Dell™ NetVault™ Backup Plug-in for PostgreSQL 3.2

User's Guide

Contents

Introducing Dell™ NetVault™ Backup Plug-in for PostgreSQL

- Dell NetVault Backup Plug-in for PostgreSQL – At a Glance

- Key benefits

- Feature summary

- Target audience

- Recommended additional reading

Dell NetVault Backup Plug-in for PostgreSQL – At a Glance

The **Dell NetVault Backup (NetVault Backup) Plug-in for PostgreSQL (Plug-in for PostgreSQL)** increases confidence in recoverability of critical data and eliminates the need for writing complex backup and recovery scripts. Through an intuitive 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 the PostgreSQL databases.

Support for database-cluster backups, individual database and table backups in different backup formats, namely Plain-text SQL Script, TAR Archive, and Custom Archive formats, gives you the flexibility to implement a backup policy that is flexible enough to account for a number of recovery scenarios without learning PostgreSQL internals. The plug-in offers granular control over restores and minimizes downtime by allowing you to restore entire database cluster, individual databases, or individual tables more reliably.

Through automatic integration with a wide range of backup devices, you can be confident that your PostgreSQL data is protected and safely stored offsite to meet your disaster recovery and business continuity goals.

Key benefits

- **Increases confidence and reduces risk while deploying PostgreSQL** – Plug-in for PostgreSQL enables total protection of the PostgreSQL databases without requiring you to write complex backup scripts. The plug-in offers the flexibility to back up the entire database cluster, individual databases, or tables in any of the supported backup formats, which can be used to reconstruct entire or parts of database in the unfortunate event of data loss.

  By relying on the plug-in to implement your backup policies, you can be confident that your data will be protected, and you will be able to recover exactly what is needed in the unfortunate event of a failure.

- **Speeds up restores and reduces downtime** – With Plug-in for PostgreSQL, DBAs are no longer required to write scripts or manually execute restore commands to recover lost data. Just select what needs to be restored from a saved backup set using the point-and-click feature, and the plug-in will recover the data without further interaction. The restores are faster due to reduced human interaction, and the chance for a syntax error in a manual execution that would require a restart is eliminated.

- **Ensures business continuity with automatic integration of backup devices** – Plug-in for PostgreSQL takes advantage of NetVault Backup's fundamental integration with a wide range of backup devices, and enables data backups to tape libraries, stand-alone devices, virtual tape libraries, or shared virtual tape libraries.

  You have the flexibility to take a quick backup on virtual tape libraries, and then later on duplicate this data to physical tape libraries. The ability to store data at offsite locations gives you the confidence that your PostgreSQL environment is protected for disaster recovery purposes.

  With Plug-in for PostgreSQL, the administrators need not be present 24x7. Even less experienced administrators can initiate restores with the certainty that they will be performed correctly and quickly, thereby reducing downtime and improving business continuity.

Feature summary

- **Backup features**
  - Full database cluster or individual database/table backups
  - Plain-text SQL Script, Tar Archive, and Custom Archive backup formats
  - Protection down to the table level
  - Global Objects-only backups
  - Template database backup
  - Data-only backups
  - Configuration files backup
- **Restore features**
  - Restore entire cluster, individual databases, or tables
  - Rename databases during restore
  - Restore to alternate PostgreSQL database cluster
- **Other features**
  - Intuitive point-and-click WebUI
  - Automatic integration with storage devices
  - Job scheduling

Target audience

While advanced DBA skills are generally not required to create and execute routine backup operations, they are required for defining the PostgreSQL Database Backup and Recovery Strategy.

Recommended additional reading

Dell recommends that you have the following documentation available for reference while setting up and using this plug-in.

- **PostgreSQL documentation**, which is available from http://www.postgresql.org/docs/.
- **NetVault Backup documentation:**
  - *Dell NetVault Backup Installation Guide* – This guide provides complete details on installing the NetVault Backup Server and Client software.

- *Dell NetVault Backup Administrator's Guide* – This guide describes how to use NetVault Backup and provides comprehensive information on all NetVault Backup features and functionality.

- *Dell NetVault Backup Command Line Interface Reference Guide* – This guide provides a detailed description of the command line utilities.

You can download these guides from http://software.dell.com/support/.

ⓘ | **IMPORTANT:** Starting with 10.0, 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 and 8.x), refer to the documentation for an earlier version of the plug-in.

Installing and removing the plug-in

- Installation prerequisites
- Reviewing the recommended configuration
- Installing the plug-in
- Removing the plug-in

Installation prerequisites

Before installing Plug-in for PostgreSQL, make sure that the following software is installed and properly configured on the machine designated as the PostgreSQL server:
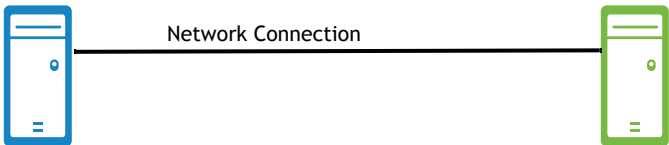
- **NetVault Backup Server/Client software (v8.6.1 or later)** – At least the Client version of NetVault Backup must be installed on the machine configured as the PostgreSQL server.

- **PostgreSQL software** – The machine must be running a supported version PostgreSQL. For more information, refer to the *Dell NetVault Backup Plug-in Compatibility Guide*.

  The **Custom Archive** backup format requires the *zlib compression library*. This library is included by default during PostgreSQL installation unless you specify the **--without-zlib** option. Do *not* use the **--without-zlib** option if you intend to use **Custom Archive** backup format.

Reviewing the recommended configuration

While you can set up a single machine as both the NetVault Backup Server and the PostgreSQL Server (that is, all software installation and configuration requirements are performed on a single machine), Dell recommends that these two entities exist on *separate* machines.

Table 1. Recommended configuration

Network Connection

PostgreSQL Server machine	NetVault Backup Server machine
Software installed/configuration	**Software installed/configuration**
PostgreSQL software (ver. 8.2 or later) **NetVault Backup Server/Client software** **NetVault Backup Plug-in for PostgreSQL**	**NetVault Backup Server software** **PostgreSQL Server added as a NetVault Backup Client** (For complete details on adding a Client to the NetVault Backup Server, see the *Dell NetVault Backup Administrator's Guide*.)

