

Quest® NetVault® Backup Plug-in *for DB2 3.3*
User's Guide



© 2017 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, Join the Innovation, and NetVault are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend

- **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.

- ⚠ **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

- i **IMPORTANT NOTE, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

NetVault Backup Plug-in for DB2 User's Guide
Updated - February 2017
Software Version - 3.3
DBG-101-3.3-EN-02

Contents

Introducing NetVault Backup Plug-in for DB2	5
NetVault Backup Plug-in for DB2: at a glance	5
Key benefits	5
Feature summary	6
Target audience	6
Recommended additional reading	6
Installing and removing the plug-in	7
Installing or upgrading the plug-in	7
Removing the plug-in	8
Configuring the plug-in	9
Configuring default settings	9
Configuring a default login account	9
Enabling Archive Logging for the database	10
Enabling the TRACKMOD parameter	10
Additional notes	11
Backing up data	12
Available backup modes	12
Available backup types	12
Defining a backup strategy	13
Performing backups	14
Selecting data for a backup	14
Setting backup options	15
Finalizing and submitting the job	16
Restoring data	17
Restore and recovery overview	17
Restoring from Full Backups	17
Selecting data for a restore	17
Setting restore options	18
Finalizing and submitting the job	20
Restoring from Incremental Backups	21
Restoring from Delta Backups	22
Advanced restore procedures	23
Renaming or relocating a database during a restore	23
Restoring a database to an alternate server	24
Working with dropped DB2 tables	25
Troubleshooting	29
About us	31
We are more than just a name	31

Our brand, our vision. Together.	31
Contacting Quest	31
Technical support resources	31

Introducing NetVault Backup Plug-in *for DB2*

- [NetVault Backup Plug-in for DB2: at a glance](#)
- [Key benefits](#)
- [Feature summary](#)
- [Target audience](#)
- [Recommended additional reading](#)

NetVault Backup Plug-in *for DB2*: at a glance

The NetVault Backup Plug-in *for DB2* (Plug-in *for DB2*) increases confidence in the recoverability of DB2 databases and eliminates the need for complex scripting. Through a web-based user interface (WebUI) and automated workflow process, the plug-in offers a centralized way to set up, configure, and define backup and restore policies for your DB2 databases. Support for online and offline backups gives you the flexibility to select the preferred backup method without requiring you to learn the DB2 internals. The plug-in offers a detailed level of control that minimizes downtime by allowing you to restore complete databases, individual tablespaces, or individual log files. Through integration with a range of backup devices, your data is protected and stored offsite to meet your disaster-recovery and business-continuity goals.

Key benefits

- **Reduce risk with flexible backup and recovery options:** Plug-in *for DB2* ensures that databases remain online and fully accessible during backup operations, assuring no user downtime. You can implement a backup policy that is flexible enough to account for many recovery scenarios without the need to understand DB2 database internals. In addition to providing options for Full, Incremental, and Delta Backups, the plug-in supports detailed recoveries. This option lets you restore complete databases, individual tablespaces, or individual archived logs up to the current time or the last backup. Restores are faster because the reliance on human interaction is reduced and the chance of syntax errors that would require a restart is eliminated to reduce downtime further.
- **Maximize information technology (IT) efficiency:** NetVault Backup's ability to manage distributed storage resources across multiple databases lets you focus on business-critical tasks rather than repetitive and manual backup operations. Less-experienced personnel can perform routine administrative options, thus empowering administrators to create comprehensive backup policies without the need to understand DB2 internals. The plug-in uses an automated workflow with point-and-click options for tasks such as defining backups and scheduling jobs.
- **Ensure business continuity:** With offsite backups being an important part of the data-protection plan for business-critical applications, the plug-in takes advantage of NetVault Backup's integration with a range of backup devices. NetVault Backup offers the flexibility to store backups on popular backup devices. You can

store the backup online in a virtual tape library (VTL). You can also duplicate the job to physical tape libraries shared by multiple DB2 databases, other proprietary databases, or even general backup files.

Feature summary

- Online and offline backup modes
- Full, Incremental, and Delta Backups
- Protection down to tablespace level
- Restore complete databases, individual tablespaces, or individual archived logs
- Recover up to the current point-in-time (PIT) or last backup
- Point-and-click WebUI
- Job scheduling
- Rename and relocate databases during restores
- Restore to an alternate server

Target audience

Advanced database administrator (DBA) skills are not required for routine backup operations. However, initial configuration and recovery operations might require database administration experience.

Recommended additional reading

The following documentation is also available:

- *Quest NetVault Backup Installation Guide*: This guide provides details on installing the NetVault Backup Server and Client software.
- *Quest NetVault Backup Administrator's Guide*: This guide explains how to use NetVault Backup and describes the functionality common to all plug-ins.
- *Quest NetVault Backup CLI Reference Guide*: This guide provides a description of the command-line utilities.

You can download these guides from <https://support.quest.com/technical-documents>.

i **IMPORTANT:** Starting with 10.0.0, NetVault Backup provides a WebUI to configure, manage, and monitor your NetVault Backup system and installed plug-ins. The procedures described in the user's guide for this version of the plug-in are intended for the new WebUI. For procedures based on the NetVault Backup Console (user interface available with NetVault Backup 9.x), see the documentation for an earlier version of the plug-in.

Installing and removing the plug-in

- [Installing or upgrading the plug-in](#)
- [Removing the plug-in](#)

Installing or upgrading the plug-in

- 1 Access the **NetVault Configuration Wizard** or **Manage Clients** page.

i | **NOTE:** If the selected clients are all the same type, you can use the configuration wizard to install the plug-in on multiple clients at the same time. When you select multiple clients, verify that the plug-in binary file is compatible with the OS and platforms of the target clients. From the **Manage Clients** page, you can only select one client for plug-in installation.

