Peer Global File Service User Guide

Updated Monday, August 15, 2022

Table of Contents

Peer Global File Service Help

Using This Help File

This help is designed to be used online. It is cross-linked so that you can find more relevant information to any subject from any location. If you prefer reading printed manuals, a PDF version of the entire help is available from our website. This may be useful as a reference, but you will probably find that the active hyperlinks, cross-references, and active index make the on-screen electronic version of this document much more useful.

Trademark Information and Copyright

Last updated: Monday, August 15, 2022

PeerGFS Version 5.0.0

Terminology

Introduction

Before getting started, it is important to have a good understanding of key concepts and terminology used throughout this help system.

Terms

Term	Definition
Active-Active	Two or more file servers that hosts data sets that are in active use, as opposed to an active-passive environment where only one file server hosts active data. Made possible by real-time file synchronization to keep all file servers in sync.
Agent	See *Peer agent*
Cloud Backup and Replication job	A job type where a single participating host has a designated set of folders and files to be replicated to a cloud storage device.
DFS (Distributed File System)	A set of client and server services that allow an organization using Microsoft Windows servers to organize many distributed SMB file shares into a distributed file system.
DFS namespace (DFS-N)	A namespace that enables you to group shared folders located on different servers into one or more logically structured namespaces.
DFS Namespaces	A Windows Server feature that allows multiple SMB shares across different file servers (and even locations) to be combined into a single unified namespace. DFS Namespaces simplifies access to files, especially in large, distributed environments. When combined with Peer file synchronization technology, DFS Namespaces can provide redundancy to file shares across file servers and locations.
DFS-N Management job	A type of job that enables the creation and management of DFS namespaces.
Event	A single operation performed by a user on a file server.

Term	Definition
Failback	The process of redirecting previously displaced users from a secondary file server back to the primary after a failure state has been resolved.
Failover	The process of redirecting users from one file server to a secondary in the event of a failure.
File access event	An event that is triggered from the opening or closing of a file.
File change event	An event that causes a file to be changed in some way, for example, file modify, file delete, file rename, file attribute change.
File Collaboration job	A type of job that combines file synchronization with distributed file locking to prevent version conflicts across multiple active file servers.
File Collaboration session	A communication session made up of two or more hosts, each with a designated root of folders and files that are to be shared or collaborated on. A collaboration session coordinates the primary functions of file locking and synchronization.
File filter	A type of filter used to include or exclude specific files from replication and locking.
File lock conflict	A file collaboration condition that exists when two users open a file at the same time, and both hold exclusive locks on the file.
File Replication job	A type of job that involves real-time and/or scheduled copying of files and folders from one file server to another.

Term	Definition
File Synchronization job	A type of job that involves multi-directional real-time replication so that two or more file servers are always up to date with each other.
Filter	Three types of filters: file, folder, and list.
Filter expression	See *list filter*.
Folder filter	A type of filter used to include or exclude specific folders (and the files they contain) from replication and locking.
Heartbeat	A communication mechanism used between Peer Management Center and all connected Peer Agents to ensure that Peer Agents are alive and responsive. Heartbeats share information about the Peer Agent host server with Peer Management Center, aid in verifying when a Peer Agents is no longer available, and signal when a disconnected Peer Agent has reconnected. All heartbeat information is sent through the Peer Management Broker.
Host	A server that a Peer Agent is installed upon.
Initialization process	The steps executed whenever a job is started in Peer Management Center. The steps for an initialization process are different for each job type.
Initial synchronization process	The background process that occurs at the start of a file collaboration session, where the directory watch set is recursively scanned on all participating hosts, file conflict resolution is performed, and any files that require updating are synchronized with the most current copy of the file.
List filter	A type of filter used to show or hide information from various views in Peer

Term	Definition
	Management Center.
Management Agent	A server running the Peer Agent. Can manage storage devices or a DFS namespace.
Master host	In file synchronization and collaboration, the master host will always win in a split-brain scenario.
Malicious Event Detector (MED)	Leverages the same real-time event detection that powers all job types to detect and alert administrators to malicious user and application behavior. For more information, see: https://kb.peersoftware.com/tb/introduction-to-peer-med.
Participant	A participant consists of an Agent and the volume/share/folder to be replicated. Applies to File Collaboration, File Replication, and File Synchronization jobs.
Peer Agent (or Agent)	A lightweight piece of software that is installed on Windows Servers to perform the storage and file management functions used by the entire Peer Global File Service solution suite. Typically installed on or alongside the file servers that will be managed by Peer Management Center.
Peer Management Center (PMC)	The focal component of Peer Global File Service. Responsible for configuration, management, and monitoring of Peer Agents and the various solutions configured in Peer Global File Service. Peer Management Center runs as three parts: a Windows Service that is always running, along with a rich client application and a web server component, both used for configuration and monitoring.

Term	Definition
Peer Management Broker	The central messaging system of Peer Global File Service. The Peer Management Broker serves to connect Peer Management Center and Peer Agents, forming a Peer Management Center "network" that can be cast over local or wide-area networks via TCP/IP. One or more Peer Management Brokers are deployed in a Peer Management Center environment.
Quarantined file	A file that has been removed from a File Collaboration or File Synchronization job as a result of a lock or replication conflict that could not be automatically resolved. This file will not be deleted from any location but will be ignored while it remains in quarantine. An administrator or help desk user must manually remove files from quarantine.
Quorum	Requirement for a minimum of two participants must be available and connected. If that number dips to one or less, quorum will not be met. Applies to File Collaboration, File Replication, and File Synchronization jobs.
Real-time event detection	A key technology that backs all job types in Peer Management Center. Peer Management Center receives notifications as end users interact with the file servers that are being monitored. These notifications will usually result in replication or locking between file servers.
Scan	The initial process of comparing data sets on two or more file servers to ensure that they match. As differences are discovered, replication will occur to bring each file server "in sync" with one another.
Seeding target	Smart data seeding helps to efficiently integrate a host that has been disconnected for a long period of time or a new host into a File Collaboration job. Such existing hosts or new hosts with pre-seeded data (using

Term	Definition
	methods like shipping a drive or server) should be set as Seeding Targets within a collaboration job. When the scan starts, non-Seeding Targets will become the masters and bring the Seeding Targets up to date. Stale updates, deletes, and renames will NOT be brought back from the Seeding Targets. All local real-time activity will be quarantined. Once that initial scan is complete, the Seeding Targets will become full participants with real-time enabled. For more information on Smart Data Seeding and its potential options, see Smart Data Seeding or contact support@peersoftware.com.
SMB/CIFS	Server Message Block or Common Internet File System, an application-layer protocol used for providing shared access to file data and other networked resources.
Source host	The file server hosting a file from which file access or change event originated.
Target host	One or more Management Agents of file servers where file access and change events will be propagated to.
TLS	Transport Layer Security, a successor to Secure Socket Layer (SSL) that secures network traffic between a client and server.
UNC Path	A UNC path can be used to access network resources and must be in the format specified by the Universal Naming Convention. A UNC path always starts with two backslash characters (\\).
View	Individual sections of Peer Management Center's user interface, each providing unique information and control.

Term	Definition
	Examples: Main view, Jobs view, Agent Summary view, Alerts view, Job Alerts view.
Volume Shadow Copy Service (VSS)	Shadow Copy is a technology included in Microsoft Windows that allows taking manual or automatic snapshots of computer files or volumes, even when they are in use. It is implemented as a Windows service called the Volume Shadow Copy service.
Watch set	The root folder and all subfolders on a file server that are being scanned and/or monitored by a File Collaboration, File Replication, File Synchronization, or Cloud Backup and Replication job.

Installation and Configuration

This topics in this section provide information about:

Requirements and Prerequisites

Installing Peer Management Center

Installing Peer Agents

Updating Peer Management Center and Peer Agents

Configuring and Managing the Web and API Services

For information about Peer Global File System licensing, see Licensing.

Requirements and Prerequisites

Before you get started, review the environmental requirements and platform prerequisites for using Peer Global File Service:

- Peer Global File Service environmental requirements
- Dell EMC Prerequisites
- NetApp Prerequisites
- Nutanix Prerequisites

Dell EMC Prerequisites

In addition to the standard Peer Global File Service environmental requirements, the following prerequisites must be met:

- Dell EMC Isilon Prerequisites
- Dell EMC Unity Prerequisites
- Dell EMC Celerra | VNX | VNX 2 Prerequisites

CEE Server Configuration

See the following guides for steps on setting up a CEE Server on which the Peer Agent will be running:

- Dell EMC Isilon Configuration Guide
- Dell EMC Unity Configuration Guide
- Dell EMC Celerra | VNX | VNX 2 Configuration Guide

NetApp Prerequisites

In addition to the standard Peer Global File Service environmental requirements, the following prerequisites must be met:

- NetApp Data ONTAP 7-Mode Prerequisites
- NetApp ONTAP | Clustered Data ONTAP Prerequisites

Nutanix Prerequisites

In addition to the standard Peer Global File Service environmental requirements, the following prerequisites must be met:

- Nutanix Files prerequisites

Installing Peer Management Center

Overview

Peer Management Center (PMC) can be installed in numerous ways based on your needs and environment. Peer Management Center installation consists of two separate installers, both of which are available for download from the Peer Software website:

- **Peer Management Center installer:** This installer installs both Peer Management Center and Peer Management Broker on the same server. Peer Management Broker handles the communication between Peer Management Center and Peer Agents. See Installing and Launching Peer Management Center for installation instructions.
- **Peer Agent installer:** This installer contains the Peer Agent installation files. You must install an Agent on each server you plan to include in any of your jobs. See Installing Peer Agents for installation instructions.

Before installing Peer Management Center, see Requirements and Prerequisites to verify that your environment satisfies the requirements and prerequisites.

Installing and Launching Peer Management Center

To install and launch Peer Management Center and Peer Management Broker:

1. Download the Peer Management Center installer (**PMC_Installer_win64.exe**) to the server you want to host Peer Management Center.
2. Run the installer and follow the installation wizard instructions.

   During the installation, you will be prompted to configure access to the **Peer Management Center Web Service** and the **Peer Management API Service**. The web service allows users to access Peer Management Center via a web browser; the API service allows access to Peer Management Center through REST API calls. For detailed

instructions on configuring access to these services, see Configuring the Web and API Services.

3. On the final page of the installation wizard, you have several options; we recommend that, at minimum, you select the first option.

   If you enabled the Peer Management Web Service and selected the **Launch the Peer Management Center Web Client**, on the final page of the installation wizard, the default username and password for accessing the web client is displayed. After logging in to the web client, you should change the password immediately.

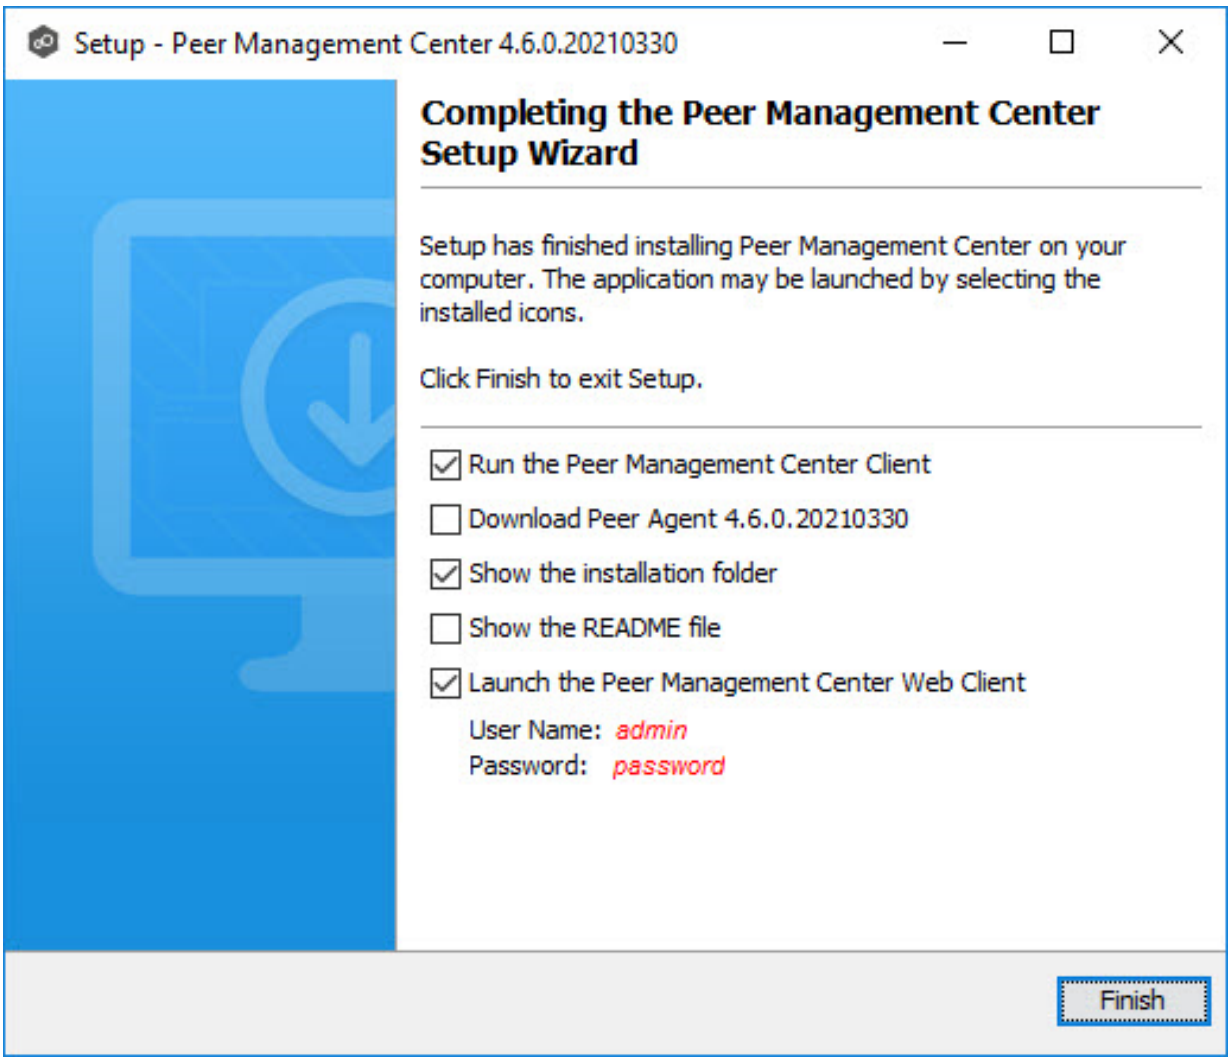

When the installation is complete, the Peer Management Center installation folder contains the following files and folders:

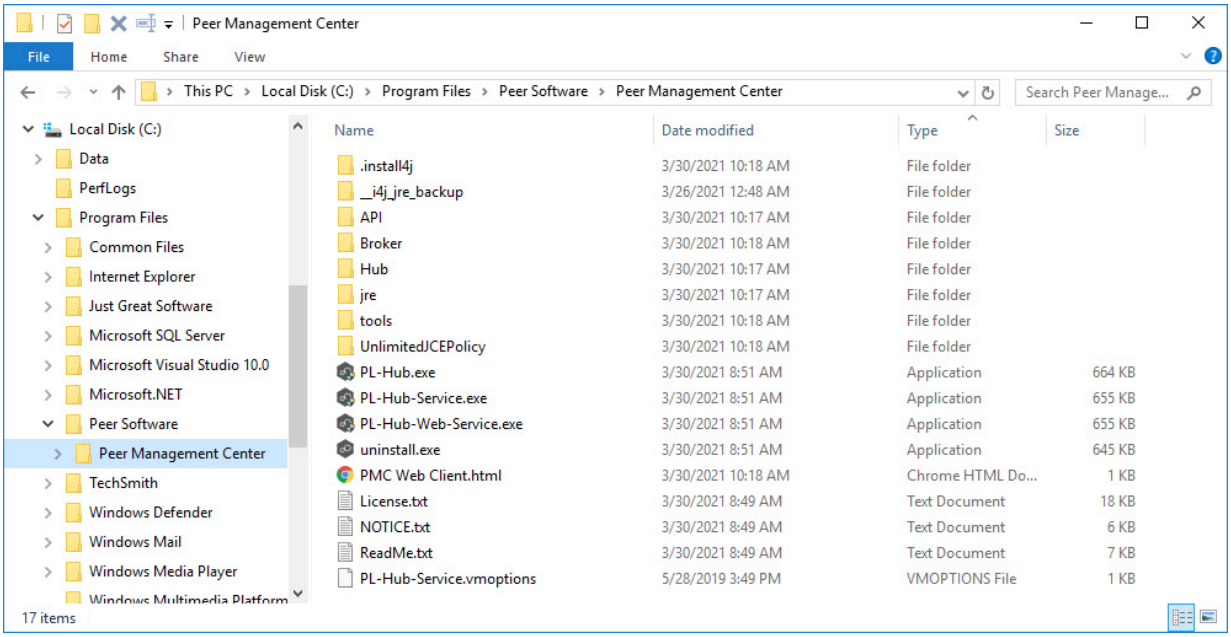

The **PL-Hub.exe** executable launches **Peer Management Center Client**, which is a Windows rich client application.

Four Windows services have been installed and are set to auto-start:

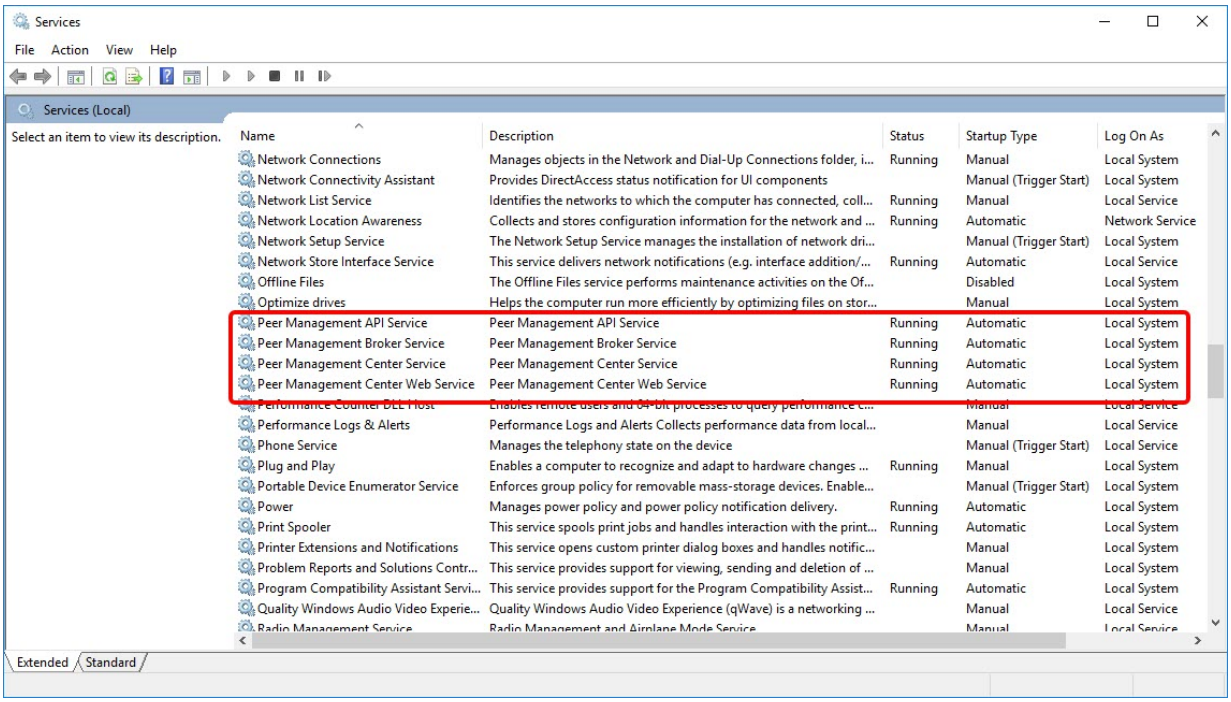

4. If you didn't select the option to launch Peer Management Center Client on the last page of the installation wizard, launch it using one of the following methods:

- Select **Peer Management Center** from the Windows **Start** menu.

- Double-click the **PL-Hub.exe** executable in the Peer Management Center installation directory.

If both the **Peer Management Center Service** and the **Peer Management Broker Service** are up and running as background services, then Peer Management Center should successfully start. If not, open the Windows Service Panel (services.msc) and start both services.

5. When launching Peer Management Center Client for the first time, you are prompted to install your license. If you haven't already done so, copy the license to the desktop of the Peer Management Center server.

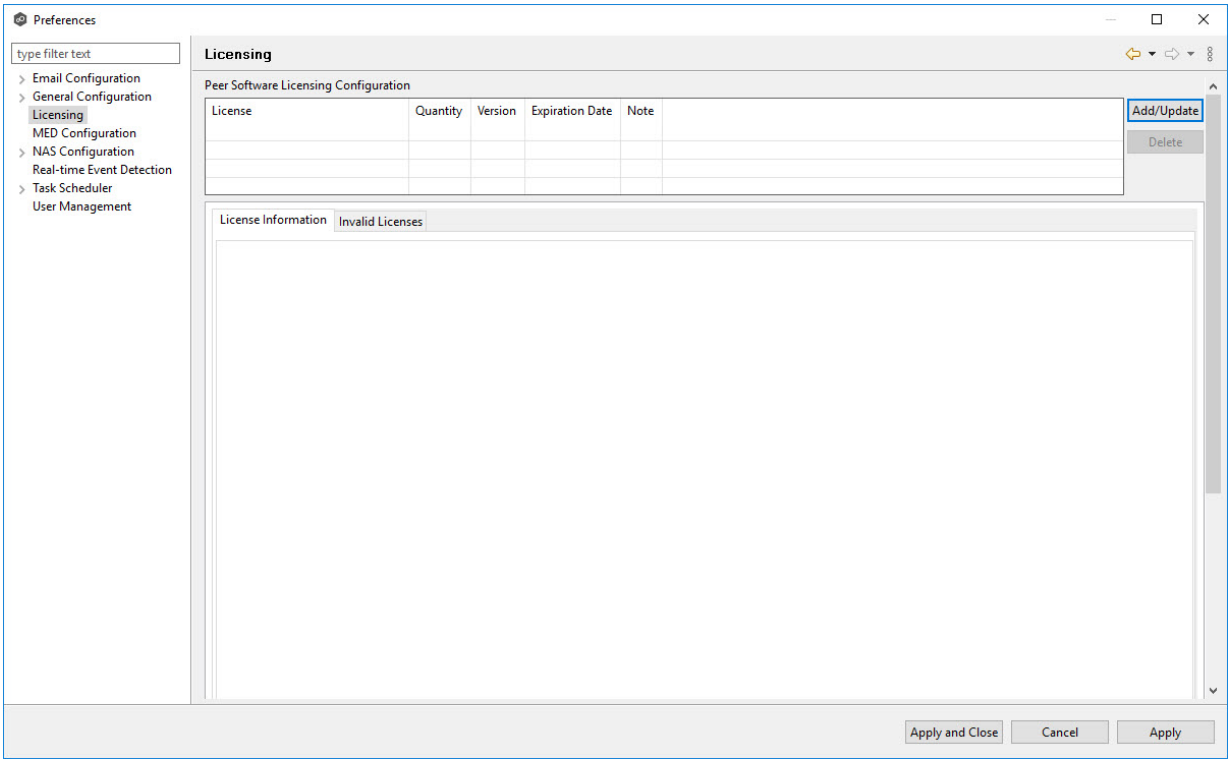

6. Click **Add/Update**.

7. Browse to where the License file is located and select the file.

8. Click **Open**.

The **License Information** tab displays your license information.

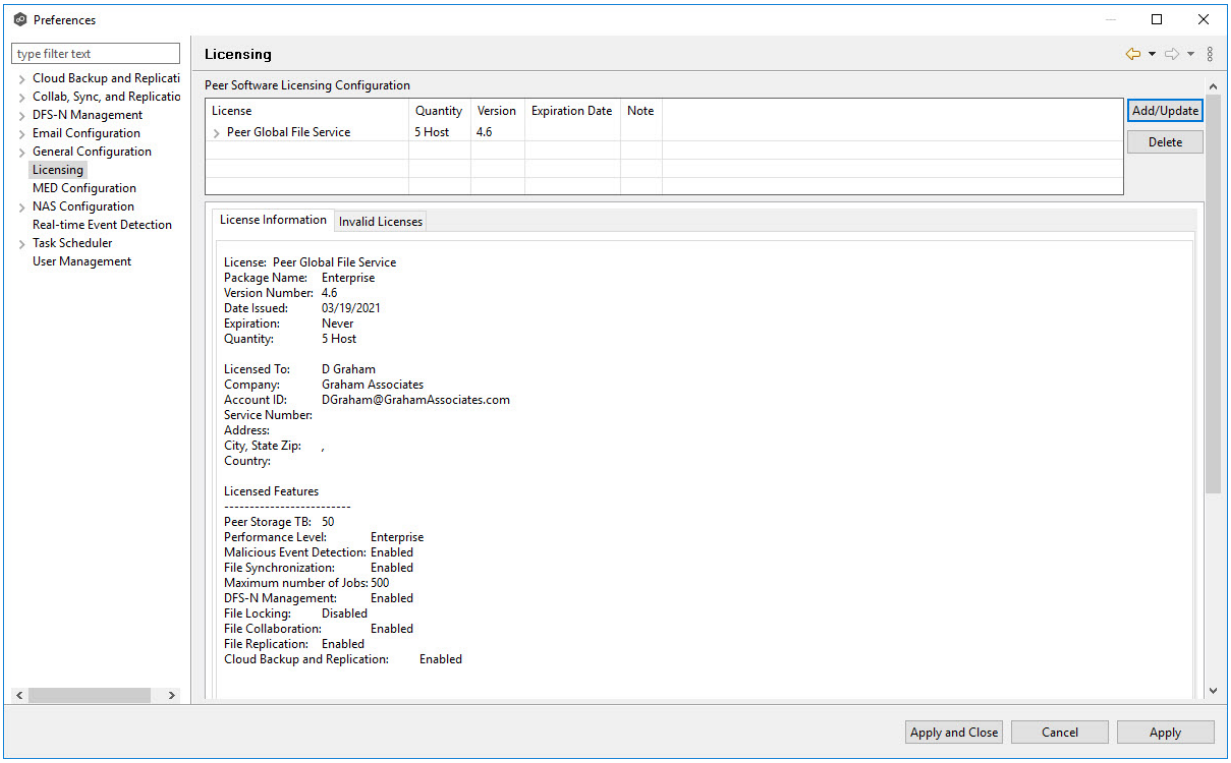

9. Click **Apply and Close**.

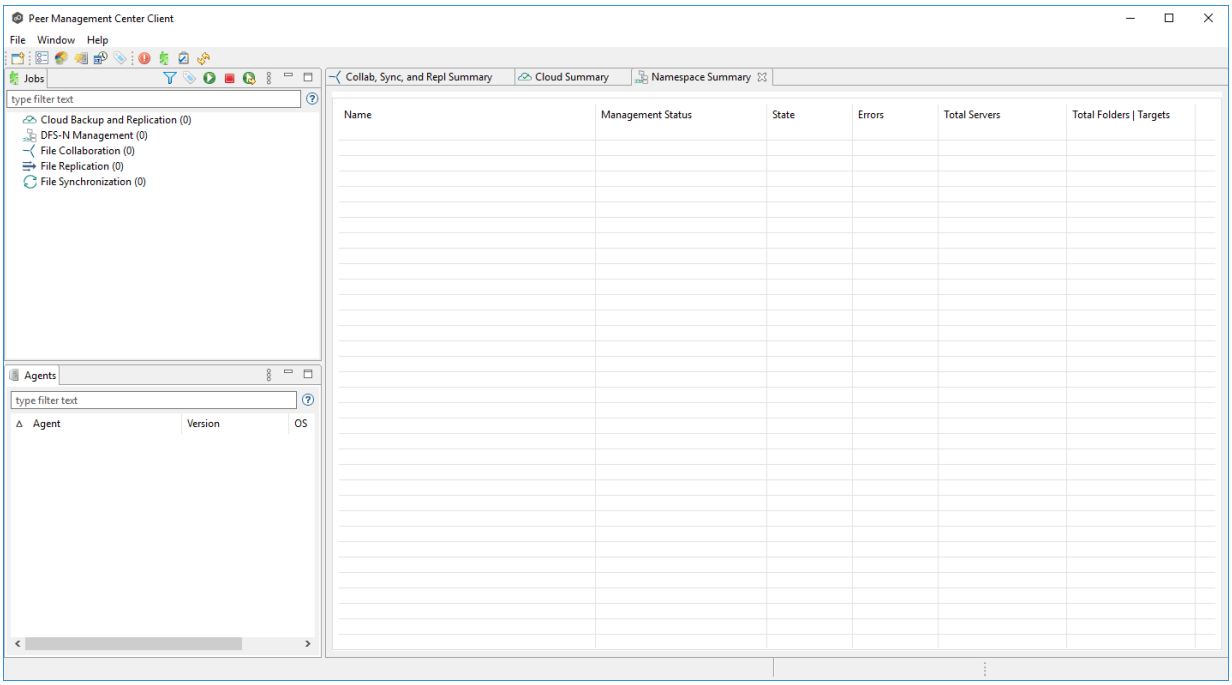

Now you are ready to [install the Peer Agents](#).

Uninstalling Peer Management Center

Peer Management Center ships with an uninstaller for the environment it is running in. Please use the standard platform-specific method for removing programs/applications to uninstall Peer Management Center.

Configuring and Managing the Web and API Services

As part of the initial installation of Peer Management Center, you are prompted to configure access to the web and API services. The web service allows users to access Peer Management Center via a web browser; the API service allows access to Peer Management Center through REST API calls.

If you enable access to the web client, you will need to secure access to the web client and manage web client user accounts.

See the following topics for more information:

- Configuring the Web and API Services
- Securing Access to the Web Client
- Accessing the Web Client
- Managing Web Client Users

You can configure access to the web and API services during the initial installation of Peer Management Center. If you do not enable access during the initial installation or want to modify settings at a later date, you can modify them through Web and API Configuration in Preferences.

Configuration Options

To configure the web and API services during the initial installation of Peer Management Center:

1. Enter the requested values when the configuration page appears in the installation wizard.

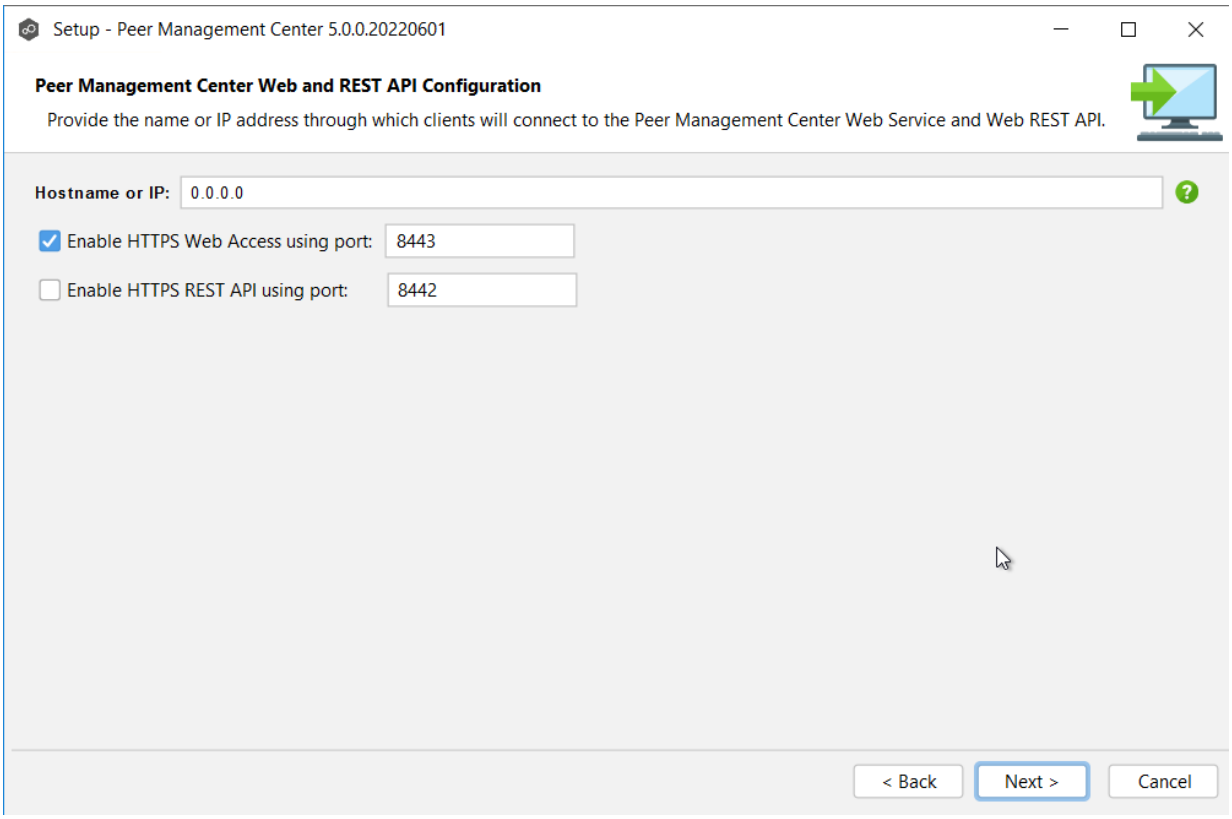

Field	Description
Hostname or IP	Enter the hostname or IP address via which the services can be accessed: • Enter localhost or 127.0.0.1 if you want the services to be accessible only to users of the local server via the loopback interface. • Enter 0.0.0.0 to make the services accessible via all network interfaces. • Enter a specific IP address to restrict access to a specific network interface.
Enable HTTPS Web Access	Select this to enable secure access to the web service, and then enter a port number.
Enable HTTPS	Select this to enable secure access to the REST API service, and then enter a port number. The REST API port cannot be the same as the web

Field	Description
REST API	service port.

2. Click **Next** to continue with the rest of the installation wizard.

Access to Peer Management Center's web client is through HTTPS, which ensures that all communication between the browser and the service hosting the web client is encrypted. However, you may want to take additional actions to secure access to Peer Management Center's web client:

- You can limit users' access to the web client when you configure the hostname or IP address for web access. For example, enter **localhost** or **127.0.0.1** if you want the web client to be accessible only to users of the local server via the loopback interface. See Configuring the Web and API Services for more details.
- While HTTPS access to the web client is secured out of the box with a built-in Transport Layer Security (TLS) certificate, this certificate can be swapped for a custom one. See TLS Certificates in Advanced Topics for information on using existing certificates and creating new certificates.
- The default password for the admin account should be changed immediately. See Editing an Internal User for information about changing the password.

Installing Peer Agents

Overview

You will need to install a Peer Agent on each server you plan to include in any of your jobs. After installing the Peer Agent software, you should verify that the **Peer Agent Service** is running and can successfully connect to the Peer Management Broker.

Note: For customers using clustered file server roles with Windows Failover Cluster, please review this Peer Software knowledge base article: Windows Failover Cluster support for the Peer Agent.

Installing a Peer Agent

To install a Peer Agent and verify its connection to Peer Management Broker:

1. Download the Peer Agent installer (**P-Agent_Installer_win64.exe**) to the server you want to host the Agent.

2. Run the installer and follow the wizard instructions.

   During installation, you will be prompted for:

   - The Peer Management Broker hostname (computer name, fully qualified domain name, or IP address) of the server where Peer Management Broker is running.

   - The TCP/IP port number of the server where Peer Management Broker is running. The default port for TLS communication is 61617.

   Enter the same values that you entered when installing Peer Management Center and Peer Management Broker.

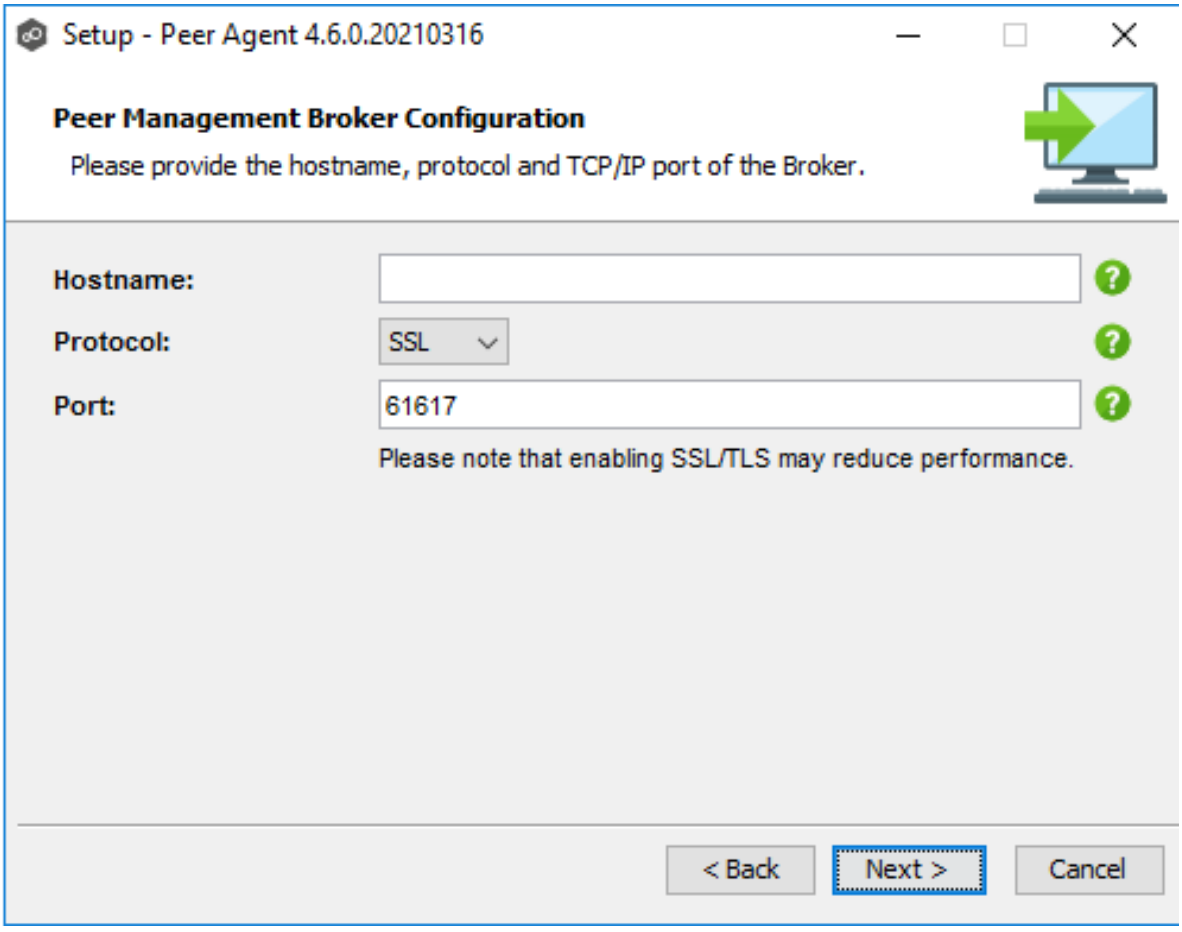

You will also need to provide the account credentials under which the Peer Agent Service will run.

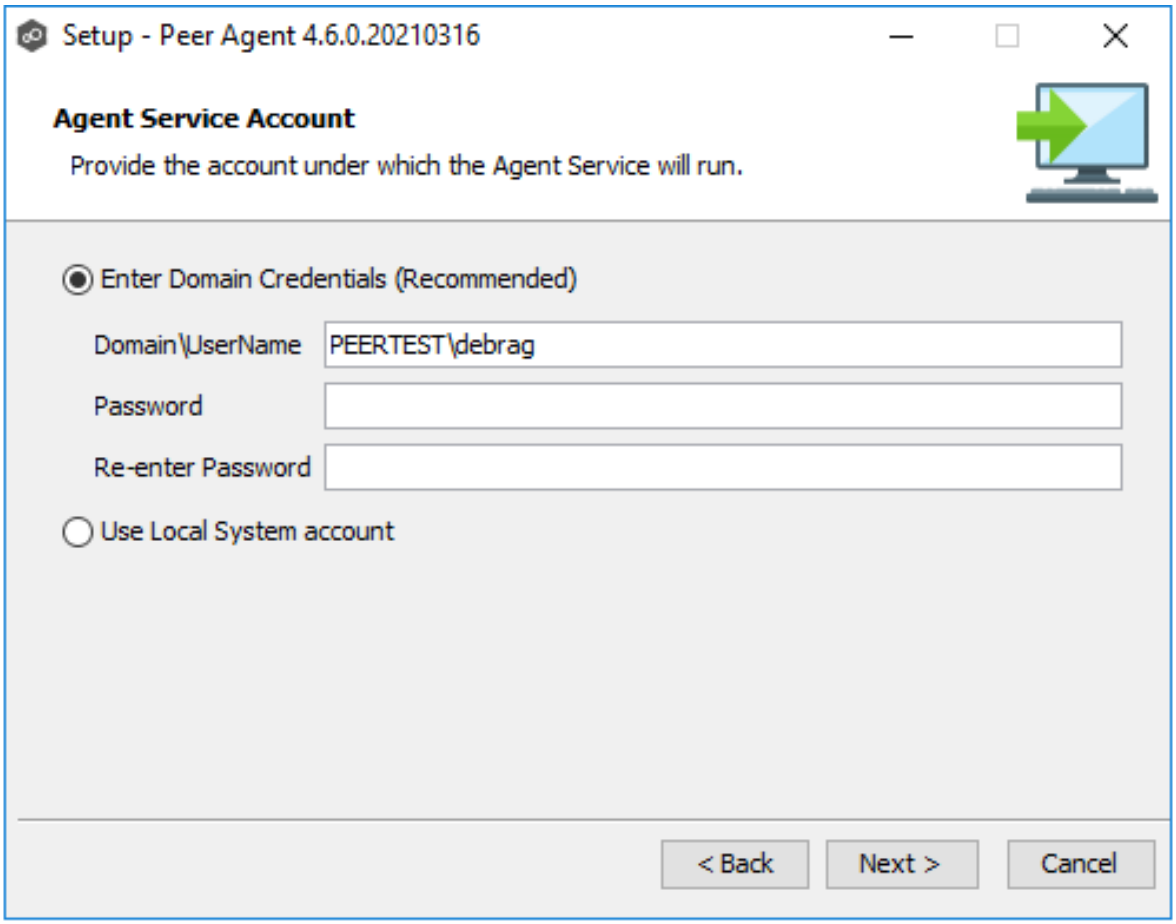

3. When the last page of the installation wizard appears, click **Finish**.

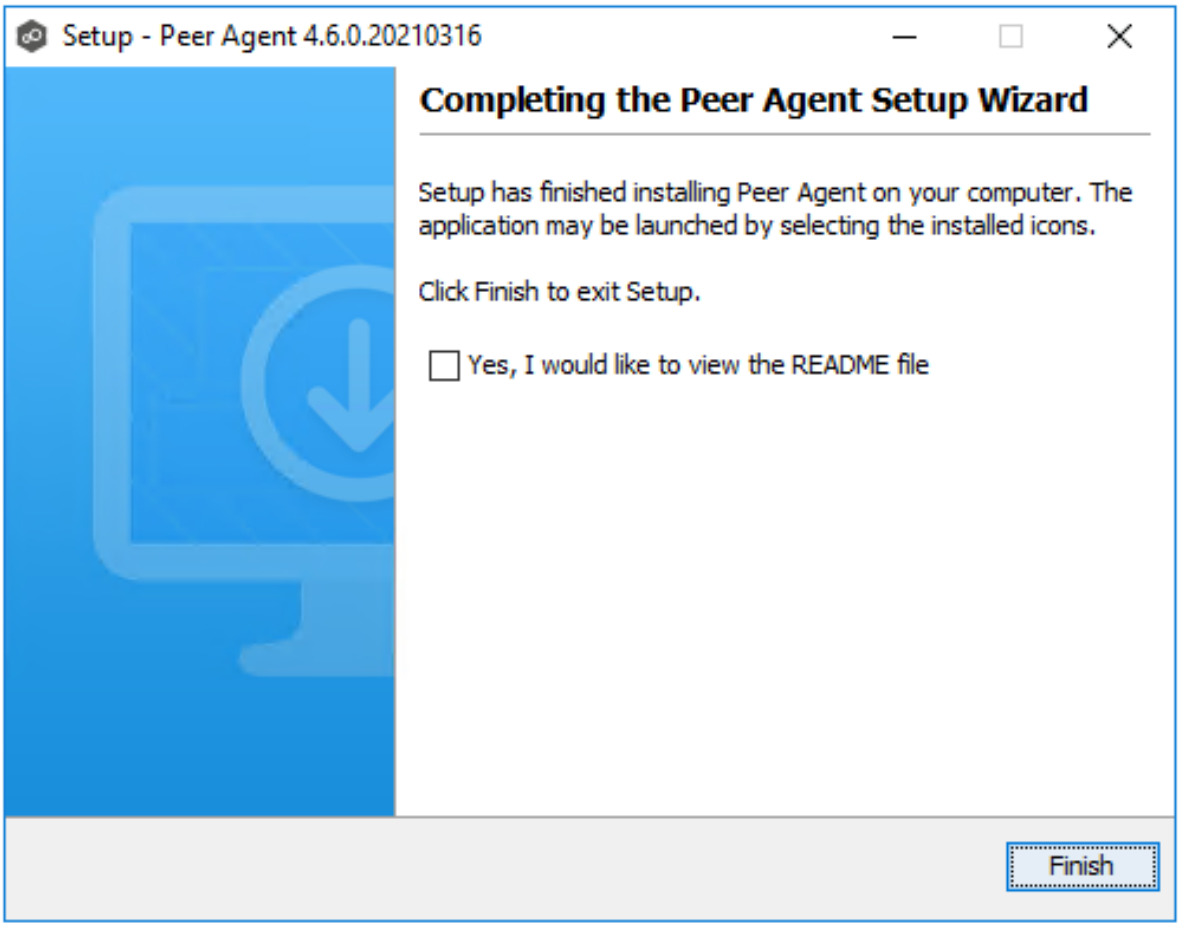

4. After the installation finishes, the Peer Agent is installed as a Windows service. You will need to verify that the **Peer Agent Service** is running, and that it was able to successfully connect to Peer Management Broker. You can do this by opening the Windows Services Panel (services.msc) and verifying that the **Peer Agent Service** has started.

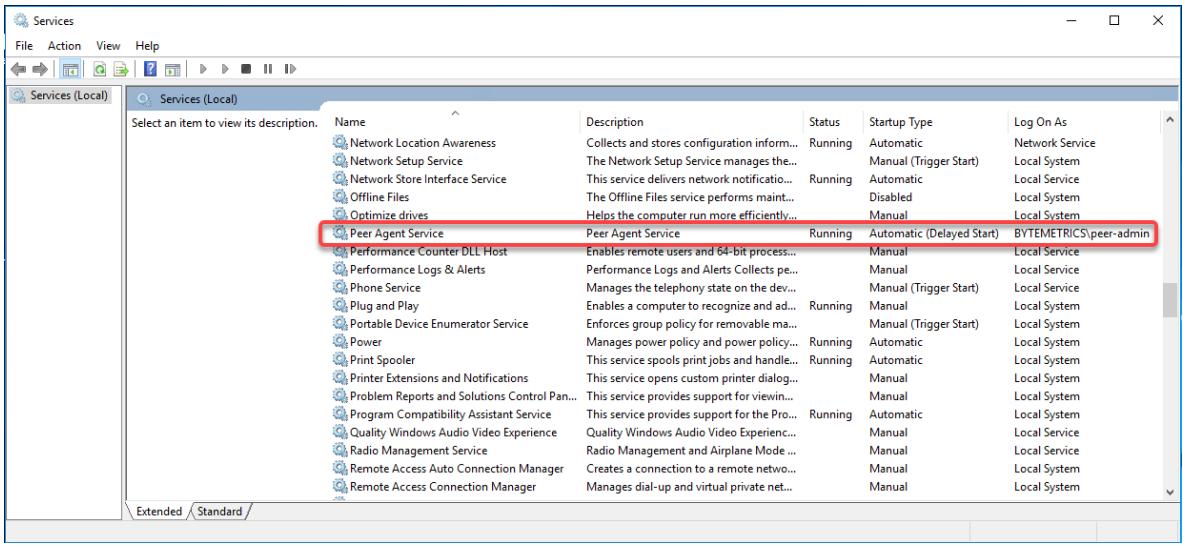

Secure Encrypted TLS Connections

By default, the Peer Agent is installed with Transport Layer Security (TLS) encryption enabled, where the Peer Agent connects to Peer Management Broker through a secure, encrypted connection. If you are running Peer Management Center on a secure LAN or via a corporate VPN, you might want to disable TLS to boost performance. For more details on disabling or enabling encryption for the Peer Agent, see Broker Configuration.

Uninstalling a Peer Agent

Peer Agent ships with an uninstaller for the environment it is running in. Please use the standard platform-specific method for removing programs/applications to uninstall the Peer Agent.

Updating Peer Management Center and Peer Agents

You can easily check for updates to the Peer Management Center software. Minor releases can be automatically downloaded and installed. Major releases require a new license key and must be requested from Peer Software Support.

For details about updating, see:

- Updating Peer Management Center
- Updating Peer Agents

Updating Peer Management Center

Overview

There are two ways to check for software updates:

- You can manually check for software updates using the **Check for Updates** command on the **Help** menu. **Note:** This command is not available in the Peer Management Center Web Client.
- You can also configure Peer Management Center to automatically check for updates and download the updates. For more information, see the Software Updates setting in Preferences.

Note: For customers running the PMC on a Windows Failover Cluster, please review this knowledge base article: https://kb.peersoftware.com/tb/windows-failover-cluster-support-for-the-pmc. The steps for upgrading the PMC on a Windows Failover Cluster are the same as installing the PMC for the first time.

Manually Checking for and Installing an Update

To manually check for an update:

1. From the **Help** menu, select **Check for Updates**.

   The **Check for Updates** dialog appears. If a minor update is available, the dialog identifies the new version (and your current version) and provides a link to the release notes. If a major update is available, the dialog presents a link to an announcement page on the Peer Software website.

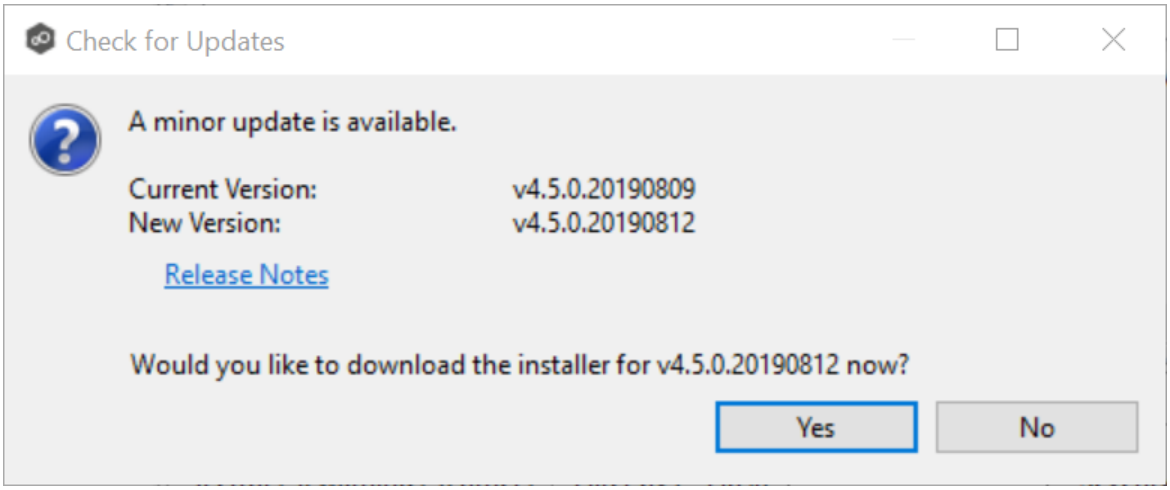

2. Click **Yes** to download the Peer Management Center installer.

As the update is downloaded, a progress bar appears in the lower right corner of the Peer Management Center window. After the download is complete, the **Check for Updates** dialog displays information about the upgrade process.

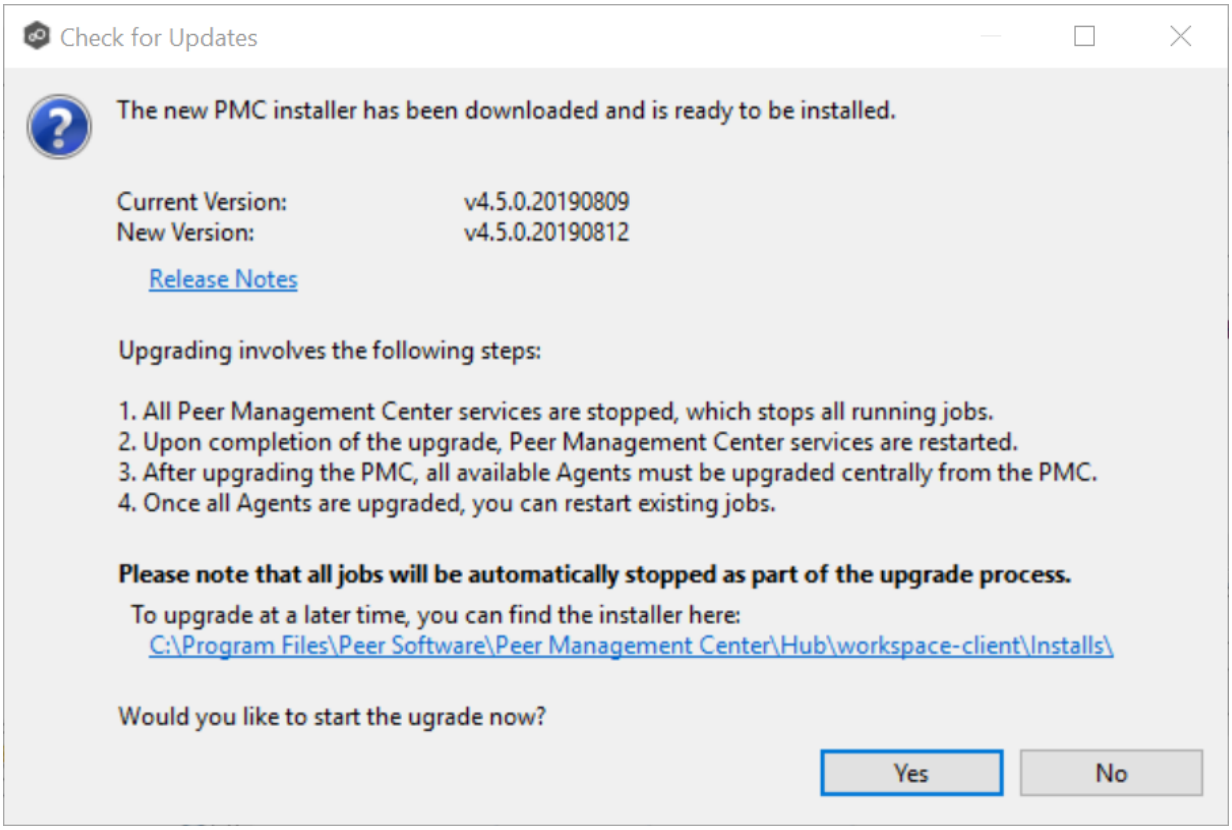

3. Click **Yes** to install the upgrade; click **No** to install the update at a later time.

   If you clicked **No**, you can install the update later by going to the folder shown in the dialog.

   If you clicked **Yes**, the **Setup** wizard appears.

4. Follow the prompts in the **Setup** wizard to install the update.

   When updating a Peer Management Center installation, you will not be prompted to specify web and API access. The settings entered previously will be used. If you wish to change those settings, you can do so by modifying them in Web and API Configuration in Preferences.

5. After the Peer Management Center upgrade is installed, update the Peer Agents. See Updating Peer Agents for details.

Updating Peer Agents

You can view the status of your Peer Agents in the Agents view. Whenever you update Peer Management Center software, you need to update the Peer Agent software before you can start any jobs managed by that Agent. When **Update Required** appears next to an Agent's name, that indicates the software needs updating.

Note: For customers using clustered file server roles with Windows Failover Cluster, please review this knowledge base article: https://kb.peersoftware.com/tb/windows-failover-cluster-support-for-the-peer-agent. The steps for upgrading Agents tied to clustered file server roles is the same as installing these Agents for the first time.

To update Peer Agents:

1. Select the Agents in the **Agents** view.

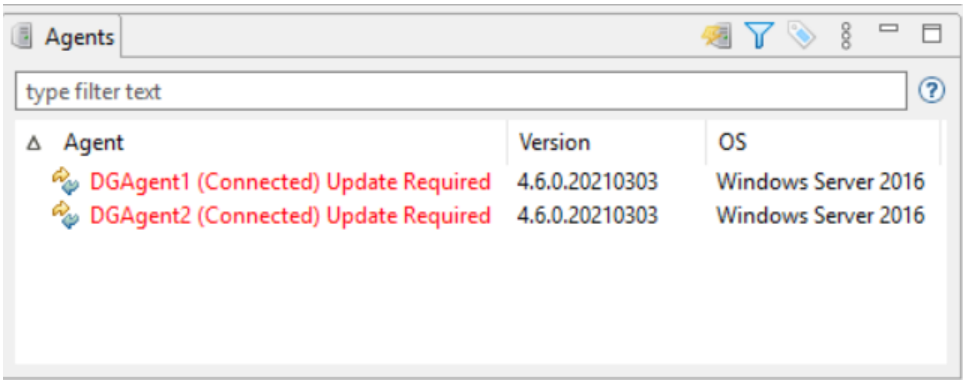

2. Right-click and select **Install Software Updates**.

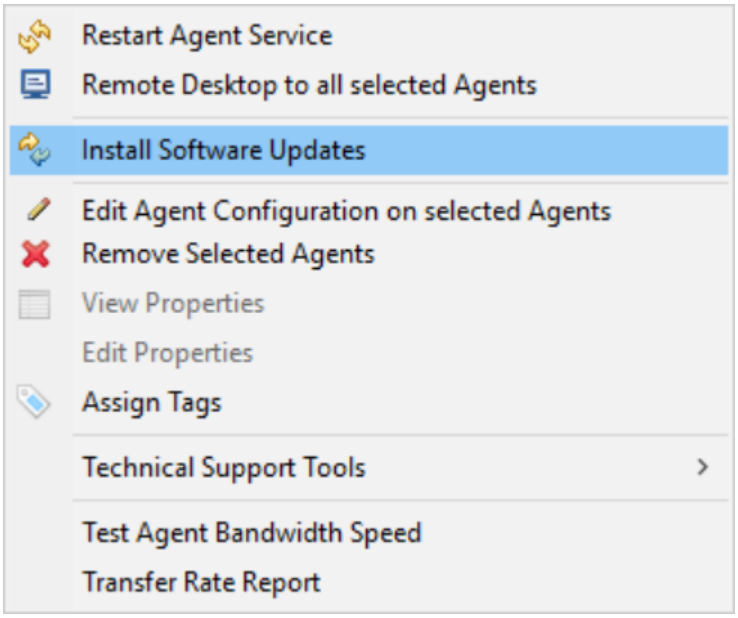

A confirmation dialog appears.

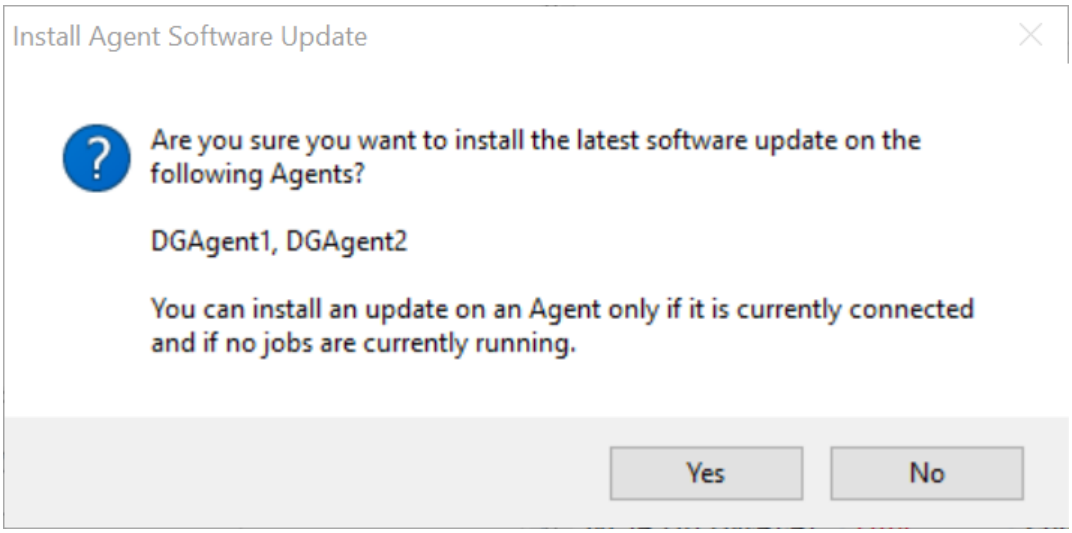

3. Click **Yes**.

4. Follow the prompts in the **Update Agent Software** dialog to complete the update.

   After the Agents are updated, the Agents appear in green. The Agents automatically restart as part of the upgrade. Any jobs set to auto-start will restart once the Agents have reconnected.

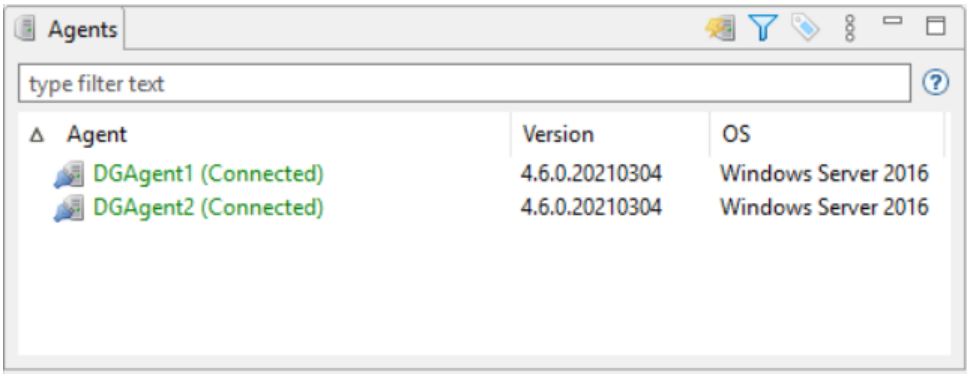

Peer Management Center User Interface

Peer Management Center is a management interface for configuring and deploying jobs, as well as view summary and runtime information for jobs. It offers two graphical user interface options:

- A **rich client** installed and run on the server running Peer Management Center.

- A **web client** that, when configured, can be accessed from remote systems via a web browser. You can manage and monitor jobs via the robust Peer Management Center web client. Unlike many other web management consoles, Peer Management Center's web client is very responsive and is built to mirror the functionality of the rich client (which is included with the Peer Management Center installer for use by system administrators). When configured, the web client allows for the management of Peer Management Center from remote clients without the need to directly log into the Peer Management Center server.

The interface can be divided into four quadrants: each quadrant displays information in panels called views. A view can contain one or more tabs. See Views for more information about the views.

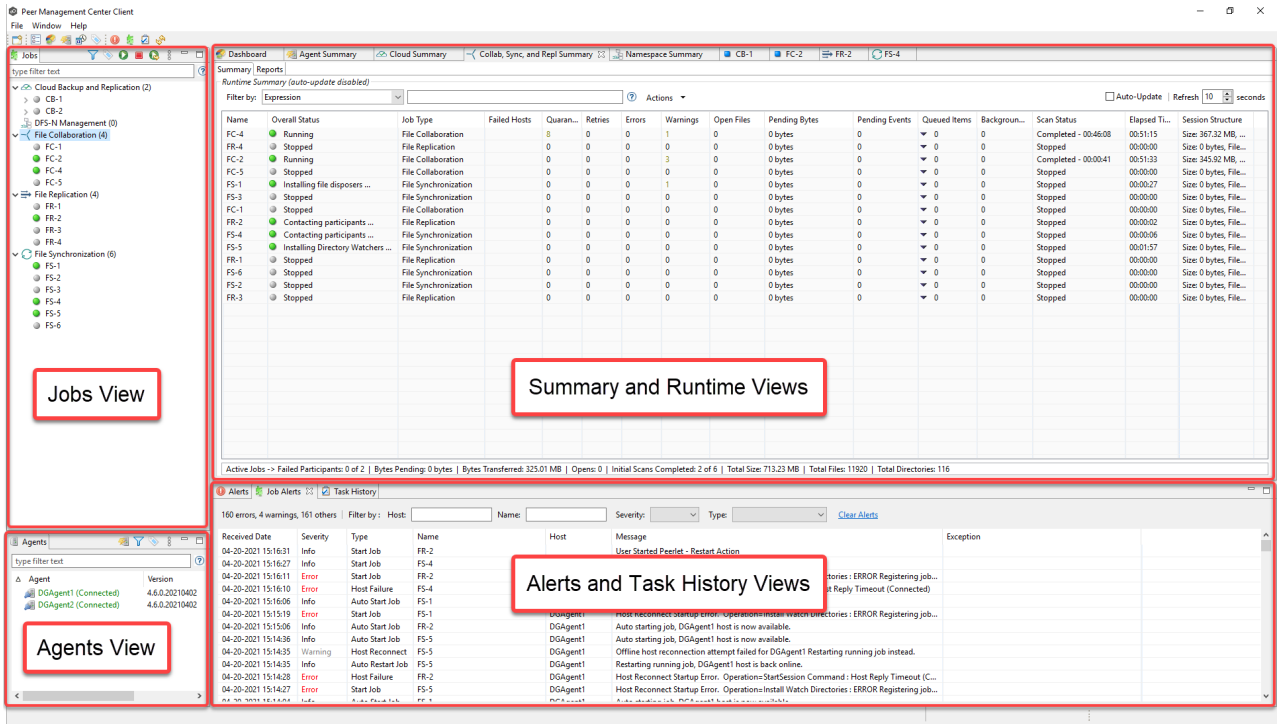

Description of Quadrants

The quadrants are described in the following table.

Quadrant	Description
Upper right	Contains one view, the Jobs view, which displays a list of all jobs, grouped by job type. The toolbar in this view allows you to start and stop jobs.
Bottom right	Contains one view, the **Agents** view. The Agents view displays a list of known Peer Agents and connection status for each. Individual Peer Agents can be updated and restarted from this view as well by right-clicking on one or more items and selecting the appropriate item from the pop-up menu.
Upper right	Several types of views are displayed in this area, including: • A dashboard that provides metrics and key performance indicators. • Summaries of jobs by job type. • An Agent Summary view, which displays a list of all known Peer Agents deployed and detailed status information that can be used to assess the health of the environment. • Runtime statistics for individual jobs.
Lower left	Contains a variety of views, including: • The Alerts view, which displays a list of Peer Management Center alerts that have occurred with detailed information about each alert. Alerts relating to Peer Agent connection status changes are reported here. • The Jobs Alerts view, which displays a list of job-specific alerts that have occurred. Alerts relating to the automatic stopping and restarting of jobs are displayed here.

For information about other aspects of the user interface, see:

- Accessing the Web Client
- Views
- Main Window Menus and Toolbar
- Tables

Accessing the Web Client

Once Peer Management Center has been installed, the Peer Management Center Web Client Service has been configured, and the necessary Peer services have been started, users can access the Peer Management Center web client in various ways:

- Log in directly through a web browser on the local Peer Management Center server.
- Remote access from another system on a network that can reach the Peer Management Center server.
- Start the Peer Management Center Client (rich client application) and select **Peer Management Center Web Client** from the **Help** menu.

Logging into the Web Client

To access the Peer Management Center web client:

1. Open a web browser.
2. Enter one of the following URLs in the address field:

If this URL was entered in the Installation Wizard	Use this URL
A specific IP address	Enter **https://** followed by that IP address and **:8443**. You cannot use **localhost** even if you are directly logged into the Peer Management Center server. For example: **https://10.0.0.1:8443**
localhost or **127.0.0.1**	Enter **https://localhost:8443** or **https://127.0.0.1:8443**.
0.0.0.0	Enter **https://** followed by the IP address of the Peer Management Center server and **:8443** For example: **https://10.0.0.1:8443**

Notes:

- The URL will depend on how the web service was configured during the installation process; you may need to contact the administrator who installed Peer Management Center to determine the correct URL
- 8443 is the default HTTPS port and should be replaced with the port used in your environment if it is different.
- You can modify the web service configuration on the General Configuration page of Preferences. For more information, see Web and API Configuration.

After you enter the URL, the Peer Management Center Home Page appears.

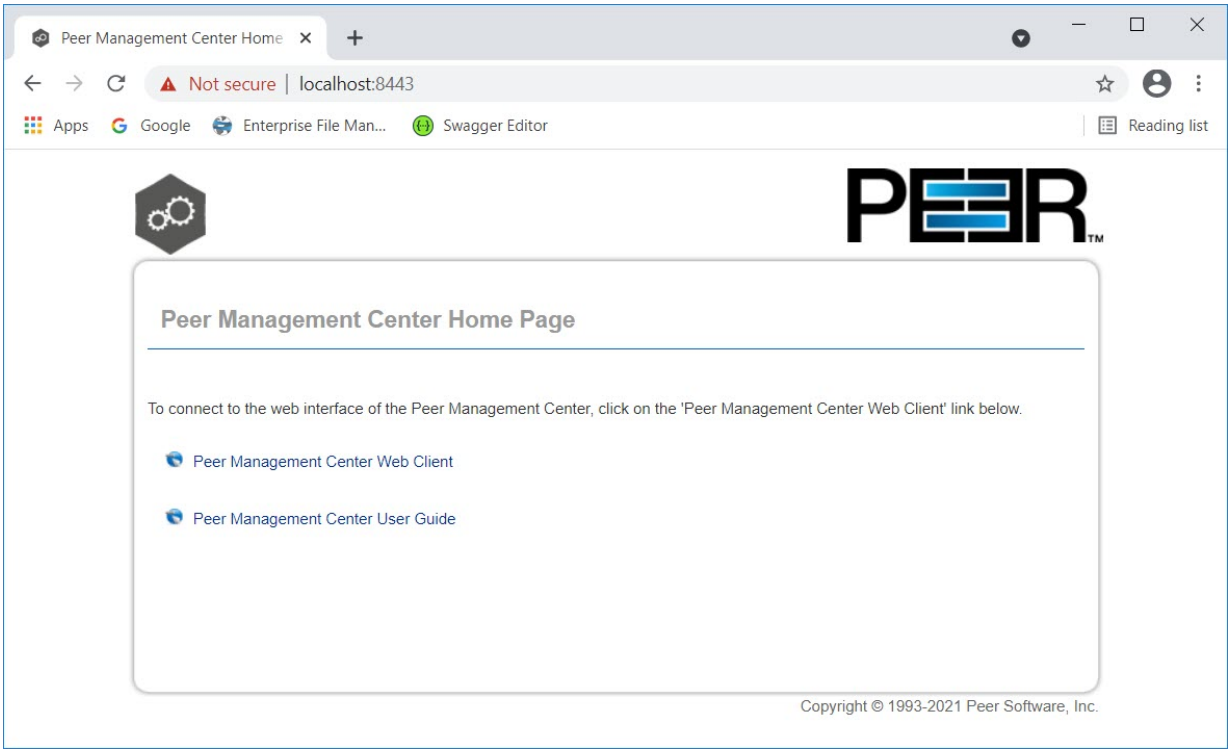

3. Select the **Peer Management Center Web Client** link.

   The login page is displayed.

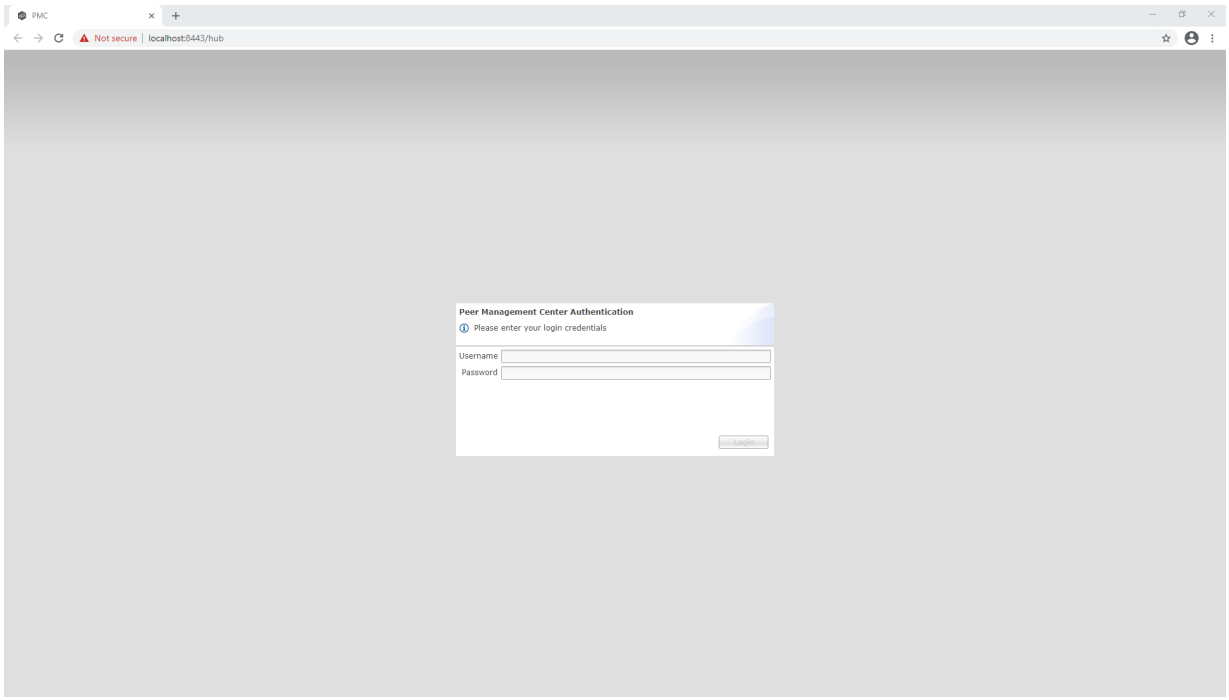

4. Enter a user name and password.

   - The default user name is **admin**; the default password is **password**. For security reasons, we highly recommend that the user immediately changes the **admin** password. See Editing an Internal User for more information on changing account passwords.

   - If logging in with an Active Directory account, enter the user name in this format: username@mydomain.local.

5. Click **Login**.

Peer Services Required for Web Client

To use the Peer Management Center web client, the following Peer services must be running on the Peer Management Center server:

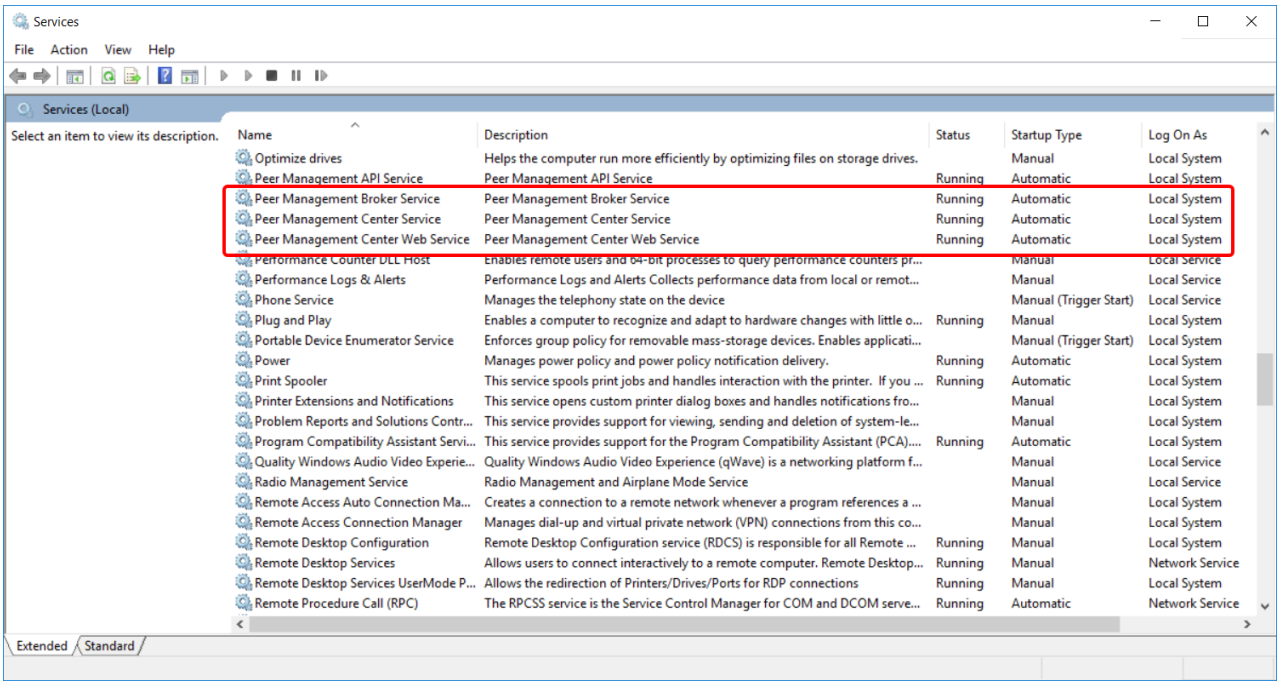

If a required service is not running, open the Windows Service Panel (services.msc) on the applicable PMC server and start the service.

Views

The Peer Management Center interface can be divided into four quadrants; each quadrant displays information in panels called views. A view can contain multiple tabs. There are various types of view. For example, some views display a combination of real-time file I/O activity, history, and configuration information for a specific job; others display a summary of information about all jobs of a specific type.

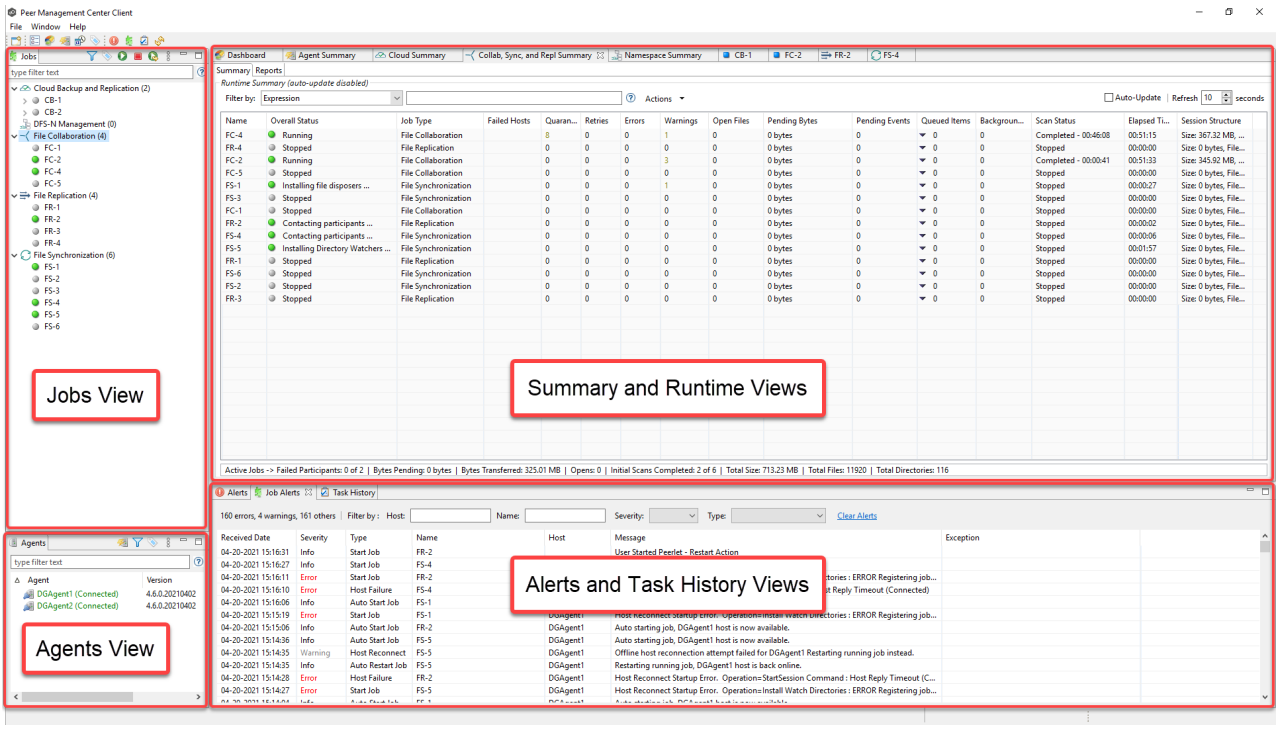

The primary views include:

- Agents View
- Jobs View
- Dashboard
- Alerts View
- Job Alerts View
- Summary Views
  - o Agent Summary View
  - o Cloud Summary View
  - o Collab, Sync, and Repl Summary View
  - o Namespace Summary View
- Runtime Views
  - o Cloud Backup and Replication Job Runtime View

- DFS-N Management Job Runtime View
- File Collaboration Job Runtime View
- File Replication Job Runtime View
- File Synchronization Job Runtime View

Displaying Views

You can open views in a variety of ways:

- Selecting a command from the **Window** menu
- Right-clicking on an item to display a context menu.
- Clicking the **View** button in a toolbar and selecting an option from the View menu.

Resizing Views

You can resize views in a variety of ways:

- Drag the separator between views.
- Click the minimize or maximize button in the toolbar.
- Reset all views to the default size by selecting the **Reset Perspective** command on the **Window** menu.

Agents View

The **Agents** view is displayed in the lower left quadrant of the Peer Management Center interface and lists all known Peer Agents installed in your environment and displays the current connection status for each. For more information, see Agent Connection Statuses. This view is automatically displayed when Peer Management Center is started.

To filter a large list of Agents, use the **Filter** field located below the Agents view toolbar. For more details on how to filter agents, see List Filters.

