log files for a complete recovery.

The Data Protection for IBM Domino command to archive the Domino server transaction log is:

```
domdsmc archivelog
```

When circular or linear loop logging is used on the Domino server, or when logging is disabled on the Domino Server, transaction log files are not archived.

Restoring an IBM Domino database with DAOS

How to use Data Protection for IBM Domino to restore an IBM Domino database with DAOS is outlined.

The restoration of a Domino NSF or NTF database is a two-step recovery process.

1. Use the Data Protection for IBM Domino **domdsmc restore** command to restore one or more databases from the Tivoli Storage Manager server backup storage to the Domino server.
2. Use the Data Protection for IBM Domino **domdsmc activate** command to bring the restored databases online for use by the Domino server and optionally apply transactions from the transaction log to update the database to the latest level.

Note: These two steps can be combined into one step by specifying the `/activate=yes` option on the restore command.

When the restored NSF database is enabled for DAOS, one or more NLO files might be missing and must be recovered. To recover the NLO files, follow these steps:

1. Determine which NLO files that are referenced by the restored NSF databases are missing.
2. If there are missing NLO files, you must use the Tivoli Storage Manager backup-archive client to restore the required NLO files.

Restore process

The restore process retrieves previously backed up copies of the databases to be restored from the Tivoli Storage Manager server and restores them to the Domino server storage. You can restore the database with the original name (replace) or with a different database file name. The database can be restored to the same name in a different Domino server directory, or to a different Domino server.

The Data Protection for IBM Domino command to restore an NSF database is:

```
domdsmc restore database_name –into=restored_database_name
```

Activation

After the NSF and NTF databases are restored, the Data Protection for IBM Domino **domdsmc activate** command is used to apply any changes to the restored databases from the recovery logs and to activate the NSF and NTF databases that are being restored. This activation step brings restored databases online for use by the Domino server.

You can optionally apply transactions from the transaction log to update the database. Transactions can be applied up to a specific point in time or up to the most recent changes that are recorded in the transaction log. If archival logging is in effect, Data Protection for IBM Domino automatically restores archived transaction log files as needed.

The Data Protection for Domino command to activate and apply logs to one or more databases is:

```
domdsmc activate –applylogs
```

Determine when there are missing NLO files

Issue the DAOS manager **tell daosmgr listnlo** command from the Domino server console to discover the names of any missing NLO files that are referenced by the restored NSF databases.

Note: If the DAOS deferred delete interval is longer than the age of the restored backup, there will be no missing NLO files. However, if the age of the backup is greater than the DAOS deferred delete interval, you might have to restore the missing NLO files.

The server command to determine the NLO files that must also be restored is:

```
Tell daosmgr listnlo –o missingnlo.txt missing nsf_database_name
```

The file `missingnlo.txt` contains a list of NLO files that are referenced by the restored NSF database and that are not found on the Domino server.

Restore the missing NLO files

Use the Tivoli Storage Manager backup-archive client **dsmc restore** command to restore the missing NLO files. The option `–latest` on the restore command specifies that the latest copy, whether active or inactive, is used. If this option is not specified and the NLO file is expired by an incremental backup, the restore operation fails and you get the message:

```
"ANS1302E No object on server match query"
```

The server command to restore the missing NLO files is:

```
dsmc restore –filelist missingnlo.txt –latest
```

Resynchronize the DAOS catalog

To ensure that DAOS has the correct reference counts after the missing NLO files are restored, you must run the **tell daosmgr resync** command from the Domino server console. If the catalog is still synchronized after the restoration, this command will close.

Restore at document level

Data Protection for IBM Domino restores Domino databases at the database level. To restore a document in a database, the entire database and the necessary NLO files must first be restored. Then, they must be copied to the "live" NSF database. A database can be restored to the production server under a temporary name, and the document you want can be copied to the appropriate database. If, for performance reasons, the production server cannot be used in the restore process, the database can be restored to an alternative server and copied to the production server.

You must run alternate server restores when possible to reduce demands on the Domino production server. Alternate server restores can be directed to an alternate partition or to a separate Domino server by using these steps:

1. Inform the user of the location of the restored NSF database.
2. The user can copy the necessary documents from the restored NSF to the live NSF.
3. When the user finishes copying the data from the restored NSF to the live NSF, delete the restored NSF.
4. If the document includes a DAOS attachment, restore any missing NLO files and resynchronize the DAOS catalog with the **tell daosmgr resync** command to repair the reference counts.

Disaster recovery for an IBM Domino database with DAOS

The disaster recovery of a Domino server with the Tivoli Storage Manager backup-archive client with Data Protection for IBM Domino is presented.

You must use the Tivoli Storage Manager backup-archive client with Data Protection for IBM Domino for a complete disaster recovery solution.

Disaster recovery of a Domino server requires:

1. Rebuilding the directory structure.
2. Restoring the non-database files with the Tivoli Storage Manager backup-archive client.
3. Recovering the database files to the latest level with Data Protection for IBM Domino.

To recover a Domino server:

1. Make sure that the new installation is configured in the same manner as the damaged installation. For example, both installations must have the same directory structure and location, and `logdir` path.

   Note: The existence of NLO files changes this step only when NLO files are recovered along with other non-database files, such as `notes.ini`.
2. Use the Tivoli Storage Manager backup-archive client to recover the non-database Domino server files, such as `notes.ini`, `cert.id` and `server.id`. If a Domino attachment and object service is present, the NLO files must also be recovered now.
3. Use a text editor to modify the `notes.ini` file for the Domino server with this setting: `TRANSLOG_Status=0`.
4. Use Data Protection for IBM Domino to restore the transaction log file to be used in the log recovery procedure. This log file is the last transaction log file to be archived before the loss of the active log.