- To access the **NetVault Configuration Wizard** page:
 - a In the Navigation pane, click **Guided Configuration**.
 - b On the **NetVault Configuration Wizard** page, click **Install Plugins**.
 - c On the next page, select the applicable clients.
 - To access the **Manage Clients** page:
 - a In the Navigation pane, click **Manage Clients**.
 - b On the **Manage Clients** page, select the machine that contains the DB2 database, and click **Manage**.
 - c On the **View Client** page, click the **Install Plugin** button (🔍).
- 2 Click **Choose Plug-in File**, navigate to the location of the “.npg” installation file for the plug-in, for example, on the installation CD or the directory to which the file was downloaded from the website.
Based on the operating system (OS) in use, the path for this software may vary on the installation CD.
 - 3 Select the file entitled “**db2-x-x-x-x.npg**,” where **xxxxx** represents the version number and platform, and click **Open**.
 - 4 To begin installation, click **Install Plugin**.

After the plug-in is successfully installed, a message is displayed.

i | **NOTE:** The 64-bit version of Plug-in *for DB2* does not support both 32-bit and 64-bit database instances on the same DB2 Server. To solve this problem, upgrade or migrate the existing 32-bit database instances to 64-bit.

Removing the plug-in

- 1 In the Navigation pane, click **Manage Clients**.
- 2 On the **Manage Clients** page, select the applicable client, and click **Manage**.
- 3 In the **Installed Software** table on the **View Client** page, select **Plug-in for DB2**, and click the **Remove Plugin** button (🗑️).
- 4 In the **Confirm** dialog box, click **Remove**.

Configuring the plug-in

- [Configuring default settings](#)
- [Configuring a default login account](#)
- [Enabling Archive Logging for the database](#)
- [Enabling the TRACKMOD parameter](#)
- [Additional notes](#)

Configuring default settings

The plug-in lets you set default options for backup and restore jobs. You can override these options on a per-job basis.

- 1 In the Navigation Pane, click **Create Backup Job**, and click **Create New** next to the **Selections** list.
- 2 In the selection tree, open the applicable client node, and select **Plug-in for DB2**.
- 3 From the **Actions** list, select **Configure**.
- 4 Enter the **DB2 Installation Directory** parameter.
If DB2 is installed in its default directory, the plug-in automatically completes this field with the corresponding path. If the software is installed in a different location, enter the correct path.
- 5 To save the settings, click **OK**.

Configuring a default login account

Plug-in *for DB2* lets you configure a default login account, which is used to connect automatically to the database instance. Configuring a default login account eliminates the need to enter the authentication details each time a connection to the instance is required. The plug-in stores the user name and the encrypted password in the configuration file. When you configure a default account, any change in the user account or password does not affect the existing job definitions.

To configure a default login account for a DB2 instance, perform the following steps.

- 1 In the Navigation pane of the NetVault Backup WebUI on the NetVault Backup Server, click **Create Backup Job**.
- 2 Next to the **Selections** list, click **Create New**.
- 3 In the selection tree, open the NetVault Backup Client on which the plug-in is installed.
- 4 In the list of plug-ins, open **Plug-in for DB2** to display the added DB2 instances, and select the applicable instance.
- 5 From the **Actions** list, select **Set Default Login**.
- 6 In the **DB2 Default Login** dialog box, configure the following details:
 - **Default Account Name:** Enter the account name for the DB2 Instance Administrator.

- **Default Password:** Enter the password for the configured user account.
- **Use Default Login:** To connect automatically to the database without waiting for user input, select this check box.

7 To save the settings, click **OK**.

Enabling Archive Logging for the database

The Circular Logging method, which is the default logging behavior for any new DB2 database, only supports Full, Offline Backups. To perform Online Backups, a database must be configured for Archive Logging. With this type of logging, the log files are retained and made available for roll-forward recovery. To enable Archive Logging, you can set either the LOGRETAIN, available only with DB2 9.1 and earlier, or the LOGARCHMETH1, available only with DB2 9.5 and later, configuration parameter.

- 1 Start a terminal session on the DB2 Server, and navigate to the "...**SQLLIB\BIN**" directory, where "..." indicates the installation path of DB2.
- 2 To access the DB2 command prompt, type:


```
DB2
```
- 3 Run the applicable command:
 - To set the LOGRETAIN configuration parameter, DB2 9.1 and earlier, to Recovery mode, type:


```
update db cfg for <databaseName> using logretain ON
```

In the Recovery mode, the active log files are retained and become online archive log files for use in roll-forward recovery. This process is called log retention logging.
 - To set the LOGARCHMETH1 configuration parameter, DB2 9.5 and later, to LOGRETAIN, type:


```
update db cfg for <databaseName> using logarchmeth1 LOGRETAIN
```

This command is equivalent to setting the LOGRETAIN configuration parameter to Recovery mode. It stores the Archived Logs in the same location as the Active Logs.
 - To set the LOGARCHMETH1 configuration parameter to DISK, type:


```
update db cfg for <databaseName> using logarchmeth1 DISK:<completePath>
```

This command lets you specify the location where the Archived Logs are stored.

Enabling the TRACKMOD parameter

The TRACKMOD configuration parameter enables tracking and recording of the database page updates. Set this parameter to YES to allow Incremental and Delta Backups. For system-managed space (SMS) and database-managed space (DMS) tablespaces, the granularity of tracking is at the tablespace level. A flag indicates whether a tablespace contains any pages that must be included in an Incremental or a Delta Backup. If there are no pages, then the backup operation can skip that tablespace.

- 1 Start a terminal session on the DB2 Server, and navigate to the "...**SQLLIB\BIN**" directory, where "..." indicates the DB2 installation path.
- 2 To access the DB2 command prompt, type:


```
DB2
```
- 3 Type the following command:


```
update db cfg for <databaseName> using trackmod ON
```

- NOTE:** After you enable the TRACKMOD parameter, run a Full Backup first before performing any Incremental or Delta Backups.

If you create a tablespace after enabling this parameter, perform a backup that includes the new tablespace. You can perform either a database-level or a tablespace-level backup.