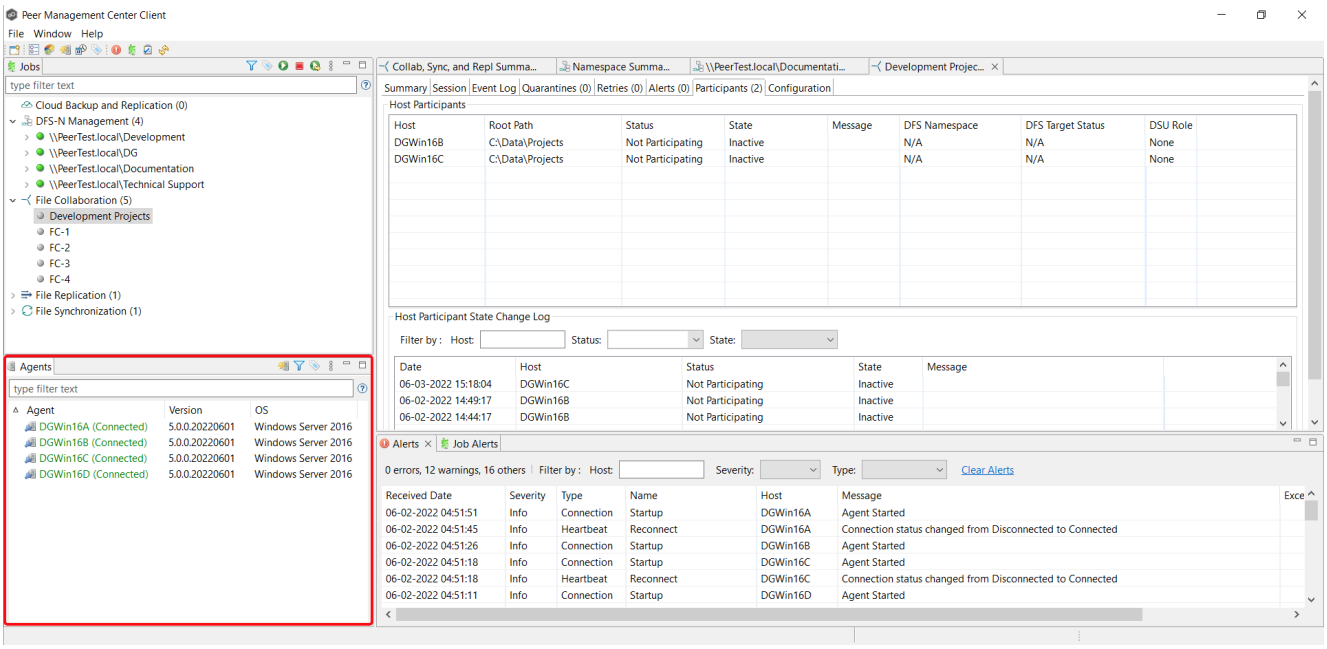

Updating Peer Agent Software

If the Peer Agent software running on a host is out of date, the host will be shown as having a pending update in this view. When right-clicking the Agent, the option to automatically update the Peer Agent software will also be available. You can update directly from the Peer Management Center; updating usually does not require any additional actions on the host server itself. See Updating Peer Agents for more information.

The following buttons are available on the toolbar in the Agents view:

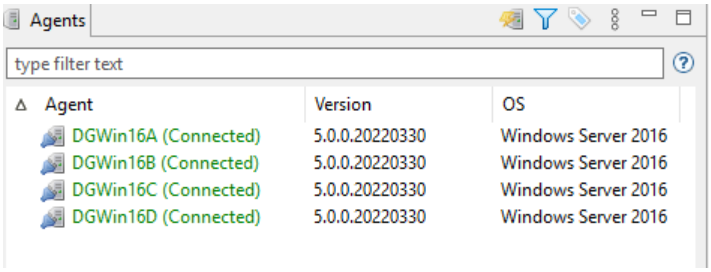

Button	Description
Show Agent Summary	Opens the Agent Summary view, which provides details for all known Agents and their status.
Manage, Save and Load Filters	Allows for the selection of predefined or user-defined filters and to save and manage filters. Default Agent filters include **Connected** and **Disconnected**.
Assign Tags	Opens the **Assign Tags** dialog where resources can be viewed and assigned to tags or categories. Tagging resources helps when managing large number of resources.

Alerts View

The **Alerts** view is displayed in the lower right quadrant of the Peer Management Center interface and displays various system alerts. The **Alerts** view is automatically displayed when a critical system alert (Error or Fatal) is received. You can also set the Alerts view to be automatically displayed when Peer Management Center is started.

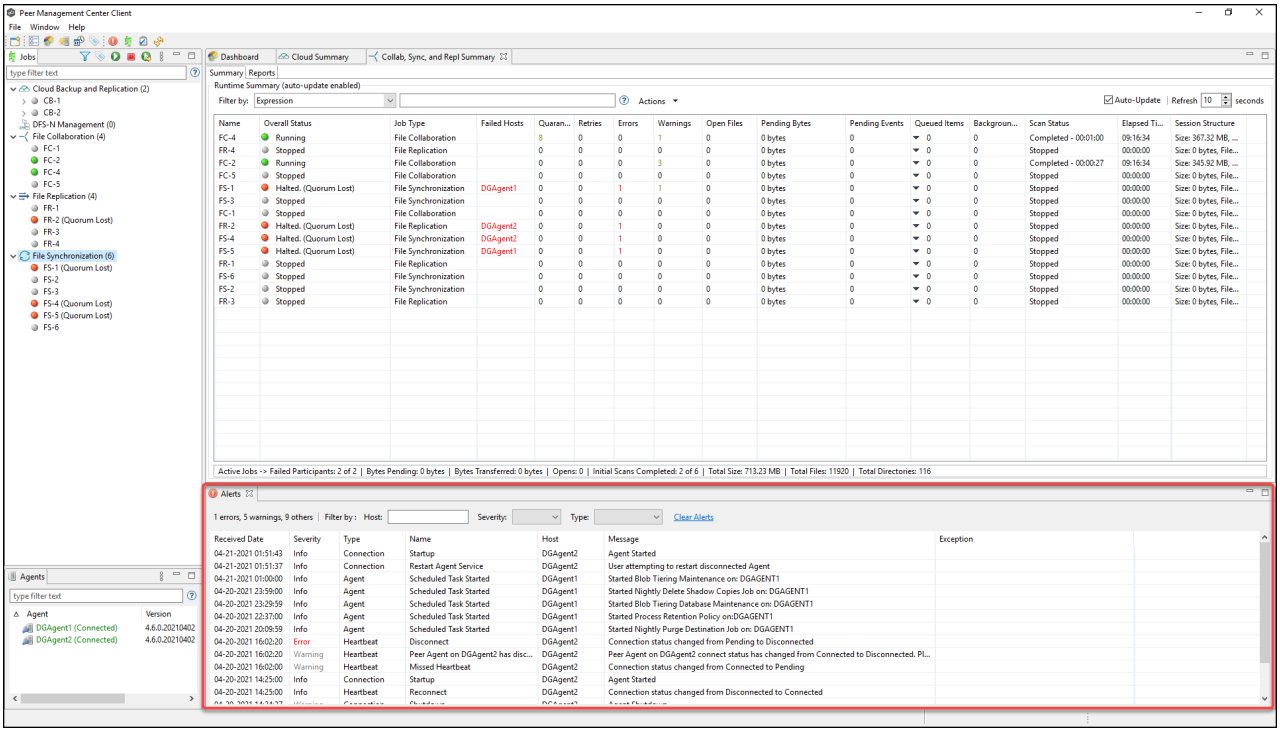

System alerts vary in their severity. The four categories of alerts are:

- Informational (containing Info, Debug, and Trace)
- Warning
- Error
- Fatal

An example of an Informational alert is when a Peer Agent connects to the Peer Management Broker. If a Peer Agent's network connection is severed, then an Error alert will be logged. All alerts are also logged to the file **hub_alert.log**, available under the **Hub\logs** subdirectory within the Peer Management Center installation directory.

You can filter alerts based on host name, severity level, or type, and you can sort alerts by clicking on a specific column header. You can also clear all alerts in the table by clicking the **Clear Alerts** link.

Displaying the Alerts View

You can open the **Alerts** view at any time by clicking the **View Alerts** button located on the Peer Management Center toolbar or by selecting **View Alerts** from the **Show View** submenu of the **Window** menu. You can close the **Alerts** view at any time by clicking on the **X** (Close) button on the **Alerts** tab.

You can resize the Alerts view by dragging the separator between the upper view and the Alerts view, or you can double-click the **Alerts** tab to maximize the view. You can restore the view to its original, non-maximized size by double-clicking the **Alerts** tab again.

Dashboard

The Dashboard is divided into two sections:

- **Collab, Sync, and Repl** - This top section displays a table of metrics and key performance indicators for all running File Collaboration, File Synchronization, and File Replication jobs. It also contains a link that opens the Collab, Sync, and Repl Summary view. Entries in the table's first column can be double-clicked to display a filtered runtime view of the selected item for additional details.
- **Agents** - The bottom section displays information about Agents. It also contains a link that opens the Agent Summary view.

Click the triangle to the left of the section name to collapse and expand the section.

For performance reasons, the Dashboard is not updated in real-time. However, you can set the table to be automatically updated every few seconds by selecting the **Auto-Update** checkbox, and then choosing the update interval.

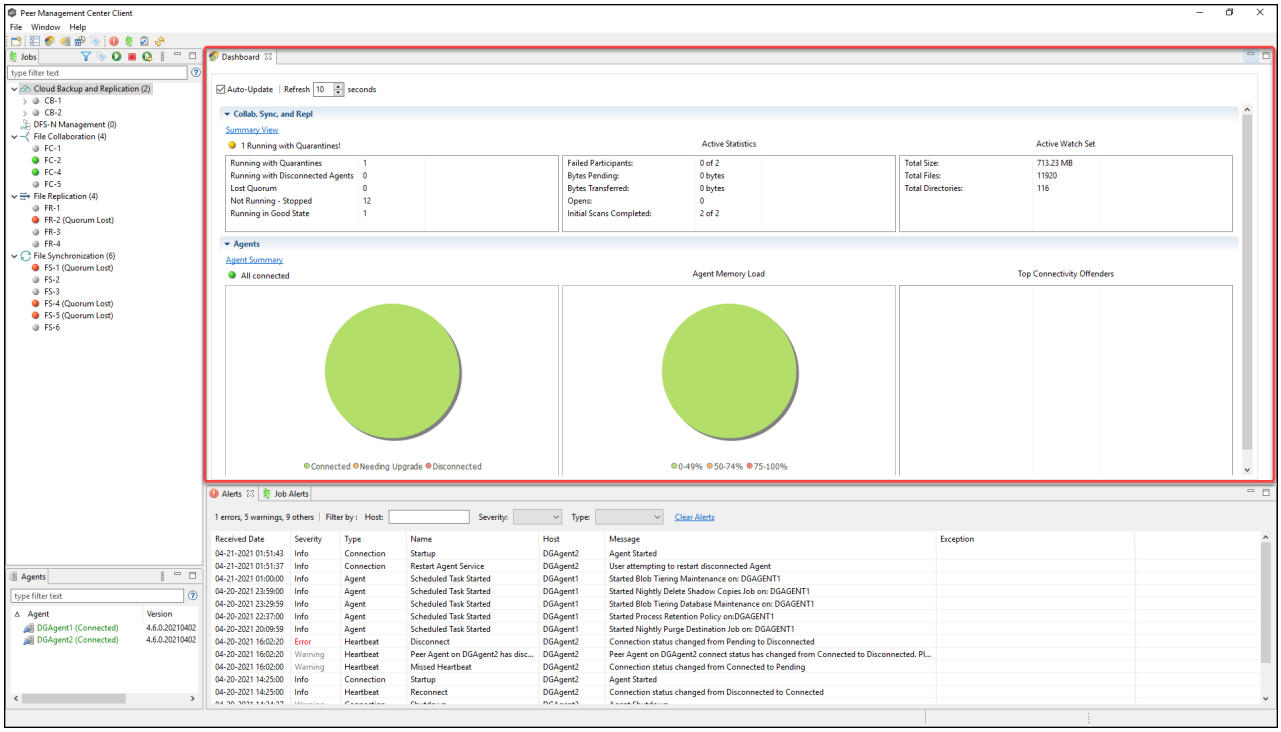

To display the Dashboard, use one of the following methods:

- Select **Show Dashboard** from the **Window** menu.
- Click the **Show Dashboard** icon in the main Peer Management Center toolbar.
- Set the Dashboard to launch automatically at start. See General Configuration.

Job Alerts View

The **Alerts** view is displayed in the lower right quadrant of the Peer Management Center interface and displays various system alerts. The **Job Alerts** view is automatically displayed when a critical job-related (Error or Fatal) alert is received.

There are four categories of alerts, distinguished by the severity of the alert:

- Informational (containing Info, Debug, and Trace information)
- Warning
- Error
- Fatal

An example of an Informational alert is when a job is started or stopped manually by the user. If a job loses one of its participating hosts and as a result, cannot keep a quorum and shuts down, then a Fatal alert will be logged. All alerts are also logged to the file **job_alert.log**, available under the **Hub\logs** subdirectory within the Peer Management Center installation directory.

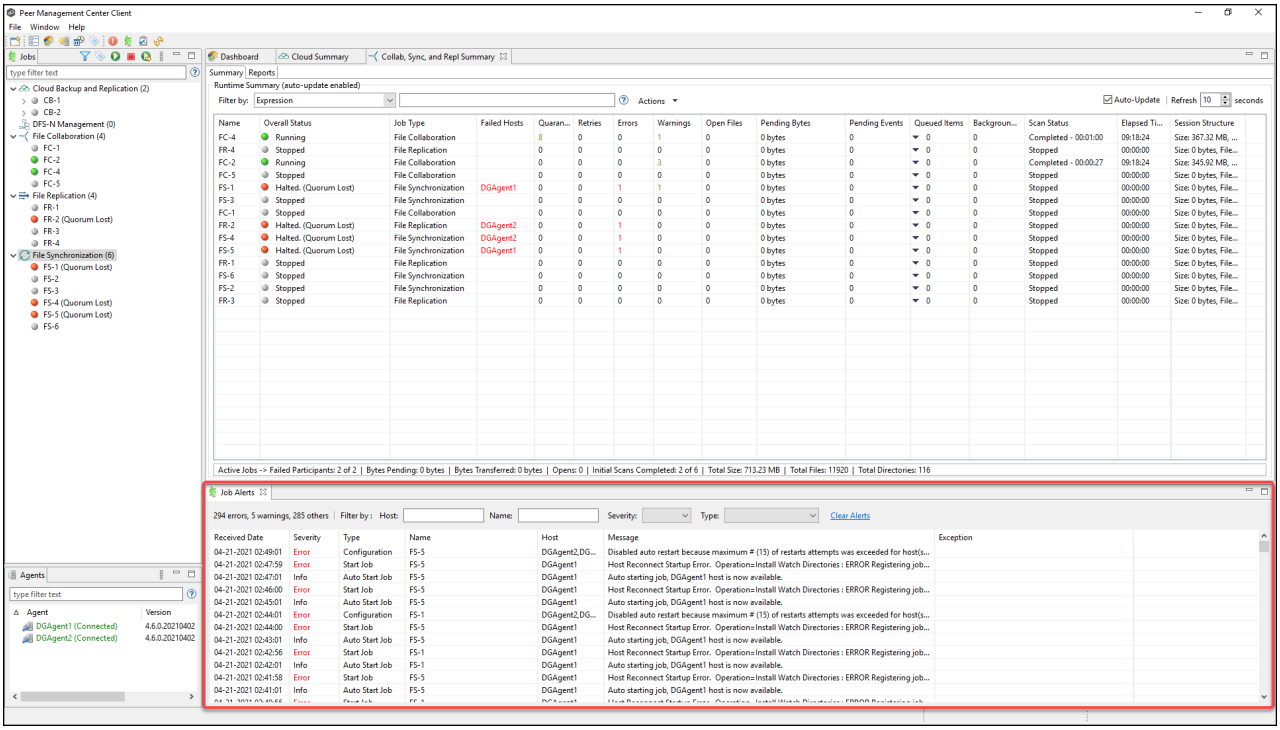

You can filter alerts based on host name, job name, severity level, or type, and you can sort alerts by clicking on a specific column header. You can also clear all alerts in the table by clicking the **Clear Alerts** link.

Displaying Job Alerts

You can open the Job Alerts view at any time by clicking the **View Job Alerts** button located on the Peer Management Center toolbar or by selecting the **Window** menu, then the **Show View** submenu, followed by the **View Job Alerts** menu item. You can close the view at any time by clicking on the **X** (Close) button on the Job Alerts tab.

You can resize the Job Alerts view by dragging the separator between the upper view and the Job Alerts view, or you can-double click the **Job Alerts** tab to maximize the view. You can restore the view to its original, non-maximized size by double-clicking the **Job Alerts** tab again.

Jobs View

The **Jobs** view is displayed in the upper left quadrant of the Peer Management Center interface and lists all the jobs, grouped by type. The number in the parentheses following the job type identifies the number of existing jobs of that type. This view is automatically displayed when Peer Management Center is started.

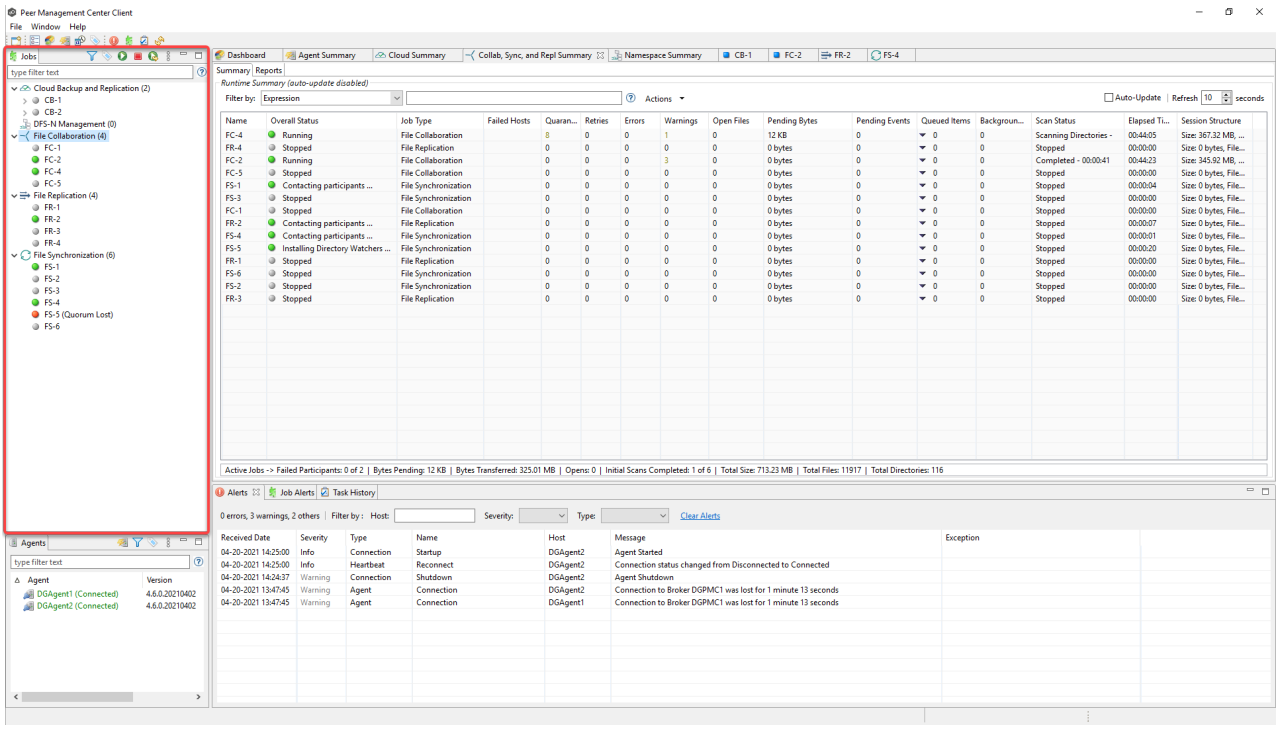

You can easily display more information about a job or job type by double-clicking a job name or job type name:

- Double-clicking any job name in the list will display a runtime view of that job.
- Double-clicking any job type name in the list will display a summary view of that job type.

To filter a large list of jobs, use the **Filter** field located below the Jobs view toolbar. For more details on how to filter jobs, see List Filters.

You can expand all or collapse all jobs by clicking the **View** button in the Jobs view toolbar and selecting an option from the **View** menu:

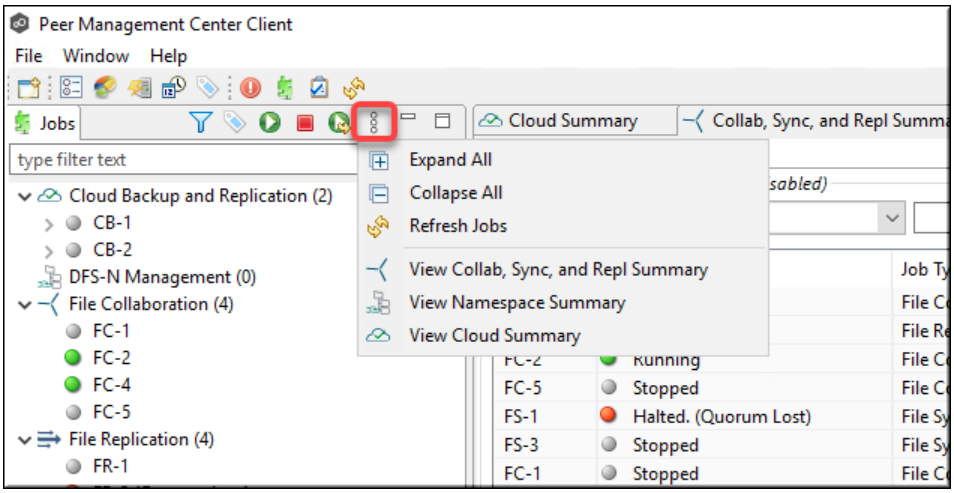

The following buttons are available on the toolbar within the **Jobs** view:

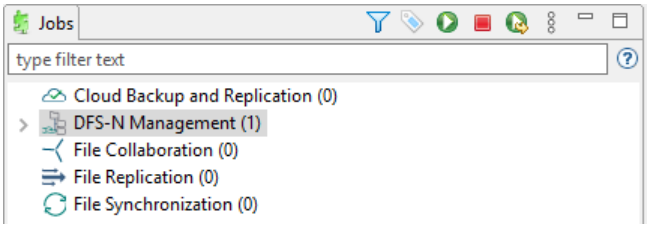

Button	Description
Manage, Save and Load Filters	Enables selection of predefined or user-defined filters and to save/manage filters. Default filters include Failed Jobs, Jobs with Backlog, and Running Scans.
Assign Tags	Opens the Assign Tags dialog where resources can be viewed, tagged, and assigned to categories. Tagging resources helps when managing large number of resources.
Start	Starts one or more selected and currently stopped jobs.
Stop	Stops one or more selected and currently running jobs.
Restart	Restart one or more selected jobs.
View	Presents options for displaying views and collapsing and expanding jobs in the **Jobs** view.

Runtime Views

Each job has a **runtime view** that show a combination of real-time file I/O activity, history, and configuration information. The job name appears as the title of the view. The **runtime** views are displayed in the upper right quadrant of the Peer Management Center interface.

A runtime view typically has several tabs. For example, in the following figure, the Cloud Backup and Replication job **CB-1** is displayed; this view contains six tabs.

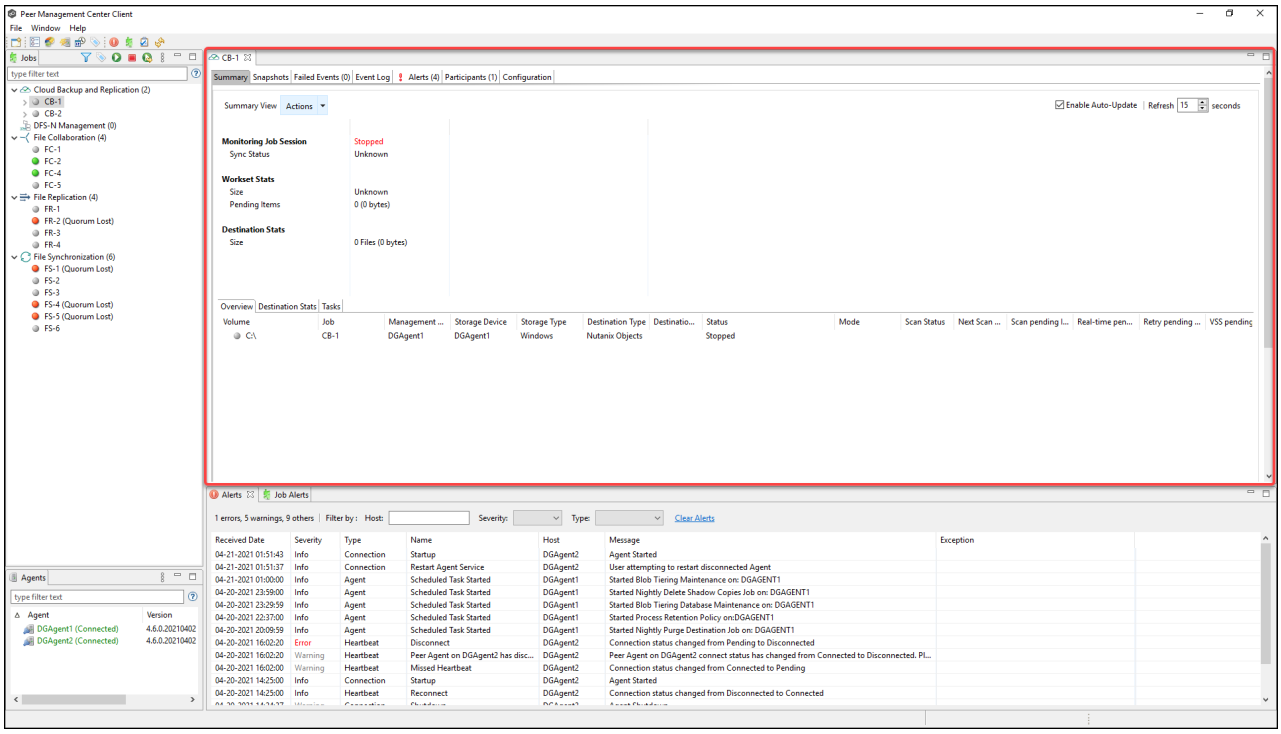

The runtime views include:

- Cloud Backup and Replication job
- DFS-N Management job
- File Collaboration job
- File Replication job
- File Synchronization job

To monitor a specific Cloud Backup and Replication job, open its runtime view.

Each Cloud Backup and Replication job has a runtime view that show a combination of real-time file I/O activity, history, and configuration. This runtime view has six tabs:

- **Summary** – Displays the status of the job, the number of and size of files uploaded in the last replication, and the size of replicated files.

- **Snapshots** – Displays a log of the snapshots taken since the job was created.
- **Failed Events** – Displays information about events that failed to successfully complete.
- **Event Log** – Displays a log of events that have occurred for the jobs – It displays the last 2500 actions that Cloud Backup and Replication has taken.
- **Alerts** – Displays a log of alerts that were issued for the job.
- **Participants** – Displays Agents that are participants in this Cloud Backup and Replication job (Currently a job can have only one participating agent.)
- **Configuration** – Displays a summary of the job configuration.

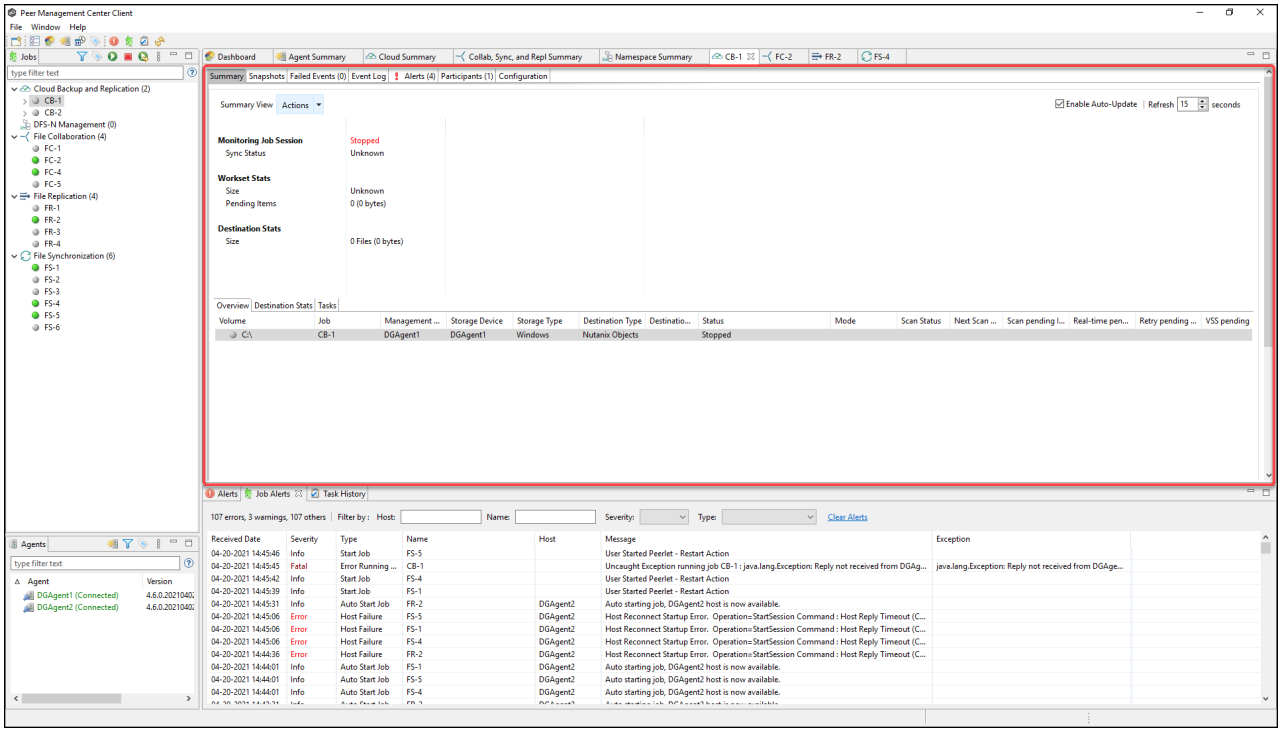

To monitor a specific DFS-N Management job, open its runtime view.

Each DFS-N Management job has a runtime view that show a combination of real-time file I/O activity, history, and configuration. This runtime view has four tabs:

- **Namespace** – The top panel of the tab displays the namespace folders is a tree structure. The namespace path is shown at the top of the tree. The bottom panel displays the folder targets linked to the selected namespace folder.

- **Namespace Servers** – Displays a list of the namespace servers and folder targets for the namespace selected in the top panel.

- **Alerts** – Displays a log of alerts that were issued for the job.

- **Configuration** – Displays a summary of the job configuration.

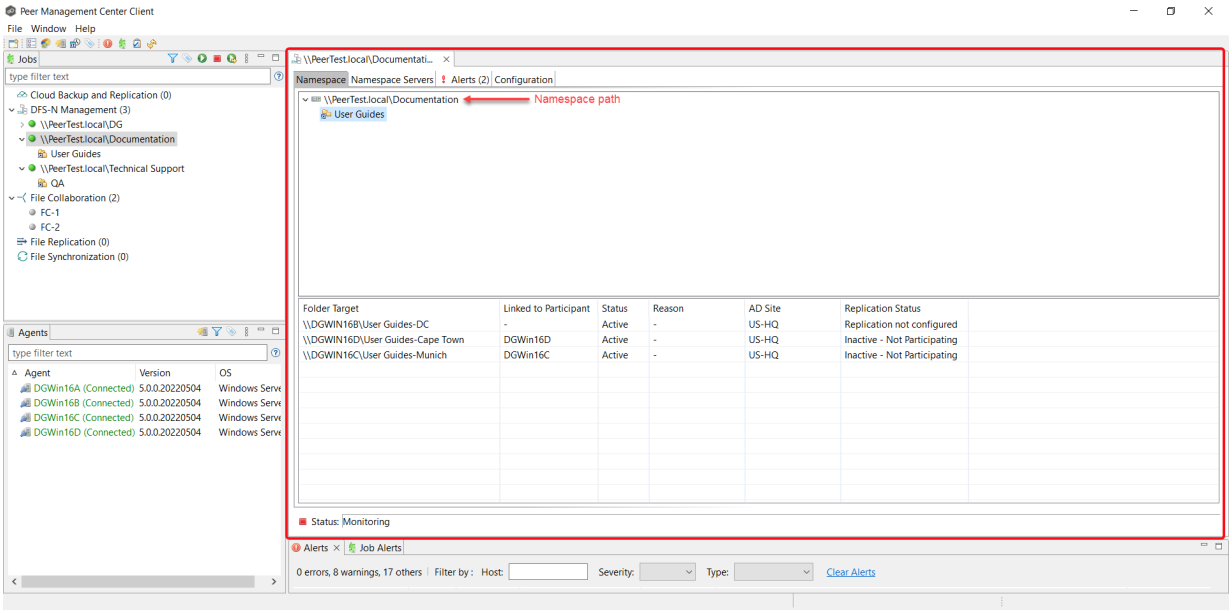

To monitor a specific File Collaboration job, open its runtime view.

Each File Collaboration job has a runtime view that show a combination of real-time file I/O activity, history, and configuration. This runtime view has eight tabs:

The view contains the following eight tabs:

- Summary - Displays overall statistics for the selected job.

- Session - Displays active open files and files that are currently in transit between participating hosts.

- Event Log - Displays a list of all runtime activity that has occurred within the selected job.

- Quarantines - Displays a list of all files that are quarantined for the session or are in conflict between two or more participating hosts.
- Retries - Displays a list of files that are currently in the Retries list.
- Alerts tab - Displays a list of all job alerts specifically tied to the selected job.
- Participants tab - Displays a list of all hosts participating in the selected job.
- Configuration tab - Displays a summary of all configurable options for the selected job.

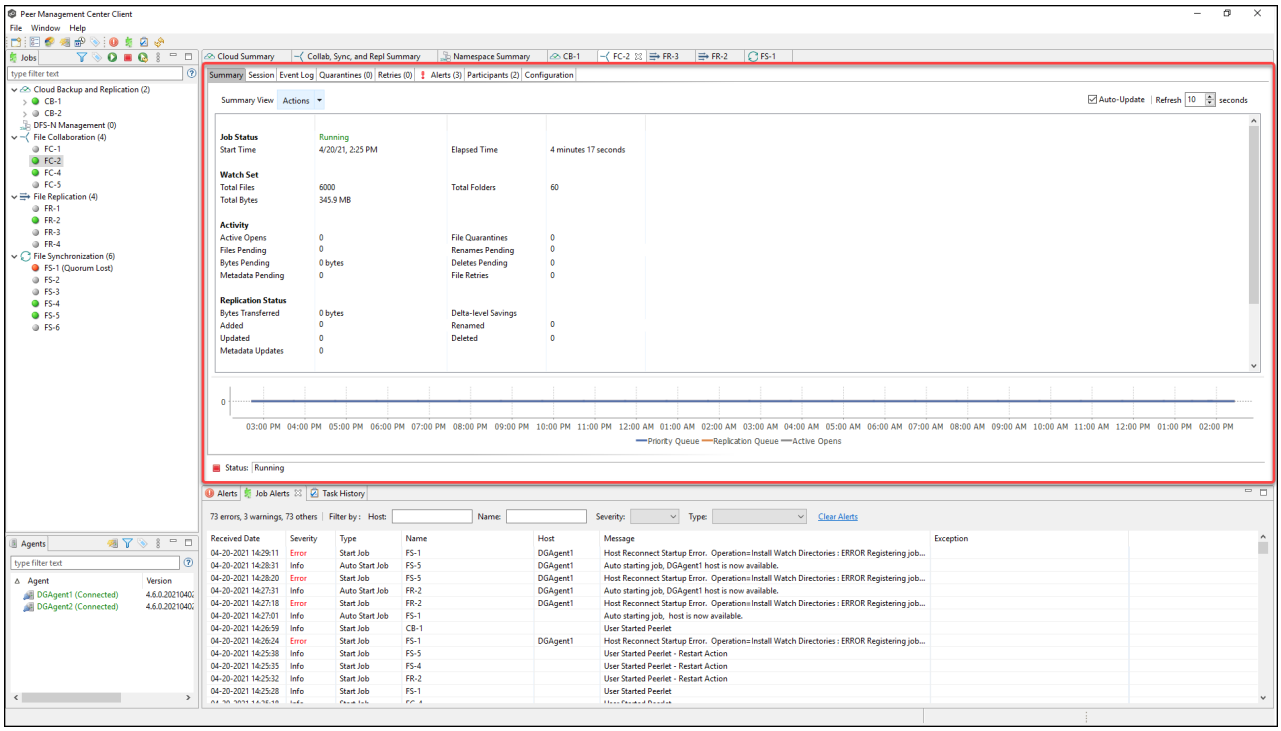

Summary Tab

The **Summary** runtime tab allows you to view current and cumulative file collaboration and synchronization statistics, as well background synchronization status. For performance reasons, this tab is not updated in real-time. However, it can be set to automatically update every few seconds. Select the **Auto-Update** checkbox to enable this functionality; set the refresh interval (in seconds) in the **Refresh** checkbox. Each refresh cycle will update the totals across all active jobs listed at the bottom of the view.

Key statistics in this view are presented in the Activity, Replication Status, and Background Scan sections. Notice that this tab is scrollable.

Activity

This section presents statistics on pending activity:

- **Files Pending** – Number of files pending synchronization, this includes queued initial scan items, bulk add files, single file adds and real-time modifies. This does not include Deletes, renames or security changes. Move your cursor over the field to see the breakdown from Adds, Updates, and Scan.
- **Bytes Pending** – Matches the Pending Bytes from the Collab, Sync, and Repl Summary view, which includes all Queued Transfers including scan works, as well as bulk adds. Note this does not track Files Pending exactly but does provide a good indication of the number of bytes currently still needing to be synchronized.
- **Metadata Pending** – Number of pending metadata changes from real-time and from initial and folder scans.
- **Renames Pending** – Total number of files and folders pending rename. Move your cursor over the field to see the breakdown for folders and files.
- **Deletes Pending** – Total number of files and folders pending delete.

Replication Status

This section presents statistics on all completed synchronization from real-time and the initial scan:

- **Bytes Transferred** – Total number of bytes transferred for all real-time Add, Bulk Add, Modify, and Scan synchronization. This does not include bulk delete, security or renames.

- **Added** – Total number of files and folders added in real-time. Move your cursor over the field to see the breakdown for folders and files.
- **Updated** – Total number of files synchronized by initial scan or real-time.
- **Deleted** – Total number of files and folders deleted.
- **Renamed** – Total number of files and folders renamed. Move your cursor over the field to see the breakdown for folders and files.
- **File Metadata Updates** – Total number of real-time and scan metadata updates for folders and files.

Background Scan

This section presents pending and completed synchronization statistics from the initial full scan.

- **Files to Replicate** – Total number of pending files synchronization queued up by initial scan.
- **Bytes to Replicate** – Total number of pending files bytes needing synchronization and queued up by initial scan.
- **Metadata to Replicate** – Total number of file and folder metadata queued up by scan.
- **Files Replicated** – Total number of completed file synchronization from the full initial scan.
- **Bytes Replicated** – Total number of bytes transferred by full initial scan.
- **Metadata Replicated** – Total number of file and folder metadata synchronized by full initial scan.

Session Tab

The **Session** tab allows you to view real-time file collaboration activity and the current session status. You can see which files are currently open in the running session, as well as any file that is currently being synchronized between hosts.

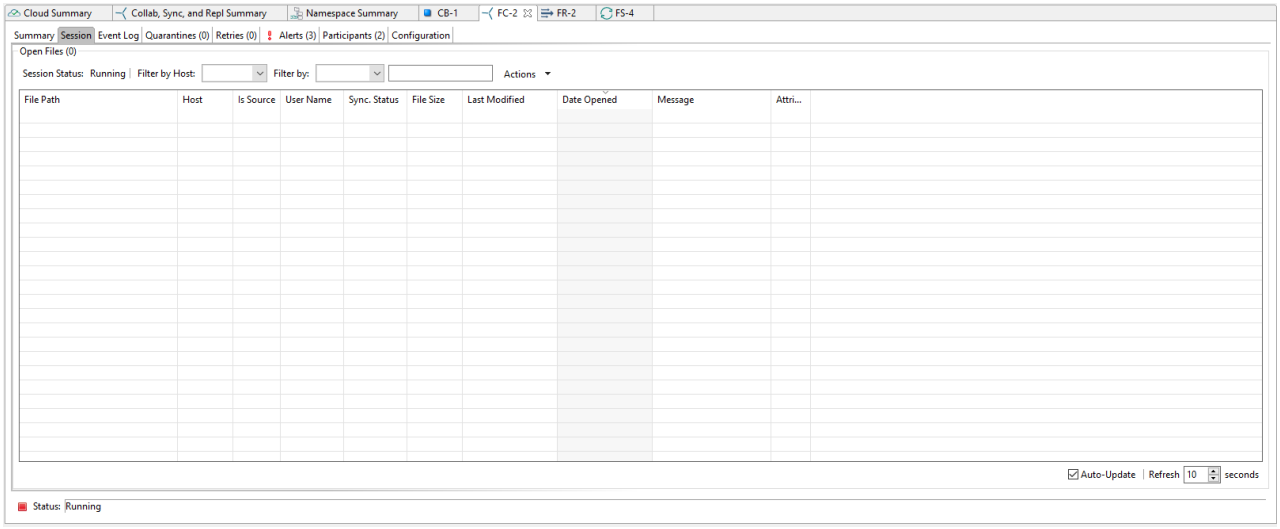

The **Session** tab has the following components:

Component	Description
Open Files table	A table showing all currently open files on the source host, any internal file locks being held by the running File Collaboration job on the target host(s), and file summary information. This table also shows all file transfers currently in progress along with file summary information, status, and overall progress. Clicking any column headers sorts by that column in ascending or descending order. All items listed in this table are grouped by file path. Each associated lock and/or transfer for each participating host will be available as a hidden child item of a root row. The root row represents the file on the source host. Pressing the + next to the root will show all associated file transfers and/or locks.
Session Status field	Field indicating the current status of the session. Valid values are: • **Stopped:** Session is stopped. • **Starting:** Session is starting up. • **Collaborating:** Real-time event detection is enabled.
Filter by Host list	A drop-down list of participating hosts to filter on. Selecting a specific host will filter the open files to show files on that host only.
Filter By Combo list	A drop-down list of additional filters that can be applied to the Open Files table, including filtering by user name (associated with the opening, adding, deleting, or modification of a file) and by file name.
Actions menu	**Refresh View:** Refreshes the entire **Open Files** table to show the latest list of file transfers and locks.

Event Log Tab

The **Event Log** tab allows you to view recent file event history for the currently running File Collaboration job based on your Logging and Alerts settings. You can specify the maximum number of events to store in the table by adjusting the Display Events spinner located in the top right corner of the panel. The maximum number of events that can be viewed is 3,000.

If you need to view more events or events from a prior session, then you can use the log files saved in the **Hub\logs** directory located in the installation directory. The event log files will start with **fc_event.log** and are written in a tab-delimited format. Microsoft Excel is a good tool to use to view and analyze a log file. See Logging and Alerts for more information about log files.

You can click any column header to sort by the column. For example, clicking the **File** column will sort by file name and you will be able to view all file events for that file in chronological order. Warnings are displayed in light gray, errors are displayed in red, and fatal errors are displayed in orange. Error records will also contain an error message in the **Message** column.

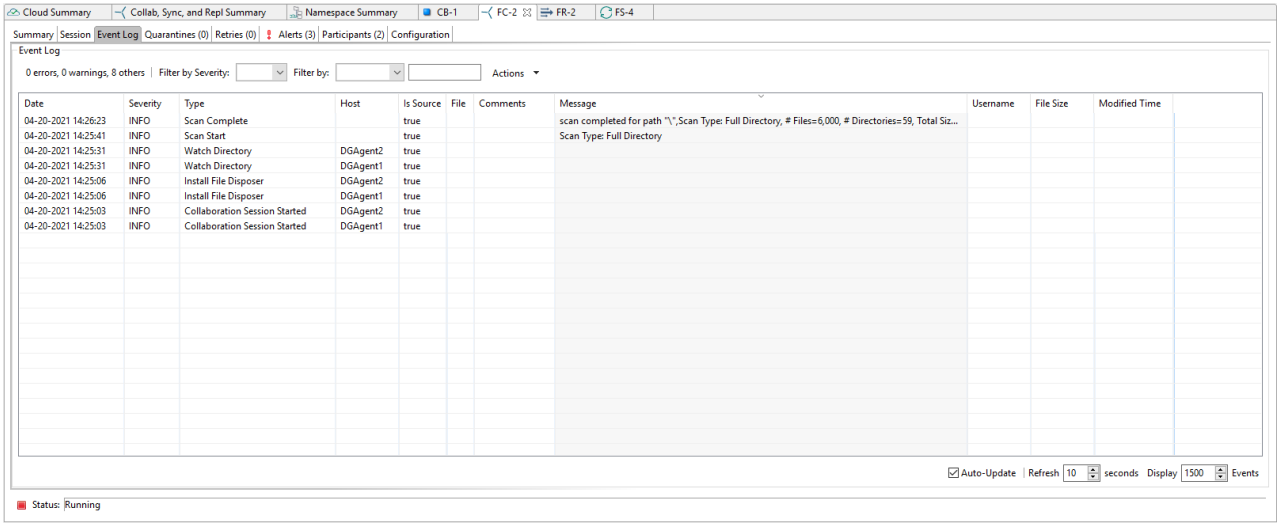

The **Actions** menu provides the following options:

Refresh View	Refresh all information provided in the table. This can also be done from the right-click context menu of the table.
Clear Events	Remove all items from the table. This can also be done from the right-click context menu of the table.

Quarantines Tab

The **Quarantines** tab displays a list of files (a) for which file conflicts could not be automatically resolved or (b) retries have failed after the maximum number of attempts. Files in this list will

no longer be synchronized or protected with file locks until a winning file is picked through the PeerGFS user interface.

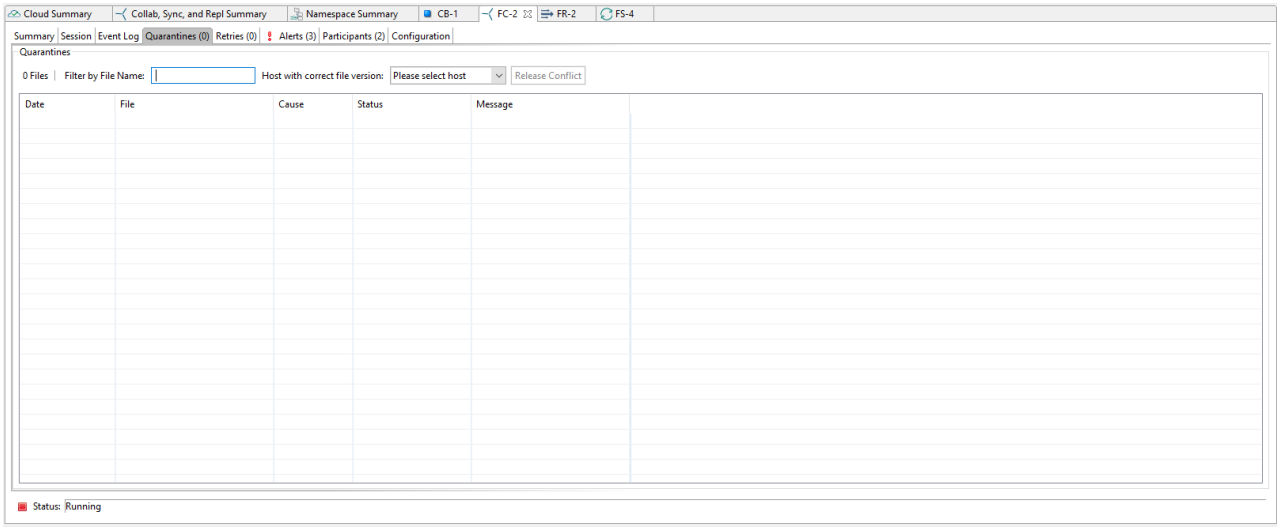

The context menu for the table contains the following actions:

Refresh View	Refresh all information provided in the table.
Purge All Quarantines	Clears all files from the quarantines list.
Copy Details	Copies the quarantine information for the selected file to your clipboard.

Retries Tab

The **Retries** tab displays the files currently in the **Retries** list. Files are put into the retry list if certain errors are thrown when trying to synchronize a file between locations. Synchronization of a file in this list will be retried every minute for a maximum of 60 attempts. The frequency of attempts and the maximum number of attempts are configurable.

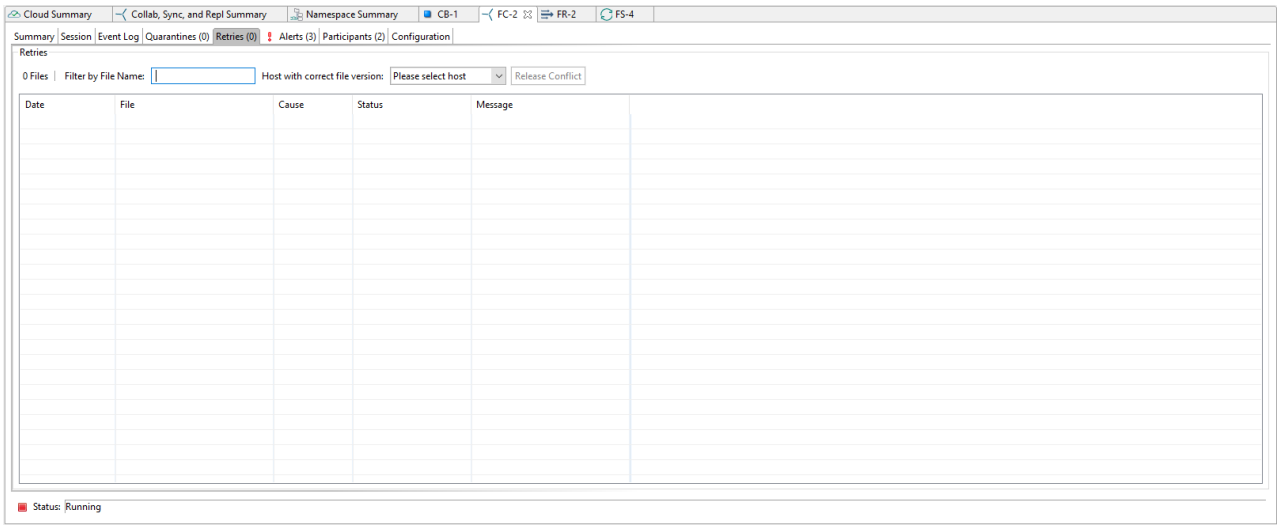

The context menu for the table contains the following actions:

Refresh View	Refresh all information provided in the table.
Purge All Quarantines	Clears all files from the **Quarantines** tab.
Copy Details	Copies the quarantine information for the selected file to your clipboard.

Alerts Tab

The **Alerts** tab allows you to view any alerts relevant to the running File Collaboration job. Items shown here are based on the configured Alerts Severity setting on the Logging and Alerts configuration page. You can specify the maximum number of alerts to store in the table by adjusting the Display Alerts spinner located in the top right corner of the panel. The alerts are also written to a tab delimited file named **fc_alert.log** within the subdirectory 'Hub/logs' within the installation directory of Peer Management Center. See the Logging and Alerts settings for more information about log files.

You can click on any column header to sort by that column. For example, clicking on the Severity column will sort by alert severity. Warnings are displayed in light gray, while errors and fatal alerts are displayed in red. In general, you should not see any alerts, but if an error or fatal alert occurs, it usually means something is wrong with the collaboration session. It may

need to be restarted or a configuration setting may need to be changed. You should consult the text in the message field for details on what occurred.

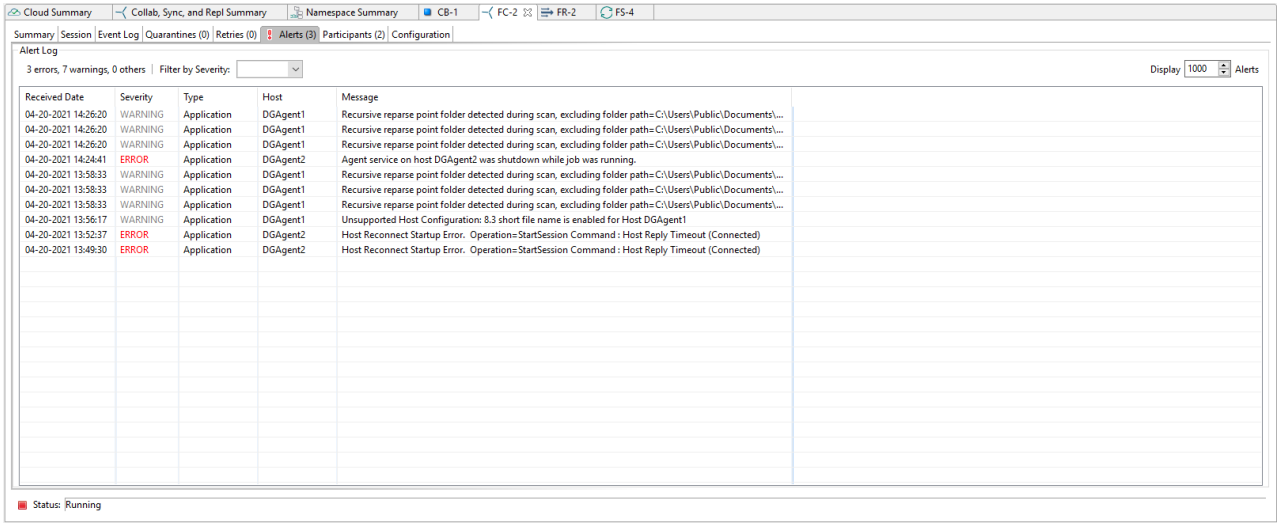

The context menu for the table contains the following actions:

Refresh View	Refresh all information provided in the table. This can also be done from the right-click context menu of the table.
Clear Events	Remove all items from the table. This can also be done from the right-click context menu of the table.

Participants Tab

The **Participants** tab is divided into two sections:

- Host Participants
- Host Participant State Change Log

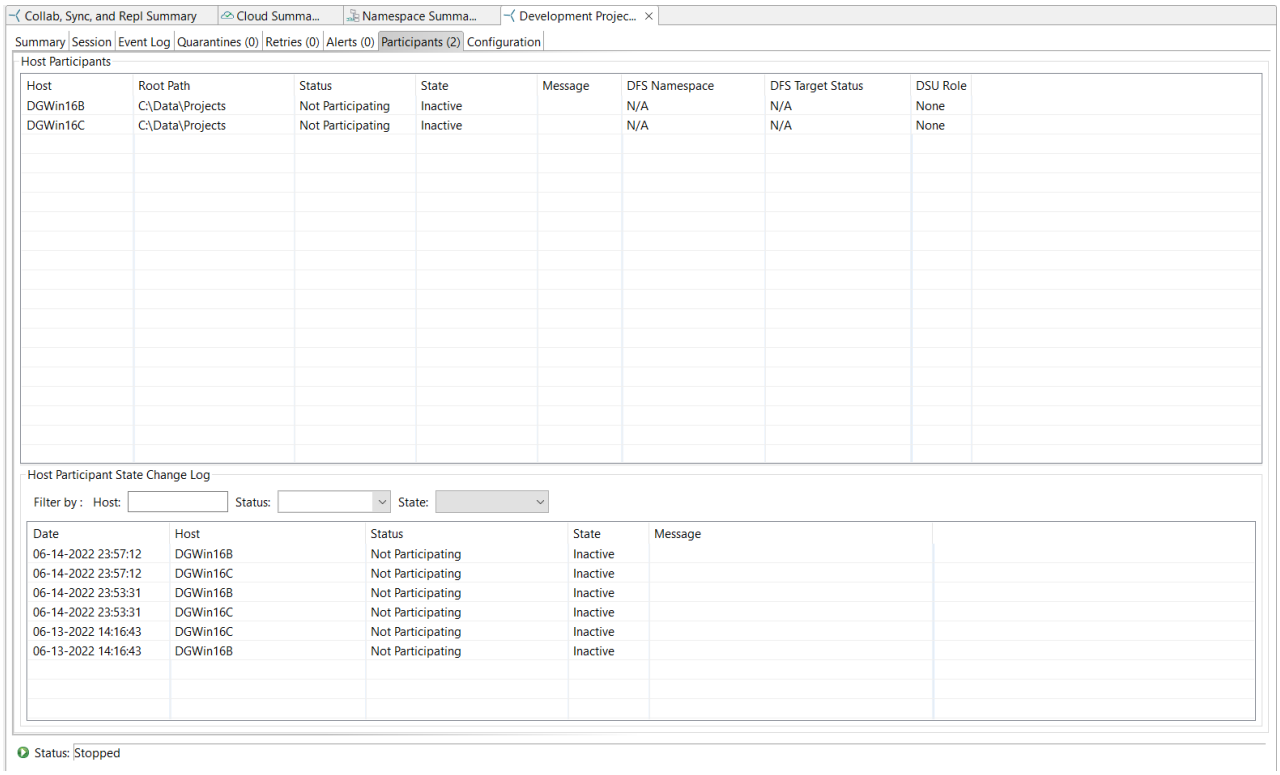

Host Participants

The **Host Participants** section contains a table that displays all the current host participants for the selected File Collaboration job. The **State** column displays activity status occurring on the hosts. If a host has become unavailable, an error message is displayed in red next to the failed host.

The following options are available in the right-click context menu for this section:

Disable Host Participant	Temporarily disables the selected participant from taking part in the File Collaboration job. You might want to do this if the host is experiencing temporary network outages.
Cancel Auto Restart	This menu item is only available if the global auto-restart functionality is enabled and the selected host has been removed from the File Collaboration job that is currently being viewed. The canceling of the auto-restart functionality for the host will only be in effect until the next time you start the File Collaboration job. If quorum has been lost for the job, canceling auto-restart on all unavailable hosts will prevent the job from automatically restarting. If quorum has not been lost, canceling auto-restart will simply prevent a host from automatically re-joining collaboration.

Host Participant State Change Log

The **Host Participant State Change Log** section contains a table that displays the most recent host participant state changes, e.g., when a host was removed from collaboration session, or when a host came back online.

The **Host Participant State Change Log** is a log of all host participant status changes (e.g., Collaborating, Not Collaborating) and/or state changes (e.g., Active, Pending Restart) of a host participant. This table is limited to 250 rows and can be filtered by host, by status, and by state.

The following options are available in the right-click context menu for this section:

Refresh View	Refresh all information provided in the table.
Clear Events	Remove all items from the table.

Configuration Tab

The **Configuration** tab displays a quick summary of all configurable items for the selected job. Each page of the File Collaboration Configuration edit wizard is represented in its own part of the view and can be collapsed if desired. Clicking **Edit this Configuration** opens the Edit Job wizard, where you can edit the current configuration.

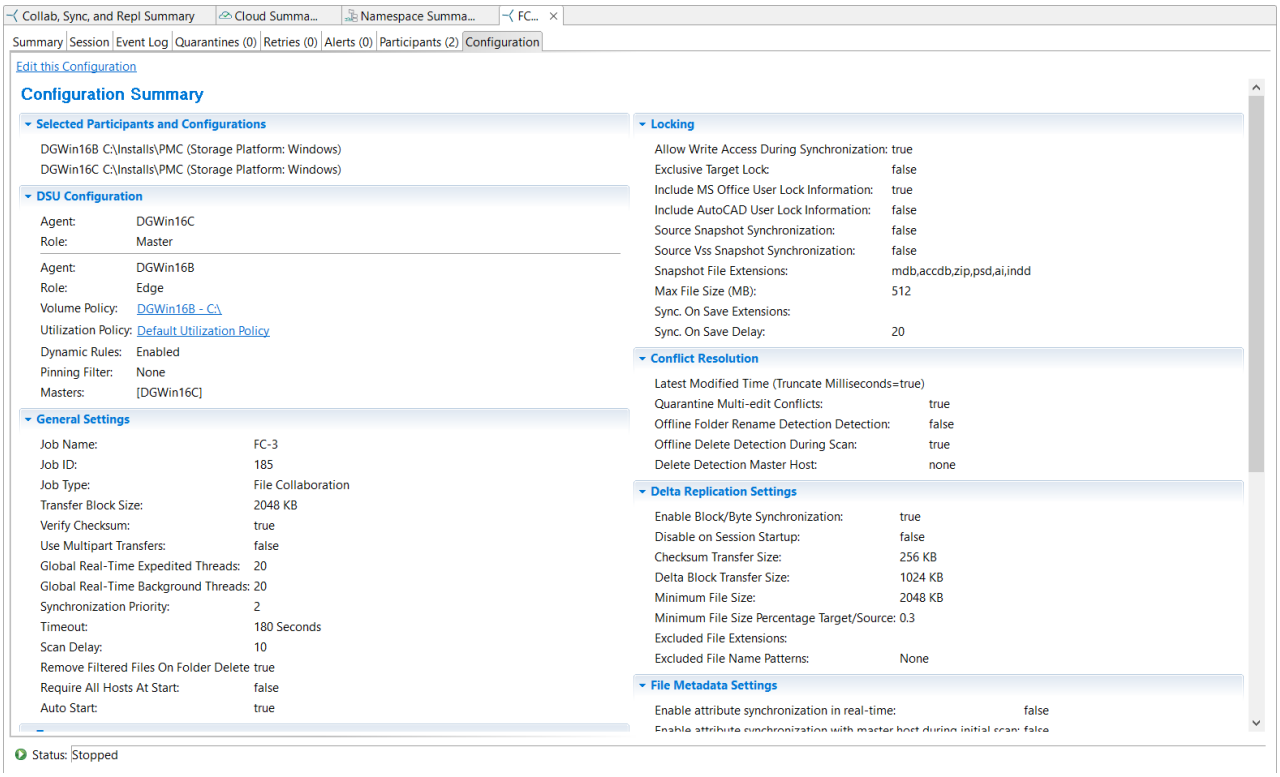

To monitor a specific File Replication job, open its runtime view.

Each File Replication job has a runtime view that show a combination of real-time file I/O activity, history, and configuration. This runtime view has six tabs:

- Summary
- Session
- Event Log
- Quarantines
- Retries
- Alerts
- Participants

- Configuration

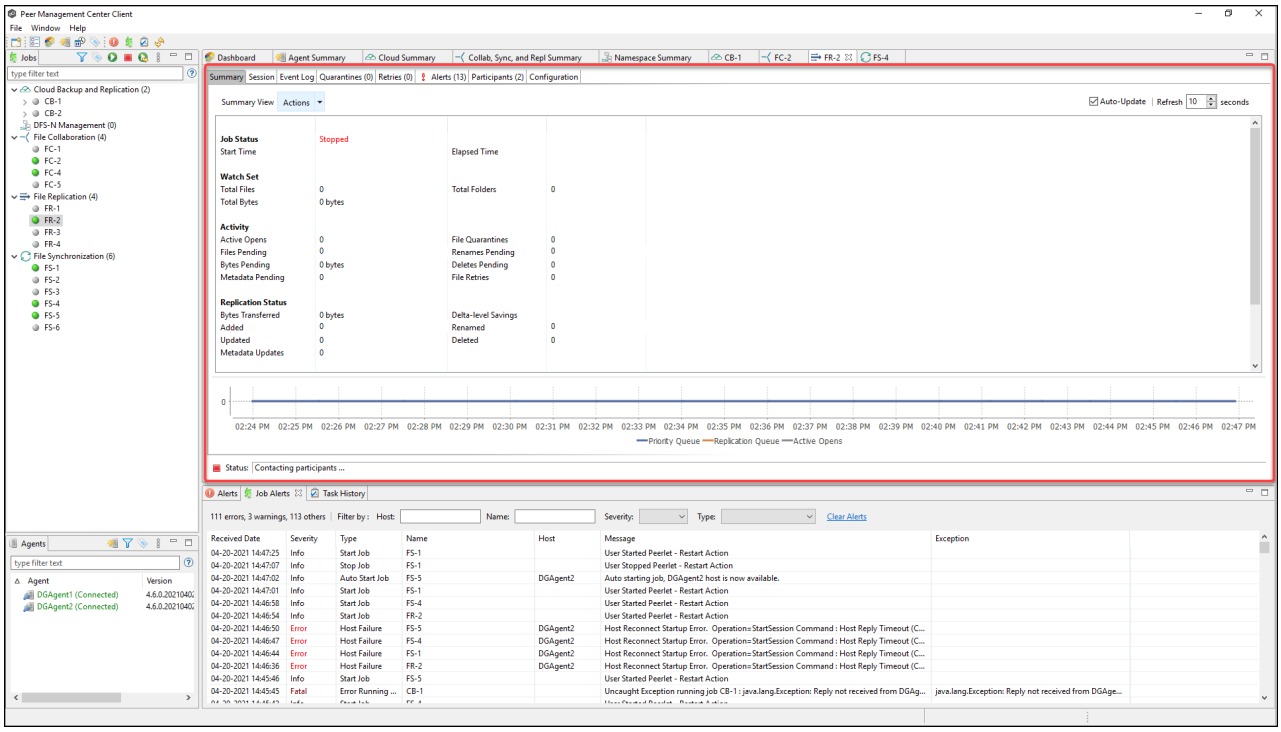

To monitor a specific File Synchronization job, open its runtime view.

Each File Synchronization job has a runtime view that show a combination of real-time file I/O activity, history, and configuration. This runtime view has six tabs:

- Summary
- Session
- Event Log
- Quarantines
- Retries
- Alerts
- Participants

- Configuration

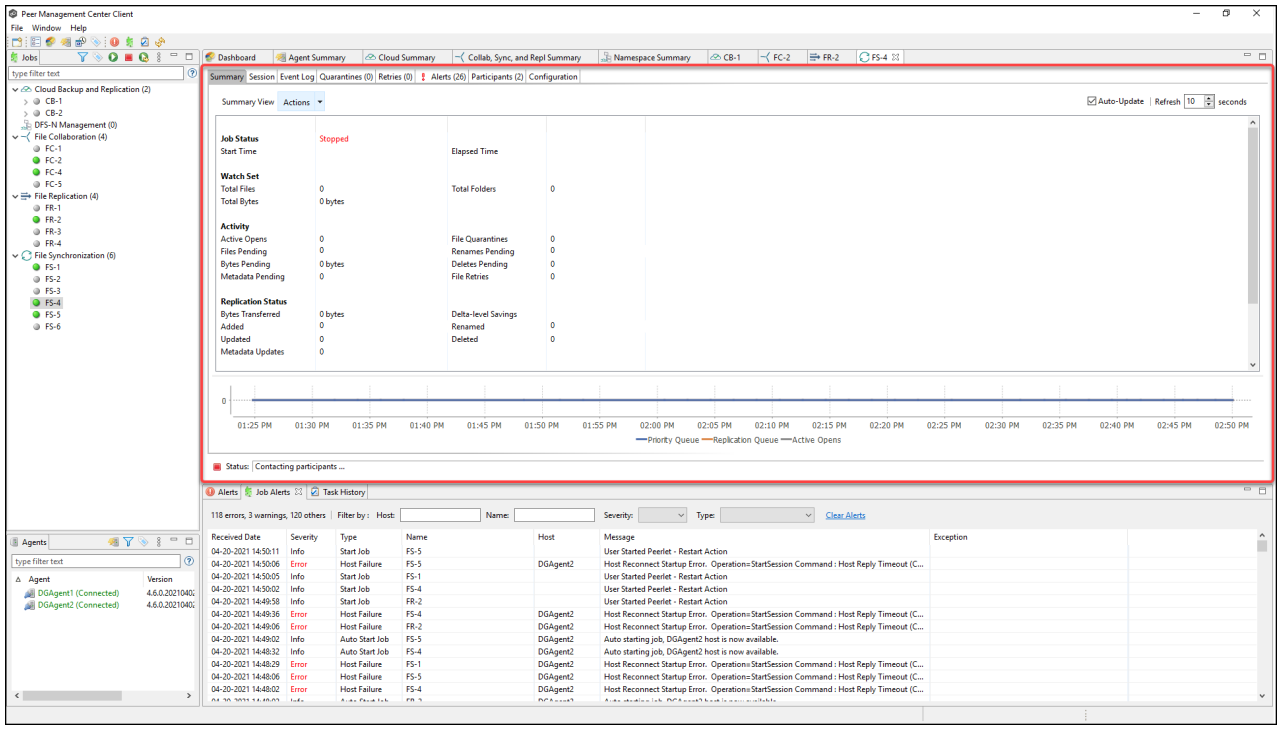

Summary Views

You can use the summary views to monitor the overall health of your jobs and Agents. You can set summary views to be automatically displayed when Peer Management Center is started. The **summary** views are displayed in the upper right quadrant of the Peer Management Center interface.

A summary view typically has several tabs. For example, in the following figure, the summary view for File Collaboration, File Synchronization, and File Replication jobs is displayed; this view contains three tabs.

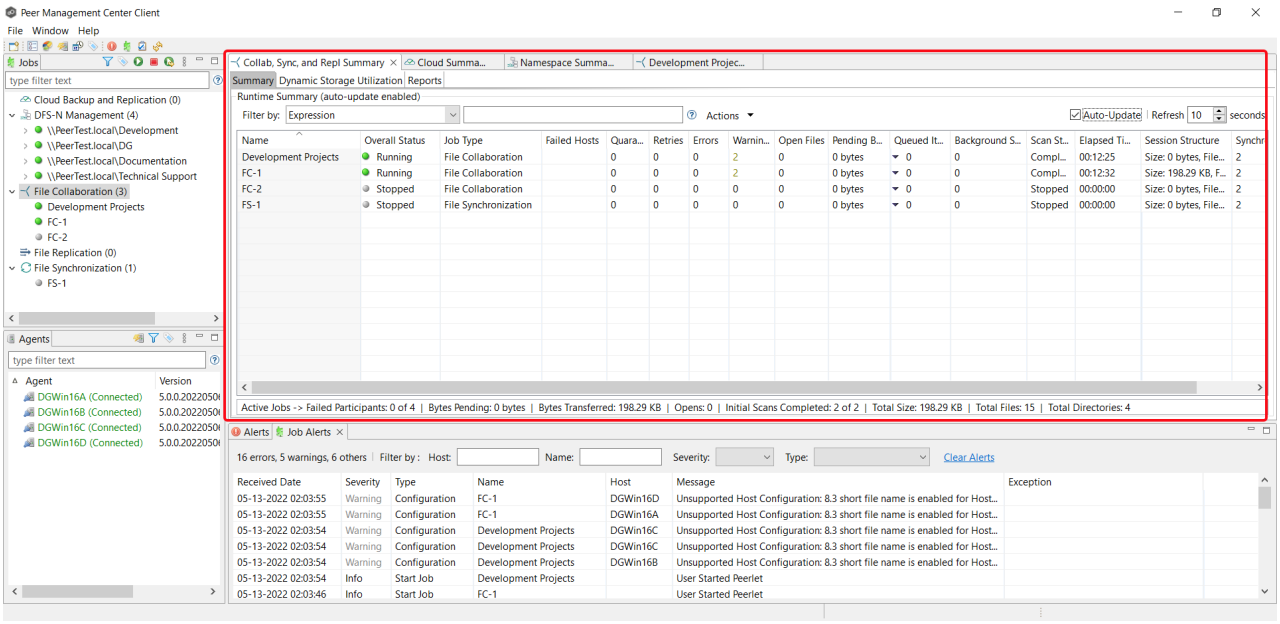

The summary views include:

- Agent Summary - Displays summary information about the Agents.
- Cloud Summary - Displays summary information about running Cloud Backup and Replication jobs.
- Collab, Sync, and Repl Summary - Displays summary information about running File Collaboration, File Synchronization, and File Replication jobs
- Namespace Summary view - Displays summary information about namespaces and DFS-N Management jobs.

The **Agent Summary** view displays a list of all known Agents deployed and their detailed status information, which can be used to assess the health of the environment. This summary view has a single tab.

The **Agent Summary** view is updated in real-time and can be filtered by using an expression or by built-in categories such as **Connected**, **Disconnected**, and **Needing Upgrade**.

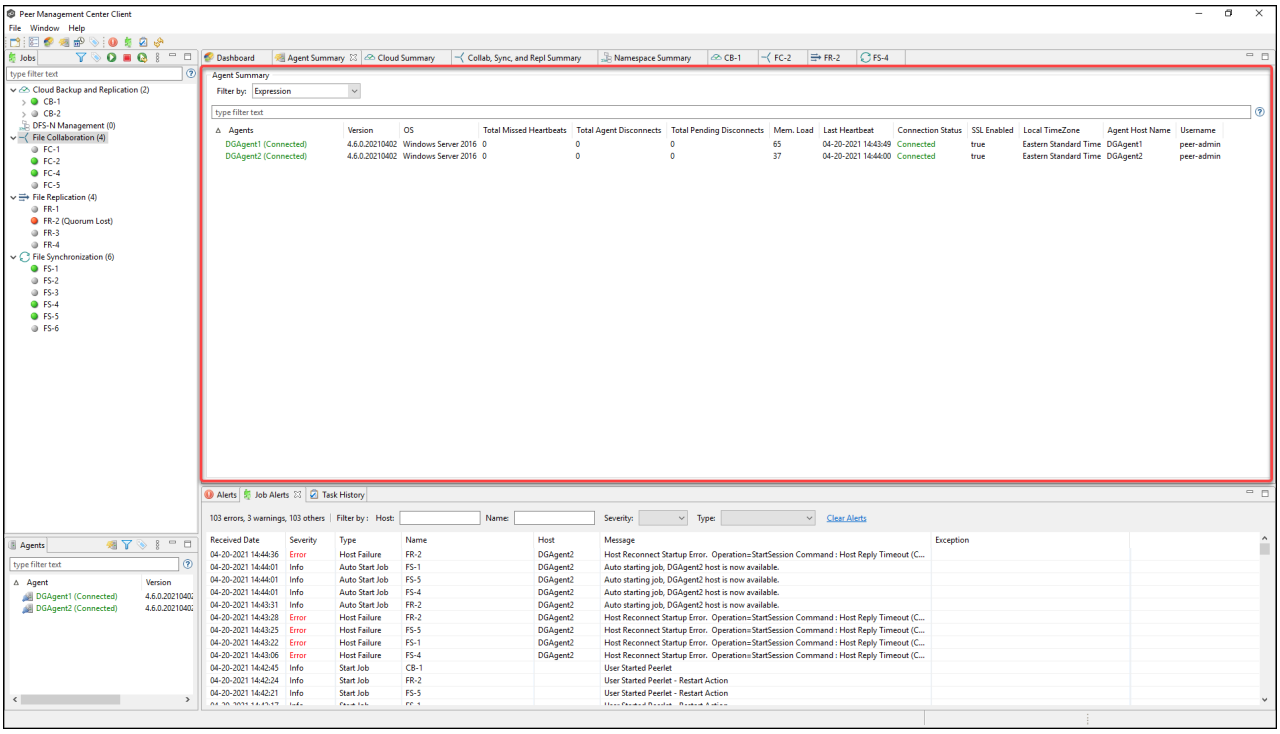

To display the Agent Summary view, use one of the following methods:

- Select **Show Agent Summary** from the **Window** menu.
- Click the **Show Agent Summary** icon in the main PMC toolbar or in the Agents view toolbar.

Use the **Cloud Summary** view to monitor the overall health of your Cloud Backup and Replication jobs. This view is the first place to check to see the status of your Cloud Backup and Replication jobs.

This view can be set to be automatically displayed when Peer Management Center is started and can be opened at any other time by double-clicking the job type name **Cloud Backup and Replication** in the Jobs view or by selecting **View Cloud Summary** from the toolbar in the **Jobs** view.

This view has four tabs:

- **Volume Summary** – Displays the volumes associated with jobs. The color of the icon next to a volume name quickly indicates the status of the job associated with that volume —a green icon indicates an active job; a gray icon indicates an inactive job, and a red icon indicates a problem with a job.

- **Job Summary** – Displays the status of all Cloud Backup and Replication jobs.
- **Destination Statistics** – Displays the total number of files that have been replicated since the first run of the jobs and other statistics.
- **Tasks** – Displays a high-level view of activities such as snapshots, and recovery processes, and background events for all Cloud Backup and Replication jobs.

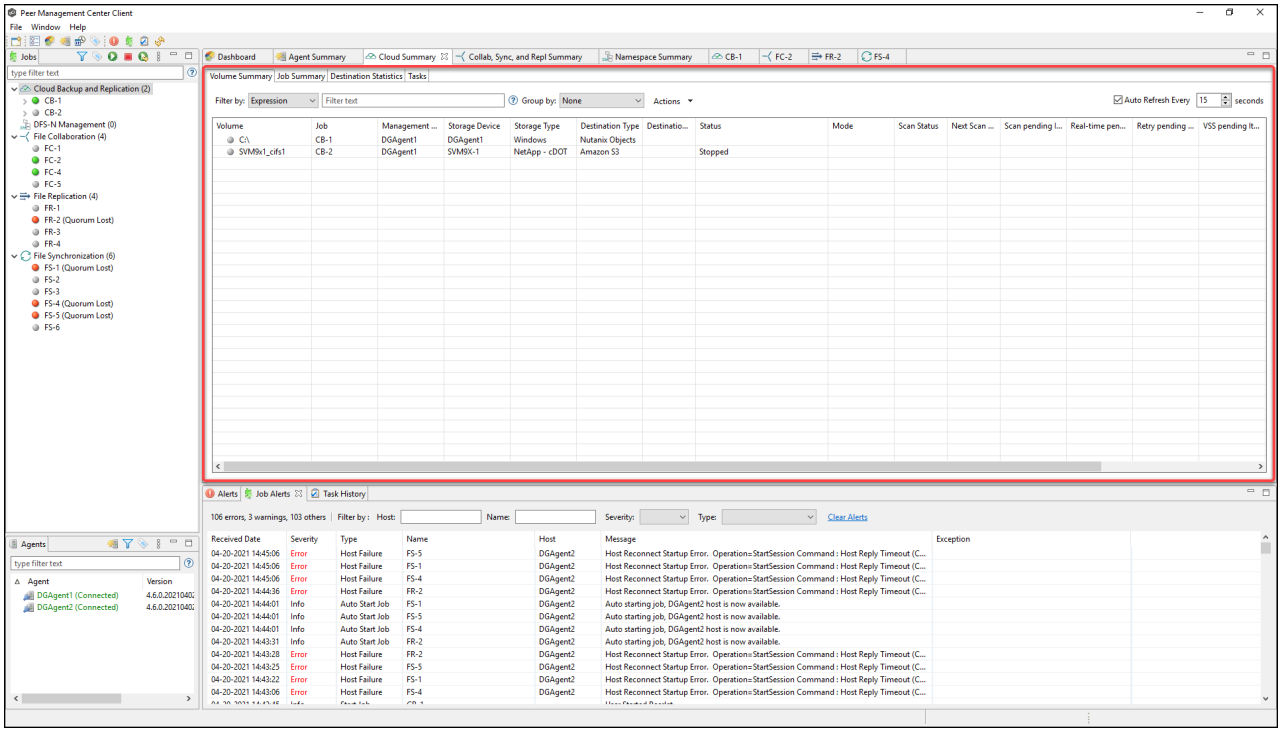

Use the **Collab, Sync, and Repl Summary** view to monitor the overall health of your File Collaboration, File Replication, and File Synchronization jobs. This view is the first place to check to see the status of your File these job types.

This view can be set to be automatically displayed when Peer Management Center is started and can be opened at any other time by double-clicking one of the job type names (**File Collaboration**, **File Replication**, or **File Synchronization**) in the Jobs view or by selecting **View Collab, Sync, and Repl Summary** from the **Jobs** view toolbar.

This view has three tabs:

- Summary

- Dynamic Storage Utilization
- Reports

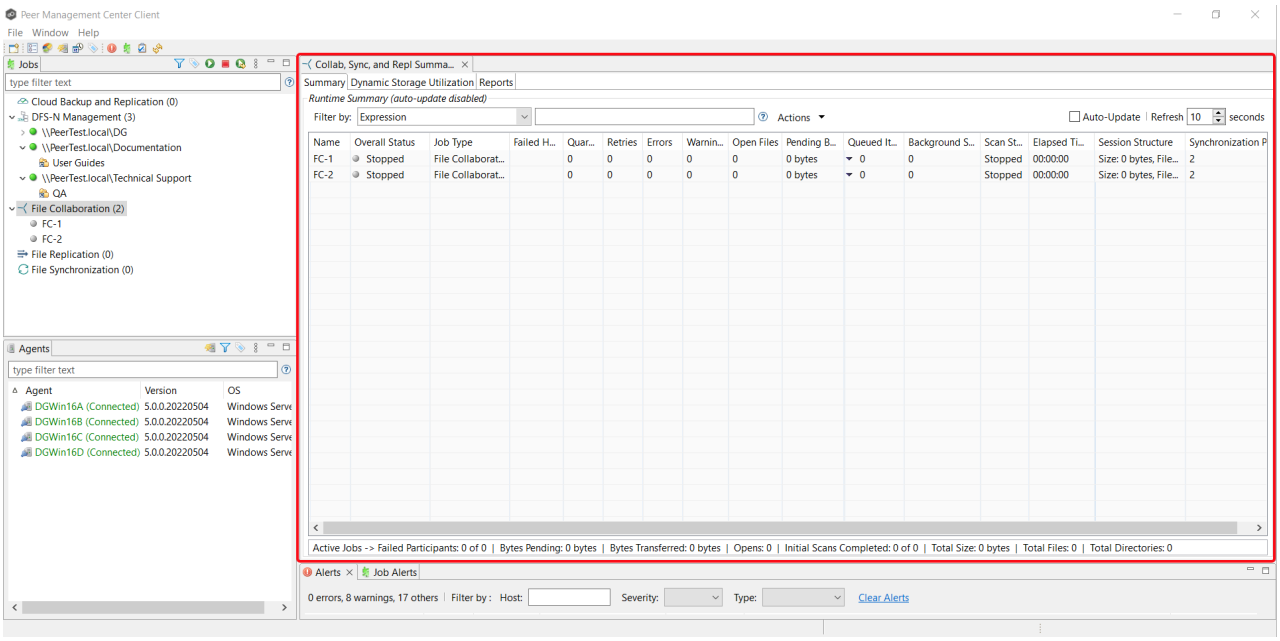

Summary Tab

The **Summary** tab aggregates critical status and statistical information from all File Collaboration, File Synchronization, and File Replication jobs in a single table. It presents overall job status, basic pending, and bytes transferred statistics. See the Reports tab for more detailed pending activity information.

Information in this view can be sorted and filtered. Operations such as starting, stopping, and editing multiple job at once are available, in addition to the ability to clear job alerts and purge quarantines from stopped jobs. Scroll the view horizontally to see all of its columns.