> **IMPORTANT:** Regardless of the configuration that you use (that is, separate NetVault Backup Sever and PostgreSQL server, vs. a single machine configured as both), the plug-in must be installed on the host on which PostgreSQL resides.
>
> Sample procedures throughout this guide assume that you are using this two-machine environment and that configuration requirements have been met.

Installing the plug-in

To install the plug-in

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

> **NOTE:** You can use the configuration wizard to install the plug-in on multiple clients at the same time (provided that the selected clients are all the same type). When you select multiple clients, make sure 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 applicable machine, and click **Manage**.

  c   On the **View Client** page, click the **Install Plug-in** button ( ).

2   Click **Choose Plug-in File**, navigate to the location of the **".npk"** installation file for the plug-in (on the installation CD or the directory to which the file was downloaded from the web site).

   Based on the operating system (OS) in use, the path for this software may vary on the installation CD.

3   Select the file entitled **"pos-x-x-x-x.npk"** (where **xxxx** represents the platform and version number), and click **Open**.

4   Click **Install Plugin** to begin installation.

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

Removing the plug-in

To remove 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 PostgreSQL**, and click the **Remove Plugin** button ( ).

4   In the **Confirm** dialog, click **Remove**.

Configuring the plug-in

- Adding a database cluster
- Setting default actions for error conditions (optional)

Adding a database cluster

To begin using Plug-in for PostgreSQL, you must first add the target database cluster to the plug-in. If you are running more than one database cluster, follow this procedure for *each* database cluster to be protected using the plug-in.

To add a cluster

1   In the Navigation pane of the NetVault Backup WebUI on the NetVault Backup Server, 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 PostgreSQL**.

3   From the **Actions** list, select **Add Database Cluster**.

4   In the **Add Database Cluster** dialog, set the following parameters:

-   **Database Cluster Name** – Enter a name for the database cluster. For easy identification, Dell recommends that you specify the actual PostgreSQL cluster name. When adding more than one cluster, ensure that a unique name is used for identification.

-   **Port** – This field is preconfigured to port 5432, which is the default listener port for PostgreSQL. Enter the correct port number in this field, if PostgreSQL is configured to listen on any other port. When adding multiple database clusters, ensure that a different listener port is configured for each cluster.

-   **Database SuperUser** – This field defaults to the superuser *postgres*. If applicable, enter a different user with superuser privileges for the plug-in's use.

-   **Password** – Enter the password associated with the user specified in the preceding field. Ensure that the password is correct. If the cluster is configured to use TRUST authentication for the HOST connection while using PASSWORD authentication for the LOCAL connection, the backups will fail if the password is incorrect or blank. To avoid this issue, either use TRUST authentication for both the HOST and LOCAL connections (in which case no password is required) or use PASSWORD authentication for both items in the **pg_hba.conf** file (and provide the password here).

-   **PostgreSQL Bin Directory** – Enter the complete path to the PostgreSQL bin directory. For example, this path will be **C:\Program Files\PostgreSQL\8.2\bin** for a default installation of PostgreSQL ver. 8.2 on the Windows platform.

-   **PostgreSQL Library Directory or Path** (Linux/UNIX-based systems only) – Enter the complete path to the directory that contains the **libpq.so** file or to the file itself (for example, **/usr/lib/libpq.so.5.0**).

-   **PostgreSQL Data Directory** – Specify the complete path to the directory that contains the configuration files postgresql.conf, pg_hba.conf, and pg_ident.conf. By default, PostgreSQL stores the configuration files in the **data** directory. For example, this path will be **C:\Program Files\PostgreSQL\8.2\data** for a default installation of PostgreSQL ver. 8.2 on the Windows platform.

- **Database Used for Initial Connection** – For the plug-in's initial connection, the *postgres* database is used by default. To use a different database for connecting to the server, specify the database name in this field.

- **Template Database Used During CREATE Database** – The template specified in this field will be used for creating an empty database during the restore process. Dell recommends that you use the default **template0**. Backups are performed relative to **template0**. This means that any languages and procedures added to **template1** will be included in the backup. As a result, if you are using customized **template1**, you must create the database from an empty template such as **template0**. Otherwise, errors due to duplicate definitions might occur during the restore process.

- **Default Character Set** – From the list, select the Character Set Encoding that the plug-in will use to display the names of databases and tables in the set locale. The following table lists the supported encodings. Note that if an incompatible encoding is configured, you will not be able to browse the tree, or perform backups or restores.

Table 2. Supported encodings

Character set encoding	Language
BIG5	Traditional Chinese
EUC_CN	Simplified Chinese
EUC_JP	Japanese
EUC_KR	Korean
EUC_TW	Traditional Chinese
GB18030	Simplified Chinese
LATIN1	• English • French • German
SQL_ASCII	Any
SJIS	Japanese
UTF8	Any

> ⓘ **IMPORTANT:** The **UTF8** and **SQL_ASCII** encodings can be used with any language, however, the support is limited to ASCII character set. If the database or table names contain non-ASCII characters, you must set a locale-specific encoding (for example, EUC_JP for Japanese). For more information, refer to http://www.postgresql.org/docs/9.2/static/multibyte.html.

5 Click **OK** to save the settings.

The plug-in will verify the values specified and add the database cluster if all parameters are properly configured.

An error message will appear if an invalid entry is found. You can choose to ignore this message and click **Save** to add the database. The database cluster will be added, however, note the following:

- If the superuser name or password *cannot* be validated, further access will be denied.

- If the **PostgreSQL Data Directory** is invalid (while all other parameters are correctly configured), the backups of databases will be completed successfully but the backups of configuration files will fail.

- If the **PostgreSQL Bin Directory** is invalid (while all other parameters are correctly configured), you will be able to submit a job but the actual job will fail.

If an error occurs while adding a database cluster and you need to reconfigure the parameters, see Correcting invalid parameters.

Correcting invalid parameters

If a parameter is not correctly configured, an error message appears. You can ignore this message and save the information. However, further access will be denied or backup jobs might not complete successfully until you correct the error.

To reconfigure the parameter

1   Click **Re-configure** in the error message that appears.

2   In the **Add Database Cluster** dialog, enter the correct values for the parameters reported as invalid.

> ⓘ | **IMPORTANT:** The password field will appear blank. If the previously configured password is correct, you do not have to re-enter the password; the plug-in will use the cached (old) password for authentication when this field is left blank. Otherwise, enter the correct password in this field.

3   Click **OK** to reconfigure the parameters.

Reconfiguring an added database cluster

To reconfigure the parameters for an existing cluster

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.