Additional notes

- To prevent users from accidentally dropping a table, you can add the **Restrict on Add** attribute as shown in the following examples:

```
create table <tableName> (<tableSpecificInformation>) in  
  <userSpace> with restrict on drop
```

– or –

```
alter table <tableName> add restrict on drop
```

- To remove the **Restrict on Add**, you can use the following command:

```
alter table <tableName> drop restrict on drop
```

Backing up data

- [Available backup modes](#)
- [Available backup types](#)
- [Defining a backup strategy](#)
- [Performing backups](#)

Available backup modes

Plug-in *for DB2* supports the following backup modes:

- **Offline Backups:** Offline backups require exclusive connection to the database. Other users and applications cannot connect to the database while the backup is in progress.
- **Online Backups:** Online backups allow other users and applications to connect to the database, and read and modify the data while the backup is in progress. For online backups, a database must be configured for Archive Logging. For details, see [Enabling Archive Logging for the database](#). During online backups, all database transactions are recorded in the logs. When the backup completes, the database manager closes the active log and archives it. After restoring an online backup, the logs must be rolled forward to recover the database to a consistent state.

Available backup types

Plug-in *for DB2* supports the following backup types:

- **Full Backup:** A Full Backup includes all the data in the specified database or tablespace. It serves as the base for the subsequent Incremental and Delta Backups.
- **Incremental Backup:** An Incremental Backup includes any data that has changed or is new since the last Full Backup. Incremental Backups duplicate the data backed up by the previous Incremental or Delta Backups. These backups are also referred to Cumulative Backups.
- **Delta Backup:** A Delta Backup, or Incremental Delta Backup, includes any data that is new or has changed since the last Full, Incremental, or Delta Backup of the specified tablespace. Delta Backups consume minimum storage space and are quicker to perform. These backups are also referred to Non-Cumulative Backups.

Defining a backup strategy

A backup strategy helps you recover from the damages caused by failure, and resume normal operations as quickly as possible. Before backing up data, ensure that a strategy is in place that safeguards data against various failure modes, such as media failure, data corruption, user error, and loss of a data center.

The following list provides guidelines on some of the backup sequences you might consider using.

- **Full Backups Only:** When requirements mandate data protection up to the previous day, performing Full Backups daily should be sufficient. If a failure occurs, the plug-in only requires that you restore a single saveset.
- **Full Backup and Archived Logs Backup:** When requirements mandate data protection up to the current time, but recovery time is not important, Full Backup and Archived Logs Backup should provide the necessary protection. For example, Full Backups can be performed every Sunday, while backups of Archived Logs can be performed Monday through Saturday.

This strategy requires the plug-in to restore and apply each Archived Logs Backup between the last Full Backup and the point-of-failure in succession. This strategy can lead to longer recovery as the week progresses. To speed up restores, you can include Incremental or Delta Backups, which reduces the number of Archived Logs Backup that must be restored.

- **Full Backup and Incremental Backup and Archived Logs Backup:** When requirements mandate data protection up to the current time, and recovery must be as fast as possible, Full Backup and Incremental Backup and Archived Logs Backup provides the best combination. For example, Full Backups can be performed every Sunday, while Incremental Backups can be performed Monday through Saturday, and backups of Archived Logs can be performed every 5 to 6 hours. Each Incremental Backup includes all changes since the last Full Backup, and the Archived Logs Backup includes the transaction logs since the last Incremental Backup.

This strategy speeds recovery because the plug-in only has to restore the Full Backup, the last Incremental Backup, and the backups of Archived Logs created between the last Incremental Backup and the point-of-failure.

- **Full Backup and Delta Backup and Archived Logs Backup:** When requirements mandate data protection up to the current time, and backups have to be as fast as possible, then Full Backup and Delta Backup and Archived Logs Backup is an ideal strategy. For example, Full Backups can be performed every Sunday, while Delta Backups can be performed Monday through Saturday, and backups of Archived Logs Backups can be performed every 5 to 6 hours. Each Delta backup includes all the changes since the last Full or Delta Backup, and the Archived Log Backup includes the transaction logs since the last Delta Backup.

This strategy slows recovery because the plug-in must restore the Full Backup, each Delta Backup since the Full Backup, and the backups of Archived Logs created between the last Delta Backup and the point-of-failure.

Performing backups

The procedure for performing backups using the plug-in includes the steps outlined in the following topics:

- [Selecting data for a backup](#)
- [Setting backup options](#)
- [Finalizing and submitting the job](#)

Selecting data for a backup

You must use sets—Backup Selection Set, Backup Options Set, Schedule Set, Target Set, and Advanced Options Set—to create a backup job.

Backup Selection Sets are essential for Incremental and Delta Backups. Create the Backup Selection Set during a Full Backup, and use it for Full, Incremental, and Delta Backups. The backup job reports an error if you do not use a Selection Set for the Incremental or Delta Backup. For more information, see the *Quest NetVault Backup Administrator's Guide*.

i | TIP: To use an existing set, click **Create Backup Job**, and select the set from the **Selections** list.

- 1 In the Navigation pane, click **Create Backup Job**.

You can also start the wizard from the Guided Configuration link. In the Navigation pane, click **Guided Configuration**. On the **NetVault Configuration Wizard** page, click **Create backup jobs**.

- 2 In **Job Name**, specify a name for the job.

Assign a descriptive name that lets you easily identify the job when monitoring its progress or restoring data. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Linux, the name can have a maximum of 200 characters. On Windows, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

- 3 Next to the **Selections** list, click **Create New**.

- 4 In the list of plug-ins, open **Plug-in for DB2** to display the added DB2 instances.

- 5 Open the target database instance.

- 6 If a default login account has not been configured for the selected instance, enter the following details in the dialog box that is displayed, and then click **Ok** to close the dialog box:

- **Account Name:** Enter the account name for the DB2 Instance Administrator.
- **Password:** Enter the password for the configured user account.

i | NOTE: This step is not applicable if you have configured a default login account for the instance. The plug-in uses the details stored in the configuration file to connect automatically to the instance.