For performance reasons, this tab is not updated in real-time. However, it can be set to automatically update every few seconds. Select the **Auto-Update** checkbox to enable this functionality; set the refresh interval (in seconds) in the **Refresh** spinner. Each refresh cycle will update the details across all jobs, as well as the active jobs totals listed at the bottom of the view.

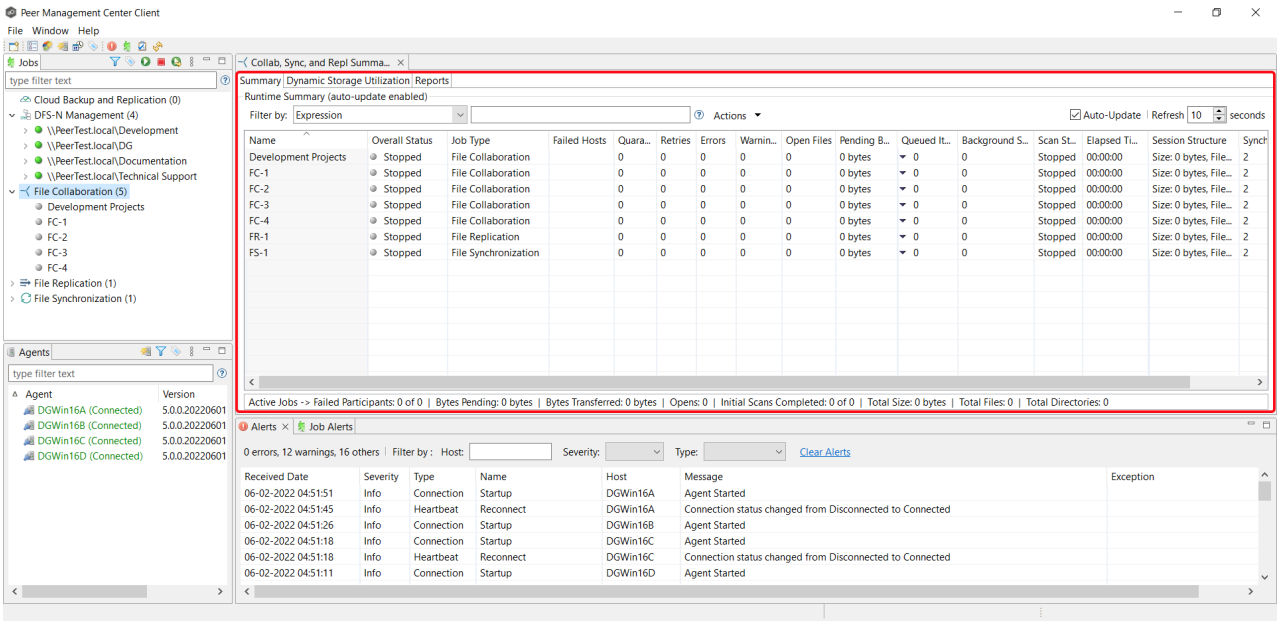

You can change which jobs are displayed in the table by filtering the list or by job state (Running in Good State, Running with Quarantines, Not Running - Stopped, Running with Disconnected Agents, Lost Quorum), Job Name, Participant, Session Status) or by tags. Select the desired filter or enter your own expression in the text field to the right of the **Filter by** drop-down list.

Column Descriptions

Key columns in this view are:

- **Pending Bytes** – Presents the number of bytes pending synchronization which includes scan work, real-time, as well as bulk adds.
- **Pending Events** – (Hidden by default) Presents the number of total pending items in Fast Queue, Slow Queue and Bulk Adds. This does not include Renames, Deletes, and Bulk Security changes. This can contain multiple events for a single file because target locks are separate operations, (e.g., if you add one file, there will be two events for this in queue.) Scan synchronization is not included and metadata synchronization is not reflected here.
- **Queued Items** – Presents the number of items in just the Fast and Slow queue (does not include bulk adds).
- **Background Sync.** – Presents the number of initial and full scan items in queue.

Additional columns can be added to and removed from the table using the right-click context menu.

Actions Menu

The **Actions** menu provides the following options:

Option	Description
Filters	Allows you to select predefined or user-defined filters and to save/manage list filters. Default job filters include Failed Jobs, Jobs with Backlog, and Running Scans.
Scheduler	Opens the Task Scheduler.
Custom Sort...	Enables you to define multi-level sort criteria for the table. This is useful for keeping important items visible at the top. For example, you may choose to create a sort level where the Overall Status column is sorted in ascending order by default.
Refresh View	Refreshes all information displayed in the table.
Copy All Filtered Statistics	Copies detailed information to the system clipboard for all items current displayed in the table, taking any filters into account. This information can then be pasted into a text editor.
Export Table Data to File	Dumps the entire contents of the table to a text file that can be viewed in any text editor.

Dynamic Storage Utilization tab

The **Dynamic Storage Utilization** tab presents about jobs using Dynamic Storage Utilization in a single table.

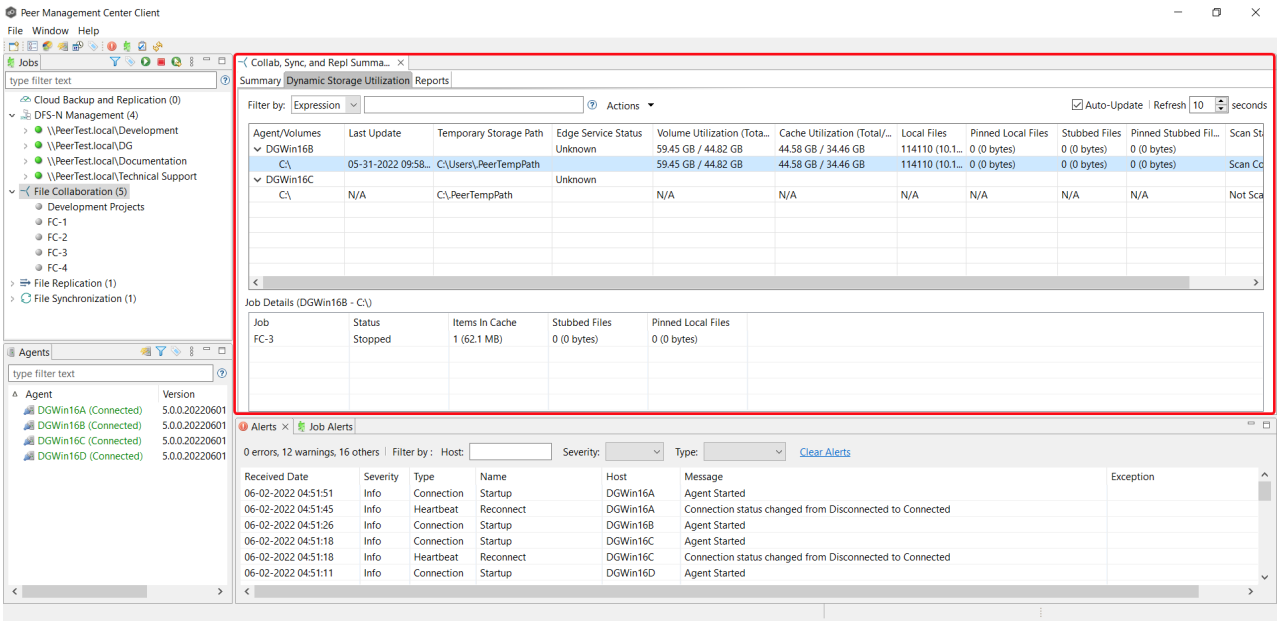

Reports Tab

The **Reports** tab presents critical statistical information from all File Collaboration, File Synchronization, and File Replication jobs in a single table. The **Reports** tab is visible when the **Enable Advanced Reporting Tab** option on the Collab, Sync, and Repl Summary page in Preferences is selected.

The **Reports** tab is especially useful to view the number of files that are in the queue waiting to be synchronized (shown in the **File Sync Queue** column). Scroll the view horizontally to see all of its columns

For performance reasons, this tab is not updated in real-time. However, it can be set to automatically update every few seconds. Select the **Auto-Update** checkbox to enable this functionality; set the refresh interval (in seconds) in the **Refresh** checkbox. Each refresh cycle will update the totals across all active jobs listed at the bottom of the view.

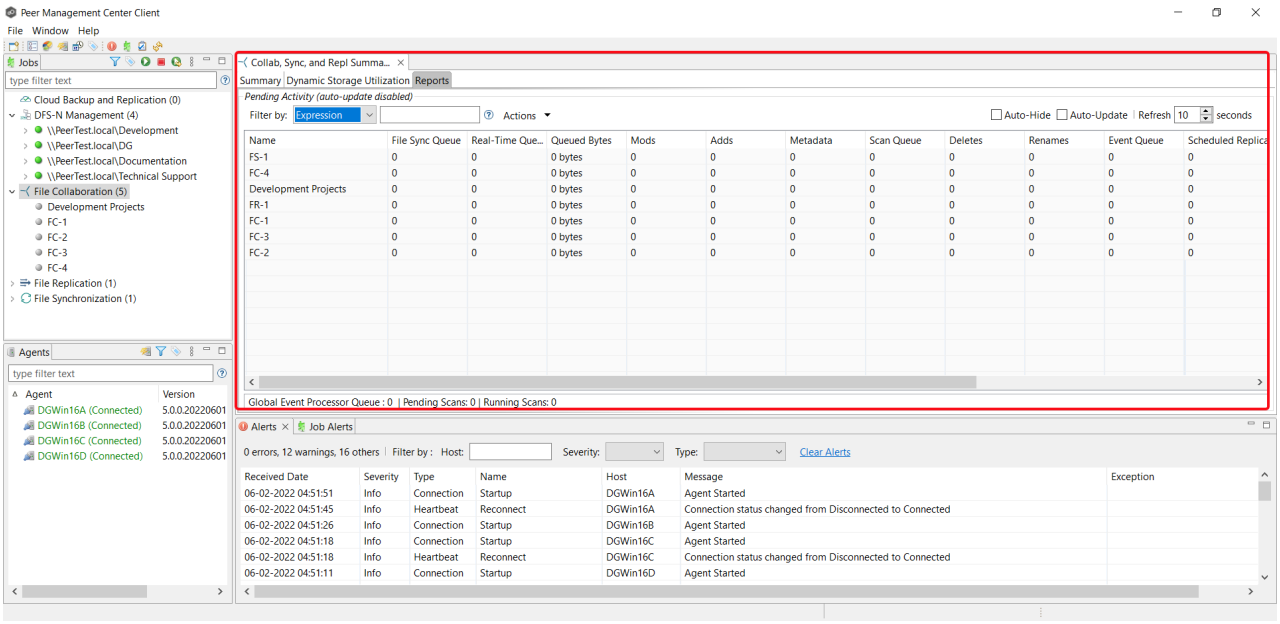

Items in the table can be filtered by a filter expression, job name, Participant, Session Status, or by tags. Select the desired filter or enter your own expression in the text field to the right of the **Filter** drop-down list. Check the **Auto-Hide** button to hide all jobs which have no pending activity.

Column Descriptions

Column	Description
Name	The name of the job.
File Sync Queue	The number of files that are in queue waiting to be processed. The number of threads available for this queue is set by the **Real-Time Background Threads** field in the Performance preferences for Collaboration, Synchronization, and Replication jobs.
Real-Time Queue	The number of open/close events that are in queue waiting to be processed. The number of threads available to process this queue is set by the **Real-Time Expedited Threads** field in the Performance preferences for Collaboration, Synchronization, and Replication jobs.
Queued Bytes	The number of bytes that are in queue waiting to be processed.

Column	Description
Mods	The number of file update events waiting to be processed for each job.
Adds	The number of file add events waiting to be processed for each job.
Metadata	The number of metadata updates waiting to be processed for each job.
Background Transfers	The number of files in the queue waiting to be synchronized as a result of a file system scan.
Deletes	The number of files deleted on a source host that are waiting to be processed.
Renames	The number of files renamed on a source host that are waiting to be processed.
Event Queue	The number of events that are queued up to run for each job.
Slow Expedited Queue	The number of events that are queued in the Slow Expedited Queue for each job.
Fast Expedited Queue	The number of events that are queued in the Fast Expedited Queue for each job.
Scheduled Replication Pending	The number of events that are queued awaiting replication at a scheduled time or interval.

Column	Description
Scheduled Replication Processing	The number of events that are queued awaiting a validation scan to make sure that the source version is correct before being released for replication.
Scheduled Replication Transfers	The number of events that are queued awaiting an available replication slot.

Actions Menu

The **Actions** menu provides the following options:

Option	Description
Filters	Allows you to select predefined or user-defined filters and to save/manage list filters. Default job filters include Failed Jobs, Jobs with Backlog, and Running Scans.
Scheduler	Opens the Task Scheduler.
Custom Sort...	Enables you to define multi-level sort criteria for the table. This is useful for keeping important items visible at the top. For example, you may choose to create a sort level where the Overall Status column is sorted in ascending order by default.
Refresh View	Refreshes all information displayed in the table.
Copy All Filtered Statistics	Copies detailed information to the system clipboard for all items current displayed in the table, taking any filters into account. This information can then be pasted into a text editor.

Option	Description
Export Table Data to File	Dumps the entire contents of the table to a file that can be viewed in any text editor.
Move Totals Row To Top	Moves the Totals row to the top of the table.
Move Totals Row To Bottom	Moves the Totals row to the bottom of the table.

Use the **Namespace** view to monitor the overall health of your DFS-N Management jobs and namespaces. This view is the first place to check to see the status of your DFS-N Management jobs. This view has a single tab.

This view can be set to be automatically displayed when Peer Management Center is started and can be opened at any other time by double-clicking the **DFS-N Management** job type name in the Jobs view or by selecting **View Namespace Summary** from the toolbar in the **Jobs** view.

The **Management Status** column shows the status of the DFS-N Management job. The **State** column shows the state of the namespace, which can be **Online**, **Offline**, **Unknown**, and **Not Found**. **Unknown** is not a common state--it typically reflects when an unexpected error has occurred or during initialization.

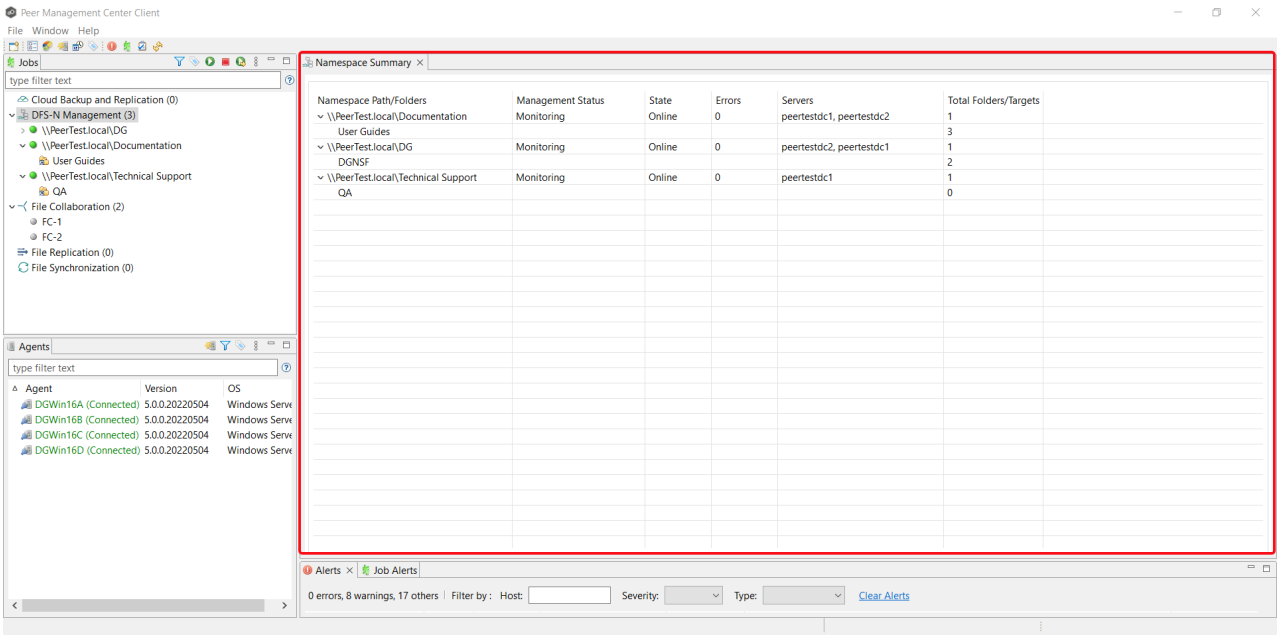

Main Window Menus and Toolbar

The main window of Peer Management Center has three menus and a toolbar:

- File
- Window
- Help

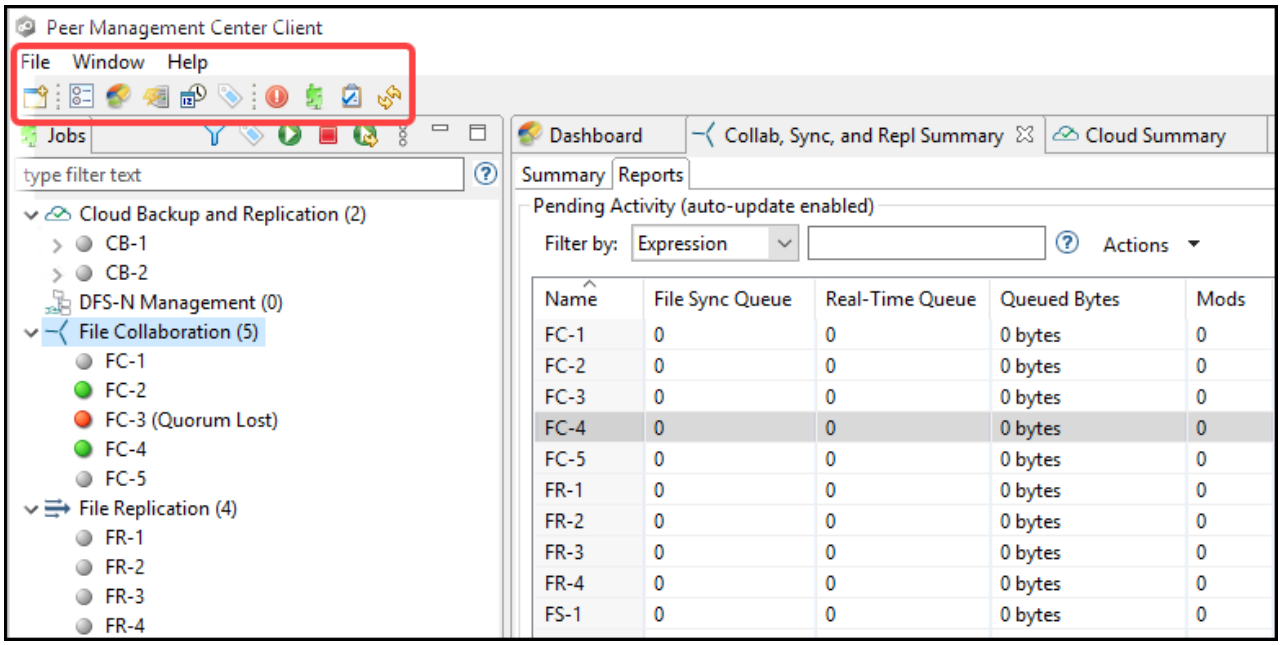

File Menu

The **File** menu in the Peer Management Center main window has the following commands:

Command	Description
New Job	Starts the Create New Job wizard.
Close	Closes the selected view.
Close All	Closes all views.
Exit	(Rich client only) Closes the Peer Management Center Client. Note that as long as the Peer Management Center Service remains running, all running jobs will continue to operate.

Command	Description
Logout	(Web client only) Logs the user out of the Peer Management Center Web client.

Help Menu

The **Help** menu in the Peer Management Center main window has the following commands:

Command	Description
User Guide	Opens the User Guide.
Support Portal	Opens the Support Portal on the Peer Software website.
Download Peer Agent Installer	Opens the Peer Software website where you can download the Peer Agent installer compatible with this version of Peer Management Center.
Peer Management Center Web Client	Opens the Peer Management Center Web Client in a web browser.
Peer Management Center API	Opens the Peer Global File Service API in a web browser.
Analytics	Opens the Set Up Analytics Wizard if Analytics has not been set up; otherwise, opens the Analytics page in Preferences.

Command	Description
Run Event Detection Analytics	Runs event detections analytics immediately. PeerGFS can perform event detection analysis every night; however, this option allows you to receive the most up-to-date analytics.
Retrieve PMC/Agent Logs	Collects and retrieve all useful log files for specified Peer Agents, Peer Management Center, and all jobs. This information is assembled into a single encrypted zip file that can optionally be uploaded to the Peer Software Technical Support. The collection and retrieval of the log and support files is performed in the background, which might take a while, depending on content size and network speed. Upon completion, you are notified and can view the zip file.
Retrieve Broker Statistics	Displays detailed statistical information about all messaging that has transpired for all connections (Peer Agents and Peer Management Center) to Peer Management Broker. Peer Technical Support can use these statistics to aid in diagnosing problems.
Generate Thread Dump File	Gives options to generate a thread dump of the running Peer Management Center Client and Service, as well as the running Peer Management Broker service. Both can be used by Peer Software technical support to debug certain issues.
Generate PMC Memory Dump File	Generates a memory dump of the running Peer Management Center Client and Service, which can be used by Peer Software Technical Support to debug certain issues.
Compress DB on Restart	(Rich client only) Compresses the database upon restart of the Peer Management Center Service. Select this option in cases where the database consumes a large amount of disk space.
Check for Updates	(Rich client only) Checks for updates to Peer Management Center. Minor releases can be automatically downloaded and installed. Major releases require a new license key and must be requested from Peer Software Technical Support.
Licenses	Displays the **Licensing** page in **Preferences**.
About Peer	Displays version information and installation details.

Command	Description
Management Center	

Window Menu

The **Window** menu in the Peer Management Center main window has the following commands:

Command	Description
Reset Perspective	Resets all current windows, views, and editors to their default size and layout.
Show Dashboard	Displays the Dashboard, which displays metrics and key performance indicators from all running File Collaboration, File Replication, and File Synchronization jobs, and Peer Agents.
Show Agent Summary	Displays the Agent Summary view, which displays a list of all known Peer Agents deployed and their detailed status information, which can be used to assess the health of the environment.
Show Summary View	Displays a submenu with the following options: • **Cloud Summary** - Displays the summary view for Cloud Back and Replication jobs. • **Namespace Summary** - Displays the summary view for DFS-N Management jobs. • **Collab, Sync, and Repl Summary** - Displays the summary view for File Collaboration, File Replication, and File Synchronization jobs.
Show View	Displays a submenu with the following options: • **Alerts** - Displays the Alerts view, which displays Peer Management Center alerts such as Peer Agent connection status changes.

Command	Description
	• **Job Alerts** - Displays the Job Alerts view, which displays alerts such as job restarts. • **Task History** - Displays the Task History view, which displays the status of tasks such as Daily Cleanup. • **Progress** - Displays the Progress view, which displays information pertaining to any running background tasks within Peer Management Center.
Preferences	Displays the Preferences page, which enables the user to configure global settings for Peer Management Center, as well as settings for individual job types.
Assign Tags	Displays the Assign Tags dialog where resources can be viewed, tagged, and assigned to categories. Tagging resources helps when managing large number of resources.
Refresh	Refreshes all open views and tabs.

Toolbar

Use the toolbar in the main Peer Management Center window to quickly launch commonly performed actions.

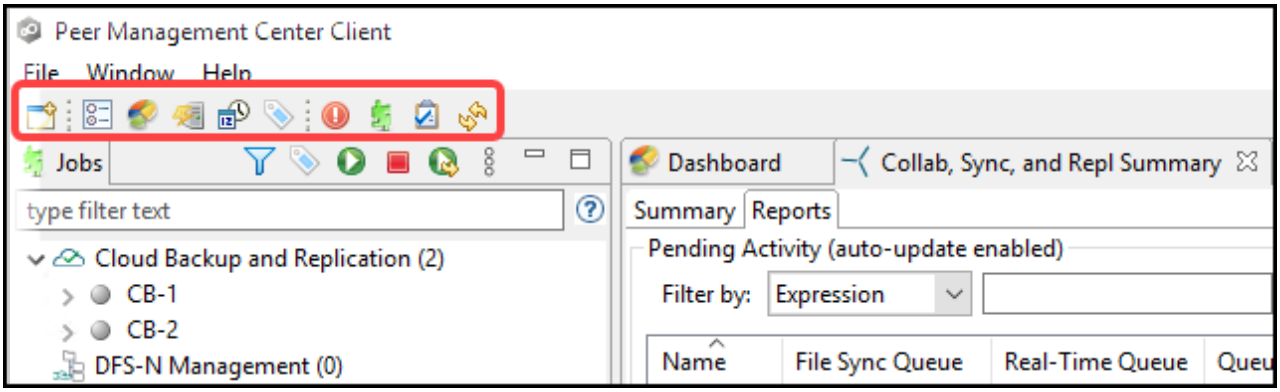

The toolbar has the following buttons:

Button	Description
New Job	Opens the New Job wizard.
Preferences	Displays the Preferences page, which enables the user to configure global settings for Peer Management Center, as well as settings for individual job types.
View Dashboard	Displays the Dashboard, which displays metrics and key performance indicators from all running File Collaboration, File Replication, and File Synchronization jobs, and Peer Agents.
Show Agent Summary	Displays the Agent Summary view, which displays a list of all known Peer Agents deployed and their detailed status information, which can be used to assess the health of the environment.
Task Scheduler	Opens the Task Scheduler, which enables a user to schedule tasks that can be carried out by the Peer Management Center at scheduled times or intervals.
Assign Tags	Displays the Assign Tags dialog where resources can be viewed, tagged, and assigned to categories. Tagging resources helps when managing large number of resources.
View Alerts	Displays the Alerts view, which displays Peer Management Center alerts such as Peer Agent connection status changes.
View Job Alerts	Displays the Job Alerts view, which displays job-related alerts such as job restarts.
View Task History	Displays the Task History view, which displays the status of tasks such as Daily Cleanup.
Refresh	Refreshes all current views and tabs.

Tables

Tables are used throughout the Peer Management Center interface to present information. For example, the Job Alerts view contains a table displaying job-specific alerts:

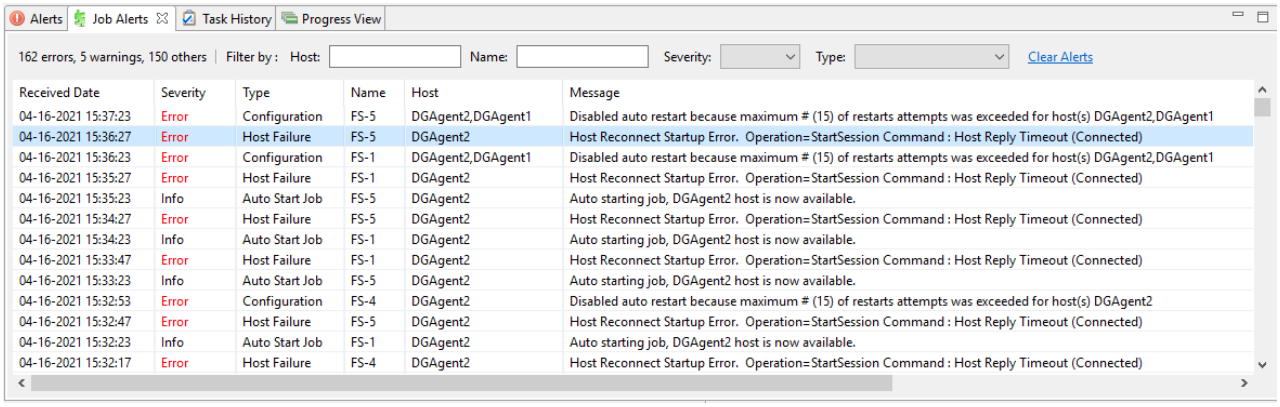

Received Date	Severity	Type	Name	Host	Message
04-16-2021 15:37:23	Error	Configuration	FS-5	DGAgent2,DGAgent1	Disabled auto restart because maximum # (15) of restarts attempts was exceeded for host(s) DGAgent2,DGAgent1
04-16-2021 15:36:27	Error	Host Failure	FS-5	DGAgent2	Host Reconnect Startup Error. Operation=StartSession Command : Host Reply Timeout (Connected)
04-16-2021 15:36:23	Error	Configuration	FS-1	DGAgent2,DGAgent1	Disabled auto restart because maximum # (15) of restarts attempts was exceeded for host(s) DGAgent2,DGAgent1
04-16-2021 15:35:27	Error	Host Failure	FS-1	DGAgent2	Host Reconnect Startup Error. Operation=StartSession Command : Host Reply Timeout (Connected)
04-16-2021 15:35:23	Info	Auto Start Job	FS-5	DGAgent2	Auto starting job, DGAgent2 host is now available.
04-16-2021 15:34:27	Error	Host Failure	FS-5	DGAgent2	Host Reconnect Startup Error. Operation=StartSession Command : Host Reply Timeout (Connected)
04-16-2021 15:34:23	Info	Auto Start Job	FS-1	DGAgent2	Auto starting job, DGAgent2 host is now available.
04-16-2021 15:33:47	Error	Host Failure	FS-1	DGAgent2	Host Reconnect Startup Error. Operation=StartSession Command : Host Reply Timeout (Connected)
04-16-2021 15:33:23	Info	Auto Start Job	FS-5	DGAgent2	Auto starting job, DGAgent2 host is now available.
04-16-2021 15:32:53	Error	Configuration	FS-4	DGAgent2	Disabled auto restart because maximum # (15) of restarts attempts was exceeded for host(s) DGAgent2
04-16-2021 15:32:47	Error	Host Failure	FS-5	DGAgent2	Host Reconnect Startup Error. Operation=StartSession Command : Host Reply Timeout (Connected)
04-16-2021 15:32:23	Info	Auto Start Job	FS-1	DGAgent2	Auto starting job, DGAgent2 host is now available.
04-16-2021 15:32:17	Error	Host Failure	FS-4	DGAgent2	Host Reconnect Startup Error. Operation=StartSession Command : Host Reply Timeout (Connected)

Most tables allow you to sort them by clicking on a column header.

Most tables support double-clicking on any row to display a dialog containing details pertaining to that row. For example, clicking a row in the Job Alerts table displays detailed information for that particular alert:

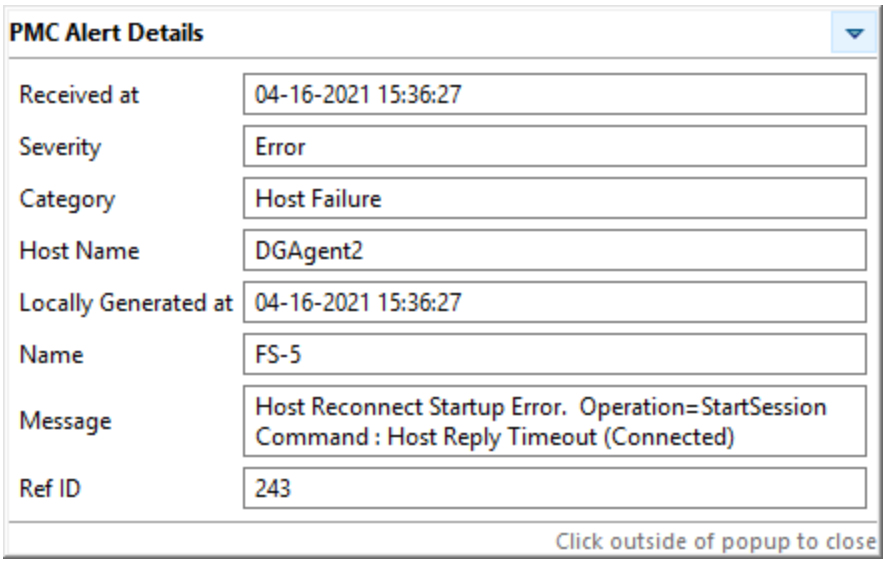

Right-clicking in a table displays a **context menu**. A context menu allows you to perform additional operations on the table. For example, you can choose which columns to hide and to display in the table. One very useful option available in many context menus is the ability to copy detailed information for one or more rows all at the same time. This information can then be pasted into any text editor.

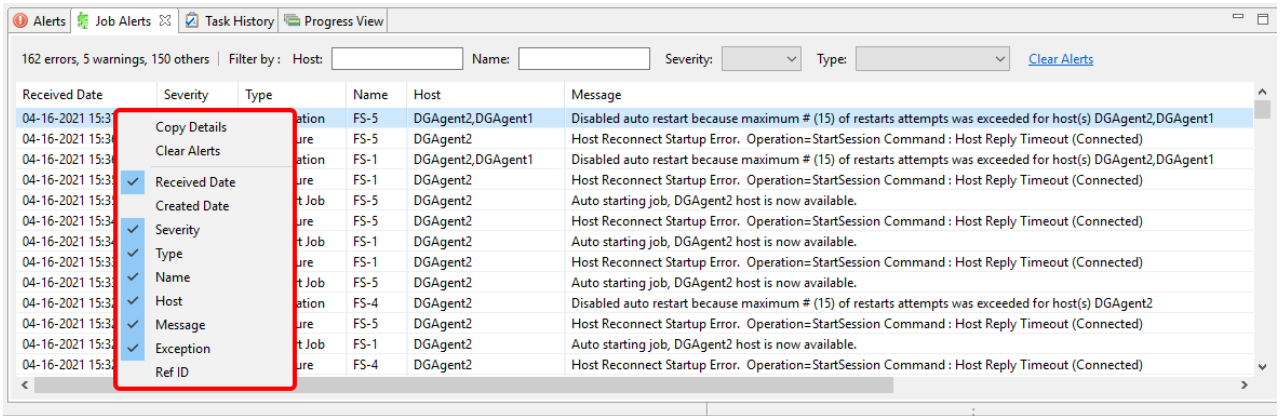

Basic Concepts

The topics in this section provide information on advanced functionality and configuration options available in Peer Management Center.

- Email Alerts
- File and Folder Filters
- List Filters
- Logging and Alerts
- SNMP Notifications
- Tags
- Web Client Users

Email Alerts

Overview

An email alert notifies recipients when a certain type of event occurs, for example, file quarantined, session aborted, host failure, system alert. When an email alert is applied to a

job, an alert is sent to all listed recipients whenever a selected event type is triggered by the job.

An email alert consists of a unique name, a selection of event types, and a list of email addresses. The available event types depend on the job type.

When you create a job, you can select an existing email alert to apply to the job or you can create a new alert and apply it to the job. Multiple email alerts can be applied to a job. You cannot modify an email alert while it is applied to a running job. You cannot delete an email alert while it is applied to any job. An alert can be applied to multiple jobs of the same type. Email alerts are defined in the preferences for a job type.

See Email Configuration for configuring an SMTP email connection. This must be configured before email alerts can be sent.

Managing Email Alerts

You can create, edit, copy, and delete alerts.

To manage email alerts:

1. Select **Preferences** from the **Window** menu.

   The **Preferences** dialog appears.

2. Select the job type from the navigation tree and expand it.

3. Select **Email Alerts** from the navigation tree.

   The **Email Alerts** page lists existing email alerts for that job type.

File and Folder Filters

Overview

A file filter enables you to specify which files (and folders) should be included and/or excluded from a job's watch set. Included files are subject to scan(s) and real-time event detection, while excluded files are not. Initially, all files are included and no files are excluded from a job, except for files matching the predefined file filters and automatically excluded file types.

Filters can also operate on folders, allowing you to include and exclude folders from a job's watch set. For more information on folder filters, see Folder Filters.

A file filter consists of a unique name and one or more filter patterns. A filter can also be based on a file's last modified time and file size. For more information on defining a filter pattern, see Defining Filter Patterns. For more information on defining a filter pattern that can be used to filter folders, see Filtering Folders.

Types of File Filters

There are three types of file filters:

- **General** - Can be applied to any job type.
- **Synchronization Only** - Can be applied to File Collaboration jobs only. Select this filter type to exclude file types from being locked when a file open is detected on a participant in a File Collaboration job.
- **Locking Only** - Can be applied to File Collaboration jobs only. Select this filter type to exclude synchronization across the entire File Collaboration job so that only opens and closes are detected and acted on without any synchronization being performed.

For more information, see Creating and Applying File Filters.

Creating and Applying File Filters

You create a file filter in the **File and Folders** page of Preferences for a job type; the filter can then be applied to individual jobs of the same type. For example, a file filter created in Cloud Backup and Replication Preferences can be applied to any Cloud Backup and Replication job; a file filter created in Collab, Sync, and Replication Preferences can be applied to any File Collaboration, File Synchronization, or File Replication job. Multiple file filters can be applied to a single job.

In addition, there are also predefined filters that are applied to jobs; some of these predefined filters are automatically applied to certain job types.

For more information about creating a file filter, see:

- Creating File Filters for a Cloud Backup and Replication Job
- Creating File Filters for File Collaboration, File Locking File Replication, and File Synchronization Jobs

Predefined File Filters

In addition to defining your own file filters, there are predefined file filters that can be applied to jobs. The predefined filters vary per job type.

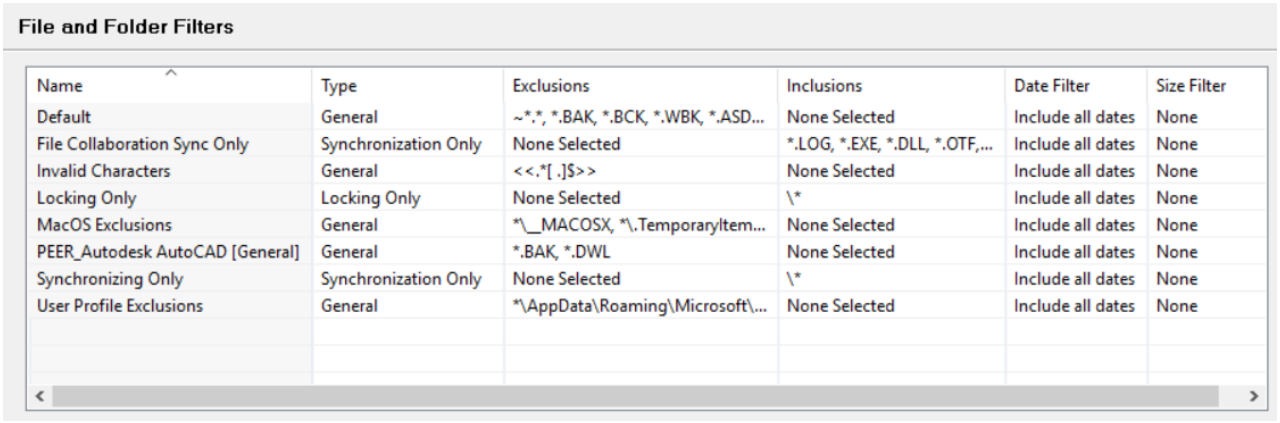

Two of the predefined filters, **Default** and **Invalid Characters**, are applied to all jobs by default. However, you can deselect a predefined filter for a specific job. Only the **Default** filter can be modified; none of the predefined file filters can be deleted.

In addition to these predefined filters, there are file types that are automatically excluded from a watch set for all job types.

To upgrade a predefined filter:

1. Select **Preferences** from the **Window** menu.

2. Expand **Cloud Backup and Replication** or **Collab, Sync, and Repl Summary** in the navigation tree, and then select **File and Folder Filters**.

   Existing file filters are listed in the **File and Folder Filters** table.

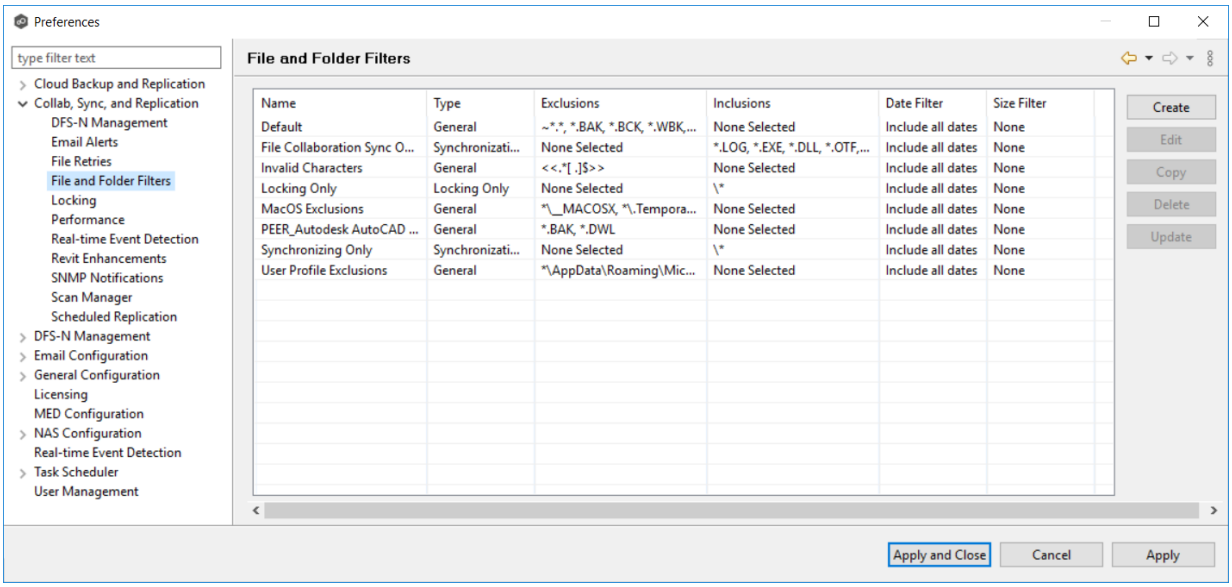

3. Select the filter to upgrade, and then click **Upgrade**.

   If an updated filter definition is available, a confirmation message lists the changes to the filter definition.

4. Click **OK** in the confirmation message.

   The new file filter is listed in the **File Filters** table and can now be applied to jobs.

Defining Filter Patterns

A **filter pattern** is a character string that defines a logical expression that is evaluated to determine which files and folders match the pattern. A file filter pattern can contain complex regular expressions and wildcards. See Folder Filters for more information about what a folder filter pattern can contain.

Files and folders that match an **exclusion pattern** are excluded from the watch set; files and folders that match an **inclusion pattern** are included the watch set. For example, in the following file filter definition, files with names ending in *.dotx are excluded and files with names ending with *.docx are included:

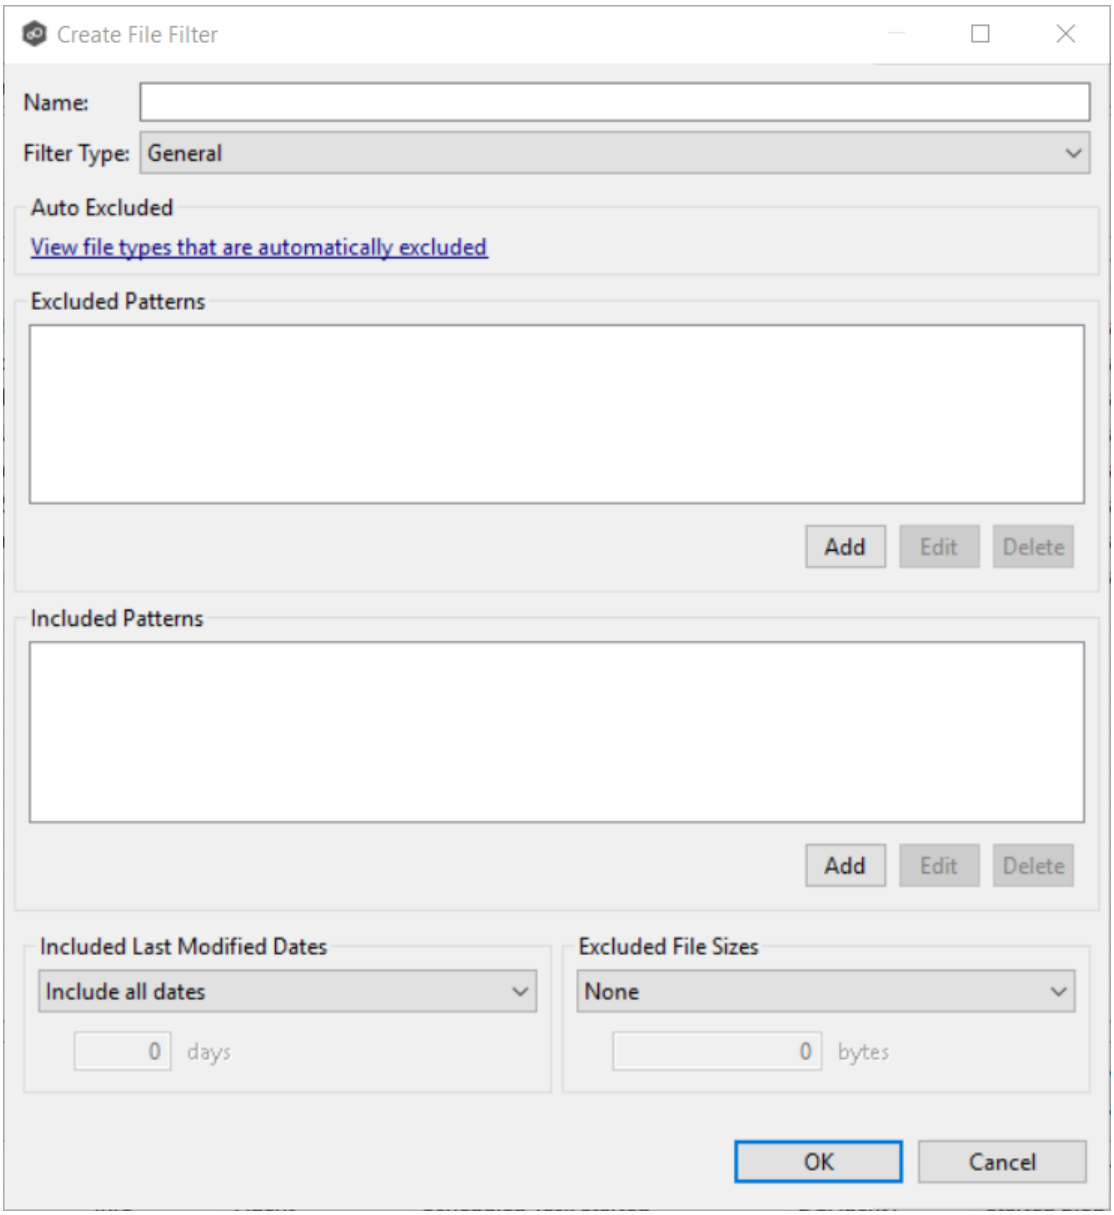

You can use the following wildcards in a file filter pattern to more easily cover well-known file extensions or names that follow established patterns.

*	Matches zero or more characters of any value

?	Matches one character of any value

The following examples show the use of a wildcard:

***.ext** Filter files that end with the **.ext** extension

ext* Filter files that begin with the string **ext**

ext Filter files that contain the string **ext**

The following expressions are automatically applied as exclusion patterns and cannot be modified.

File Type	Exclusion Pattern
Temporary files generated by common applications	~$*.* *.tmp *.$$$ Any file without a file extension, e.g., abcdefg
Explorer System Files	desktop.ini, thumbs.db, and Windows shortcut file, e.g., *.lnk

You will generally want to exclude all temporary files created by the applications you use so they are not propagated to the target hosts. For example, if your watch set contains files created by AutoCAD applications, you should create a file filter to exclude the temporary files created by these applications.

Typically, AutoCAD files have the following extensions:

.AC$

.SV$

.DWL

.BAK

To create a file filter that excludes these temporary AutoCAD files, you would add these extensions (with wildcards) to the **Excluded Patterns** field:

1. Click the **Add** button under the **Excluded Patterns** field.

   The **Add Exclusion Pattern** dialog appears.

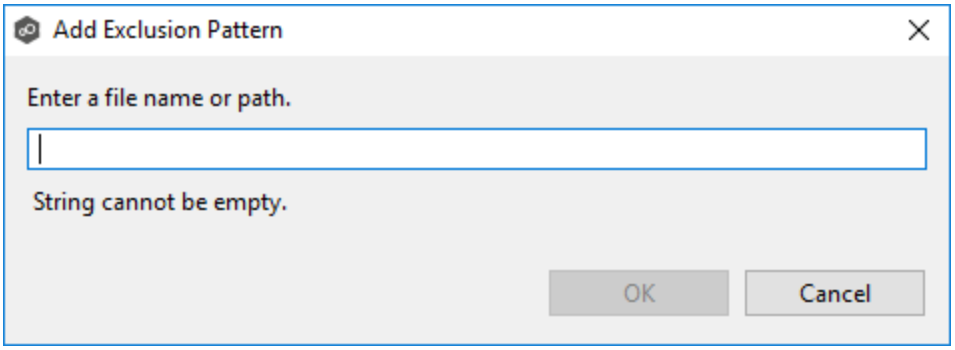

2. Enter ***.AC$**, and then click **OK**.

3. Repeat Step 2 to add ***.BAK**, ***.DWL*** and ***.SV$**.

   The patterns are listed in the **Excluded Patterns** field.

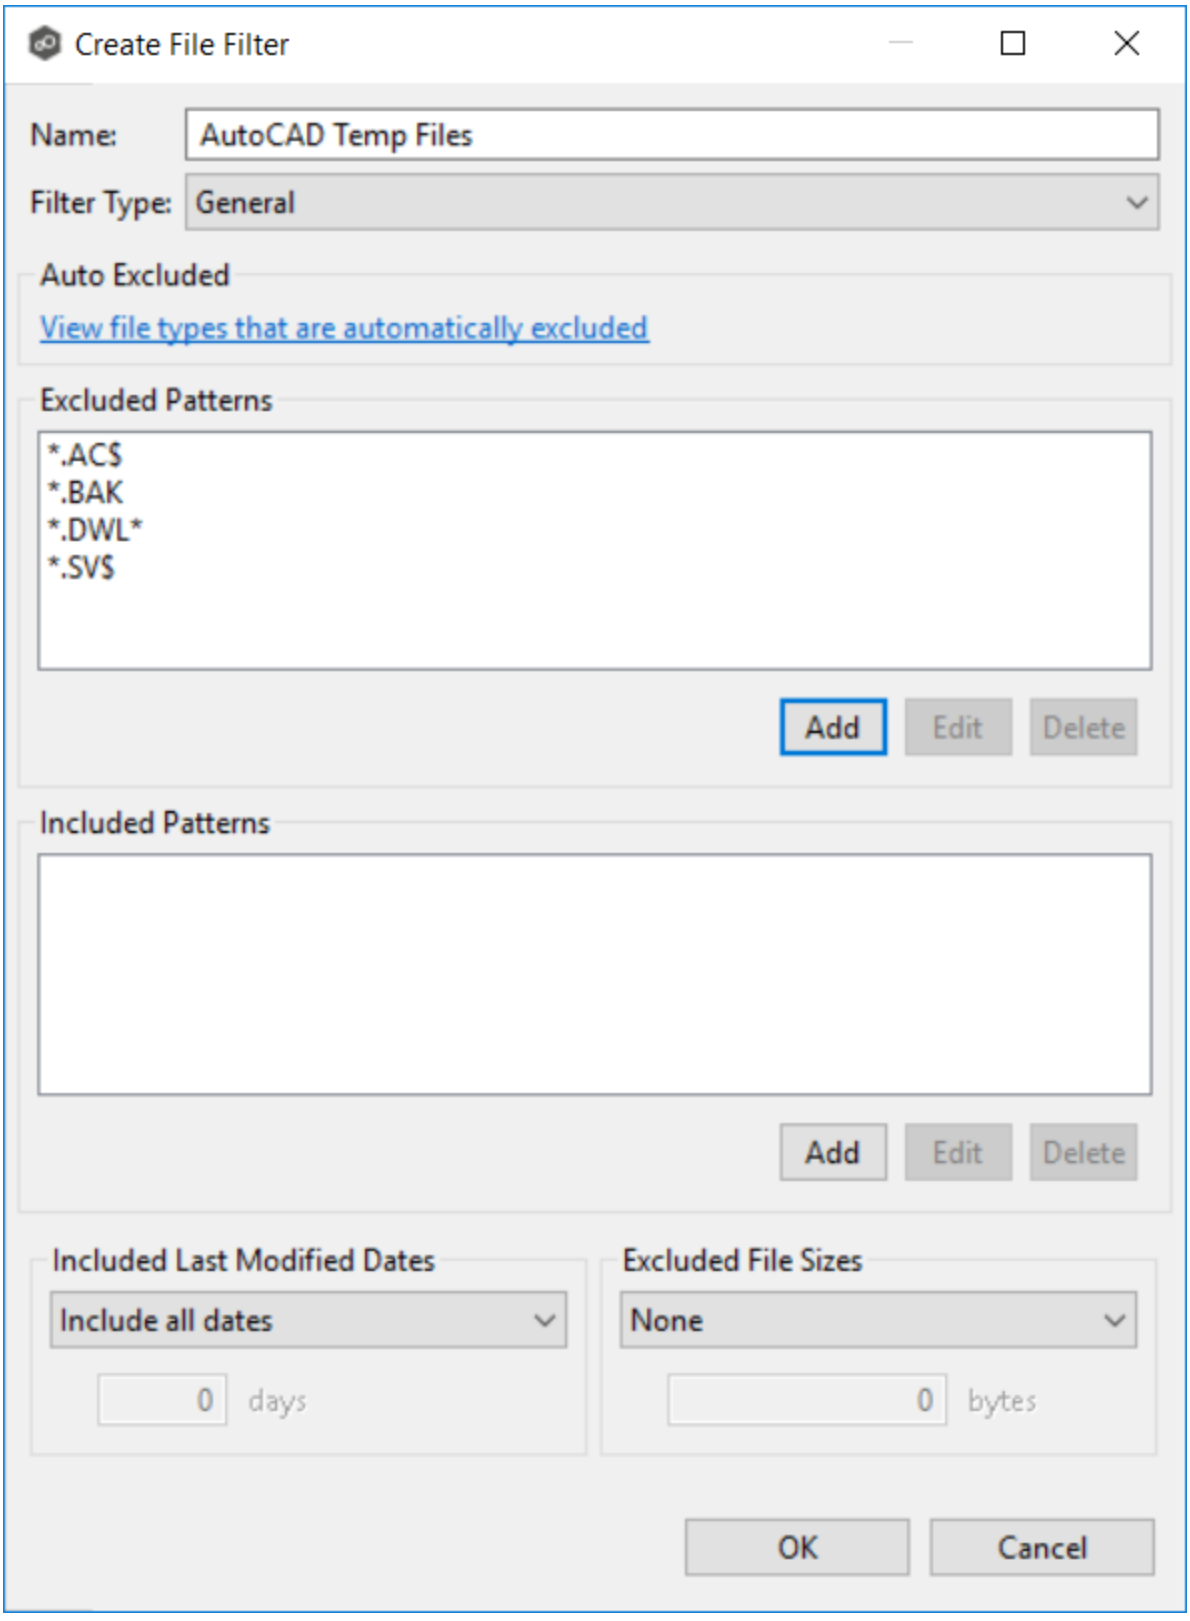

You have now created a file filter that excludes temporary AutoCAD files—all files ending in *.AC$, *.BAK, *.DWL*, or *.SV$ will be excluded from any running job that uses this filter.

Using Complex Regular Expressions in Filter Patterns

You can use complex regular expressions in filter patterns. Use the following format for a regular expression:

```
<<regEx>>
```

For example, the following filter pattern contains a regular expression that finds AutoCAD temporary files (atmp files):

```
<<^.*\\atmp[0-9]{4,}$>>
```

Using the following regular expression in an exclusion pattern excludes any path containing a folder **XX** that also contains a child folder **YY**:

```
<<^.*\\XX\\YY(\\.*$|$)>>
```

The following files and folders MATCH the above expression:

```
\projects\xx\yy
\accounting\projects\xx\yy\file.txt
\accounting\projects\xx\yy\zz\file.txt
```

The following files and folders DO NOT MATCH the above expression:

```
\projects\accounting\file.txt
\projects\xx\y
\projects\xx\yyy\file.txt
\accounting\projects\xx\file.txt
\accounting\projects\yy\xx\zz\file.txt
```

For a good reference on regular expressions, see http://www.regular-expressions.info/reference.html

In addition to filtering on file names, file extensions, folder paths, or partial path wildcard pattern matching, you can filter based on a file's last modified date:

- Peer Management Center supports filtering on a file's last modified date but does not support filtering on a folder's last modified date.

- If you have a folder hierarchy that contains files that are all being filtered based on last modified date, then all folders will still be created during the initial scan process on all hosts.

- If a file is excluded from collaboration based on its last modified date, then the initial scanning process will not synchronize the file even if the file's last modified time and size do not match, or the file does not exist on all hosts. However, the file will be synchronized, if and when the file is modified in the future, and if a user deletes or renames the file on any host, the file will be deleted or renamed from all other hosts where the file exists.

- A file filter cannot combine filtering on last modified date with inclusion or exclusion patterns or file size. The last modified date is the sole criteria used to identify matching files.

Options for Included Last Modified Date Filter

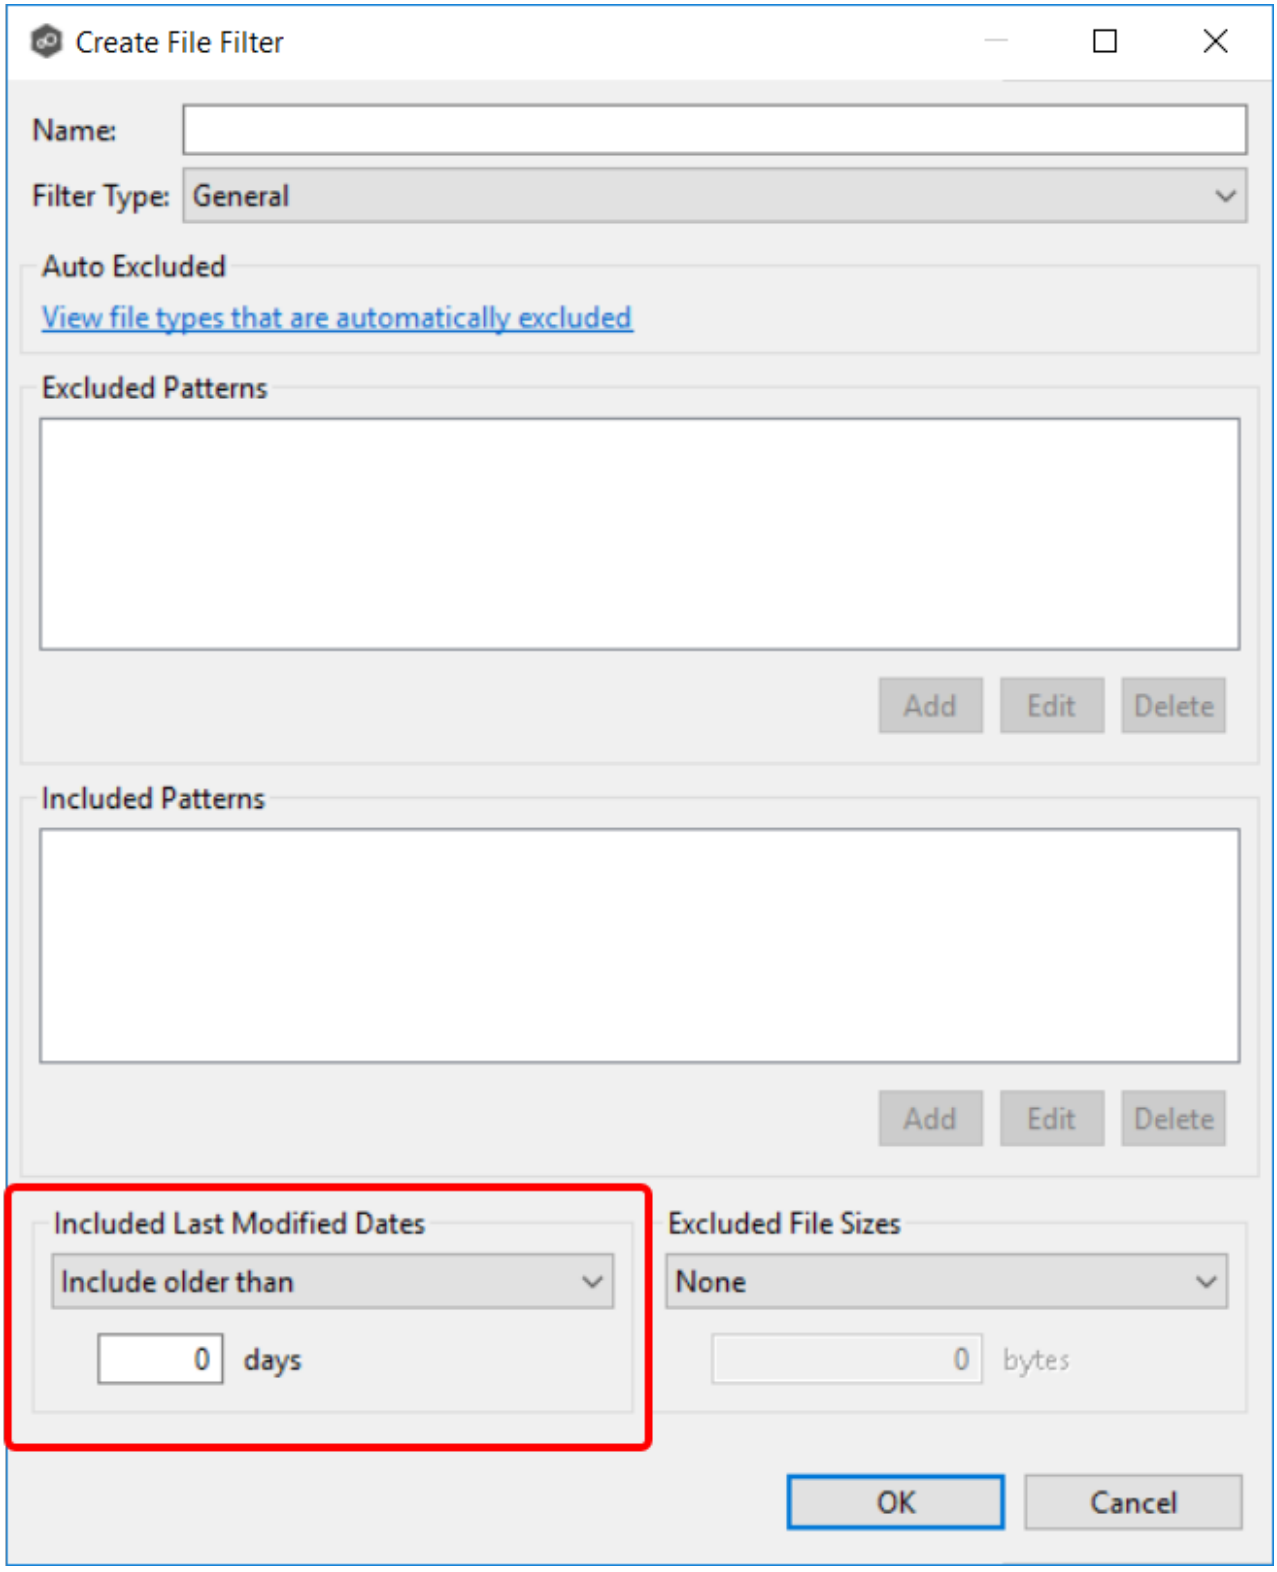

Include all dates	This is the default option and will include all files regardless of last modified date.
Include today and past	Includes all files whose last modified date are more recent than the specified number days. For example, you can exclude all files that have not been modified within the last year (365 days).
Include older than	Includes all files whose last modified date are older than the specified number days.

In addition to filtering on file names, file extensions, folder paths, or partial path wildcard pattern matching, you can filter based on the size of an individual file, excluding files that are greater or less than a specified size:

- Peer Management Center does not support filtering on a folder's total size.
- If you have a folder hierarchy that contains files that are all being filtered based on size, then all folders will still be created during the initial scan process on all hosts.
- If a file is excluded from collaboration based on size, then the initial scanning process will not synchronize the file even if the file's last modified time and size do not match, or the file does not exist on all hosts. However, if a user deletes or renames the file on any host, the file will be deleted or renamed from all other hosts where the file exists.
- You cannot define a file filter that combines filtering on file size with inclusion or exclusion patterns or last modified date. The file size is the sole criteria used to identify matching files.

Options for Excluded File Sizes

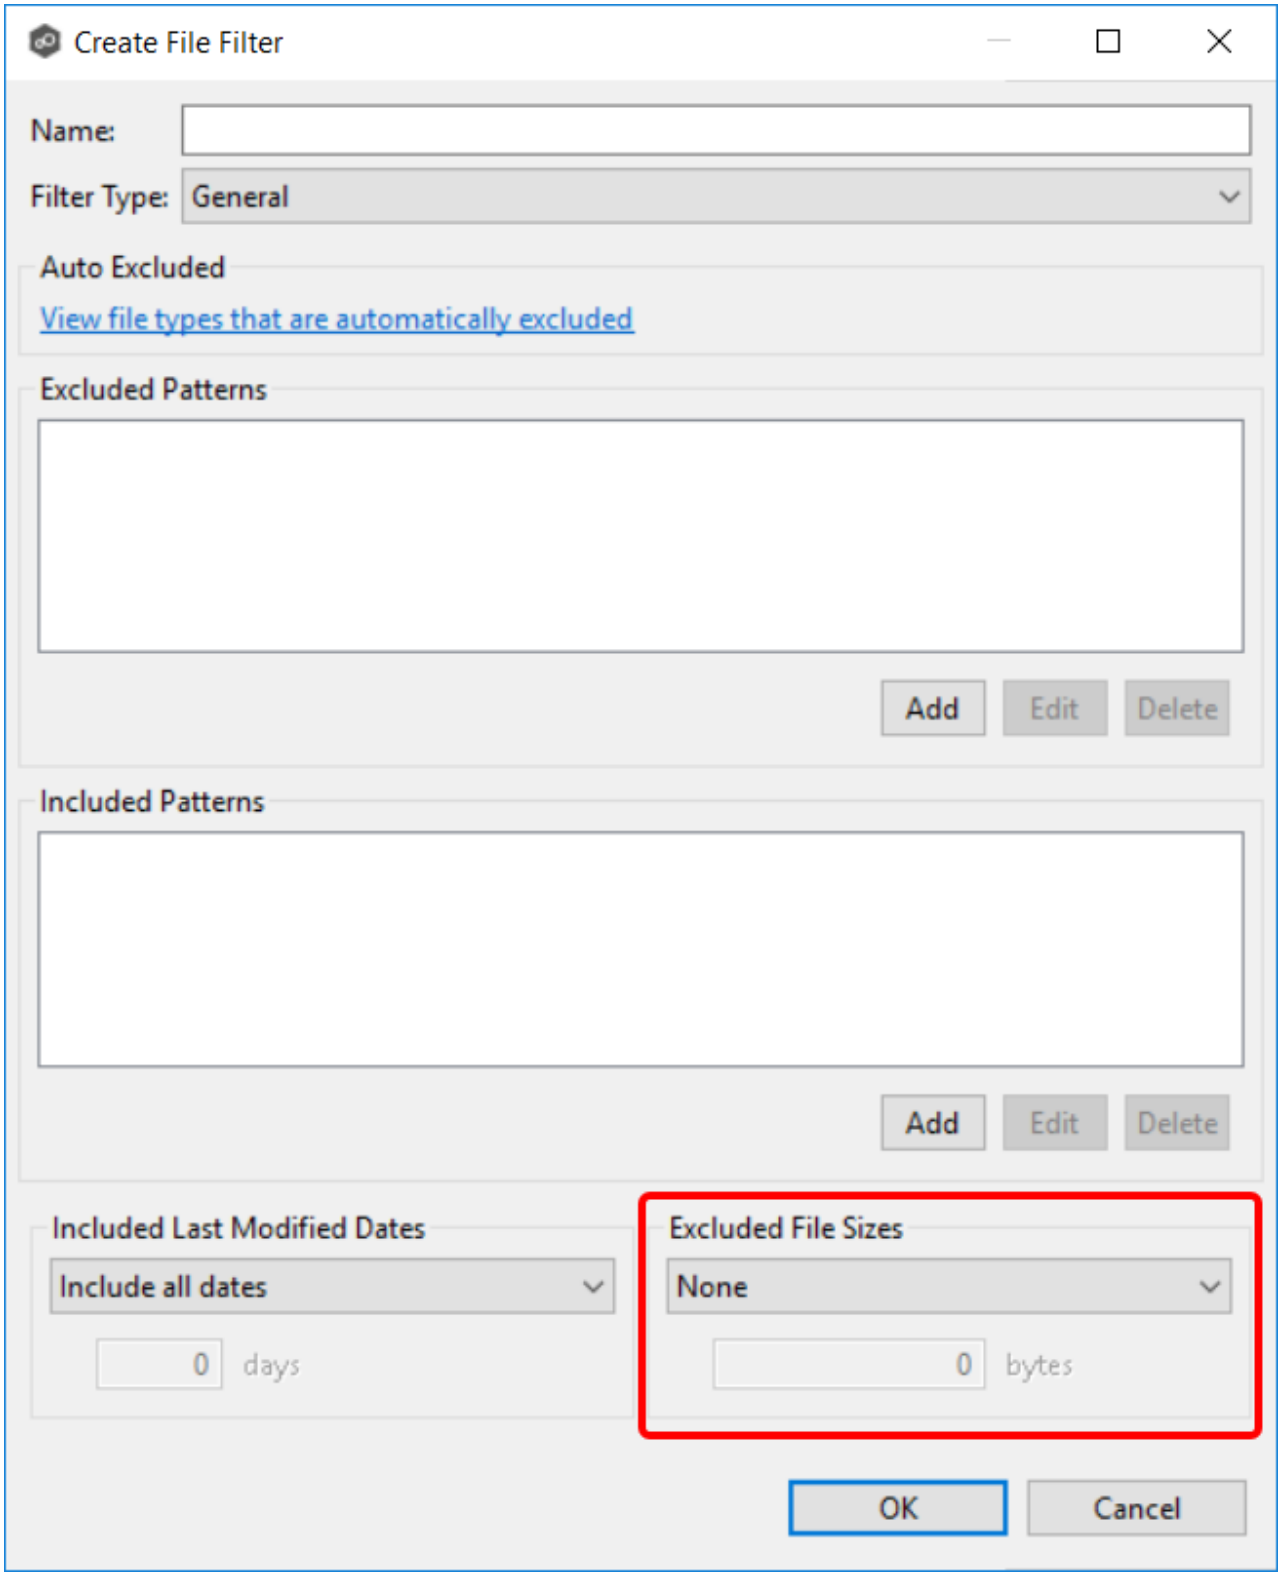

None	Default option. Select this option to include all files regardless of file size.
Exclude files greater than or equal to	Select this option to exclude all files whose size is greater than or equal to the specified number of bytes. For example, you can configure a job to exclude all files greater than 1 GB.
Exclude files less than	Select this option to exclude files whose size is less than the specified number of bytes.

Filtering Folders

In addition to creating file filters, you can create folder filters. Folder filters allow you to include and exclude folders from a job's watch set. See Folder Filter Examples for examples of folder filters. Folder filters are created in the same way as file filters.

Reduce the Number of Jobs Using Folder Filtering

For management purposes, we recommend keeping the total number of jobs as low as possible. Using folder filters, you can reduce the total number of jobs without sacrificing efficiency. This process involves analyzing all existing jobs, identifying all the folders and hosts that will be collaborating, and consolidating them into fewer jobs by watching a few root folders at a higher level. Filters will then be added to include or exclude only the folders of interest.

Folder Filter Syntax

When defining a filter pattern to use on folders, use the following syntax:

\Folder or **\Folder*** or **\Folder***

Presently, Peer Management Center supports included expressions for a full folder path only and does not support wildcard matching on parent paths. For example, the following expression is not valid:

\Folder*\Folder

Example of a Simple Folder Filter

The following example reduce the number of existing jobs from four to two:

		Server 1		Server 2	
		Drive D	Drive E	Drive D	Drive F
Old Jobs	Job 1	D:\General		D:\General	
	Job 2		E:\Common		F:\Common
	Job 3	D:\Projects		D:\Projects	
	Job 4		E:\Documents		F:\Documents

After consolidation:

				Filter Option 1	Filter Option 2
		Server 1	Server 2	INCLUDE	EXCLUDE
New Jobs	Job 1	D:\	D:\	\General*	All other files
				\Projects*	
	Job 2	E:\	F:\	\Common*	All other files
				\Documents*	

Jobs 1 and 3 were merged into a single job watching the root D drive on both servers while using Filter Option 1 or 2.

Jobs 2 and 4 were merged into a single job watching the root E drive on Server 1 and the root F drive on Server 2 while using Filter Option 1 or 2.

Please note the following regarding regular expressions:

- Peer Management Center does not support the ability to use regular expressions for multi-level folder inclusions such as \Level1\Level2\FolderName.
- Peer Management Center does not currently support the ability to filter on certain parts of a path, like \Folder*\Folder and \Folder*\.

Additional Examples of Folder Filters

To exclude a specific folder from anywhere within the watch set:	*\FolderName *\FolderName\FolderName
To exclude a specific folder from the ROOT of the watch set:	\FolderName \FolderName\FolderName
To exclude folders that end with a specific name from anywhere within the watch set:	*FolderName\
To include a specific folder from the root of the watch set:	\FolderName \FolderName\FolderName

File Filter Usage Notes

Conflicting Patterns

Since inclusions and exclusions patterns are expressed separately, it is possible to submit conflicting patterns. The pattern evaluator addresses this by exiting when a file is determined to be excluded. Therefore, exclusions patterns override inclusion patterns.

Rename Operations

Rename operations may subject files to an inclusion status change. Renaming a file out of the watch set will trigger a target deletion, while renaming into the Rename operations may subject files to an inclusion status change. Renaming a file out of the watch set triggers a target addition.

Folder Deletions

Folder deletions only affect included files, possibly leading to folder structure inconsistencies. When a session participant deletes a folder, the target outcome will vary depending on whether excluded files are present. Folder deletions are propagated in detail to the targets as to the exact files that have been affected.

List Filters

Peer Management Center provides the ability to filter lists throughout the Peer Management Center interface. List filters can help you quickly find jobs, Agents, and sort through summary reports

To use a list filter, enter a filter expression in the filter expression box. The search results of your filter are displayed in the window below the expression.

You can save the list filters and reuse them. For more information, see Saving and Managing List Filters. This is useful when you frequently use the same list filter or when you create complex list filters.

Use the **Ctrl + Space** keyboard shortcut to list all possible list filters and predefined labels, which can be selected to refine your search quickly.

Basic Filter Expressions

The simplest filter expressions contain words you are looking for. For example, to find all items related to sales, simply type the word *sales* in the filter expression box. All items from the list that contain the word *sales* in their name, tag names, or tag categories will be displayed, and all other items will be hidden. The agent attribute fields (see attr below) are not included in generic searches.

If you want an exact word match or the words contain a space, enclose the terms in double quotes. For example, if you want to search for the words *North America*, the two words must be contained in double quotes. If you want to search for the word *agent* only without showing *USAgent* or *Agent2015* in results, the word *agent* must be contained in double quotes.

For information about creating more complex filter expressions using operators and labels, see Creating Complex Filter Expressions.

Predefined List Filters

- Default job filters include **Failed Jobs**, **Jobs with Backlog**, and **Running Scans**.
- Default Agent filters include **Connected** and **Disconnected** (e.g., filter:"Running Scans").

Creating Complex Filter Expressions

You can create more sophisticated list filters by using operators and labels.

Using Operators

Operators allow you to combine multiple simple expressions into a single compound expression. Supported operators are: OR, AND, and NOT. For example, typing tag:Americas AND sales in the Filter Expression will show only Agents with the word *Americas* in their tag(s) AND the word sales in their name, tags, or tag categories. Parentheses can be used to build more complex expressions by grouping simple expressions.

Using Labels

Use predefined labels to specify in which field your filter word should appear. Use the following format to take advantage of labels in your filter expression:

<label>:<search string>.

List of possible labels include:

name	List only items that match the string (e.g., name:"Design Data")
tag	Show only items with the word specified in their tag(s) (e.g., tag:Americas)
cat	Search for items that have been assigned a specific category (e.g., search for Jobs that were categorized as Design - cat:Design)
host	Filter through Jobs and list only those that contain the host in the list of job participants (e.g., host:WIN12R2A)
attr	Search for the specified string in the following Agent fields: Connection Status, Operating System, JVM Architecture, and Agent Version (e.g., attr:x86)
filter	List items that have been assigned a default or user-created filter.

Examples

Example 1: Show all Agents with the word *Sales* in their name, tag name, or tag category:

Sales

Example 2: Show all Agents with a tag that has *North America* in the tag name and *Location* in the tag category:

cat:Location AND tag:"North America"

Example 3: Show all Agents with the word *Sales* in their name, tag name, and tag category and with a tag that has *North America* in the tag name and *Location* in the tag category.

```
Sales AND (cat:Location AND tag:"North America")
```

Saving and Managing List Filters

Throughout the Peer Management Center interface, you will have the opportunity to save your filter expression by clicking the **Manage, Save, and Load filters** button, usually located above the **Filter Expression** field or in the **Actions** drop-down menu. The **Manage, Save, and Load filters** button is available in the Jobs view panel, the Agent Summary view, the and the Collaboration Summary panel.

To remove a list filter and show all items in the list, click the pencil icon to the right of the filter expression.

Logging and Alerts

PeerGFS performs an extensive amount of logging to track events and activities processed by PeerGFS. The results are stored in log files that are useful for troubleshooting and analytics. PeerGFS tracks and logs many types of information and activities, including file events, preferences, job-specific configuration files, and analytics files.

Many of the log files have an .log extension; these are text files that can be opened in a text editing application. Other log files are stored in other file formats such as .xml, .csv, and .prefs. Log files are stored in the **workspace** folder in the Peer Management Center and Agent installation directories:

PMC Log Files	Agent Log Files
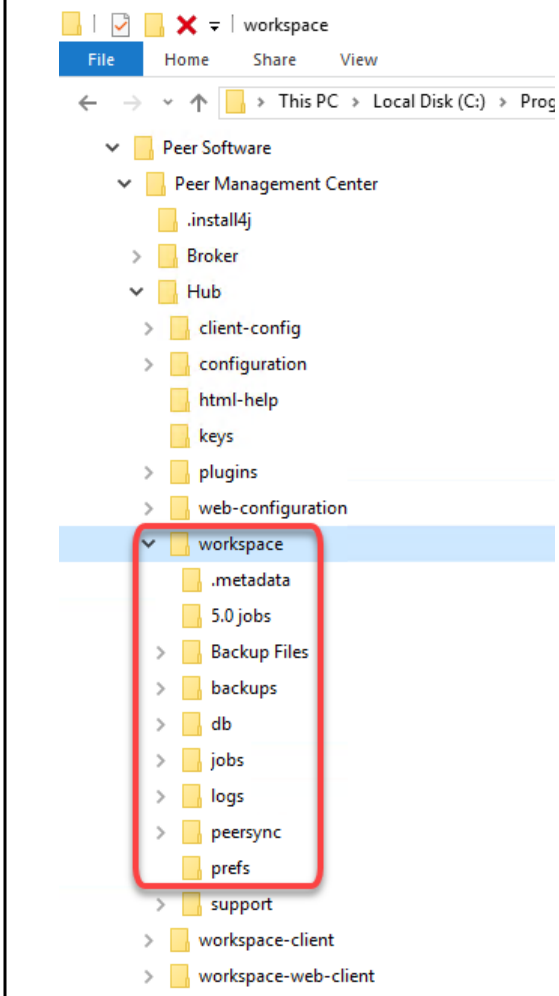	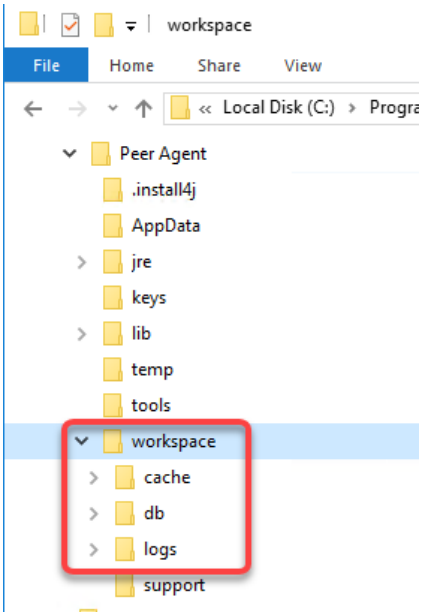

If you want to review log files for troubleshooting or analytical purposes, you can retrieve them as a single, compressed file, which is then stored in the **support** folder in the Peer Management Center installation directory. The retrieval process compiles the various log files into a single zip file that is easy to review and send to others for review. When retrieving log files, you have various options, such as choosing which log files are included, whether to encrypt log files (which may contain sensitive information), and whether to have the zip file automatically sent to Peer Software Support.

Retrieving Log Files

To retrieve log files:

1. Open Peer Management Center.
2. From the **Help** menu, choose **Retrieve PMC/Agent Logs**.

   The **Retrieve PMC/Agent Log Files** dialog is displayed.