3   Open the **Plug-in for PostgreSQL** node, and select the applicable database-cluster node.

4   From the **Actions** list, select **Configure**.

The **Configure** dialog appears with all previous settings revealed, allowing you to make any necessary modifications.

5   On the **Configure Database Cluster** dialog that appears, update the parameters as required.

For more information, see Adding a database cluster.

Note the following:

- You cannot change the **Database Cluster Name**.

- The password field will appear blank. If the previously configured password is still valid, you do not have to re-enter the password; the plug-in will use the cached (old) password for authentication when this field is left blank. Otherwise, enter the correct password in this field.

6   Click **OK** to save the changes.

> ⓘ | **IMPORTANT:** If the password for the configured superuser is changed in the PostgreSQL database, you must update the password in the plug-in.

Removing an added database cluster

To remove a cluster

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.

3   Open the **Plug-in for PostgreSQL** node, and select the applicable database-cluster node.

4   From the **Actions** list, select **Remove**.

5   In the confirmation dialog that appears, click **OK**.

> **(i) | IMPORTANT:** The Plug-in for PostgreSQL node will remain in an Open state even after the last database cluster is removed. To close this node and access the other plug-ins in the selection tree, double-click the plug-in node.

Setting default actions for error conditions (optional)

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

To configure default settings

1. In the Navigation pane, click **Change Settings**.

2. On the **Configuration** page, click **Server Settings** or **Client Settings**, as applicable.

3. If you selected **Client Settings**, select the applicable client, and click **Next**.

4. On the **NetVault Server Settings** or **Client Settings** page, click **Plugin Options**.

   In the **Plug-in for PostgreSQL** section, the following items are listed:

   - **Manually Selected Database Unavailable** – This occurs when an individual database that was manually selected (that is, explicitly clicked) for inclusion in the backup is unavailable for backup for any reason, such as being dropped since the backup job was defined.

   - **Manually Selected Table Unavailable** – This occurs when an individual table that was manually selected (that is, explicitly clicked) for inclusion in the backup is unavailable for backup for any reason, such as being dropped since the backup job was defined.

   - **Failed to Backup Configuration Files** – This occurs when a configuration file is not found or cannot be backed up for any other reason. The pg_hba.conf file is used for user authentication; if this file is missing during job execution, the backup will fail because of logon failure (regardless of the action set for the **Failed to Backup Configuration Files** error condition).

5. For each of these conditions, select one of the following settings:

   - **Complete with Warnings - Saveset Retained** – The job returns a status of **"Backup Completed with warnings"** and a backup saveset is created that includes the items that were successfully backed up.

   - **Complete without Warnings - Saveset Retained** – The job completes and returns a status of **"Backup Completed."** While errors are logged in the NetVault Backup binary logs, the errors are essentially ignored in the **Job Status** page and a backup saveset is created that includes the items that were successfully backed up.

   - **Fail - Saveset Retained** – The job returns a status of **"Backup Failed."** However, a backup saveset is generated that includes the items that were successfully backed up.

   - **Fail - No Saveset Retained** – The job returns a status of **"Backup Failed"** and no saveset of backed up objects is kept. That is, even if some of the objects were successfully backed up, the saveset is discarded.

6. Click **Apply** to save the settings.

Backing up data

- Defining a backup strategy
- Performing a backup
- Backing up Global Objects only
- Backing up Configuration Files

> ⓘ **IMPORTANT:** On Windows, you must use appropriate encoding during database cluster initialization (that is, when you execute **initdb**). Backups will fail if you use **UTF-8** encoding and the database or table names contain non-ASCII characters.

Defining a backup strategy

Before commencing with database backups, ensure that you have a backup strategy that safeguards data against media failure, data corruption, user error, and loss of the database server. The following sections provide information that will help you devise a backup strategy for use with Plug-in for PostgreSQL.

- Supported backup formats
- Available backup types
- Backup strategy considerations
- Examples of backup sequences

Supported backup formats

The plug-in supports backups in **Plain-text SQL Script File**, **TAR Archive**, and **Custom Archive** backup formats, which can be used to back up an entire database cluster or individual databases and tables.

- **Plain-text SQL Script File** – This format generates a plain-text file that contains the SQL commands required to reconstruct the database to the state it was in at the time of backup.
- **TAR Archive File** – This format generates a tape archive (TAR) file that is not compressed.
- **Custom Archive File** – This format generates an archive file that is compressed by default. It lets you select the compression level and can be used for tables that are larger than the maximum file size supported by the OS.

  The custom archive format requires the **zlib compression library** and can be used only on PostgreSQL installations built with this library. This library is included by default during PostgreSQL installation unless you use the **--without-zlib** option to disable support for **Custom Archive**.

While similar backup options are available with all three formats, there are minor differences. All three formats let you restore only what is needed, as well as restore data to a different cluster or server.

Available backup types

The plug-in provides the following backup types:

- **Full Database Cluster** – This option backs up all the databases within the cluster in the selected backup format. It also includes a backup of Global Objects, which will be in the form of a Plain-text Script file, regardless of the backup format specified. The plug-in automatically includes the configuration files when the database cluster is selected; however, you can exclude it from the backup if desired.

- **Individual Database/Table** – This option backs up selected databases or tables in the selected backup format. You can use it to back up databases that are updated more frequently than others or use it as a source when creating test environments. This option does *not* include Global Objects when the databases or tables are backed up individually.

- **Global Objects Only** – This option backs up only the Global Objects, which you can use to back up the updates made to the database users, groups, and access permissions.

- **Configuration Files Backup** – You can back up the configuration files with any of the preceding backup types, or separately. The following files are included when you select this option:

  - **postgresql.conf** – This is the primary server configuration file that contains all the parameter settings and directives.

  - **pg_hba.conf** – This is the host-based authentication (HBA) or client authentication file.

  - **pg_ident.conf** – This file stores ident maps (excluding **sameuser**) and is used for ident-based authentication.

  This backup type is similar to a file-system backup. You can use it to restore corrupt or lost configuration files, or to restore the cluster settings to another server.

Backup strategy considerations

When defining the backup strategy, consider the following:

- Database size.

- Frequency of updates and changes.

- Off-peak period for full backups.

- Whether some databases are updated more frequently than others.

- Whether changes are confined to a small or large number of tables in a database.

- Available storage space (and future expansion).

Examples of backup sequences

Following are a few sequences that might meet your PostgreSQL database protection requirements.

