

IBM Spectrum Protect

# ***Migrating IBM Spectrum Protect Servers from One Operating System to Another***

---

Document Version: 1.07

*Daniel Eisenberg, Larry Schuh, Ken Hannigan*

**IBM**

© Copyright IBM Corporation 2014, 2020.

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

## Contents

Introduction.....	1
Prerequisites.....	1
Database backup and restore.....	1
DSMSERV EXTRACTDB and DSMSERV INSERTDB utilities.....	2
Migration Process.....	3
Planning the migration.....	3
Preparing the target system for the migration.....	4
Preparing the source system for the migration.....	6
Migrating the server by using database backup and restore operations.....	7
Migrating the server by using database extract and insert utilities.....	8
Migrate by using the network method.....	9
Migrate by using the media method.....	10
Extracting the data to media.....	10
Loading the extracted data into the target database.....	11
Taking the first steps after migration.....	12
Appendix A. Utilities for server migration.....	15
DSMSERV EXTRACTDB (Extract data from a server database).....	15
Syntax.....	15
Parameters.....	15
Examples.....	16
The manifest file for data extraction to media.....	17
DSMSERV LOADFORMAT (Format a database).....	17
Syntax.....	17
Parameters.....	18
DSMSERV INSERTDB (Move a server database into an empty database).....	19
Requirements for insertion by using media.....	19
Syntax.....	20
Parameters.....	20
Notices.....	23
Trademarks.....	24

## Introduction

You can migrate an IBM Spectrum Protect™ (formerly IBM Tivoli® Storage Manager) server from one operating system to another.

**Tip:** The product now known as IBM Spectrum Protect was named IBM Tivoli Storage Manager in releases earlier than Version 7.1.3. To learn more about the rebranding transition, see [technote 1963634](#).

Depending on the version of the server, and the operating systems that are involved in the migration, different migration methods are available:

- [Tivoli Storage Manager V5 servers on IBM z/OS® systems](#) can be migrated to IBM AIX® or Linux on System z® as part of an upgrade to V6.3.0, V7.1.0, or later.
- [Tivoli Storage Manager V5 servers on AIX, HP-UX, and Solaris operating systems](#) can be migrated to Linux x86\_64 as part of an upgrade to V6.3.4, V7.1.0, or later.
- V6.3 or V7.1 servers can be migrated by using [node replication](#) to [move one or more nodes to another server](#) on any supported operating system.
- Any server version can be migrated by using [export](#) and [import](#) operations to move one or more nodes to another server on any supported operating system.

This paper describes two additional cross-platform migration methods to migrate a server's database to another operating system:

- The first relies on database backup and restore operations.
- The second uses a variation of the V5 upgrade utilities, which have been adapted to work with a V7 database.

As with most other methods, a limited set of server levels and operating systems support these methods.

## Prerequisites

Your system must meet the hardware and software requirements. The different migration methods have different prerequisites.

### Database backup and restore

You can use database backup and restore operations to migrate a server from one AIX, HP-UX, Linux, or Oracle Solaris operating system to a different AIX, HP-UX, Linux, or Solaris operating system if both systems have the same endianness.

The following table lists operating systems with the same endianness.

Big endian	AIX, Solaris SPARC, HP-UX, Linux on System z, and Linux on IBM Power Systems™ (big endian)
Little endian	Linux x86_64, Solaris x86_64, and Linux on Power Systems (little endian)

The source server must be at one of the following levels:

- V6.2.5 or later
- V6.3.4 or later
- V7.1.0 or later

In all cases, you migrate a server to a target system that has the same server level installed. After you migrate

the server, you can upgrade it to any level that is supported on the target system.

### DSMSERV EXTRACTDB and DSMSERV INSERTDB utilities

You can use the DSMSERV EXTRACTDB and DSMSERV INSERTDB utilities to migrate a server from one operating system to another. Depending on the operating systems and the server levels involved, different migration paths are possible.

Source server level	Target server level	Operating system of source server	Operating system of target server	Is it possible to have different levels of IBM Spectrum Protect software installed on the source and target servers? <sup>1</sup>
V7.1.5 or later	V7.1.5 or later	AIX	Linux x86_64	No
V7.1.7 or later	V7.1.7 or later	HP-UX	Linux x86_64	No
V7.1.7 or later	V7.1.7 or later	Solaris	Linux x86_64	No
V7.1.7 or later	V7.1.7 or later	Linux on Power Systems (big endian)	Linux x86_64	No
V7.1.8 or later	V7.1.8 or later	AIX	Linux on Power Systems (little endian)	No
V7.1.8 or later	V8.1.4 or later	Linux on Power Systems (big endian)	Linux on Power Systems (little endian)	Yes
V7.1.8 or later	V8.1.5 or later	Linux x86_64	AIX	Yes
V8.1.1 or later	V8.1.5 or later	Linux on Power Systems (little endian)	AIX	Yes
V8.1.9 or later	V8.1.9 or later	Linux on System Z	Linux x86_64	Yes

<sup>1</sup> The target server must have the same level of IBM Spectrum Protect software installed as the source server unless different levels are explicitly supported. After you migrate the server, you can upgrade it to any level that is supported on the target system.

## Migration Process

The migration process consists of several steps:

1. Planning the migration
2. Preparing the source and target systems
3. Moving the server database from the source system to the target system, either by
  - Backing up the server database on the source system and restoring it on the target system when the two systems have the same endianness, or by
  - Extracting the contents of the database on the source system and inserting the contents into an empty database on the target system when moving from AIX, HP-UX, Oracle Solaris, or Linux on Power Systems to Linux x86\_64.
4. Reconfiguring the server on the target system to account for differences between the source and target systems

Most of the planning, preparation, and reconfiguration steps are required to ensure that the server can continue to access its storage from the target system. For servers that use tape storage, you might have to install device drivers, and you must ensure that the target system can access the devices. You might have to update configurations to account for changes in device names.