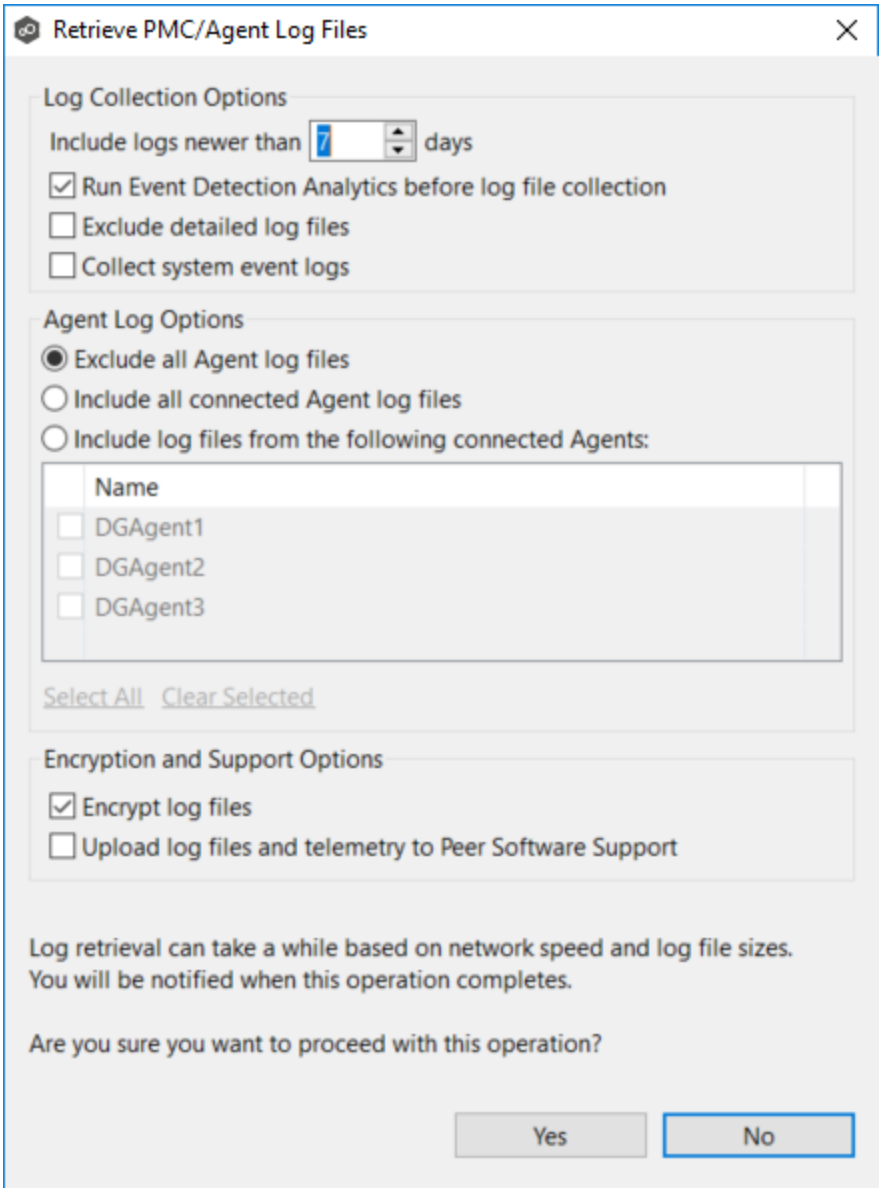

3. Select the logging options.

   There are three sets of options:

- Log Collection Options
- Agent Log Options
- Encryption and Support Options

4. Enter your contact information and a description of the problem:

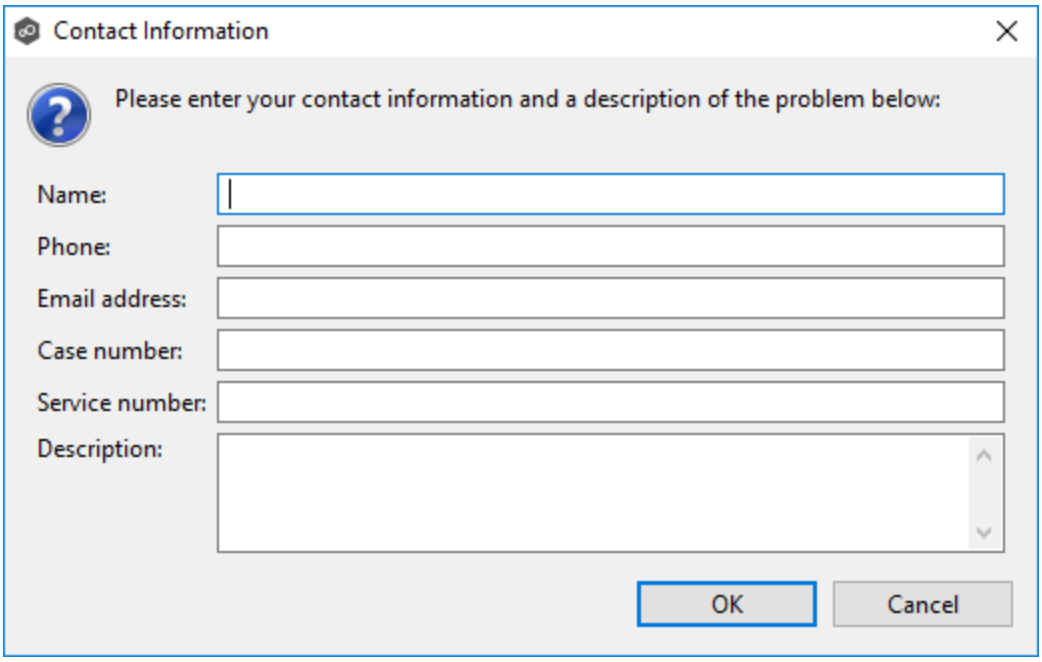

This information will be sent to Peer Software Technical Support.

5. Click **Yes** to start the log retrieval process.

It may take some time for the log files to be collected and compiled into a single, compressed file. When the retrieval is finished, a message is displayed.

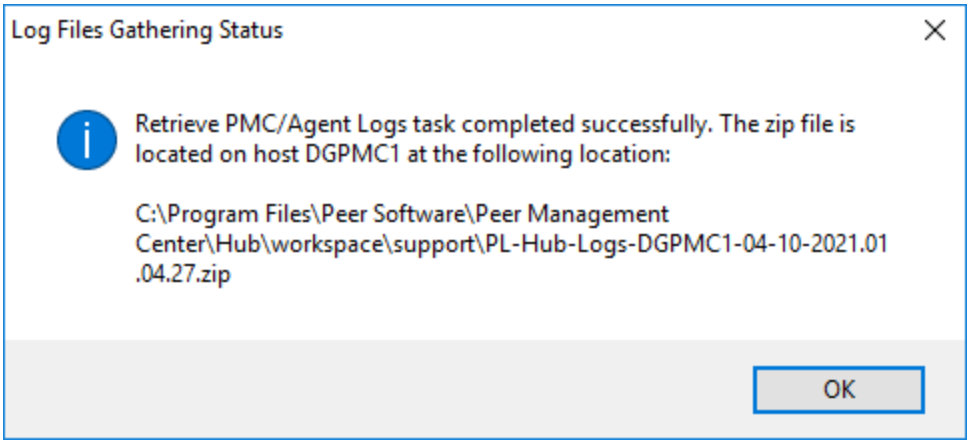

6. Click **OK**.

   The retrieved log file is stored as a zip file in the **workspace/support** subfolder in the Peer Management Center installation directory.

Log Collection Options

Option	Description
Include logs newer than X days	Use this option to restrict the logs retrieved to a certain time period.
Run Event Detection Analytics before log file collection	Select this option to run event detections analytics immediately before the log files collected. PeerGFS can perform event detection analysis every night; however, this option ensures that the log bundle contains the most up-to-date analytics.
Exclude detailed log files	Select this option to reduce size of log uploads. By default, log collection includes all Peer-generated log files (for example, event log files, activity log files, and Agent output logs if Agents are selected in the Agent Log Options section). If this option is selected, detailed log files will not be uploaded, thus reducing Peer Technical Support's ability to troubleshoot using log files.
Collect system event logs	Select this option to retrieve Windows event logs.

Agent Log Options

Option	Description
Do not include any agent log files	Select this option if you do not want to retrieve log files for any Agent.
Include all connected agent log files	Select this option if you want to retrieve log files for all connected Agents.
Include the log files from the connected agents selected below	Select this option if you want to retrieve log files for selected connected Agents.

Encryption and Support Options

Option	Description
Encrypt log files	Select this option if you want to encrypt the log files in the zip file. We suggest checking this option if you are uploading the log bundle to Peer Software Support.
Automatically upload log files and telemetry to Peer Software Support	Select this option if you want to automatically upload the zip file containing the log files and telemetry information to Peer Software Support. No file data will be uploaded—only Peer-specific configuration, logs, etc.

Job Logs and Alerts

You can configure the logging and alert settings for a job when you edit a job. By default, all file collaboration and synchronization activity are logged for all severity levels. You can enable or disable file event logging, as well as select the level of granularity.

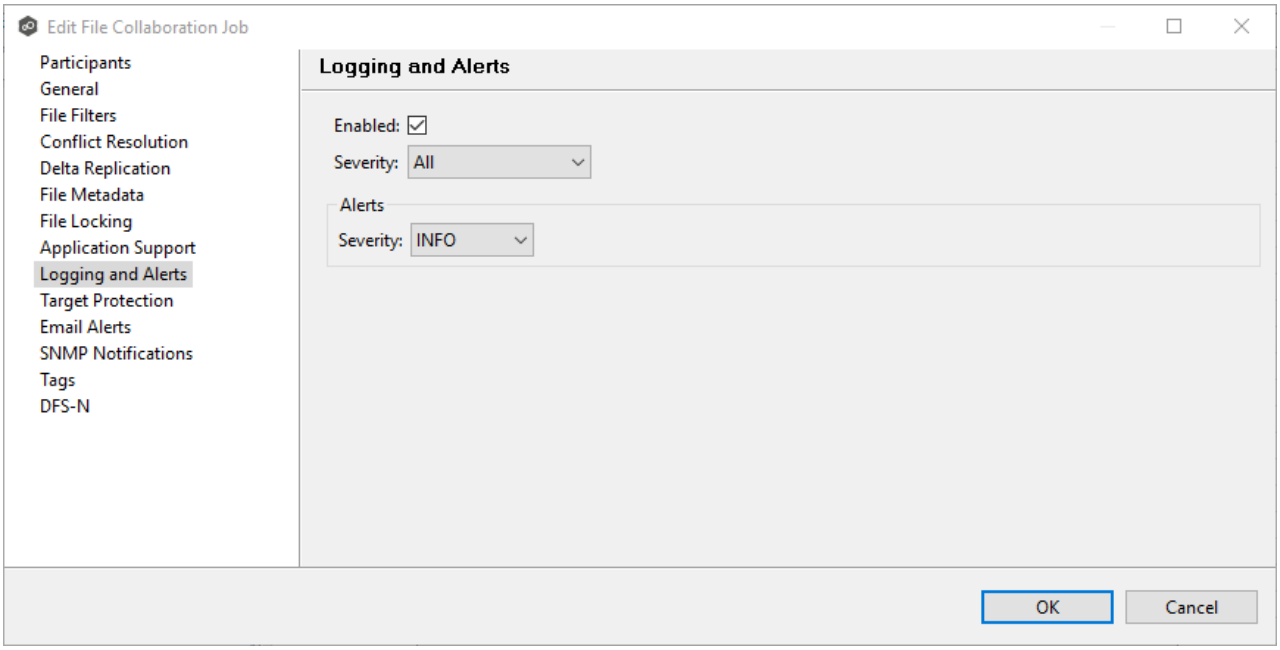

Log Entry Severity Levels

Level	Description
Informational	Informational log entry, e.g., a file was opened.
Warning	Some sort of warning occurred that did not produce an error but was unexpected or may need further investigation.
Error	An error occurred performing some type of file activity.
Fatal	A fatal error occurred that caused a host to be taken out of the session, a file to be quarantined, or a session to become invalid.

Job Alerts

Various types of alerts will be logged to a log file and to the **Alerts** table located within the job's runtime summary view for the selected job. Each job will log to the **fc_alert.log** file located in the **Hub\logs** subdirectory within the Peer Management Center installation directory.

The default log level is WARNING, which will show any warning or error alerts that occur during a running session. Depending on the severity of the alert, the job may need to be restarted.

SNMP Notifications

Overview

Peer Management Center provides support for SNMP v1 messaging. A SNMP notification notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. When an SNMP notification is applied to a job, a SNMP trap is sent to the destination IP address or hostname whenever a selected notification type is triggered by the job. The available notification types depend on the job type.

When you create a job, you can select an existing SNMP notification to apply to the job or you can create a new notification and apply it to the job. You cannot modify or delete a notification while it is applied to a job. An SMNP notification can be applied to multiple jobs of the same type. SNMP notifications are defined in the preferences for a job type. An SNMP notification can be applied to all job types except File Replication.

Note that before Peer Management Center can send SMNP notifications on behalf of any job, you must configure some SNMP settings.

Managing SMNP Notifications

To manage SMNP notifications:

1. Select **Preferences** from the **Window** menu.

   The **Preferences** dialog appears.

2. Select the job type from the navigation tree and expand it.

3. Select **SMNP Notifications** from the navigation tree.

   The **SMNP Notifications** page lists existing SMNP notifications for that job type. You can create, edit, copy, and delete notifications.

Tags

Tags can be used to categorize resources and customize a user's workspace or perspective. Tagging helps when managing large number of resources.

You can assign tags to:

- Jobs
- Resources
- Web roles
- Agents

You can also assign resources to tags. See Using Tags to Filter Resources.

Creating Tags and Categories

Tags and categories are created in Tags Configuration in **Preferences**. The **Assign Tags** dialog also offers the option to create tags and categories.

Assigning Tags

You can:

- Assign tags during job creation
- Assign tags while editing an existing job
- Assign tags to one or more resources
- Assign tags to web roles
- Assign resources to one or more tags

Assigning Tags to Jobs

- During job creation - You can assign tags during the creation of a job from the Tags page of the job creation wizard.
- During job editing - You can assign tags to individual jobs by right-clicking on the job, selecting **Edit Job(s)**, and navigating to the **Tags** page of the job editing wizard.

Assigning Tags to Resources

To assign tags to one or more resources:

1. Click the **Assign Tags** button from the main view, Jobs view, or Agent Summary view toolbars.
2. In the **Assign Tags** dialog, click the **Tags** radio button.
3. Select the tag that needs to be assigned to one or more resources.
4. Click the **Edit** button.

   The **Assign/Unassign resources** dialog appears.
5. In the **Unassigned Resources** table, select the resources to be assigned the selected tag, and then click the right-arrow button (Add One) to move it to the table on the right side.

   Tip: To select multiple resources, press the Shift key on the keyboard when selecting resources.
6. Click **Save**.
7. Repeat the preceding steps for all the tags that need to be assigned to one or more resources.

Assigning Tags to Web Roles

Web roles can be assigned tags that customize a user's Jobs view when they log in via the web client. For example, in a very large deployment scenario, a user that is part of the Help Desk role can be assigned tags that limit their view to only jobs that are part of their region.

To assign tags to user roles:

1. Create tags and categories as outlined in Step 1 above.
2. Assign tags to one or more jobs as outlined in Step 2 above.

3. Go to User Management in the Preferences page.

4. Select the desired role to which you wish to assign specific job tags.

5. Click the **Edit** button.

6. In the **Tags** window, from the **Unassigned Tags** table on the left side, select the tags that need to be assigned the selected role, and then click the right-arrow button (Add one) to move it to the table on the right side (hold down the Shift key on the keyboard when selecting tags to select more than one).

7. Click **OK** to commit your changes and close the dialog, and then close the **Preferences** page.

   The user will see only the jobs that were tagged in the user's role.

Assigning Resources to One or More Tags

To assign resources to one or more tags:

1. Click the **Assign Tags** button from a summary view, Jobs view, or Agent Summary view toolbar.

2. In the **Assign Tags** dialog, click the **Resources** radio button.

3. On the left-hand side, click inside the **Resource Name Filter** or **Type Filter** fields and press the CTRL + Space keys on the keyboard to list all possible filters and predefined labels, which can be selected to refine your search quickly.

4. Select the resource that needs to be assigned to one or more tags.

5. Click the **Edit** button to the right.

6. From the **Unassigned Tags** table on the left side, select the tags that need to be assigned the selected Resource, and then click the right-arrow button (Add one) to move it to the table on the right side (hold down the Shift key on the keyboard when selecting tags to select more than one).

7. Click the **Save** button to commit your changes, and then close the dialog.

8. Repeat the preceding steps for all the resources that need to be assigned to one or more tags.

Using Tags to Filter Resources

You can use tags to filter resources:

- Filter jobs
- Filter agents

To filter resources using tags, use the tag label in any list filter field throughout the Peer Management Center interface.

Filter Jobs

To filter through a large list of jobs, use the filter field located below the toolbar buttons in the **Jobs** view. For more details on how to filter through resources, see List Filters.

Example:

Show all jobs with a tag that has "North America" in the tag name and "Location" in the tag category:

```
tag:"North America" AND cat:Location
```

Filter Agents

To filter through a large list of Agents, use the **Filter** field located below the toolbar buttons in the Agent Summary View panel. For more details on how to filter through resources, see Filter Expressions.

Web Client Users

Peer Management Center offers two interfaces:

- A rich client interface: Rich client users have access to all Peer Management Center functionality. The rich client is accessible only on the server where Peer Management Center is installed
- A web client interface: Web client users access to Peer Management Center functionality is controlled by their web role.

Web client users can be divided into two categories based on how their access to the web client is authenticated:

- Internal users

- Active Directory users and groups

For information about managing web client users, see Managing Web Client Users.

Internal Users

An **internal user** is one whose access to the Peer Management Center web client is authenticated by an internal Peer Management Center database rather than through Active Directory.

For information about managing internal users, see Managing Internal Users.

Internal Users and Web Roles

When you add internal users to Peer Management Center, you need to specify their access to Peer Management Center functionality by assigning them a standard web role or a custom web role. A **web role** is a set of permissions that specifies the appropriate level of access to Peer Management Center functionality. For example, some users will need the ability to create and edit jobs, while other users may need only to view job summaries. Assign the most suitable role to each user, giving them the most appropriate level of control and not more.

For more information about web roles, see Overview of Web Roles.

Default Internal User: The admin User

There is a default internal user who has access to all Peer Management Center functionality available in the web client: the **admin** user. This user does not need to be created. This internal user has the following properties:

Username	admin
Password	password This should be changed immediately upon first log-in.
Web Role	Administrator

Unlike other internal users, the admin user cannot be renamed or deleted, nor can its role be changed. However, for security reasons, the password should be changed immediately.

Active Directory Users and Groups

An Active Directory (AD) user or group is one whose access to Peer Management Center is authenticated through Active Directory. Adding an Active AD user or group authenticates and authorizes that user or group members to use Peer Management Center. The AD user or group must already exist in Active Directory prior to adding the user or group to Peer Management Center.

Active Directory users won't be able to access the web client until Active Directory authentication is configured in Peer Management Center.

For information about managing Active Directory users and groups, see Managing Active Directory Users.

Active Directory Users and Web Roles

When you add Active Directory users and groups to Peer Management Center, you need to specify their access to Peer Management Center functionality by assigning them a standard web role or a custom web role. A **web role** is a set of permissions that specifies the appropriate level of access to Peer Management Center functionality. For example, some users will need to the ability to create and edit jobs, while other users may need only to view job summaries. Assign the most suitable role to each user, giving them the most appropriate level of control and not more.

For more information about web roles, see Overview of Web Roles.

Overview of Web Roles

All users that access Peer Management Center through the web client must have an assigned **web role**. A web role is a set of permissions that specifies the appropriate level of access to Peer Management Center functionality. For example, some users will need to the ability to create and edit jobs, while other users may need only to view job summaries.

Web client users can have a standard web role or a custom web role. In contrast, a user who accesses Peer Management Center through the rich client does not have a web role. All Peer Management Center functionality is accessible to a rich client user.

For more information about web roles, see Managing Web Roles.

There are three standard web roles, each with a predefined set of permissions:

- **Administrator** - This role has complete access to all functionality found in the Peer Management Center's rich client.
- **Power User** - This role has view-only access to jobs and the **Agent Summary** view; this role cannot create, edit, or delete jobs, access settings in Preferences, or assign tags.
- **Help Desk** - This role has view-only access to jobs. Specifically, Help Desk users are limited to view-only access to the following:
  - The Jobs view
  - The runtime views
  - The **Summary** and **Session** tabs of each job.

  In addition, Help Desk users have read-write access to the **Quarantines** tab of each job, with the ability to release conflicts for any running jobs.

Standard web roles cannot be modified or deleted, with one exception: tags can be assigned to standard roles. For a list of the permissions associated with standard web roles, see Standard Web Role Permissions.

Standard Web Role Permissions

Each of the three standard web roles (Administrator, Power User, and Help Desk) has permission to access the resources shown in the following table.

Functionality	Administrator	Power User	Help Desk
Advisory Alert View	Edit	Edit	
Broker Statistics Action	Edit	Edit	
Collaboration Summary View	Edit	Edit	
Configuration Interface	Edit	Edit	

Functionality	Administrator	Power User	Help Desk
Event Analyzer Configuration Interface	Edit	Edit	
Event Analyzer Log View	Edit	Edit	View-only
Event Analyzer Participant view	Edit	Edit	
Event Analyzer Runtime Summary Interface	Edit	Edit	View-only
Event Log View	Edit	Edit	
Expression List Dialog	Edit	Edit	
File Conflict View	Edit	Edit	Edit
File Sync Advisory Alert View	Edit	Edit	
Folder Analyzer View	Edit	Edit	View-only
Job Alert View	Edit	Edit	
Job View	Edit	Edit	View-only
Log Dump Action	Edit	Edit	
Memory Dump Action	Edit	Edit	
New Job Action	Edit		
Participant View	Edit	Edit	

Functionality	Administrator	Power User	Help Desk
Permission Mode	Edit	Edit	
PMC Alert View	Edit	Edit	
PMC Download Agent	Edit	Edit	
PMC Refresh Perspective	Edit	Edit	
PMC View Progress	Edit	Edit	
Preferences	Edit		
Runtime Summary Interface	Edit	Edit	View-only
Session View	Edit	Edit	View-only
Status Agent Tree View	Edit	View-only	
Tag Resources Dialog	Edit		
Thread Dump Action	Edit	Edit	

PeerSync Management Job Permissions

The following table outlines the permissions for PeerSync Management jobs.

Functionality	Administrator	Power User	Help Desk
PeerSync Summary View	Edit	Edit	Edit

Functionality	Administrator	Power User	Help Desk
PeerSync Job Stats View Part View	Edit	Edit	
PeerSync Configuration Interface	Edit	Edit	View-only
PeerSync Job Stats View	Edit	Edit	Edit
PeerSync Update Log View	Edit	Edit	Edit
PeerSync Add Log View	Edit	Edit	Edit
PeerSync File Conflict View	Edit	Edit	Edit
PeerSync Runtime Summary Interface	Edit	Edit	View-only
PeerSync Participant View	Edit	Edit	
PeerSync Advisory Alert View	Edit	Edit	
PeerSync Messages Log View	Edit	Edit	Edit
PeerSync Delete Log View	Edit	Edit	Edit
PeerSync Event Log View	Edit	Edit	

A custom web role allows you to customize and fine-tune the access that a user has to Peer Management Center resources. This is useful if you have multiple types or levels of users that need different types of access. For example, if you have multiple tiers of help desk staff,

creating custom roles based on the standard Help Desk role allows you to provide them varying levels of access to Peer Management Center.

A custom role is based on one of the three standard web roles (Administrator, Power User, and Help Desk); the custom role starts with the same set of permissions as the role it is based on. However, during the process of creating the custom role, you modify the permissions associated with the new role.

For more information, see Creating a Custom Web Role.

Custom Web Role Permissions

You can create custom web roles in User Management and specify the permissions you want associated with the role.

When creating a custom web role, you select the permissions for the web role in the **Permissions** table, which has three columns:

- **Category** - Identifies the general area of the user interface that the permission applies:
  - **Cloud Backup and Replication UI** - Applies to Cloud Backup and Replication jobs.
  - **Collab/Sync/Repl UI** - File Collaboration, File Synchronization, and File Replication jobs.
  - **DFS-N UI** - Applies to DFS-N Management jobs.
  - **PMC UI** - Applies to Agent Summary view, statistics, task scheduling and task history, logs, memory dumps, and thread dumps.
- **Name** - Identifies the specific area of the user interface.
- **Access** - Identifies the level of access:
  - **Full access** - Has complete access.
  - **View-only access** - Can view but not create, edit, or delete.
  - **No access** - No access.

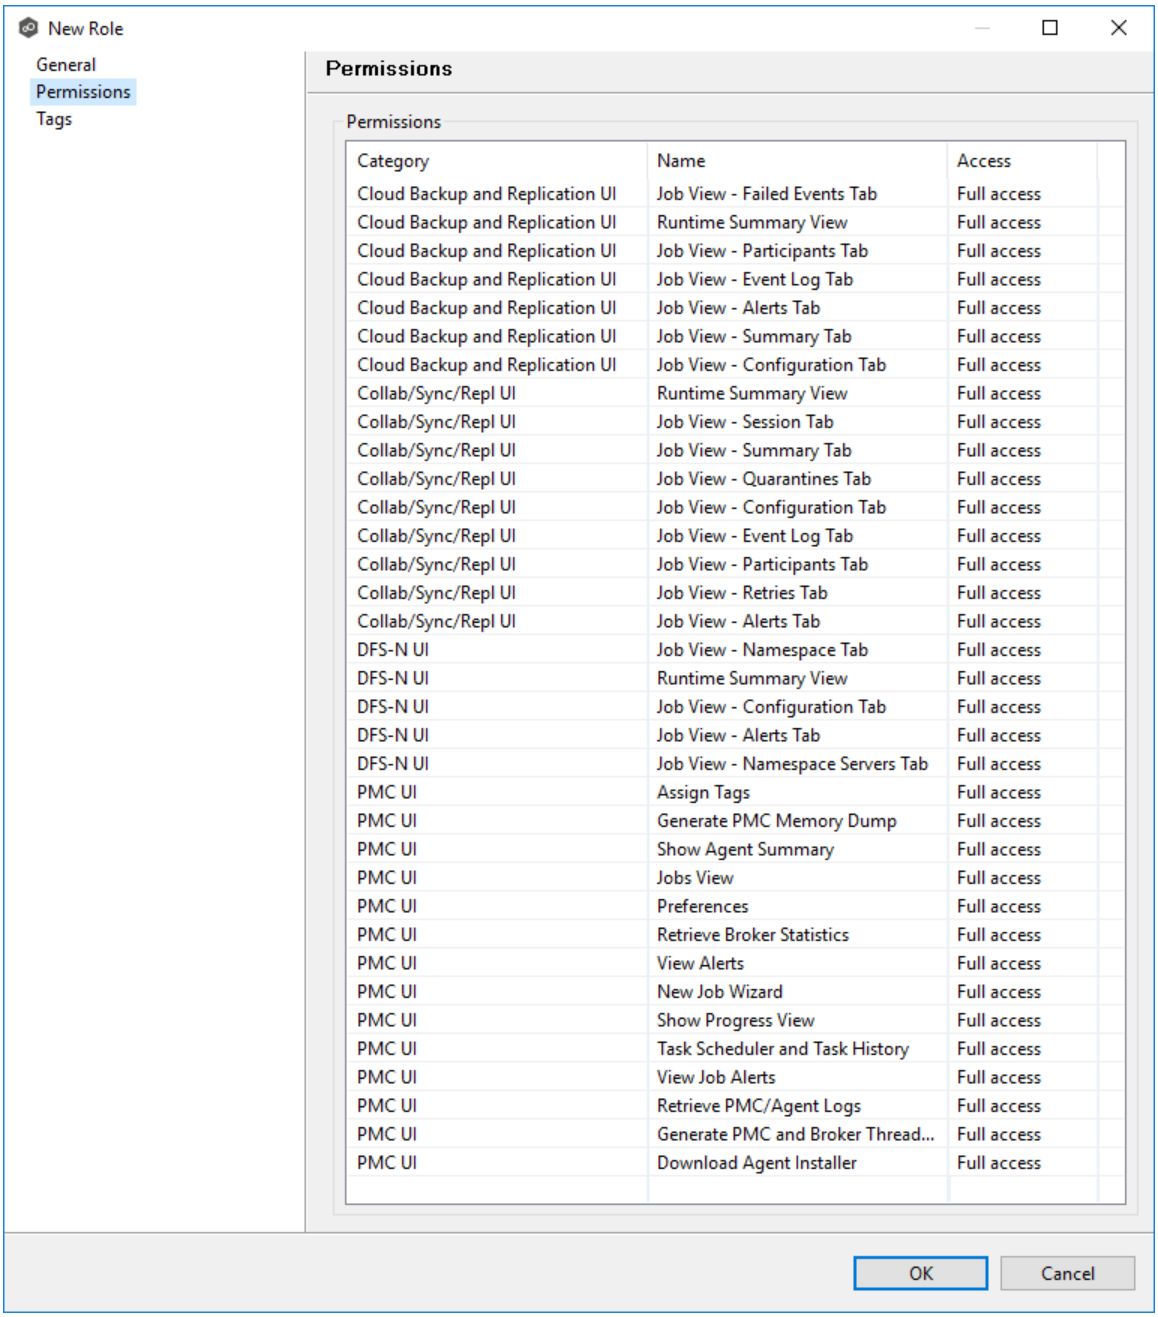

Advanced Topics

This section discusses the following topics:

- Analytics

- Conflicts, Retries, and Quarantines
- DFS Namespaces
- Dynamic Storage Utilization
- File Metadata Synchronization
- Managing Peer Agents
- PeerGFS API
- Scheduled Replication
- Smart Data Seeding
- Storage Capacity
- TLS Certificates

Analytics

Overview

Peer Software provides a Microsoft Power BI-based dashboard to our subscription customers that system administrators can use to monitor the health and performance of PeerGFS and the replication environment. The Analytics Dashboard provides intelligent insights into how PeerGFS and your storage environment are performing.

The Analytics Dashboard is viewable online in a web browser or on a mobile device. The Analytics Dashboard is easy to access--Peer Software sends you a link to view it in a web browser. If you want to view on a mobile device, you will need the Power BI app. It is downloaded automatically when you click the link to access the dashboard.

The Analytics Dashboard combines information from multiple sources, making it easy for you to view the status of PeerGFS and your environment in one place. The dashboard provides a visual and interactive interface that displays telemetry that is updated automatically every 30 minutes by default. You can display details on a chart by hovering over a data point on the chart. You can also subscribe to daily reports that are sent to you via email.

The Analytics Dashboard is currently geared to provide information such as:

- Agent information, including disk space and memory utilization of all Agents.

- Job information, including watch set growth, overall performance, and how long it takes to replicate files.
- Overall PMC information, including disk space and memory utilization, queue backlogs, and quarantine counts.

Some of the information displayed in the Analytics Dashboard is available in the Peer Management Center client interface; however, the dashboard brings the information together for easy access and provides historical data not found in the PMC client.

Health Checker

The Analytics Dashboard integrates information from Peer Software's Health Checker tool. The Health Checker is a standalone service designed to alert on outages and backlog spikes, as well as track overall performance of the replication environment. To view replication performance details in the Analytics Dashboard, the Health Checker must be installed on either a standalone server in the environment (if both alerting and performance monitoring are desired) or on the PMC server itself (if only performance monitoring is desired).

Analytics in the Future

In the future, the Analytics Dashboard will provide even more information and in greater detail. For example:

- Proactive information about your PeerGFS environment, how well it is performing, and details about what it is replicating.
- Trends and insights around what is stored on the file servers across your storage infrastructure.
- Trends and insights around how your users and applications are using the data on the file servers across your storage infrastructure.

For more information about Analytics:

- Prerequisites
- Setting Up Analytics
- Enabling and Disabling Analytics

Prerequisites

Analytics Dashboard Prerequisites

- Must be a subscription-based Peer Global File Service customer.
- Because the Analytics Dashboard relies on information in Microsoft Azure Blob Storage, the PMC and Health Checker servers must be able to get beyond any firewalls that would prevent access to Microsoft Azure.
- For mobile access, the Power BI app must be installed on your mobile device.

Health Checker Prerequisites

While Peer Software's Health Checker is a lightweight application, there are some minimum requirements that must be met to use the tool:

- Windows Server 2012 operating system or later is required.
- Virtual servers running on enterprise-class hypervisors are sufficient if they have a minimum of 2 processor cores, 2 GB of RAM, and 20 GB of free disk space.
- SMB network access is required from the Health Checker server to all participating file servers and shares.
- The Health Checker service must be run under a domain user account with the ability to create folders and files on all participating file servers and shares.
- To provide failure and backlog alerting, the Health Checker should not be installed on any server that is already running the Peer Management Center or the Peer Agent. The Health Checker may be run on servers hosting other infrastructure (such as Active Directory Domain Controllers), provided that the requirements above are met.

Setting Up Analytics

This topic describes the process of setting up Analytics.

Overview

The setup process consists of the following steps:

1. Set up the Analytics engine to upload data to Microsoft Azure by running the **Set Up Analytics** wizard.

2. Request access to the Analytics Dashboard by completing the onboarding form. Peer Software will email you a link to the Analytics Dashboard.

3. Click the link to log in to the Analytics Dashboard:

   - If you are logging in through a web browser, the Analytics Dashboard will appear immediately after entering your credentials.

   - If you are logging in through a mobile device, you will be asked to install the Power BI app on your device. After the installation is complete, you may be asked to log in again to the Power BI app. Once logged in, you will be able to see the mobile version of the dashboard.

Starting the Analytics Setup Process

To start the Analytics set-up process:

1. From the **Help** menu, select **Analytics**.

   You can also initiate the set up process by opening the Analytics page in Preferences and selecting **Enabling Analytics**.

2. If you have not previously set up Analytics or if you have disabled Analytics, the **Set Up Analytics Wizard** will open.

   Follow the prompts in the wizard. Once you finish the wizard, Analytics will run until you disable it.

   If you have previously set up Analytics and it is enabled, the Analytics Preferences page will appear. Modify the settings as desired.

The **Basic Configuration** page is where you enter your Subscription ID and other basic information.

1. Enter your Subscription ID if the field is not auto-filled.

   Contact Peer Software if you do not know your Subscription ID.

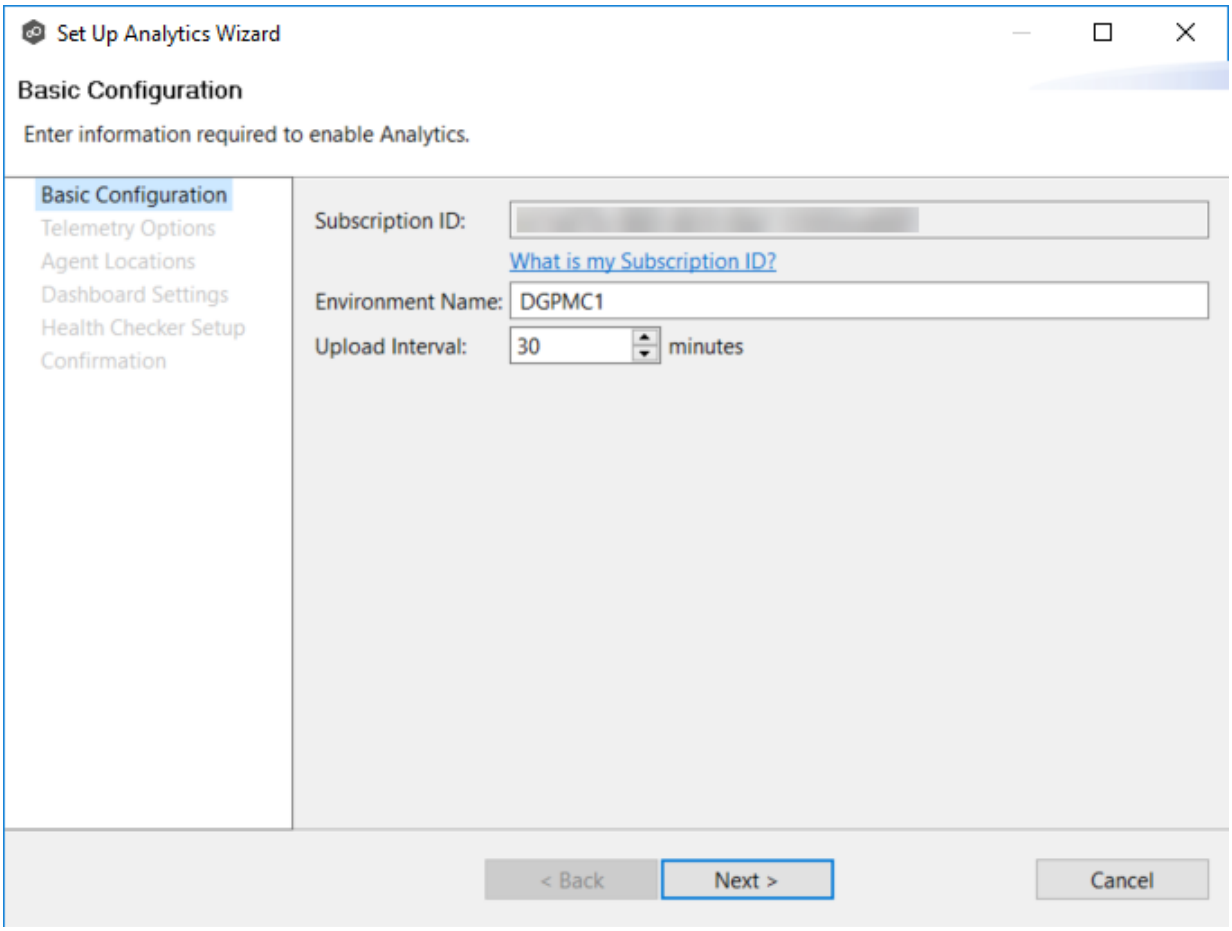

2. In the **Environment Name** field, change the name if the auto-filled value doesn't match the name of the server or environment where the PMC server is installed.

   Note: Changing the name here will change the name in the General Configuration page in Preferences.

3. In the **Upload Interval** field, enter the number of minutes to wait between uploads of data to the Analytics Dashboard.

   The default upload interval is 30 minutes. The minimum interval is 15 minutes; the maximum interval is 180 minutes.

4. Click **Next**.

   The Telemetry Options page appears.

The **Telemetry Options** page allows you to decide whether to upload detailed data, in addition to the basic telemetry data, to the Analytics Dashboard.

1. Select the data to be uploaded:

   The telemetry data is divided into three categories:

   - PMC Details
   - Agent Details
   - Job Details

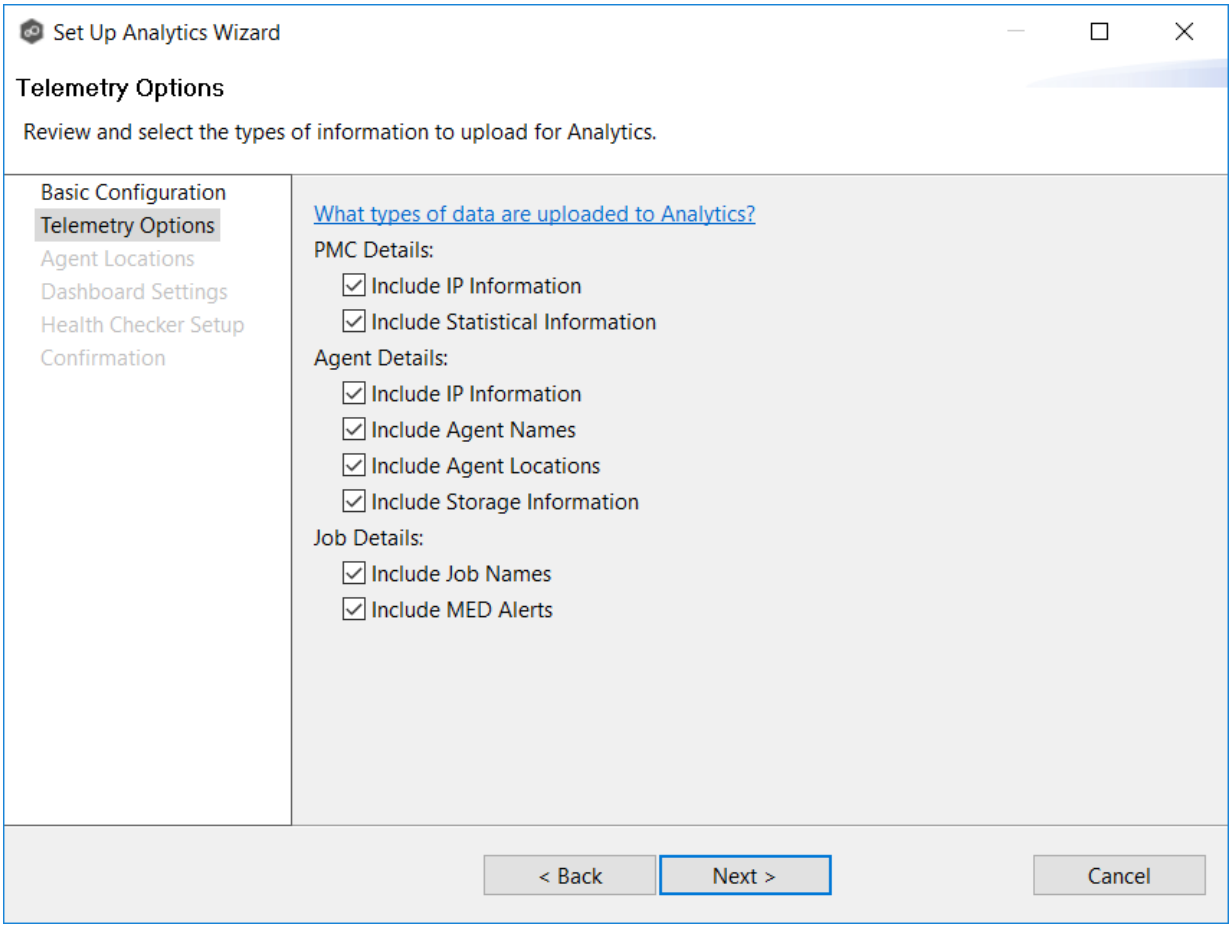

   The table below describes what data is included in each category. Data in the **Standard Data Upload** column is uploaded when Analytics is enabled. Data in the **Include Optional Data** column is uploaded only if you select that option.

Category	Standard Data Upload	Include Optional Data
PMC Details	Includes details about this PMC deployment. Details include service memory consumption, replication backlog, quarantines, license consumption, and watch set size.	• **IP Information** - The IP address of the server that this PMC is installed on. • **Statistical information** - In-depth statistics about the queues and performance of PeerGFS's replication engine.
Agent Details	Includes the details about the Agents that are connected to this PMC. Details include service and server memory consumption, replication throughput, uptime, operating system, and disconnect counts.	• **IP information** - The IP addresses of the servers that the Agents are installed on. • **Agent Names** - The names assigned to the Agents (typically the name of each Agent's Windows Server). If this option is not checked, random strings will be used in the Analytics Dashboard to represent each Agent. • **Agent Locations** - The locations (the latitude, longitude, city, state, and country) of the Agents. Agent locations can be displayed in a map showing replication traffic by geographic location. If you choose to include Agent location data, you will be prompted to enter Agent location information on the next page of this wizard (the **Agent Locations** page of this wizard). No location information is automatically determined - it must be manually entered. If you don't select this option, the **Agent Locations** wizard page will not be displayed. • **Storage Information** - Information specific to the storage platforms that each Agent is managing, including available and used disk space.
Job Details	Includes the details about the file collaboration, synchronization and/or replication jobs configured within this PMC. Per-job details include replication	• **Job Names** - The names of the file collaboration, synchronization, and replication jobs configured in this PMC. If this option is not checked, random strings will be used in the Analytics Dashboard to represent each job.

2. Click **Next**.

   If you selected **Include Agent Locations**, the Agent Locations page appears; otherwise the Dashboard Settings page appears.

The **Agent Locations** page appears only if you selected the **Include Agent Locations** option on the previous wizard page.

The **Agent Locations** page allows you to set location details for each Agent for use with dashboard maps. If you do not want to display Agent locations on dashboard maps, you do not need enter data in this page.

If an Agent's location has already been set through its Agent Configuration, those values will automatically appear on this page.

1. For each Agent that you want to appear on a map, enter the following information:

   - **Location** - Enter the city, state, and country of where the Agent server is installed.
   - **Longitude** - Enter the longitude of where the Agent server is installed.
   - **Latitude** - Enter the longitude of where the Agent server is installed.

   Agents do not automatically self-detect their locations. You can look up location coordinates using free online geographic tools such as https://www.latlong.net/ or by using Google Maps.

   To use Google Maps to look up coordinates

   (a) Open Google Maps, navigate to the Agent server's location and right-click this location.

   (b) On the menu that appears, click the coordinates (the first item in the menu).

   The coordinates are automatically copied to the Clipboard.

   (c) Paste the coordinates into a text editor window (such as Notepad).

   (d) Paste each coordinate into the proper coordinate field.

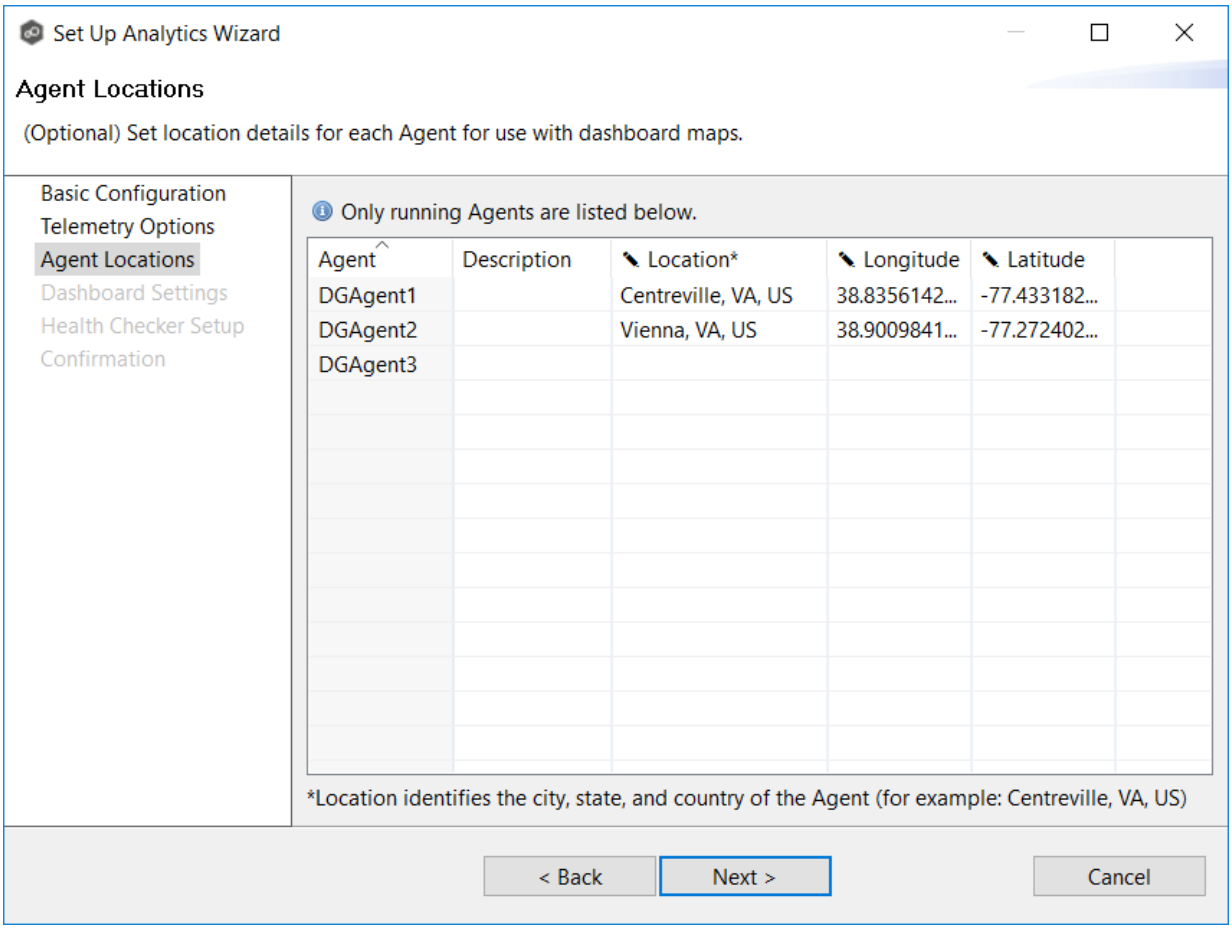

Note: The **Description** column is read-only. Its value is set through Microsoft Windows.

2. Click **Next**.

   The Dashboard Settings page appears.

This step is optional.

The **Dashboard Settings** page allows you to customize the desired real-time delta value and other dashboard settings.

1. Customize the Analytics Dashboard settings:

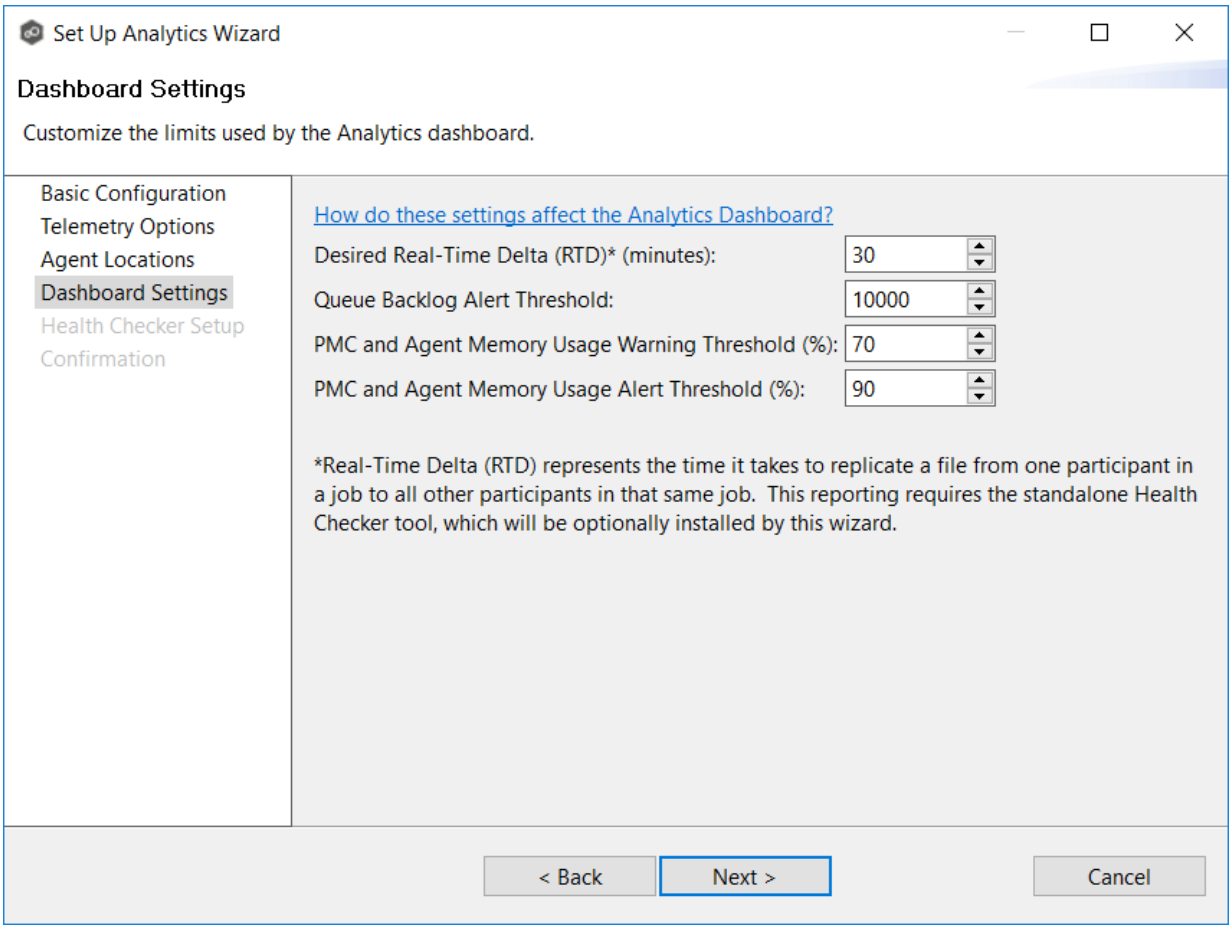

Setting	Description
Desired Real-Time Delta (RTD)* (minutes)	Enter the maximum amount of time that is acceptable for a file change to replicate from one participant in a job to all other participants in that same job. Anytime a job is found to be taking more than the desired RTD, the Analytics Dashboard will display warnings. The default is 30 minutes. The minimum is 5; the maximum is 1440. Note: This only applies to real-time activity. It does not apply to scans or bulk activity.
Queue Backlog Alert Threshold	Enter the desired threshold for alerts about queue backlog. Once the queues inside PeerGFS hit this threshold, the Analytics Dashboard will display warnings. The default is 10000 items in the queue before an alert is sent. The minimum is 1000; the maximum is 1000000.
PMC and Agent Memory Usage Warning Threshold (%)	Enter the desired threshold for warnings about both Agent and PMC service memory usage. If either the PMC or any Agent hits this threshold, the Analytics Dashboard will display warnings. The default is 70 percent of memory usage. The minimum is 50 percent; the maximum is 95 percent.
PMC and Agent Memory Usage Alert Threshold (%)	Enter the desired threshold for alerts about both Agent and PMC service memory usage. If either the PMC or any Agent hits this threshold, the Analytics Dashboard will display errors. The default is 90 percent of memory usage. The minimum is 75 percent; the maximum is 100 percent.

2. Click **Next**.

   The Health Checker Setup page appears.

The **Health Checker Setup** page is where you decide whether you want to set up the Health Checker, which is a standalone monitoring and reporting tool. The Health Checker alerts on replication issues and tracks the **Real-Time Delta (RTD)**. The RTD represents how many minutes out of sync the replication is between participants. It is similar to a **Recovery Point Objective (RPO)** but focuses only on the real-time replication engine and not the scan engine.

You can install the Health Checker **locally** (on the same server that Peer Management Center is installed on) or **externally** (on a server that neither Peer Management Center nor Agents are running on). To receive the full benefit of the Health Checker, we recommend that you install it externally. If you install it locally, you will not get failure alerts if the PMC server goes down.

To set up Health Checker:

1. Select a setup option.

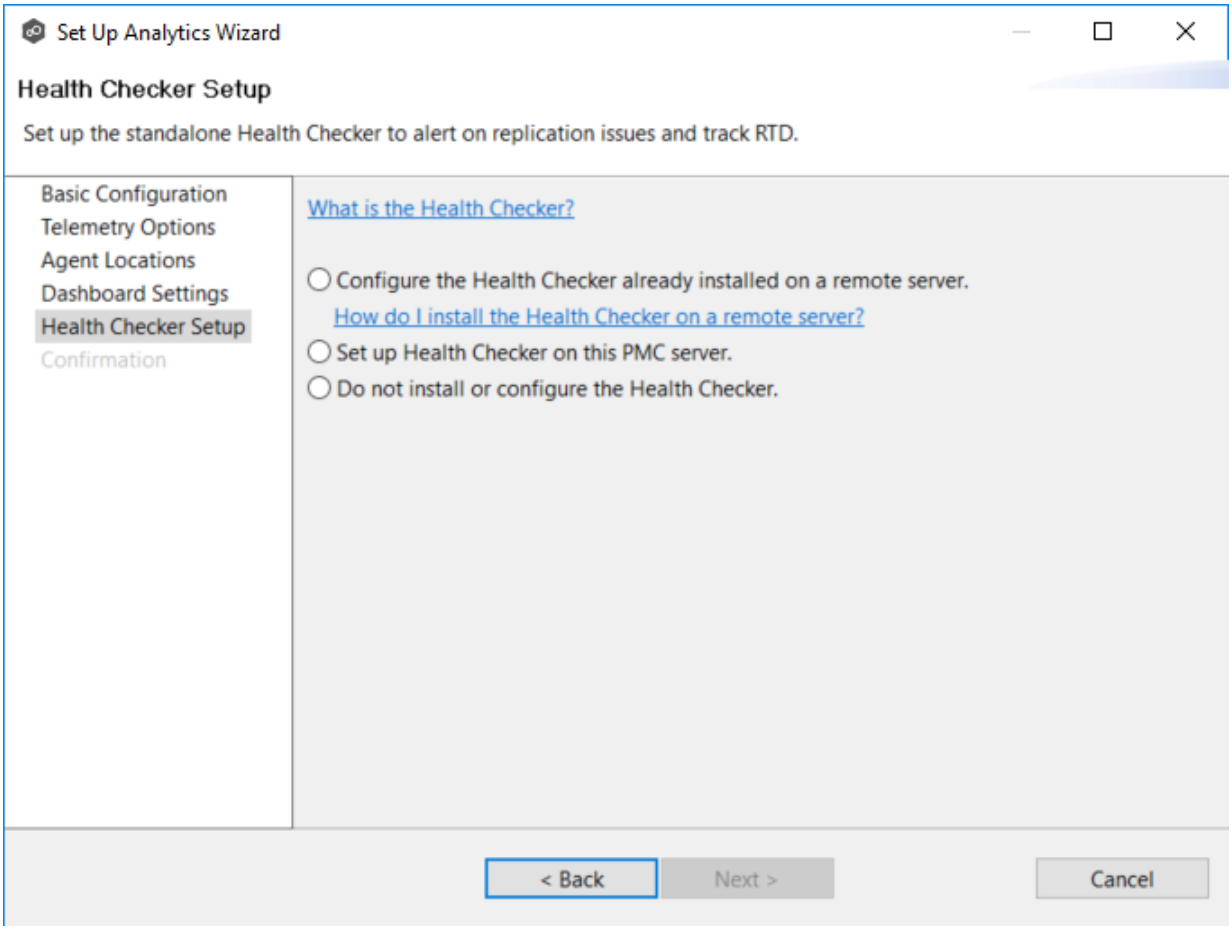

Option	Description
Configure the Health Checker already installed on a remote server.	Recommended option for receiving the full benefit of Health Checker. If installed on remote server, Health Checker does both failure as well as performance monitoring. The remote server can be a domain controller or some other infrastructure server. It should not be running PMC or Agent software.
Set up Health Checker on this PMC server.	Select this option if you want performance statistics but not full failure monitoring. If the PMC server fails or is shut down, replication will stop but no failure alerts will be sent.
Do not install or configure the Health Checker.	Select this option if you do not want performance statistics or failure monitoring.

2. Click **Next**.

   Your next step depends on the option you selected in Step 1:

Selected Option	Continue on this wizard page
Configure the Health Checker already installed on a remote server.	Select Jobs
Set up Health Checker on this PMC server.	Set up the Health Checker
Do not install or configure the Health Checker.	Confirmation

Set Up the Health Checker

The **Set up the Health Checker** page appears only if you selected option 2 on the previous wizard page.

From this page, you install Health Checker on the PMC server, initialize the Health Checker components, and run the Health Checker service. If you have previously installed Health Checker on the PMC server, it will verify that Health Checker was successfully installed and initialized and is currently running.

1. Click the **Start** button.

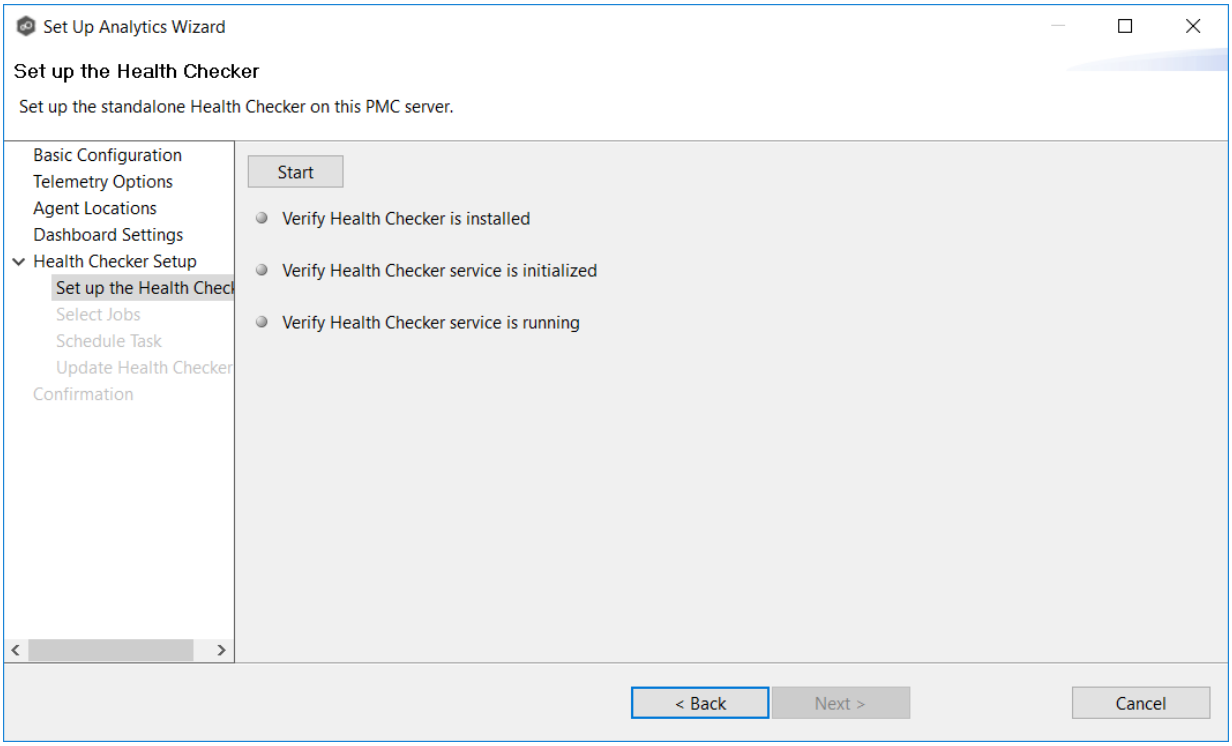

   As the wizard performs the setup process, it checks to see whether Health Checker was previously installed on this server. It communicates the results via colored dots. A green dot indicates that successful completion of that setup stage. A red dot indicates that action is needed.

2. If the dot next to **Verify Health Checker is installed** turns red, Health Checker has not yet been installed. Click the **Install** button that appears. It runs a silent packaged installer for the Health Checker. Click **OK** in the Success message that appears.

3. If the dot next to **Verify Health Checker service is initialized** turns red, the Health Checker service has not yet been initialized. Click the **Initialize** button that appears. In the **Service Account Info** dialog, enter the user name and password for the service, and then click **OK**. Click **OK** in the Success message that appears.

4. If the dot next to **Verify Health Checker service is running** turns red, click the **Start** button that appears.

   The Health Checker service will start. If this is the first time the service has run, the following message appears; click **OK**.

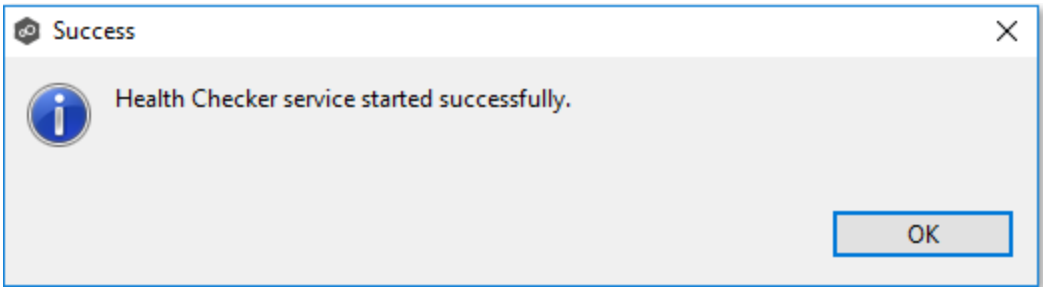

5. Once all dots are green, click **Next**.

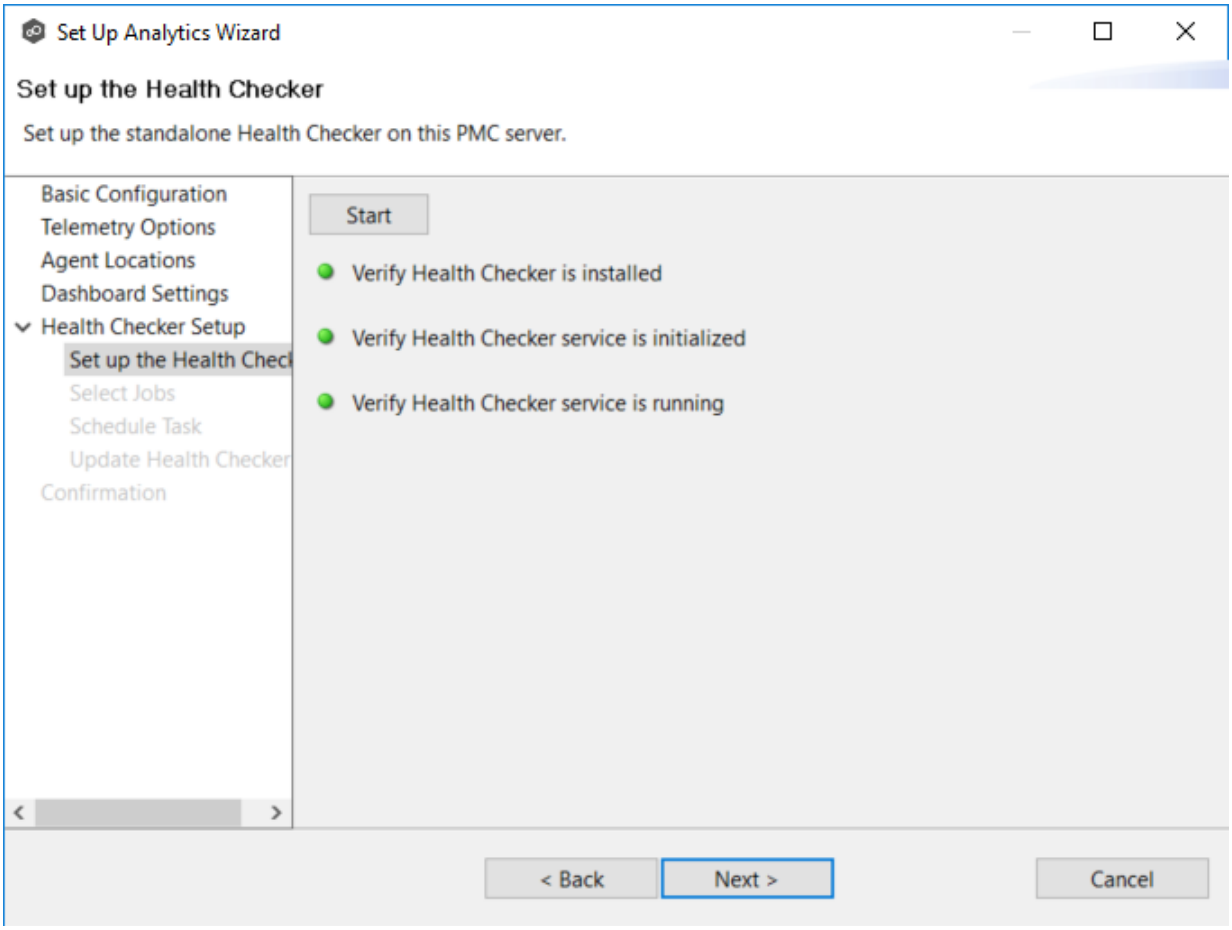

The Select Jobs page appears.

Select Jobs

The **Select Jobs** page allows you to specify which jobs are monitored by Health Checker. By default, all jobs are selected.

To select jobs:

1. Click the **Apply for all jobs** to monitor all current and future jobs. Otherwise, use the **Add** and **Remove** buttons to move jobs from between the **Selected** and **Available** lists.

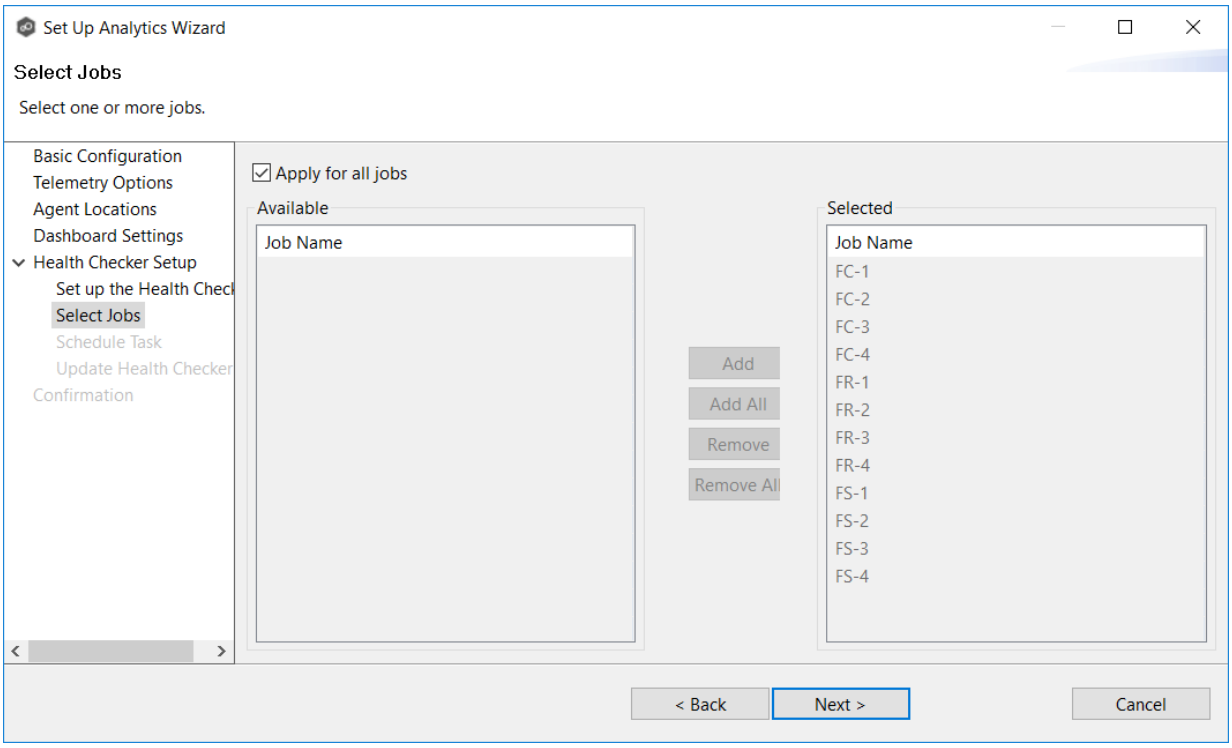

2. Click **Next**.

   The Schedule Task page appears.

Schedule Task

The **Schedule Task** page allows you to specify when and how often Health Checker configuration information is updated.

By default, the task is set to run every day at 4-hour intervals.

1. In **Settings**, select a frequency as well as the start date and time.

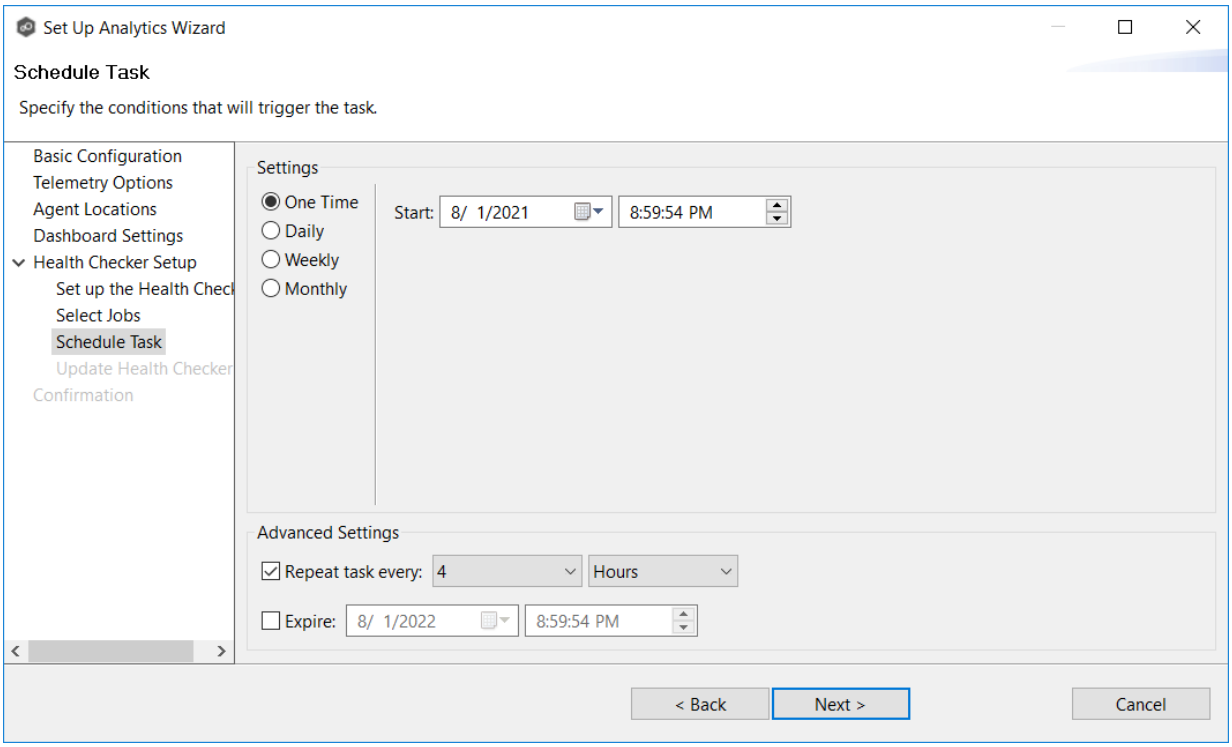

2. In **Advanced Settings**, select whether you want the task repeated and the frequency of the repetition. We recommend repeating this task every 1 to 4 hours.

3. In **Advanced Settings**, select when you want the task to expire.

   If you don't select an expiration date, the task will run indefinitely.

4. Click **Next**.

   The Configure External Health Checker page appears.

Update Health Checker

The **Update Health Checker** page is where you specify where the Health Checker was installed and specify the types of files that Health Checker uses as test files. In addition, you specify which jobs Health Checker should monitor, based on job state.

To configure the Health Checker:

1. In the **Path to Health Checker Configuration** field, enter the path to the **workspace** subfolder under the Health Checker installation folder.

Note: If Health Checker is installed on an external server, you must enter the path in UNC format.

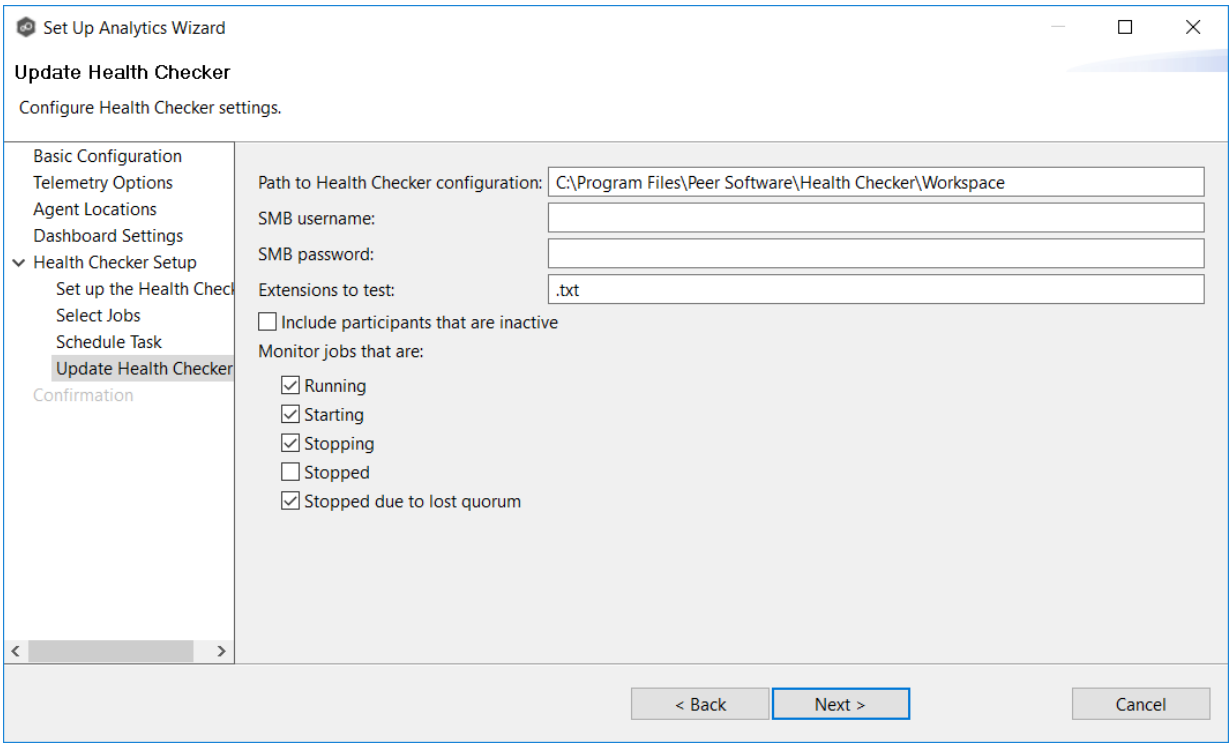

2. In the **SMB Username** field, enter the user name if the server hosting the Health Checker requires account credentials. In most cases, a locally installed Health Checker will not need a user name.

3. in the **SMB Password** field, enter the password if the server hosting the Health Checker requires account credentials. In most cases, a locally installed Health Checker will not need a password.

4. In the **Extensions to Test** field, enter the extensions for the file types that you want Health Checker to monitor. Separate the extensions with a semicolon.

5. Select the **Include Participants that are inactive** checkbox if you want the Health Checker to monitor inactive participants.

6. In the **Monitor jobs that are** section, select which job states should be included in Health Checker monitoring.

   We recommend checking all but the **Stopped** category. The **Stopped** category includes jobs that are stopped for any reason other than a quorum lost. For example, if you have manually stopped a job, you may not want Health Checker monitoring it.

7. Click **Next**.

The Confirmation page appears.

The **Confirmation** page displays a summary of the settings you have selected.

1. Review the configuration.

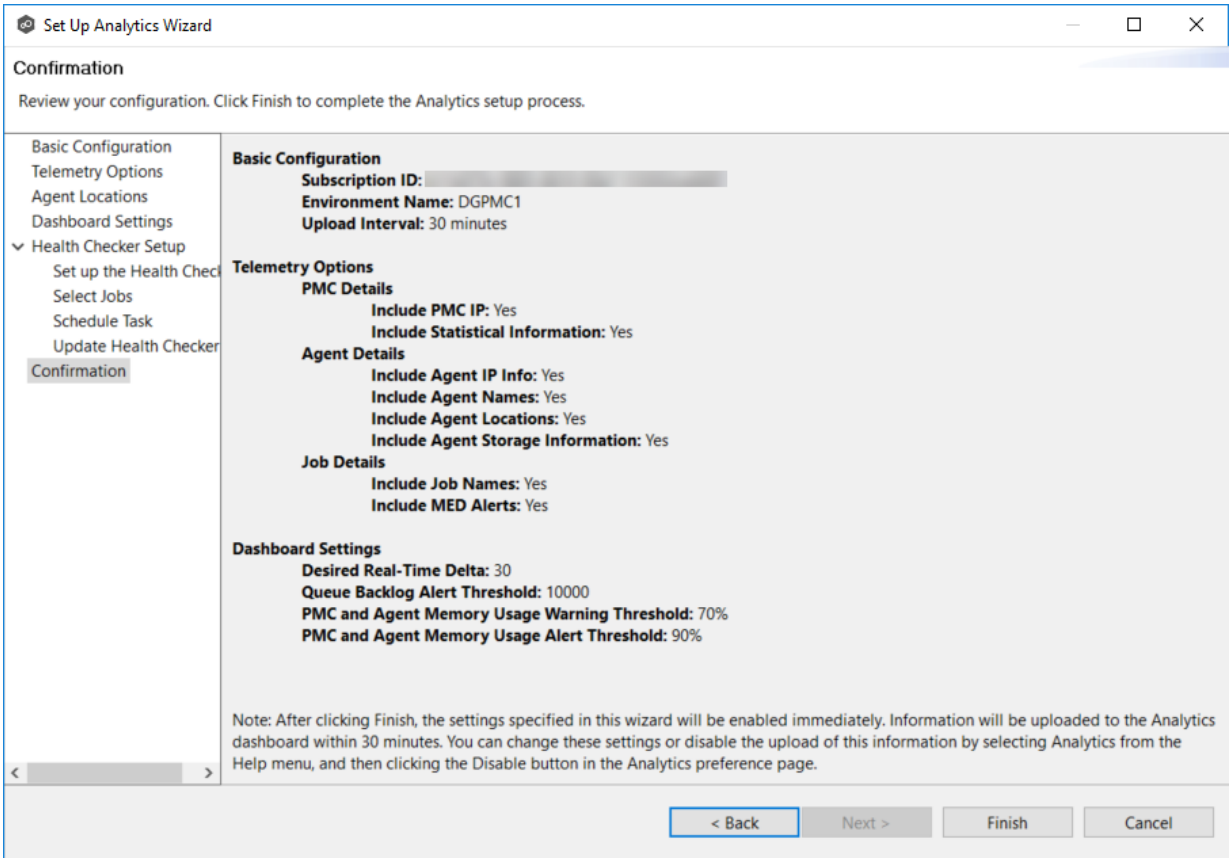

2. Click **Finish** to complete the setup.

   If you modified the location of an Agent, the following message appears:

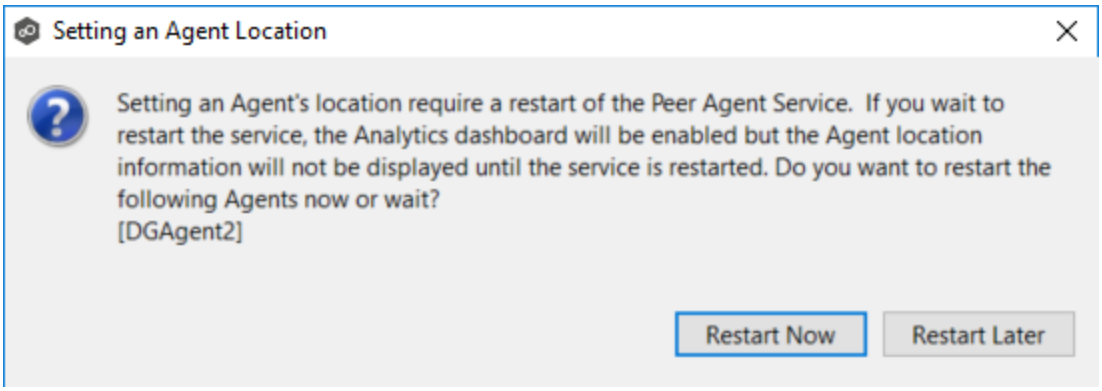

Click **Restart Later** if you do not want to restart the Agent services because you have other jobs running; otherwise, click **Restart Now**.

3. Continue by submitting completing the onboarding form described in Setting Up Analytics.

Disabling and Re-enabling Analytics

Disabling Analytics

The effects of disabling Analytics are:

- Your environment will no longer be monitored for performance or failures.
- Your Analytics Dashboard will not be updated.
- Telemetry data will no longer be uploaded to Peer Software. If you disable Analytics, you have the option of having your telemetry data deleted by Peer Software.

To disable Analytics:

1. Select **Preferences** from the **Window** menu.

   The **Preferences** dialog appears.