- **Full Database Cluster backups only** – If update characteristics are similar across multiple databases, **Full Database Cluster** backups scheduled every night ensures data protection up to the previous day.

  A **Full Database Cluster** backup consumes large amounts of storage space. However, restores are easier and quicker as only a single job is required to restore the entire cluster. Using the **Custom Archive** backup format can reduce the backup size.

- **Full Database Cluster backups and Individual Database/Table backups** – If some databases are updated more frequently than others, regular **Individual Database/Table** backup coupled with periodic **Full Database Cluster** backup will be an ideal strategy. For example, **Full Database Cluster** backups can be performed every Sunday night at 11:00 p.m., while **Individual Database** backups for the databases that are updated frequently can be performed Monday through Saturday at 11:00 p.m. This ensures data protection up to the previous day. You can increase backup frequency to reduce the amount of data loss,

and use the **Custom Archive** format with a higher compression level for large databases to reduce the backup size.

Databases that are updated less frequently can be backed up two or three times a week or on demand after an update is made.

During restores, the databases that have not been updated since the last backup can be restored from the **Full Database Cluster**, while the others can be restored from their last individual backup. Even though the backups are quicker, the restores can take longer due to the additional intervention that is required to execute multiple restore jobs.

- **Global Objects Only backup** – Because the individual database backups exclude the global objects, you can use **Global Objects Only** after updates are made to the database users, groups, and access permissions. This ensures that these recent backups can be used instead of **Full Database Cluster** backups.

- **Individual Database/Table backups and Global Objects Only backups** – You could use a backup strategy of **Individual Database/Table** backups plus **Global Objects Only** backup where you back up the individual databases according to their update frequency. You might back up databases that updated less frequently on a weekly basis and back up the remaining databases updated on a daily basis. Additionally, you might run **Global Objects Only** backups on a weekly basis or on demand after an update to the database groups, users, and access permissions is made.

Performing a backup

A backup using Plug-in for PostgreSQL includes the steps outlined in the following sections.

- Selecting data for backup
- Setting backup options
- Finalizing and submitting the job

Selecting data for backup

You must use sets (Backup Selection Set, Backup Options Set, Schedule Set, Target Set, and Advanced Options Set) to create a backup job. For more information, refer to the *Dell NetVault Backup Administrator's Guide*.

To create a Backup Selection Set

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 title that lets you easily identify the job for monitoring its progress or restoring data. The job name can contain alphanumeric and non-alphanumeric characters, but it cannot include non-English characters. On Linux, the names 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 Click **Create New** next to the **Selections** list.

4 On the **NetVault Backup Selections** page, enter a name for the set in the **Backup Selection Set** box, and then open the NetVault Backup Client on which the plug-in is installed.

5 In the list of plug-ins, open **Plug-in for PostgreSQL**.

   The plug-in attempts to connect to the database server with the name and password specified in the **Configure** dialog. After successful authentication, the added database clusters will be displayed.

6 Open the target database-cluster node, and you will see two nodes: **Configuration Files** and **All Databases**.

- To make **Full Database Cluster** backups, do one of the following:

  - To back up all the databases (including Template Databases) *and* the configuration files contained in the cluster, select the database-cluster node.

  - To back up *only* the databases (including Template Databases), select the **All Databases** node (or clear the **Configuration Files** node after selecting the database-cluster node).

- To select individual databases or tables, do the following:

  - Open the **All Databases** node, and select the required databases. To include Template Databases in the backup, select the node (for example, template1).

  - For individual table selection, open the applicable database node. Select the database node, which will automatically include all the tables, and then clear the check boxes for the tables that you want to exclude from the backup.

  The Template Database nodes cannot be opened further.

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

To use an existing Backup Selection Set

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 the **Selections** list, select an existing Backup Selection Set.

Setting backup options

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

To create a Backup Options Set

1 Click **Create New** next to the **Plugin Options** list.

2 In the **Backup Type** section, select the applicable option:

- **Full Database Cluster**

- **Individual Database/Table**

For more information, see Available backup types.

3 In the **Backup Format** section, select the applicable option:

- **Plain-Text SQL Script File**

- **TAR Archive File (Uncompressed)**

- **Custom Archive File (Compressed)**

For more information, see Supported backup formats.

> ⓘ **IMPORTANT:** **NetVault Backup limitation on Windows** – NetVault Backup resets the selection in **Backup Format** to the default (**Plain Text SQL Script**) if you change the **Backup Type** *after* selecting the file format. To avoid this, select the desired **Backup Format** explicitly after selecting the **Backup Type**.
>
> A similar problem also occurs when you save a **Backup Options** set. NetVault Backup resets the **Backup Format** to the default after you click **OK**. However, the saved **Backup Options** set will contain your actual selection. Similarly, when you try to load a **Backup Options** set, the **Plain Text SQL Script** format will appear as selected, which is incorrect. Internally, the plug-in will use the actual selected **Backup Format**.

4　If you selected a **Backup Format** of **Plain-Text SQL Script File**, select the applicable **Backup Options**:

- **Include Data Only (No Schema Definitions)** – Select this option to back up only the data without the schema. Data-only restores can only be performed for backups performed with this option selected.

    ⓘ | **IMPORTANT:** Template Databases are *not* included if you select this option.

- **Disable Triggers** – This option is only relevant when creating a *data-only* dump. Select this option if there are referential integrity checks or other triggers on the tables. This option temporarily disables triggers on the target tables during restore. After a successful restore, the triggers are enabled.

- **Include BLOBs** – Select this option to back up binary large objects (BLOBs). When this option is selected, all BLOBs within the selected database are backed up *regardless of whether the associated table is included*.

- **Include Object Identifiers** – This option will dump object identifiers (OIDs) as part of the data for every table. Use this option if the OID columns are referenced (for example, in a foreign key constraint). Otherwise, do not select this option.

- **Include Drop Object Commands** – Select this option to drop and re-create objects during the restore process.

- **Exclude Access Privileges** – Select this option to prevent dumping of access privileges (grant/revoke commands).

- **Disable $ Quoting for Function Bodies** – Select this option to disable the use of dollar quoting for function bodies and force them to be quoted using SQL-standard string syntax.

5　If you selected a **Backup Format** of **TAR Archive File**, select the applicable **Backup Options**:

- **Include BLOBs** – Select this option to back up BLOBs. When this option is selected, all BLOBs within the selected database are backed up *regardless of whether the associated table is included*.