Ensuring access to data for servers that primarily use disk storage can be more difficult. Few AIX, HP-UX, Linux, or Solaris operating systems have a file system in common with a different operating system, so it is usually not possible to connect an existing disk LUN to the target system and use it as-is.

If the source and target operating systems do not share a common file system, use one of the following methods:

- Move all data from disk to tape before the migration
- Leave the disk data on the source system, and make the data available to the target system via a Network File System (NFS) server

On the other hand, if the file system can be used by both the source and target operating systems (such as IBM Spectrum Scale on both AIX and Linux), you can simply mount the file system on the target system after you complete the migration.

**Using documentation links:** The following procedures include links to supporting information in IBM Knowledge Center. Most of the links are to Tivoli Storage Manager V7.1.1 topics. These topics are similar to topics from earlier versions of the server and can be used as-is, even if the server that is being migrated is at V6.2 or V6.3. However, if you prefer to use documentation that was specifically written for the server version that is being migrated, go to IBM Knowledge Center:

V6.2: [http://www.ibm.com/support/knowledgecenter/SSGSG7\\_6.2.0/com.ibm.itsm.ic.doc/welcome.html](http://www.ibm.com/support/knowledgecenter/SSGSG7_6.2.0/com.ibm.itsm.ic.doc/welcome.html)

V6.3: [http://www.ibm.com/support/knowledgecenter/SSGSG7\\_6.3.4/KC\\_ditamaps/product\\_welcome.html](http://www.ibm.com/support/knowledgecenter/SSGSG7_6.3.4/KC_ditamaps/product_welcome.html)

V7.1: [http://www.ibm.com/support/knowledgecenter/SSGSG7\\_7.1.0/com.ibm.itsm.ic.doc/welcome.html](http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.0/com.ibm.itsm.ic.doc/welcome.html)

V8.1: [http://www.ibm.com/support/knowledgecenter/en/SSEQVQ\\_8.1.0/tsm/welcome.html](http://www.ibm.com/support/knowledgecenter/en/SSEQVQ_8.1.0/tsm/welcome.html)

**Tip:** In IBM Knowledge Center, you can click the displayed product name and version at the top left to display a different version of the same topic.

### *Planning the migration*

Plan the migration carefully to minimize server downtime. Thoroughly review this document before you start the server migration. If you are migrating multiple servers on multiple systems, plan to upgrade one server first in a test environment.

1. Ensure that each server that you plan to migrate is at V6.2.5 or later, V6.3.4 or later, or V7.1.0 or later.

- If the server is being migrated from AIX to Linux x86\_64, it must be at level V7.1.5 or later.
- If the server is being migrated from HP-UX, Oracle Solaris, or Linux on Power Systems to Linux x86\_64, it must be at level V7.1.7 or later.
- If the server is being migrated from AIX to Linux on Power Systems (little endian), it must be at level V8.1.1 or later.
- If the server is being migrated from Linux on Power Systems (big endian) to Linux on Power Systems (little endian), the source server must be at level V7.1.8 or later, and the target server must be at level V8.1.4 or later.

If the server is not at the appropriate level, upgrade it before you start the migration process. For instructions about upgrading the server, see the online documentation:

[Upgrading from Tivoli Storage Manager V6.1](#)

[Upgrading to Tivoli Storage Manager V6.3 or later](#)

[Upgrading to Tivoli Storage Manager V7.1](#)

[Upgrading to V8.1](#)

2. Review the hardware and software requirements for running the server on the target system. For more information, see [Prerequisites](#) and [technote 1243309](#).
3. Plan data movement and ensure that the devices that you plan to use are supported on both operating systems. For information about supported devices, see [Supported Devices for AIX, HP-UX, Solaris, and Windows® systems](#) and [Supported Devices for Linux systems](#). Device support can differ, depending on the operating system, for example:
  - Device classes with DEVTYPE=GENERICTAPE are not supported on Linux. If the target system is installed on Linux, all data on generic tape devices must be migrated or backed up to supported tape devices before migration.
  - If the source and target operating systems do not share a common file system, all data in DISK or FILE storage pools must be migrated or backed up to tape devices that are supported on the target system. Alternatively, you can export the storage pools via an NFS server so that the data can be read from the target system. For requirements and restrictions on the use of NFS file systems with the server, see [technote 1470193](#). DISK storage pools that use raw logical partitions cannot be exported over NFS, and must be migrated or backed up to tape.
4. Provision at least as much disk storage for the target system as is used by the source system.
5. If you plan to back up large storage pools to tape, consider starting this process before you migrate the server. In this way, you might be able to reduce the time that you spend backing up these storage pools while the server is unavailable.

### ***Preparing the target system for the migration***

By completing these steps before you migrate the server, you can reduce the time that the server is unavailable.

1. Install the server on the target system. If you are using database backup and restore, the target server must be at the same level as the source system. For example, if a V6.2.5 server is installed on a source system, install V6.2.5 on the target system. For information about installing a V6.2 or V6.3 server, see [Installing the server](#). For V7.1, see [Installing Tivoli Storage Manager](#). For V8.1, see [Installing and upgrading the server](#).  
If you use the DSMSEV EXTRACTDB and DSMSEV INSERTDB utilities to migrate the server, beginning in V8.1.4, the target server can be at a later level than the source server.
2. Install the device drivers for the devices that you plan to use on the target system. Configure the system to ensure that the devices can be accessed from the target system.
3. Complete the following steps on the target system for each server instance that is running on the source system:

- a. Create a user account (a user ID and group) that will own the server instance. Ensure that this account has the same user name as on the source system. If the target system will use NFS to access DISK or FILE storage pool volumes remotely, the account must have the same user ID and group ID as on the source system.