2. Select **Analytics** in the navigation tree.

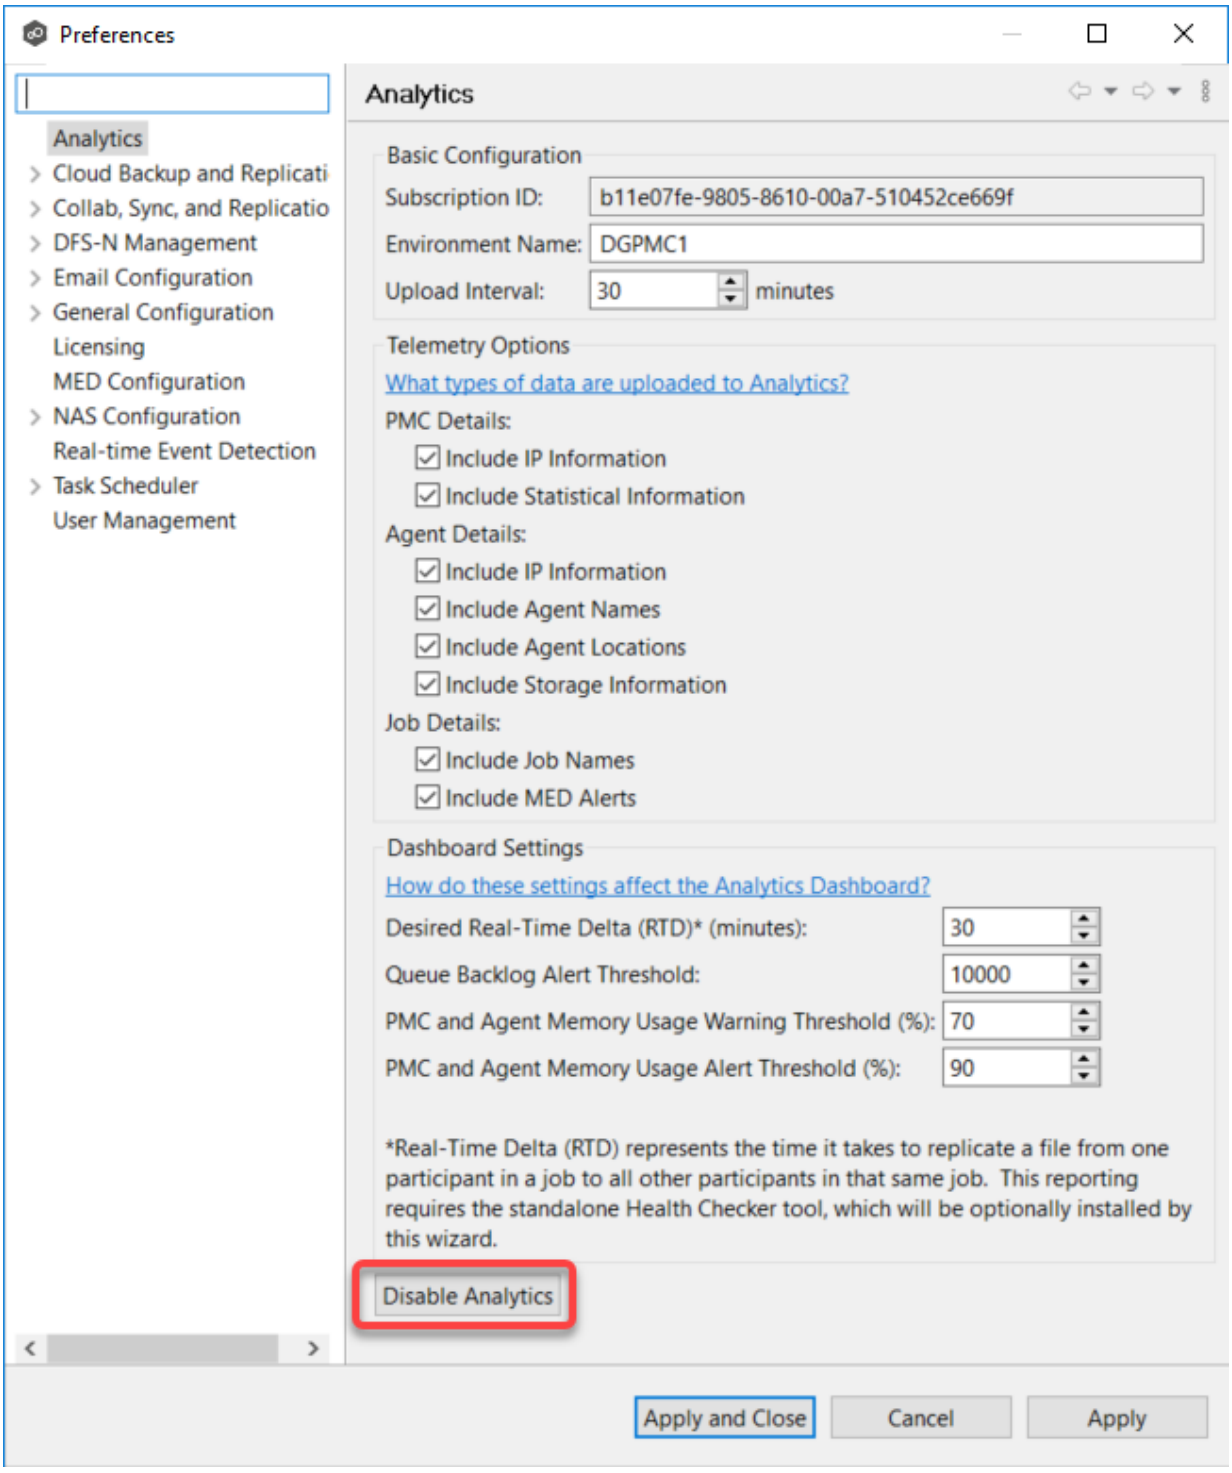

3. Click the **Disable Analytics** button.

4. In the **Disable Analytics** dialog that appears, click **Yes** to confirm that you want to disable Analytics.

Note: If you click **Yes**, Analytics will be disabled immediately. Clicking the **Cancel** button will not re-enable Analytics.

5. In the **Remove Analytics Data** dialog that appears, click **Yes** if you want your telemetry data deleted. Peer Software will be notified to delete your data from Microsoft Azure. This process will remove all of your current and historical information. Once deleted, this data will not be recoverable.

6. If Health Checker was installed locally, select an option in **Health Checker** dialog that appears:

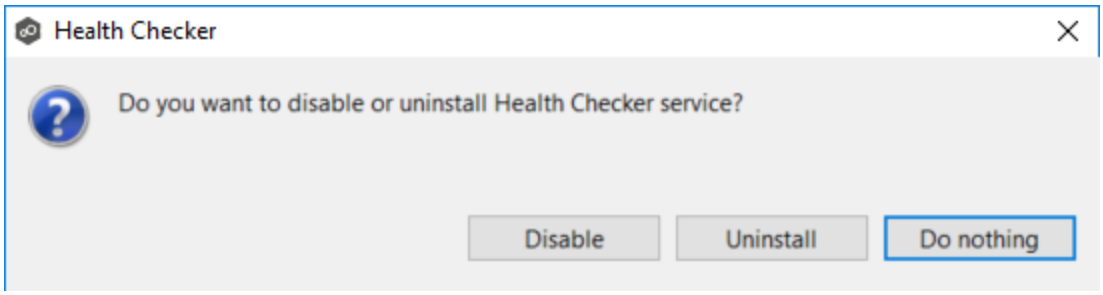

   - Click **Disable** if you want Health Checker disabled but not uninstalled. The Health Checker Service will be stopped and prevented from starting automatically.
   - Click **Uninstall** if you want Health Checker uninstalled. An uninstaller will start automatically to remove the Health Checker service and all installed files.
   - Click **Do Nothing** if you want the Health Checker to remain installed and running.

7. Click **Apply and Close** or **Apply**.

Re-enabling Analytics

To re-enable Analytics, you must rerun the Set up Analytics wizard.

You can access the wizard by selecting **Analytics** from the **Help** menu or by clicking the **Enable Analytics** button in the Analytics preferences page.

Conflicts, Retries, and Quarantines

Making unstructured data active at multiple locations introduces the chance of users making conflicting changes to different copies of the same file. The real-time synchronization and locking engines built into Peer Global File Service are designed to prevent these conflicts by ensuring that only one user can modify a file at a time while also making sure that all locations always have the most up to date version of a file. There are scenarios, however, where the

synchronization and locking engines may not be able to prevent version conflicts. Such scenarios include network outages and file system issues.

The conflict resolution engine in Peer Global File Service is designed to handle these circumstances with a three-tiered approach backed by a combination of scans and real-time activity:

- **File Conflicts** – The initial state of detection of a potential version conflict. Depending on user activity, these can often be resolved automatically.
- **File Retries** – If certain errors are thrown when trying to synchronize a file between locations, this file will be automatically put into a retry list. Synchronization of this file will be retried every minute for a maximum of 60 attempts. The frequency of attempts and the maximum number of attempts are configurable.
- **File Quarantines** – These are file conflicts that could not be automatically resolved, as well as file retries that have failed after the maximum number of attempts. Files in the quarantine list will no longer be synchronized or protected with file locks until a winning file is picked through the PeerGFS user interface.

File conflicts (and potentially quarantines) can occur for any of the following reasons:

- Two users open a file at the same time or in-and-around the same time.
- A file is open at the start of a job and has been modified on a host where the configured conflict resolution strategy selects a different host as the winner.
- Two or more users have the same file open on different hosts when a collaboration job is started.
- A file was modified on two or more hosts between job restarts or network outages.
- Peer Management Center is unable to obtain a lock on a target host file for various reasons.
- Peer Management Center may conflict a file when an unexpected error occurs, or a file is in an unexpected state.

File retries can occur for any of the following reasons:

- The transfer of a file between locations is interrupted for any reason.
- The renaming of a temp file after a successful file transfer is blocked for any reason.

An example of a file conflict versus a file quarantine is as follows:

Two users have the same file open at two different locations prior to a Peer Global File Service job being enabled. When starting the job, PeerGFS will track this file as a potential conflict. If

only one or no users make a change to the file, this conflict will automatically be resolved. If both users make a change, the conflict will become a quarantine.

DFS Namespaces

Overview

A DFS namespace enables you to group shared folders that are located on different servers into one or more logically structured namespaces. Each namespace appears to users as a single shared folder with a series of subfolders. However, the underlying structure of the namespace can consist of numerous file shares that are located on different servers and in multiple sites.

The elements that make up a DFS namespace are:

- **Namespace server** - A namespace server hosts a namespace. The namespace server can be a member server or a domain controller.
- **Namespace root** - The namespace root is the starting point of the namespace. For example, if you have a namespace path of \\Domain.local\MyNamespace, the root is MyNamespace. This is a domain-integrated namespace, meaning that its metadata is stored in Active Directory Domain Services.
- **Folders** (also referred to as **namespace folders**)- Namespace folders without folder targets add structure and hierarchy to the namespace, while folders with folder targets provide users with actual content. When users browse a folder that has folder targets, the client computer receives a referral that transparently redirects the client computer to one of the folder targets.
- **Folder targets** - A folder target is the UNC path of a shared folder or another namespace that is associated with a folder in a namespace. The folder target is where data and content are stored. For example, if a user navigates to **\\Domain.local\MyNamespace\MyFolder**, the user is transparently redirected to **\\NYC-FS.Domain.local\MyFolder** or **\\LA-FS.Domain.local\MyFolder**, depending on which site the user is currently accessing. Adding multiple folder targets increases the availability of the folder in the namespace.

For more information about DFS namespaces, see DFS Namespaces overview on Microsoft's website.

Managing DFS Namespaces through PeerGFS

PeerGFS enables you to create a namespace and manage various activities related to it, such as creating namespace folders, adding folder targets, and linking the namespace to a File Collaboration or File Synchronization job. You could manage DFS namespace using Microsoft

tools; however, you can manage DFS namespaces through a dedicated job type in Peer Management Center, the DFS-N Management job.

The benefits of creating and managing a DFS namespace within Peer Management Center are:

- **Ease of managing a namespace** - You can create and manage a namespace within the same interface that manages PeerGFS synchronization and replication technologies. This removes the need to use two different tools to manage the key elements of multi-site and multi-vendor file services.
- **Integration with PeerGFS collaboration and synchronization** - When linked to file collaboration and synchronization jobs, DFS namespaces can provide redundancy to file shares across file servers and locations.
- **Automating failover and failback** - If a file server goes offline, Peer Management Center can disable the associated folder target in the DFS namespace. This automatically redirects users to another available file server. When the original file server comes back, Peer Management Center will automatically make sure it is brought back in sync, and then enable the associated folder target so users can once again connect to it. See DFS Namespace Failover and Failback for more information.

Note: Although Microsoft provides two types of namespaces, a stand-alone namespace or a domain-based namespace, you can manage only a domain-based namespace in PeerGFS.

For more information about using DFS namespaces in PeerGFS, see:

- Using DFS Namespaces with Jobs
- DFS Namespace Failover and Failback
- DFS-N Management Jobs
  - o Creating a DFS-N Management Job
  - o Managing DFS Namespaces
  - o Linking a DFS Namespace to File Collaboration or File Synchronization Job

Using DFS Namespaces with Jobs

If you want to use a DFS namespace with a File Collaboration or File Synchronization job, you can create a DFS-N Management job to manage the namespace from within PeerGFS.

PeerGFS is very flexible and lets you proceed in various ways. For example:

- You can create a new namespace or import an existing one by first creating a DFS-N Management job and then later linking the namespace to a File Collaboration or File Synchronization job.

- You can create a new namespace or import an existing one when creating a File Collaboration or File Synchronization job, thus automatically linking namespace folder targets to the watchsets of the collaboration or synchronization participants. A DFS-N Management job is automatically created during this process.

- You can also import an existing namespace by right-clicking on the Namespace Summary view which will guide you through the import of a namespace into PeerGFS. Importing an existing namespace will automatically create a DFS-N Management job which can then be linked to a File Collaboration or File Synchronization job.

Before creating jobs that use namespaces, you may want to configure DFS preferences.

See Managing DFS Namespaces for information about adding namespace servers, namespace folders, or folder targets to a DFS namespace.

DFS Namespace Failover and Failback

One of the primary benefits of using DFS Namespaces with PeerGFS is that Peer Management Center can control failover and failback by automatically disabling and enabling DFS namespace folder targets.

Failover

Peer Management Center and Agents are constantly looking for connectivity issues and other failures across linked file servers, the Peer Agents themselves, and entire sites. If Peer Management Center detects a failure, Peer Management Center can be set to automatically disable a linked DFS namespace folder target from a namespace folder. This will prevent end users from accessing the associated folder target. For details about enabling and disabling automatic failover to another folder target, see DFS-N Management in Collaboration, Replication, and Synchronization Job Preferences.

Failback

When Peer Management Center determines that a file server, Peer Agent, or entire site is back online, it automatically runs the following process to re-integrate that file server:

1. Kicks off a rescan to ensure the disconnected site or file server is brought back in sync with the others.

2. Re-enables the associated folder target once the re-scan is complete. Once this is done, DFS Namespaces begins to direct end users back to this file server.

For details about enabling and disabling automatic failback to another folder target, see DFS-N Management in Collaboration, Replication, and Synchronization Job Preferences. Automatic failback is enabled by default.

Dynamic Storage Utilization

Overview

Dynamic Storage Utilization (DSU) allows you to save storage space on edge storage devices (for example, storage devices used in branch offices) where only a small subset of files are used on a frequent basis. Files that are used less frequently are replaced with stub files on the edge storage device so that it appears to have a complete set of files. When a user accesses a stub file, DSU retrieves the full version of the file from a master storage device. The benefit of using DSU is that it allows you to efficiently utilize storage capacity on edge devices while preserving fast access performance on files that are used most often.

DSU offers flexible edge storage management with:

- The ability to assign an an amount or percentage of available storage to be used on the edge storage device.
- Dynamic adjustments of the time periods used to determine whether to stub or rehydrate a file, allowing DSU to keep the assigned storage space as full as possible (best experience for the end user).
- Direct integration with our file collaboration and file synchronization job types.
- Point-to-point data transfer capability between one edge and one or more masters.
- The flexibility to mix and match master and edge roles across different jobs.
- The ability to pin files or folders to always be local or always be stubbed on the edge storage device.
- Alerting to ensure you stay ahead of potential storage capacity limits.

Fundamental Concepts

A master participant has a complete set of hydrated files and no stub files. An edge participant contains a subset of the complete, hydrated files on a master participant, while the rest of the files will be stub files that don't take up any space. Users can retrieve stubbed files directly

from a master participant as needed. The goal of dynamic storage utilization (DSU) is to keep as much as possible cached locally on edge participants for rapid access.

Every edge participant must have at least one master participant assigned to it. When a stub file needs to be rehydrated, DSU will retrieve the file from a master participant.

User-defined business rules (volume and utilization policies) manage the storage capacity on edge devices. DSU scans edge participants on a set basis (typically at least once daily) and uses these policies to determine whether adjustments are needed, i.e., whether to stub files to free up space or to rehydrate files. This ensures that the storage capacity is being used at optimum efficiency.

DSU Glossary

This glossary presents some of the most important terms used in conjunction with DSU.

Term	Definition
Stub file	A file that appears to the user to be stored on the local disk and immediately available for use but is actually held either in part or entirely on a different storage medium.
Local file	A file that is fully available without network access to a master participant; all of its bytes are present (stored locally) on the participant.
Rehydrated file	A file that was stubbed but has been fully reconstituted on the edge participant.
Master participant	Always has a complete set of files for the job. None of its files are stubbed; they are always stored physically on that device.
Edge participant	A subset of the files stored on a master participant are physically stored on an edge participant; the rest of the files on an edge participant are stub files that take up minimal space but can be rehydrated as needed.
Master Data Service	A service that handles requests from edge participants for files on a master participant. The Master Data Service is installed on the Peer Agent server as part of the Peer Agent installation process.

Term	Definition
Volume policy	Specifies how much of the available space on the volume monitored by the Agent/edge participant to be assigned for local (hydrated) files.
Temporary storage space	Space that is used to temporarily store the content of stub files that are being rehydrated.
Utilization policy	Defines when a file should be stubbed versus fully hydrated across all volumes of this edge participant. Parameters are based on the size of the files to be potentially stubbed and when they were last accessed and modified. A utilization policy enables you to balance getting the best performance while keeping the cache as full as possible.
Pinning filter	Specifies whether specific files or files in a particular directory are always stubbed or always local on the edge participant. A pinning filter similar is to a utilization policy—it can be applied to multiple jobs. If there is a conflict between a pinning filter and utilization policy (where, for example, you might have something set to be always stubbed), the pinning filter will take precedence.

File Metadata Synchronization

Overview

File metadata is additional information stored as part of a file. The primary component of file metadata is Security Descriptor Information, also known as access control levels (ACLs).

The Security Descriptor Information elements that can be synchronized are:

- **Owner**: NFTS Creator-Owner. By default, the owner is whomever created the object. The owner can modify permissions and give other users the right to take ownership.
- **DACL**: Discretionary Access Control List. It identifies the users and groups that are assigned or denied access permissions to a file or folder.
- **SACL**: System Access Control List. It enables administrators to log attempts to access a secured file or folder and is used for auditing.

File Metadata Conflict Resolution

File metadata conflict resolution occurs only the first time a file is synchronized during the initial scan, and only when one or more security descriptors do not match the designated master host.

If the file does not exist on the designated master host, then no conflict resolution is performed. If a master host is not selected, then no file metadata synchronization is performed during the initial scan.

ACL Requirements

- Enabling ACL synchronization requires that all participants be members of any referenced domains that are configured in the ACL(s) or as the owner of the file. Failure to do so may render the file unreadable on the offending target host.
- All Peer Agents must be run under a domain Administrator account and cannot be run under a local or System account.
- To ensure accurate and consistent ACL propagation, the security settings for the watch set must match EXACTLY across all the participants. The best and easiest way to ensure the security settings match is to compare the permissions in the Microsoft **Advanced Security Settings** dialog for the root folder being watched.

Managing Peer Agents

The ability to remotely manage the configuration for connected Peer Agents is available from within Peer Management Center. Right-clicking one or more agent names in the **Agents** view displays the following context menu:

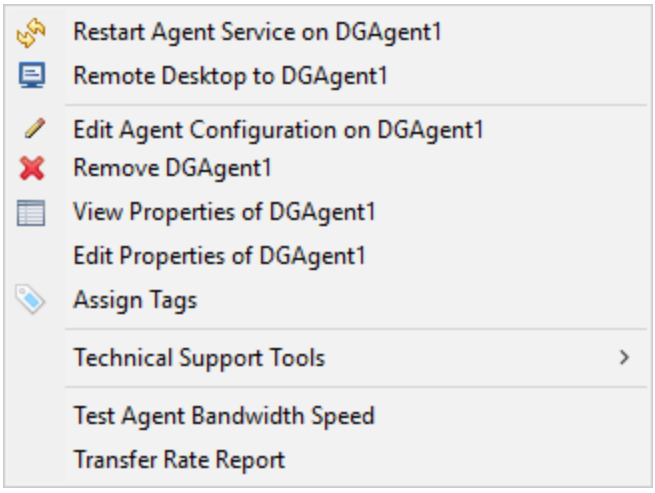

Options

Option	Description
Restart Agent Service	Restarts the Peer Agent Windows service running on the corresponding host if the selected Peer Agent is connected. If the Peer Agent is not connected to the Peer Management Broker, an attempt is made to restart the Peer Agent Windows service using the Windows **sc** command. Note that this works only if the user running the Peer Management Center can access the remote Peer Agent system and has the appropriate domain permissions to start and stop services on the remote Peer Agent system.
Remote Desktop to Agent	Launches a Windows Remote Desktop connection to the selected Peer Agent.
Edit Agent Configuration	Displays a dialog through which the selected Peer Agent can be configured. Configurable options include Peer Management Center connectivity, Peer Agent logging, Peer Agent memory usage, among others. For more information, see Editing an Agent Configuration.
Remove Agent	Remove the selected Peer Agent(s) from the **Agents** view, but if the Peer Agent is still running or connects again, then it will be added back to the list when the next heartbeat is received.
View Properties of Agent	Displays properties for the selected Peer Agent, for example, heartbeat information, host machine configuration, messaging statistics, performance statistics. See Viewing Agent Properties for more details.
Edit Properties of Agent	Allows you to edit the connection type, preferred host, and RDP connection string.
Assign Tags	Displays a dialog where you can view and assign tags to resources.

Option	Description
Technical Support Tools	Displays a list of tools that can be used to assist Peer Software Technical Support.
Test Agent Bandwidth Speed	Runs a bandwidth speed test to be performed in the background if the selected Peer Agent is connected. You are notified at completion with the results of the test.
Transfer Rate Report (not available on Web Client)	(Rich client only) Displays a time series performance chart of average transfer rate for the selected Peer Agent over the last 24 hours.

Technical Support Tools

The options on the **Technical Support Tools** submenu are:

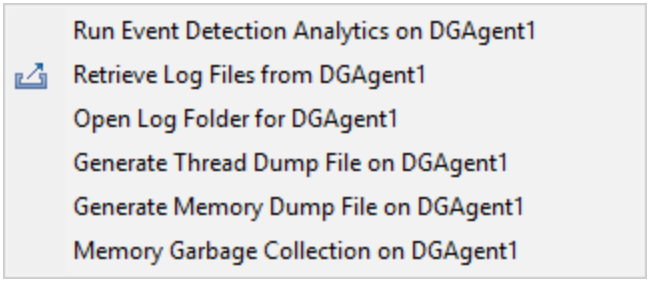

Command	Description
Run Event Detection Analytics	Runs the Event Detection Analytics tool for the selected job, which looks at real-time activity that has been occurring on that specific agent.
Retrieve Log Files	Retrieves log files for the selected Peer Agent. The log files contain information that the Peer Software Technical Support uses in debugging issues. The log files are encrypted and are located in the support folder of the Peer Management Center installation directory. They can optionally be uploaded to the Technical Support team.

Command	Description
Open Log Folder for Agent	Opens the log folder.
Generate Thread Dump File on Agent	Generates a thread dump file for the selected Peer Agent, which can be used by Peer Software technical support to debug certain issues. The debug file is located in the Peer Agent installation directory.
Generate Memory Dump File on Agent	Generates a memory dump file for the selected Peer Agent, which can be used by Peer Software technical support to debug certain issues. The debug file is located in the Peer Agent installation directory.
Memory Garbage Collection on Agent	Forces a garbage collection operation to attempt to reclaim memory that is no longer used within the Peer Agent's JVM.

Peer Agent Connection Statuses

A connection status indicates the state of the Peer Agent's connection to the Peer Management Broker. The Peer Management Broker serves to connect Peer Agents to Peer Management Center.

Peer Agent connection statuses are displayed in the **Agents** view in the Peer Management Center:

- The status of an Agent is displayed in parentheses after the Agent name.
- The color of an Agent is a visual aid that allows users to quickly identify the status.

Agent can have the following statuses:

Status	Description
Con	Indicates Peer Agent is currently connected to the Peer Management

Status	Description
nected	Broker.
Disconnected	Indicates that Peer Agent has disconnected from the Peer Management Broker. This can be a result of stopping the Peer Agent, or if the network connection between the Peer Agent and the Peer Management Broker was severed.
Pending	This indicates that a heartbeat for the Peer Agent was not received within the configured threshold and that the Peer Agent is in the process on being disconnected if a heartbeat is not received soon. This status can also occur if the Peer Agent does not respond to a pending ping.
Unknown	If no connection status is displayed, then either the Peer Agent was not running on that host when Peer Management Center was started, or the first heartbeat message has not been received from that host.

Re-enabling a Disabled Agent Within a Job

Once disabled within a job, an Agent will not be involved in replication or locking. After the malicious activity that triggered MED is investigated and it is safe to re-enable the afflicted Agent, it will need to be re-enabled on a per job basis.

To review the status of an Agent within a job and to re-enable it, navigate to the **Participants** tab in the job's Runtime Summary view.

If an error is disabled because of a MED action, the message will be similar to the following:

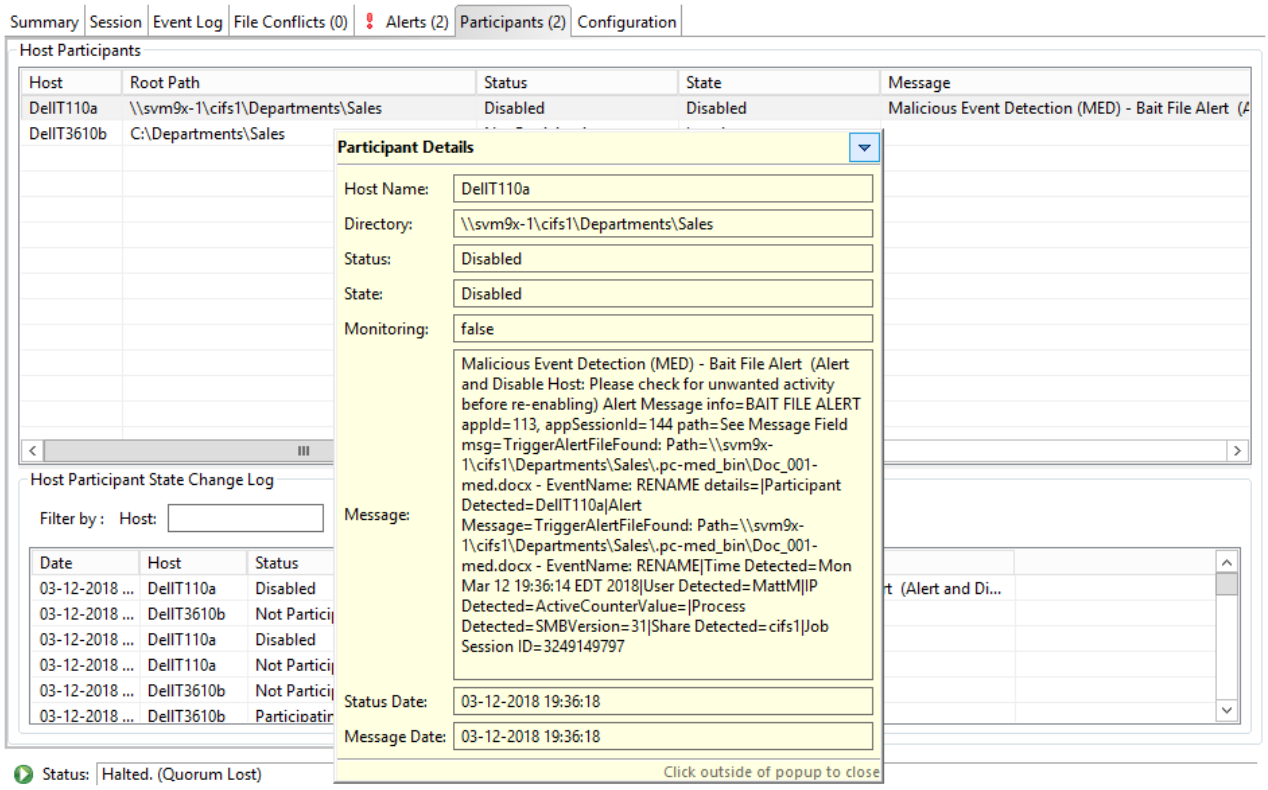

To re-enable the Agent, right-click it within this view, and select **Enable Host Participant**.

Editing an Agent Configuration

The ability to remotely manage the configuration of connected Peer Agents is available from within Peer Management Center.

Note: For customers using clustered file server roles with Windows Failover Cluster, review this knowledge base article: Windows Failover Cluster support for the Peer Agent. Custom Agent settings must be applied to each potential node in the cluster that may host the Peer Agent. Contact Peer Software Support for more information.

To edit an Agent's configuration:

1. Right-click the connected Peer Agent in the **Agents** view:

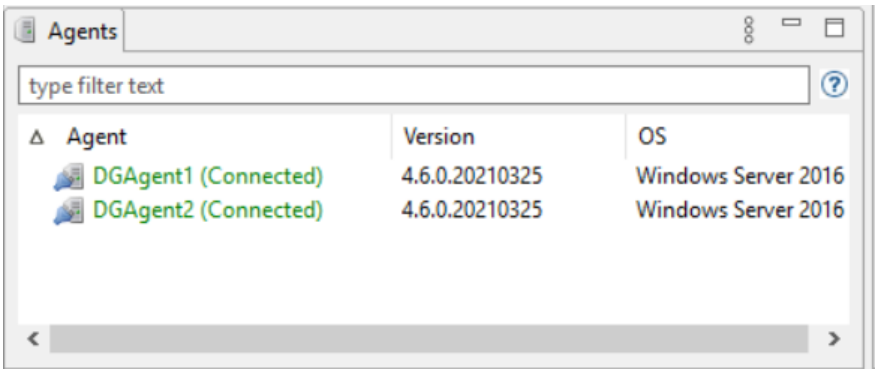

2. Select **Edit Agent Configuration**.

3. If the following dialog appears, Click **OK**:

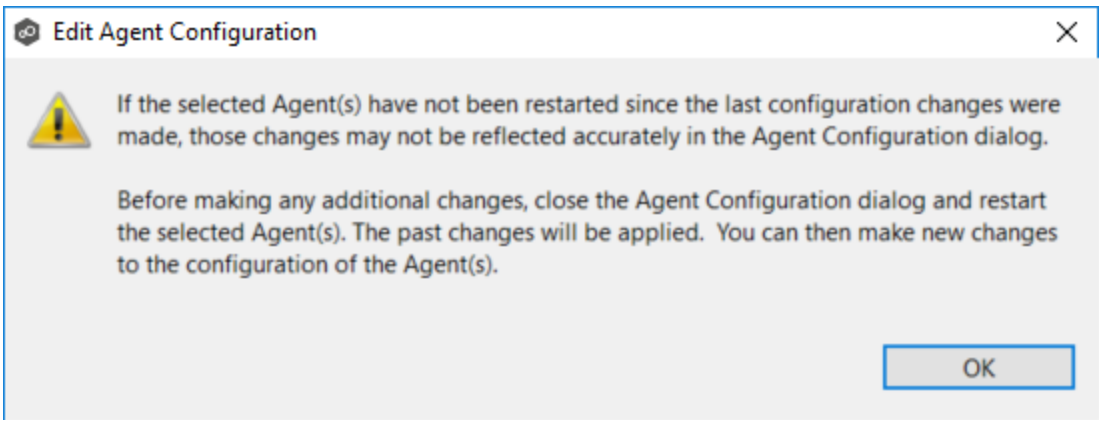

The **Broker Configuration** page appears.

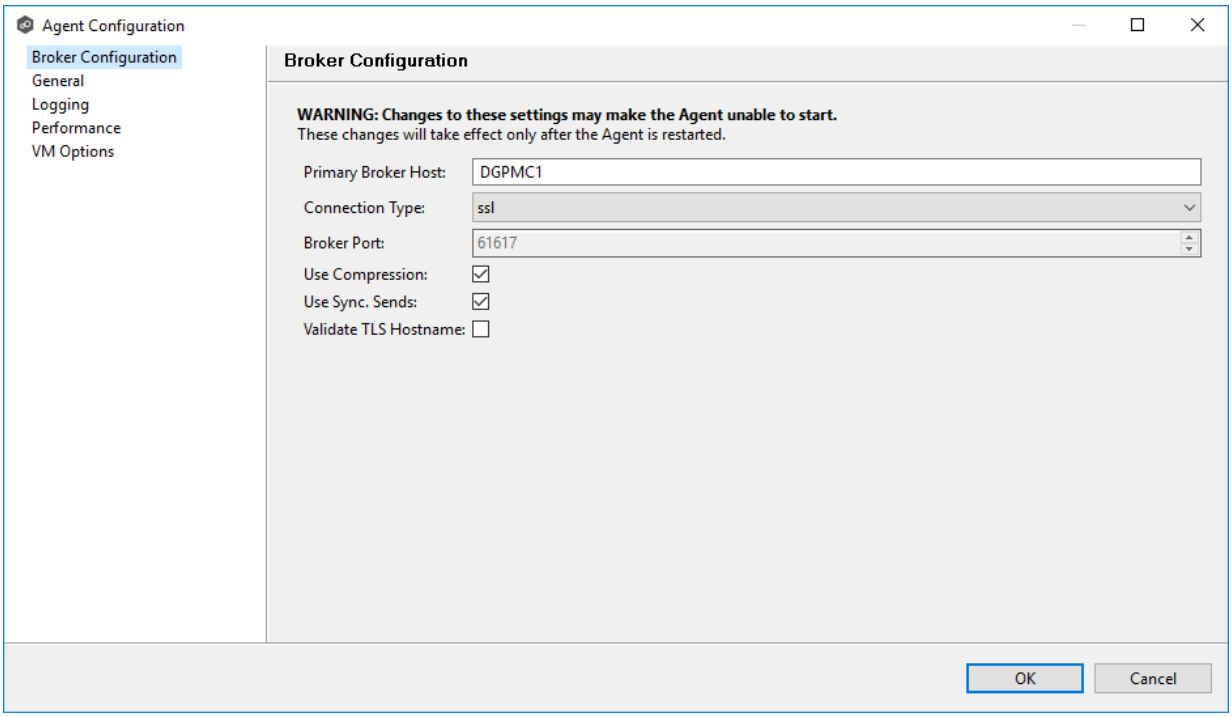

4. Select a page to edit and make the desired changes:

   - Broker Configuration General
   - General
   - Logging
   - Performance
   - VM Options

5. Click **OK**.

   For any configuration change to take effect, the selected Peer Agent must be restarted. If no jobs are running, you will have the option of restarting the Peer Agent at the close of the dialog.

Warning: Changes to any of these options may result in problems when the Peer Agent restarts. Ensure all settings are correct before closing the dialog and restarting the selected Peer Agent.

T
he settings in **Broker Configuration** apply to communication between the selected Peer Agent and Peer Management Broker only. It does not apply to communication between Peer Management Center and Peer Management Broker.

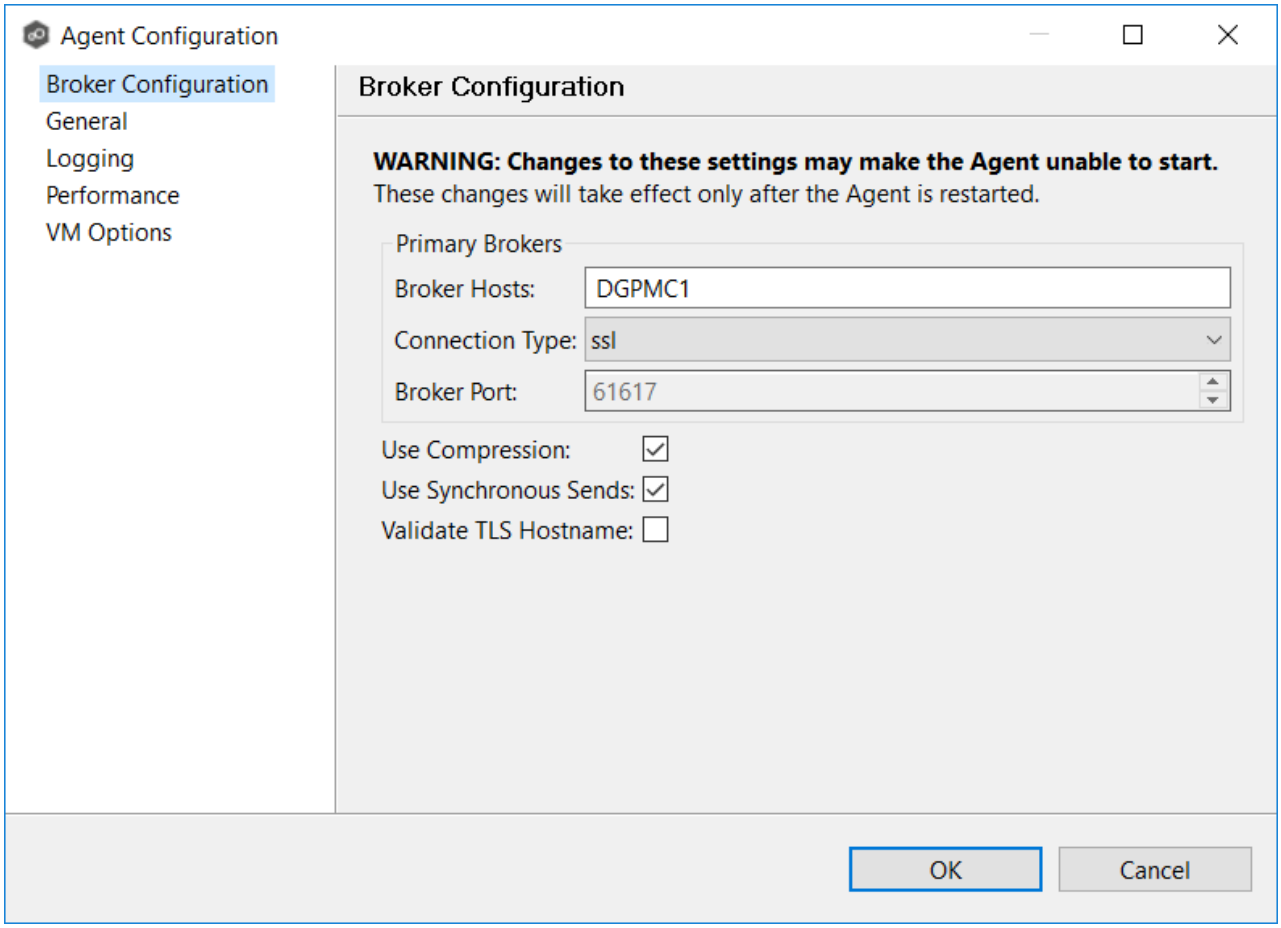

Options

Option	Description
Primary Broker Host	Enter the IP address or fully qualified host name of the server running Peer Management Center Broker. Can enter a comma-separated list of hosts. Agent will connect to any of the primary brokers in the order that they in listed. Will try to failover to a primary broker first before trying the failover brokers.

Option	Description
Connection Type	Select the type of connection to use when communicating with Peer Management Center Broker. Types include SSL (encrypted) and TCP (not encrypted).
Broker Port	The port on which to communicate with Peer Management Broker.
Use Compression	Enable to compress all communication between the selected Peer Agent and Peer Management Broker.
Use Synchronous Sends	Enable to always send messages from an Agent to Peer Management Broker in synchronous mode. If not enabled, then messages between Agent and Broker will always be sent asynchronously. Note: Enabling this will affect the performance of communication between the Broker and the Agent, especially over connections with high latency.
Validate TLS Hostname	Enable if you are using your own certificates and would like certificate hostnames to be validated between an Agent and Peer Management Center Broker.

The General page has two sets of options:

- Workspace
- Location Information

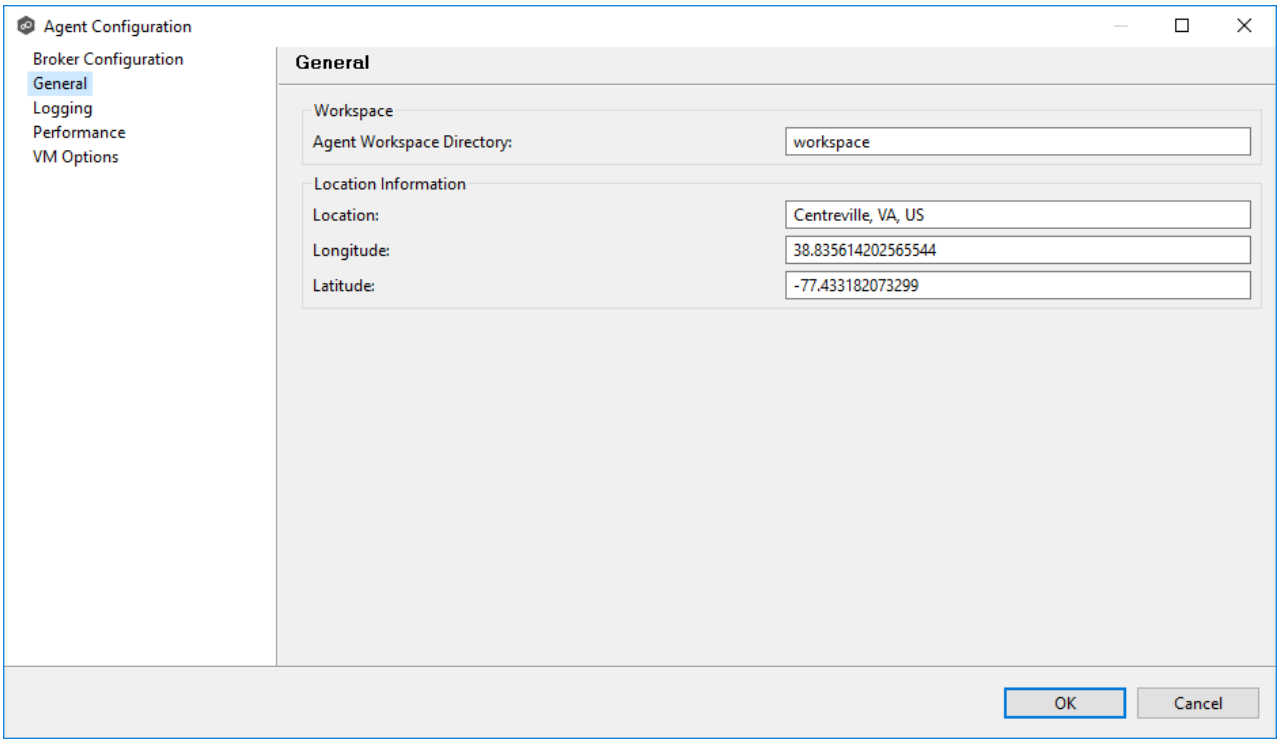

Options

Workspace

Option	Description
Agent Workspace Directory	Enter the directory where log files and other application data is stored. This path is relative to the Peer Agent's installation directory. It can also be set to an explicit full path.

Location Information

If you are using Peer Analytics, Agent location information is used by Analytics.

Option	Description
Location	Enter the city, state, and country where this Agent is located.

Option	Description
Longitude	Enter the longitude coordinates of this Agent.
Latitude	Enter the latitude coordinates of this Agent.

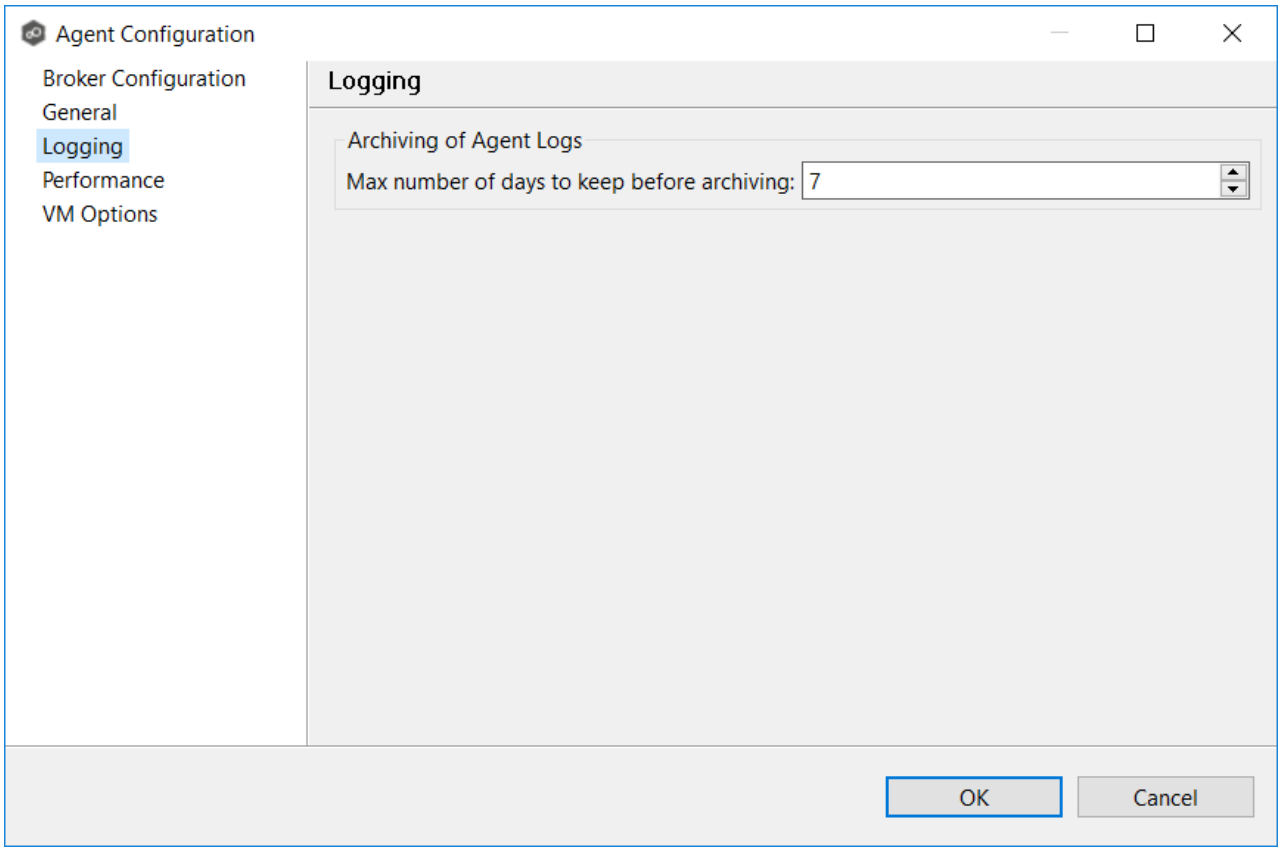

Option

Option	Description
Max number of days to keep before archiving	Log files that are older than this date will be relocated automatically to an archive folder, potentially reducing required space on disk. If the default Agent output log files rollover in less than the number of days selected, log bundles sent to Peer Software support may have gaps.

The **Performance** page offers control over settings that affect Agent performance. This page has two sets of options:

- Blob Transfers - These settings control the number of parallel streams of data that can be sent between the Agent and the Broker. In active, latent environments, adjusting these settings can improve performance or limit the data throughput between the Agent and the Broker.

- Processor Affinity - Allows you to specify the number of processors that the Agent should use.

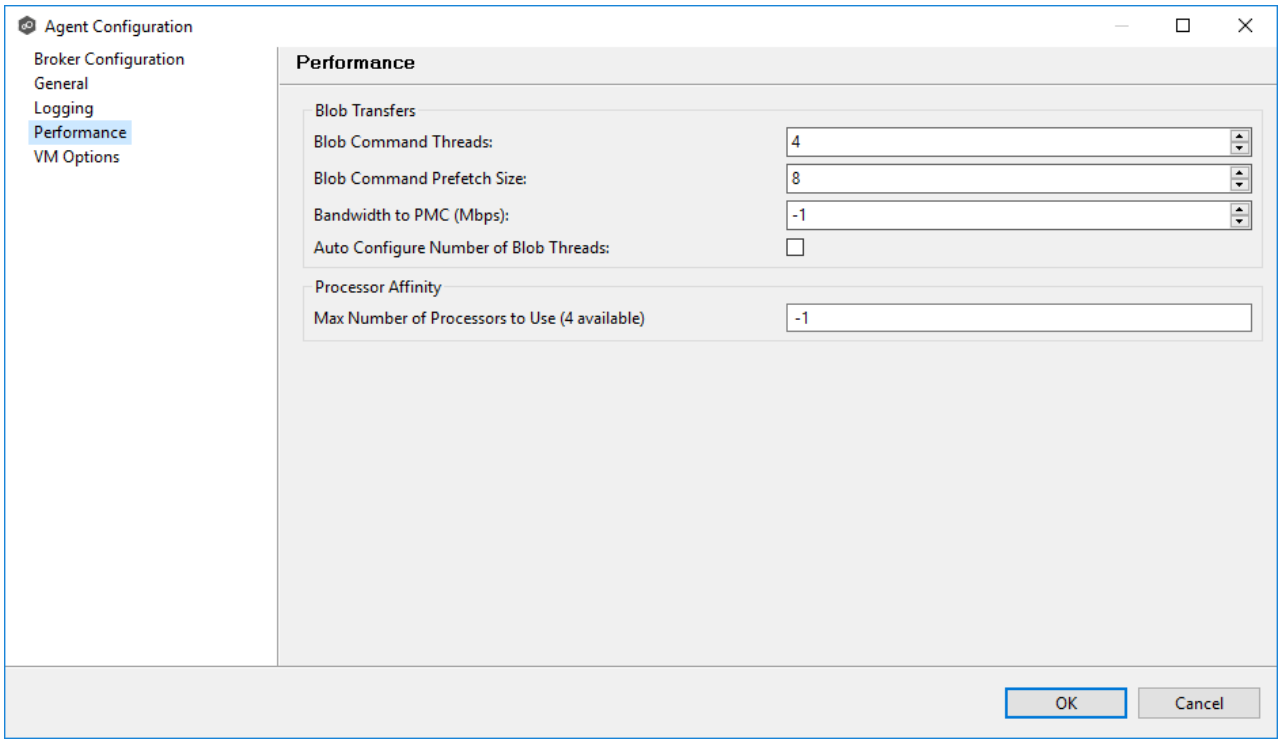

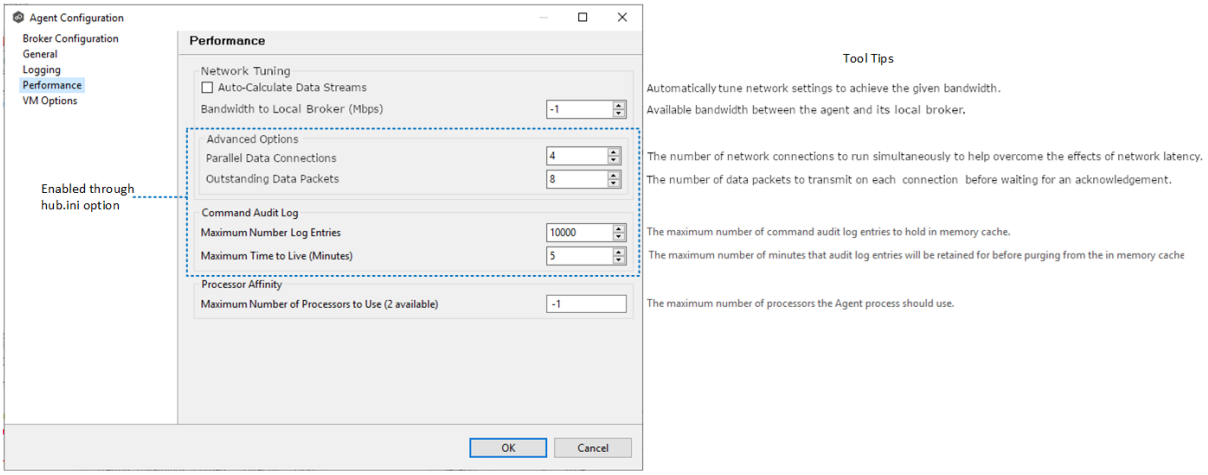

Options

Blob Transfers

Field	Description
Blob Command Threads	Enter the maximum number of threads to handle data transfer between each Agent and the Peer Management Broker. Increasing this typically improves replication performance but also increases memory consumption. The default value is 4 threads. The minimum is 1; the maximum is 100.
Blob Command Prefetch Size	Modify this setting only at the instruction of the Peer Software Technical Support Team as it can lead to increased memory consumption. Enter the maximum number of blocks of data to be buffered to be sent to the Agent. The default number is 8 commands; the maximum size is 100 commands.
Bandwidth to PMC (Mbps)	Enter the bandwidth in megabits per second that you want to use for the connection between the Agent and Peer Management Broker. The default value is -1, which means use all available bandwidth.
Auto Configure Number of Blob Threads	Select this checkbox if you want the number of blob command threads to be calculated rather than using the value in the **Blob Command Threads** field. The optimum number of Agent Blob Connections is calculated based on network performance, using the value of **Bandwidth to PMC (Mbps)** in the calculation along with latency between the Broker and Agent.

Processor Affinity

Field	Description
Max Number of Processors to Use (x available)	Enter the number of processors that the Agent process can use on the server where it is installed. This number should be less than or equal to the number of processors available on the server.

Field	Description
	The default value is -1, which means use all available processors.

The option on the page allows you to tune the maximum amount of system memory that the Peer Agent service will use on the server where it is installed. We strongly recommend that this value be set to no lower than 2 GB.

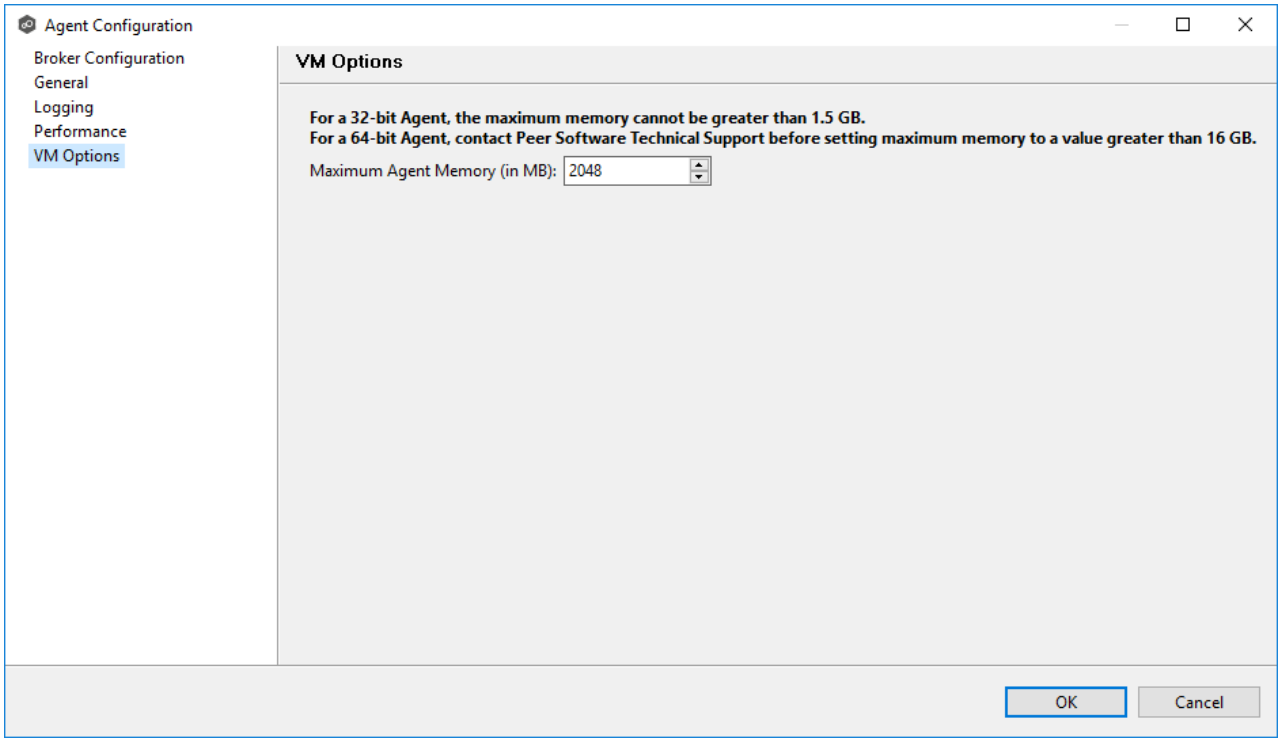

Options

Field	Description
Maximum Agent Memory (in MB)	Enter the maximum amount of memory in megabytes that the JVM portion of the Agent service can use. We recommend a minimum value of 2048 MB on 64-bit Agent servers with a recommended maximum of 16384 MB.

Viewing Agent Properties

To view the properties of an Agent:

1. Right-click the Agent in the **Agents** view.
2. Select **View Properties**.

   The **View Agent Properties** dialog opens.

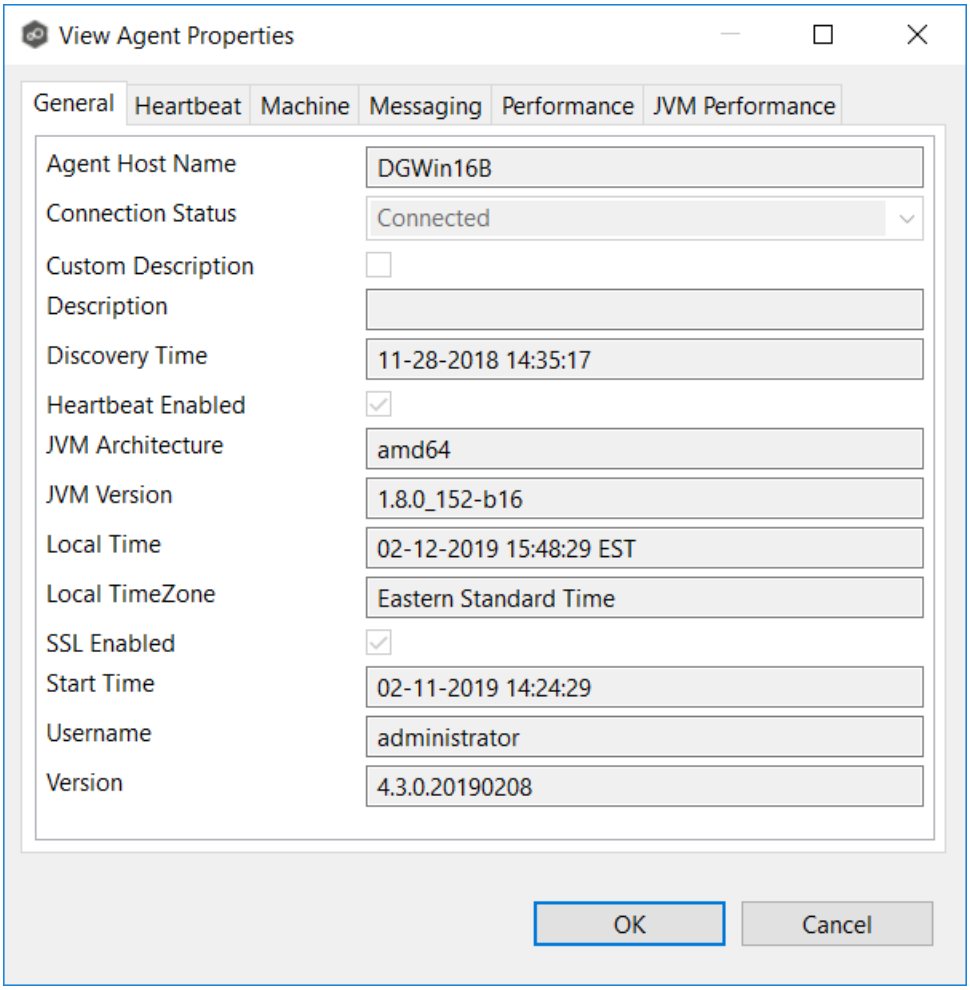

This dialog displays Peer Agent and host machine information across the following tabs:

Tab	Description
General	Displays general Peer Agent run-time information such as discovery time, local time, TLS use, Peer Agent start up time, Peer Agent version, and the user name Peer Agent service is running as.
Heartbeat	Displays heartbeat information and statistics such as heartbeat frequency, average heartbeat time, last heartbeat time, total Peer Agent disconnects, total missing heartbeats.
Machine	Displays machine information of the host that the Peer Agent is running on such as number of processors, computer name, domain name, IP address, installed memory, O/S.
Messaging	Displays general Peer Management Center Broker messaging statistics for the selected host, such as total messages received, total messages sent, # errors.
Performance	Displays general performance statistics for the underlying host machine such as available virtual memory, available physical memory, memory load.
JVM Performance	Displays JVM performance statistics for the running Peer Agent application such as active number of threads, heap memory used, non-heap memory used.

Editing Agent Properties

Selecting **Edit Properties** menu item for a selected agent will result in the opening of the following Peer Agent **Properties** dialog:

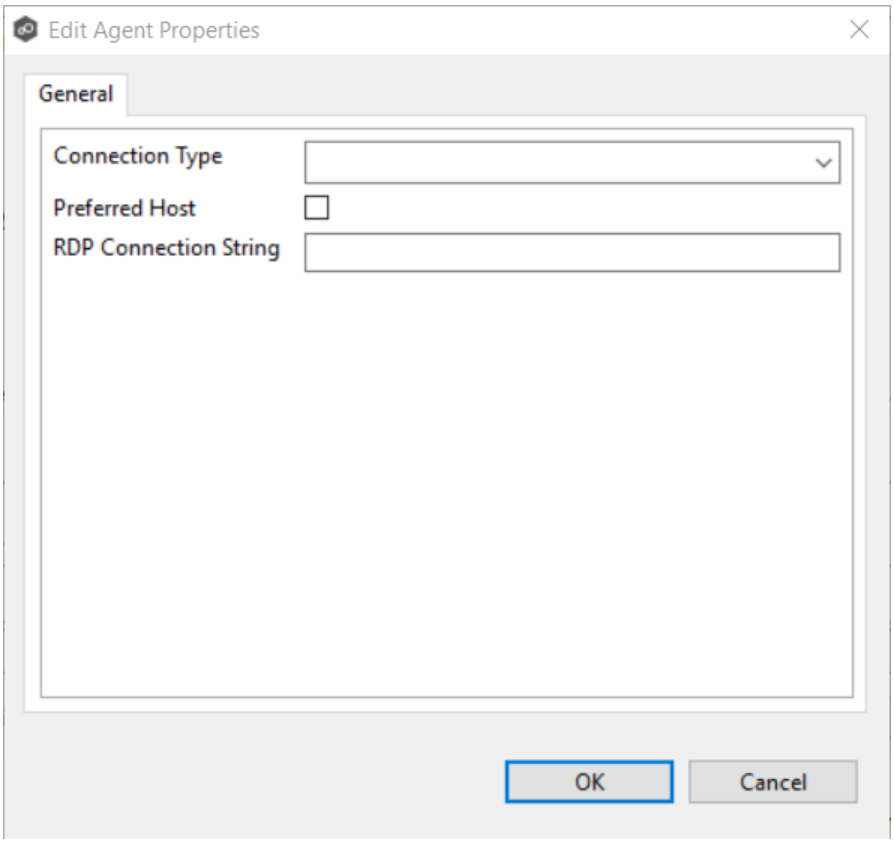

This dialog displays the following configurable Peer Agent and host machine options:

Option	Description
Connection Type	Allows for the selection of a connection type between the selected Peer Agent and the associated Peer Management Broker. When set, optimizations are made to the communication between the two parties based on the selected connection type.
Preferred Host	A best practice optimization for selecting which Peer Agent has the fastest connection to the Peer Management Broker (or in appropriate cases, for selecting which Peer Agent are on the same subnet as the Peer Management Broker).
RDP Connection String	The connection string to use when activating a **Remote Desktop Protocol** (**RDP**) session to this Peer Agent.

Updating a Peer Agent

If the Peer Agent software running on a host is out of date, the host is shown as having a pending update in the Agents view.

When right-clicking the host, the option to automatically update the Peer Agent software is also available. This process can be done from Peer Management Center and usually does not require any additional actions on the host server itself.

PeerGFS API

The PeerGFS API is a RESTful API. It allows system administrators to monitor PeerGFS activity and developers to integrate PeerGFS functionality into their own application.

Currently, the API allows users to:

- Get information about running jobs, such as open files as well as files in the process of being synchronized; statistical info about watch set, queue sizes, replication metrics; scan status, alerts, and quarantined files.
- Start and stop jobs.
- View and restart agents.
- View scheduled tasks.
- Trigger log uploads.

Additional functionality, such as the ability to create and edit jobs, will be provided in future versions of the API.

Accessing the PeerGFS API

Access to the PeerGFS API is available as a combination of two elements:

- A web URL hosted by the API service. This URL is defined as a combination of a PMC server name or IP and a port, as specified by the Web and API Configuration settings in Preferences.

- Local (aka basic) authentication with a user name and password that is passed into a script or the API web interface. This user name and password is used to authenticate the user with the PeerGFS API service.

If you are authenticated, you are authorized to access the entire API. Role-based access will be added in future versions of the API.

Testing the PeerGFS API

One way to test the PeerGFS API is to use the API web interface.

To access the web interface, open a browser, go to the API endpoint (e.g., https://<PMC IP or name> or <8442>), and try the API calls.

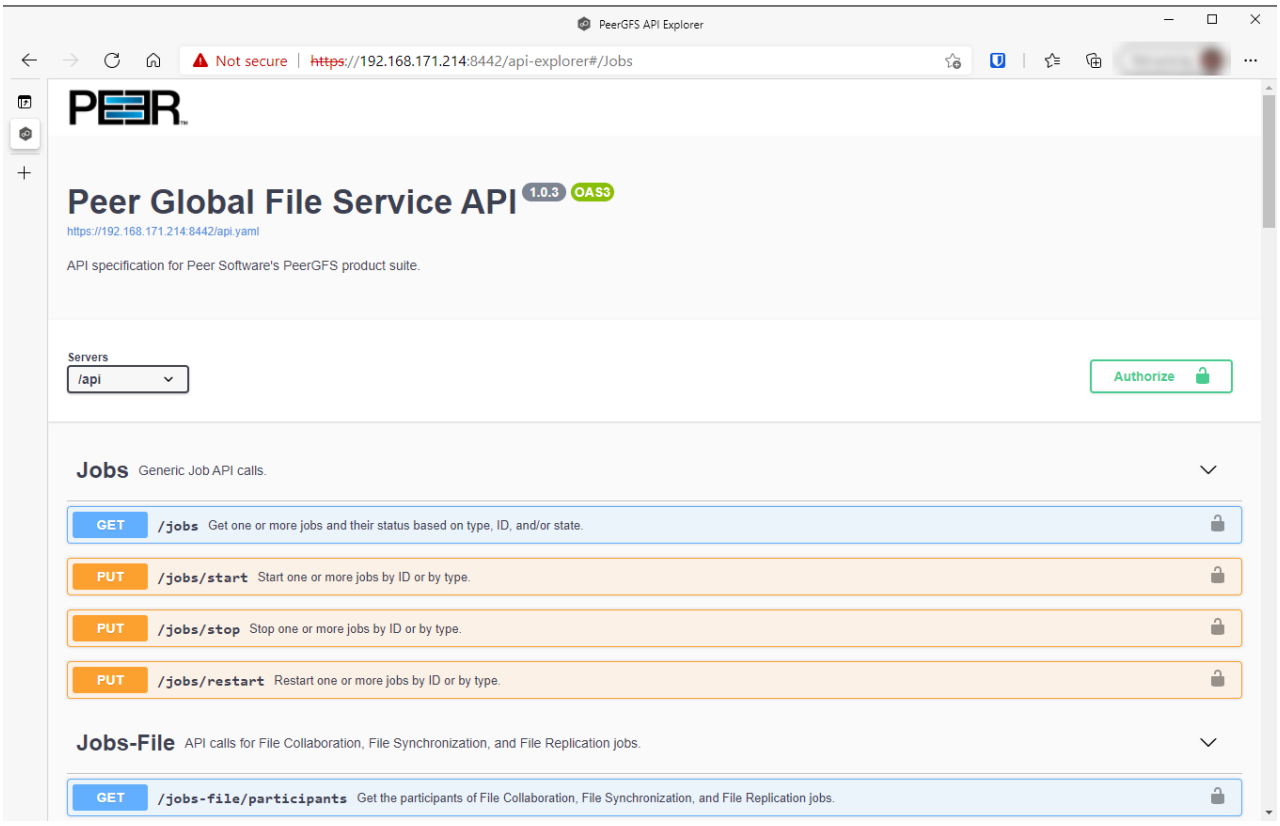

Integrating Your Own Tools and Scripts with the PeerGFS API

The PeerGFS installation folder contains PowerShell and Bash toolkits in the **tools** subfolder of the PMC's installation folder. If you need a different language, contact Peer Technical Support for our latest YAML file.

If you would like to access the API through a client in a language other than PowerShell or BASH, you can use the Swagger Editor to convert our YAML file to the appropriate client code:

1. Save the PeerGFS YAML file to your desktop.
2. Open a web browser tab and point it to https://editor.swagger.io/#/.
3. Go to the **File** menu inside the web interface, and then select **Import file**.
4. Select the PeerGFS YAML file on the desktop.

   The manifest should appear on the left with the front-end mockup on the right.
5. Use the Swagger Editor to generate client code.

API Quick Reference

The PeerGFS API REST specifications are documented using OpenAPI (also known as Swagger). This documentation is visible via the PMC API's web interface. To access the web interface, see Testing the PeerGFS API.

Within the API web interface, you can also send test requests and view responses as well as see REST calls that can be made to the API service.

The PeerGFS is divided into four types of API calls:

- Jobs - Generic job-related calls
- Jobs-File - Job-specific calls
- Agents - Agent-related calls
- PMC - Calls related to alerts, tasks, and logs

The PeerGFS API has three status codes:

200 - Success

401 - Unauthorized

404 - Job(s) not found

Scheduled Replication

Use a scheduled replication filter to identify files and folders that you don't want replicated in real-time--instead, they will be replicated at a scheduled time or interval. If a file or folder meets the filter criteria, instead of being processed in real-time, it will be placed in a queue for synchronization and processed as scheduled.

Scheduled replication filters can be used with file collaboration, file replication, and file synchronization jobs. For information on defining a scheduled replication filter, see Scheduled Replication in Preferences.

Note: When a scheduled replication filter is used in a file collaboration job, files that meet the filter criteria will not be locked.

Smart Data Seeding

Overview

Smart data seeding applies to File Collaboration, File Replication, and File Synchronization jobs.

Occasionally, a new host or a host which has been removed from the session for a long time, needs to be introduced into an existing collaboration. Smart Data Seeding supports integrating new hosts into a collaboration seamlessly. Conventional seeding methods take a long time over typically slow WAN connections and require a cut-over with a final scan to get the data synchronized. With Smart Data Seeding's default settings, real-time events are processed from the Smart Data Seeding hosts while the initial one-way background scan ensures the target(s) have all the files in place.

Smart Data Seeding provides the ability to set one or more participants in a Smart Data Seeding mode. Smart Data Seeding hosts are considered the hosts from where files will be copied to all the other participants in the session. When a host is in Smart Data Seeding mode, it follows the rules of the job’s Smart Data Seeding Mode configuration (see below). Initial

scans run in a one-way mode to avoid bringing back deleted files. It is not recommended to have active (Active-Active) users on the target hosts. Once the initial scan is completed, the Smart Data Seeding host(s) are set back to their default full collaboration mode with no user interaction or final scan.

To enable advanced settings in the Conflict Resolution window, add the following fc.ini option and restart Peer Management Center:

```
fc.scan.enable.preseeding.ui=true
```

Smart Data Seeding Options

From the **Conflict Resolution** window, select from one of the following Smart Data Seeding modes:

Mode	Description
PASSIVE (Default)	Initial scan will be one-way only with any host in Smart Data Seeding mode: • Real-time activity on Smart Data Seeding host is disabled. • Real-time events on that host will be quarantined. • Renamed files will be restored.
PASSIVE_WITH_RESTORE	Initial scan will be one-way only with any host in Smart Data Seeding mode: • Real-time activity on Smart Data Seeding host is disabled. • Any activity on that host will be restored to its original state.
ACTIVE_LIMITED	Initial scan will be one-way only with any host in Smart Data Seeding mode: • Real-time activity on Smart Data Seeding host is enabled in a limited mode (real-time file adds are processed). • Unsynchronized file updates will be quarantined. • Unsynchronized file renamed will be restored. • Unsynchronized file deletes will be restored.

Mode	Description
ACTIVE_FULL	Initial scan will be one-way only with any host in Smart Data Seeding mode except for updates (updates will be processed as Latest Modified wins): • Real-time activity on Smart Data Seeding host is enabled with latest modified file wins, regardless of whether the latest file is on the Smart Data Seeding host.
REACTIVATION	Initial Scan will be one-way only with any host in Smart Data Seeding mode: • Real-time activity on Smart Data Seeding host is enabled with Quarantine (Added and Updated Files will be quarantined during the scan). • Unsynchronized file updates will be quarantined during real-time. • Unsynchronized file renames will be restored. • Unsynchronized deletes will be restored.

The default setting is ACTIVE_LIMITED, which will initiate a one-way scan with any host in Smart Data Seeding mode. During the scan, new files will be deleted, newer files will be overwritten, and deleted files will be restored on the Target(s). During real-time activity, add events will be processed, but updates will be quarantined if the files are unsynchronized. Renames and deletes will be restored if the files are unsynchronized.

The ACTIVE_LIMITED setting is recommended in most cases in which a new host or a host which has been removed from the session for a long time needs to be introduced into an existing collaboration.

Storage Capacity

The storage capacity available for your jobs is based on your Peer Global File Service license. Automated alerts will notify you when you close to reaching your licensed storage capacity. If you exceed your licensed storage capacity, contact your Peer Software sales representative.

Total capacity consumed is defined by the total number of unique TBs under management across all participants rather than the total capacity used by all participants. In this unique TB model, a 1 TB file that is synchronized across 10 participants only counts as 1 TB and not 10 TBs. For example, if your licensed storage capacity is 100 TB and you have a job with 5

participants totaling 20 unique TBs, you have used total of 20% of your storage capacity, not 100%.

TLS Certificates

You can use custom or private Transport Layer Security (TLS) certificates to connect a Peer Agent to Peer Management Broker. The Keytool certificate management utility will be used to store the key and certificate into a keystore file, which protects the private keys with a password.

Note the paths in the following topics reference a default install directory for both Peer Management Center and Peer Agent.

For step-by-step instructions, see:

- Creating New Certificates
- Using Existing Certificates

For additional information, please contact Peer Software's support team via email: support@peersoftware.com.

Creating New Certificates

Keytool Utility

Perform the necessary commands using the Keytool certificate management utility bundled with your Peer Management Center or Peer Agent installation. The location of the utility is:

- Peer Management Center system: PMC_INSTALLATION_FOLDER\jre\bin
- Peer Agent system: PEER_AGENT_INSTALLATION_FOLDER\jre\bin

Broker Keystore Generation

Step 1. Using the Keytool utility, create a certificate for Peer Management Center.

```
keytool -genkey -alias broker -keyalg RSA -keystore broker.ks -storepass
plBroker4321 -validity 3000
```

broker	The alias of the new broker keystore containing the new certificate.
broker.ks	Destination broker keystore that will be created containing the new certificate.
plBroker4321	The password you assign to the new broker keystore.

Note: The broker.ks file will be created in the \jre\bin folder.

Example:

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -genkey -alias
broker -keyalg RSA -keystore broker.ks -storepass plBroker4321 -validity 3000
What is your first and last name?
  [Unknown]:  Monika Cuellar
What is the name of your organizational unit?
  [Unknown]:  Peer Software, Inc.
What is the name of your organization?
  [Unknown]:  Peer Software, Inc.
What is the name of your City or Locality?
  [Unknown]:  Centreville
What is the name of your State or Province?
  [Unknown]:  VA
What is the two-letter country code for this unit?
  [Unknown]:  US
Is CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=VA, C=US
correct?
  [no]:  yes

Enter key password for <broker>
        (RETURN if same as keystore password):

C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>
```

Step 2: Export the broker's certificate so it can be shared with clients.

```
keytool -export -alias broker -keystore broker.ks -file broker.cer
```

broker	The alias of the new broker keystore containing the new certificate.
broker.ks	Destination broker keystore that will be created containing the new certificate.
broker.cer	The name of the broker's certificate to be created.

Note: The broker.cer file will be created in the \jre\bin folder.

Example:

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -export -alias
broker -keystore broker.ks -file broker.cer
Enter keystore password:  plBroker4321
Certificate stored in file <broker.cer>

C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>
```

Step 3: Create a certificate/keystore for the client.

```
keytool -genkey -alias client -keyalg RSA -keystore client.ks -storepass
plClient4321 -validity 3000
```

client	The alias of the new client keystore containing the new certificate.
client.ks	Destination keystore for the client that will be created containing the new certificate.
plClient4321	The password you assign to the new client keystore.

Note: The client.ks file will be created in the \jre\bin folder.

Example:

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -genkey -alias
client -keyalg RSA -keystore client.ks -storepass plClient4321 -validity 3000
What is your first and last name?
  [Unknown]:  Monika Cuellar
What is the name of your organizational unit?
  [Unknown]:  Peer Software, Inc.
What is the name of your organization?
  [Unknown]:  Peer Software, Inc.
What is the name of your City or Locality?
  [Unknown]:  Centreville
What is the name of your State or Province?
  [Unknown]:  VA
What is the two-letter country code for this unit?
  [Unknown]:  US
Is CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville, ST=VA,
C=US
correct?
  [no]:  yes

Enter key password for <client>
        (RETURN if same as keystore password):

C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>
```

Step 4: Create a truststore for the client and then import the broker's certificate. This establishes that the client "trusts" the broker.

```
keytool -import -alias broker -keystore client.ts -file broker.cer -
storepass plClient4321
```

broker	The alias of the broker keystore created in step 1.
client.ts	Destination truststore for the client that will be created containing the broker's certificate.
broker.cer	The broker's certificate created in step 2.
plClient4321	The password assigned to the client keystore in Step 3.

Example:

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -import -alias
broker -keystore client.ts -file broker.cer -storepass plClient4321
Owner: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=VA, C
=US
Issuer: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=VA,
C=US
Serial number: 4fa7f34f
Valid from: Mon May 07 12:07:43 EDT 2012 until: Fri Jul 24 12:07:43 EDT 2020
Certificate fingerprints:
         MD5:  2C:18:DD:B5:CD:C5:3D:B2:9B:E3:93:50:D6:74:2B:64
         SHA1: 30:77:94:9B:34:63:6C:DE:2C:98:9C:00:C2:B9:F6:21:AE:22:D7:DE
Trust this certificate? [no]:  yes
Certificate was added to keystore

C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>
```

Optional: List the certificates in the broker keystore.

```
keytool -list -v -keystore broker.ks -storepass plBroker4321
```

Example:

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -list -v -
keystore broker.ks -storepass plBroker4321

Keystore type: jks
Keystore provider: SUN

Your keystore contains 1 entry

Alias name: broker
Creation date: May 7, 2012
Entry type: keyEntry
Certificate chain length: 1
Certificate[1]:
Owner: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=VA, C=US
Issuer: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=VA, C=US
Serial number: 4fa7f34f
Valid from: Mon May 07 12:07:43 EDT 2012 until: Fri Jul 24 12:07:43 EDT 2020
Certificate fingerprints:
         MD5:  2C:18:DD:B5:CD:C5:3D:B2:9B:E3:93:50:D6:74:2B:64
         SHA1: 30:77:94:9B:34:63:6C:DE:2C:98:9C:00:C2:B9:F6:21:AE:22:D7:DE


*******************************************
*******************************************

C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>
```

Verify Client Certificate

If you want to verify client certificates, you need to take a few extra steps.

Step 1: Export the client's certificate so it can be shared with broker.

```
keytool -export -alias client -keystore client.ks -file client.cer -
storepass plClient4321
```

Note: The client.cer file will be created in the \jre\bin folder.

Example:

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -export -alias
client -keystore client.ks -file client.cer -storepass plClient4321
Certificate stored in file <client.cer>

C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>
```

Step 2: Create a truststore for the broker and import the client's certificate. This establishes that the broker "trusts" the client:

```
keytool -import -alias client -keystore broker.ts -file client.cer -
storepass plBroker4321
```

Example:

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -import -alias
client -keystore broker.tx -file client.cer -storepass plBroker4321
Owner: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=VA, C
=US
Issuer: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=VA,
C=US
Serial number: 4fa7f982
Valid from: Mon May 07 12:34:10 EDT 2012 until: Fri Jul 24 12:34:10 EDT 2020
Certificate fingerprints:
         MD5:  A7:D9:6E:78:8B:A9:AD:32:96:2D:51:6B:53:0B:E4:BD
         SHA1: 16:05:7C:C4:D5:AB:E7:D3:7D:5B:2E:02:B5:3B:69:54:D1:C3:53:52
Trust this certificate? [no]:  yes
Certificate was added to keystore

C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>
```

Optional: List the certificates in the client keystore.

```
keytool -list -v -keystore client.ks -storepass plClient4321
```

Example:

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -list -v -
keystore client.ks -storepass plClient4321

Keystore type: jks
Keystore provider: SUN

Your keystore contains 1 entry

Alias name: client
Creation date: May 7, 2012
Entry type: keyEntry
Certificate chain length: 1
Certificate[1]:
Owner: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=NY,
C=US
Issuer: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=NY,
 C=US
Serial number: 4fa80618
Valid from: Mon May 07 13:27:52 EDT 2012 until: Fri Jul 24 13:27:52 EDT 2020
Certificate fingerprints:
         MD5:  06:11:97:71:D6:23:91:63:2F:19:F4:05:EA:2F:9D:14
         SHA1: A7:26:80:9E:18:2B:46:8E:92:BB:AD:89:44:0A:8A:9C:8C:1F:62:38


*******************************************
*******************************************



C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>
```

Copy the Generated Keystore Files into Their Appropriate Location

On the Peer Management Center system: Copy the following files from the "C:\Program Files\Peer Software\Peer Management Hub\jre\bin" directory into the "C:\Program Files\Peer Software\Peer Management Hub\Broker\keys" directory on the Peer Management Center system. Overwrite the existing files.

broker.ks

broker.ts

On the Peer Agent system: Copy the following files from the "C:\Program Files\Peer Software\Peer Management Hub\jre\bin" directory into the "C:\Program Files\Peer Software\Peer Agent\keys" directory on the Peer Agent systems. Overwrite the existing files.

client.ks

client.ts

Restart All Peer Management Center Services for the Changes to Take Effect

We recommend you create a folder outside the Peer Management Center/Peer Agent installation directories in which to store the keystore files. This will ensure that upgrades will not clear/overwrite these files. The steps outlining this process will be posted shortly.