- **Disable $ Quoting for Function Bodies** – Select this option to disable the use of dollar quoting for function bodies and force them to be quoted using SQL-standard string syntax.

6　If you selected a **Backup Format** of **Custom Archive File**, select the applicable **Backup Options**:

- **Include BLOBs** – Select this option to back up BLOBs. When this option is selected, all BLOBs within the selected database are backed up *regardless of whether the associated table is included*.

- **Disable $ Quoting for Function Bodies** – Select this option to disable the use of dollar quoting for function bodies and force them to be quoted using SQL-standard string syntax.

- **Compression Level** – Set the compression level between 0 and 9 to use for the custom archive format.

7　If you want to override the default action for an error condition, select the applicable options:

A default action for all backup jobs can be set as outlined in Setting default actions for error conditions (optional).

- For a **Full Database Cluster** backup, the action can be set for the following scenario: **Failed to Backup Configuration Files**. This occurs when a configuration file is not found or cannot be backed up for any other reason. The pg_hba.conf file is used for user authentication; if this file is missing during job execution, the backup will fail because of logon failure (regardless of the action set for the **Failed to Backup Configuration Files** error condition).

- For an **Individual Database/Table** backup, the action can be set for the following additional scenarios:

    - **Manually Selected Database Unavailable** – This occurs when an individual database that was manually selected for inclusion in the backup is unavailable for backup for any reason.

- **Manually Selected Table Unavailable** – This occurs when an individual table that was manually selected for inclusion in the backup is unavailable for backup for any reason.

The plug-in can be configured to take one of following possible courses of action when any of the preceding error conditions occur during the job execution:

- **Complete with Warnings - Saveset Retained** – The job returns a status of **"Backup Completed with warnings"** and a backup saveset is created that includes the items that were successfully backed up.

- **Complete without Warnings - Saveset Retained** – The job completes and returns a status of **"Backup Completed."** While errors are logged in the NetVault Backup binary logs, the errors are essentially ignored in the **Job Status** page and a backup saveset is created that includes the items that were successfully backed up.

- **Fail - Saveset Retained** – The job returns a status of **"Backup Failed."** However, a backup saveset is generated that includes the items that were successfully backed up.

- **Fail - No Saveset Retained** – The job returns a status of **"Backup Failed"** and no saveset of backed up objects is kept. That is, even if some of the objects were successfully backed up, the saveset is discarded.

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

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

The name can contain alphanumeric and non-alphanumeric characters, but it cannot include non-English characters. On Linux, the names 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.

To use an existing Backup Options Set

In the **Plugin Options** list, select the existing Backup Options Set that you want to use.

Finalizing and submitting the job

To finalize and submit a backup job

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.

> ⓘ | **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.**

Backing up Global Objects only

You can back up Global Objects in two ways:

- **Via the Full Database Cluster backup type** – A **Full Database Cluster** backup performed in any format includes the Global Objects. Because these are cluster-wide and not specific to any individual database, backups of type **Individual Database/Table** do not include Global Objects.

- **Via the Global Objects Only backup type** – You can also back up the Global Objects separately using the **Global Objects Only** backup type provided by the plug-in.

To create a Global Object Only backup

1   Following the steps outlined in Selecting data for backup, select the target database-cluster node.

2   If you want to exclude the configuration files, which are automatically selected when you select the database-cluster node, clear the **Configuration Files** node in the selection tree.

3   Following the steps outlined in Setting backup options, select **Global Objects Only** in the **Backup Type** section.

4   If you included the configuration files, select the required action to take if an error occurs during the backup:

   - **Complete with Warnings - Saveset Retained**

   - **Complete without Warnings - Saveset Retained**

   - **Fail - Saveset Retained**

   - **Fail - No Saveset Retained**

   Because the **pg_hba.conf** file is used for user authentication, backup will fail if this file is missing during job execution. This logon failure causes the backup to fail regardless of the action set for the **Failed to Backup Configuration Files** error condition.

5   Following the steps outlined in Finalizing and submitting the job, complete and submit the job.

Backing up Configuration Files

You can select configuration files for backup with any of the backup types. Regardless of the format selected, the plug-in performs a file-system backup of the following files:

- **postgresql.conf**

- **pg_hba.conf**

- **pg_ident.conf**

To back up these files separately

1   Following the steps outlined in Selecting data for backup, select the **Configuration Files** node in the selection tree.

2   Because the **Backup Type** and **Backup Format** are not considered, select any of the options in **Backup Options** when you create the Backup Options Set by following the steps outlined in Setting backup options.

3   Following the steps outlined in Finalizing and submitting the job, complete and submit the job.

Restoring data

- Selecting data for restore
- Setting restore options
- Finalizing and submitting a job
- Performing advanced restore procedures

Selecting data for restore

To select data to restore

1 In the Navigation pane, click **Create Restore Job**.

2 On the **Create Restore Job - Choose Saveset** page, select **Plug-in for PostgreSQL** 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 saveset to display the database cluster that was the target of the backup.

   When you select a saveset, the following details are displayed in the **Saveset Information** area: Job ID, Job Title, name of the NetVault Backup Server, name of the client from which the data was backed up, plug-in used to create the saveset, saveset creation date and time, saveset retirement setting, whether it is an Incremental Backup, whether it is an Archive, and saveset size.

   The following table outlines the backup type identifiers:

Table 3. Backup type identifiers

Backup type identifier	Backup type and format
SQL FULL DB CLUSTER ARCHIVE	• Full Database Cluster • TAR/Custom Archive File
SQL FULL DB CLUSTER PLAIN	• Full Database Cluster • Plain-text SQL Script File
SQL INDIVIDUAL ARCHIVE	• Individual Database/Table • TAR/Custom Archive File
SQL INDIVIDUAL PLAIN	• Individual Database/Table • Plain-text SQL Script File
SQL GLOBAL OBJECTS ONLY	• Global Objects Only

5 To continue, click **Next**.

6 On the **Create Selection Set** page, select the data that you want to restore.

Based on the backup type, opening the database cluster will display the following:

- **Full Database Cluster** – The following nodes are displayed for this type:

    - **Configuration Files** (if included in the backup)

    - **Global Objects**

    - **All Databases**

- **Individual Database Cluster** – The following nodes are displayed for this type:

    - **Configuration Files** (if included in the backup)

    - **All Databases**

7   To view the databases included in the backup, open **All Databases**.

8   For next-level selections, drill-down further.

9   Select the database cluster, databases, or tables as required.

> ⓘ **IMPORTANT:** BLOBs included in the backup appear as a separate node (and not beneath the associated table). Selecting this node restores all BLOBs contained in the saveset and not just objects associated with the selected table. Also note that, to restore BLOBs, this node must be selected explicitly (that is, either this node or the parent database node must be selected). Selecting only the table will *not* restore the associated BLOBs.