- 7 When the instance opens, select the items that you want to back up:

- **Full Database Backup:** Select the database node.
- **Table space Backup:** Select one of the following:
 - To back up all tablespaces, open the database node, and select the **Tablespaces** node.
 - To back up individual tablespaces, open the **Tablespaces** node, and select one or more tablespaces.
- **Archived Logs:** To back up the Archived Logs, select the **Archived Logs** node. Archived logs are required for a roll-forward recovery, and must be backed up regularly. You can back up the archived logs separately or together with the tablespaces. The plug-in provides the option to delete the archived logs from the disk after backing them up.

- 8 Click **Save**, enter a name in the **Create New Set** dialog box, and click **Save**.

The name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Linux, the name can have a maximum of 200 characters. On Windows, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

Setting backup options

The next step involves creating the Backup Options Set or selecting an existing one.

i | TIP: To use an existing set, in the **Plugin Options** list, select the set that you want to use.

1 Next to the **Plugin Options** list, click **Create New**.

2 Configure the following options

- **Backup Mode:** Select one of the following options: **Online** or **Offline**. For details, see [Available backup modes](#).
- **Backup Type:** Select one of the following options: **Full**, **Incremental**, or **Delta**. For details, see [Available backup types](#).
- **DB2 Buffers Setting:** Configure the following parameters:
 - **Number of Buffers:** If you are using multiple buffers and I/O channels, specify at least twice as many buffers as channels to ensure that the channels do not have to wait for data. Increasing the number of buffers results in the backup operation completing faster.
 - **Buffer Size:** To change the buffer size, enter the value here. The value must be a multiple of the tablespace extent size. The default value for this parameter is 1024 KB.

i | NOTE: Increasing the value specified for either option might cause the restore operation to use more memory and more processing time.

- **Number of Sessions:** This parameter determines the number of tablespaces that are backed up in parallel. Each session is assigned one tablespace at a time until all tablespaces are backed up.

The default value for this parameter is 1. For backup jobs created with plug-in v3.0 or any earlier versions, it is implicitly set to 1.

If you set this parameter to a value larger than the number of tablespaces being backed up, the plug-in starts the requested number of sessions and each session creates a backup stream. However, some of the backup streams contain just a header and a trailer generated by DB2, and no tablespace data. This issue results in inefficient backups and uses more media than necessary.

If some of the selected tablespaces are small while others are large, reduce the number of sessions accordingly. For example, to back up two small and two large databases, you can set this parameter to 2.

i | IMPORTANT: If you are creating an Incremental or Delta Backup job, set this field to **1**. If you enter a different number, the plug-in limits the number to **1**, and displays the following message upon completion with warnings: **Backup Option 'Number of Sessions' = '<configuredNumber>' changed to '1' for 'INCREMENTAL' or 'DELTA' backup**

- **Purge Logs:** Under **Purge Logs**, configure the following parameters
 - **Purge Archived Logs:** To remove the archived logs automatically after they are backed up, select this check box. The plug-in deletes all archived logs except the Active Logs. Also, the plug-in does not delete the number of logs specified in the **Logs to Keep Before Active Log** option.
 - **Logs to Keep before Active Log Option:** To specify the number of logs preceding the current active database log in the server that should *not* be purged, use this option with the **Purge Archived Logs** option. Enter a number from 1 to 999.

3 Click **Save** to save the set.

- 4 In the **Create New Set** dialog box, specify a name for the set, and click **Save**.

The name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Linux, the name can have a maximum of 200 characters. On Windows, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

Finalizing and submitting the job

The final steps include setting additional options on the Schedule, Target Storage, and Advanced Options pages, submitting the job, and monitoring the progress through the Job Status and View Logs pages. These pages and options are common to all NetVault Backup Plug-ins. For more information, see the *Quest NetVault Backup Administrator's Guide*.

- 1 Use the **Schedule**, **Target Storage**, and **Advanced Options** lists to configure any additional required options.
- 2 Click **Save** or **Save & Submit**, whichever is applicable.

i | **TIP:** To run a job that you have already created and saved, select **Manage Job Definitions** in the Navigation pane, select the applicable job, and click **Run Now**.

You can monitor progress on the **Job Status** page and view the logs on the **View Logs** page. For more information, see the *Quest NetVault Backup Administrator's Guide*.

Restoring data

- [Restore and recovery overview](#)
- [Restoring from Full Backups](#)
- [Restoring from Incremental Backups](#)
- [Restoring from Delta Backups](#)
- [Advanced restore procedures](#)

Restore and recovery overview

Restore recreates a full database or a subset of tablespaces from a backup copy. A database restore operation requires an exclusive connection, but a tablespace restore operation can be completed online. To restore the tablespaces, you can use either a tablespace-level backup or select the required tablespaces from a full database backup. Offline backups allow you to restore a database to the same state that it was in at the time of backup. With online backups, you can use the Archived Logs and Active Logs to recover the database to the current time or last backup.

Restoring from Full Backups

The procedure for restoring a Full Backup includes the steps outlined in the following topics:

- [Selecting data for a restore](#)
- [Setting restore options](#)
- [Finalizing and submitting the job](#)

Selecting data for a restore

- 1 In the Navigation pane, click **Create Restore Job**.
- 2 On the **Create Restore Job—Choose Saveset** page, select **Plug-in for DB2** from the **Plugin Type** list.
- 3 To filter the items displayed in the saveset table further, use the **Client**, **Date**, and **Job ID** lists.
The table displays the saveset name (job title and saveset ID), creation date and time, and size. By default, the list is sorted alphabetically by saveset name.
- 4 In the saveset table, select the applicable item.
When you select a saveset, the following details are displayed in the **Saveset Information** area: Job ID, job title, server name, client name, plug-in name, saveset date and time, retirement setting, Incremental Backup or not, Archive or not, saveset size, and snapshot-based backup or not.
- 5 Click **Next**.