Using Existing Certificates

Keytool Utility

Perform the necessary commands using the Keytool certificate management utility bundled with your Peer Management Center or Peer Agent installation. The location of the utility is:

- Peer Management Center system: PMC_INSTALLATION_FOLDER\jre\bin
- Peer Agent system: PEER_AGENT_INSTALLATION_FOLDER\jre\bin

Peer Management Broker and Peer Agent Keystore Generation

You will need to have two custom/private certificates. One for the Peer Management Broker and one for all the participating Peer Agents. You may select different algorithms and encryption key size (i.e., RSA, DSA with 1024 or 2048 key size).

Step 1. Using the Keytool utility, list the contents of the custom/private certificates. Perform these steps for both certificates (Peer Management Broker and Peer Agent. Make a note of the Alias of the certificate, if it exists.

```
keytool -list -v -keystore HubCert.pfx -storetype pkcs12
```

HubCert.pfx	Represents the custom/private certificate for Peer Management Center Broker.
AgentCert.pfx	Represents the custom/private certificate for the Peer Agents.

Note: The command will prompt you to enter the password you set on your custom certificate, if applicable.

Step 2. Add the custom/private Peer Management Center Broker certificate into the Peer Management Center Broker keystore.

```
keytool -importkeystore -deststorepass plBroker4321 -destkeypass
plBroker4321 -destkeystore broker.ks -srckeystore HubCert.pfx -
srcstoretype PKCS12 -srcstorepass PASSWORD -alias ALIAS -destalias broker
```

plBroker4321	The password you assign to the new Broker keystore.
broker.ks	Destination keystore that will be created containing the custom/private certificate.
HubCert.pfx	Custom/private certificate being imported into the new keystore.
PASSWORD	The password of the custom/private certificate, if it exists. If you omit the -srcstorepass command, you will be prompted for the certificate password if needed.
ALIAS	The Alias of the custom/private certificate you discovered in Step 1 above.

broker	The Alias of the new keystore containing the custom/private.

Note: The broker.cer and broker.ks files will be created in the \jre\bin folder where the keytool utility resides.

Step 3. Add the custom/private Peer Agent certificate into the Client keystore.

```
keytool -importkeystore -deststorepass plClient4321 -destkeypass
plClient4321 -destkeystore client.ks -srckeystore AgentCert.pfx -
srcstoretype PKCS12 -srcstorepass PASSWORD -alias ALIAS -destalias client
```

plClient 4321	The password you assign to the new Broker keystore.
client.ks	Destination keystore that will be created containing the custom/private certificate.
AgentCert.pfx	Custom/private certificate being imported into the new keystore.
PASSWORD	The password of the custom/private certificate, if it exists. If you omit the -srcstorepass command, you will be prompted for the certificate password if needed.
ALIAS	The Alias of the custom/private certificate you discovered in Step 1 above.
client	The Alias of the new keystore containing the custom/private.

Note: The client.cer and client.ks files will be created in the \jre\bin folder where the keytool utility resides.

Step 4. Export the broker's certificate so it can be shared with clients.

```
keytool -export -alias broker -keystore broker.ks -file broker.cer
```

broker	The Alias of the broker keystore containing the custom/private certificate created in Step 2 above.
broker.ks	The keystore file created in Step 2 above containing the custom/private certificate for the Broker.
broker.cer	The certificate file created in Step 2 above.

The command will prompt you to enter the password for the broker keystore (e.g., plBroker4321).

Step 5. Export the client's certificate so it can be shared with broker.

```
keytool -export -alias client -keystore client.ks -file client.cer
```

client	The Alias of the client keystore containing the custom/private certificate created in Step 3 above.
client.ks	The keystore file created in Step 3 above containing the custom/private certificate for the Peer Agents.
client.cer	The certificate file created in Step 3 above.

The command will prompt you to enter the password for the client keystore (e.g., plClient4321).

Step 6. Create a truststore for the broker and import the client's certificate. This establishes that the broker "trusts" the client:

```
keytool -import -alias client -keystore broker.ts -file client.cer
```

client	The Alias of the client keystore containing the custom/private certificate created in Step 3 above.

broker.ts	The broker trust store to be created.
client.cer	The certificate file created in Step 3 above.

The command will prompt you to enter the password for the broker keystore (e.g., plBroker4321).

Step 7. Create a truststore for the client and import the broker's certificate. This establishes that the client "trusts" the broker.

```
keytool -import -alias broker -keystore client.ts -file broker.cer
```

broker	The Alias of the client keystore containing the custom/private certificate created in Step 3 above.
client.ts	The client truststore to be created.
client.cer	The certificate file created in Step 2 above.

The command will prompt you to enter the password for the client keystore (e.g., plClient4321).

Copy the Generated Keystore Files into Their Appropriate Location

On the Peer Management Center system:

Copy the following files from the **Peer Management Center_INSTALLATION_FOLDER\jre\bin** directory into **the Peer Management Center_INSTALLATION_FOLDER\Broker\keys** directory on the Peer Management Center system. Overwrite the existing files.

broker.ks

broker.ts

On the Peer Agent system:

Copy the following files from **Peer Management Center_INSTALLATION_FOLDER\jre\bin** directory into the **PEER_AGENT_INSTALLATION_FOLDER\keys** directory on the Peer Agent systems. Overwrite the existing files.

client.ks

client.ts

Restart All Peer Management Center Services for the Changes to Take Effect

We recommend you create a folder outside the Peer Management Center/Peer Agent installation directories in which to store the keystore files. This will ensure that upgrades will not clear/overwrite these files.

Preferences

The **Preferences** dialog enables you to configure global settings, as well as settings specific to a job type. Before creating any jobs or configuring individual aspects of a job, Peer Software recommends first configuring a number of settings. Some settings are global and apply program-wide and/or to all job types; others are specific to a job type.

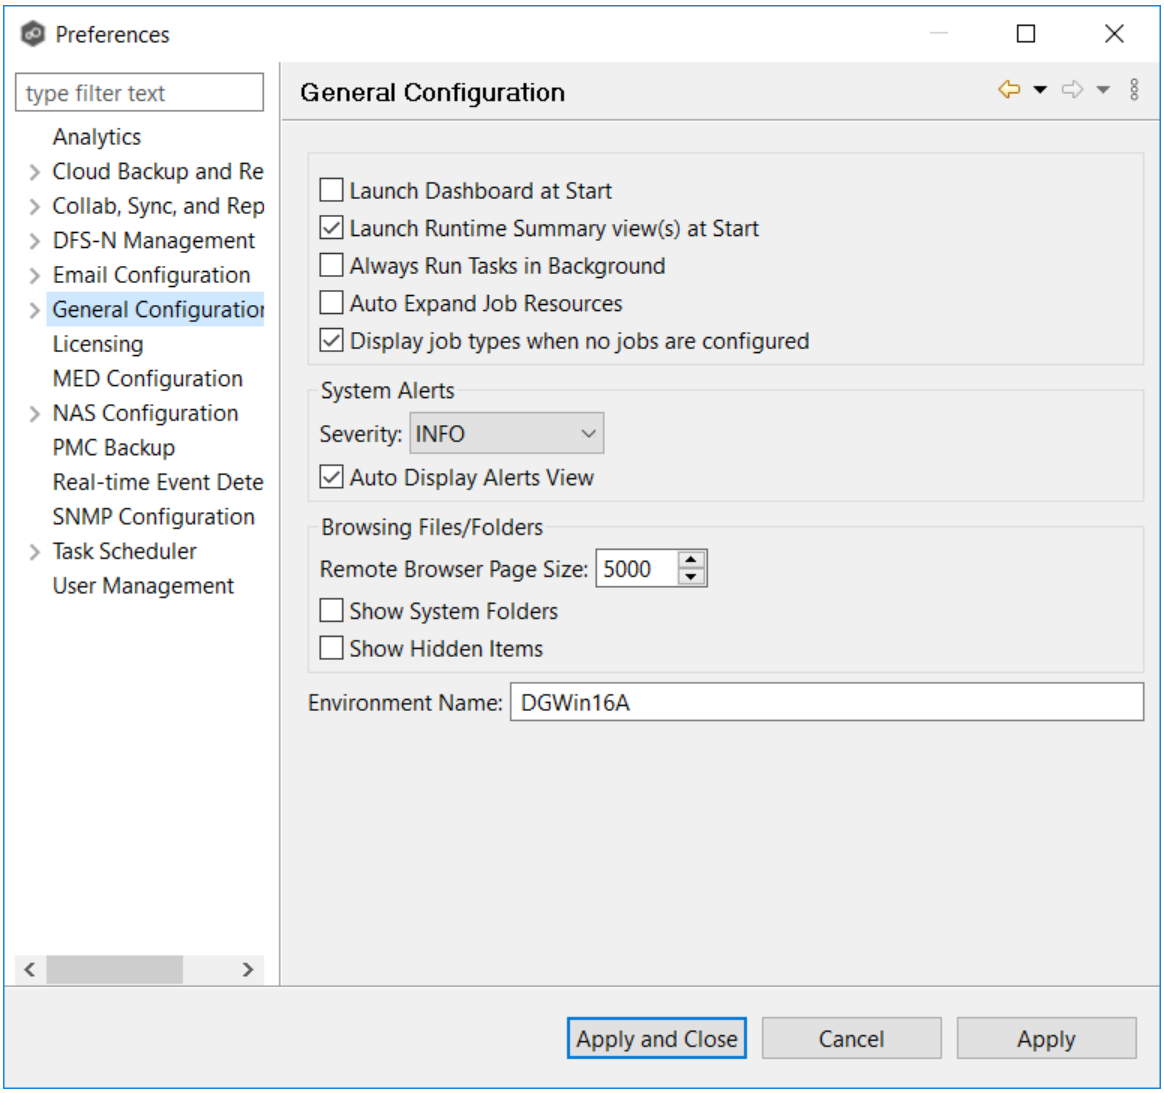

Configuring Global Settings

Peer Software strongly recommends configuring the following settings before creating any jobs:

- Email Configuration
- Contacts and Distribution Lists
- System Alerts

Modify other global settings as needed. You may want to consult with Peer Software Technical Support when modifying the other global settings.

Configuring Job Type Specific Settings

Job Type	Setting
Cloud Backup and Replication	• Email Alerts • File and Folder Filters • Proxy Configuration.
File Collaboration, File Replication, and File Synchronization	• Email Alerts • File and Folder Filters
DFS-N Management	• Email Alerts • File and Folder Filters

Configuring Preferences

To modify settings:

1. Click a category on the left to see its corresponding options appear on the right side of the dialog.

   For example, click the **General Configuration** category to view and configure general program-wide settings.

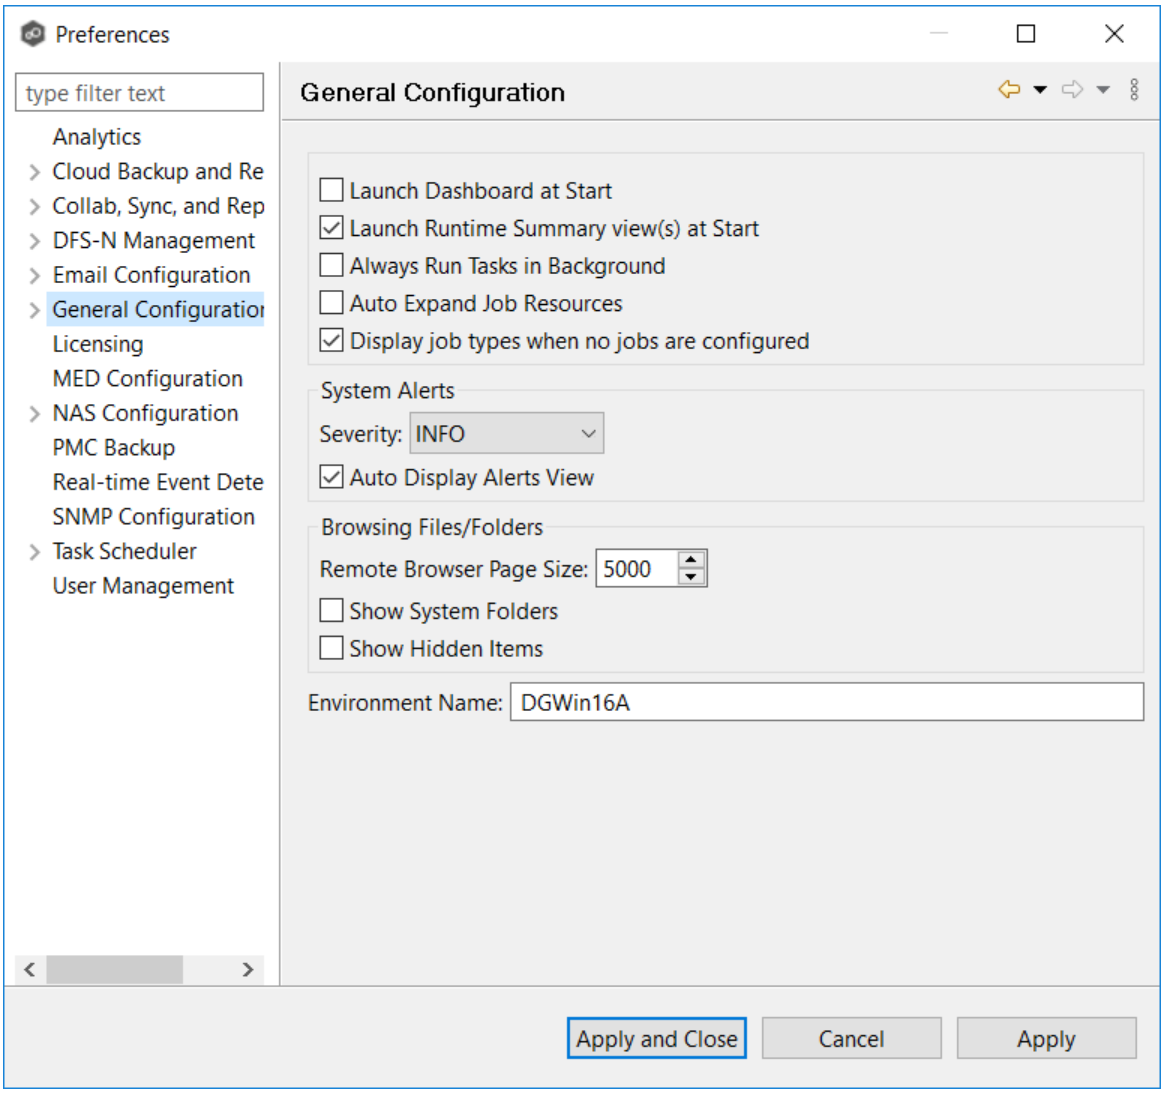

2. Make as many changes as you like to the category settings, and then click:

   - **Apply and Close** to save the new settings and return to the program.
   - **Cancel** to close the dialog without saving your changes.
   - **Apply** to save your changes and keep the **Preferences** dialog open.

Analytics Preferences

The **Analytics Preferences** page allows you to enable and disable the Analytics engine and modify Analytics settings. If you disable Analytics, you will need to rerun the Set Up Analytics wizard, which you can access by selecting **Analytics** from the **Help** menu.

For information about setting up Analytics, see Setting Up Analytics in Advanced Topics.

To modify Analytics settings:

1. Select **Preferences** from the **Window** menu.

2. Select **Analytics** in the navigation tree.

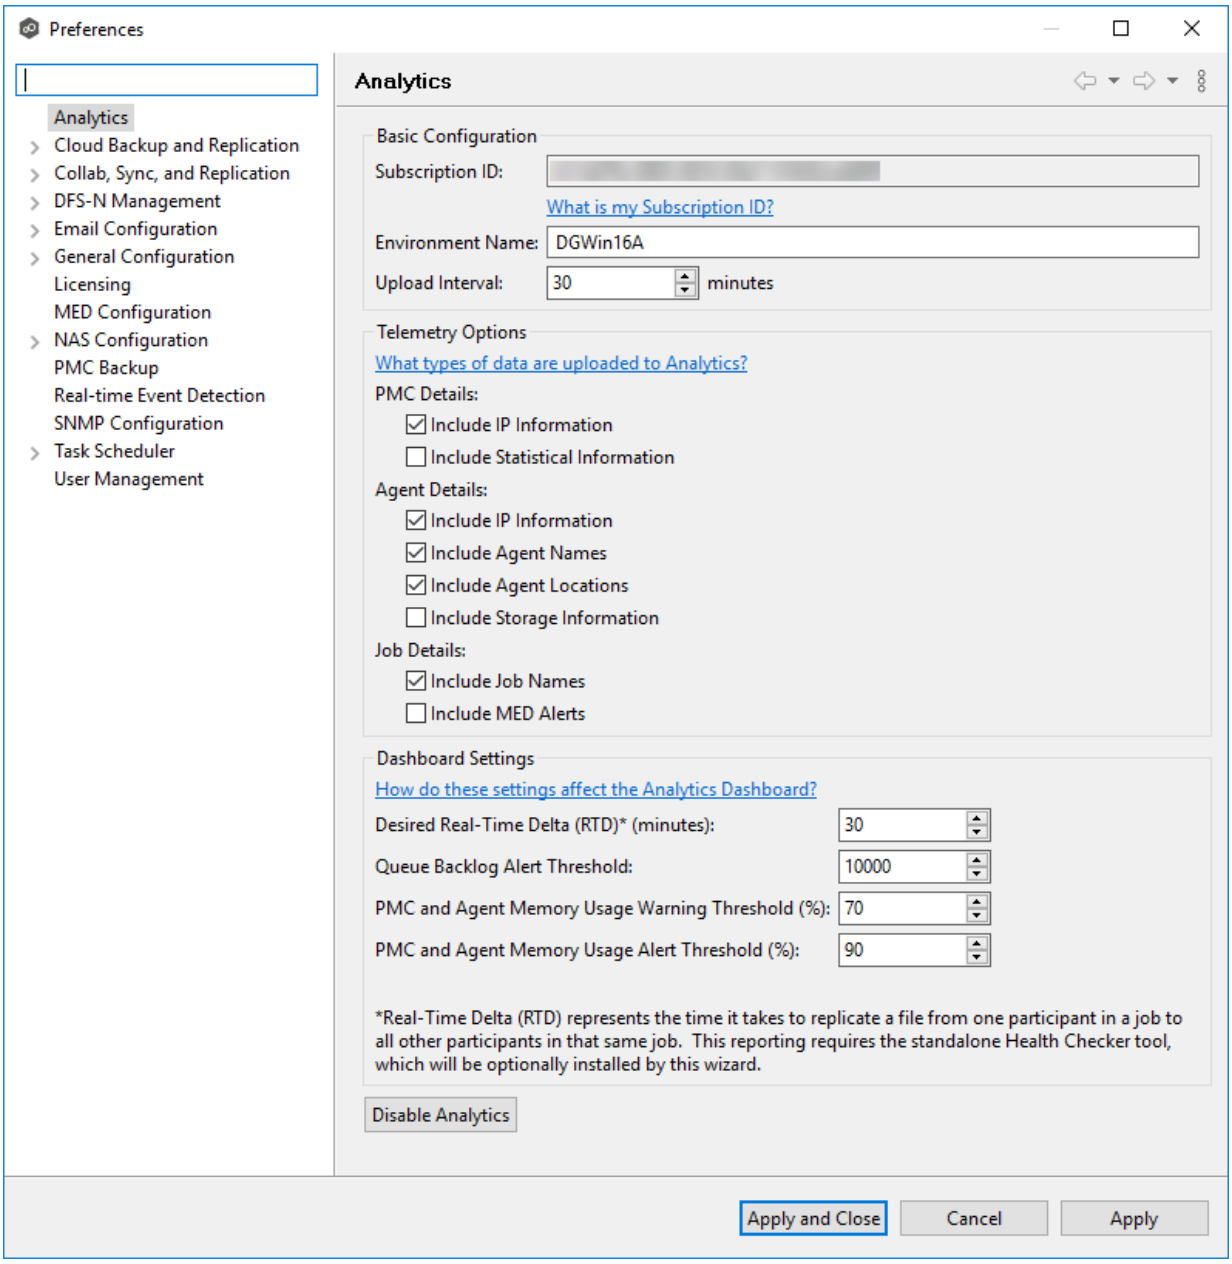

3. Select options in the **Basic Configuration** section.

Setting	Description
Subscription ID	Enter your subscription ID if the field is not auto-filled. Contact Peer Software if you do not know your Subscription ID. Once you enter a value, it cannot be changed.
Environment Name	In the **Environment Name** field, change the name if the auto-filled value doesn't match the name of the environment where the PMC server is installed. Note: Changing the name here will change the name in the General Configuration page in Preferences.
Upload Interval	In the **Upload Interval** field, enter the number of minutes to wait between uploads of data to the Analytics Dashboard. The default upload interval is 30 minutes. The minimum interval is 15 minutes; the maximum interval is 180 minutes

4. Select options in the **Telemetry Options** section.

Category	Standard Data Upload	Include Options
PMC Details	Includes details about this PMC deployment. Details include service memory consumption, replication backlog, quarantines, license consumption, and watch set size.	• **IP Information** - The IP address of the server that this PMC is installed on. • **Statistical information** - In-depth statistics about the queues and performance of PeerGFS's replication engine.
Agent Details	Includes the details about the Agents that are connected to this PMC. Details include service and server memory consumption, replication throughput, uptime, operating system, and disconnect counts.	• **IP information** - The IP addresses of the servers that the Agents are installed on. • **Agent Names** - The names assigned to the Agents (typically the name of each Agent's Windows Server). If this option is not checked, random strings will be used in the Analytics Dashboard to represent each Agent. • **Agent Locations** - The locations (the latitude, longitude, city, state, and country) of the Agents. This information is entered either through the Set Up Analytics wizard or in the Agent Configuration dialog. Agent locations can be displayed in a map showing replication traffic by geographic location. • **Storage Information** - Information specific to the storage platforms that each Agent is managing, including available and used disk space.
Job Details	Includes the details about the file collaboration, synchronization and/or replication jobs configured within this PMC. Per-job details include replication backlog, replication throughput, scan details, and watch set growth.	• **Job Names** - The names of the file collaboration, synchronization, and replication jobs configured in this PMC. If this option is not checked, random strings will be used in the Analytics Dashboard to represent each job. • **MED Alerts** - If MED alerts are enabled in this PMC, this option will include any alerts for use in the Analytics Dashboard.

5. Select options in the **Dashboard Settings** section.

Setting	Description
Desired Real-Time Delta (RTD)* (minutes)	Enter the maximum amount of time that is acceptable for a file change to replicate from one participant in a job to all other participants in that same job. Anytime a job is found to be taking more than the desired RTD, the Analytics Dashboard will display warnings. The default is 30 minutes. The minimum is 5; the maximum is 1440. Note: This only applies to real-time activity. It does not apply to scans or bulk activity.
Queue Backlog Alert Threshold	Enter the desired threshold for alerts about queue backlog. Once the queues inside the PMC hit this threshold, the Analytics Dashboard will display warnings. The default is 10000 items in the queue before an alert is sent. The minimum is 1000; the maximum is 1000000.
PMC and Agent Memory Usage Warning Threshold (%)	Enter the desired threshold for warnings about both Agent and PMC service memory usage. If either the PMC or any Agent hits this threshold, the Analytics Dashboard will display warnings. The default is 70 percent of memory usage. The minimum is 50 percent; the maximum is 95 percent.
PMC and Agent Memory Usage Alert Threshold (%)	Enter the desired threshold for alerts about both Agent and PMC service memory usage. If either the PMC or any Agent hits this threshold, the Analytics Dashboard will display errors. The default is 90 percent of memory usage. The minimum is 75 percent; the maximum is 100 percent.

6. Click **Apply and Close** or **Apply**.

Cloud Backup and Replication Job Preferences

You can modify the following Cloud Backup and Replication settings:

- Cloud Backup and Replication
- Database Connections

- Destination Credentials
- Email Alerts
- File and Folder Filters
- Performance
- Proxy Configuration
- File Retries and Source Snapshots
- Replication and Retention Policies
- SNMP Notifications
- Scan Manager

Cloud Backup and Replication

Cloud Backup and Replication settings control the overall performance of all Cloud Backup and Replication jobs.

To modify these settings:

1. Select **Preferences** from the **Window** menu.
2. Select **Cloud Backup and Replication** in the navigation tree.

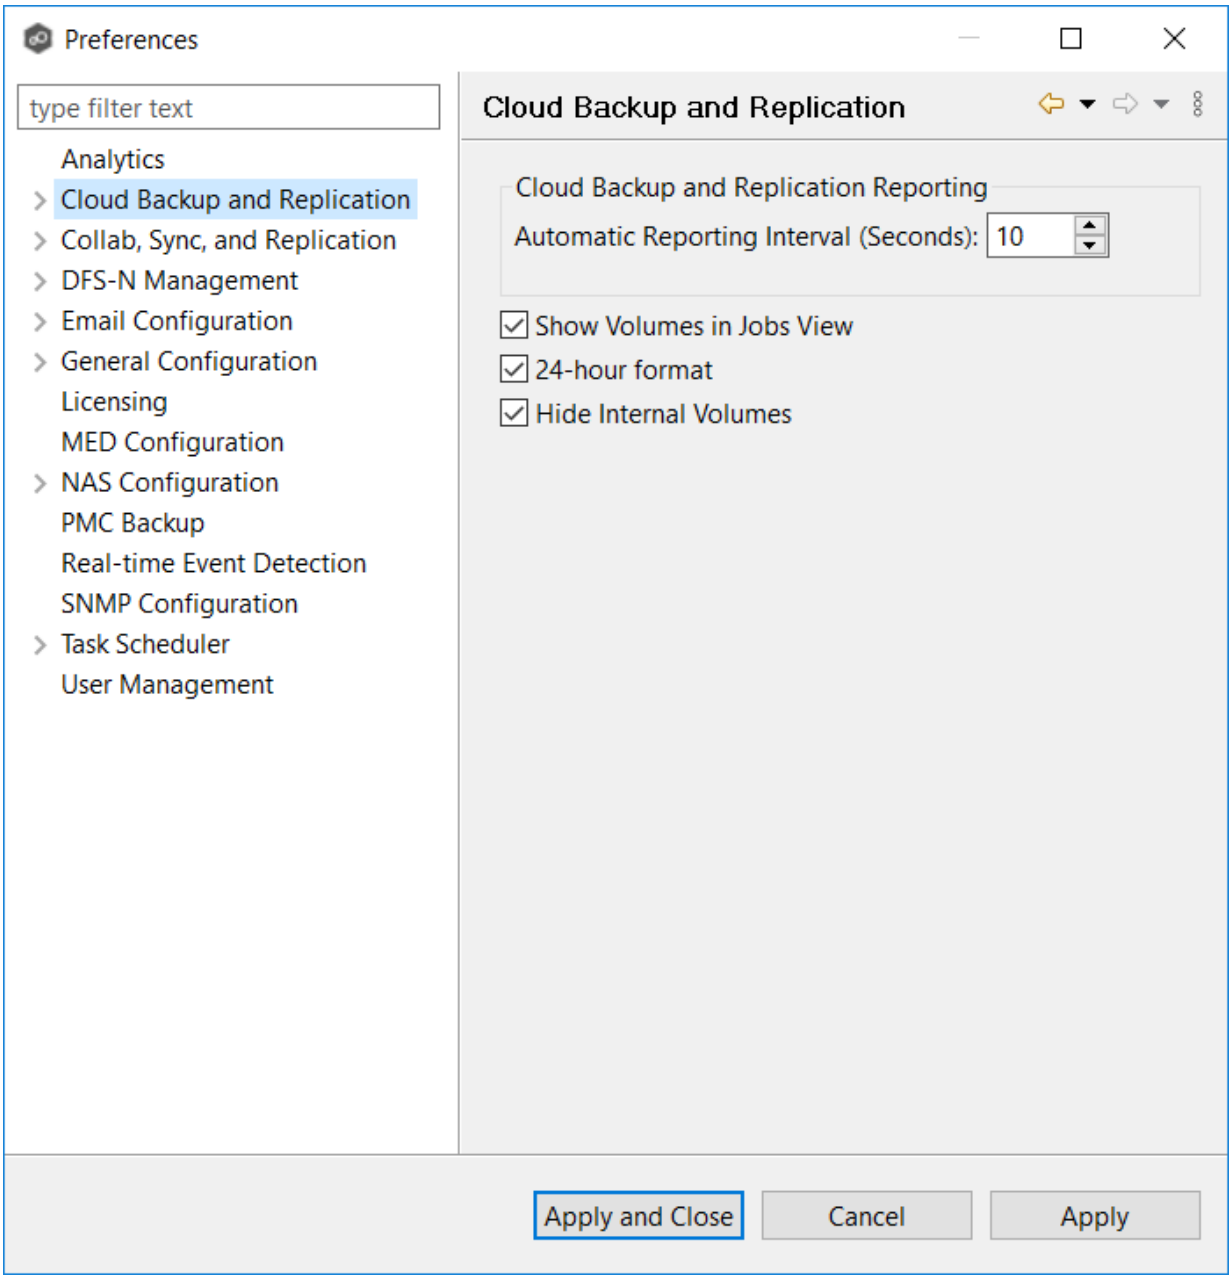

3. Modify the settings as needed.

Automatic Reporting Interval (Seconds)	Each Peer Agent automatically reports its statistics to Peer Management Center at regular intervals. Select the number of seconds between these intervals. The default is 10 seconds.

Show Volumes in Jobs View	Select this checkbox if you want volumes to be displayed in the Jobs view.
24-hour format	Select this checkbox if you want times to be displayed in a 24-hour format rather than a 12-hour format.
Hide Internal Volumes	Select this checkbox if you don't want internal volumes displayed when choosing which volumes to replicate.

4. Click **Apply and Close** or **Apply**.

Database Connections

Cloud Backup and Replication uses a Microsoft SQL Server or SQL Server Express database to track files and folders that have been replicated, individual file versions, and snapshots. When creating a Cloud Backup and Replication job, the Management Agent that you select for the job must have a connection to your SQL Server. You can set up the connection in advance on this page; otherwise, you will be prompted to set up the connection when you create a job.

You cannot modify or delete a database connection while a job using the connection is run.

To create a new database connection:

1. Select **Preferences** from the **Window** menu.

2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Database Connections**.

   Any existing database connections are listed in the **Database Connections** table.

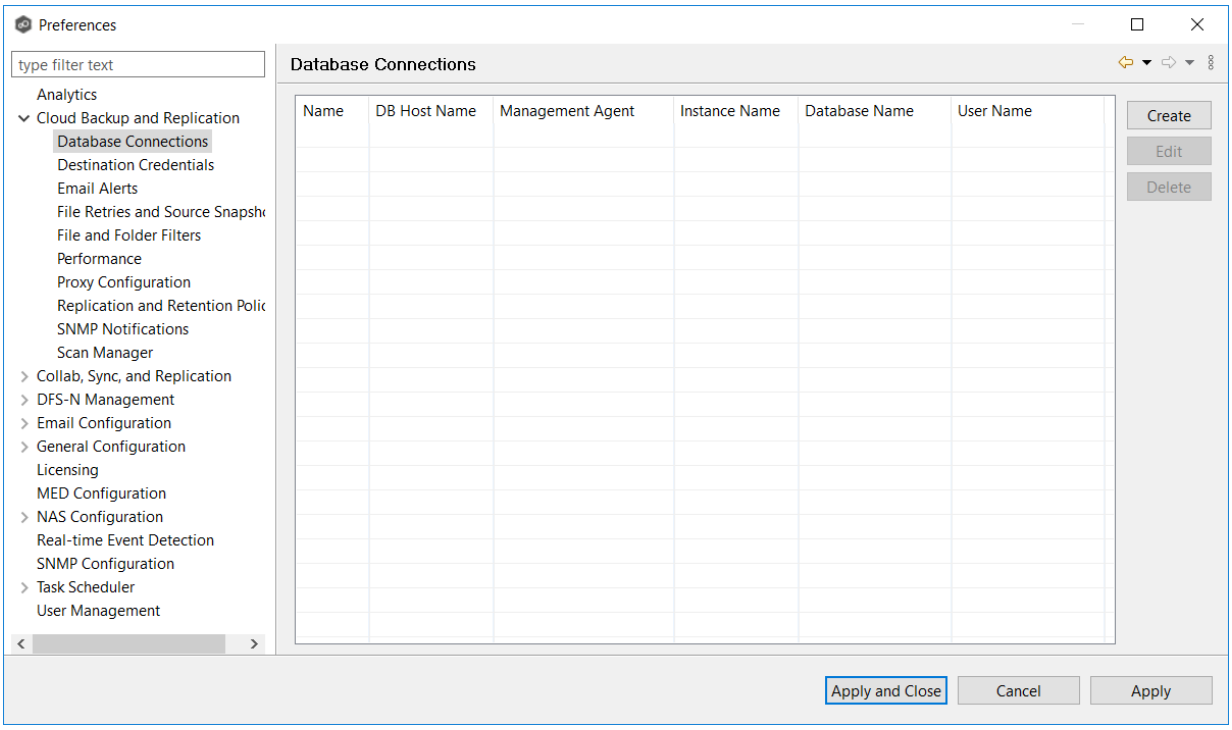

3. Click the **Create** button.

   The **Create Database Connection** dialog appears.

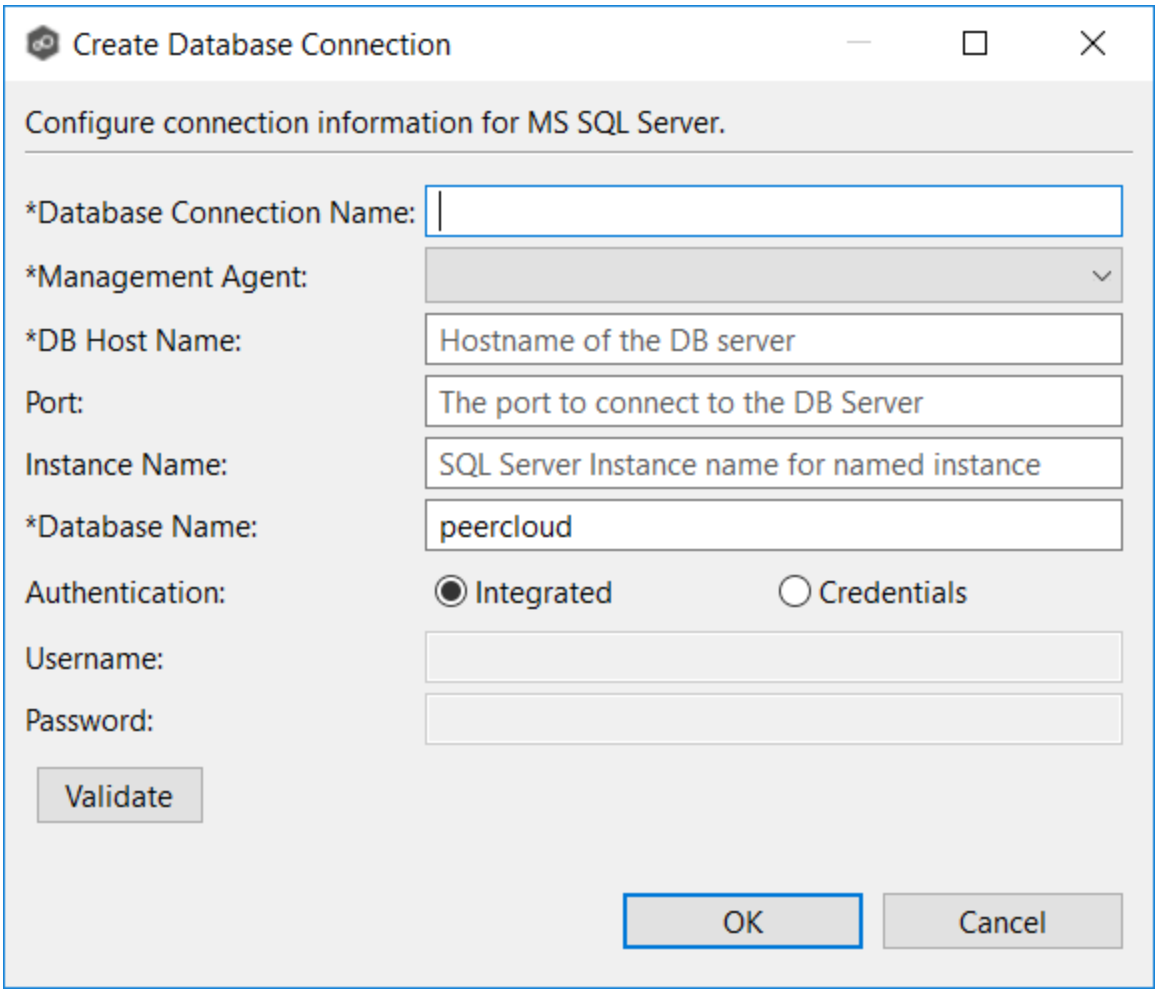

4. Enter the required values.

Field	Description
Database Connection Name	Enter a name for this database connection.
Management Agent	Select the Management Agent that will use this connection. The Agent must be the same one as managing the job.
DB Host Name	Enter the name of the SQL Server hosting the database. If the database is installed on the Agent server itself, enter the name of the Agent server.

Field	Description
Port	Optional. Enter the port to be used to communicate with the specified SQL Server. If not defined, the connection defaults to port 1433.
Instance Name	Optional. Enter the database instance name to use on the specified SQL Server. If no named instances are installed on the specified SQL Server, leave this blank.
Database Name	Enter the name of the database that Cloud Backup and Replication will create. The default name is *peercloud*, but it can be changed to a name that follows your company's naming conventions.
Authentication	Select **Integrated** if the Agent service account is granted admin rights on the selected SQL instance. Otherwise, select **Credentials** to enter the user name and password of a database administrator.
Username	Required when **Credentials** is selected for **Authentication**. Enter the user name of an account to be used to connect to the database. This can be a locally defined account such as "sa" or a domain account. The account must have adequate privileges to manage the database, such as database owner.
Password	Required when **Credentials** is selected for **Authentication**. Enter the password for account being used to connect to the database.

5. Click **Validate** to test the connection, and then click **OK** in the confirmation message that appears.

6. Click **OK** to close the dialog.

   The new database connection is listed in the **Database Connections** table.

7. Click **Apply and Close** or **Apply**.

Destination Credentials

When you create a Cloud Backup and Replication job, you can select existing destination storage account credentials to apply to the job or you can create new credentials and apply them to the

job. This Preferences page lists the existing credentials. From this page, you can view, create, edit, and delete credentials. However, you cannot edit or delete credentials while they are applied to a job.

To create new destination storage account credentials:

1. Select **Preferences** from the **Window** menu.

2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Destination Credentials**.

   Any existing credentials are listed in the **Destination Credentials** table.

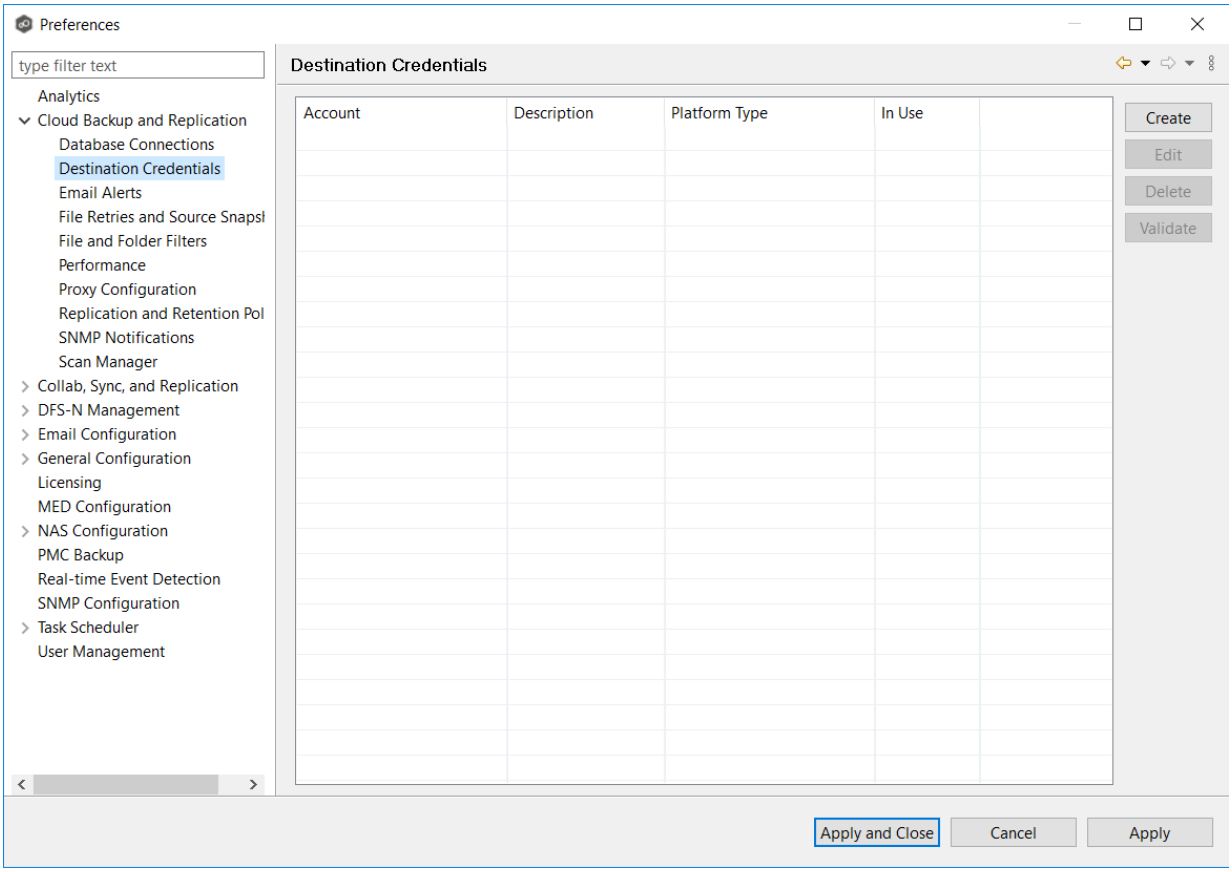

3. Click the **Create** button.

   The **Storage Account** dialog appears.

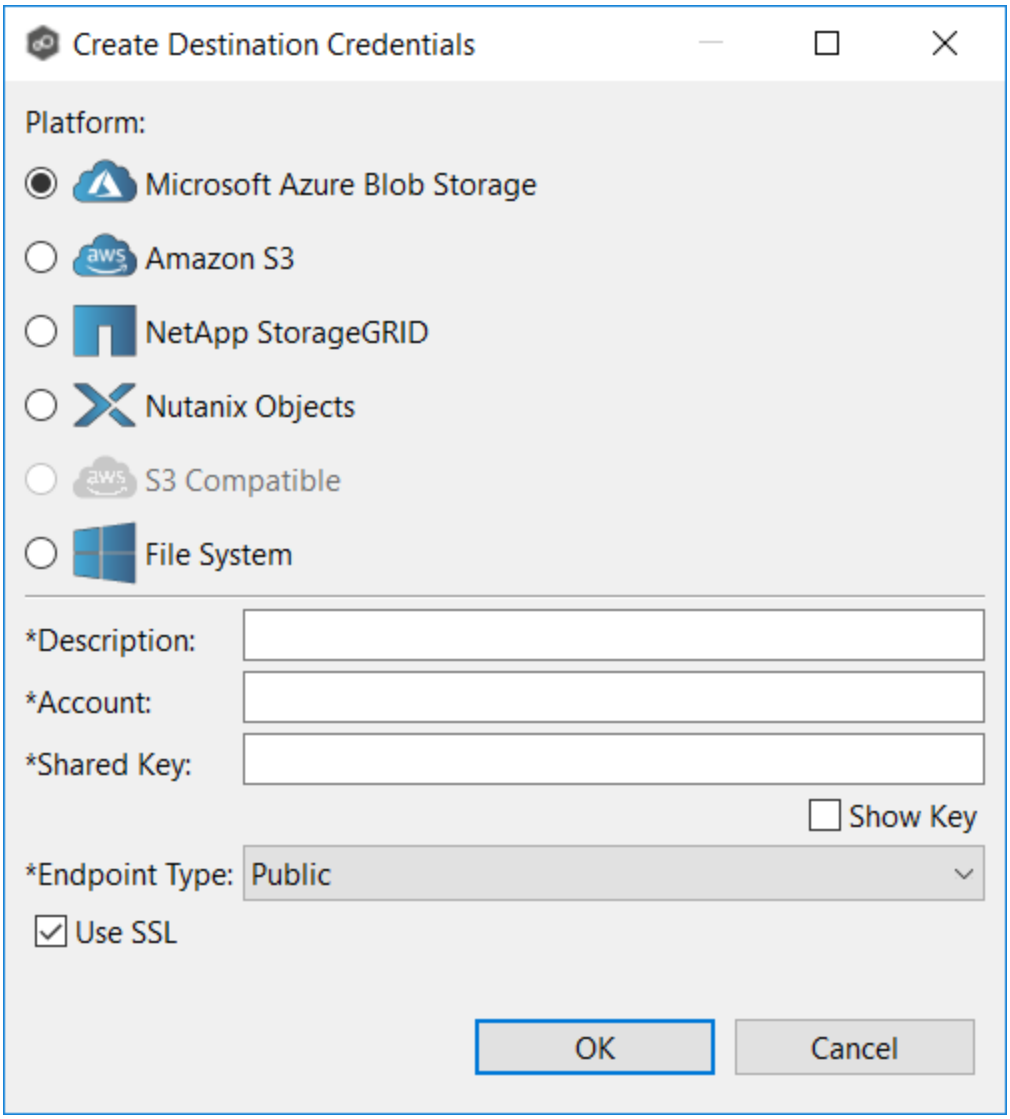

4. Enter the required values. For information about the required values, see Step 8: Destination Credentials in Creating a Cloud Backup and Replication Job.

5. Click **OK**.

   The new credential is listed in the **Destination Credentials** table and can now be applied to jobs.

6. Click **Apply and Close** or **Apply**.

Email Alerts

When you create a job, you can select existing email alerts to apply to the job or you can create new email alerts and apply them to the job. This Preferences page lists the existing email alerts. From this page, you can view, create, edit, and delete email alerts. However, you cannot edit or delete an email alert while it is applied to a job. See Email Alerts in the Basic Concepts section for more information about email alerts.

Note: An SMTP email connection must be configured before email alerts can be sent. See Email Configuration for information about configuring SMTP email settings.

To create an email alert:

1. Select **Preferences** from the **Window** menu.

2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Email Alerts**.

   Any existing Cloud Backup and Replication email alerts are listed in the **Email Alerts** table.

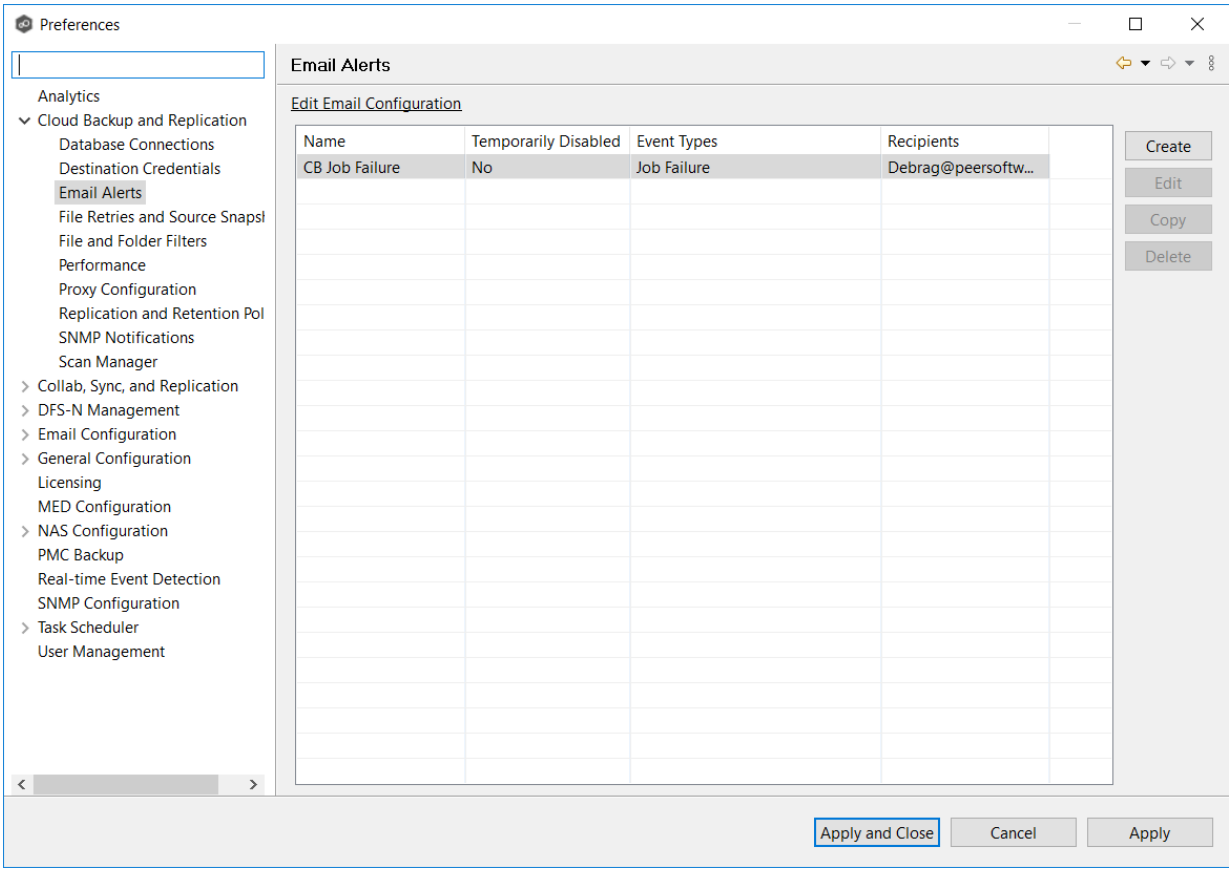

3. Click the **Create** button.

The **Create Email Alert** dialog appears.

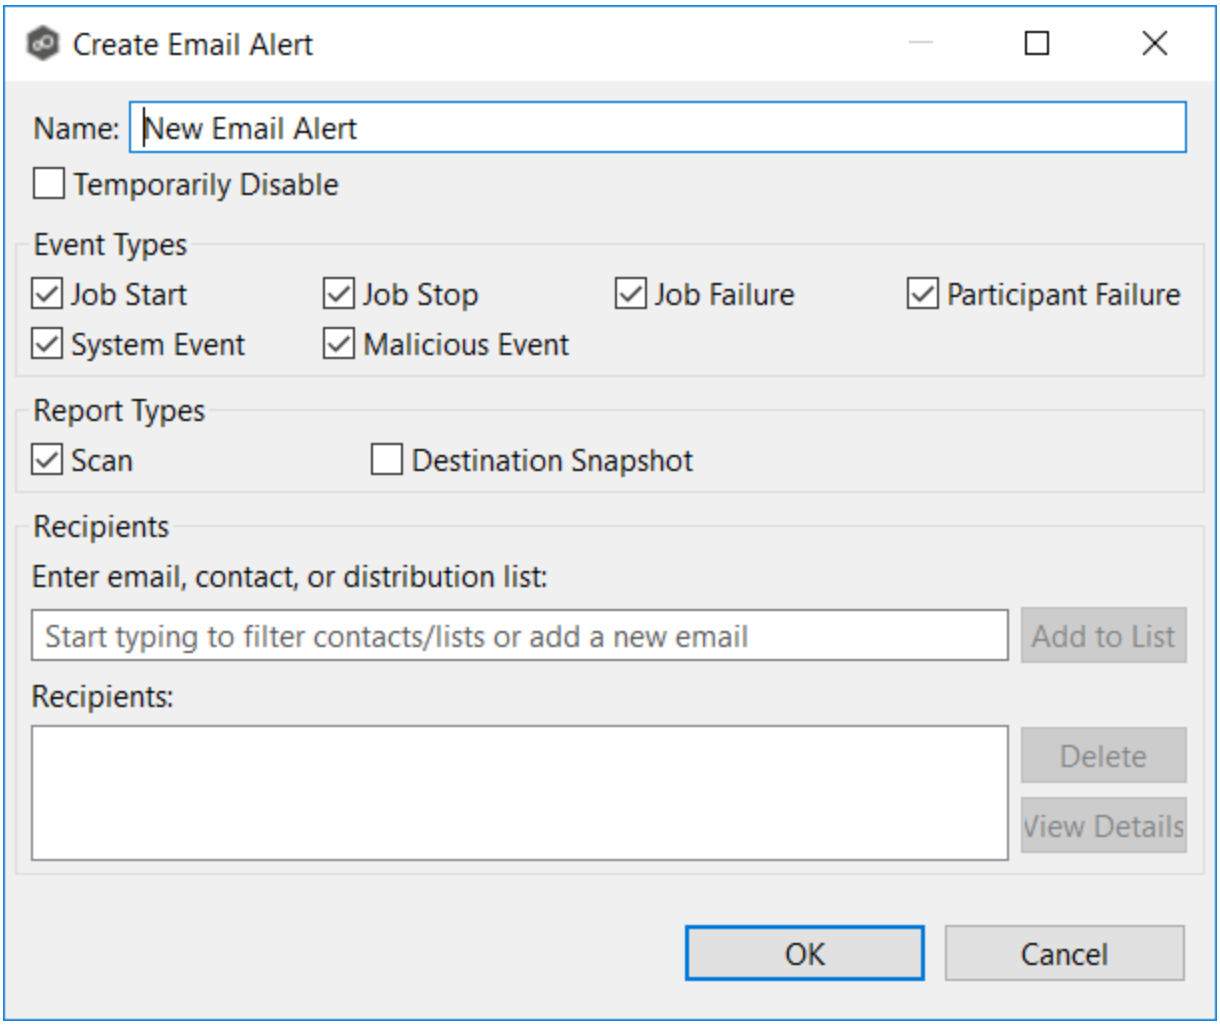

4. Enter a name for the alert.
5. Select the event types to be alerted.

   The event type determines what will trigger the email alert to be sent.

Event Type	Description
Job Start	Sends an alert when the job starts.

Event Type	Description
Job Stop	Sends an alert when the job stops.
Job Failure	Sends a notification when job stops unexpectedly.
Participant Failure	Sends an alert when the Management Agent disconnects or stops responding.
System Event	Sends an alert when a system event such as low memory or low hub disk space occurs.
Malicious Event	Sends an alert when Peer MED detects potentially malicious activity. For more information, see MED Configuration.

6. Select the report types to be sent.

Report Type	Description.
Scan	Sends scan statistics after a scan has completed.
Destination Snapshot	Sends information about the snapshot after the snapshot is taken.

7. Enter alert recipients, and then click **Add to List**.

   The recipients are listed in the **Recipients** field.

8. Click **OK**.

   The new email alert is listed in the **Email Alerts** table and can now be applied to jobs.

9. Click **Apply and Close** or **Apply**.

File Retries and Source Snapshots

This page allows you to specify two sets of options:

- **File Retries** - Settings that are used when retry issues that arise while replicating a file or folder.
- **Source Snapshot Replication** - Settings that control how and when source snapshots are used.

To modify these options:

1. Select **Preferences** from the **Window** menu.
2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **File Retries and Source Snapshots**.

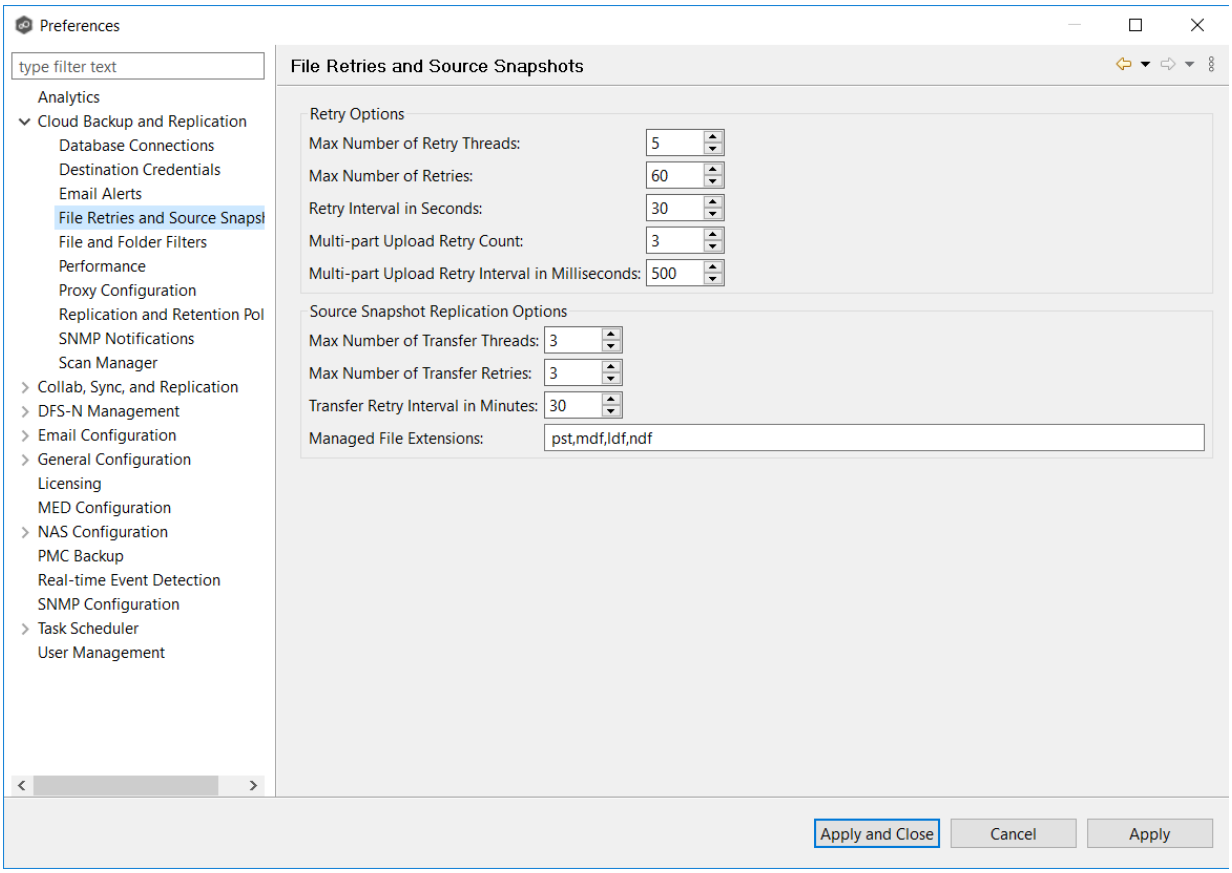

3. Modify the **Retry Options** as needed:

Option	Description
Max Number of Retry Threads	Enter the maximum number of threads available for handling retries of failed file or folder transfers.
Max Number of Retries	Enter the maximum number of retries to perform on a file or folder that has failed to be replicated. If the number of retries is exceeded, the file or folder will be added to the Failed Events view and will need to be manually processed.
Retry Interval in seconds	Enter the number of seconds to wait in between retries of the failed replication of a file or folder.

Option	Description
Multi-part Upload Retry Count	Enter the maximum number of retries when performing multi-part upload.
Multi-part Upload Retry Interval in Milliseconds	Enter the number of minutes to wait between retries of multi-part uploads.

4. Modify the **Source Snapshot Replication Options** as needed:

Option	Description
Max Number of Transfer Threads	Enter the maximum number of threads available for replicating files from a source snapshot.
Max Number of Transfer Retries	Enter the maximum number of retries to perform on a file or folder that has failed to be replicated from a source snapshot. If the number of retries is exceeded, the file or folder will be added to the Failed Events view and will need to be manually processed.
Transfer Retry Interval in Minutes	Enter the number of minutes to wait between retries of the failed replication of a file or folder from a source snapshot.
Managed File Extensions	Enter the extensions for managed files that should be read from a source snapshot.

5. Click **Apply and Close** or **Apply**.

File and Folder Filters

When you create a Cloud Backup and Replication job, you can select existing file filters to apply to the job or you can create new file filters and apply them to the job. This Preferences page lists the existing file filters. From this page, you can view, create, edit, and delete file filters. However, you cannot edit or delete a file filter while it is applied to a job. See File and Folder Filters in the Basic Concepts section for more information about file and folder filters.

To create a file filter:

1. Select **Preferences** from the **Window** menu.

2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **File and Folder Filters**.

   Any existing Cloud Backup and Replication file filters are listed in the **File Filters** table.

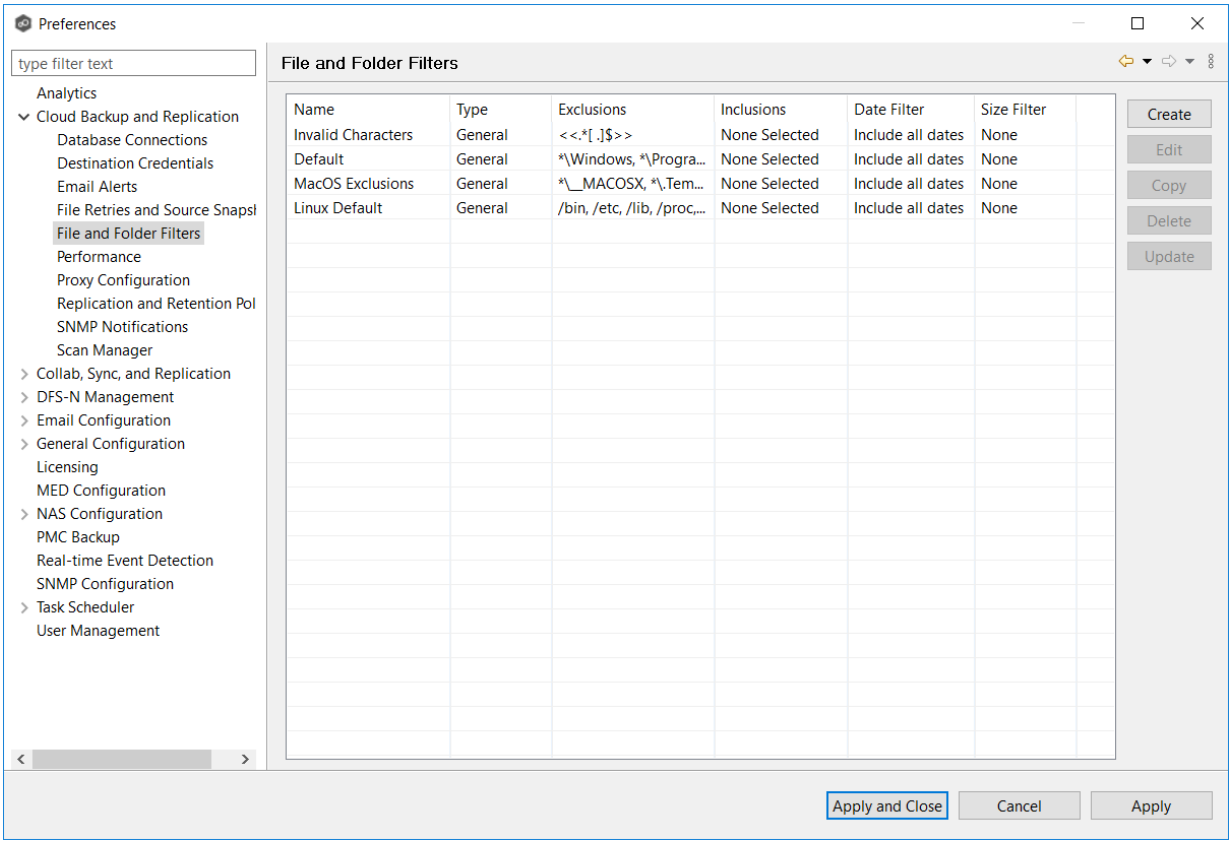

3. Click the **Create** button.

   The **Create File Filter** dialog appears.

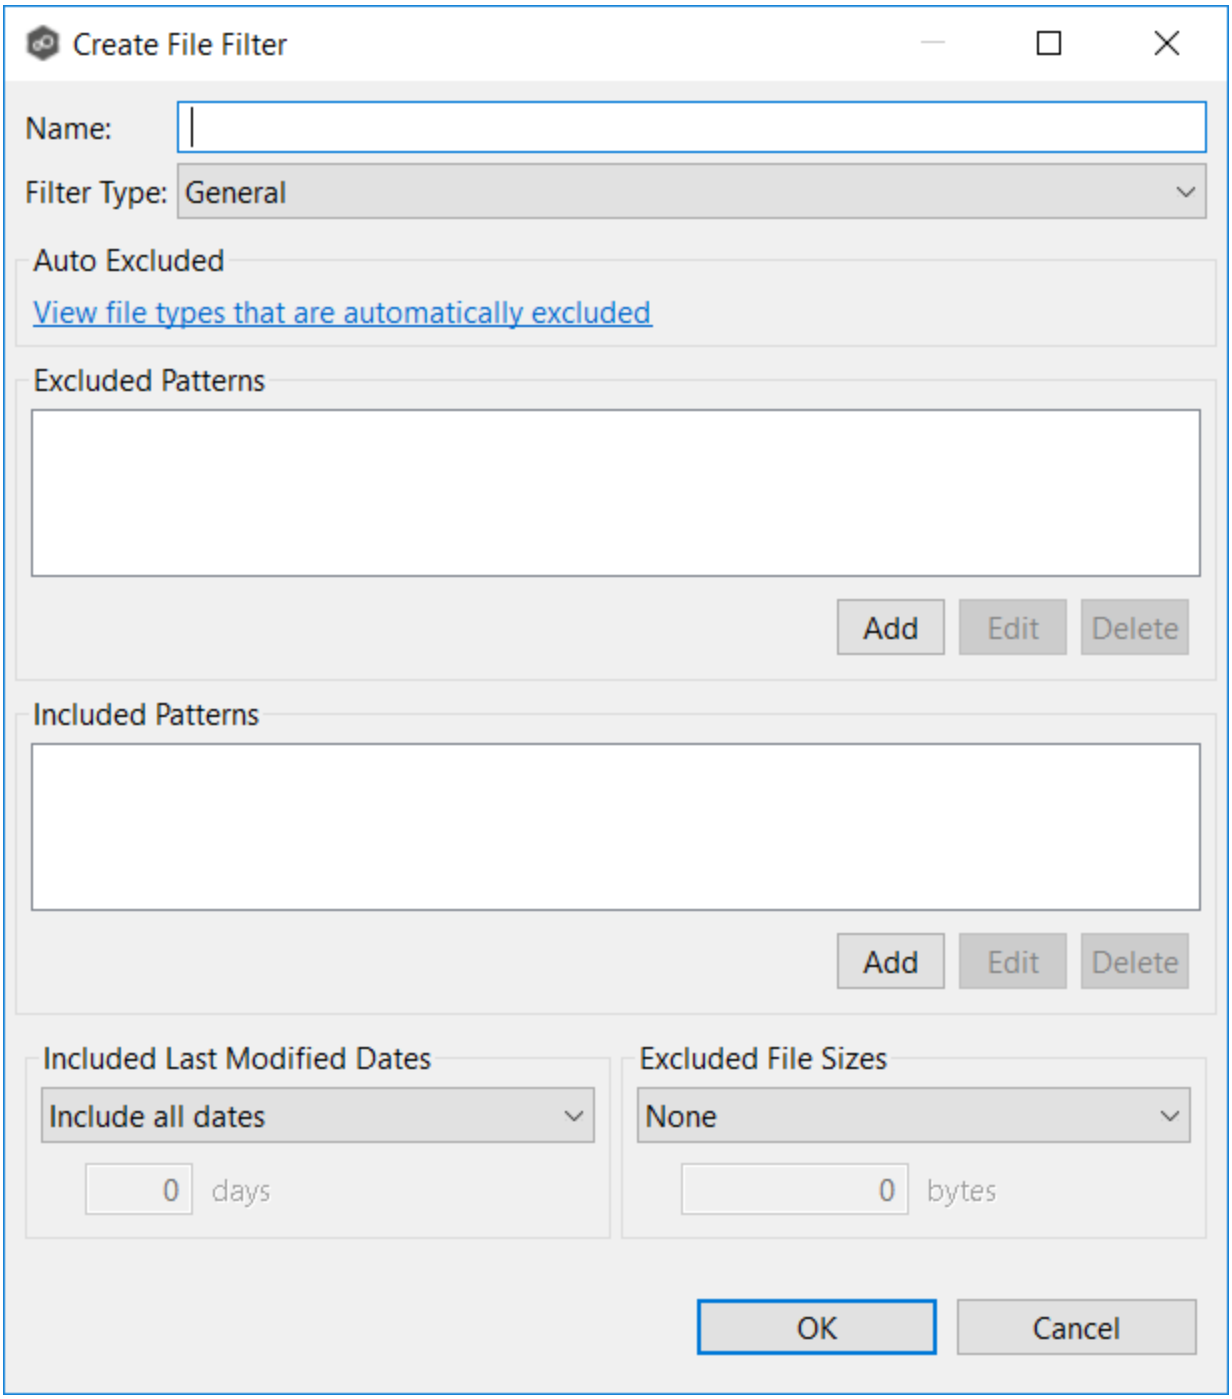

4. Enter a unique name for the filter.

5. Select the filter type.

6. (Optional) In the **Excluded Patterns** section, click the **Add** button to enter a filter pattern for files that you want excluded from the job. Repeat to add more filter patterns.

   See Defining Filter Patterns for information about filters patterns.

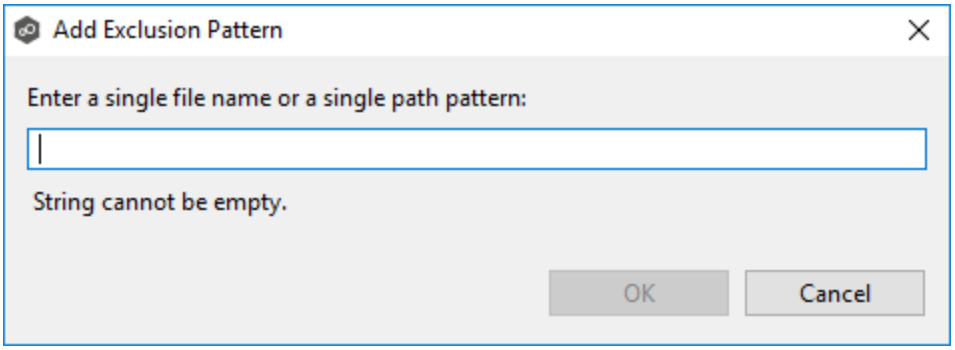

7. (Optional) In the **Included Patterns** section, click the **Add** button to enter a filter pattern for files that you want included in the job. Repeat to add more filter patterns.

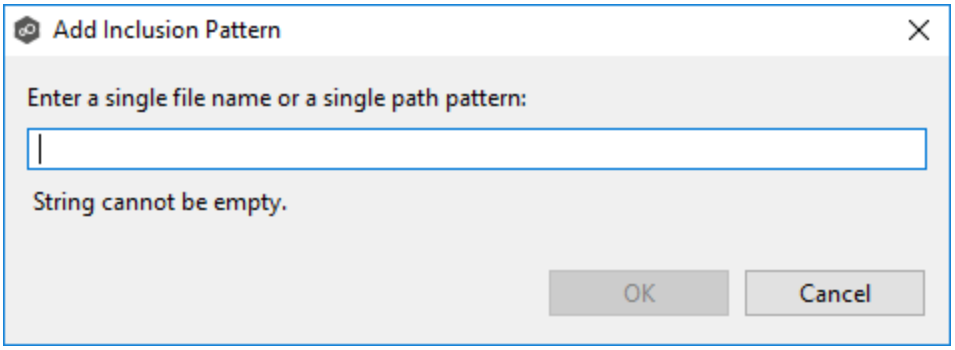

8. (Optional) Select a value for Included Last Modified Dates.

   Note: A filter cannot use **Included Last Modified Dates** in conjunction with any other excluded or included patterns.

9. (Optional) Select a value for Excluded File Sizes.

   Note: A filter cannot use **Excluded File Sizes** in conjunction with any other excluded or included patterns.

10. Click **OK**.

    The new file filter is listed in the **File Filters** table and can now be applied to jobs.

11. Click **Apply and Close** or **Apply**.

Performance

Performance settings allow you to adjust the performance of Cloud Backup and Replication jobs.

To modify the Cloud Backup and Replication performance settings:

1. Select **Preferences** from the **Window** menu.
2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Performance**.

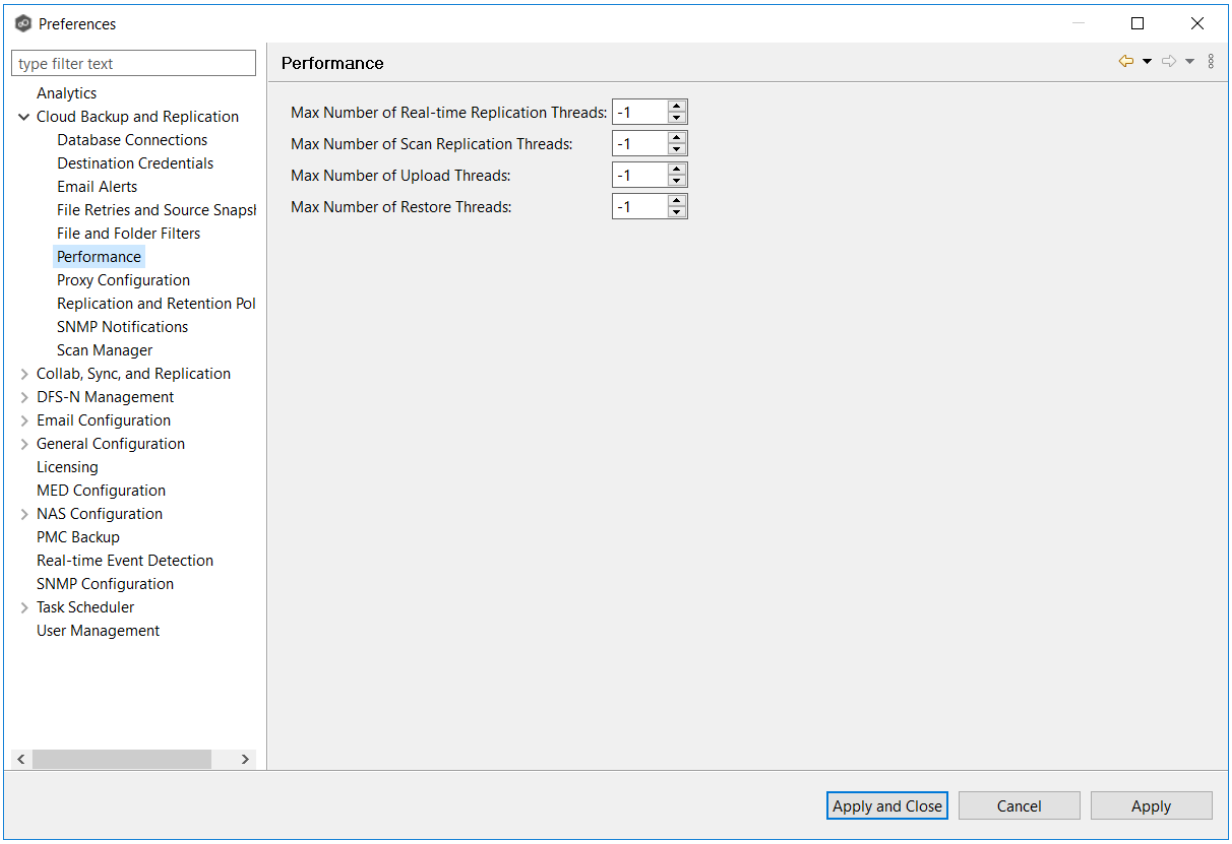

3. Modify the settings as needed:

Setting	Description
Max Number of Real-time Replication Threads	Enter the maximum number of threads available for replicating files as they are updated in real-time on the source storage device.

Setting	Description
Max Number of Scan Replication Threads	Enter the maximum number of threads available for replicating files during scheduled and on-demand scans of the source storage device.
Max Number of Upload Threads	Enter the maximum number of threads available for uploading files to the destination storage device.
Max Number of Restore Threads	Enter the maximum number of threads available for restoring from the destination storage device.

4. Click **Apply and Close** or **Apply**.

Proxy Configuration

The **Proxy Configuration** page allows you to create a proxy to be used with Microsoft Azure Blob Storage, Amazon S3, and S3 Compatible storage accounts.

To create a proxy configuration:

1. Select **Preferences** from the **Window** menu.

2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Proxy Configuration**.

   Any existing proxies are listed in the **Proxy Configuration** table.

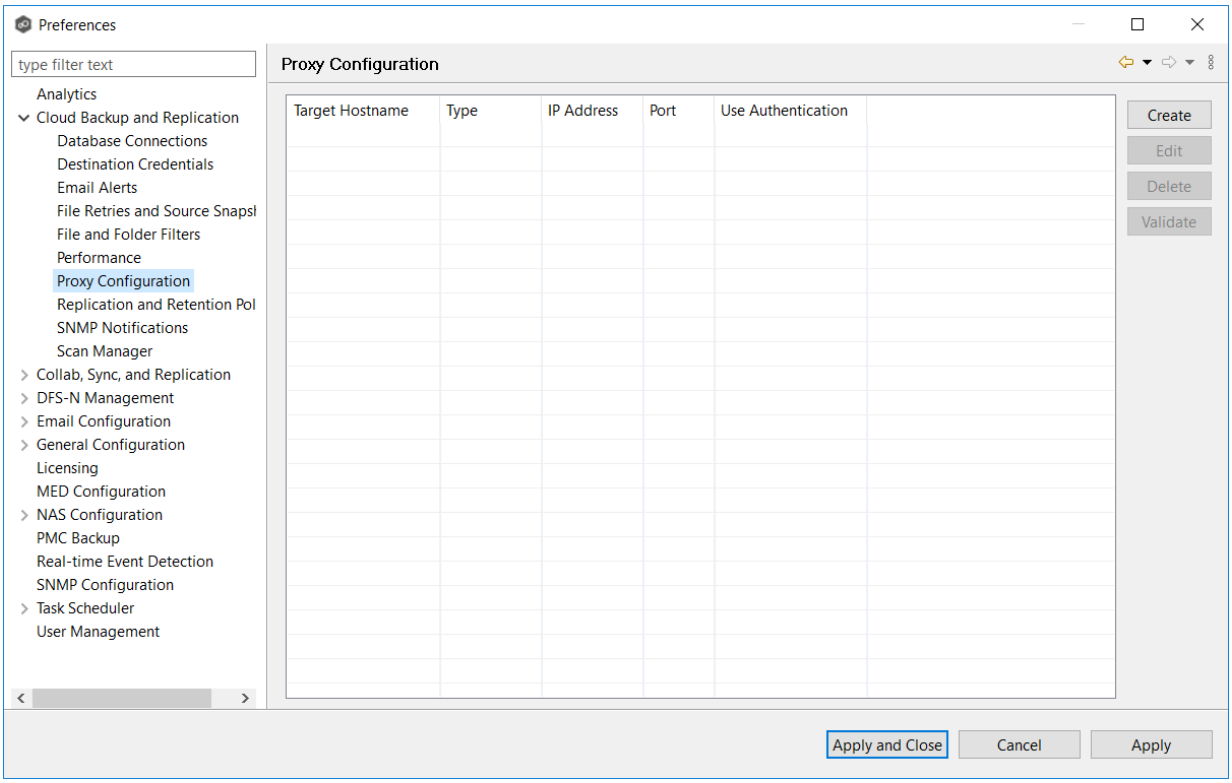

3. Click the **Create** button.

   The **Create Proxy Configuration** dialog appears.

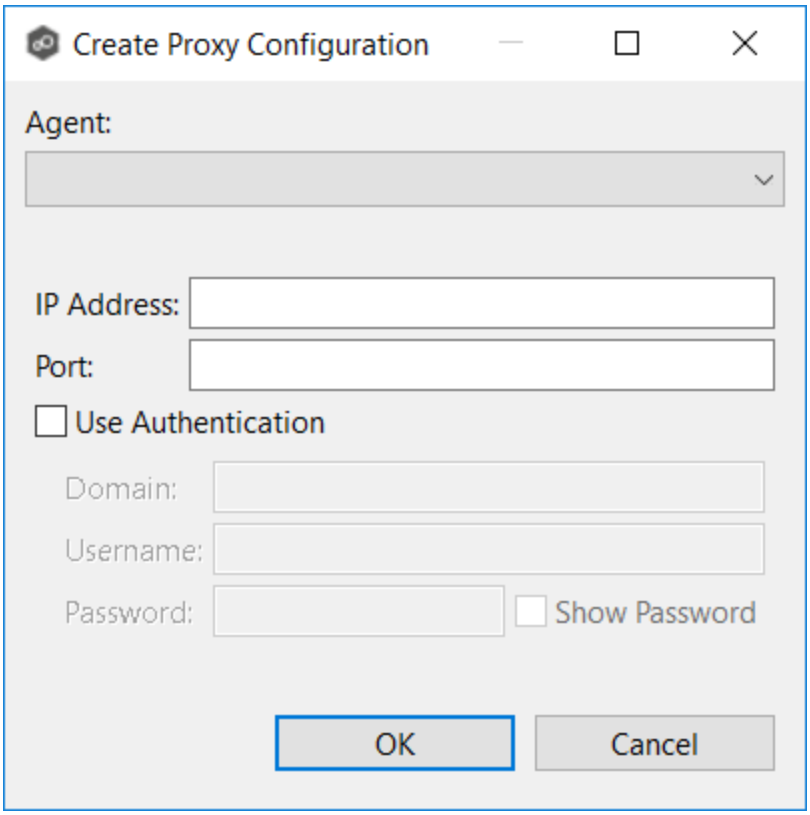

4. Select the Agent that that manages your storage device.

5. Enter values for the following fields:

Field	Description
Address	Enter the IP address or fully qualified domain name of the proxy server.
Port	Enter the port number.
User Authentication	Select if your proxy server requires authentication. This option does not apply for proxy servers connecting to an Azure storage device

6. If your proxy server requires authentication, select the **Use Authentication** checkbox, and then supply the necessary values:

Field	Description
Domain	Enter the domain name on the proxy server.
Username	Enter the user name for the proxy server.
Password	Enter the password for the proxy server.