Setting restore options

To set the restore options

1   On the **Create Selection Set** page, click **Edit Plugin Options**, and configure the applicable parameters on the **Options** tabs.

2   If you are restoring a **Plain-Text SQL Script File** backup, select the applicable options in the **Pre-Restore Options** section to specify what actions should be taken before the plug-in begins the restoration of a database:

- **Drop and Recreate Database** – Select this option if you want to drop an existing database before restoring it from the backup. The database will be created using the template specified in the **Configure** dialog (**Template Database Used during CREATE Database** option). Note that all existing data is discarded if this option is selected.

    > ⓘ **IMPORTANT:** Regardless of whether the **Drop and Recreate Database** option is selected, the plug-in performs this action if you are restoring Template Databases.

- **Create Database Only** – Select this option if you want to create a non-existent database or to write data to the existing database. If the database is not found during restore, a new database is created. If the database is present, tables that are not part of the current restore job are not modified in any manner. If the selected tables are present in the database, the action is based on whether the **Include Drop Object Commands** option was selected during backup:

    - **Include Drop Object Commands was selected** – The table is dropped and re-created. The restored table will only contain records that were present at the time of backup.

    - **Include Drop Object Commands was not selected** – For existing tables, you can use the **Truncate Table(s)** option to truncate existing records before restoring.

        If **Truncate Table(s)** is not enabled, the plug-in attempts to insert all the backed up records into the existing table. For tables on which the unique or primary keys have been defined, this action might result in constraint errors.

- **Truncate Table(s)** – Select this option if you want to delete all existing records from the tables and restore the records present in the backup saveset. After restoration, the table will only contain records that were present at the time of backup. Be aware of the following:

  - The database must exist.

  - For existing tables, the **"create table"** command during **psql** execution will fail. If the **Stop on Error** option is selected, the restore stops immediately on encountering this error. Select the **Continue with Warnings** option if the selected tables are present to ignore these errors and continue with table truncation and data restoration.

  - No errors will be encountered for tables that do not exist.

3 If you are restoring a **Plain-Text SQL Script File** backup, select the applicable options in the **Restore Options** section:

- **Stop on Error** – Select this option if you want the restore to stop immediately when an error occurs. The changes made prior to the error are not rolled back, which leaves the database in an inconsistent state.

  Note that the entire restore job will not be aborted; only the restoration of the current database will stop. Restoration of other databases included in the job will continue normally.

  If the **Disable Triggers** option was enabled for backup and the restoration of the database stops in between, the plug-in will *not* enable triggers defined for the last table that caused the failure.

- **Restore as Single Transaction** – Select this option if all commands in the database backup must complete successfully; otherwise, no changes are applied. The changes made are rolled back if an error occurs.

  Note that the commit or roll-back action applies to a specific database and not the entire restore job. The databases for which restoration has completed successfully are not affected. Similarly, this will not affect restoration of remaining databases in the job.

- **Continue Restore with Warnings** – Select this option if you want the plug-in to ignore the errors and continue with the restore. The errors encountered during job execution are recorded as warning messages in the NetVault Backup logs and the job status is reported as **"Completed with Warnings."**

  You can use this option to restore non-existing or non-duplicate records into existing tables without truncating the existing records.

  > ⓘ | **IMPORTANT:** If you select this option, review the logs after the job is complete to confirm that they only contain warnings pertaining to the following:
  >
  > **Primary or Unique Key ConstraintsDatabase Object (like table, language, user, role, and others) Already Exists**
  >
  > You can ignore these errors. However, if any other type of error message is noted, you must take the necessary corrective action, such as restoring the database again by selecting **Drop and Recreate Database or Truncate Table(s)**.

4 If you are restoring a **TAR Archive File** or **Custom Archive File** backup, select the applicable options:

- **Stop on Error** – Select this option if you want the restore to stop immediately when an error occurs. The changes made prior to the error are not rolled back, which leaves the database in an inconsistent state.

  Note that the entire restore job will not be aborted; only the restoration of the current database will stop. Restoration of other databases included in the job will continue normally.

> **IMPORTANT:** If this option is not selected, the plug-in ignores the errors and continues the job. The errors encountered during job execution are recorded as warning messages in the NetVault Backup logs and the job status is reported as **"Completed with Warnings."** Review the logs after the job is complete to confirm that they only contain warnings pertaining to the following:
>
> **Primary or Unique Key ConstraintsDatabase Object (like table, language, user, role, and others) Already Exists**
>
> You can ignore these errors. However, if any other type of error message is noted, you must take the necessary corrective action, such as restoring the database again by selecting Drop and **Recreate Database** or **Do Not Restore Data for Failed Create Table Commands**.

- **Drop and Recreate Database** – Select this option if you want to drop an existing database before restoring it from the backup. The database will be created using the template specified in the **Configure** dialog (**Template Database Used during CREATE Database** option).

  You can also use the **Exclude Access Privileges** option in conjunction with **Create Database Only**. Select this option to prevent restoration of access privileges (grant/revoke commands). Without this option selected, the restoration will fail if the required users/roles do not exist.

  > **IMPORTANT:** Regardless of whether the **Drop and Recreate Database** option is selected, the plug-in performs this action if you are restoring Template Databases.

- **Create Database Only** – Select this option if you want to create a non-existent database or to write data to the existing database. If the database is not found during restore, a new database is created. If the database is present, tables that are not part of the current restore job are not modified in any manner. If the selected tables are present in the database, the action is based on the following options:

  - **Drop Database Objects** – Select this option to drop the existing table and restore it from the backup. You will lose the existing records. The restored table will only contain records that were present at the time of backup.

  - **Restore Data for Failed Create Table Commands** – Select this option if you want the plug-in to insert all the backed up records into the existing table. For tables on which the unique or primary keys have been defined, this action might result in constraint errors.

  - If **Stop on Error** option is selected, the restoration of the database will stop.

  - **Do Not Restore Data for Failed Create Table Commands** – With this option enabled, the existing tables will not be restored.