- 6 On the **Create Selection Set** page, select the data that you want to restore.
 - **Restore Full Database:** Select the database node.
 - **Restore Table spaces:** To restore all tablespaces, open the database node, and select the **Tablespaces** node. To restore individual tablespaces, open the **Tablespaces** node, and select one or more tablespaces.
 - **Apply Archived Logs:** Select the **Archived Logs** node, or open this node, and select the required log files.

Setting restore options

This step involves the steps outlined in the following topics:

- [Setting restore type and other restore options](#)
- [Configuring Destination Details](#)
- [Viewing backup information](#)

Setting restore type and other restore options

On the **Create Selection Set** page, click **Edit Plugin Options**, and configure the following parameters on the **Options** tab:

- **Restore Type:** Select one of the following:
 - **Restore all table spaces in the database:** To restore all tablespaces in a selected database, select this option.
 - **Restore only table space-level backup images:** When performing a restore of one or more selected tablespaces, select this option.
- **Restore Mode:** Select one of the following:
 - **Offline:** To prevent other applications and users from connecting to the database during a restore, select this option.
 - **Online:** This option is applicable only when performing a tablespace-level restore. It allows other users and applications to connect to the database, and read or modify the data in other tablespaces while the specified tablespaces are being restored.
- **Incremental Restore:** To alert NetVault Backup that a Full Backup being recovered is a portion of an incremental sequence, rather than a standalone Full Backup recovery, select this option. For more information, see [Restoring from Incremental Backups](#) and [Advanced restore procedures](#).
- **Archived Logs Options:** Select the applicable options:
 - **Rollforward Archived Logs:** To apply the Archived Logs, select this check box.
 - **Soft Recovery:** To restore the existing logs to their most up-to-date state, select this option.
 - **Hard Recovery:** To restore the logs to the point at which the backup was performed, thus overwriting the existing archived logs files, select this option.
 - **Delete existing logs:** To delete all existing log files, select this check box.
 - **Rollforward Only (will not Restore any Files):** To perform only a roll-forward recovery without restoring any files, select this check box.
- **DB2 Buffers Setting:** Configure the following parameters:
 - **Number of Buffers:** If you are using multiple buffers and I/O channels, specify at least twice as many buffers as channels to ensure that the channels do not have to wait for data. Increasing the number of buffers results in the backup operation completing faster.

- **Buffer Size:** To change the buffer size, enter the value here. The value must be a multiple of the tablespace extent size. The default value for this parameter is 1024 KB.
 - **NOTE:** Increasing the value specified for either option might cause the restore operation to use more memory and more processing time.
- **Number of Sessions:** This parameter determines the number of tablespaces that are restored in parallel. Each session is assigned one tablespace at a time until all tablespaces are restored. This parameter is read-only. It is set to the number of sessions used to create the selected backup.

Configuring Destination Details

To restore the data to a different DB2 instance, click the **Destination Details** tab, and configure the following parameters:

- **DB2 Instance Name:** Enter the name of the instance to which you want to restore the data.
- **DB2 Instance Admin Account Name:** Enter the account name for the DB2 Instance Administrator.
- **DB2 Instance Admin Password:** Enter the password for the configured user account.

Viewing backup information

To view the contents of various backup pieces in the selected backup set, click the **Backup Information** tab. This tab provides the following details:

- **Backup done with:** The plug-in version number
- **Backup Type:** Full, Incremental, or Delta
- **Backup Mode:** Online or Offline
- **DB2 Version**
- **Backup Includes Tables from the following Databases:** This table includes the following columns:
 - **Database Name**
 - **Database Alias**
 - **DB2 Instance**
 - **DB2 Version ID**
 - **Backup Timestamp**
- **Backup Includes Archived Logs from the following Databases:** This table includes the following columns:
 - **Database Alias**
 - **DB2 Instance**
 - **Archive Log**
 - **Archived Log Path**

Finalizing and submitting the job

The final steps include setting additional options on the Schedule, Source Options, and Advanced Options pages, submitting the job, and monitoring the progress through the Job Status and View Logs pages. These pages and options are common to all NetVault Backup Plug-ins. For more information, see the *Quest NetVault Backup Administrator's Guide*.

1 To save the settings, click **Ok**, and then click **Next**.

2 In **Job Name**, specify a name for the job if you do not want to use the default setting.

Assign a descriptive name that lets you easily identify the job when monitoring its progress. The job name can contain alphanumeric and nonalphanumeric characters, but it cannot contain non-Latin characters. On Linux, the name can have a maximum of 200 characters. On Windows, there is no length restriction. However, a maximum of 40 characters is recommended on all platforms.

3 In the **Target Client** list, select the machine on which you want to restore the data.

i | **TIP:** You can also click **Choose**, and then locate and select the applicable client in the **Choose the Target Client** dialog box.


4 Use the **Schedule**, **Source Options**, and **Advanced Options** lists to configure any additional required options.

5 Click **Save** or **Save & Submit**, whichever is applicable.

You can monitor progress on the **Job Status** page and view the logs on the **View Logs** page. For more information, see the *Quest NetVault Backup Administrator's Guide*.

Restoring from Incremental Backups

Figure 1. Example of steps necessary to perform recovery using Incremental Backups



- 1 Restore the most recent Incremental Backup.

Use the following table to select the correct items on the **Create Selection Set** page and the **Options** tab.

Create Selection Set page		Options tab	
Selected	Cleared	Selected	Cleared
Tablespaces only	None necessary	Incremental Restore	<ul style="list-style-type: none"> • Rollforward Archived Logs • Delete Existing Logs

- 2 Restore the last Full Backup.

Use the following table to select the correct items on the **Create Selection Set** page and the **Options** tab.