7. Click **OK**.

   The new proxy configuration is listed in the **Proxy Configuration** table.

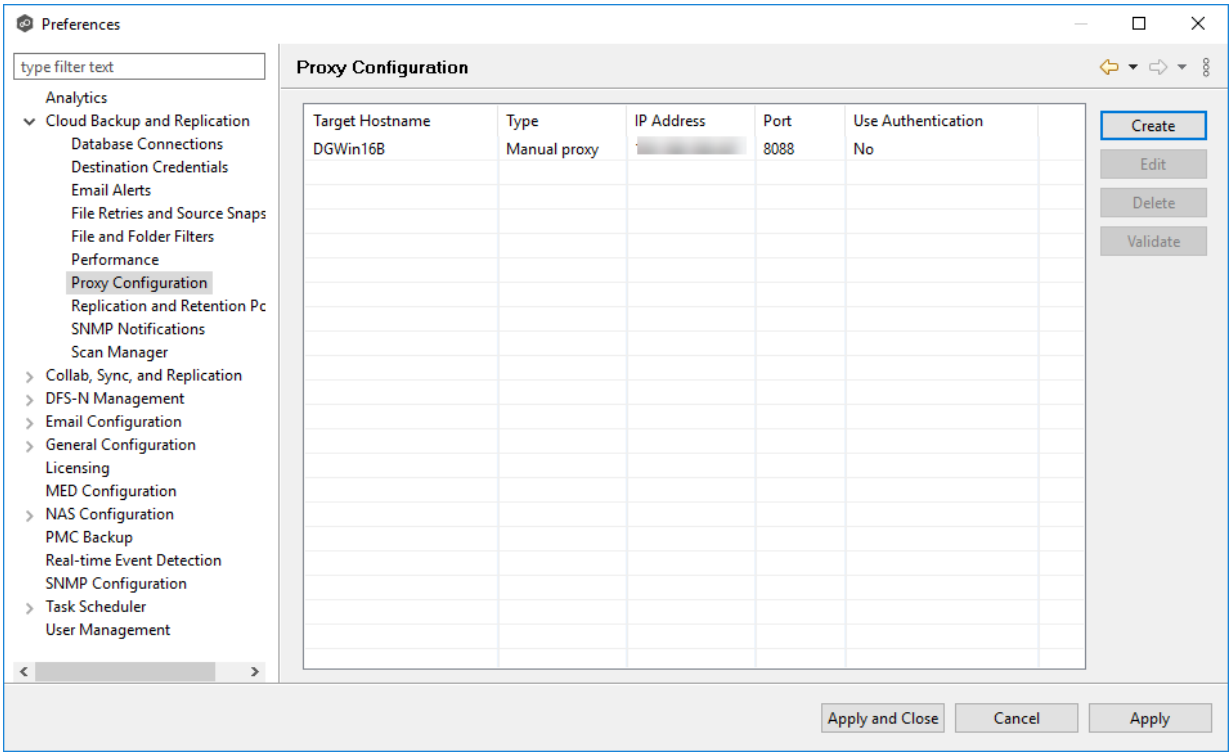

8. Click **Apply and Close** or **Apply**.

Replication and Retention Policies

Each Cloud Backup and Replication job must have a Replication and Retention Policy applied to it. When you create a job, you can select an existing policy to apply to the job or you can create a new policy and apply it to the job.

You can modify and delete a policy, however, you cannot:

- Modify a policy while it is applied to a running job.
- Delete a policy while it is applied to any job.

For more information about policies, see Step 10: Replication and Retention Policy.

To create a new replication and retention policy:

1. Select **Preferences** from the **Window** menu.
2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Replication and Retention Policies**.

   Any existing policies are listed in the table.

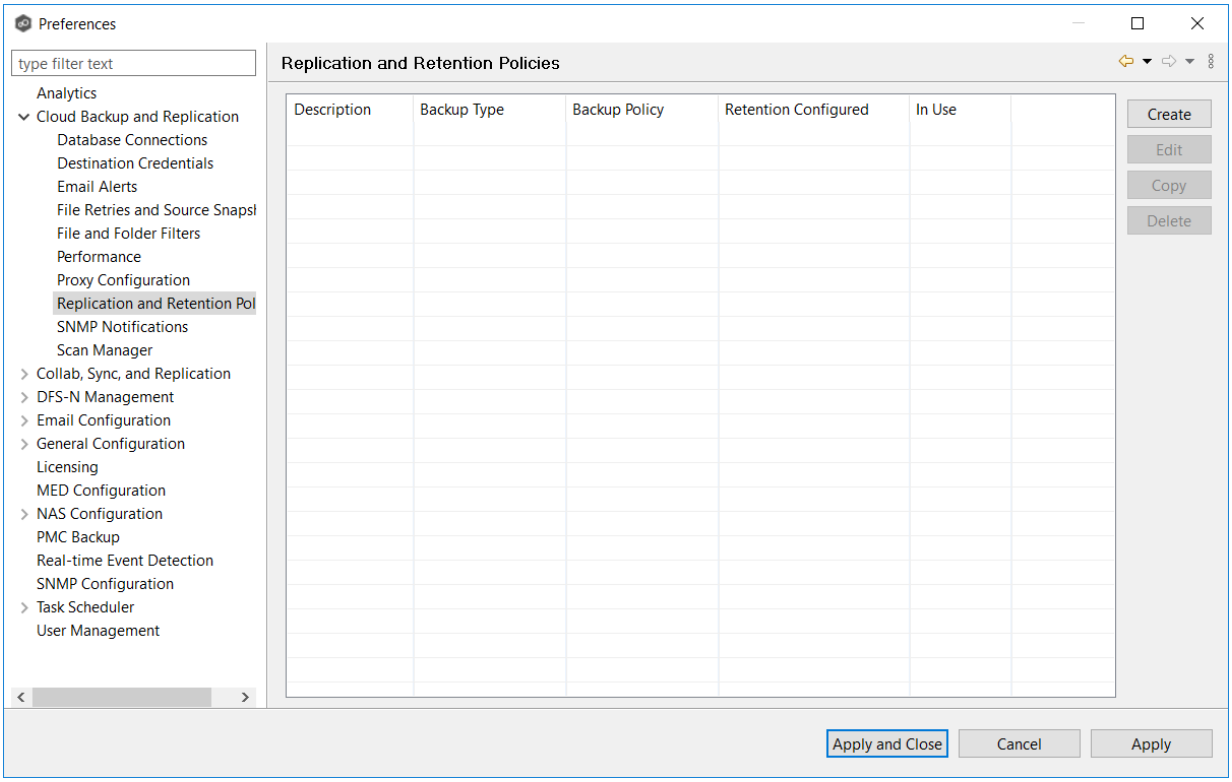

3. Click the **Create** button.

The **Replication and Retention Policy Wizard** opens.

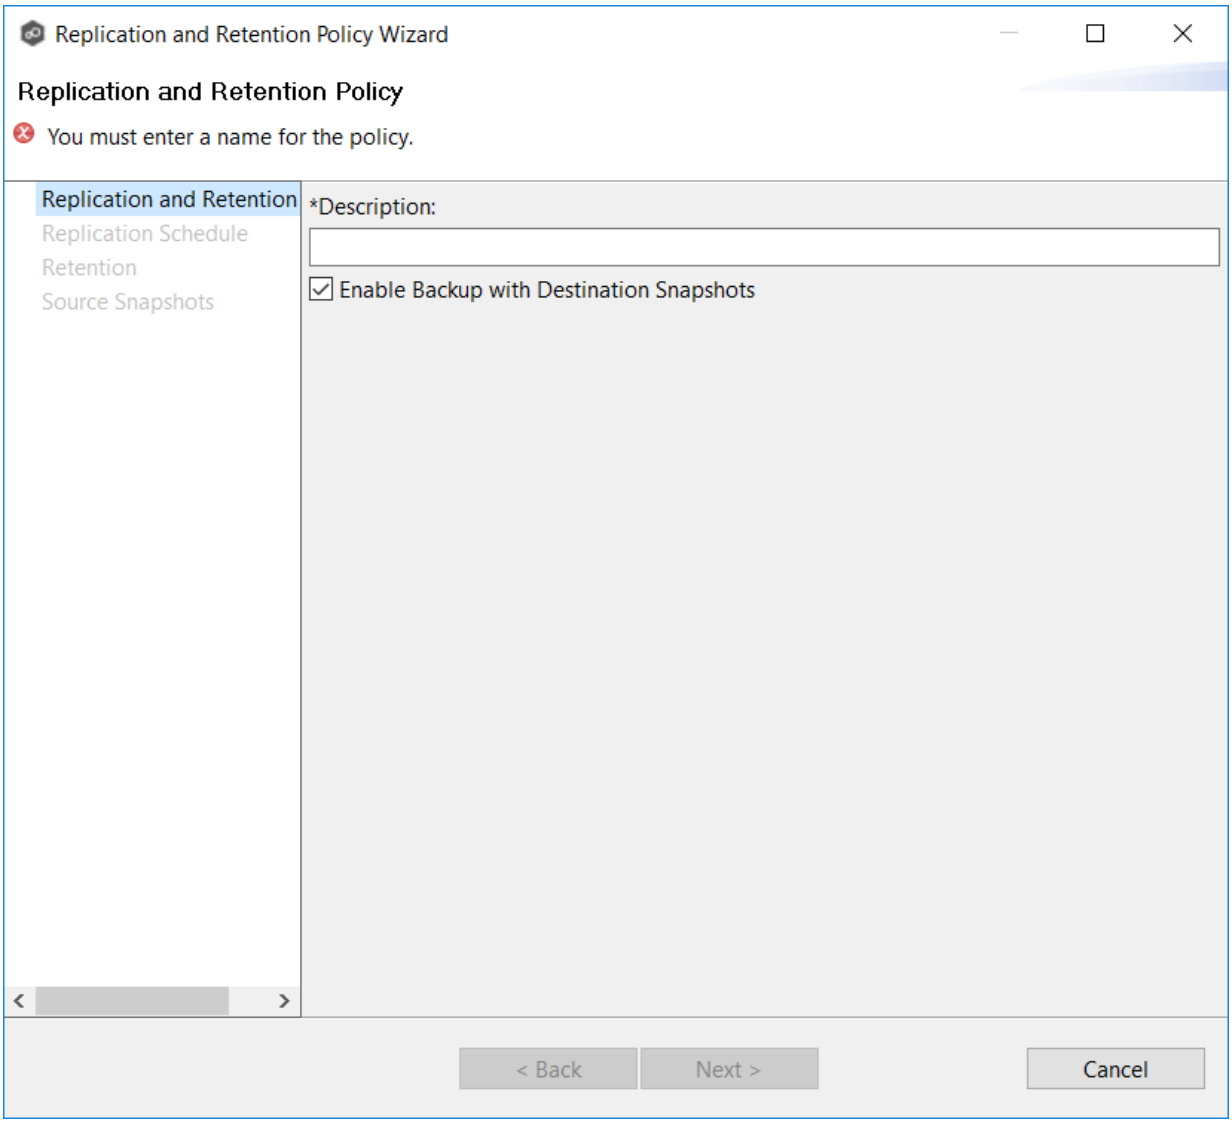

4. Enter the required values, and then click **Finish**.

   See Step 10: Replication and Retention Policy in Creating a Cloud Backup and Replication Job for assistance in completing the wizard.

SNMP Notifications

When you create a job, you can select an existing SMNP notification to apply to the job or you can create a new notification and apply it to the job. You cannot edit or delete an SMNP

notification while it is applied to a job. See SMNP Notifications in the Basic Concepts section for more information about SMNP notifications.

To create an SMNP notification:

1. Select **Preferences** from the **Window** menu.

2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **SNMP Notifications**.

   Any existing SNMP notifications are listed in the **SNMP Notifications** table.

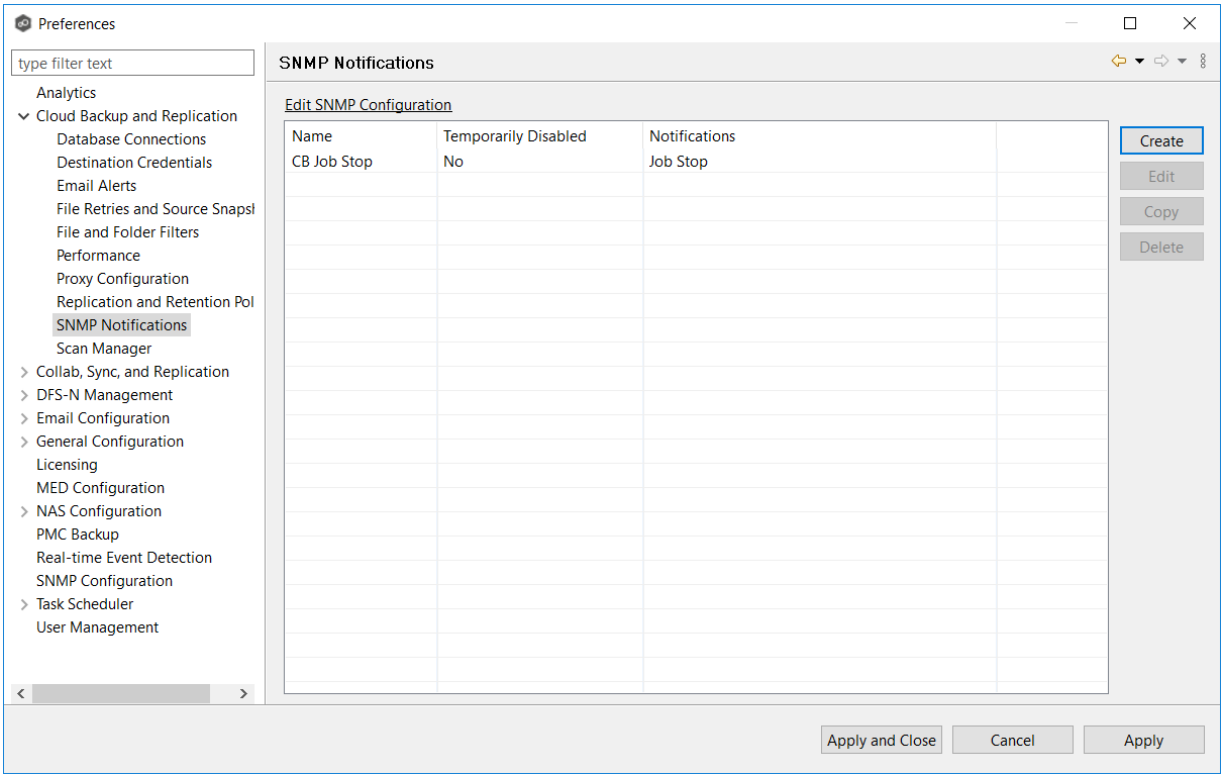

3. Click the **Create** button.

   The **Add SNMP Notification** dialog appears.

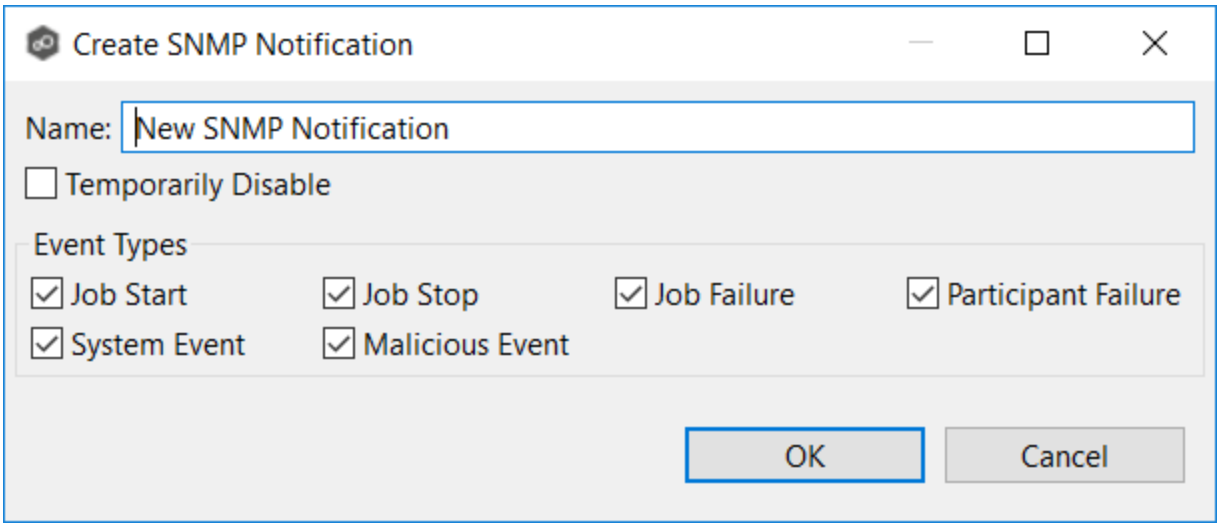

4. Select the types of events that will trigger the generation of an SNMP trap:

Event Type	Description
Job Start	Sends a notification when the job starts.
Job Stop	Sends a notification when the job stops.
Job Failure	Sends a notification when job stops unexpectedly.
Participant Failure	Sends a notification when the Management Agent job disconnects or stops responding.
System Event	Sends a notification when a system event such as low memory or low hub disk space occurs.

Event Type	Description
Malicious Event	Sends a notification when a malicious event is detected. For more information, see MED Configuration.

5. Click **OK**.

   The new notification is listed in the **SNMP Notifications** table and can now be applied to jobs

6. Click **Apply and Close** or **Apply**.

Scan Manager

The Cloud Backup and Replication Scan Manager is responsible for handling all scheduled and on-demand scans of the source storage device.

To modify the Scan Manager settings for Cloud Backup and Replication jobs:

1. Select **Preferences** from the **Window** menu.

2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Scan Manager**.

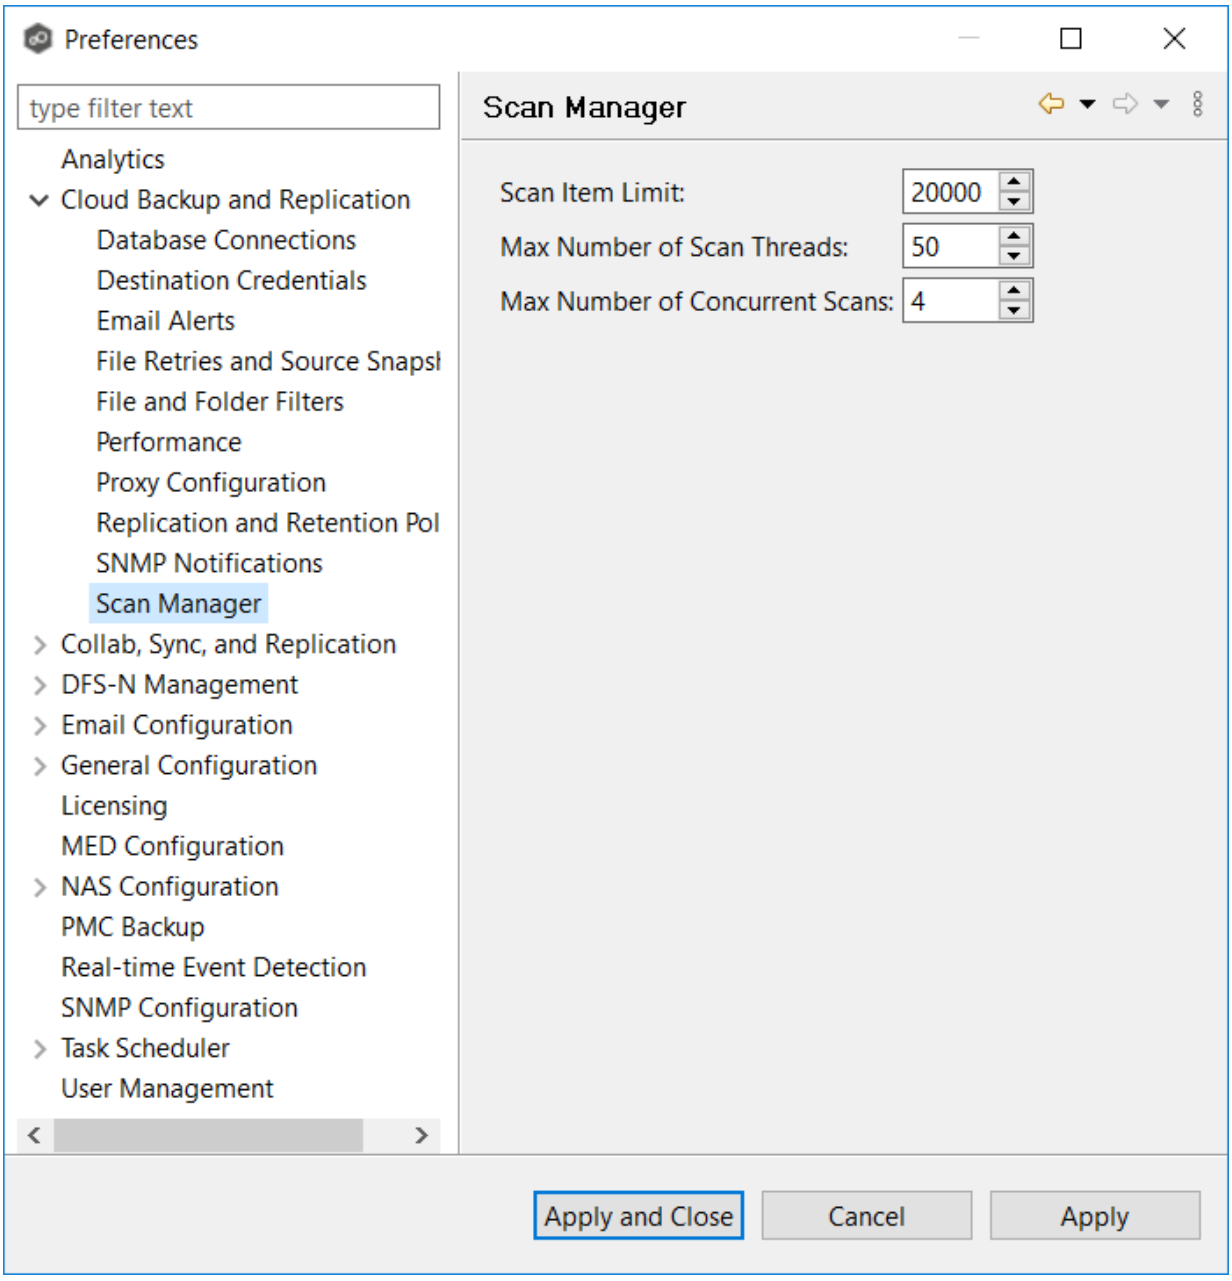

3. Modify the settings as needed.

Setting	Description
Scan Item Limit	Enter the maximum number of files and folders to obtain from a folder structure at a time during a scan.

Setting	Description
Max Number of Scan Threads	Enter the maximum number of threads available for scanning files and folders. Set the number to at least the maximum number of jobs running on any single Management Agent.
Max Number of Concurrent Scans	Enter the maximum number of scans that can run in parallel. If the number of active scan threads is greater than this number, scan threads will process on a rotating basis. Increasing this number can increase scan performance but will also increase system memory and CPU utilization.

4. Click **Apply and Close** or **Apply**.

Collaboration, Replication, and Synchronization Job Preferences

You can modify the following settings for File Collaboration, File Synchronization, and File Replication jobs:

- Collab Sync, and Replication
- DFS-N Management
- Dynamic Storage Utilization
- Email Alerts
- File and Folder Filters
- File Retries
- Locking
- Performance
- Real-time Event Detection
- Revit Enhancements

- SMNP Notifications
- Scan Manager

Collab, Sync, and Replication

These settings control basic GUI and reconnect settings for all File Collaboration, File Synchronization, and File Replication jobs.

To modify these settings:

1. Select **Preferences** from the **Window** menu.
2. Select **Collab, Sync, and Replication** in the navigation tree.

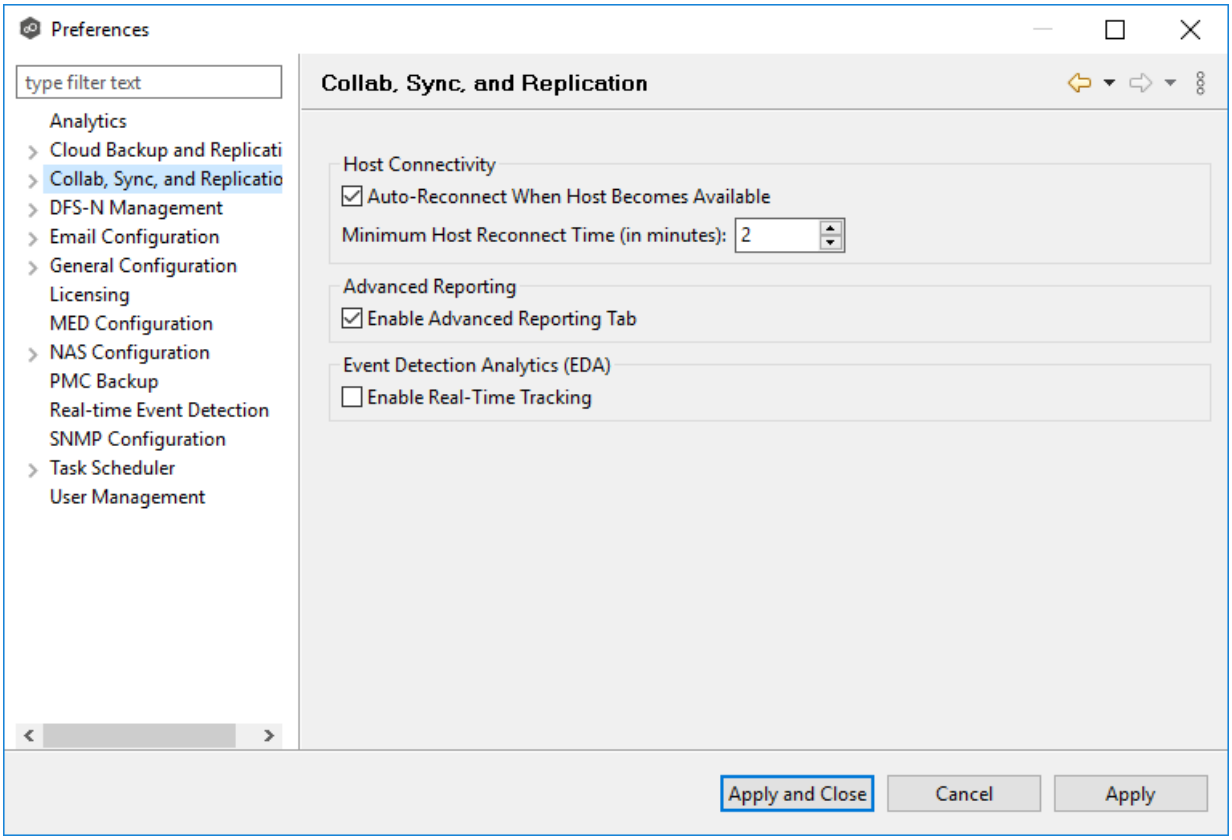

3. Modify the settings as needed.

Option	Description
Auto Reconnect When Host Becomes Available	When an Agent reconnects to Peer Management Center after a failure, automatically re-enables it in any associated jobs. Highly recommended.
Minimum Host Reconnect Time (in minutes)	Enter the minimum number of minutes to wait after an Agent reconnects before re-enabling it in any associated jobs.
Enable Advanced Reporting Tab	Enables the **Reporting** tab in the global **Collab, Sync, and Repl Summary** view.
Enable Real-Time Tracking	Enables **Event Detection Analytics** to track and report common activity processed by Peer Global File Service. If enabled, every 24 hours, an Excel-based report will be written to disk that shows top folders, files, extensions, and users by total processed activity over the previous 24-hour window. These reports are stored under the installation folder of Peer Management Center and can be reviewed by Peer Software Technical Support when uploading log files.

4. Click **Apply and Close** or **Apply**.

DFS-N Management

These settings control the failover and failback capabilities for PeerGFS-managed namespaces that are linked to File Collaboration and File Synchronization jobs.

To modify these settings:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **DFS-N Management**.

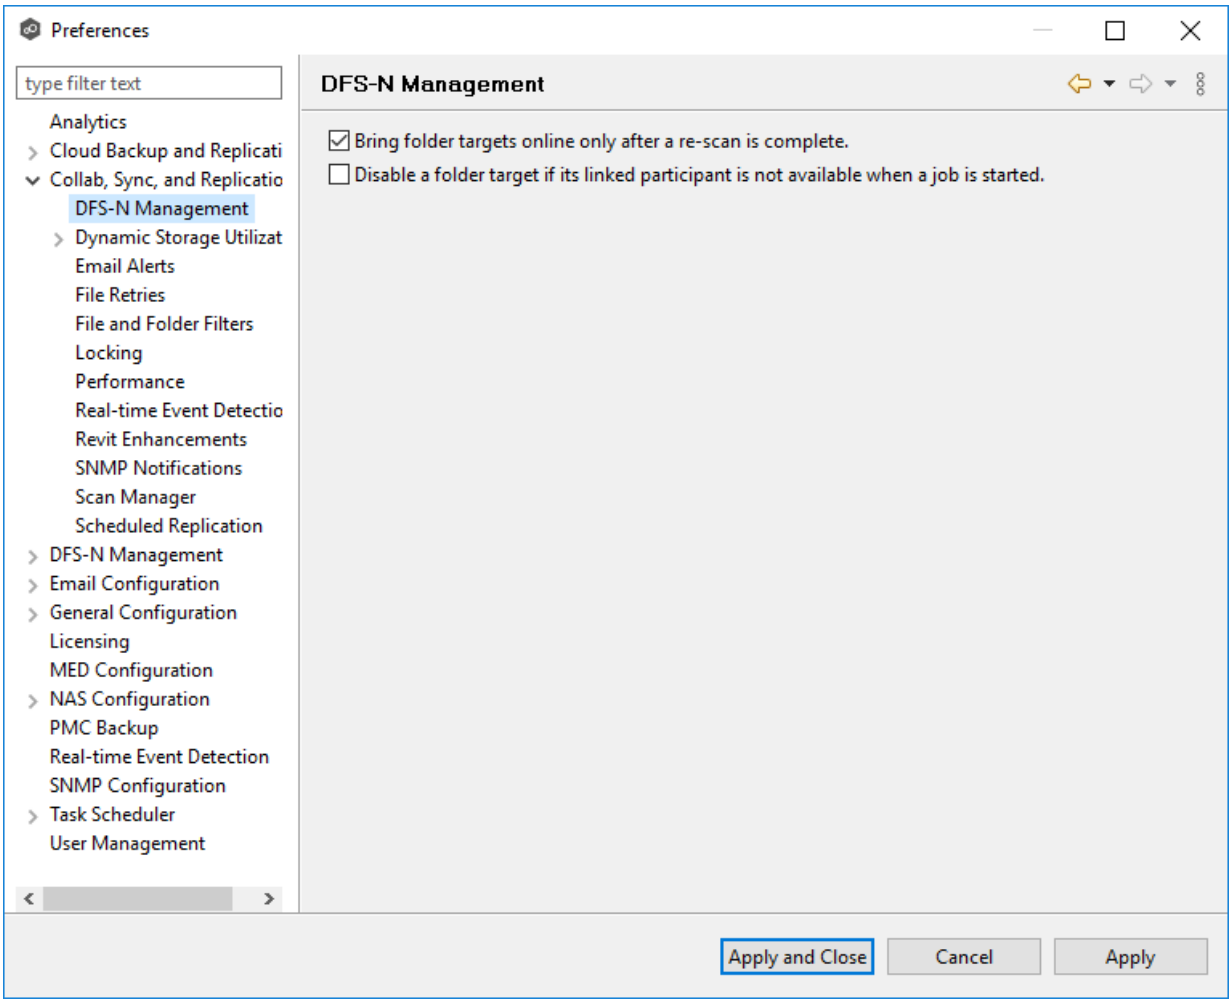

3. Select options as needed.

Option	Description
Bring folder targets online only after a re-scan is complete	Re-enable a disabled folder target in a PeerGFS-managed DFS namespace only when it has been rescanned and is back in sync after an outage. Highly recommended.
Disable a folder target if its linked participant is not available when a job is started	If a File Collaboration or File Synchronization job is started and a participant is not available, automatically disable its associated folder target in a managed DFS namespace.

4. Click **Apply and Close** or **Apply**.

Dynamic Storage Utilization

These settings control the following aspect of jobs that use Dynamic Storage Utilization:

- Dynamic Storage Utilization
- Email Alerts
- Master Data Service
- Pinning Filters
- Utilization Policies
- Volume Policies

The Peer Master Data Service and Peer Edge Service are used by Dynamic Storage Utilization. The Peer Master Data service is a web service that handles requests from edge participants for files on a master participant. it runs on Master participants and is configurable. The primary job of the Peer Edge service is to service reparse requests for Peer stub files and read and/or rehydrate stub files.

Use this page to set the frequency that these services are checked to see if they are operational.

To change the frequency:

1. Select **Preferences** from the **Window** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Dynamic Storage Utilization**.
3. Change the frequency.

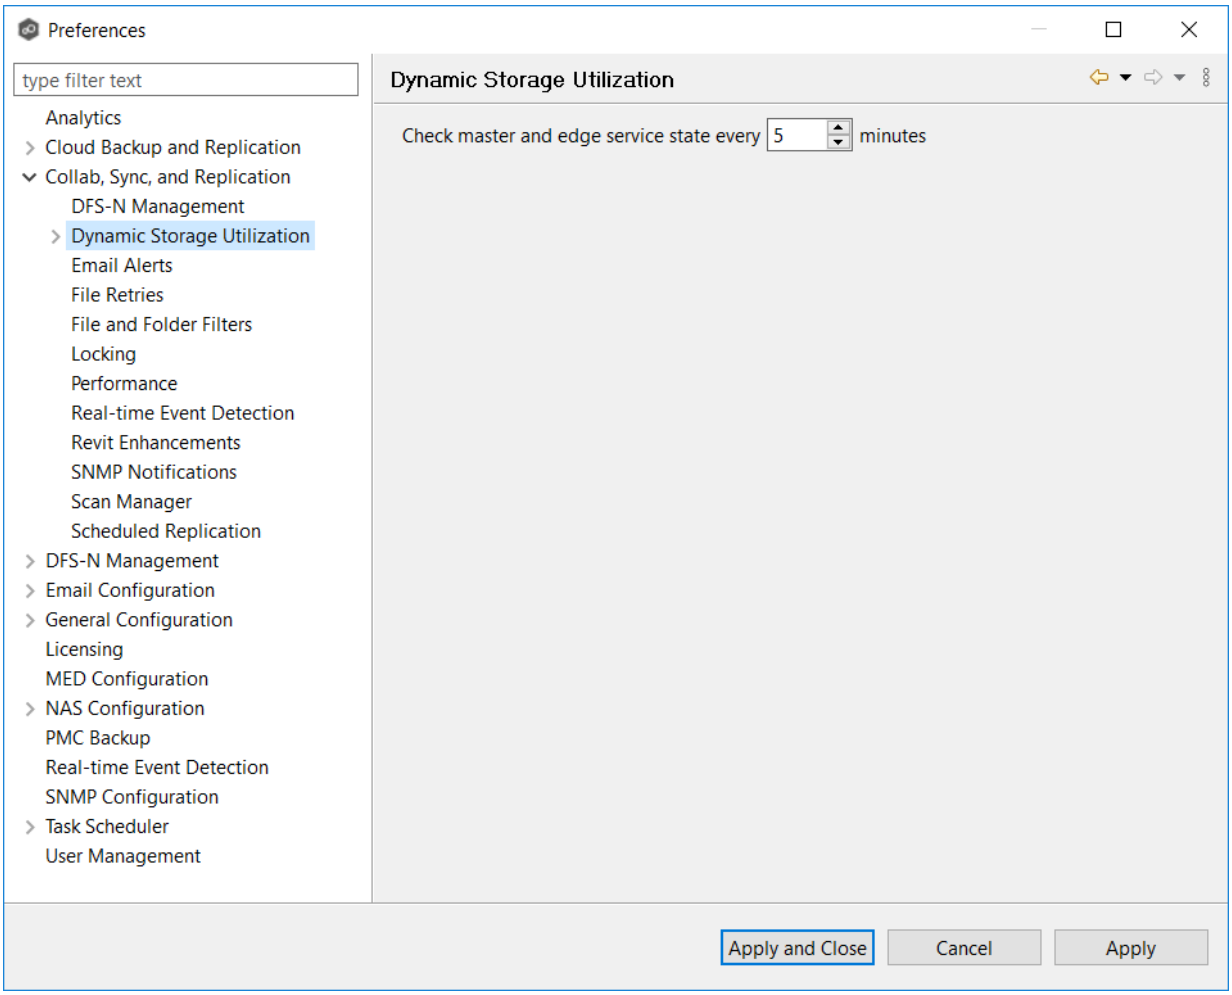

4. Click **Apply and Close** or **Apply**.

When you create a DSU-enabled job, you can select existing DSU email alerts to apply tdo the job or you can create new email alerts and apply them to the job. This Preferences page lists the existing email alerts for DSU-enabled jobs. From this page, you can create, edit, and delete DSU alerts. However, you cannot edit or delete an email alert while it is applied to a job. See Email Alerts in the Basic Concepts section for more information about email alerts.

Note: An SMTP email connection must be configured before email alerts can be sent. See Email Configuration for information about configuring SMTP email settings.

To create a DSU email alert:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Dynamic Storage Utilization**.

3. Select **Email Alerts**.

   Any existing DSU email alerts are listed in the **Email Alerts** table.

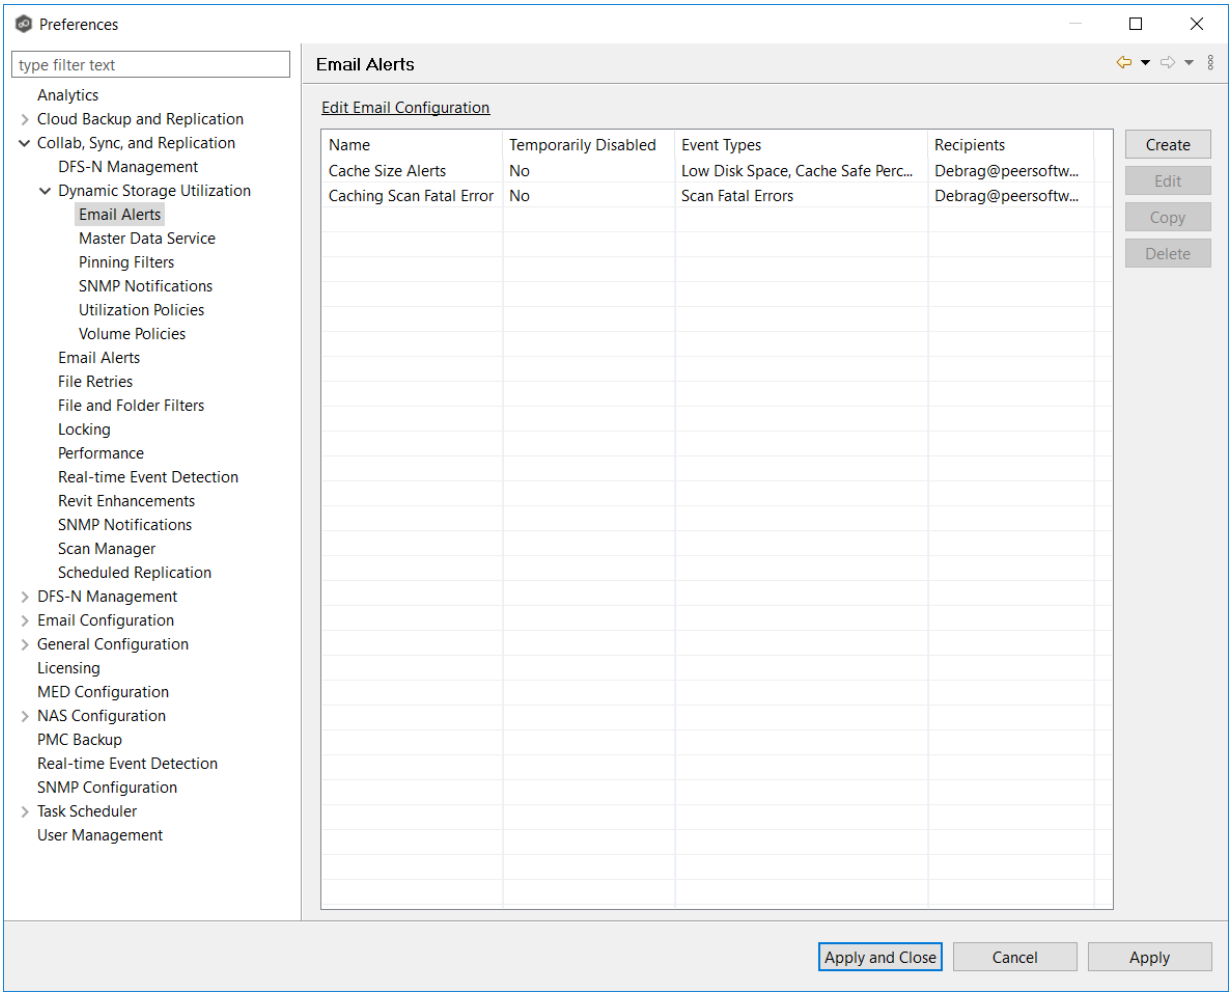

4. Click **Create**.

   The **Create Email Alert** dialog appears.

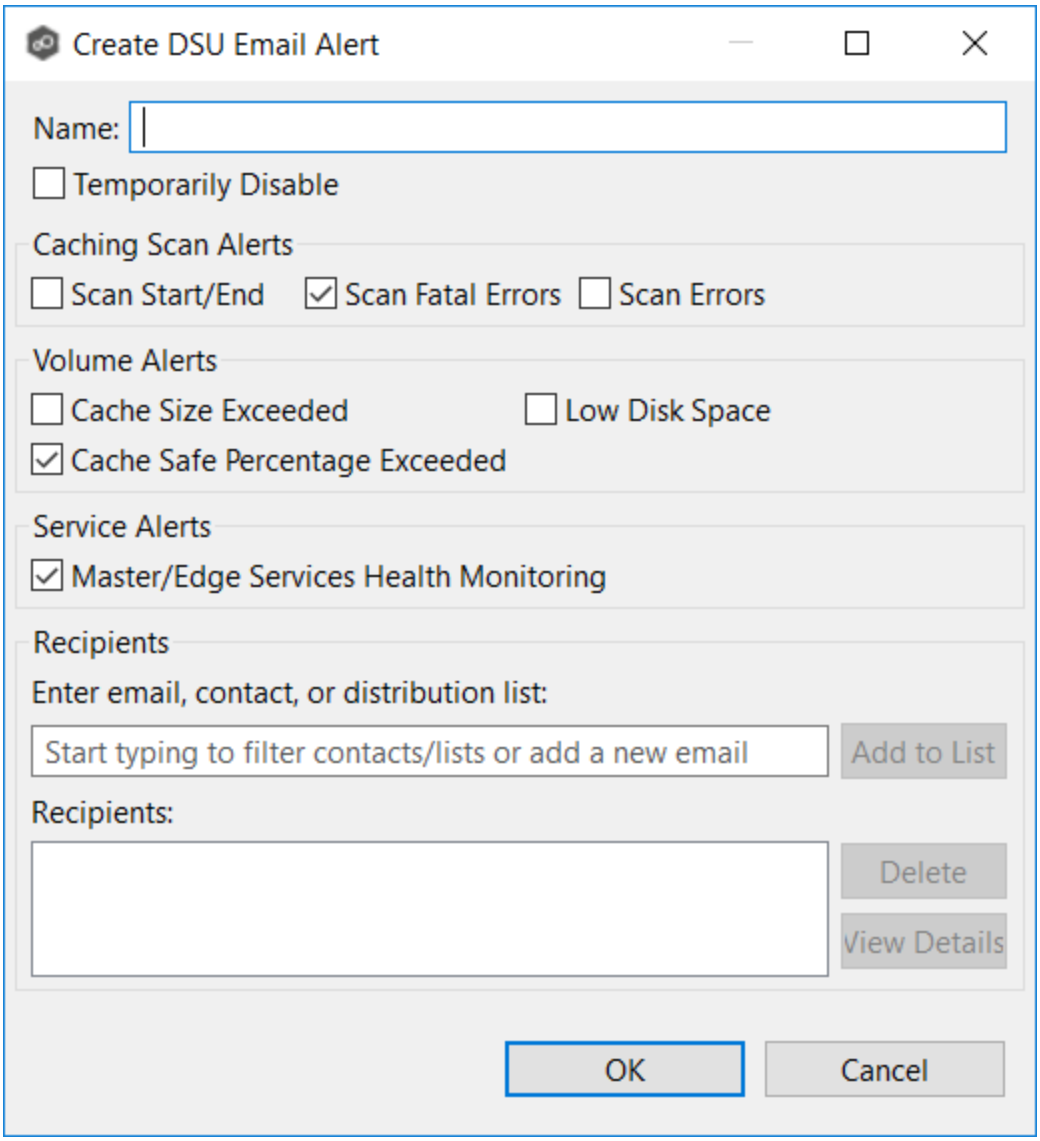

5. Enter a name for the alert.

6. Select the caching scan event types to be alerted.

Event Type	Description
Scan Start/End	Sends a notification when a caching scan is started or stopped.
Scan Fatal Errors	Sends a notification when a fatal error occurs during a caching scan.

Event Type	Description
Scan Errors	Sends a notification when errors occur during a caching scan.
Cache Size Exceeded	Sends a notification when the amount of volume disk space used by DSU exceeds the size specified by the **Cache Size** option in the volume policy.
Low Disk Space	Sends a notification the volume disk space falls below the size specified by the **Disk space is less than X** option in the volume policy.
Cache Safe Percentage Exceeded	Sends a notification when the percentage specified by the **Cache usage exceeds X% of cache size** option in the volume policy.
Master/Edge Services Health Monitoring	Sends a notification if either the Peer Master Data Service or the Peer Edge Service goes down.

7. Enter alert recipients, and then click **Add to List**.

   The recipients are listed in the **Recipients** field.

8. Click **OK**.

   The new email alert is listed in the **Email Alerts** table and can now be applied to DSU-enabled jobs.

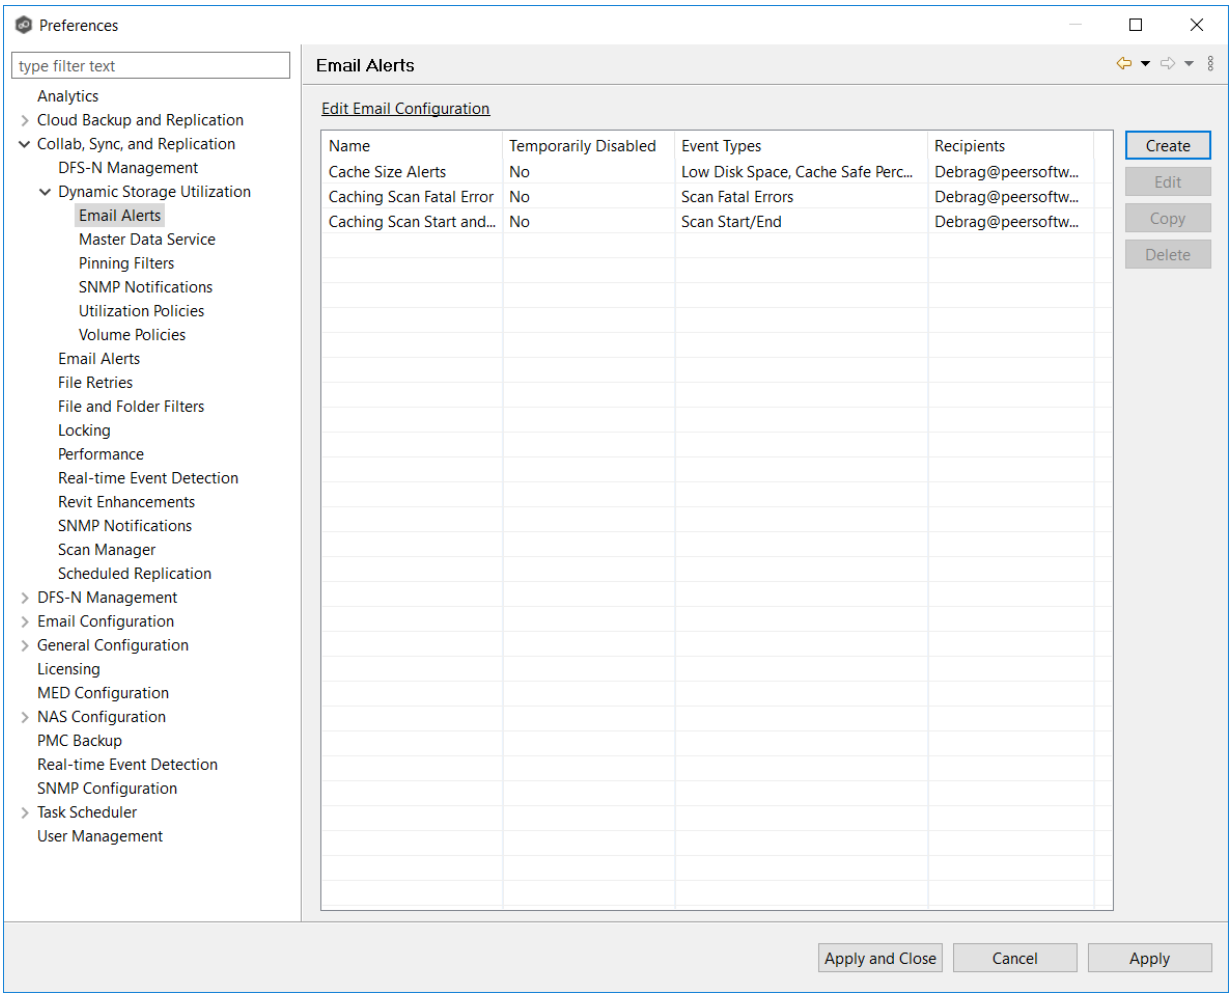

9. Click **Apply and Close** or **Apply**.

When you create a job that is DSU-enabled, you specify the settings to be used for the Peer Master Data Service. The Peer Master Data Service handles requests from edge participants for files on a master participant. The Peer Master Data Service is installed as part of the Peer Agent installation process. The **Master Data Service** page displays the existing parameters for the Master Data Service.

To edit the Master Data Service configuration:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Dynamic Storage Utilization**.

3. Select **Master Data Service**.

   The **Master Data Service** table lists master participants. The **In Use** column identifies whether the participant is currently being used as a master participant in a job.

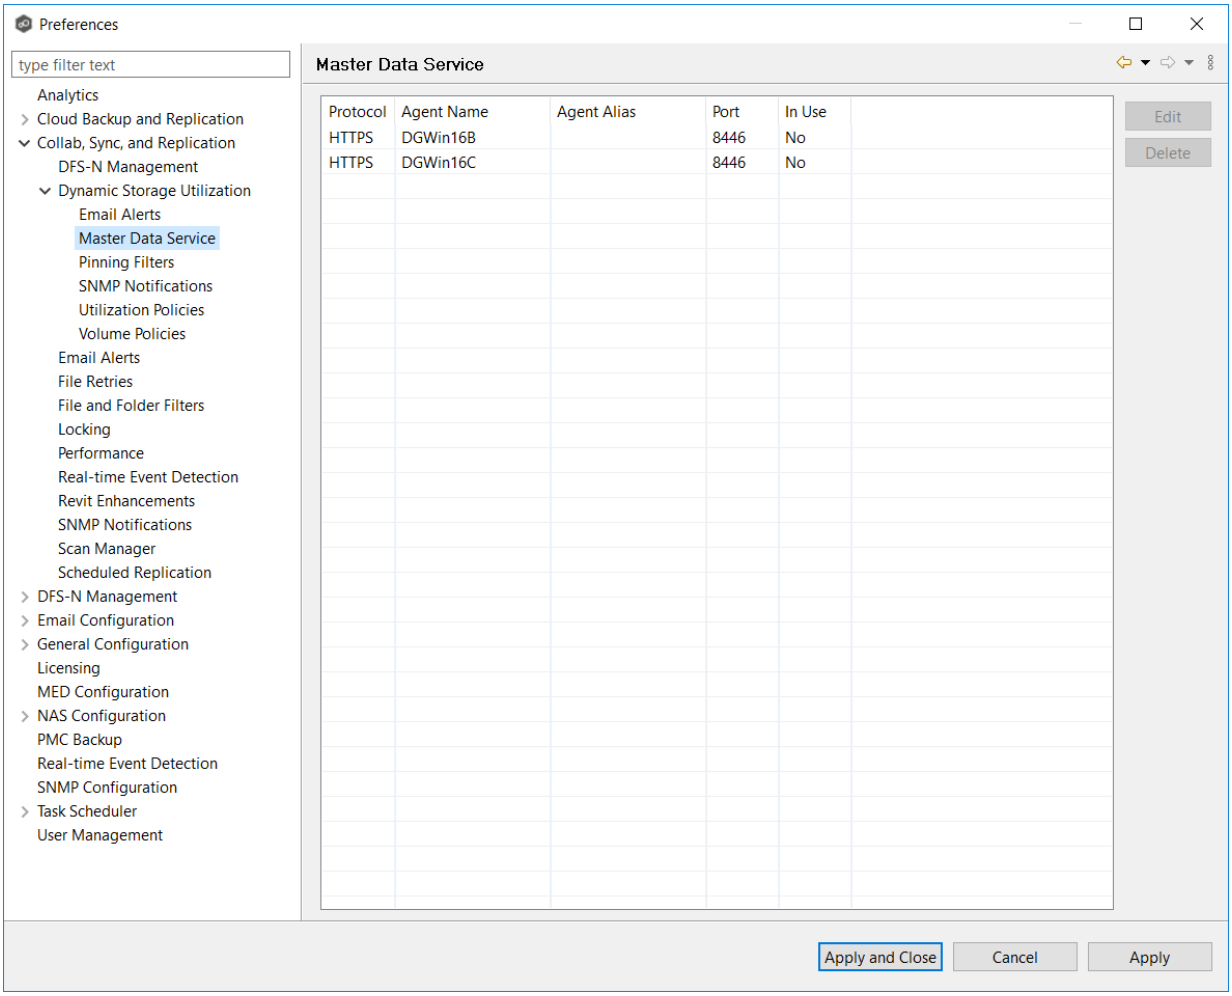

4. Select a participant, and then click **Edit**.

   If you selected a master participant currently being used in a job, the **In Use Configuration** dialog appears. Click **OK** to close the dialog.

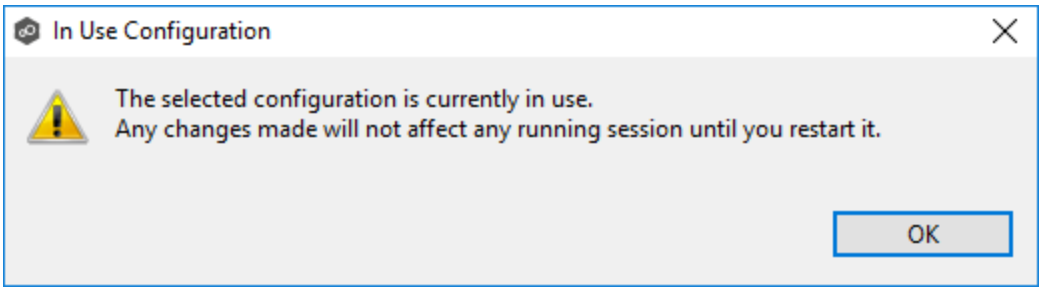

Otherwise, the **Edit Master Data Service** dialog appears. The first two fields on this page are automatically populated:

- **Protocol**: This field lists the protocol that will be used to transfer file content between master participants and edge participants. HTTPS is currently the only available option as it requires only a single open firewall port on each master participant and does a better job of handling latency over WAN-based connections.
- **Agent Name**: This field lists the name of the Agent.

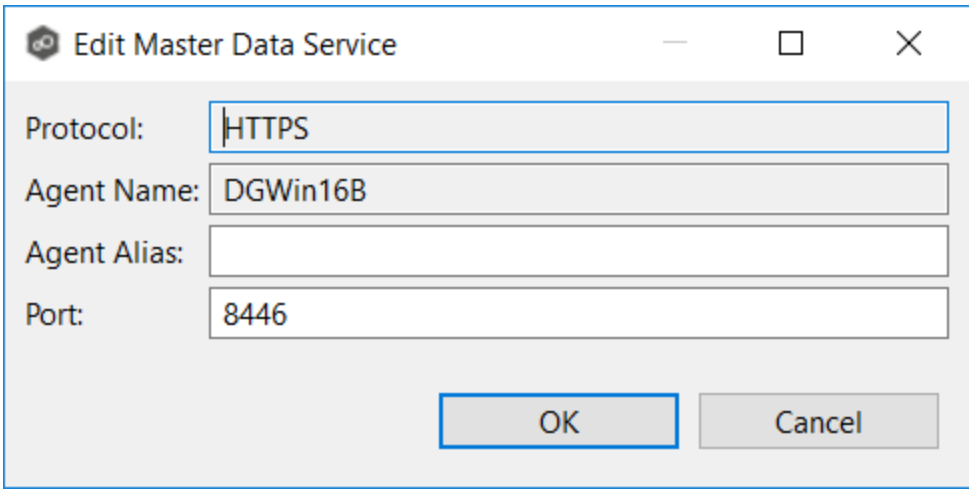

5. (Optional) Enter a value for **Agent Alias**; the value can be a hostname, FDQN, or IP address.

   A value for **Agent Alias** is required only if the name of the Agent cannot be converted to an IP address via DNS. If an alias name is entered, it will be used by the Edge Service on edge participants to connect with the Master Data Service. If no alias name is entered, the name of the Agent will be used.

6. (Optional) Modify the default value for **Port** if you are use a different port.

   If you modify the port number, the Master Data Service will be restarted and the new port number will take effect immediately.

7. Click **OK**.

The revised Master Data Service is listed in the Master Data Service table.

8. Click **Apply and Close** or **Apply**.

   The new settings will be applied to all DSU-enabled jobs.

When you create a DSU-enabled job, you can select existing pinning filters to apply to the job or you can create new pinning filters and apply them to the job. This Preferences page lists the existing pinning filters. From this page, you can create, edit, and delete pinning filters. However, you cannot edit or delete a pinning filter while it is applied to a job.

A pinning filter specifies whether specific files or files in a particular directory are always stubbed or always local on the edge participant. A pinning filter similar is to a utilization policy —it can be applied to multiple jobs. If there is a conflict between a pinning filter and utilization policy (where, for example, you might have something set to be always stubbed), the pinning filter will take precedence.

To create a pinning filter:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Dynamic Storage Utilization**.

3. Select **Pinning Filters**.

   Any existing pinning filters are listed in the **Pinning Filters** table.

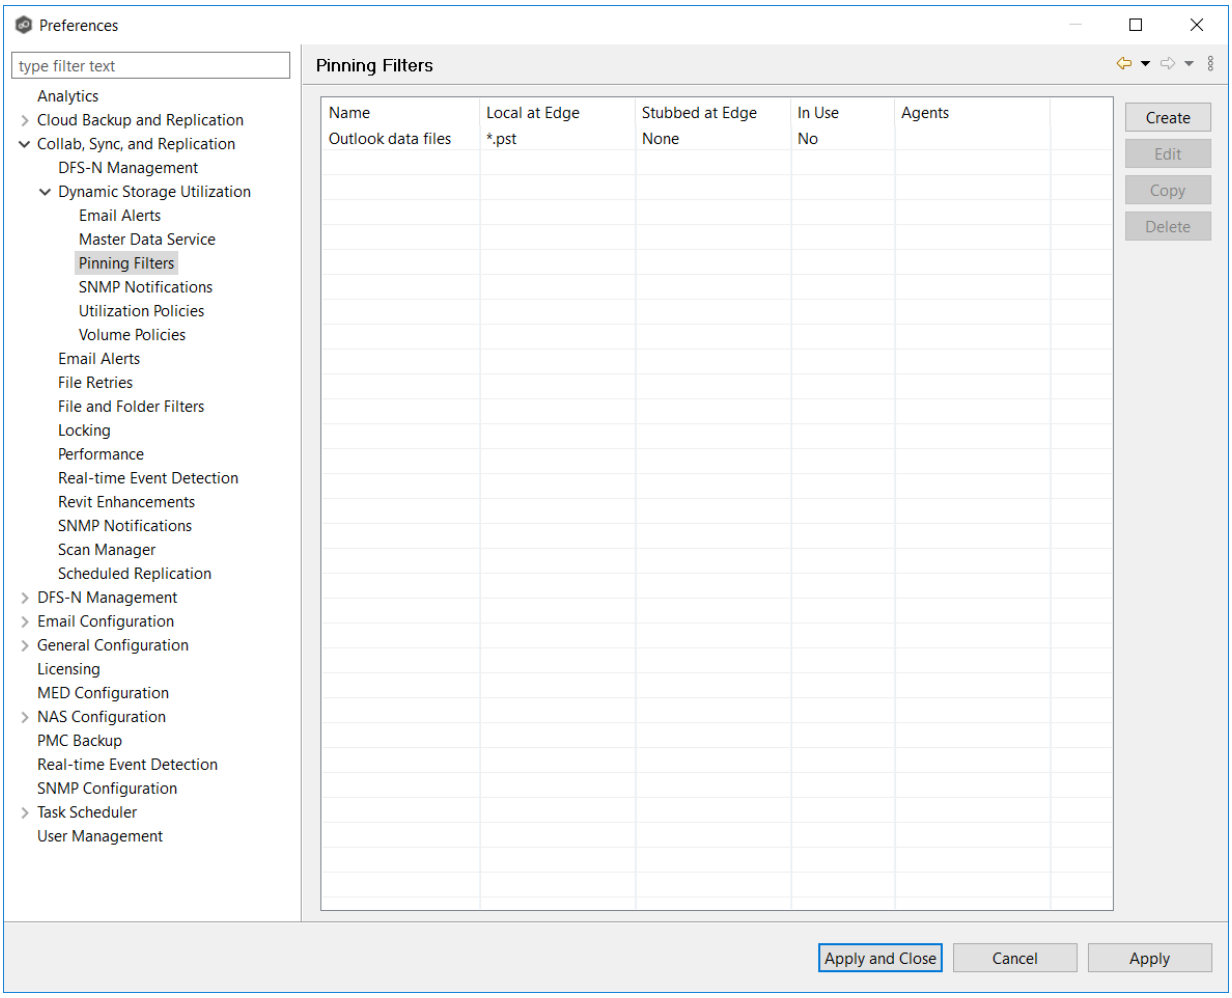

4. Click **Create**.

   The **Create Pinning Filter** dialog appears.

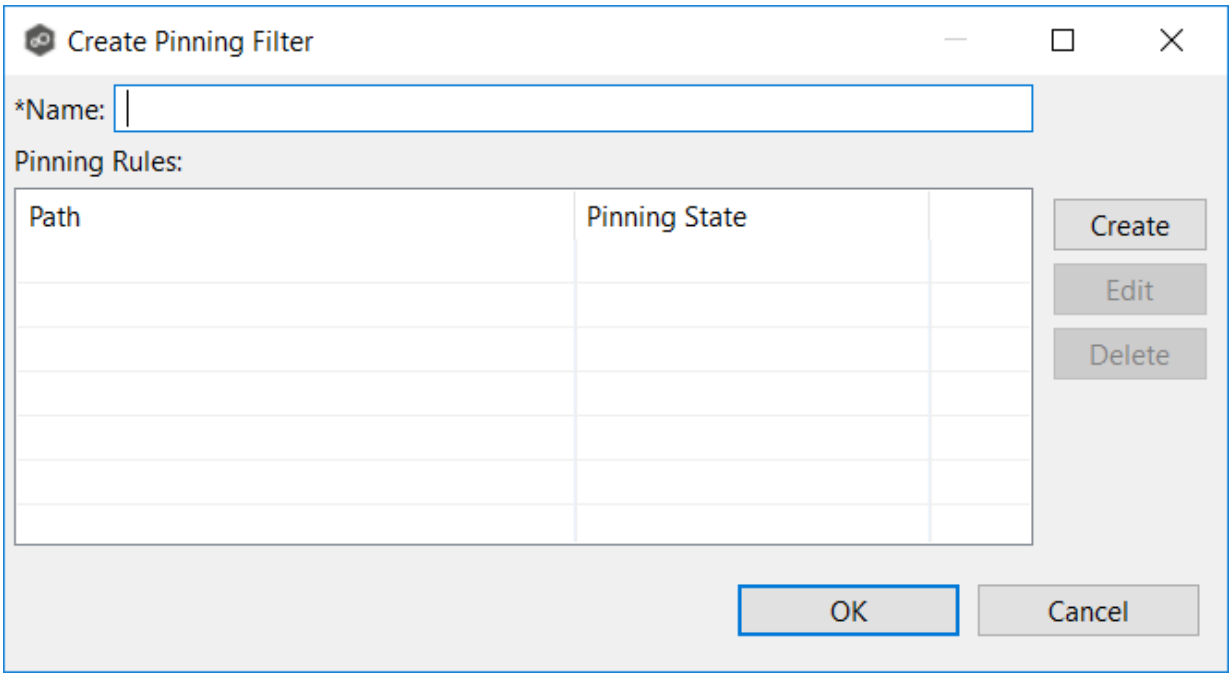

5. Enter a name for the pinning filter.
6. Click **Create** to add a pinning rule to the filter.

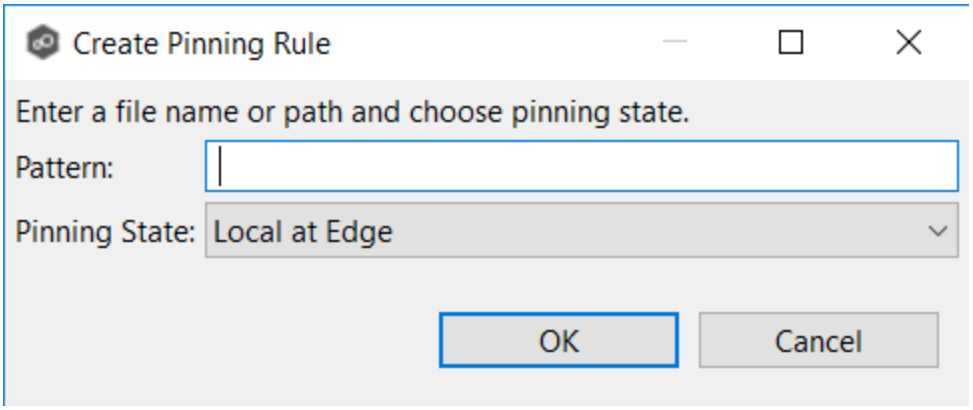

7. Enter a file name or enter a path.
8. Choose a pinning state:
   - Select **Local at Edge** if you want the specified files or path to always be local and never stubbed.
   - Select **Stubbed at Edge** if you want the specified files or path to always be stubbed at edge.
     AuA

9. Click **OK**.

   The rule appears in the **Pinning Rules** table.

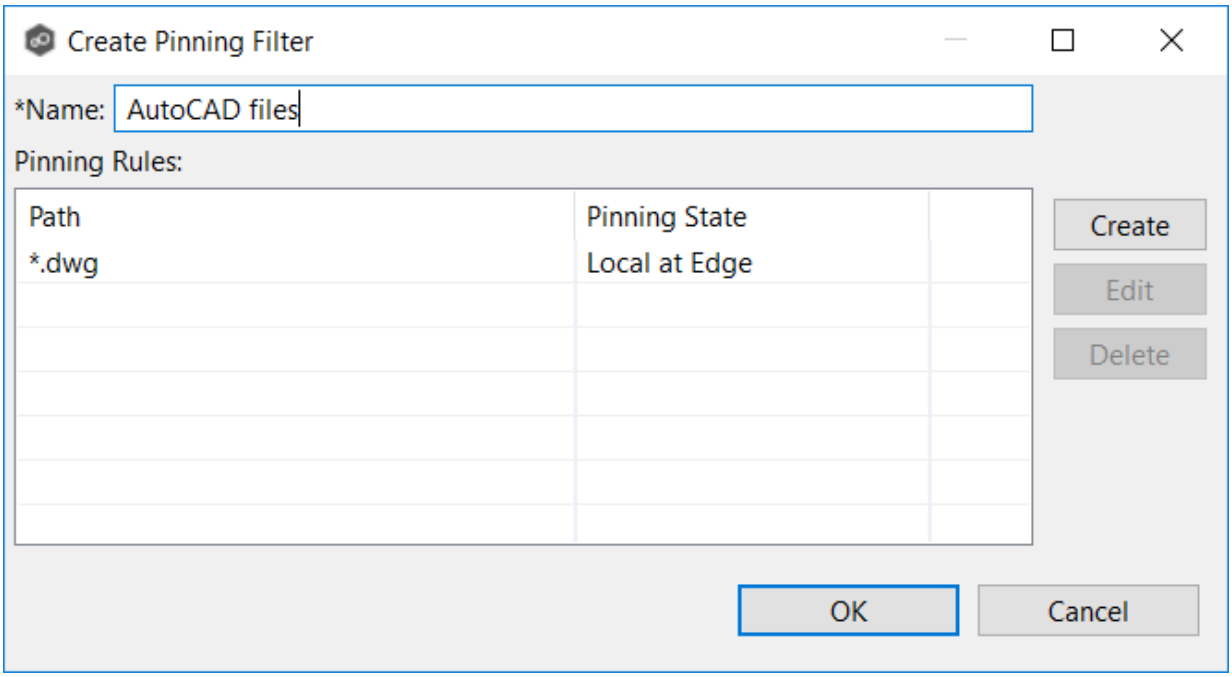

10. If you want to add additional rules to the pinning filter, repeat Steps 6-9.

11. Click **OK** to close the **Create Pinning Filter** dialog.

    The new filter is listed in the **Pinning Filters** table and can now be applied to DSU-enabled jobs.

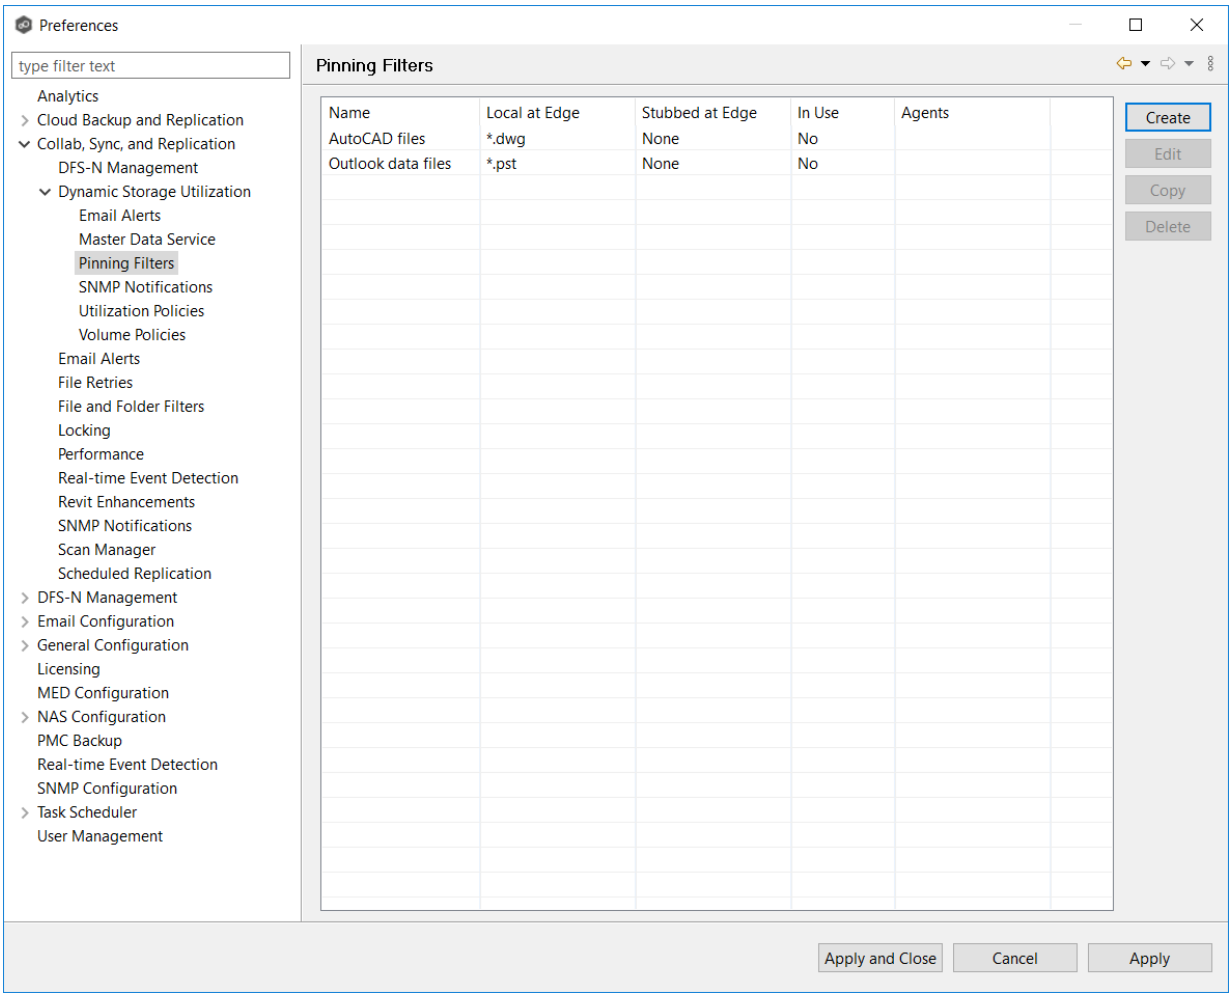

12. Click **Apply and Close** or **Apply**.

When you create a job, you can select an existing SMNP notification to apply to the job or you can create a new notification and apply it to the job. You cannot modify or delete an SMNP notification while it is applied to a job. See SMNP Notifications in the Basic Concepts section for more information about SMNP notifications.

To create an SMNP notification:

1. From the **Window** menu, select **Preferences**.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Dynamic Storage Utilization**.

3. Select **SNMP Notifications**.

   Any existing SNMP notifications are listed in the **SNMP Notifications** table.

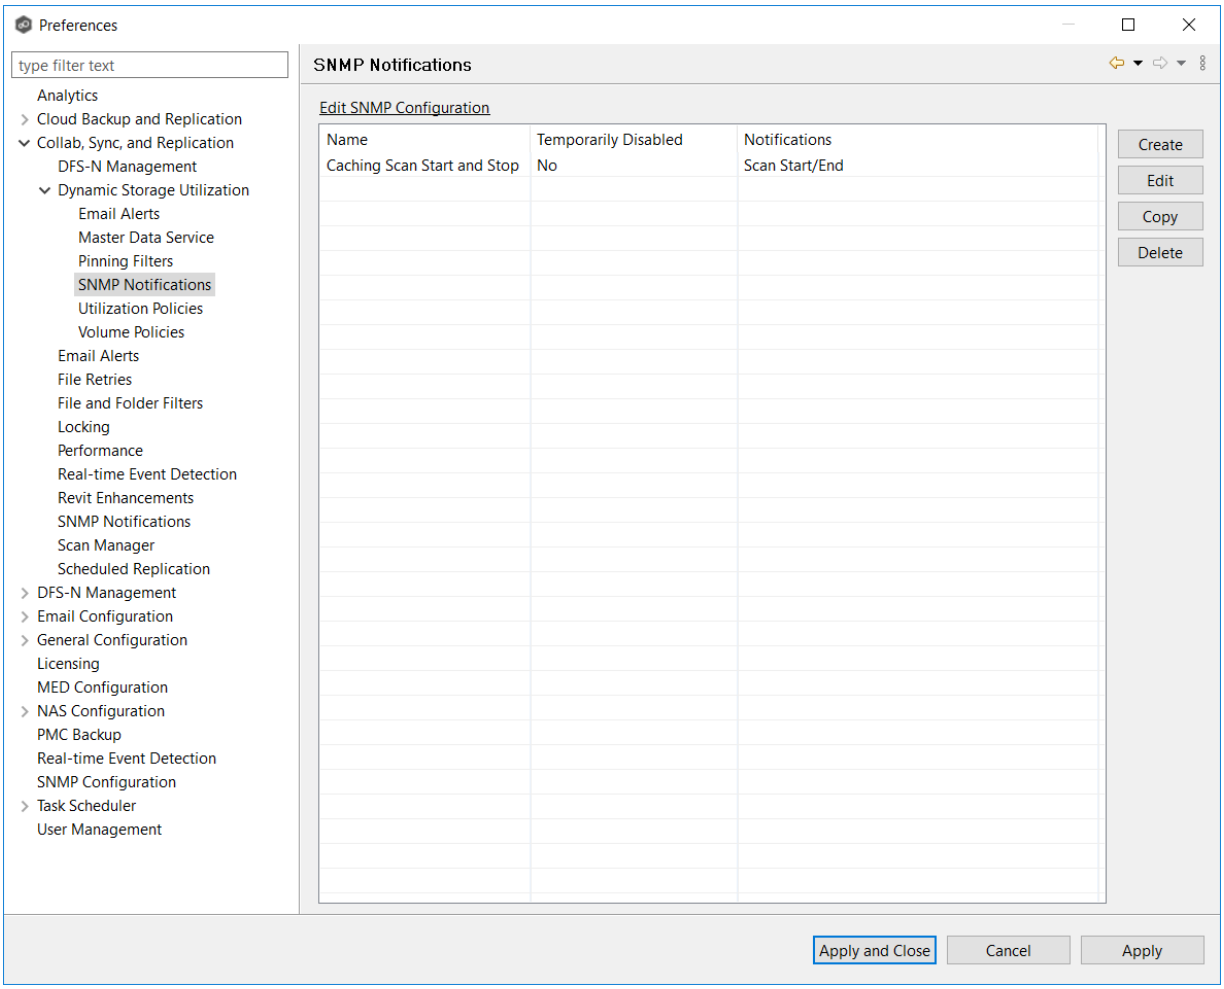

4. Click the **Create** button.

   The **Create DSU SNMP Notification** dialog appears.

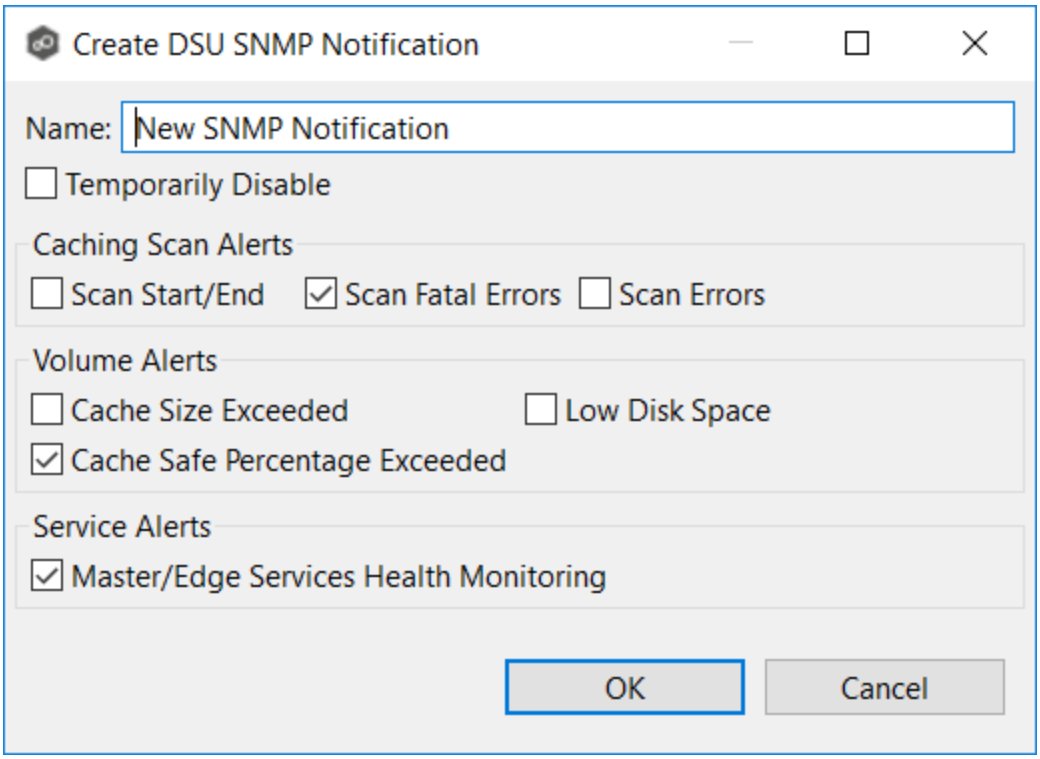

5. Select the types of events that will trigger the generation of an SNMP trap:

Event Type	Description
Scan Start/End	Sends a notification when a caching scan is started or stopped.
Scan Fatal Errors	Sends a notification when a fatal error occurs during a caching scan.
Scan Errors	Sends a notification when errors occur during a caching scan.
Cache Size Exceeded	Sends a notification when the amount of volume disk space used by DSU exceeds the size specified by the **Cache Size** option in the volume policy.
Low Disk Space	Sends a notification the volume disk space falls below the size specified by the **Disk space is less than X** option in the volume policy.

Event Type	Description
Cache Safe Percentage Exceeded	Sends a notification when the percentage specified by the **Cache usage exceeds X % of cache size** option in the volume policy.
Master/Edge Services Health Monitoring	Sends a notification if either the Peer Master Data Service or the Peer Edge Service goes down.

6. Click **OK**.

   The new notification is listed in the **SNMP Notifications** table and can now be applied to DSU-enabled jobs.