  You can also use the **Exclude Access Privileges** option in conjunction with **Create Database Only**. Select this option to prevent restoration of access privileges (grant/revoke commands). Without this option selected, the restoration will fail if the required users/roles do not exist.

- **Restore Data Only (No Schema Definitions)** – Select this option if you want to restore only the data from the backup and exclude the schema definitions. The restore will fail if the target database or table does not exist.

  > **IMPORTANT:** Template Databases are *not* restored if you select this option.

  You can also use the **Disable Triggers** option in conjunction with **Restore Data Only**. Select this option if there are referential integrity checks or other triggers on the tables. This option will temporarily disable triggers on the target tables during restore. After successful data restoration, the triggers will be enabled. If **Stop on Error** is selected and the restoration of the database stops in between, the plug-in will *not* enable triggers defined for the last table that caused the failure.

Finalizing and submitting a job

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

To finalize and submit a restore job

1   Click **Ok** to save the settings, 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 title that lets you easily identify the job for monitoring its progress. The job name can contain alphanumeric and non-alphanumeric characters, but it cannot include non-English characters. On Linux, the names 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.

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

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

5   Click **Submit** to submit the job for scheduling.

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

Performing advanced restore procedures

You can complete the following advanced procedures:

- Restoring Global Objects
- Restoring configuration files
- Renaming a database during restore
- Restoring data to a different database cluster
- Restoring to a database cluster on an alternate server

Restoring Global Objects

You can restore Global Objects from a T Plain-Text SQL Script File version of a Full Database Cluster or Global Objects Only backup.

To restore Global Objects

1   Select the **Global Objects** node from the backup saveset by following the instructions outlined in Selecting data for restore.

2   Set the **Restore Options** by following the instructions outlined in **Setting restore options** and completing the following parameters:

    - **Stop on Error** – Select this option if you want the restore to stop immediately when an error occurs.

- **Restore as Single Transaction** – Select this option if all commands in the database backup must complete successfully; otherwise, no changes are applied. The changes made are rolled back if an error occurs.

- **Continue Restore with Warnings** – Select this option if you want the plug-in to ignore the errors and continue with the restore. The errors encountered during job execution are recorded as warning messages in the NetVault Backup logs and the job status is reported as **"Completed with Warnings."**

   > ⓘ **IMPORTANT:** PostgreSQL does not allow granular restores of individual Global Objects, nor does it allow you to delete global objects before restoring. Therefore, when restoring Global Objects, all the Global Objects are restored. As a result you will encounter error messages about duplicate items during restore. Dell recommends that you restore Global Objects to an empty database cluster during disaster recovery or use the **Continue Restore with Warnings** option. If the Global Objects are being restored to recover dropped roles or users and the **Continue Restore With Warnings** option has been selected, warnings will be logged in the NetVault Backup binary logs for roles or users that exist.

3   Continue with the restore procedure as explained in Finalizing and submitting a job.

To restore Global Objects from a **TAR Archive File** or **Custom Archive File** backup, no additional **Restore Options** are available. Perform Step 1 and Step 3 to restore Global Objects from these backup formats.

For restoring the data to a different database cluster, see Restoring data to a different database cluster.

Restoring configuration files

When restoring configuration files, you do not have to set any restore options (unless you are restoring to a different cluster or server; for more information, see Restoring data to a different database cluster). The plug-in does not let you select individual configuration files for restoration.

To restore the configuration files

1   Select the **Configuration Files** node from the backup saveset by following the instructions outlined in Selecting data for restore.

2   Set the **Restore Options** by following the instructions outlined in **Setting restore options**.

3   Continue with the restore procedure as explained in Finalizing and submitting a job.

4   After restoring the configuration files, restart the PostgreSQL Server to apply the settings.

   > ⓘ **IMPORTANT:** You can restore configuration files when the database server is not running. However, you should not include database, table, and Global Objects in the same job. If you attempt to restore any of these items together with the configuration files when the database server is not running, the restore job will fail as a result of logon failure.

Renaming a database during restore

Renaming a database can be useful if you do not want to overwrite the existing version and want to create a copy of the database.

To rename the database during restore

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

   For more information, see Selecting data for restore.

2   On the **Create Selection Set** page, select the database that you want to rename.

3   From the **Actions** list, select **Rename**.

4   In the **Rename/Relocate** dialog, enter the new name in the **Rename to** box.

You can enter the name of an existing database to restore the selected tables to this database.

> ⓘ **IMPORTANT:** If the cluster is configured to use UTF-8 encoding at **initdb**, you cannot use a name that contains non-ASCII characters.

5 Click **Ok**.

The database item is accompanied by renaming information in parenthesis.

6 Continue with the restore procedure as explained in Setting restore options and Finalizing and submitting a job.

Restoring data to a different database cluster

You can restore a database to a different database cluster running on the database server. The following sections outline the procedure for performing this type of a restore.

Prior to initiating this type of restore, make sure that the following prerequisites are met:

- **Same version of PostgreSQL installed** – The PostgreSQL version must be same as the version running on the existing database server.

- **Database cluster exists on the PostgreSQL Server** – A database cluster must exist on the PostgreSQL server and this cluster *must* be added to the plug-in. To add the cluster, see Adding a database cluster.

To perform the restore

1 Select the data to be restored from the backup saveset by following the instructions outlined in Selecting data for restore.

2 Set the **Restore Options** by following the instructions outlined in **Setting restore options** and completing the following parameter:

Specify the name of the database cluster to which the selected database is to be restored using the **Database Cluster Name** box in the **Restore Target Details** section.

3 Complete the procedure by following the instructions outlined in Finalizing and submitting a job.

> ⓘ **IMPORTANT:** You can also Rename a database while performing this type of restore. For more information, see Renaming a database during restore.

Restoring to a database cluster on an alternate server

You can restore a database cluster to a different PostgreSQL Server. Prior to initiating this type of restore, make sure that the following prerequisites are met on the new restore target:

- **Same OS** – The OS running on the target system must be same as the source system.

- **Same version of PostgreSQL installed** – The PostgreSQL version must be same as the version running on the existing database server.

- **NetVault Backup software and Plug-in for PostgreSQL installed** – The same version of NetVault Backup software and the plug-in must be installed and configured on the new restore target.

- **Client added to the NetVault Backup Server** – The target machine must be added to the NetVault Backup Server via the NetVault Backup WebUI.