For instructions about creating a user account, see [Creating the user ID and directories for the server instance](#).

Change to the user account that you created.

- b. Create the directories that the server requires, starting with the server instance directory. Create the database, active log, and archive log directories for the server instance by using the same paths as on the source system, or by using different paths. Also, create the log mirror and secondary archive log directories, if these exist on the source system, by using the same paths as on the source system. For instructions, see [Creating the user ID and directories for the server instance](#).

- c. Configure a new server instance by using the user account and directories that you created in the previous steps. Take one of the following actions:

- o To configure the new server instance by using the configuration wizard, follow the instructions for your operating system:

AIX: [Configuring Tivoli Storage Manager using the configuration wizard](#)

HP-UX: [Configuring Tivoli Storage Manager using the configuration wizard](#)

Linux: [Configuring Tivoli Storage Manager using the configuration wizard](#)

Solaris: [Configuring Tivoli Storage Manager using the configuration wizard](#)

- o To configure the new server instance manually, follow the instructions for your operating system:

AIX: [Configuring the server instance manually](#)

HP-UX: [Configuring the server instance manually](#)

Linux: [Configuring the server instance manually](#)

Solaris: [Configuring the server instance manually](#)

- d. To prepare the new server instance for the migration, clean up the temporary database and other files that could interfere with the migration process:

- i. Use the DSMSEV REMOVEVDB utility to remove the temporary database. Issue the following command:

```
# dsmserv removedb TSMDB1
```

For more information about using this command, see [DSMSEV REMOVEVDB \(Remove a database\)](#).

- ii. Manually delete the archived log volumes in the archive log directory. The DSMSEV REMOVEVDB utility does not remove these volumes, and you cannot format a new database by using the DSMSEV LOADFORMAT utility while the volumes exist.
- iii. Manually delete the SSL key database, exported certificates, and password stash files that were created when the instance was created. These files are associated with the temporary server database, and cannot be used with the migrated database:

- cert.kdb
- cert256.arm
- cert.arm
- cert.crl
- cert.rdb
- cert.sth

- e. If you run the DSMSEV EXTRACTVDB and DSMSEV INSERTVDB utilities to migrate a server from AIX, HP-UX, Oracle Solaris, or Linux on Power Systems to Linux x86\_64, use the DSMSEV



LOADFORMAT utility to create a new, empty database into which the source database contents will be loaded. For more information about using this utility, see [DSMSERV LOADFORMAT \(Format a database\)](#).

4. If the target system will use NFS to access DISK and FILE storage pool volumes remotely, export the file systems from the source system and mount them on the target system. Similarly, if the target system will use IBM Spectrum Scale to access DISK and FILE storage pools, configure IBM Spectrum Scale and mount the volumes on the target system. Ensure that all volumes can be accessed by using the same path as on the source system.
5. Copy the following configuration files from the source server to the target server:
  - The server options file, `dsmserv.opt`.
  - The device configuration file.
  - The volume history file.
  - The server password file, `dsmserv.pwd`, if available. This file is created only by servers at level V7.1.3 or later.
  - The server password files `dsmkeydb.kdb` and `dsmkdb.sth`, if available. These files are created in lieu of the `dsmserv.pwd` file by server levels V7.1.8 or later and V8.1.2 or later.

### ***Preparing the source system for the migration***

When you are ready to begin the migration, complete the following steps on the source system:

1. [Disable sessions](#).
2. [Back up or migrate data stored on DISK and FILE devices](#). This step is not required if all of the source server's DISK and FILE storage pools are to be accessed remotely via NFS.
3. If the target system runs on a Linux operating system, [migrate data that is stored on GENERICTAPE devices](#).
4. If the backup set volumes of the source server will not be accessed remotely via NFS, [move backup sets that are stored on FILE devices](#).
5. To prevent potential issues during and after the migration, make the following adjustments to settings on your server and clients. These adjustments must be completed so that you can revert to using the source server after the migration, if a problem occurs.
  - a. For each sequential-access storage pool on tape, use the UPDATE STGPOOL command to set the REUSEDELAY parameter. Set the parameter to the number of days during which you want to be able to revert to the source server, if necessary. For example, if you want to be able to revert to the source server for up to 30 days after the migration, set the REUSEDELAY parameter to 31 days by issuing the following command:

```
update stgpool sequential_access_storage_pool reusedelay=31
```

where `sequential_access_storage_pool` specifies the name of the storage pool. For more information about setting the REUSEDELAY parameter, see [UPDATE STGPOOL \(Update a storage pool\)](#).

- b. For each copy storage pool on tape, use the QUERY STGPOOL command to determine the value of the RECLAIM parameter. Note the value for future reference. If you must revert to the source server at any time during the migration process, it is useful to know the value of the RECLAIM parameter so that you can restore the setting. Issue the following command:

```
query stgpool sequential_access_storage_pool format=detailed
```

where `sequential_access_storage_pool` specifies the name of the copy storage pool. In the system output, the value of the RECLAIM parameter can be found in the **Reclamation Threshold** field. For

more information about determining the value of the RECLAIM parameter, see [QUERY STGPOOL \(Query storage pools\)](#).