5. Delete the contents of the Domino transaction log directory, except for the log file that is restored in Step 4.
6. Use a text editor to modify the `notes.ini` file for the Domino server with these settings: `TRANSLOG_Recreate_Logctrl=1TRANSLOG_Status=1`
7. Use Data Protection for IBM Domino to restore (but not to activate) the databases you want to recover to the latest level within the archived log extents.
8. Use the Tivoli Storage Manager backup-archive client to recover the latest NLO files.
9. Use Data Protection for IBM Domino to activate the databases you are recovering and apply transaction logs. The **`TRANSLOG_Recreate_Logctrl`** parameter in the `notes.ini` file is automatically reset to 0.
10. Start the Domino server. With the disaster recovery complete, it is now safe to start the Domino server and run the server tasks and functions.
11. Use the selective backup function in Data Protection for IBM Domino to run full backups of all databases. This ensures correct recoverability with subsequent transaction log files.
12. Use Data Protection for IBM Domino to archive the transaction log. The transaction log file that is used in the recovery procedure is modified and available for archiving. This transaction log has the ID of the current logger.

Chapter 5. Reference information

Frequently asked questions and best practices that contain troubleshooting information for Data Protection for IBM Domino are presented.

Frequently asked questions

Here are some frequently asked questions for Data Protection for IBM Domino.

Why do I receive a "DB2 CONNECTION ERROR" message when I am running an existing restore for a DB2 database that is enabled for `rollforward` recovery?

When a DB2 database is enabled for `rollforward` recovery, and the database is used for an existing restore, the Domino server cannot connect to the DB2 database until after the `rollforward` operation completes. As a result, the command output displays this message text:

```
Starting Domino DB2 database rollforward...
Initializing Domino connection...
Restart Analysis (0 MB): 100%
04/21/2007 12:02:57 AM  A RM error occurred.: An error occurred accessing the db
2 datasource.

DB2 CONNECTION ERROR:  Domino unable to connect to DB2 database 'DOMDB2' as user
 'db2admin'...
[IBM][CLI Driver] SQL1117N  A connection to or activation of database "DOMDB2" c
annot be made because of ROLL-FORWARD PENDING.  SQLSTATE=57019

DB2 CONNECTION ERROR:  set DEBUG_DB2CONNECT=0 to suppress this message.
04/21/2007 12:02:57 AM  Unable to initialize DB2 services.  DB2-based nsfs will
be unusable.: An error occurred accessing the db2 datasource.
```

There is no DB2 connection error and therefore, this message text can be ignored.

Why does my backup session timeout even though I used the `sessions` option?

This situation can occur when the number of specified backup sessions exceeds the number of available mount points. Each session requests a mount point from the Tivoli Storage Manager server when backup processing begins. If a mount point is in use, then the mount point is not released for use by a new session until the backup is complete. Because of this behavior, it is possible that a session that is waiting for an available mount point might timeout, causing the backup attempt to fail. To avoid this situation, make sure that the number of available mount points from the Tivoli Storage Manager server is equal to the number of sessions that are specified with the `sessions` option. It is the responsibility of the user to determine the number of available mount points as Data Protection for IBM Domino does not determine this information.

How can I avoid being prompted for a Domino server password when I am backing up encrypted databases?

Use the Domino server Administrator to select the `Don't prompt for a password from other Notes-based programs` option in the Domino Server ID file.

Can I run multiple `domdsmc` instances?

You can run multiple instances of `domdsmc` for backup processing. This option can improve performance when you are backing up many

databases. It can also improve resource utilization if your data is in sequential access storage pools on multiple drives. Running multiple instances of `domdsmc` is controlled with the `SESSION` option. This option is described in the `DOMDSMC SET SESSION` section.

The `SESSION` option specifies the number of sessions to open to the Tivoli Storage Manager server. This option applies to NSF database backups only.

Can I restore an individual document?

To restore an individual document, the entire database must first be restored and then the document copied. A database can be restored to the production server under a temporary name and the document can be copied to the appropriate database. If for performance reasons, the production server cannot be used in the restore process, the database can be restored to a different server and copied to the production server. Run alternate server restores when possible to reduce the demands on the Domino Production server. Alternate server restores can be saved to an alternate partition or to a separate Domino server. For more information about restoring documents, see "NSF databases restore to alternate server and alternate partition" on page 167.

Can I back up and restore private folders with Data Protection for IBM Domino?

Data Protection for IBM Domino runs backup and restore processing at the database level. The contents of the entire database is processed. As a result, a private folder is processed if it is stored in the database. Data Protection for IBM Domino does not back up or restore private folders in a desktop file.

What are the .nsf and .nsf.DATA Tivoli Storage Manager server objects?

The `nsf.DATA` object contains the actual file data from the Domino server database. The `.nsf` object contains information about the `.nsf` file but no actual file data. Both files are created on the Tivoli Storage Manager server during Data Protection for IBM Domino backup processing.

When you are issuing an include statement, make sure to include both files. For example:

```
include dbname.nsf
include dbname.nsf.data
```

or

```
include dbname.nsf*
```

What is a `.pdb` file and where is it located?

The `.pdb` file tracks the Domino databases that are in a state of pending activation. It is used during `query pendingdbs` processing. Do not attempt to edit this file.

Can I manually edit the Data Protection for IBM Domino preferences file `domdsm.cfg`?

You must use the **`set`** command (or use the Preferences Editor in the GUI) to edit the preferences file. If edited manually, hidden characters can be introduced that negatively affect parsing.

How do I encrypt my backups?

You can encrypt your Domino databases during Data Protection for IBM Domino backup and restore processing by specifying *enableclientencryptkey=yes* in the Data Protection for IBM Domino options file `dsm.opt` and adding an include statement with the `include.encrypt` option.

- See "More configuration options" on page 33 for more information about the `enableclientencryptkey` option.
- See "Include and exclude processing" on page 171 for examples of `include.encrypt` statements.