- **Database cluster exists on the new restore target** – A database cluster must exist on the new restore target and the cluster *must* be added to the plug-in. To add the cluster, see Adding a database cluster.

To perform the restore

ⓘ | **IMPORTANT:** If you are not restoring from a Full Database Cluster backup, you must restore the Global Objects first.

1   Select the database to be restored from the backup saveset by following the instructions outlined in Selecting data for restore.

2   Set the **Restore Options** by following the instructions outlined in **Setting restore options**.

   In addition, specify the name of the database cluster to which the selected database is to be restored using the **Database Cluster Name** box in the **Restore Target Details** section.

3   Click **Ok** to save the settings, and then click **Next**.

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

   Assign a descriptive title that lets you easily identify the job for monitoring its progress. The job name can contain alphanumeric and non-alphanumeric characters, but it cannot include non-English characters. On Linux, the names 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.

   ⓘ | **IMPORTANT:** Do not use special characters that are not supported in a file name on the target OS. For example, the characters /,\,*,@, and so on, should not be used on Windows. (This is because Plug-in for PostgreSQL tries to create a folder with the same name as the Job Title for restoring data temporarily.)

5   In the **Target Client** list, select the new restore target.

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

6   Use the **Schedule** and **Advanced Options** lists to configure the any additional required options.

7   Click **Submit** to submit the job for scheduling.

Troubleshooting

This section describes some common errors and their solutions. In those cases where an error occurs and is not described in this table, obtain the PostgreSQL error number from the NetVault Backup Logs, and then refer to the relevant PostgreSQL documentation for the resolution.

Table 4. Troubleshooting

Symptom	Error	Explanation
NetVault Backup 10.x service (netvault-pgsql) will not start on Windows	Check the Windows Event Viewer to see if it displays 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, refer to https://support.software.dell.com/netvault-backup/kb/122475.
Log on failed	Cannot connect to the database server.	The PostgreSQL Server is not running. Start the services.
	Password Authentication Failed for the User *<username>*.	The superuser password set in the **Configure** dialog is invalid. Update the password in the plug-in by following the instructions in Reconfiguring an added database cluster.
	Authentication request failed.	The **pg_hba.conf** is not available for user authentication.

Table 4. Troubleshooting

Symptom	Error	Explanation
Backup failed	Login Failed.	• **Cannot connect to the database server** – The PostgreSQL Server is not running. Start the services. • **Fatal: Password Authentication Failed for the User *<username>*** – The superuser password was changed in the database server after job submission. Update the password in the plug-in by following the instructions in Reconfiguring an added database cluster. • **Authentication request failed** – The **pg_hba.conf** is not available for user authentication.
	Backup in **TAR Archive File** format fails with the following message: **Database found, but failed to backup Database *<dbname>*.**	Check the **pg_dump** command with the **TAR Archive** format from a command prompt or terminal session. If the following error appears, **"pg_dump: [tar archiver] actual file length *<xxxx>* does not match expected *<xxxx>*,"** you cannot use the **TAR Archive File** backup format on the current installation. Use the **Custom Archive File** format instead.
	On Windows, the backup fails with the following message: **Database was found, but failed to backup *<locale_name>* database.**	Check the **pg_dump** command from the command prompt. If the following error appears, **"pg_dump: Cannot access "??-UTF8-C" [database to be dumped]: FATAL: database "??-UTF8-C" does not exist pg_dump: *** Stopped by error,"** the problem is in the character-set encoding configured for the database cluster. The backups of databases or tables named with non-ASCII characters is not supported if the encoding is set to UTF-8 at **initdb**. Appropriate encoding must be set for the database cluster (for example, EUC_JP for Japanese or EUC_KR for Korean) to back up the databases and tables.
	Backup fails with the following message: **Database found, but failed to backup Database XXXXX.**	Ensure that the password provided in the **Add Database Cluster/Configure Database Cluster** dialog of the plug-in is correct. If the cluster is configured to use TRUST authentication for HOST connection while PASSWORD authentication for LOCAL connection, the backups will fail if the password is incorrect or blank. Note that you will be able to browse the tree and submit a job with a blank or incorrect password. However, the actual job will fail. Dell recommends the following in this scenario: • In the **Configure Database Cluster** dialog, enter the correct password, and run the job again. • To avoid this issue, either use TRUST authentication for both HOST and LOCAL connection (in which case no password is required) or use PASSWORD authentication for both in the pg_hba.conf file (and provide the correct password in the **Configure Database Cluster** dialog).

Table 4. Troubleshooting

Symptom	Error	Explanation
Backup Completed with warnings	If the backup job shows this status, check the NetVault Backup Log entries for the job to see if they display one or both of the following messages: • **Failed to add backup record** • **Failed to write index of backup to the database** These messages indicate that the selected data was actually backed up, but the job's index information was not properly added by NetVault Backup to its database. Without this index information, the data cannot be properly restored.	**Method 1:** Access the **Manage Devices** page of the NetVault Backup WebUI, and perform a scan of the media targeted by the job. NetVault Backup stores index information for backup jobs in two locations: in the NetVault Database and on the media targeted by the backup. Performing this scan adds the index information to the NetVault Database. To verify this, open the **Create Restore Job - Choose Saveset** page and locate the specific job. If you can browse it and set up a restore job, the scan process has corrected the problem. **Method 2:** If the previous method failed, you must run the backup job again.
Restore failed	Login Failed.	• **Cannot connect to the database server** – The PostgreSQL Server is not running. Start the services. • **Fatal: Password Authentication Failed for the User *<username>*** – The superuser password was changed in the database server after job submission. Update the password in the plug-in by following the instructions in Reconfiguring an added database cluster. • **Authentication request failed** – The **pg_hba.conf** is not available for user authentication.

Dell listens to customers and delivers worldwide innovative technology, business solutions and services they trust and value. For more information, visit www.software.dell.com.

Contacting Dell

Technical support:
Online support

Product questions and sales:
(800) 306-9329

Email:
info@software.dell.com

Technical support resources

Technical support is available to customers who have purchased Dell software with a valid maintenance contract and to customers who have trial versions. To access the Support Portal, go to https://support.software.dell.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. In addition, the portal provides direct access to product support engineers through an online Service Request system.

The site enables you to:

- Create, update, and manage Service Requests (cases)
- View Knowledge Base articles
- Obtain product notifications
- Download software. For trial software, go to Trial Downloads.
- View how-to videos
- Engage in community discussions
- Chat with a support engineer