- c. For each copy storage pool on tape, use the UPDATE STGPOOL command to set the RECLAIM parameter to 100, meaning 100%. For example, issue the following command:

```
update stgpool copy_storage_pool reclaim=100
```

where *copy\_storage\_pool* specifies the name of the copy storage pool.

- d. For each tape, DISK, or FILE volume that was used by the server on the source system, specify the read-only access mode. Issue the following command:

```
update volume volume_name access=readonly
```

where *volume\_name* specifies the name of the tape, DISK, or FILE volume. For more information about specifying read-only access mode, see [UPDATE VOLUME \(Change a storage pool volume\)](#).

- e. For each client that plays an essential role in your system, verify that the value for the schedlogretention client option is set to retain the client schedule log for a sufficient time.

The client schedule log can be useful if you must revert the system to the source server. If the retention period for the schedule log is too short, the schedule log information might be deleted too soon.

For example, to prune the log every 45 days and save the log entries, add the following option:

```
schedlogretention 45 S
```

Add the schedlogretention client option to the dsm.sys file within the server stanza for the server that is being migrated. For more information about pruning the log, see [Schedlogretention](#).

6. Determine whether any server processes are running. Either cancel them or allow them to finish. Use the following commands:

```
query process
```

```
cancel process process_number
```

where *process\_number* specifies the number of the process. Allow time for the processes to be stopped. Some processes might take a considerable amount of time to stop.

7. [Create a summary of database contents](#).

## ***Migrating the server by using database backup and restore operations***

Complete the following steps:

1. On the source system, back up the server database to tape. Issue the following command:

```
backup db type=full devclass=tape_device
```

where *tape\_device* specifies the name of a tape device class.

Beginning with V7.1.1, you can compress database backups. Do not use compression for this backup operation because compressed database backups cannot be restored on a different operating system. If you have V7.1.1 or later installed, issue the following command to back up the server database:

```
backup db type=full devclass=tape_device compress=no
```

For more information about backing up the database, see [BACKUP DB \(Back up the database\)](#).

2. Unmount any tapes that are mounted. Use the following commands:

```
query mount
```

```
dismount volume volume_name
```

where *volume\_name* specifies the name of the volume. For more information about unmounting tapes, see [QUERY MOUNT \(Display information on mounted sequential access volumes\)](#) and [DISMOUNT VOLUME \(Dismount a volume by volume name\)](#).

3. Back up the volume history and device configuration files to another directory. Issue the following commands:

```
backup devconfig filenames=devconfig_file_name
backup volhistory filenames=volhist_file_name
```

where *devconfig\_file\_name* specifies the file in which to store device configuration information, and *volhist\_file\_name* specifies the file in which to store volume history information.

For more information about using the commands, see [BACKUP DEVCONFIG \(Create backup copies of device configuration information\)](#) and [BACKUP VOLHISTORY \(Save sequential volume history information\)](#).

4. To stop the server, issue the following command:

```
halt
```

5. If you did not complete the steps in [Preparing the target system for the migration](#), complete them now.
6. Move copies of the volume history file, device configuration file, and server options file to the target system.
7. On the target system, edit the device configuration file to update the device path information for any devices that are required to restore the database. For each relevant DEFINE PATH entry in the device configuration file, update the DEVICE parameter to the corresponding path on the target system.
8. Restore the server database by using the following command:

```
# dsmserv restore db
```

For more information about restoring the database, see [DSMSERV RESTORE DB \(Restore a database to its most current state\)](#). If you set up the target system by using different directories than the source server, specify the ON= and ACTIVELOGDIRECTORY= parameters on the DSMSERV RESTORE DB command to identify the new directories that you created.

## ***Migrating the server by using database extract and insert utilities***

The database extract and insert utilities are used to extract the contents of the source server's database and insert the data into an empty database on the target system. Two utilities are available:

- DSMSERV EXTRACTDB
- DSMSERV INSERTDB

Both utilities are based on the Tivoli Storage Manager V5 Upgrade Utilities. Starting with IBM Spectrum Protect V7.1.5, the utilities were enhanced to work with data from an IBM DB2® database that is created with an IBM Spectrum Protect V7.1.5 or later server.

The database extract and insert utilities support the following two methods of operation:

### **Network method**

The database contents are extracted from the source database and transmitted over a network to the target server, where they are simultaneously inserted into the target database.

### **Media method**

The database contents are extracted from the source database and stored on tape or FILE media. The contents are inserted into the target database later.

**Tip:** The database extract and insert utilities run in the foreground and write all output to the console. Before you

stop the server to run the DSMSEV EXTRACTDB utility, update any event logging rules for the console to allow all messages to be displayed. Complete the following steps:

1. Display the current event logging rules by issuing the following commands. Save the output so that you can use it to reinstate the rules after the migration process is completed.
  - QUERY ENABLED CONSOLE
  - QUERY EVENTRULES CONSOLE
2. To allow all events to be written to the console, issue the following command:
  - ENABLE EVENTS CONSOLE ALL

After you successfully run the DSMSEV INSERTDB utility, use the output of the queries in Step 1 to reestablish the event rules that you want to have in place in production mode.

### Migrate by using the network method

Move the database by starting the insertion process for the target server to accept the server database. Then, start the process for the source server to extract and send the database.

#### **Before you begin**

Ensure that the source and target servers are not running.

#### *Procedure*

To move the server database over a network, complete the following steps:

1. Verify that there is a good network connection between the two systems.
2. Start the insertion process on the source server to accept the database. Use the DSMSEV INSERTDB command. To monitor the process, direct the output of the process to a file. For example, to start the server, allow the default time of 60 minutes for the source server to contact the source server, and direct the process output to the insert.out file, use this command:

```
nohup /opt/tivoli/tsm/server/bin/dsmserv insertdb \  
sesswait=60 >insert.out 2>&1 &
```

For more information about inserting data, see [DSMSEV INSERTDB \(Move a server database into an empty database\)](#).