Can I restore a database to an operating system that is different from the one it was backed up from?

Data Protection for IBM Domino does not support restore processing across operating systems. For example, you cannot restore a backup that runs on an AIX system to a Windows system. You must restore Data Protection for IBM Domino backups to the same platform from which it was backed up.

Can I run Data Protection for IBM Domino through Windows terminal service?

Windows Terminal Services cannot remotely administer Data Protection for IBM Domino. A new session opens for every connection. As a result, you are not connecting to the current server session. Use a third-party solution that allows direct connection to an active session.

Can I run multiple scheduled backups simultaneously?

The Tivoli Storage Manager client scheduler allows only one scheduler process at a time. Other schedules must wait for the first scheduled backup to complete before they can begin. This wait can be an issue when a scheduled backup is processing at the same time an `archivelog` backup is scheduled to begin. See "The Tivoli Storage Manager client scheduler and associated log files" on page 162 for an example of how to maintain a scheduled database backup and scheduled `archivelog` backup.

How do I automate (schedule) a backup?

See "Automating backups" on page 159.

How do I run a silent installation?

See "Silent installation of Data Protection for IBM Domino" on page 23.

How do I recover from the loss of a Domino transaction log?

See "Recovering from loss of Domino transaction logs for NSF databases" on page 166

How do I run an alternate Domino server restore?

See "Restoring NSF databases to an alternate server" on page 167.

How do I run an alternate Domino partition restore?

See "Restoring NSF databases to an alternate partition" on page 169.

Where do I find error information?

See "Problem determination" on page 176.

How do I include and exclude files?

See "Include and exclude processing" on page 171.

How do I migrate backups from earlier versions?

See "Migration" on page 177.

How do I use Data Protection for IBM Domino with multiple Domino server partitions?

See "Using multiple Domino server partitions" on page 175.

How do I use Data Protection for IBM Domino with multiple Tivoli Storage Manager servers?

See "Multiple Tivoli Storage Manager servers" on page 175.

How do I compress backups?

Add the line `COMPRESS YES` to the `dsm.opt` file.

How do I enable data deduplication for Data Protection for Domino?

The Tivoli Storage Manager server must be enabled for data deduplication and the node must be enabled. In the `dsm.opt` file, add the line `DEDUPLICATION YES`

For more information about data deduplication, see Data deduplication in Tivoli Storage Manager V6.2 and V6.1

Where can I find out more about data deduplication with Tivoli Storage Manager?

More information is available in these documents:

- Data deduplication in Tivoli Storage Manager V6.2 and V6.1
- Data deduplication best practices for Tivoli Storage Manager V6.2

What happens when the connection is broken between the Data Protection for IBM Domino client and the Tivoli Storage Manager server?

When the connection is broken between the Data Protection for IBM Domino client and the Tivoli Storage Manager server, the client attempts to reconnect to the server. The Domino client attempts to reconnect during the **COMRESTARTDURATION** timeframe.

Note: In certain circumstances, some connection errors might not be recoverable. For example, when the network adapter is disabled and re-enabled on an AIX server the session might not reconnect successfully.

How do I cancel a session between a Data Protection for IBM Domino client and a Tivoli Storage Manager server?

This task must be carried out by a Tivoli Storage Manager administrator with access to the Tivoli Storage Manager administrator password.

1. Log on to the server console from the Tivoli Storage Manager administrator command-line interface.
2. Enter the **query session** command to list all the active sessions.
3. Locate the Domino client session that you want to cancel by looking for the node name under the client name column. Take note of the session number.

   Note: There might be more than one session that is listed for the Domino client. In TCPIP communication, canceling any one of these sessions cancels all of the sessions. In LAN-free communication, you must cancel the session that has the session stage `RecW`.
4. Cancel the session by entering the **cancel session <session number>** command.

Best practices for optimizing Data Protection for IBM Domino performance

The best practices that are listed, help you to achieve the best use of Data Protection for IBM Domino.

For more information about Data Protection for IBM Domino performance, see "Performance" on page 13.

Restore data to an alternative Domino Server

When you restore databases to a different Domino server, you can reduce the demands that are placed on the Domino production server. When you use a different Domino server, you can run restore processing and apply

transaction logs on that server. The use of the alternate server reduces the processing demands on the production server.

Specify an alternate restore path with the TRANSLOG_RECOVER_PATH variable in the NOTES.INI file to restore transaction logs to a different path. For instructions about restoring NSF databases to an alternate server, see "Restoring NSF databases to an alternate server" on page 167. For instructions about restoring NSF databases to an alternate partition, see "Restoring NSF databases to an alternate partition" on page 169.

Run parallel sessions to improve performance
Use the SESSIONS option to control the number of parallel sessions during backup operations. Running parallel sessions can improve performance. For more information about performance, see "Performance" on page 13.

Set up multiple schedules to increase performance
Setting up multiple schedules can help Data Protection for IBM Domino performance. For more information about running multiple scheduled backups, see the *Parallel Session Scheduled Backups of Domino on Windows* technote http://www.ibm.com/support/docview.wss?uid=swg21194688.

Reduce query processing time
When you are searching for Tivoli Storage Manager server database names, find the name with letters and a wildcard character (*) in the **By Database Name** field. Using this search technique can reduce query processing time. For example, when you type in a* all the databases in the selected folder that begin with the letter a are shown. Make sure to click **Update** after you enter the database query.

Maintain a scheduled database backup and a scheduled archivelog backup
With the Tivoli Storage Manager client scheduler, only one scheduler process runs at a time. Other schedules must wait for the first scheduled backup to complete before they can begin. This wait can be an issue when a scheduled backup is processing at the same time as an archivelog backup is scheduled to begin. For instructions about scheduling backups, see "Defining a Tivoli Storage Manager server schedule" on page 160.