Create Selection Set page		Options tab	
Selected	Cleared	Selected	Cleared
Tablespaces only	Archived Logs	Incremental Restore	<ul style="list-style-type: none"> • Rollforward Archived Logs • Delete Existing Logs


- 3 Restore the most recent Incremental Backup a second time.

Use the following table to select the correct items on the **Create Selection Set** page and the **Options** tab. In this step, the plug-in applies the Archived Logs to perform a Soft or Hard Recovery.

Create Selection Set page		Options tab	
Selected	Cleared	Selected	Cleared
<ul style="list-style-type: none"> • Tablespaces only • Archived Logs 	None necessary	<ul style="list-style-type: none"> • Incremental Restore • Rollforward Archived Logs • Soft Recovery or Hard Recovery, as applicable • Delete Existing Logs 	None necessary

Restoring from Delta Backups

Figure 2. Example of steps necessary to perform recovery using Delta Backups



- 1 Restore the last Delta Backup.
- 2 Restore the last Full Backup.
- 3 Restore the first Delta Backup.
- 4 Restore the subsequent Delta Backups.

For the preceding steps, use the following table to select the correct items on the **Create Selection Set** page and the **Options** tab.

Create Selection Set page		Options tab	
Selected	Cleared	Selected	Cleared
Tablespaces only	<ul style="list-style-type: none"> • Full Database • Archived Logs 	Incremental Restore	<ul style="list-style-type: none"> • Rollforward Archived Logs • Delete Existing Logs

- 5 Restore the last Delta Backup a second time.

Use the following table to select the correct items on the **Create Selection Set** page and the **Options** tab. In this step, the plug-in applies the Archived Logs to perform a Soft or Hard Recovery.

Create Selection Set page		Options tab	
Selected	Cleared	Selected	Cleared
<ul style="list-style-type: none"> • Tablespaces only • Archived Logs 	None necessary	<ul style="list-style-type: none"> • Incremental Restore • Rollforward Archived Logs • Soft Recovery or Hard Recovery, as applicable • Delete Existing Logs 	None necessary

i **NOTE:** You can mix the different types of Incremental Backups, for example, restoring both Delta and Incremental Backups, in a recovery scheme to achieve the required result. However, if you use a combination, the restore order and the data items included in each phase are crucial to achieving the wanted result.

Advanced restore procedures

This topic describes the steps outlined in the following topics:

- [Renaming or relocating a database during a restore](#)
- [Restoring a database to an alternate server](#)
- [Working with dropped DB2 tables](#)

Renaming or relocating a database during a restore

During a restore, you can rename a database or relocate all or individual tablespaces to a specified directory. This option lets you create a copy of the database instead of overwriting the existing version.

Prerequisites (Windows only)

On Windows, the following requirements must be met before you attempt to relocate the tablespaces:

- Ensure that the required path exists on the client. If it does not, create the path.
- Ensure that the `DB2_CREATE_DB_ON_PATH` parameter is enabled.

Note the following:

- If the plug-in cannot find the specified path, the restore job fails with the following message in the NetVault Backup Logs:

```
SQL1052N The database path <path> does not exist.
```

- The same message is displayed in the NetVault Backup Logs if you do not enable the `DB2_CREATE_DB_ON_PATH` parameter. Also, the following message is displayed in the DB2 Instance diagnostic file, `db2diag.log`:

```
An attempt is made to create a database on path without first enabling this feature using DB2_CREATE_DB_ON_PATH registry variable. Please enable it first, then retry again.
```

In DB2 9.x and later, the `DB2_CREATE_DB_ON_PATH` parameter lets you create databases under a path. In pre-DB2 9.x versions, only the root directory was supported. For example:

- Create Database statement in pre-9.x DB2 versions:

```
db2 create database mydb on c
```

- Create Database statement in 9.x and later versions:

```
db2 create database mydb on c:\mydb
```

To set the `DB2_CREATE_DB_ON_PATH` parameter from the command prompt, perform the following steps.

- 1 At the prompt, type:

```
db2set DB2_CREATE_DB_ON_PATH=YES
```

- 2 Restart the DB2 instance:

```
db2stop  
db2start
```

Restore procedure

To rename or relocate data, or both, during a restore, perform the following steps.

- 1 In the Navigation pane, click **Create Restore Job**, select **Plug-in for DB2** from the **Plugin Type** list, select the applicable saveset, and click **Next**.

For more information, see [Selecting data for a restore](#).

- 2 To rename a database, do the following:
 - a Select the database on the **Create Selection Set** page.
 - b From the **Actions** list, select **Rename**.
 - c In the **Rename/Relocate** dialog box, enter the new name in the **Rename** box, and click **Ok**.

The plug-in updates the database node to display the new name in parentheses.

i | **NOTE:** The plug-in does not let you rename individual tablespaces.

- 3 To relocate the data, do the following:
 - To relocate all tablespaces to the same directory, select the **Tablespaces** node. From the **Actions** list, select **Rename**. In the **Rename/Relocate** dialog box, enter the new location in the **Relocate** box, and click **Ok**.
 - To relocate individual tablespaces, open the **Tablespaces** node, and select the applicable tablespace. From the **Actions** list, select **Rename**. In the **Rename/Relocate** dialog box, enter the new location in the **Relocate** box, and click **Ok**.

i | **IMPORTANT:** When you relocate the individual tablespaces, also specify a path for the parent **Tablespaces** node. This directory is used to restore the individual tablespaces and Archived Log files for which you have not specified any relocation path.

The plug-in updates the tablespace node to display the new location in parentheses.

i | **NOTE:** You cannot relocate tablespaces that are managed by automatic storage. This issue is a functional limitation of DB2.

- 4 Configure the applicable restore options.

For more information, see [Setting restore options](#).
- 5 Complete and submit the job.

For more information, see [Finalizing and submitting the job](#).

Restoring a database to an alternate server

During a restore, you can also relocate a database to an alternate DB2 Server, which can be useful while migrating a server or performing a disaster recovery operation.