The target server starts and waits up to 60 minutes to be contacted by the source server. Some time might pass during which no messages are issued. During this time, DB2 operations are running in the background.

**Optional:** To verify that operations are continuing as expected, monitor the processor and I/O usage for the server process and the corresponding DB2 process.

3. Monitor the output of the DSMSEV INSERTDB process. Verify that the DSMSEV INSERTDB process issues the following message before you continue to the next step:

```
ANR1336I INSERTDB: Ready for connections from the source server
```

Issue the following command to monitor the process output in the insert.out file:

```
tail -f insert.out
```

4. Start the data extraction from the source server by using the DSMSEV EXTRACTDB command. Issue the command from the source server's instance directory. Specify the TCP/IP address and port for the target server. Direct the output of the process to a file for monitoring. For example, enter the following command on one line:

```
nohup /opt/tivoli/tsm/server/bin/dsmserv extractdb \  
hladdress=9.11.25.124 lladdress=1500 >extract.out 2>&1 &
```

For more information about extracting data, see [DSMSERV EXTRACTDB \(Extract data from a server database\)](#).

5. Monitor the processes for errors and warning messages, and for items that might require attention. From the instance directory for the server that you are upgrading, issue the following command to monitor the extraction process:

```
tail -f extract.out
```

The length of time that the process runs depends on the size of the database, the hardware, and the network.

6. Examine the process outputs for the extraction and insertion processes to find the messages that indicate the success or failure of the operations.

Process	Success message	Failure message
Extraction	ANR3402I EXTRACTDB: Process 1, database extract, has completed.	ANR3414E EXTRACTDB: Process 1, database extract, has completed with errors.
Insertion	ANR1395I INSERTDB: Process 1, database insert, has completed.	ANR1396E INSERTDB: Process 1, database insert, has completed with errors.

## Migrate by using the media method

### Extracting the data to media

Use the DSMSERV EXTRACTDB utility to extract data from the source server database and store it on a tape device.

#### *Procedure*

1. Log in by using the server's instance ID on the source system.
2. Ensure that the storage device is available. The server database and the device configuration file must contain a valid device class definition for the device.
3. From the instance directory for the server that you are migrating, issue the command to start the extraction. Direct the output of the process to a file for monitoring. For example, issue the following command on one line:

```
nohup /opt/tivoli/tsm/server/bin/dsmserv extractdb \  
devclass=tape manifest=./manifest.txt >extract.out 2>&1 &
```

For more information about extracting data, see [DSMSERV EXTRACTDB \(Extract data from a server database\)](#).

**Tip:** Messages that are issued during the extract operation are not saved in the server activity log. Direct the output of the utility to a file, as shown in the examples, to record the messages.

4. Monitor the process for errors and warning messages, and for items that might require attention. A message near the end of the process output indicates the success or failure of the operation:

Success message:

```
ANR3402I EXTRACTDB: Process 1, database extract, has completed.
```

Failure message:

```
ANR3414E EXTRACTDB: Process 1, database extract, has completed with  
errors.
```

For example, from the instance directory for the server that you are upgrading, issue the following command to monitor the process:

```
tail -f extract.out
```

The length of time that the process runs depends on the database size. The time is approximately the length of time that is required for a full database backup.

### Loading the extracted data into the target database

After you format an empty database by using the DSMSEV LOADFORMAT utility, load the data that you extracted from the source server database.

#### *Before you begin*

Ensure that the following requirements are met before you begin to load the data:

- The manifest file from the DSMSEV EXTRACTDB operation must be available.
- The server options file must contain an entry for the device configuration file.
- The device configuration file must have information about the device class that is specified in the manifest file.
- The media that contains the extracted database must be available to the target server. The device must be physically attached to the system. The permissions must be set to grant access to the media for the instance user ID.

#### *Procedure*

Complete the following steps:

1. Verify that the target server can access the extracted data. The tape drive that is used for the extracted data must be physically attached to the target system.
2. Ensure that the instance user ID has ownership or read/write permission for the manifest file that was created by the extraction process.
3. Log in with the instance user ID on the system where you installed the target server.
4. Copy the manifest file that was created by the extraction process to the target system.
5. On the target server, complete the following steps:
  - a. Verify that the server options file from the source server includes the DEVCONFIG option, and that the option specifies the full path of the device configuration file.
  - b. Verify that the device configuration file from the source server is available in the location that is specified by the DEVCONFIG option.
  - c. Verify that the permissions on the device configuration file allow read access for the instance user ID.
6. Verify that the contents of the device configuration file are correct. The device class that was used for the extraction step is recorded in the manifest file, and that device class must exist and be valid on the target system. Verify entries for tape. For example, the device names might have changed.
7. Verify the contents of the manifest file and edit the file if necessary:
  - a. Ensure that the device names in the manifest file are valid for the target system. Device names for the same device might be different on source and target systems.
  - b. Ensure that the manifest file contains a list of volumes to be used when the extracted data is loaded into the new database.
8. To load an extracted server database into the prepared, empty target database, issue the DSMSEV INSERTDB command. Direct the output of the process to a file for monitoring. For example, enter the

following command on one line:

```
nohup /opt/tivoli/tsm/server/bin/dsmserv insertdb \  
manifest=./manifest.txt >insert.out 2>&1 &
```

For more information about loading an extracted database into a new database, see [DSMSERV INSERTDB \(Move a server database into an empty database\)](#).