Check log files for the completion of processes
Regularly check the domdsm.log to ensure that scheduled processes are successfully completed. For more information, see "Domdsmc query preferences" on page 95.

Avoid backing up identical data in a clustered environment
Use the exclude option to exclude databases to prevent them being backed up twice in a clustered environment. In a Domino clustered environment, Data Protection for IBM Domino is installed on multiple nodes. It is not necessary to back up identical replicated databases from multiple nodes. For more information about using the exclude statement, see "Include and exclude processing" on page 171.

Naming conventions
Use different names for your Data Protection for IBM Domino nodes, management classes, and policy domains than the names that are used for your client nodes, management classes, policy domains, and Domino partitions. Using different names greatly reduces the possibility of confusion and error throughout your Data Protection for IBM Domino environment. For information about required options for Data Protection for IBM Domino, see "Required options" on page 32. For information about registering a node name, see "Registering with the Tivoli Storage Manager server" on page 29.

Appendix A. Tivoli support information

You can find support information for Tivoli and other IBM products from various sources.

From the IBM Support Portal at http://www.ibm.com/support/entry/portal/, you can select the products that you are interested in and search for a wide variety of relevant information.

Communities and other learning resources

In addition to product documentation, many forms of assistance are available to help you get started as you deploy and use the Tivoli Storage Manager family of products. These resources can also help you to solve problems that you might have.

You can use forums, wikis, and other social media tools to ask questions, talk to experts, and learn from others.

User groups

Tivoli Global Storage Virtual User Group

Access this user group at http://www.tivoli-ug.org/storage.

This group makes it possible for individuals from many different industries and types of organizations to share information and work directly with the IBM product experts. Local chapters also exist where members meet in person to share experiences and hear from guest speakers.

ADSM.ORG

Access this mailing list at http://adsm.org.

This independently managed Storage Management discussion forum started when Tivoli Storage Manager was known as ADSTAR Distributed Storage Manager (ADSM). The members of this forum have many years of experience with Tivoli Storage Manager in almost every type of IT environment.

To subscribe to the forum, send an email to listserv@vm.marist.edu. The body of the message must contain the following text: `SUBSCRIBE ADSM-L` *your_first_name your_family_name*.

Tivoli Storage Manager community on Service Management Connect

Access Service Management Connect at http://www.ibm.com/developerworks/servicemanagement. In the Storage Management community of Service Management Connect, you can connect with IBM in the following ways:

- Become involved with transparent development, an ongoing, open engagement between users and IBM developers of Tivoli products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and the Tivoli Storage Manager community.
- Read blogs to benefit from the expertise and experience of others.

- Use wikis and forums to collaborate with the broader user community.

Tivoli Storage Manager wiki on developerWorks®

Access this wiki at https://www.ibm.com/developerworks/servicemanagement/sm/index.html.

Find the latest best practices, white papers, and links to videos and other resources. When you log on, you can comment on content, or contribute your own content.

Tivoli Support Technical Exchange

Find information about upcoming Tivoli Support Technical Exchange webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html. Replays of previous webcasts are also available.

Learn from technical experts who share their knowledge and then answer your questions. The sessions are designed to address specific technical issues and provide in-depth but narrowly focused training.

Other social media sites

LinkedIn

You can join groups on LinkedIn, a social media site for professionals. For example:

- **Tivoli Storage Manager Professionals**: http://www.linkedin.com/groups/Tivoli-Storage-Manager-Professionals-54572
- **TSM**: http://www.linkedin.com/groups?gid=64540

Twitter

Follow `@IBMStorage` on Twitter to see the latest news about storage and storage software from IBM.

Tivoli education resources

Use these education resources to help you increase your Tivoli Storage Manager skills:

Tivoli Education and Certification website

View available education at http://www.ibm.com/software/tivoli/education.

Use the Search for Training link to find local and online offerings of instructor-led courses for Tivoli Storage Manager.

Education Assistant

Access resources at http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp.

Scroll to view the list of available training videos. Recorded product demonstrations are also available on a YouTube channel.

Searching knowledge bases

If a problem occurs while you are using one of the Tivoli Storage Manager family of products, you can search several knowledge bases.

Begin by searching the Tivoli Storage Manager Information Center at http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1. Within the information center, you can enter words, phrases, or message numbers in the **Search** field to find relevant topics.

Searching the Internet

If you cannot find an answer to your question in the Tivoli Storage Manager information center, search the Internet for the information that might help you resolve the problem.

To search multiple Internet resources, go to the IBM support website at http://www.ibm.com/support/entry/portal/. You can search for information without signing in.

Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources:

- IBM technotes.
- IBM downloads.
- IBM Redbooks® publications.
- IBM Authorized Program Analysis Reports (APARs). Select the product and click **Downloads** to search the APAR list.

Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that can help you with problem determination. It is available for some Tivoli Storage Manager and Tivoli Storage FlashCopy Manager products.

IBM Support Assistant helps you gather support information when you must open a problem management record (PMR), which you can then use to track the problem. The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

You can find more information and download the IBM Support Assistant web page at http://www.ibm.com/software/support/isa.

You can also install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use. Find add-ons for specific products at http://www.ibm.com/support/docview.wss?uid=swg27012689.

Finding product fixes

A product fix to resolve a software problem might be available from the IBM software support website.

Procedure

Determine what fixes are available by checking the IBM software support website at http://www.ibm.com/support/entry/portal/.

If you previously customized the site based on your product usage:

1. Click the link for the product, or a component for which you want to find a fix.
2. Click **Downloads**, and then click **Search for recommended fixes**.

If you have not previously customized the site:

Click **Downloads** and search for the product.

Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

Procedure

1. From the support page at http://www.ibm.com/support/entry/portal/, click **Sign in** and sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
2. Click **Manage all my subscriptions** in the Notifications pane.
3. Click the **Subscribe** tab, and then click **Tivoli**.
4. Select the products for which you want to receive notifications and click **Continue**.
5. Specify your notification preferences and click **Submit**.

Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM subscription and support contract, and if you are authorized to submit problems to IBM.

Procedure

1. Ensure that you have completed the following prerequisites:
   a. Set up a subscription and support contract.
   b. Determine the business impact of the problem.
   c. Describe the problem and gather background information.
2. Follow the instructions in "Submitting the problem to IBM Software Support" on page 198.

Setting up and managing support contracts

You can set up and manage your Tivoli support contracts by enrolling in IBM Passport Advantage®. The type of support contract that you need depends on the type of product you have.

Procedure

Enroll in IBM Passport Advantage in one of the following ways:

- **Online:** Go to the Passport Advantage website at http://www.ibm.com/software/lotus/passportadvantage/, click **How to enroll**, and follow the instructions.
- **By telephone:** For critical, system-down, or high-severity issues, you can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html and click **Contacts**.

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must understand and assess the business impact of the problem you are reporting.

Severity level	Description
Severity 1	**Critical** business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	**Significant** business impact: The program is usable but is severely limited.
Severity 3	**Some** business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	**Minimal** business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

Describing the problem and gathering background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.
- Are you using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submitting the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by telephone.

Online

Go to the IBM Software Support website at http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone

For critical, system-down, or severity 1 issues, you can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html and click **Contacts**.

Appendix B. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision to use information technology products successfully.

Accessibility features

The IBM Tivoli Storage Manager family of products includes the following accessibility features:

- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:

- Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager Server can be installed in console mode, which is accessible.

The Tivoli Storage Manager Information Center is enabled for accessibility. For information center accessibility information, see "Accessibility features in the information center" (http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/topic/com.ibm.help.ic.doc/iehs36_accessibility.html).

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (http://www.ibm.com/able) for information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

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

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

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

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

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

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

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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 want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.

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

The licensed program described in this information 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.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

COPYRIGHT LICENSE:

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

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample 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_.

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

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml.

Intel, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other product and service names might be trademarks of IBM or other companies.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Glossary

This glossary provides terms and definitions for Tivoli Storage Manager, Tivoli Storage FlashCopy Manager, and associated products.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM Terminology website at www.ibm.com/software/globalization/terminology.

A

absolute mode
In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also mode, modified mode.

access control list (ACL)
In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights.

access mode
An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume.

ACK See acknowledgment.

acknowledgment (ACK)
The transmission of acknowledgment characters as a positive response to a data transmission.

ACL See access control list.

activate
To validate the contents of a policy set and then make it the active policy set.

active-data pool
A named set of storage pool volumes that contain only active versions of client backup data. See also server storage, storage pool, storage pool volume.

active file system
A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. See also inactive file system.

active policy set
The activated policy set that contains the policy rules currently in use by all client nodes assigned to the policy domain. See also policy domain, policy set.

active version
The most recent backup copy of a file stored. The active version of a file

cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. See also backup version, inactive version.

activity log
A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup
A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client
A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the server. See also backup-archive client.

administrative command schedule
A database record that describes the planned processing of an administrative command during a specific time period. See also central scheduler, client schedule, schedule.

administrative privilege class
See privilege class.

administrative session
A period of time during which an administrator user ID communicates with a server to perform administrative tasks. See also client node session, session.

administrator
A person responsible for administrative tasks such as access authorization and content management. Administrators can also grant levels of authority to users.

agent node
A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

aggregate
An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also logical file, physical file.

aggregate data transfer rate
A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

application client
A program that is installed on a system to protect an application. The server provides backup services to an application client.

archive
To copy programs, data, or files to another storage media, usually for long-term storage or security. See also retrieve.

archive copy
A file or group of files that was archived to server storage

archive copy group
A policy object containing attributes that control the generation, destination, and expiration of archived files. See also copy group.

archive-retention grace period
The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also bind.

association
The defined relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

audit To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

authentication rule
A specification that another user can use to either restore or retrieve files from storage.

authority
The right to access objects, resources, or functions. See also privilege class.

authorization rule
A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user
A user who has administrative authority for the client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS
See automounted file system.

automatic detection
A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration
The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also demand migration, threshold migration.

automounted file system (AutoFS)
A file system that is managed by an automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B

backup-archive client
A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. See also administrative client.

backup copy group
A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class. See also copy group.

backup retention grace period
The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set
A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection
A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version
A file or directory that a client node backed up to storage. More than one backup version can exist in storage, but only one backup version is the active version. See also active version, copy group, inactive version.

bind To associate a file with a management class name. See also archive-retention grace period, management class, rebind.

C

cache To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file
A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD See client acceptor daemon.

central scheduler
A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See also administrative command schedule, client schedule.

client A software program or computer that requests services from a server. See also server.

client acceptor
A service that serves the Java applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX, UNIX, and Linux systems, the client acceptor is run as a daemon.

client acceptor daemon (CAD)
See client acceptor.

client domain
The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node
A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session
A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. See also administrative session.

client option set
A group of options that are defined on the server and used on client nodes in conjunction with client options files.

client options file
An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client-polling scheduling mode
A method of operation in which the client queries the server for work. See also server-prompted scheduling mode.

client schedule
A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also administrative command schedule, central scheduler, schedule.

client/server
Pertaining to the model of interaction in distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file
A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. See also client user-options file, options file.

client user-options file
A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the dsm.opt file. For AIX, UNIX, or Linux systems, see also client system-options file. See also client system-options file, options file.

closed registration
A registration process in which only an administrator can register workstations as client nodes with the server. See also open registration.

collocation
The process of keeping all data belonging to a single-client file space, a

single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group
A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point
A point in time when data is considered to be consistent.

communication method
The method by which a client and server exchange information. See also Transmission Control Protocol/Internet Protocol.

communication protocol
A set of defined interfaces that permit computers to communicate with each other.

compression
A function that removes repetitive characters, spaces, strings of characters, or binary data from the data being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for data.

configuration manager
A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also enterprise configuration, managed server, profile.

conversation
A connection between two programs over a session that allows them to communicate with each other while processing a transaction. See also session.

copy backup
A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted.

copy group
A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also archive copy group, backup copy group, backup version,management class.

copy storage pool
A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also destination, primary storage pool, server storage, storage pool, storage pool volume.