Before you start the restore procedure, review the following notes and set up the target system:

- **Install the DB2 Server software:** Install the DB2 Server software on the target system. The software version must be the same as that running on the original server.
- **Install NetVault Backup and the Plug-in for DB2:** Install the NetVault Backup Server or Client software and the plug-in on the target system. The software version must be the same as the version on the original server.
- **Add the NetVault Backup Client to the NetVault Backup Server:** If the target system is a NetVault Backup Client, add it to the NetVault Backup Server. For details, see the *Quest NetVault Backup Administrator's Guide*.

- **Add the DB2 Instance to the plug-in:** In the **Configure** dialog box, specify the installation path. For more information, see [Configuring default settings](#).

To restore a database to an alternate server, perform the following steps.

- 1 Start the NetVault Backup WebUI on the NetVault Backup Server.
- 2 In the Navigation pane, click **Create Restore Job**, select **Plug-in for DB2** from the **Plugin Type** list, select the applicable saveset, and click **Next**.
- 3 Open the saveset that contains that database backup, and select the database.
For more information, see [Selecting data for a restore](#).
- 4 Configure the applicable restore options, and click **Next**.
For more information, see [Setting restore options](#).
- 5 In the **Target Client** list, select the target DB2 Server.
- 6 Complete and submit the job.
For more information, see [Finalizing and submitting the job](#).

Working with dropped DB2 tables

A standard restore of a DB2 backup using Plug-in *for DB2* does not recover any tables that have been dropped. To restore a dropped table, use the procedures described in this topic.

- [Setting the Dropped Table Recovery option](#)
- [Recovery method 1: soft recovery](#)
- [Recovery method 2: hard recovery](#)

Setting the Dropped Table Recovery option

To recover a dropped table, the **Dropped Table Recovery** option must be set to “Y” for the tablespace that contained the dropped table. For tablespaces created using the DB2 command prompt, this option is enabled by default. For tablespaces created using the DB2 Control Center graphical user interface (GUI), the Create Tablespace wizard lets you enable this option.

To verify the status of this option and to enable it if it has not been enabled, perform the following steps.

- 1 Start a terminal session on the DB2 Server, and navigate to the “**...SQLLIB\BIN**” directory, where “...” indicates the installation path of DB2.
- 2 To access the DB2 command prompt, type:

```
DB2
```
- 3 To connect to the local instance of the target database, type:

```
connect to <localDatabaseName>
```
- 4 To determine the status of Dropped Table Recovery option for the tablespaces in the database, type:

```
select tbspace, drop_recovery from syscat.tablespaces
```
- 5 If the output displays “N” in the Drop_Recovery column for the target tablespace, type:

```
alter tablespace <tablespaceName> dropped table recovery on
```

Recovery method 1: soft recovery

With a soft recovery, the data in other tables in the tablespace is rolled forward to the most current transaction log. This process prevents any current data that exists in other tables from being overwritten when the backup is restored.

To perform a soft recovery of a dropped database table, perform the following steps.

- 1 To connect to the local instance of the database, repeat [Step 1](#) through [Step 3](#) described in [Setting the Dropped Table Recovery option](#).

- 2 To list the dropped tables, type:

```
list history dropped table all for db <databaseName>
```

Note the Backup ID, name of the dropped table, tablespace, and the DDL statement for the table.

```
DB2=>list history dropped table all for db database_1
List History File for sample
Number of matching file entries = 1
Op  Obj  Timestamp+Sequence  Type  Dev  Earliest Log  Current Log  Backup ID
D    T    20120202114619    000000000000bb000002000d
"ADMINISTRATOR"."EMPLOYEES" resides in 1 tablespace(s):
00001 USERSPACE1
Comment: DROP TABLE
Start Time: 20120202114619
End Time: 20120202114619
00001
DDL: CREATE TABLE "ADMINISTRATOR"."EMPLOYEES" ("NAME" CHAR(20) NOT NULL,
"PROFESSION" VARCHAR(30)) IN "USERSPACE1";
DB2=>
```

- 3 Start the NetVault Backup WebUI on the NetVault Backup Server.
- 4 In the Navigation pane, click **Create Restore Job**, select **Plug-in for DB2** from the **Plugin Type** list, select a saveset that contains the backup of the target tablespace, and click **Next**.

Also select the Archived Logs if they were removed from the database directory on the DB2 Server.

For more information, see [Selecting data for a restore](#).

- 5 Click **Edit Plugin Options**, and configure the following parameters on the **Options** tab:

- **Restore Mode:** Select **Online**.
- **Restore only table space-level backup images:** Select this option.
- **Rollforward Archived Logs:** Clear this check box.

For more information, see [Setting restore type and other restore options](#).

- 6 Complete and submit the job.

For more information, see [Finalizing and submitting the job](#).

- 7 To verify that the Roll Forward state for the target tablespace is set to "Pending," type:

```
list tablespaces
```

Verify that the **Detailed Explanation** column displays "Roll forward pending" status.

```
DB2=>list tablespaces
Tablespaces for Current Database
Tablespace ID      = 0
Name               = SYSCATSPACE
Type               = System managed space
Contents           = Any data
State              = 0x0000
Detailed explanation:
Normal
```

```

Tablespace ID      = 1
Name               = USERSPACE1
Type              = System managed space
Contents          = Any data
State             = 0x0080
Detailed explanation:
    Roll forward pending
DB2=>

```

- 8 On the DB2 Server, create a directory—recovery directory—to store the recovered table, and the note the complete path.
- 9 To perform the roll-forward recovery for the database, connect to the database, and type:


```

rollforward database <databaseName> to end of logs and complete tablespace
<tablespaceName> recover dropped table <backupID> to
"<completePathToRecoveryDirectory>"

```
- 10 To recreate the table, type:


```

create table <DDLstatementNotedEarlier>

```
- 11 To import the table data from the recovery directory, type:


```

import from "<recoveryDirectoryPath>\NODE0000\data" of del messages
<anyFileOutsideTargetDatabase> insert into <tableName>

```
- 12 To verify that the table data has been restored correctly, type:


```

select * from <tableName>

```
- 13 Repeat [Step 7](#) through [Step 12](#) for each table that you want to restore.