9. Monitor the process for error messages, warning messages, and any items that might require your attention. The system displays interim statistics about the process of loading the database. However, there might be periods when no messages are issued. During this time, DB2 operations are running in the background. The length of time that the process runs depends on the size of the database. For more information, see [Example: Estimating the upgrade time based on the database size](#).

**Optional:** Verify that the database is being loaded by monitoring the processor and I/O usage for the server process and the corresponding DB2 process. For example, issue the following command to monitor the process:

```
tail -f insert.out
```

A message in the output of the DSMSERV INSERTDB command indicates the status of the operation:

Success message:

```
ANR1395I INSERTDB: Process 1, database insert, has completed.
```

Failure message:

```
ANR1396E INSERTDB: Process 1, database insert, has completed with  
errors.
```

10. After the data is loaded into the database, remove or check out from the library the tape that holds the extracted data. Prevent the tape from being reused until you are sure that you do not need to run the database-loading operation again.

## ***Taking the first steps after migration***

1. Verify the migration results:
  - a. Start the target server. Monitor the messages that the server issues as it starts. Watch for error and warning messages.
  - b. Set the ownership of, or permissions to, the device special files to ensure that the storage devices of the source server that are supported by the target system are accessible to the target server. If the server storage area network (SAN) discovery function is enabled, the target server should update device information automatically. Compare the device names on the target system with the device parameters of the path definitions on the source server. To see the server path definitions, issue the following command:

```
query path format=detailed
```

Update server device path definitions manually as necessary. For more information about this task, see [QUERY PATH \(Display a path definition\)](#), [UPDATE PATH \(Change a path\)](#), and [SANDISCOVERY](#).

- c. Rerun commands to [get a summary of information in the database](#). Compare the summary with the results for the same commands before the migration.
  - d. Update the network address, and if necessary the port number, that is used by backup-archive clients, administrative clients, storage agents, library client servers, and other servers for communicating with the target server. Use the [SET SERVERHLADDRESS](#) command and, if necessary, the [SET SERVERLLADDRESS](#) command.

Alternatively, consider whether you can use the network address that was used by the source system as the address of the target system. You might also be able to update the domain name

service (DNS) to point to the target system instead of the source system. Consult your network administrator.

- e. Verify that you can connect to the target server by using an administrative client as you did for the source server.
- f. Back up several client nodes and verify that the backup operations work as expected. Ensure that the clients can connect to the target server, and that the data was written to the appropriate storage devices.
- g. Verify that operations such as LAN-free data movement and library sharing work correctly. In some cases, it might be necessary to reset the server verification keys of both the library manager and library client servers by using the [UPDATE SERVER](#) command and specifying the FORCESYNC=YES parameter.

After you are satisfied that the target server is operating as expected and you do not have to revert to the source system, restore any settings that you changed to prepare for the migration back to their original values.

2. [Register licenses.](#)
3. If you did not use the instance configuration wizard to configure the target server, ensure that you have completed the steps to manually configure the system for database backup operations. Then, perform a full database backup.
4. Back up the volume history file.
5. If you must change the host name of the target server, follow the instructions in [Changing the host name for the Tivoli Storage Manager server.](#)
6. If the system uses NFS to access DISK, FILE, and directory-container storage pools, update the storage pools that use NFS-mounted file systems from the source system to use local non-NFS storage for new data. For each random access storage pool, define replacement DISK volumes on non-NFS storage. This step is not required if the system uses IBM Spectrum Scale to access DISK, FILE, and directory-container storage pools.
7. If DISK or FILE data was migrated to tape before the migration, migrate the data from tape back to DISK or FILE devices.
8. Redefine backup sets as necessary by using the [DEFINE BACKUPSET command.](#)
9. Configure the target server for communications by using Secure Sockets Layer (SSL) and Transport Layer Security (TLS) following the instructions at [Setting up Secure Sockets Layer and Transport Layer Security:](#)
  - You might be able to use the same certificates on the target server as were used on the source server if the target server address is covered by the fields in the source server certificate.
  - For third-party certificates, if the certificate authority (CA) certificates for the target server already appear in partner key databases, only the target certificates must be replaced. If the new certificates use a CA whose certificates are not in partner key databases, distribute the certificates to clients and other partner systems so that the new target certificates can be trusted.
  - When using self-signed certificates, distribute the newly generated target certificates to clients and other partner systems for inclusion in their key databases.

You can now upgrade the target server to any level that is supported on the target system. For more information, see the appropriate topic:

- [Upgrading from Tivoli Storage Manager Version 6.2 to 6.3 or later](#)
- [Upgrading from Tivoli Storage Manager V6.2 or V6.3 to V7.1.1](#)
- [Upgrading from V6.3 to V8.1](#)



## Appendix A. Utilities for server migration

Utilities are available to migrate a server from AIX, HP-UX, Oracle Solaris, or Linux on Power Systems (big endian) to Linux x86\_64.

### ***DSMSERV EXTRACTDB (Extract data from a server database)***

Use the DSMSERV EXTRACTDB utility to extract data from a V7.1.5 or later server database. The data can be inserted into a target server database either during or after the extraction process.