D

daemon
A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file
A physical file in which read errors have been detected.

database backup series
One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series. See also database snapshot, full backup.

database snapshot
A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also database backup series, full backup.

data center
In a virtualized environment, a container that holds hosts, clusters, networks, and data stores.

data deduplication
A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage media. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server
A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover
A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-programming interface (DSMAPI)
A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

data store
In a virtualized environment, the location where virtual machine data is stored.

deduplication
The process of creating representative records from a set of records that have been identified as representing the same entities.

default management class
A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

demand migration
The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated. See also automatic migration, selective migration, threshold migration.

desktop client
The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination
A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated. See also copy storage pool.

device class
A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file
1. For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.
2. For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

disaster recovery manager (DRM)
A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan
A file that is created by the disaster recover manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and hardware that is used by the server, and the location of recovery media.

domain
A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See also policy domain.

DRM See disaster recovery manager.

DSMAPI
See data storage-management application-programming interface.

dynamic serialization
Copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive. See also shared dynamic serialization, shared static serialization, static serialization.

E

EA See extended attribute.

EB See exabyte.

EFS See Encrypted File System.

Encrypted File System (EFS)
A file system that uses file system-level encryption.

enterprise configuration
A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also configuration manager, managed server, profile, subscription.

enterprise logging
The process of sending events from a server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also event.

error log
A data set or file that is used to record error information about a product or system.

estimated capacity
The available space, in megabytes, of a storage pool.

event An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process. See also enterprise logging, receiver.

event record
A database record that describes actual status and results for events.

event server
A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)
For processor, real and virtual storage capacities and channel volume, 2 to the power of 60 or 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude
The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup, from space management, or from both backup and space management.

exclude-include list
See include-exclude list.

expiration
The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file
A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated

file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend
To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)
Names or value pairs that are associated with files or directories. There are three classes of extended attributes: user attributes, system attributes, and trusted attributes.

external library
A collection of drives that is managed by the media-management system other than the storage management server.

F

file access time
On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age
For migration prioritization purposes, the number of days since a file was last accessed.

file device type
A device type that specifies the use of sequential access files on disk storage as volumes.

file server
A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space
A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)
A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state
The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also migrated file, premigrated file, resident file.

file system migrator (FSM)
A kernel extension that intercepts all file system operations and provides any space management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

file system state
The storage management mode of a file system that resides on a

workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency
A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See file space ID.

FSM See file system migrator.

full backup
The process of backing up the entire server database. A full backup begins a new database backup series. See also database backup series, database snapshot, incremental backup.

fuzzy backup
A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy
A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified.

G

GB See gigabyte.

General Parallel File System (GPFS™)
A high-performance shared-disk file system that can provide data access from nodes in a clustered system environment. See also information lifecycle management.

gigabyte (GB)
For processor storage, real and virtual storage, and channel volume, 10 to the power of nine or 1,073,741,824 bytes. For disk storage capacity and communications volume, 1,000,000,000 bytes.

global inactive state
The state of all file systems to which space management has been added when space management is globally deactivated for a client node.

Globally Unique Identifier (GUID)
An algorithmically determined number that uniquely identifies an entity within a system. See also Universally Unique Identifier.

GPFS See General Parallel File System.

GPFS node set
A mounted, defined group of GPFS file systems.

group backup
The backup of a group containing a list of files from one or more file space origins.

GUID See Globally Unique Identifier.

H

hierarchical storage management (HSM)
A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity. See also hierarchical storage management client, recall, storage hierarchy.

hierarchical storage management client (HSM client)
A client program that works with the server to provide hierarchical storage management (HSM) for a system. See also hierarchical storage management, management class.

HSM See hierarchical storage management.

HSM client
See hierarchical storage management client.

I

ILM See information lifecycle management.

image A file system or raw logical volume that is backed up as a single object.

image backup
A backup of a full file system or raw logical volume as a single object.

inactive file system
A file system for which space management has been deactivated. See also active file system.

inactive version
A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. See also active version, backup version.

include-exclude file
A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also include-exclude list.

include-exclude list
A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services. See also include-exclude file.

incremental backup
The process of backing up files or directories, or copying pages in the database, that are new or changed since the last full or incremental backup. See also selective backup.

individual mailbox restore
See mailbox restore.

information lifecycle management (ILM)
A policy-based file-management system for storage pools and file sets. See also General Parallel File System.

inode The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An inode contains the node, type, owner, and location of a file.

inode number
A number specifying a particular inode file in the file system.

IP address
A unique address for a device or logical unit on a network that uses the Internet Protocol standard.

J

job file
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface. See also migration job.

journal-based backup
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K

KB See kilobyte.

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 2 to the power of 10 or 1,024 bytes. For disk storage capacity and communications volume, 1,000 bytes.

L

LAN See local area network.

LAN-free data movement
The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network.