Recovery method 2: hard recovery

In a hard recovery, data in the tables of the target tablespace is rolled forward to the point at which the backup was taken. Changes to the tables within the target tablespace that occurred between the time the backup was taken and the time this recovery is performed, are lost. However, this procedure requires the least number of steps and time to perform.

To perform a hard recovery of a dropped database table, perform the following steps.

- 1 To connect to the local instance of the database, repeat [Step 1](#) through [Step 3](#) described in [Setting the Dropped Table Recovery option](#).
- 2 To list the dropped tables, type:


```

list history dropped table all for db <databaseName>

```

Note the names of tablespaces that contain any dropped tables.
- 3 To take the database offline, which is required for a hard recovery, type:


```

disconnect all

```
- 4 Start the NetVault Backup WebUI on the NetVault Backup Server.
- 5 In the Navigation pane, click **Create Restore Job**, select **Plug-in for DB2** from the **Plugin Type** list, select a saveset that contains the backup of the target tablespace, and click **Next**.

Also select the Archived Logs if they were removed from the database directory on the DB2 Server.

For more information, see [Selecting data for a restore](#).
- 6 To avoid overwriting the existing data, use the procedure described in [Renaming or relocating a database during a restore](#) to restore the data to an alternate location.
- 7 Click **Edit Plugin Options**, and configure the following parameters on the **Options** tab:
 - **Restore all table spaces in the database:** Select this option.

- **Rollforward Archived Logs:** Select this check box.
- **Hard Recovery:** Select this option.

For more information, see [Setting restore type and other restore options](#).

8 Complete and submit the job.

For more information, see [Finalizing and submitting the job](#).

9 Connect to the database

10 To list the available tablespaces, type the following command:

```
list tablespaces
```

11 To verify that the table data has been restored correctly, type:

```
select * from <tableName>
```

Troubleshooting

This topic describes some common errors and their solutions. In cases where an error occurs and is not described in this table, obtain the DB2 error number from the NetVault Backup Logs, and then see the relevant DB2 documentation for help.

Table 1. Troubleshooting

Symptom	Error	Explanation or solution
NetVault Backup 10.x service (netvault-pgsql) does not start on Windows.	Check the Windows Event Viewer for the following message: PDT FATAL: lock file "postmaster.pid" already exists.	NetVault Backup 10.x uses a PostgreSQL database. If the PostgreSQL database does not start, NetVault Backup cannot start. To correct this issue, delete the "postmaster.pid" file from the location referenced in the log and restart the NetVault Backup Server. For more information, see https://support.questcom/netvault-backup/kb/122475 .
Backup fails to run	The database is in use. Offline backup is not allowed.	No other application or user should connect to the database.
Backup fails to run	Online backup is not allowed.	The "userexit" or "logretain" database configuration parameter is not enabled.
Backup error	Unable to Get Contact Addresses for "host_name" when running a remote backup.	The Plug-in for DB2 Client cannot obtain a connection with the NetVault Backup Server. Add the IP address and the host information of the NetVault Backup Server to the "...etc\hosts" file.
Backup Completed with warnings	<p>Check the NetVault Backup Logs for the job to see if they display one or both of the following messages:</p> <ul style="list-style-type: none"> Failed to add backup record Failed to write index of backup to the database <p>These messages indicate that the selected data was backed up, but the job's index information was not properly added to NetVault Database. Without this index information, the data cannot be properly restored.</p>	<p>Method 1:</p> <p>Open the Manage Devices page, select the backup media, and click Scan. NetVault Backup stores index information for backup jobs in the NetVault Database and the backup media. When you scan the backup media, the index information is added to the NetVault Database. To verify that the information was added, open the Manage Job Definitions page, and locate the specific job. If you can run the job now, the scan process has corrected the problem.</p> <p>Method 2:</p> <p>If the scan has failed, run the backup job again.</p>

Table 1. Troubleshooting

Symptom	Error	Explanation or solution
Cannot browse a database in the selection tree	SQL0332N Character conversion from the source code page "<codePageNumber>" to the target code page "<codePageNumber>" is not supported. SQLSTATE=57017	<p>When you install NetVault Backup, the system is set in locale US, codeset=ISO8859-1, LANG=C. NetVault Backup uses a constant code page 819.</p> <p>A DB2 database is created using code-set IBM-eucJP Database, code page 954.</p> <p>Result: The system cannot translate between 819 and 954 as they are not compatible. This issue results in the error message "SQL0332N Character conversion from the source code page "819" to the target code page "UNKNOWN" is not supported. SQLSTATE=57017"</p> <p>While installing a non-English version of NetVault Backup, verify that your system is running the correct to prevent any problems related to localization.</p>
Restore job fails	DB2 Server error message 'SQL2043N Unable to start a child process or thread	<p>Unable to start the child processes or threads required during the processing of a database task. There may not be enough available memory or system resources to create the process or thread. Close applications that are not being used or increase your system resources. Update your DB2 version to the latest IBM DB2 fix pack available for such version. Ensure that the system limit for number of processes or threads has not been reached by increasing the limit or by reducing the number of processes or threads already running. Ensure that there is enough virtual memory for new process and threads. Resubmit the restore job.</p>
Relocating Tablespaces completes with warnings for automatic storage tablespaces	The SET TABLESPACE CONTAINERS command is not allowed on an automatic storage tablespace.	<p>Containers associated with automatic storage tablespaces are under the control of the database. In DB2 version 9, the CREATE DATABASE command creates automatic storage enabled databases by default. However the user is allowed to create tablespaces with or without automatic storage enabled.</p>

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece—you—to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit <https://www.quest.com/company/contact-us.aspx> or call +1-949-754-8000.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.