#### **Syntax**

```
>>-DSMSERV-----+-----+-----+-----EXTRACTDB----->
    '- -quiet-' '- -o -options_file-'

>---+| A: Extract to media |-----+-----><
    '-| B: Extract and transmit by using a network |-'
```

#### **A: Extract to media**

```
|--DEVclass---device_class_name--MANifest---file_name----->

    .-SCRatch---Yes-----.
>---+-----+-----+-----+----->
    '-SCRatch---+Yes--+'
        '-No--'

>---+-----+-----+-----+-----|
    |                               .-,-----|. |
    | (1)                           V           | |
    '|-----VOLumenames-----volume_name--+'
```

#### **B: Extract and transmit by using a network**

```
|--HLAddress---ip_address--LLAddress---tcp_port-----|
```

#### **Notes:**

1. You must specify VOLUMENAMES if you specify SCRATCH=NO.

#### **Parameters**

**-quiet**

Specifies that messages to the console are suppressed. This parameter is optional.

**-o options\_file**

Specifies an options file to use. This parameter is optional.

**DEVclass**

Specifies a sequential device class to use for storing the extracted data. This parameter is required if you want to extract the database to media.

**Restriction:** You cannot use a device class with a device type of NAS or CENTERA.

**MANifest**

Specifies the location of the manifest file. Use a fully qualified file name, or place the file in a local directory. For example: ./manifest.txt

This parameter is required if you want to extract the database to media.

**SCRatch**

Specifies whether scratch volumes can be used to store the data. This parameter is optional. The default value is YES. Possible values are:

Yes

Specifies that scratch volumes can be used.

If you specify SCRATCH=YES and also specify volumes with the VOLUMENAMES parameter, the specified volumes are used first. Scratch volumes are used only if the specified volumes do not have enough space to contain the extracted data.

If the device type for the device class is FILE, the names for the scratch volumes are generated based on a time stamp.

No

Specifies that scratch volumes cannot be used. If you specify this value, you must also specify volumes by using the VOLUMENAMES parameter.

If the volumes that you specify with the VOLUMENAMES parameter do not have enough space to contain the extracted data, the process fails.

**VOLumenames**

Specifies the volumes to use to store the extracted database. To specify multiple volumes, separate the names with commas and without intervening spaces.

**HLAddress**

Specifies either the numeric IP address or the domain name of the target server. This parameter is required if you want to extract and simultaneously transmit the data to the target server by using the network.

**LLAddress**

Specifies the low-level address of the target server. This address is the same as the value that is specified with the TCPPOINT server option for the target server. This parameter is required if you want to extract and simultaneously transmit the data to the target server by using the network.

**Examples**

Extract and transmit the data to a server by using the network:

```
dmserv extractdb hladdress=xyz.company.com lladdress=1555
```

Extract and store the data on media:

```
dsmserv extractdb devclass=tapeclass manifest=keepthis.txt
```

### The manifest file for data extraction to media

During data extraction to media, the DSMSERV EXTRACTDB utility creates a manifest file. The manifest file contains information about the volumes and device classes that are used for storing the extracted data.

The manifest file is required by the data insertion process when you are using media to move the database to the target server. You specify the name of the manifest file when you run the DSMSERV INSERTDB utility.

In the manifest file, the device configuration section contains information about the configuration settings of the source server device. You might need to edit settings, such as device names and drive element numbers, before you use the manifest file with the DSMSERV INSERTDB utility, especially if the utility is run on a different system.

A manifest file has content that is similar to the following example:

```
* Version 2 manifest file for EXTRACTDB
DEVCLASS=3584L2
STREAMS=1
VOLUMENAMES000=L12345
* Device Configuration Information
DEFINE DEVCLASS 3584L2 DEVTYPE=LTO FORMAT=DRIVE MOUNTLIMIT=DRIVES MOUNTWAIT=60
MOUNTRETENTION=60 PREFIX=ADSM LIBRARY=3584L2 WORM=NO DRIVEENCRYPTION=ALLOW
SET SERVERNAME MOKSHA
DEFINE LIBRARY 3584L2 LIBTYPE=SCSI SERIAL="00000002267304AB" SHARED=YES
AUTOLABEL=YES RESETDRIVE=YES
DEFINE DRIVE 3584L2 3584L2_DRV1 ELEMENT=265 ONLINE=Yes WWN="500507630F18BA09"
SERIAL="0007859130"
DEFINE PATH MOKSHA 3584L2 SRCTYPE=SERVER DESTTYPE=LIBRARY DEVICE=/dev/IBMchanger5
ONLINE=YES
DEFINE PATH MOKSHA 3584L2_DRV1 SRCTYPE=SERVER DESTTYPE=DRIVE LIBRARY=3584L2
DEVICE=/dev/IBMtape20 ONLINE=YES
```

### ***DSMSERV LOADFORMAT (Format a database)***

Use the DSMSERV LOADFORMAT utility to format an empty database in preparation for inserting an extracted database into the empty database.

#### **Syntax**

```
>>-DSMSERV -+-----+-----+-----+-----+----->
           |               | |               |
           '- -u--user_name-'   '- -i--instance_dir-'
```

```

>+-----+-----+-----+-----+-----+----->
'- -o--options_file-' '- -noexpire-' '- -quiet-'

                .-,-----.
                v          |
>--LOADFORMAT---DBDir---directory--+----->
                '-DBFile---file-----'

                .-ACTIVELOGSize---16384-----.
>+-----+-----+-----+-----+-----+----->
                '-ACTIVELOGSize---megabytes-'

>--ACTIVELOGDirectory---directory----->

>--ARCHLogdirectory---directory----->

>+-----+-----+-----+-----+-----+----->
                '-ARCHFailoverlogdirectory---directory-'

>+-----+-----+-----+-----+-----+-----><
                '-MIRRORlogdirectory---directory-'

```

## Parameters

### **-u user\_name**

Specifies a user name to switch to before you initialize the server. This parameter is optional.

### **-i instance\_dir**

Specifies an instance directory to use. This directory becomes the current working directory of the server. This parameter is optional.