LAN-free data transfer
See LAN-free data movement.

leader data
Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library

1. A repository for demountable recorded media, such as magnetic disks and magnetic tapes.
2. A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client
A server that uses server-to-server communication to access a library that is managed by another storage management server. See also library manager.

library manager
A server that controls device operations when multiple storage management servers share a storage device. See also library client.

local

1. Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line. See also remote.
2. For hierarchical storage management products, pertaining to the destination of migrated files that are being moved. See also remote.

local area network (LAN)
A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volume
Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See loopback virtual file system.

logical file
A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also aggregate, physical file, physical occupancy.

logical occupancy
The space that is used by logical files in a storage pool. This space does not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy. See also physical occupancy.

logical unit number (LUN)
In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

logical volume
A portion of a physical volume that contains a file system.

logical volume backup
A back up of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)
Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)
A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LUN See logical unit number.

LVSA See Logical Volume Snapshot Agent.

M

macro file
A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. See also Tivoli Storage Manager command script.

mailbox restore
A function that restores Microsoft Exchange Server data (from IBM Data Protection for Microsoft Exchange backups) at the mailbox level or mailbox-item level.

managed object
A definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server.

managed server
A server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also configuration manager, enterprise configuration, profile, subscription.

management class
A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also bind, copy group, hierarchical storage management client, policy set, rebind.

maximum transmission unit (MTU)
The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

MB See megabyte.

media server
In a z/OS® environment, a program that provides access to z/OS disk and tape storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

megabyte (MB)
For processor storage, real and virtual storage, and channel volume, 2 to the 20th power or 1,048,576 bytes. For disk storage capacity and communications volume, 1,000,000 bytes.

metadata
Data that describes the characteristics of data; descriptive data.

migrate
To move data to another location, or an application to another computer system.

migrated file
A file that has been copied from a local file system to storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on

the local file system. On Windows systems, creation of the stub file is optional. See also file state, premigrated file, resident file, stub file.

migration
The process of moving data from one computer system to another, or an application to another computer system.

migration job
A specification of files to migrate, and actions to perform on the original files after migration. See also job file, threshold migration.

migration threshold
High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring
The process of writing the same data to multiple disks at the same time. The mirroring of data protects it against data loss within the database or within the recovery log.

mode A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See also absolute mode, modified mode.

modified mode
In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also absolute mode, mode.

mount limit
The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also mount point.

mount point
A logical drive through which volumes are accessed in a sequential access device class. For removable media device types, such as tape, a mount point is a logical drive associated with a physical drive. For the file device type, a mount point is a logical drive associated with an I/O stream. See also mount limit.

mount retention period
The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period
The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See maximum transmission unit.

N

Nagle algorithm
An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe
A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS file server
See network-attached storage file server.

NAS file server node
See NAS node.

NAS node
A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system
A file system that is locally added to the file server and is not added for space management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

native format
A format of data that is written to a storage pool directly by the server. See also non-native data format.

NDMP
See Network Data Management Protocol.

NetBIOS (Network Basic Input/Output System)
A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

network-attached storage file server (NAS file server)
A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System
See NetBIOS.

Network Data Management Protocol (NDMP)
A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate
A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

node name
A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class
A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also privilege class.

non-native data format
A format of data that is written to a storage pool that differs from the format that the server uses for operations. See also native format.

O

offline volume backup
A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup
A backup in which the volume is available to other system applications during the backup operation.

open registration
A registration process in which users can register their workstations as client nodes with the server. See also closed registration.

operator privilege class
A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also privilege class.

options file
A file that contains processing options. See also client system-options file, client user-options file.

originating file system
The file system from which a file was migrated. When a file is recalled, it is returned to its originating file system.

orphaned stub file
A file for which no migrated file can be found on the server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is modified to contact a server that is different than the one to which the file was migrated.

P

packet In data communication, a sequence of binary digits, including data and control signals, that are transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode
A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation
A process that creates and stores a new password in an encrypted

password file when the old password expires. Automatic generation of a password prevents password prompting.

path
An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character
See wildcard character.

physical file
A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also aggregate, logical file, physical occupancy.

physical occupancy
The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also logical file, logical occupancy, physical file.

plug-in
A separately installable software module that adds function to an existing program, application, or interface.

policy domain
A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain. See also active policy set, domain.

policy privilege class
A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also privilege class.

policy set
A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also active policy set, management class.

premigrated file
A file that has been copied to server storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and in server storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. See also file state, migrated file, resident file.

premigrated files database
A database that contains information about each file that has been premigrated to server storage.

premigration
The process of copying files that are eligible for migration to server storage, but leaving the original file intact on the local file system.

premigration percentage
A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool
A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also copy storage pool, server storage, storage pool, storage pool volume.

privilege class
A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also authority, node privilege class, operator privilege class, policy privilege class, storage privilege class, system privilege class.

profile
A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also configuration manager, enterprise configuration, managed server.

profile association
On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

Q

quota

1. For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the amount of data that can be migrated and premigrated from a file system to server storage.
2. For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization
The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

raw logical volume
A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

rebind
To associate all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. See also bind, management class.

recall To copy a migrated file from server storage back to its originating file system using the hierarchical storage management client. See also selective recall.

receiver
A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the server console and activity log. See also event.

reclamation
The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold
The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation
The process of ensuring consistency between the original data repository and the larger system where the data is stored for backup. Examples of larger systems where the data is stored for backup are storage servers or other storage systems. During the reconciliation process, data that is identified as no longer needed is removed.

recovery log
A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register
To define a client node or administrator ID that can access the server.

registry
A repository that contains access and configuration information for users, systems, and software.

remote
For hierarchical storage management products, pertaining to the origin of migrated files that are being moved. See also local.

resident file
On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in server storage. On a UNIX or Linux system, a complete file on a local file system that has not been migrated or premigrated, or that has been recalled from server storage and modified. See also file state.

restore
To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention
The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve
To copy archived information from the storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool. See also archive.

root user
A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN See storage area network.

schedule
A database record that describes client operations or administrative commands to be processed. See also administrative command schedule, client schedule.

scheduling mode
The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume
A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use. See also volume.

script A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as they are run. See also Tivoli Storage Manager command script.

Secure Sockets Layer (SSL)
A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup
The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. See also incremental backup.

selective migration
The process of copying user-selected files from a local file system to server storage and replacing the files with stub files on the local file system. See also demand migration, threshold migration.

selective recall
The process of copying user-selected files from server storage to a local file system. See also recall, transparent recall.

serialization
The process of handling files that are modified during backup or archive processing. See also shared dynamic serialization, shared static serialization, static serialization.

server A software program or a computer that provides services to other software programs or other computers. See also client.

server options file
A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

server-prompted scheduling mode
A client/server communication technique where the server contacts the client node when tasks must be done. See also client-polling scheduling mode.

server storage
The primary, copy, and active-data storage pools that are used by the server to store user files such as backup versions, archive copies, and files

migrated from hierarchical storage management client nodes (space-managed files). See also active-data pool, copy storage pool, primary storage pool, storage pool volume, volume.

session
A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data for the duration of the session. See also administrative session.

session resource usage
The amount of wait time, processor time, and space that is used or retrieved during a client session.

shadow copy
A snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

shadow volume
The data stored from a snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

shared dynamic serialization
A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. The backup-archive client retries the backup or archive operation a number of times; if the file is being modified during each attempt, the backup-archive client will back up or archive the file on its last try. See also dynamic serialization, serialization, shared static serialization, static serialization.

shared library
A library device that is used by multiple storage manager servers. See also library.

shared static serialization
A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. The client attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also dynamic serialization, serialization, shared dynamic serialization, static serialization.

snapshot
An image backup type that consists of a point-in-time view of a volume.

space-managed file
A file that is migrated from a client node by the hierarchical storage management (HSM) client. The HSM client recalls the file to the client node on demand.

space management
See hierarchical storage management.

space monitor daemon
A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file
A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file
On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL See Secure Sockets Layer.

stabilized file space
A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Stanzas are usually separated by blank lines or colons, and each stanza has a name.

startup window
A time period during which a schedule must be initiated.

static serialization
A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, the backup-archive client cannot back up or archive the file. See also dynamic serialization, serialization, shared dynamic serialization, shared static serialization.

storage agent
A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)
A dedicated storage network tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy
A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also storage pool.

storage pool
A named set of storage volumes that is the destination that is used to store client data. See also active-data pool, copy storage pool, primary storage pool, storage hierarchy.

storage pool volume
A volume that has been assigned to a storage pool. See also active-data pool, copy storage pool, primary storage pool, server storage, volume.

storage privilege class
A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also privilege class.

stub A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file
A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary

to recall a migrated file from server storage. It also contains additional information that can be used to eliminate the need to recall a migrated file. See also migrated file, resident file.

stub file size
The size of a file that replaces the original file on a local file system when the file is migrated to server storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription
In a storage environment, the process of identifying the subscribers to which the profiles are distributed. See also enterprise configuration, managed server.

system privilege class
A privilege class that gives an administrator the authority to issue all server commands. See also privilege class.

T

tape library
A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix
The high-level-qualifier of the file name or the data set name in the standard tape label.

target node
A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See trusted communications agent.

TCP/IP
See Transmission Control Protocol/Internet Protocol.

threshold migration
The process of moving files from a local file system to server storage based on the high and low thresholds that are defined for the file system. See also automatic migration, demand migration, migration job, selective migration.

throughput
In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout
A time interval that is allotted for an event to occur or complete before operation is interrupted.

Tivoli Storage Manager command script
A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can

run from any interface to the server. The script can include substitution for command parameters and conditional logic. See also macro file, script.

tombstone object
A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)
An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types. See also communication method.

transparent recall
The process that is used to automatically recall a migrated file to a workstation or file server when the file is accessed. See also selective recall.

trusted communications agent (TCA)
A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See Universal Naming Convention.

Unicode
A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus many classical and historical texts.

Unicode-enabled file space
Unicode file space names provide support for multilingual workstations without regard for the current locale.

Universally Unique Identifier (UUID)
The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier. See also Globally Unique Identifier.

Universal Naming Convention (UNC)
The server name and network name combined. These names together identify the resource on the domain.

UTF-8 Unicode Transformation Format, 8-bit encoding form, which is designed for ease of use with existing ASCII-based systems. The CCSID value for data in UTF-8 format is 1208. See also UCS-2.

UUID See Universally Unique Identifier.

V

validate
To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For example, the validation process checks whether the policy set contains a default management class.

version
A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space
A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual mount point
A directory branch of a file system that is defined as a virtual file system. The virtual file system is backed up to its own file space on the server. The server processes the virtual mount point as a separate file system, but the client operating system does not.

virtual volume
An archive file on a target server that represents a sequential media volume to a source server.

volume
A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also scratch volume, server storage, storage pool, storage pool volume.

volume history file
A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service (VSS)
A set of Microsoft application-programming interfaces (APIs) that are used to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See Volume Shadow Copy Service.

VSS Backup
A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Microsoft Exchange data. This copy can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore
An operation that restores data from a local snapshot. The snapshot is the VSS backup that resides on a local shadow volume. The restore operation retrieves the data by using a file-level copy method.

VSS Instant Restore
An operation that restores data from a local snapshot. The snapshot is the

VSS backup that resides on a local shadow volume. The restore operation retrieves the data by using a hardware assisted restore method (for example, a FlashCopy operation).

VSS offloaded backup
A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Microsoft Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore
A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on Tivoli Storage Manager server storage to their original location.

W

wildcard character
A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workload partition (WPAR)
A partition within a single operating system instance.

workstation
A terminal or personal computer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name (WWN)
A 64-bit, unsigned name identifier that is unique.

WPAR See workload partition.

WWN See worldwide name.

Index

Special characters

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

V

W

Product Number: 5608-E06

Printed in USA