### **-o options\_file**

Specifies an options file to use. This parameter is optional.

### **-noexpire**

Specifies that expiration processing is suppressed when the server starts. This parameter is optional.

### **-quiet**

Specifies that messages to the console are suppressed. This parameter is optional.

### **DBDir**

Specifies the relative path names of one or more directories that are used to store database objects. Directory names must be separated by commas but without spaces. You can specify up to 128 directory names. You must specify either the DBDIR or the DBFILE parameter.

### **DBFile**

Specifies the name of a file that contains the relative path names of one or more directories that are used to store database objects. Each directory name must be on a separate line in the file. You can specify up to 128 directory names. You must specify either the DBDIR or the DBFILE parameter.

#### **ACTIVELOGSize**

Specifies the size of the active log in megabytes. This parameter is optional. The minimum value is 2048 MB (2 GB); the maximum is 131,072 MB (128 GB). If you specify an odd number, the value is rounded up to the next even number. The default is 16384 MB.

#### **ACTIVELOGDirectory (Required)**

Specifies the directory in which the server writes and stores active log files. There is only one active log location. The name must be a fully qualified directory name. The directory must exist, it must be empty, and it must be accessible by the user ID of the database manager. The maximum number of characters is 175.

#### **ARCHLogdirectory (Required)**

Specifies the directory for the archive log files. The name must be a fully qualified directory name. The maximum number of characters is 175.

#### **ARCHFailoverlogdirectory**

Specifies the directory to be used as an alternative storage location if the ARCHLOGDIRECTORY directory is full. This parameter is optional. The maximum number of characters is 175.

#### **MIRRORlogdirectory**

Specifies the directory in which the server mirrors the active log (those files in the ACTIVELOGDIRECTORY directory). This parameter is optional. The directory must be a fully qualified directory name. The maximum number of characters is 175.

#### **Example: Format a database**

```
dsmserv loadformat dbdir=/tsmdb001 activelogsiz=8192
activelogdirectory=/activelog archlogdirectory=/archlog
archfailoverlogdirectory=/archfaillog mirrorlogdirectory=/mirrorlog
```

### ***DSMSERV INSERTDB (Move a server database into an empty database)***

Use the DSMSERV INSERTDB utility to move a server database into a new database. The database can be extracted from the source server and inserted into a new database on the new server by using a network connection between the two servers. The database can also be inserted from media that contains the extracted database.

Before you use the DSMSERV INSERTDB utility, complete the planning and preparation tasks, such as backing up the database and saving configuration information. Ensure that you meet all requirements before you move the server database.

#### **Requirements for insertion by using media**

Before you run the utility to insert the server database into an empty database, ensure that your system meets the following requirements.

- The manifest file from the DSMSERV EXTRACTDB operation must be available.
- The media that contains the extracted database must be available to the target server. Also, the permissions must be set to grant access to the media for the user ID that owns the target server instance.

## Syntax

```
>>-DSMSERV -+-----+--+-----+----->
          '- -u--user_name-' '- -i--instance_dir-'
```

```
>+-----+--+-----+--+-----+----->
          '- -o--options_file-' '- -noexpire-' '- -quiet-'
```

```
>--INSERTDB---+| A: Insert from media |----->
          '-| B: Insert over a network |-'
```

```
.-PREview---No-----.
```

```
>+-----+--+-----+-----><
          '-PREview---+Yes-+-'
          '-No--'
```

A: Insert from media

```
|+-----+--+-----+-----+--MANifest---file_name --|
          '-DEVclass---device_class_name-'
```

B: Insert over a network

```
.-SESSWait---60-----.
```

```
|+-----+--+-----+-----+-----|
          '-SESSWait---minutes-'
```

## Parameters

**-u user\_name**

Specifies a user name to switch to before you initialize the server. This parameter is optional.

**-i instance\_dir**

Specifies an instance directory to use. This directory becomes the current working directory of the server. This parameter is optional.

**-o options\_file**

Specifies an options file to use. This parameter is optional.

**-noexpire**

Specifies that expiration processing is suppressed when starting. This parameter is optional.

**-quiet**

Specifies that messages to the console are suppressed. This parameter is optional.

**DEVclass**

Specifies a sequential-access device class. You can specify any device class except for the DISK device class. The definition for the device class must exist in either the manifest file or the device configuration file.

This parameter is optional and is used only when the database that you want to insert into the empty database was extracted to media. If the database is on media and you do not specify a device class, the device class that is identified in the manifest file is used.

**Restriction:** You cannot use a device class with a device type of NAS or CENTERA.

**MANifest**

Specifies the location of the manifest file. Use a fully qualified file name, or place the file in a local directory. For example: ./manifest.txt

This parameter is required when the database that you want to insert into the empty database was extracted to media.

**SESSWait**

Specifies the number of minutes that the target server waits to be contacted by the source server. The default value is 60 minutes.

Use this parameter only if the data that is inserted into the empty target database is transmitted from the source server with a network connection.

**PREview**

Specifies whether to preview the insertion operation. This parameter is optional. The default value is NO.

Use the PREVIEW=YES parameter to test a database. When you use this parameter, the operation includes all steps of the process, except for the actual insertion of data into the new database. When you preview the insertion operation, you can quickly verify that the source database is readable. You can also identify any data constraint violations that might prevent an upgraded database from being put into production.

## Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US

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

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 jurisdictions 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 provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
US

Such information may be available, subject to appropriate terms and conditions, including in some cases,



payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises 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.

### **Trademarks**

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Linear Tape-Open and LTO are trademarks of HP, IBM Corp. and Quantum in the U.S. 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.