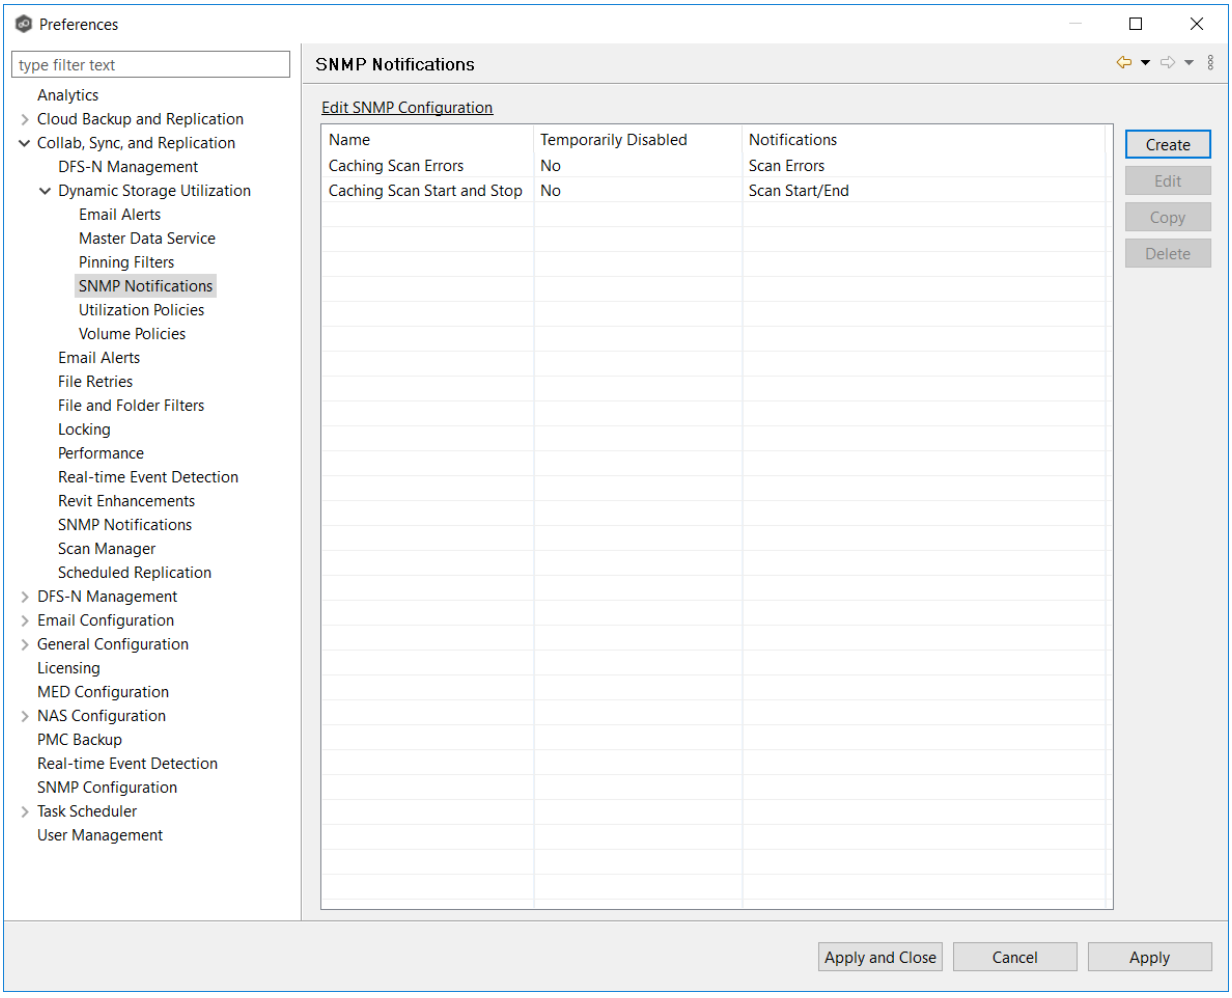

7. Click **Apply and Close** or **Apply**.

When you create a DSU-enabled job, you can select an existing utilization policy to apply to the job or create a new utilization policy and apply it to the job. This Preferences page lists the existing utilization policies. From this page, you can create, edit, and delete utilization policies. However, you cannot edit or delete a utilization policy while it is applied to a job.

A **utilization policy** defines when a file should be stubbed versus fully hydrated across all volumes of an edge participant. The policy parameters are based on the size of the files to be potentially stubbed and when they were last accessed and modified. A utilization policy enables you to balance getting the best performance while keeping the cache as full as possible.

To create a utilization policy:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Dynamic Storage Utilization**.

3. Select **Utilization Policies** in the navigation tree.

   Any existing utilization policies are listed in the **Utilization Policies** table.

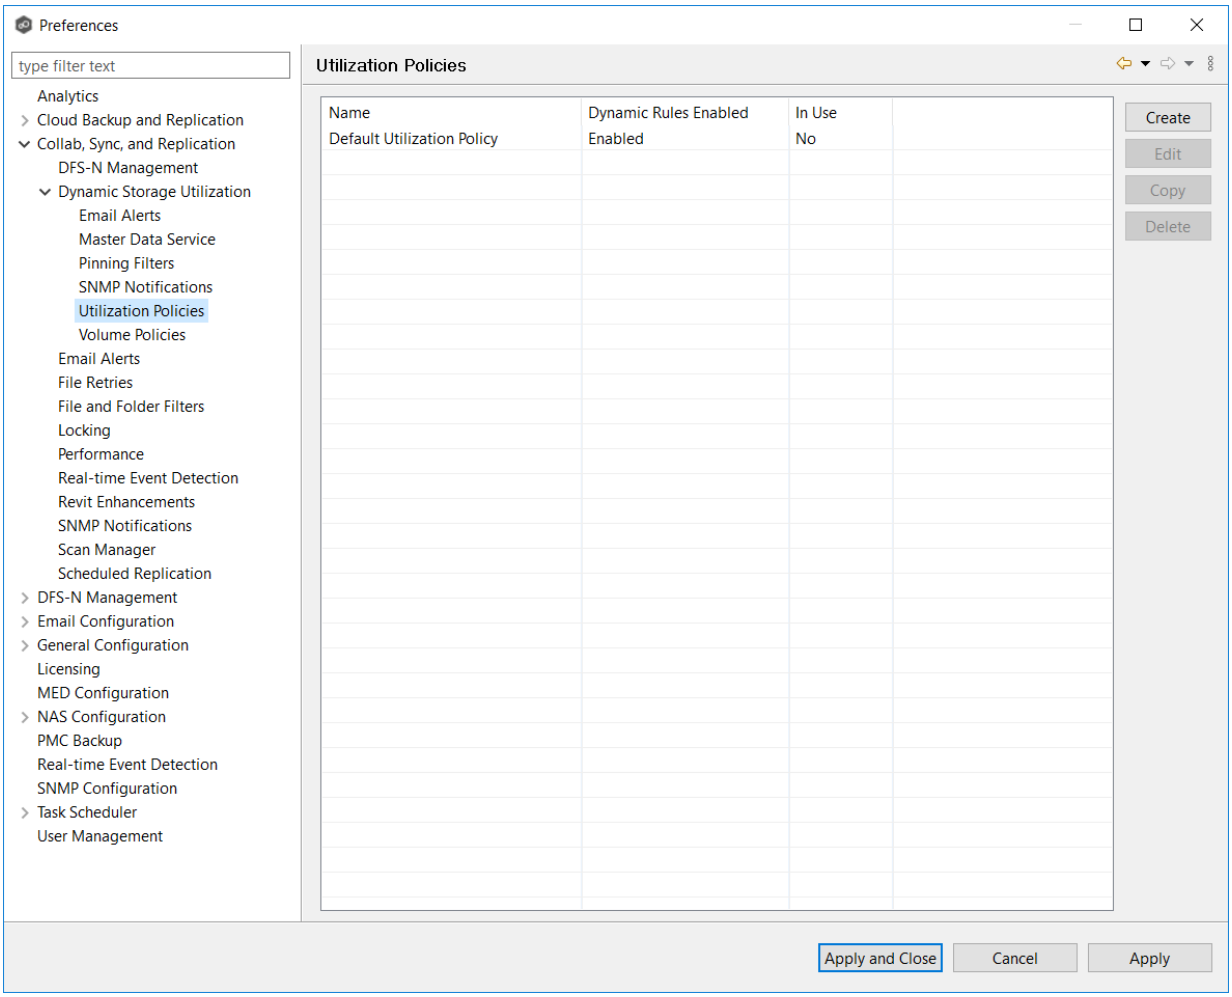

4. Click **Create**.

   The **Create Utilization Policy** dialog appears.

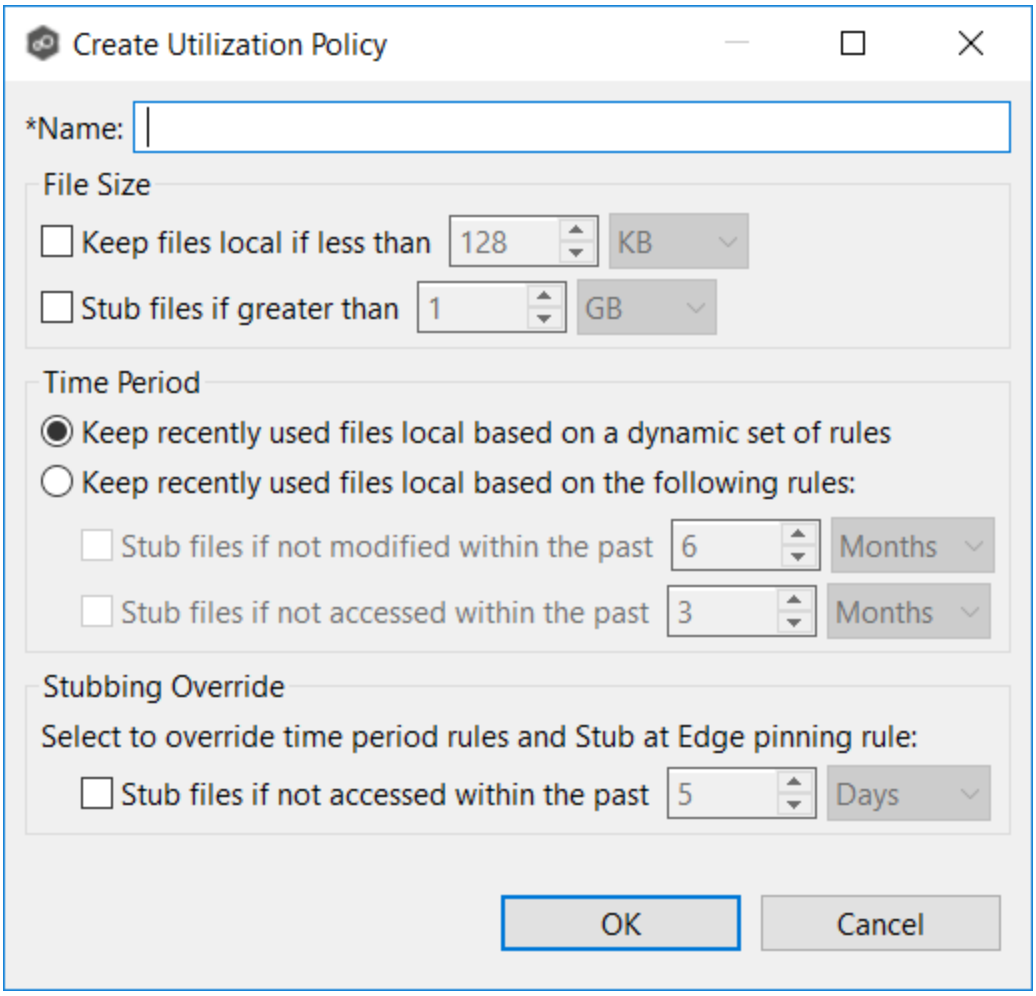

5. Enter a name for the policy.

6. (Optional) In the **File Size** section, select one or both options:

Field	Description
Keep files local if less than X size	Select this option if you want files under a specified size to remain local .
Stub files if greater than X size	Select this option if you want files over a specified size to be stubbed.

7. (Optional) In the **Time Period** section, select one of the options:

Field	Description
Keep recently used files local based on a dynamic set of rules	Select this option if you want DSU to control when to stub files based on last accessed and last modified times. DSU dynamically adjusts the accessed and modified time periods it uses based on the total amount of space that DSU is actively using on a volume.
Keep recently used files local based on the following rules	Select this option if you want to specify the dates used when stubbing files based on the last accessed and last modified.

8. If you selected the **Keep recently used files local based on the following rules** option in **Time Period**, two additional options are enabled; select the options to be used and specify the time periods to be used:

Field	Description
Stub files if not modified within the past X time period	Select this option and specify a time range to use the last modified date of a file to determine if it should be kept local or stubbed.
Stub files is not accessed within the past X time period	Select this option and specify a time range to use the last accessed date of a file to determine if it should be kept local or stubbed.

9. (Optional) in the **Stubbing Override** section, select the checkbox and a time period if you want to use the last accessed date of all recently hydrated files to determine if they should be kept local or re-stubbed.

10. Click **OK**.

    The new policy is listed in the **Utilization Policies** table and can now be applied to jobs.

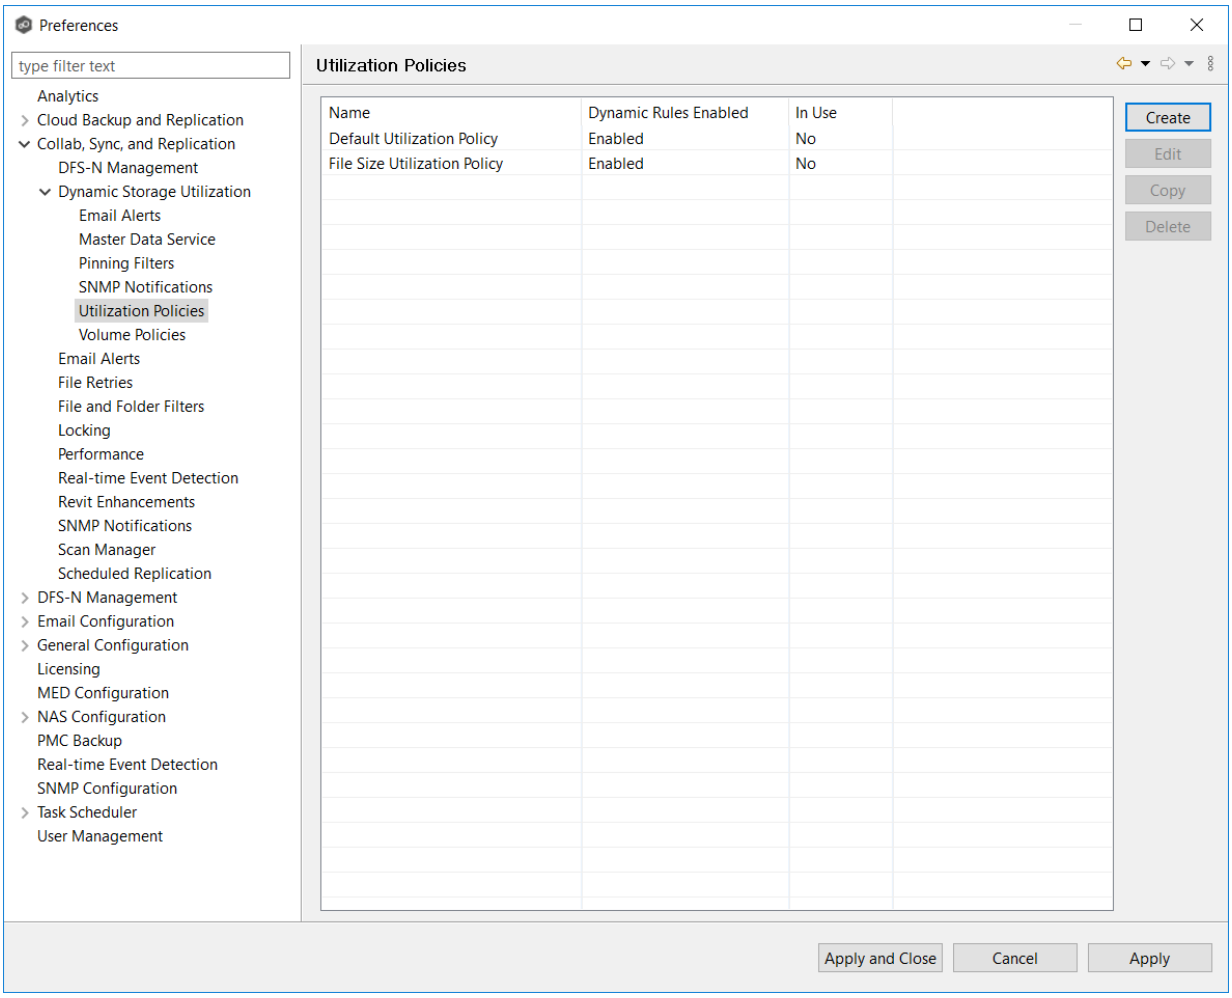

11. Click **Apply and Close** or **Apply**.

When you create a DSU-enabled job, you specify a volume policy for an edge participant. A **volume policy** specifies how much of the available space on the volume monitored by the edge participant to assign for local (hydrated) files. This space is referred to as a **cache**. The volume refers to the drive letter of the path set on the **Path** page of the wizard (for example, if the participant is configured to monitor D:\Data, the volume policy for this participant would apply to the D volume).

If the Agent you selected is already being used as an edge participant in another job utilizing DSU, the existing volume policy will be displayed on this page. You can edit the existing volume policy; however, be aware that any modifications to the volume policy will be applied to every other job that use this Agent as an edge participant and "touch" the same volume.

From this page, you can edit and delete volume policies. However, note that:

- Any changes made to a volume will not be applied to a running job until the job is restarted.
- You cannot delete a volume policy while it is being used by an edge participant.

To edit a volume policy:

1. Select **Preferences** from the **Window** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Dynamic Storage Utilization**.
3. Select **Volume Policies** in the navigation tree.

   Any existing volume policies are listed in the **Volume Policies** table.

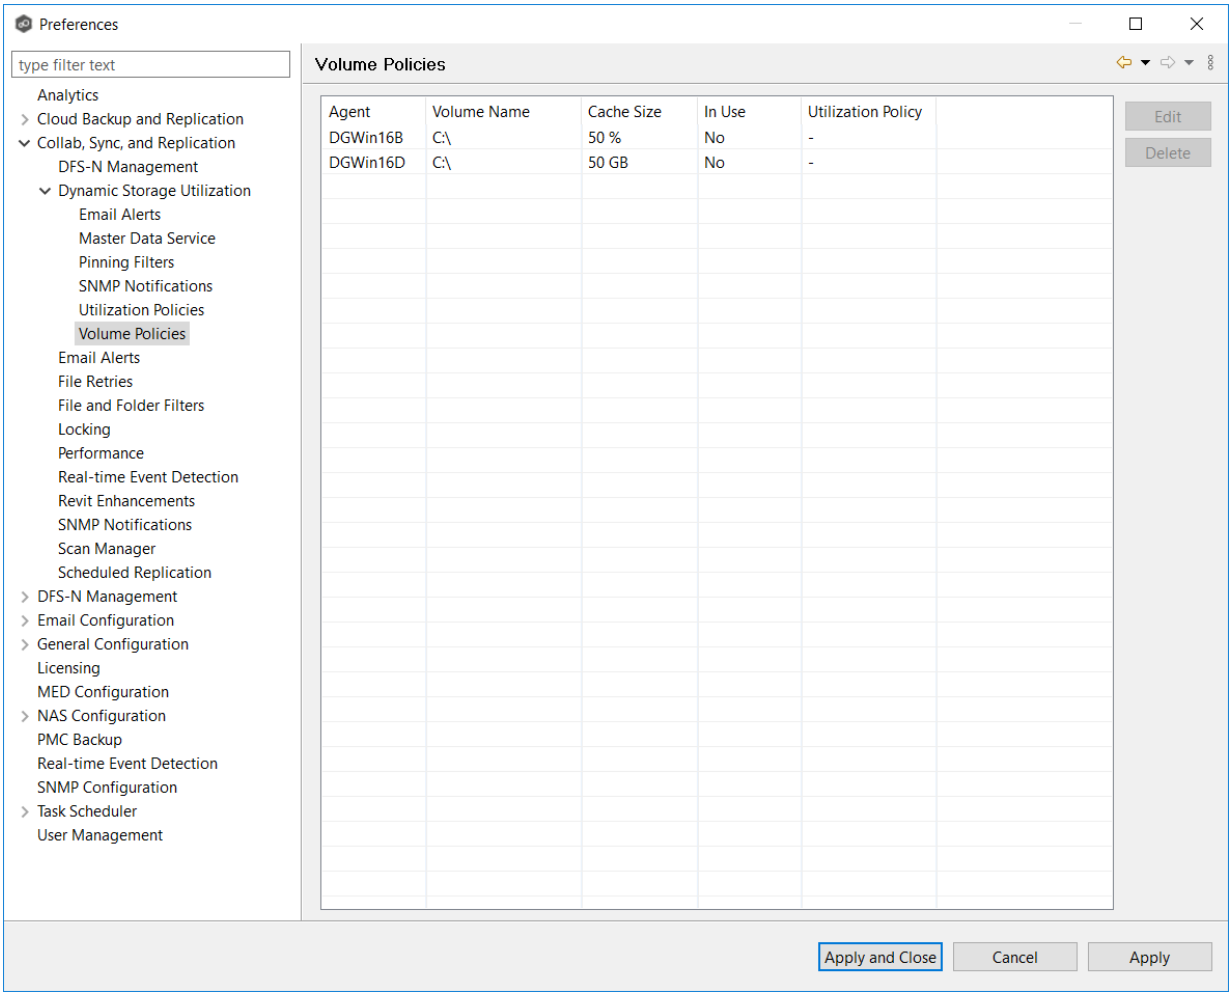

4. Select the policy you want to modify, and then click **Edit**.

5. If you have created a DSU email alert or an SNMP notification but not saved it, you will be prompted to save it.

6. If you selected a policy currently being used in a job, the **In Use Configuration** message appears.

7. Click **OK** to close the dialog.

   The **Edit Volume Policy** dialog appears.

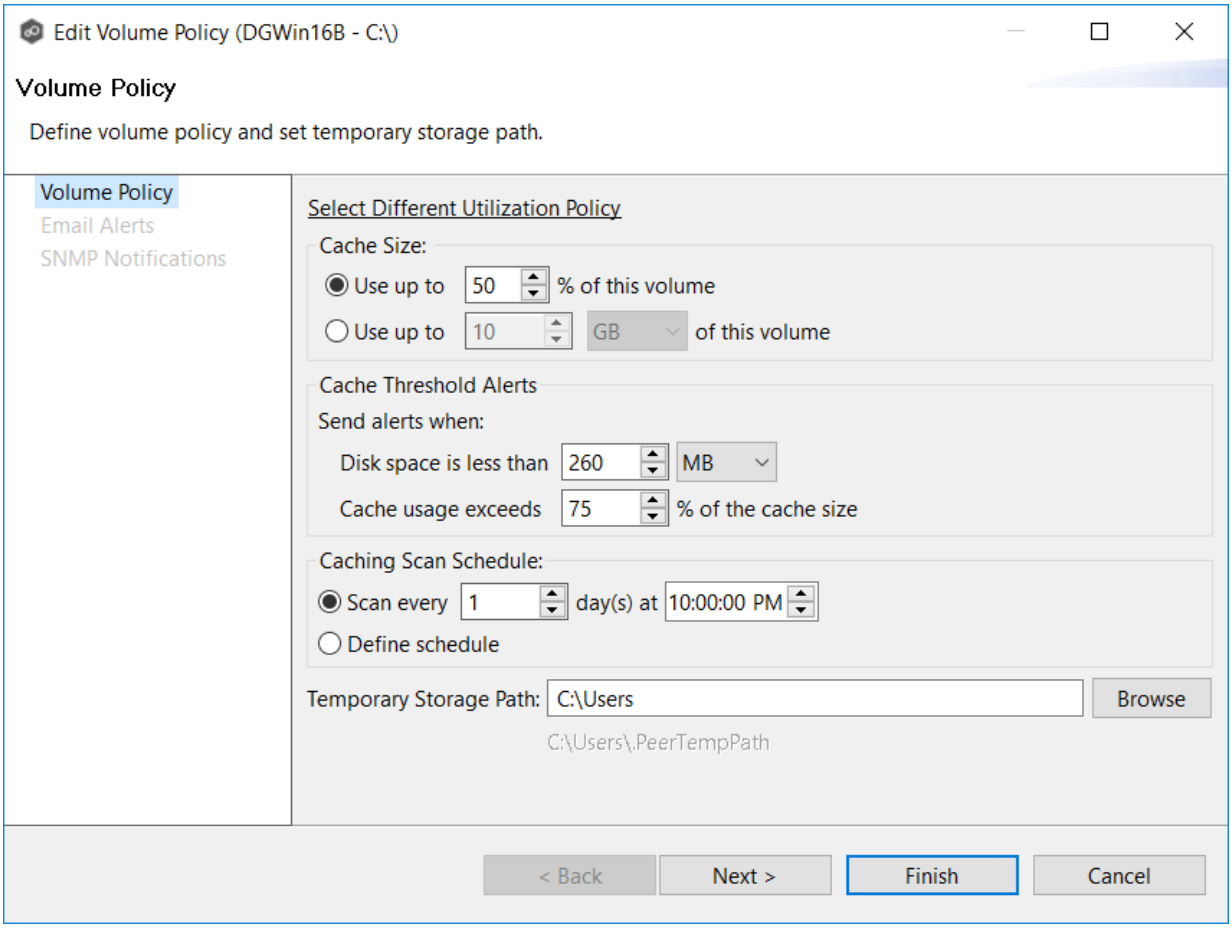

8. If you want to associate a different utilization policy with this volume policy, click **Select Different Utilization Policy** link, select the policy, and then click **OK**.

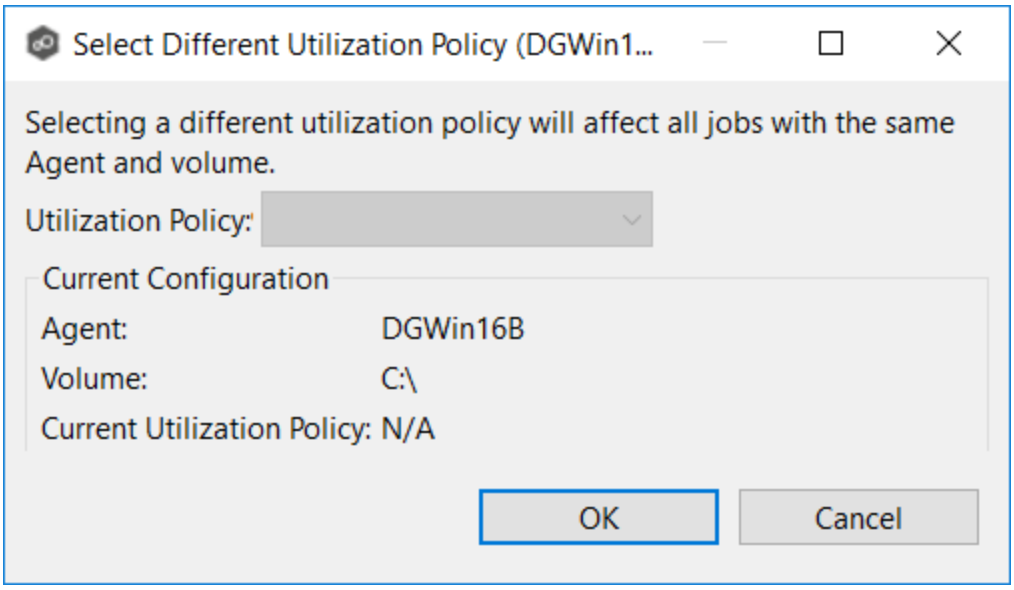

9. In the **Cache Size** section, choose an option for setting the cache size.

   - Use up to X % of this volume
   - Use up to X size of this volume

10. In the **Cache Threshold Alerts** section, enter values for automatic alerts about free disk space and cache usage.

    Peer Management Center will automatically display alerts in the **Alerts** tab and send alerts via email (if configured) when:

    - The amount of free disk space on the volume falls below the specified value. For example, if a 1TB volume has 500MB of free space and the threshold is set to 512MB, an alert will be sent.
    - Cache usage on the volume exceeds the specified percentage of the cache size. For example, if the cache size is set to 80%, equating to to 750 GB, DSU will start sending alerts when it has used 600 GB.

11. In the **Caching Scan Schedule**, set the frequency that the volume should be scanned to optimize the balance of fully hydrated vs stubbed files.

    This scan can be run daily at a specified time or you can define a more customized schedule.

12. In the **Temporary Storage Path** field, enter a path or browse to the location you want to be used as temporary storage space.

The temporary storage space will be used to store the content of stub files as they are are being rehydrated. The content of files undergoing rehydration are referred to as **file blocks**. File blocks are fixed-length chunks of data that are read into memory when requested by an application. DSU will create a subfolder named **.PeerTempPath** under the location that you specify and temporarily store the file blocks in that subfolder.

For optimal performance, we recommend that the temporary storage space be on the same volume as the watch set. If that is not possible, it should be on a high performance disk.

13. (Optional) Select one or more email alerts to be associated with this volume policy, and then click **Next**.

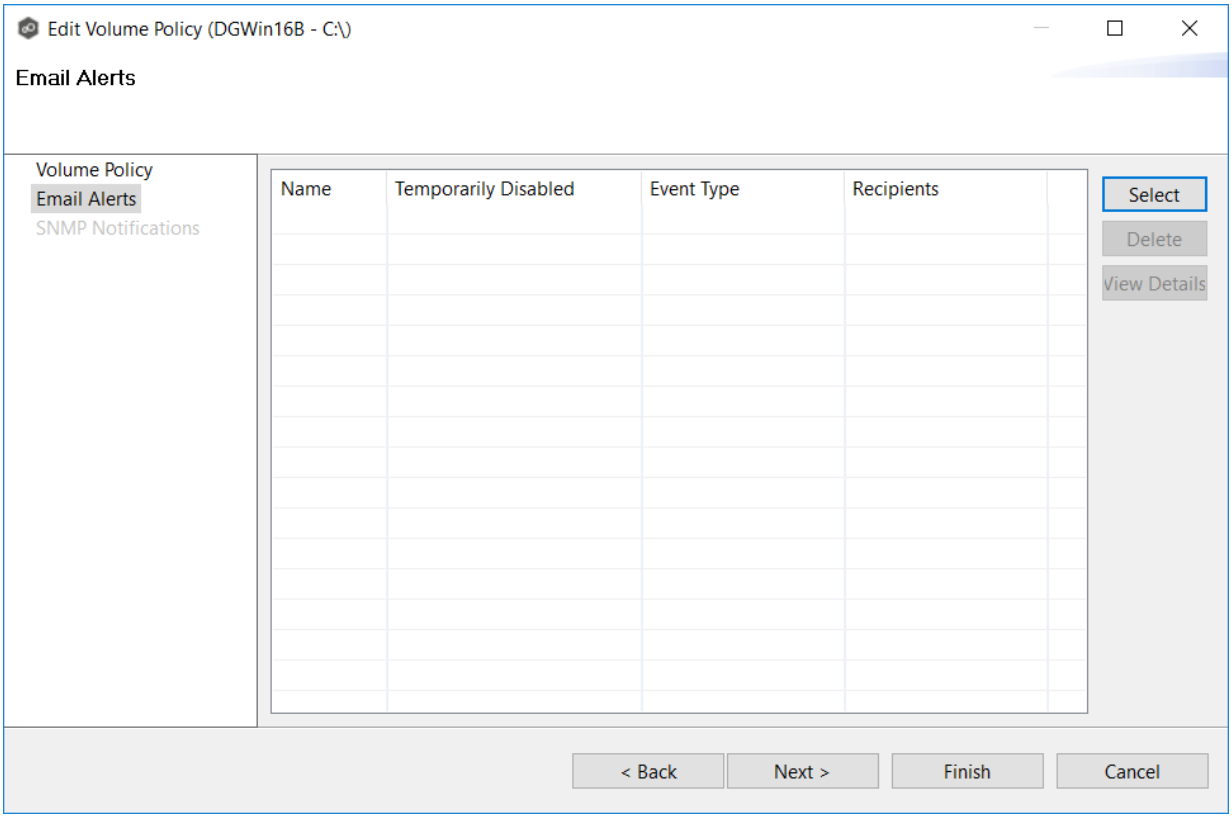

14. (Optional) Select one or more SNMP notifications to be associated with this volume policy.

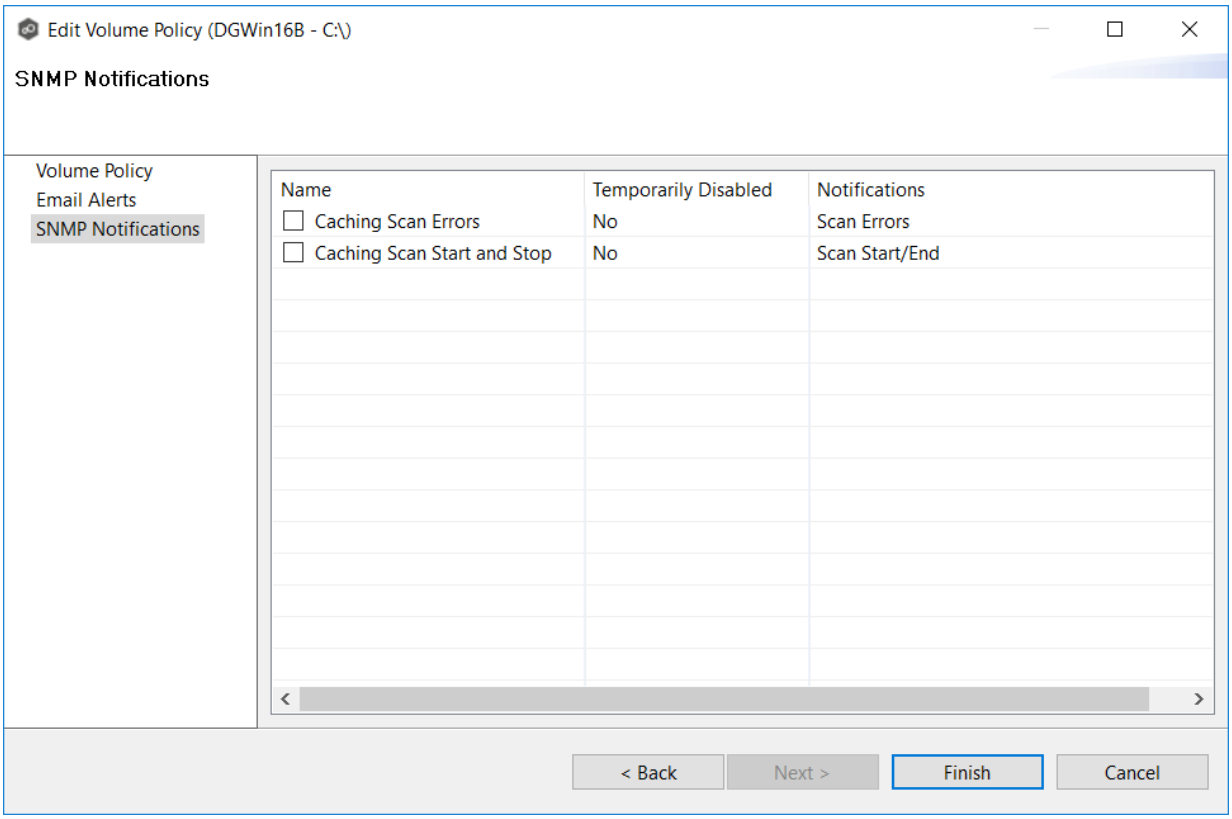

15. Click **Finish**.

    The following message is displayed:

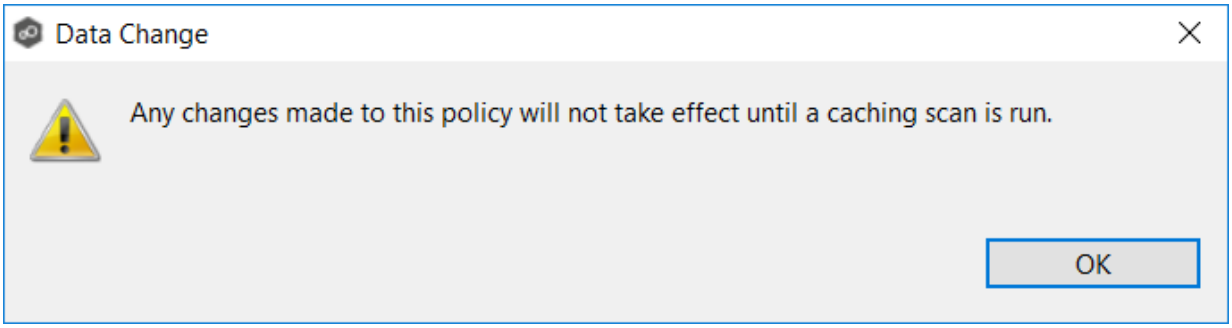

16. Click OK.

    The revised volume policy is listed in the **Volume Policies** table. It will be used after jobs using the policy are restarted.

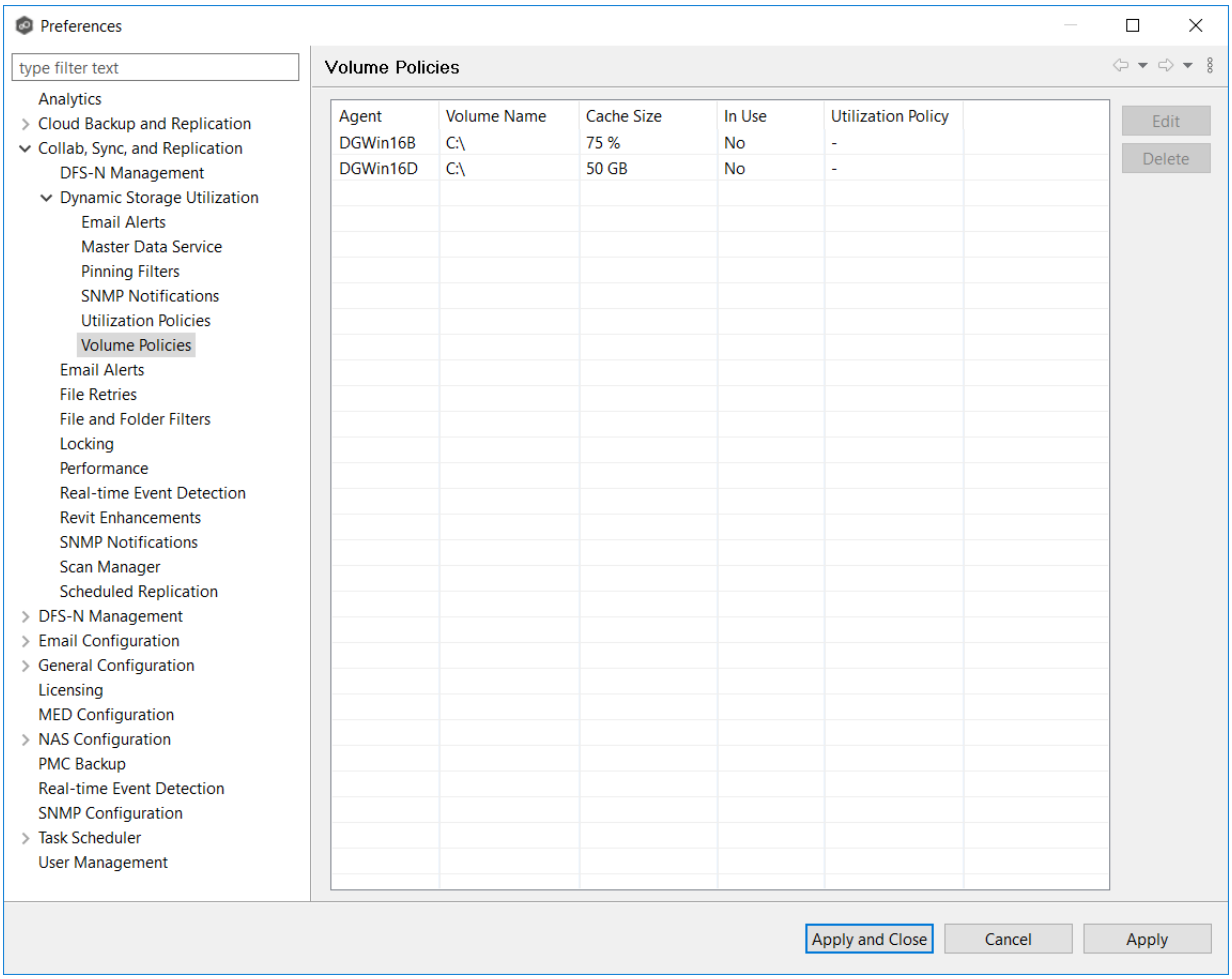

17. Click **Apply and Close** or **Apply**.

Email Alerts

When you create a job, you can select existing email alerts to apply to the job or you can create new email alerts and apply them to the job. This Preferences page lists the existing email alerts. From this page, you can view, create, edit, and delete email alerts. However, you cannot edit or delete an email alert while it is applied to a job. See Email Alerts in the Basic Concepts section for more information about email alerts.

Tip: If you are performing maintenance, you can temporarily disable an email alert by clicking in the **Temporarily Disabled** column in the **Email Alerts** table and selecting **Yes**. No email alerts will be sent for that type of alert until you reenable the alert.

Note: An SMTP email connection must be configured before email alerts can be sent. See Email Configuration for information about configuring SMTP email settings.

To create an email alert:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Email Alerts**.

   Any existing email alerts are listed in the **Email Alerts** table.

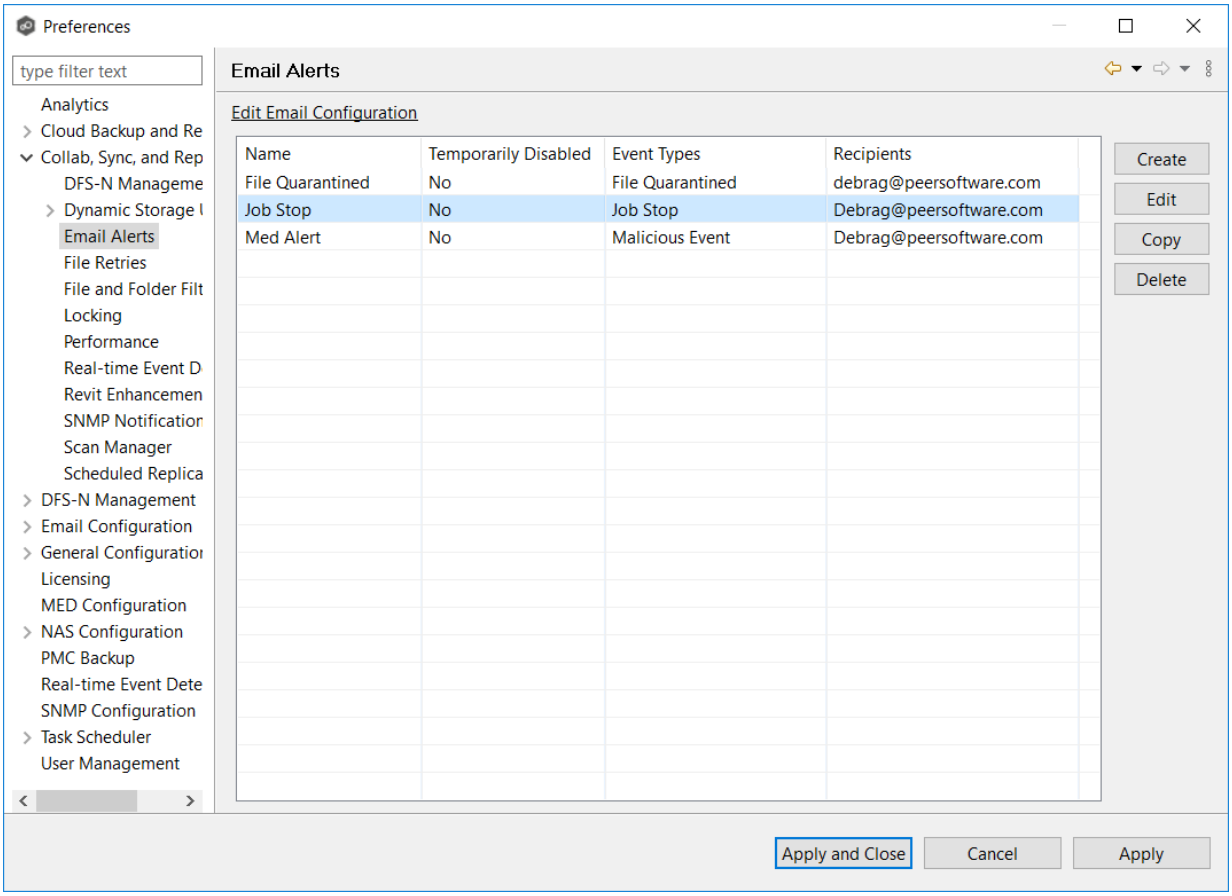

3. Click **Create**.

   The **Create Email Alert** dialog appears.

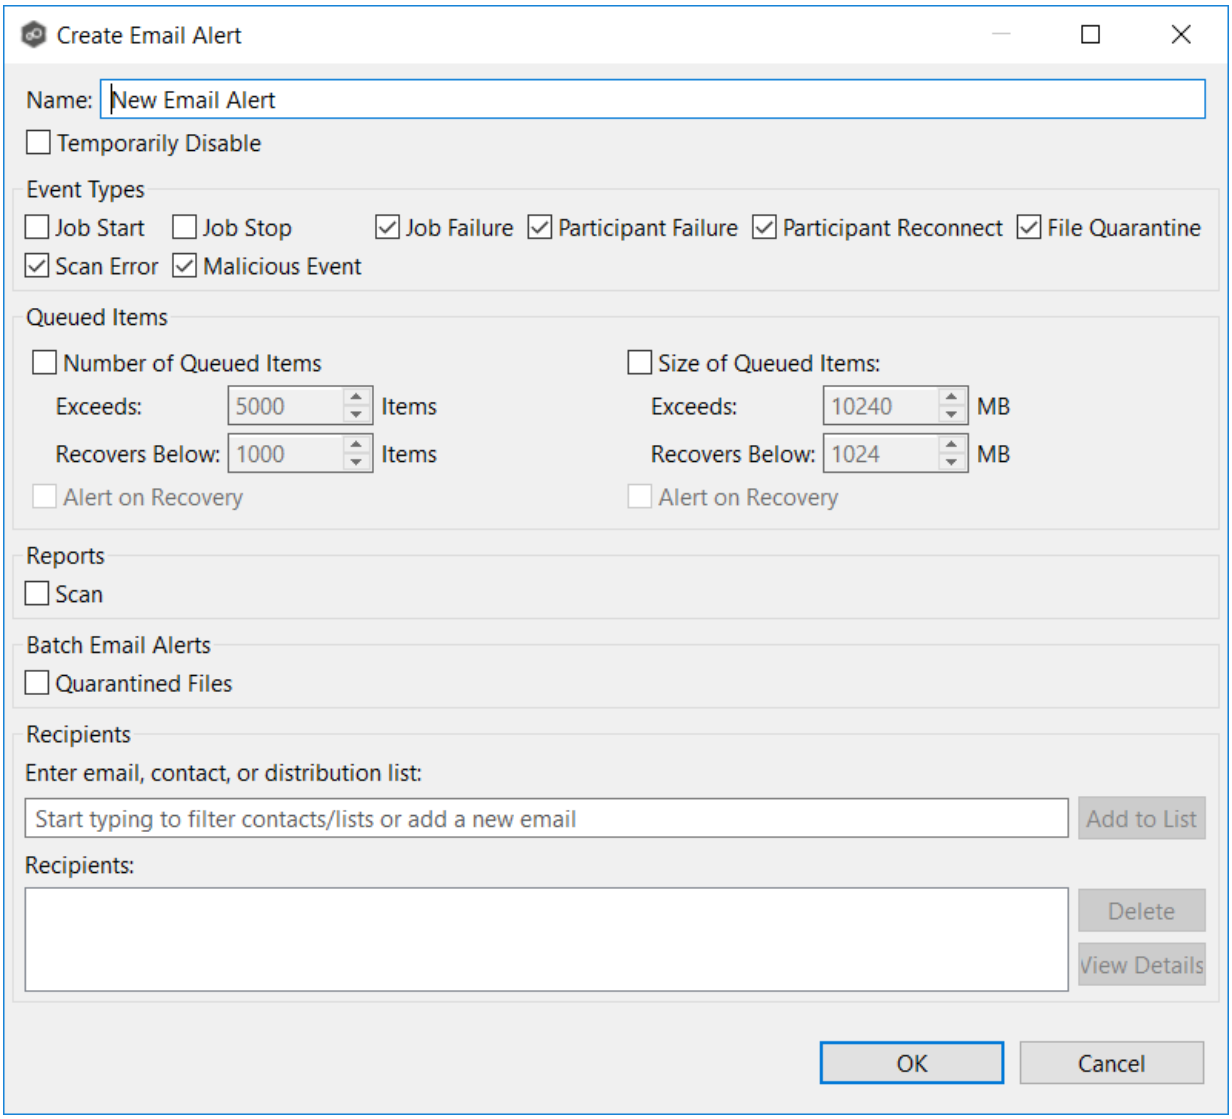

4. Enter a name for the alert.

5. Select the types of events that will trigger an email alert to be sent.

Event Type	Description
Job Start	Sends an alert when the job starts.
Job Stop	Sends an alert when the job stops.
Job Failure	Sends an alert when the job is aborted because of lack of quorum due to one or more failed participants.
Participant Failure	Sends an alert when a participant timeout occurs, and the participant is taken out of the running job.
Participant Reconnect	Sends an alert when a participant reconnects to the job and the job resumes with the reconnected participant.
File Quarantine	Sends an alert when a file is marked as quarantined because a file conflict was not able to be resolved.
Scan Error	Sends an alert when an error occurs during the initial synchronization process.
Malicious Event	Sends an alert when Peer Malicious Event Detection (MED) detects potentially malicious activity. For more information, see MED Configuration.

6. Select options in the **Queued Items** section.

Option	Description
Number of Queued Items section	Select this checkbox if you want alerts to be sent regarding number of items in the queue. This is useful to notify you about when there is a queue backlog potentially due to latency issues. If you select this checkbox, you need to enter two values that work in tandem: • **Exceeds X Items** - Enter the highest number of queued items before an email alert is sent. The default value is 5000. • **Recovers Below X Items** - Enter a value. The default value is 1000. An alert is sent the first time that the **Queued Items** counter has items is greater than the value set in **Exceeds X Items**. The counter's value is displayed in the **Queued Items** column in the in the **Collab, Sync, and Repl Summary** view. The counter's value is a combination of the **Real-time** and **File Sync** queues. Another alert will not be sent until the **Queued Items** counter has dropped below the **Recovers Below x Items** value and then exceeds the **Exceeds X Items** value again. This prevents multiple or redundant alerts from being sent.
Alert on Recovery	Select this option if you want an alert to be sent when the number of queued items has fallen below the **Recovers Below** value.
Size of Queued Items section	Select this if you want alerts to be sent based on the total data size of queued items for a job. This is useful to notify you about when there is a queue backlog potentially due to bandwidth issues. If you select this checkbox, you need to enter two values that work in tandem: • **Exceeds X MB** - Enter the highest number of queued items before an email alert is sent. The default value is 10240 MB. • **Recovers Below X MB** - Enter a value. The default value is 1024 MB. An alert is sent the first time that the **Pending Bytes** for a job has items is greater than the value set in **Exceeds X MB**. The counter's value is displayed in the **Pending Bytes** column in the **Collab, Sync, and Repl Summary** view. Another alert will not be sent until the **Pending Bytes** counter has dropped below the **Recovers Below x MB** value and then exceeds the **Exceeds X MB** value again. This prevents multiple or redundant alerts

7. Select the **Scan** checkbox in the **Reports** section if you want scan statistics emailed to you after a scan has completed.

8. Select the **Quarantined Files** checkbox in the **Batch Email Alerts** section if you want email alerts about quarantined files sent to you in batches.

   The default is a maximum of 1000 alerts with a quiet period of 60 seconds. To change the number of alerts and quiet period time, modify the values for **Batch Email Alerts for Quarantined Files** in Email Configuration.

9. Enter alert recipients, and then click **Add to List**.

   The recipients are listed in the **Recipients** field.

10. Click **OK**.

    The new email alert is listed in the **Email Alerts** table and can now be applied to jobs.

11. Click **Apply and Close** or **Apply**.

    The new alert is listed in the **Email Alerts** table and can now be applied to jobs.

File Retries

File retries settings enable you to configure the frequency of attempts and the maximum number of attempts. These settings apply to all File Collaboration, File Replication, and File Synchronization jobs. For more information about file retries, see Conflicts, Retries, and Quarantines.

To modify the file retries settings:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **File Retries**.

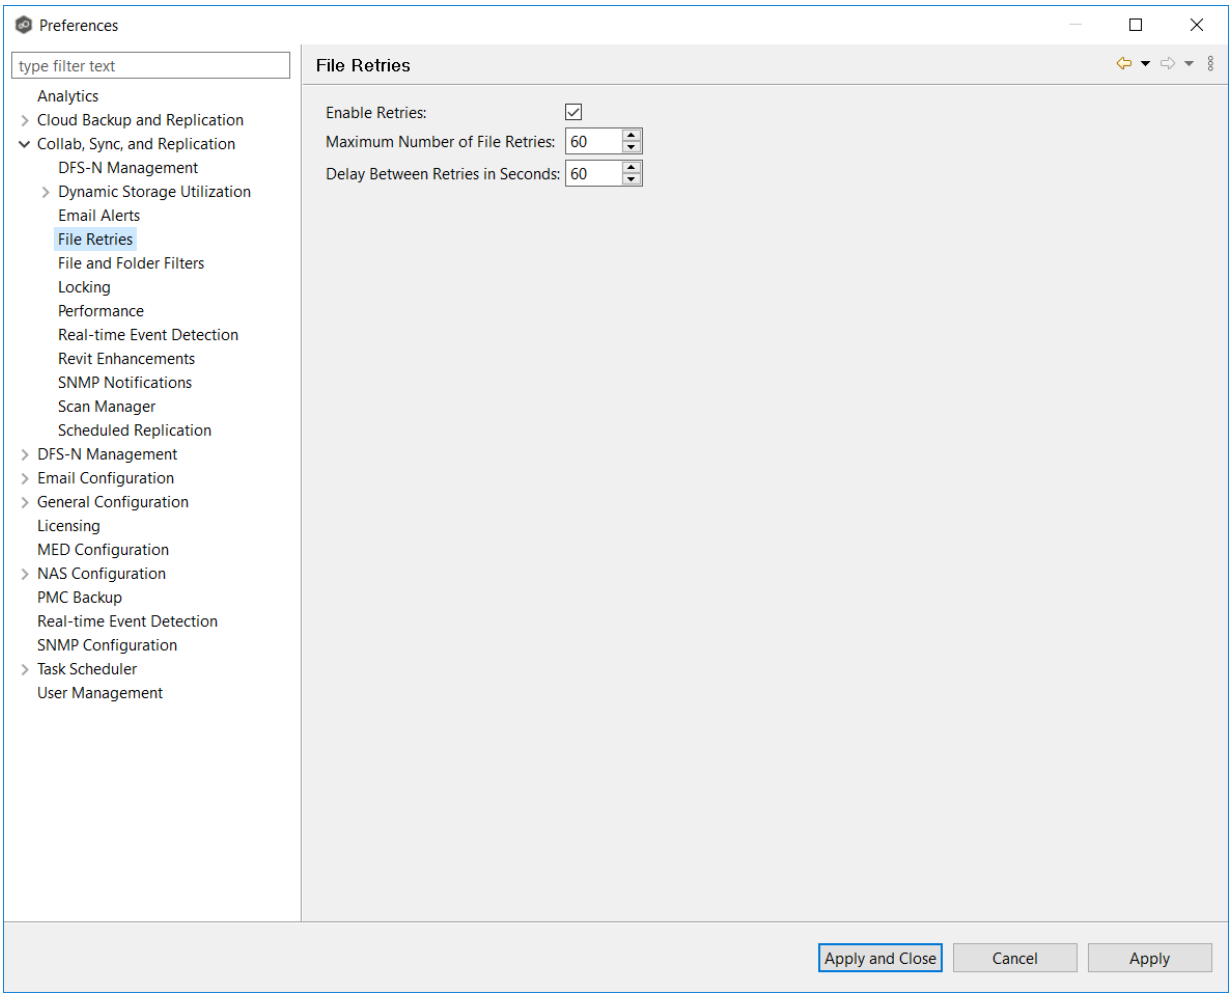

3. Modify the settings as needed.

Setting	Description
Enable Retries	Select this checkbox to enable the retry of failed file transfers. If this option is not enabled, files that would have been candidates for retries will be automatically quarantined.
Maximum Number of File Retries	Enter the maximum number of attempts to retry a failed file transfer before it is quarantined.
Delay Between Retries in Seconds	Enter the number of seconds to wait between retries of a failed file transfer.

4. Click **Apply and Close** or **Apply**.

File and Folder Filters

When you create a File Collaboration, File Synchronization, or File Replication job, you can select existing file and folder filters to apply to the job or you can create new file filters and apply them to the job. This Preferences page lists the existing file and folder filters. From this page, you can view, create, edit, and delete file filters. However, you cannot edit or delete a file filter while it is applied to a job. See File and Folder Filters in the Basic Concepts section for more information about file and folder filters.

To create a file and folder filter:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **File and Folder Filters**.

   Any existing file filters are listed in the **File and Folder Filters** table.

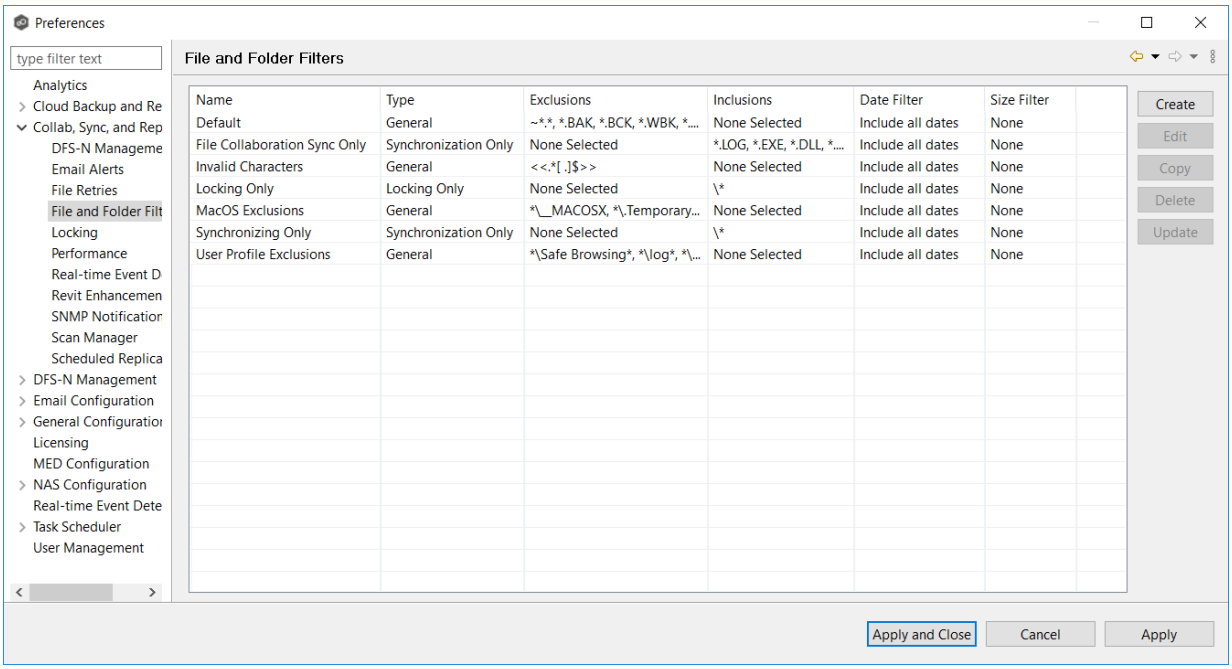

3. Click **Create**.

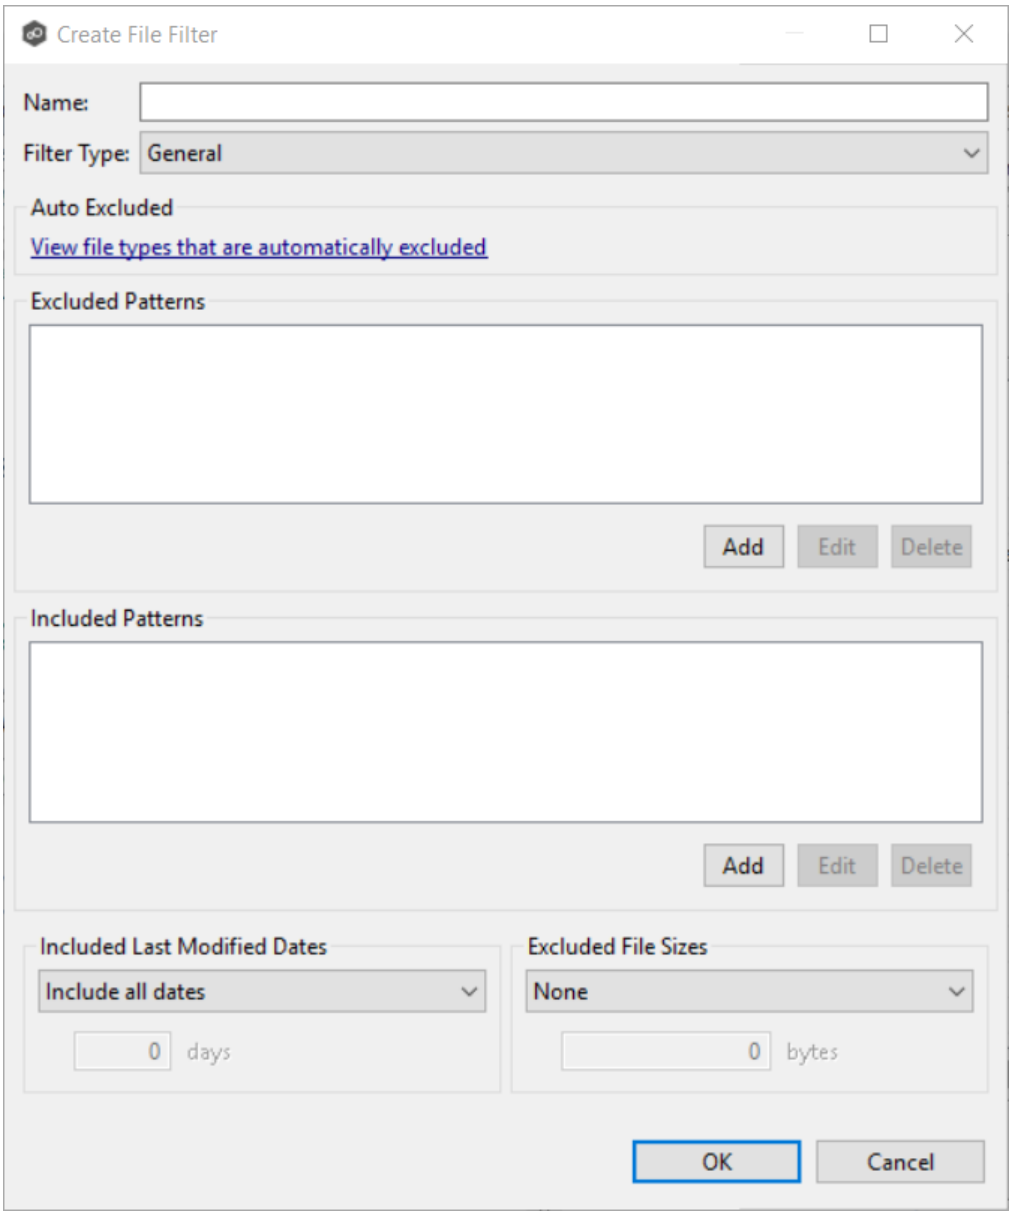

4. Enter a unique name for the filter.

5. Select the filter type.

6. (Optional) Click **Add** to enter a filter pattern for files that you want excluded from the job. Repeat to add more filter patterns.

   See Defining Filter Patterns for information about filter patterns.

7. (Optional) Click **Add** to enter a filter pattern for files that you want included in the job. Repeat to add more filter patterns.

8. (Optional) Select a value for Included Last Modified Dates.

   Note: A filter cannot use **Included Last Modified Dates** in conjunction with any other excluded or included patterns.

9. (Optional) Select a value for Excluded File Sizes. Note: This cannot be combined with any other filter criteria

   Note: A filter cannot use **Excluded File Sizes** in conjunction with any other excluded or included patterns.

10. Click **Apply and Close** or **Apply**.

    The new file filter is listed in the **File and Folders Filters** table and can now be applied to jobs.

Locking

An option is available to mark certain file types as non-collaborative, changing the way locks on the specified file types are handled. These settings apply to all File Collaboration, File Synchronization, and File Replication jobs. These settings are critical for certain file types so that the job can correctly read these files, ensuring that managed file types are synchronized in a consistent and usable state.

To modify the locking settings:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Locking**.

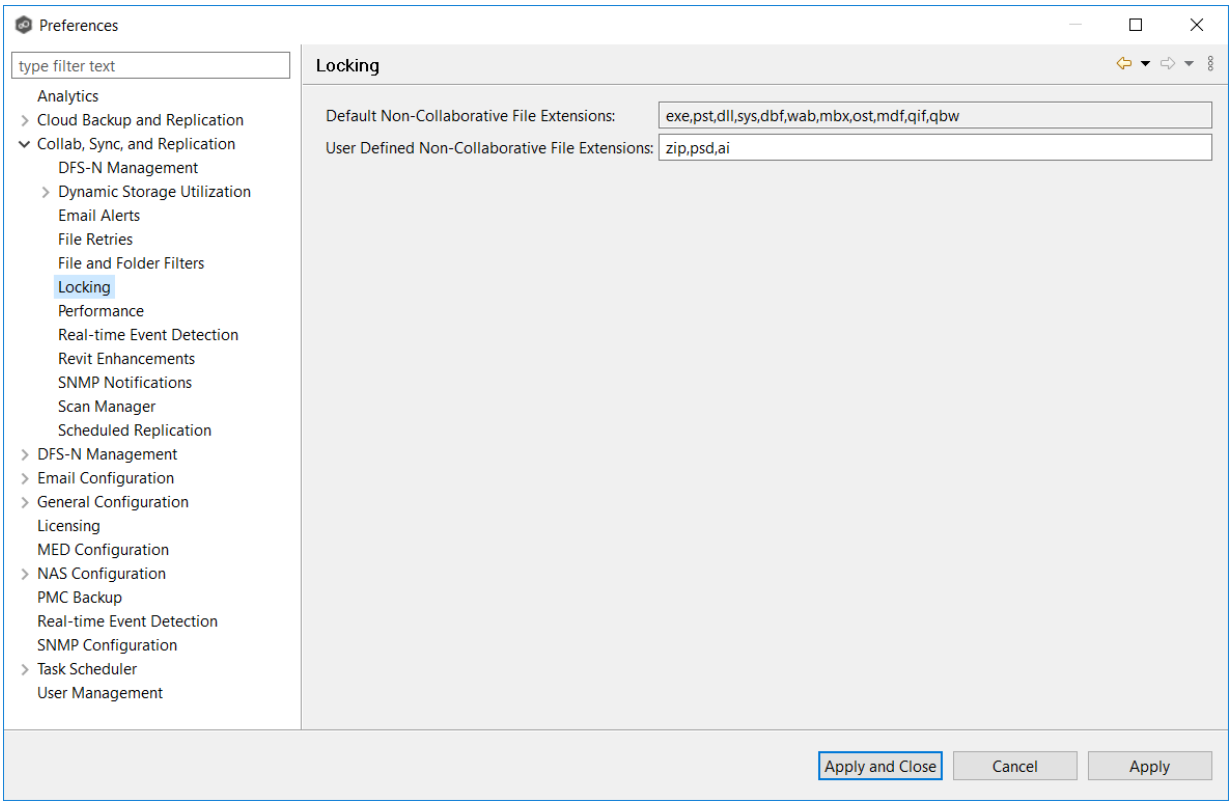

3. Modify the options as needed.

Option	Description
Default Non-Collaborative File Extensions	Non-editable. Displays the default, comma-separated list of file extensions of non-collaborative file types (e.g., database files). Write access to source files of these types is denied while the files are being synchronized.
User Defined Non-Collaborative File Extensions	Displays an editable, comma-separated list of file extensions of non-collaborative file types (e.g., database files). Write access to the source files of these types is denied while the files are being synchronized.

4. Click **Apply and Close** or **Apply**.

Performance

To customize the performance settings of File Collaboration, File Synchronization, and File Replication jobs:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Performance**.

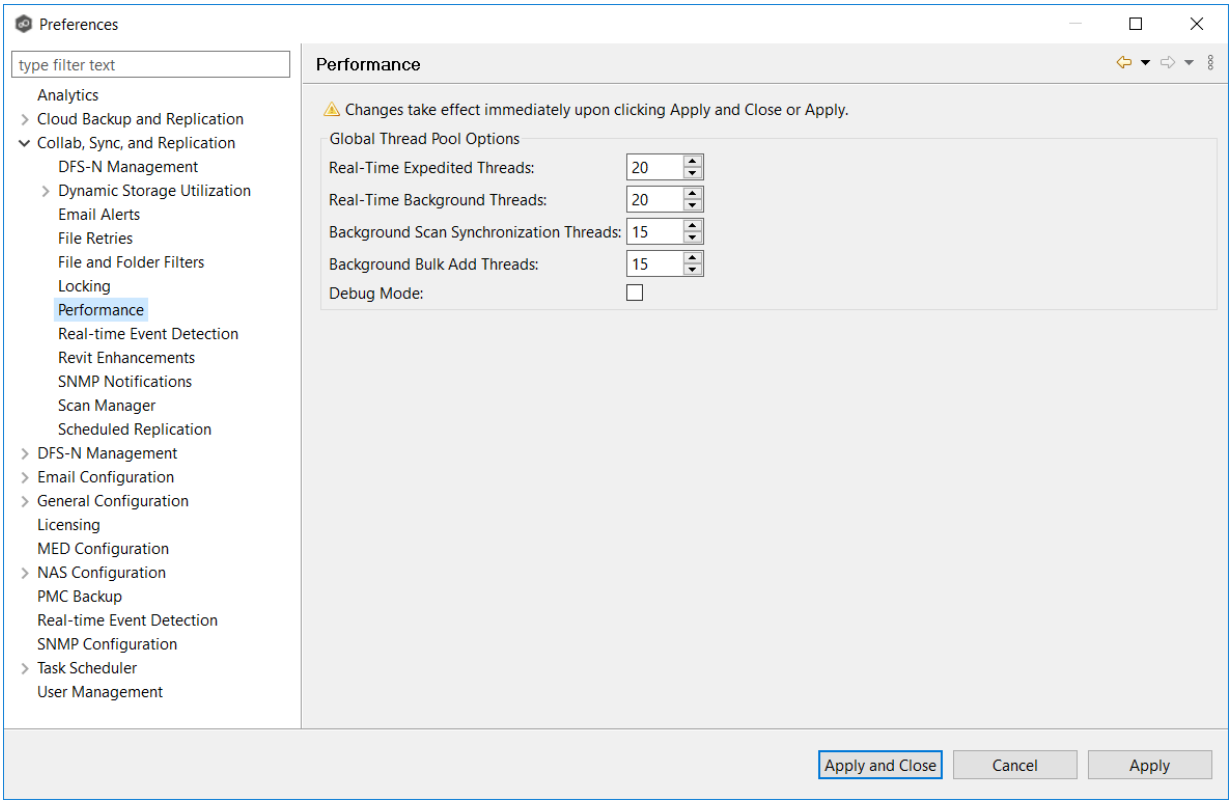

3. Modify the options as needed.

Option	Description
Real-Time Expedited Threads	Enter the maximum number of threads for controlling file locking and renames.
Real-Time Background Threads	Enter the maximum number of threads for controlling the replication of file content.
Background Scan Synchronization Threads	Enter the maximum number of threads for processing the differences found by background scans.
Debug Mode	Select to enable debug mode for the various types of threads.

4. Click **Apply and Close** or **Apply**.

Real-time Event Detection

To modify the File Collaboration real-time detection settings:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Real-time Event Detection**.

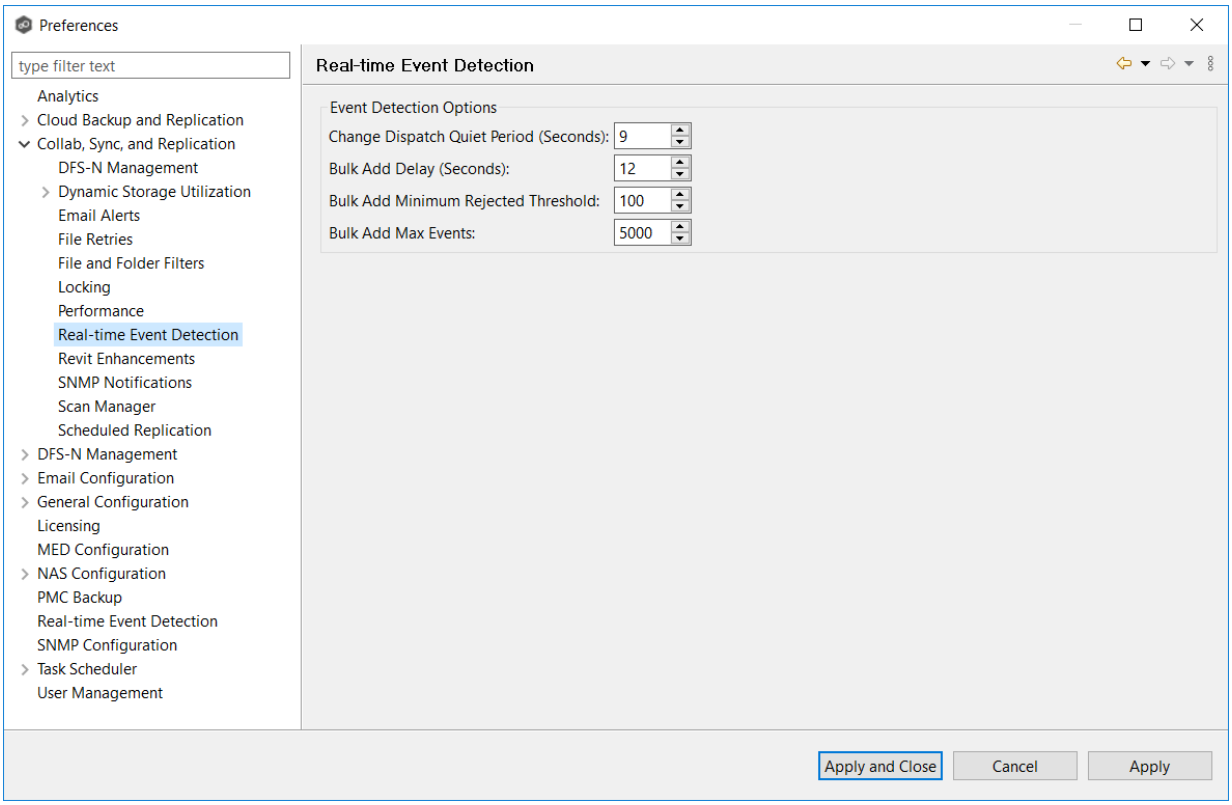

3. Modify the options as needed.

Option	Description
Change Dispatch Quite Period (Seconds)	The number of seconds to wait before acting on a file modification, rename, or delete.
Bulk Add Delay (Seconds)	Controls when the bulk add logic is triggered. This is used to help deprioritize mass copying or adding of files to a directory.
Bulk Add Minimum Rejected Threshold	The minimum number of file adds that must occur within the Bulk Add Delay for bulk add logic to be triggered.
Bulk Add Max Events	The maximum number of file adds to lump together in one batch.

4. Click **Apply and Close** or **Apply**.

Revit Enhancements

Revit Enhancements enable the Expedited Sync Queue for files specified in the Expedited Sync Queue File List.

To set advanced settings for Revit Enhancements:

1. Select **Preferences** from the Window menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Revit Enhancements**.

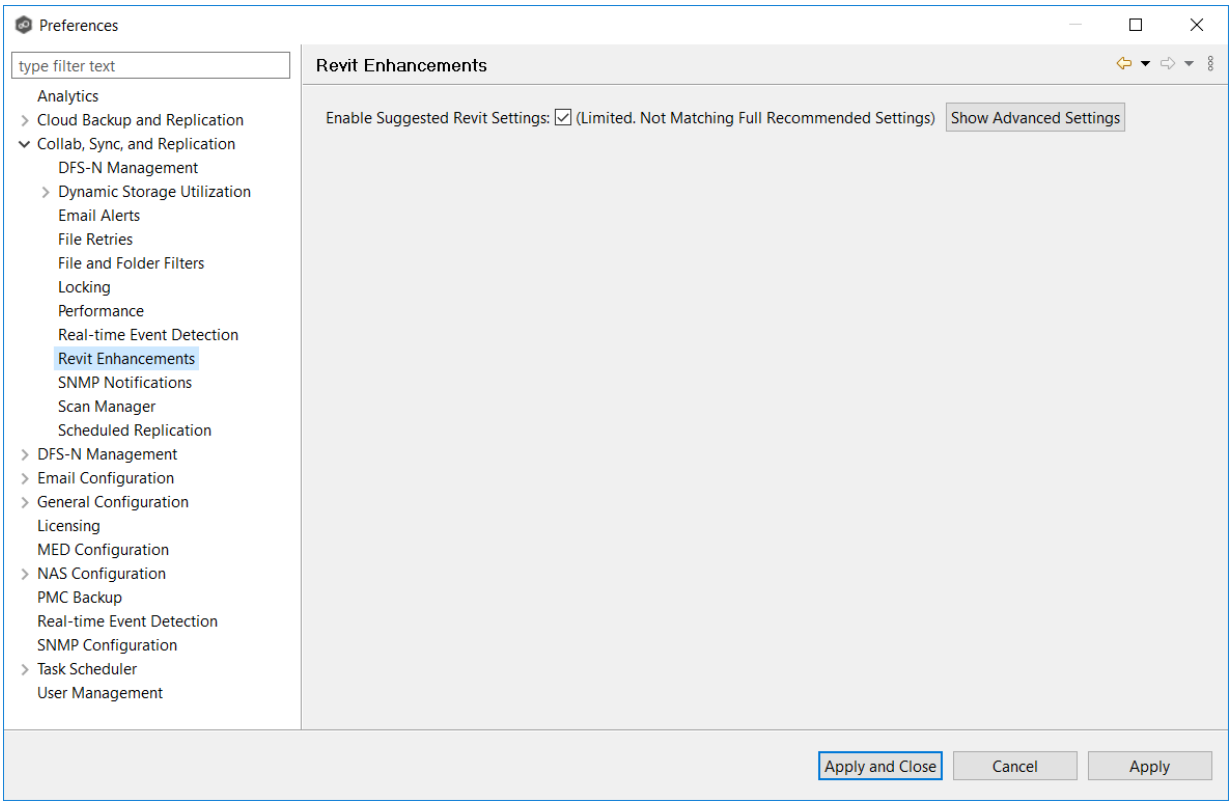

3. Click **Show Advanced Settings**.

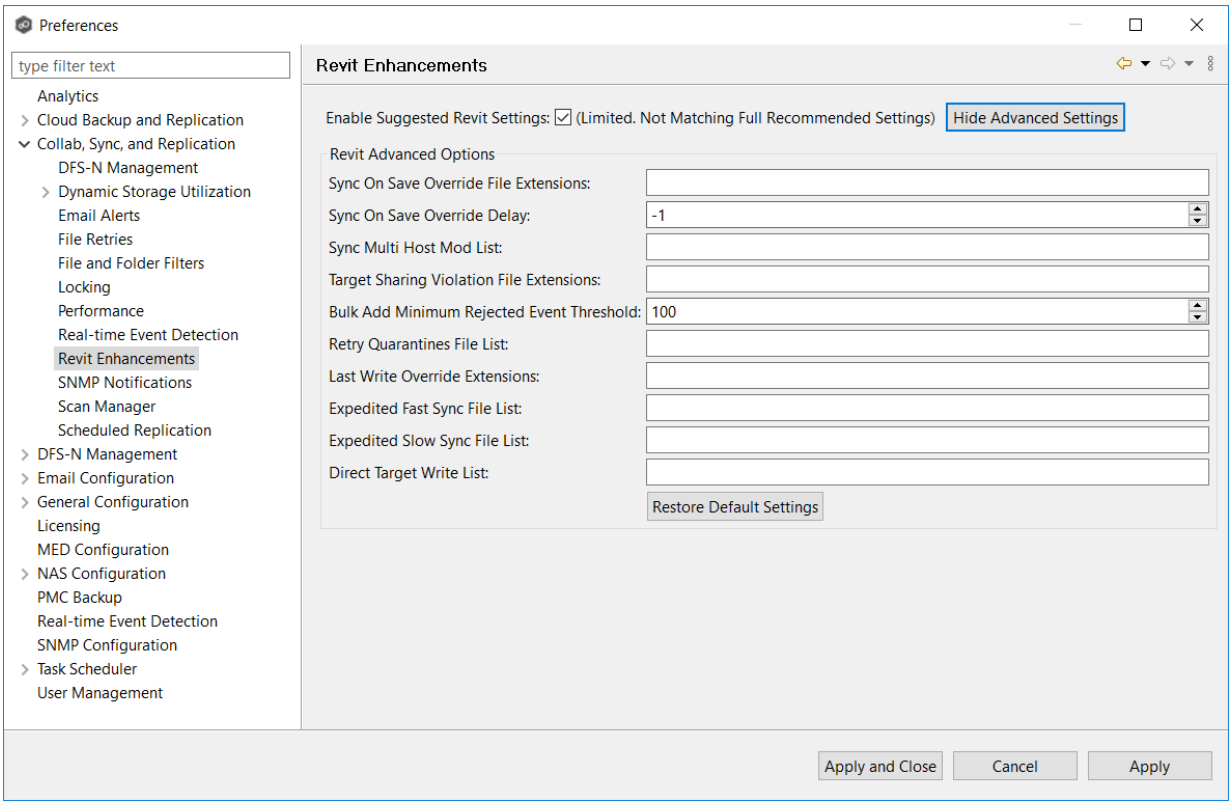

Y

4. Modify the options as needed.

Option	Description
Sync On Save Override File Extensions	Extensions configured here will overwrite the **Sync. On Save** values configured in the interface for the job. In addition, these extensions use the delay value in **Sync On Save Override Delay** setting instead of the delay value configured in the interface. If no delay value is set, it will default to using a one second delay. Extensions configured in this list will still be processed via **Sync. On Save** even if they also exist in the user defined non-collaborative extension list (under the Window > Preferences menu option). Extensions in the normal **Sync. On Save** list that also exist in this list will not be processed.
Sync On Save Override Delay	The **Sync. On Save** delay value in seconds that applies only to the internal list of extensions listed in the **Sync On Save Override File Extension** field.

Option	Description
Sync Multi Host Mod List	Extensions configured here will not be quarantined if they are modified on two hosts simultaneously. The file with the latest modified time stamp will win.
Target Sharing Violation File Extensions	This is an option to retry setting the target lock when receiving error code 32 for the specified list of extensions. This may be useful for file types such as .one (OneNote), .rvt (Revit), and .dat (associated Revit files) that don't sustain a handle when the user has the file open.
Bulk Context Minimum Rejected Event Threshold	The number of bulk add files that can process immediately before batching the remainder of the files and process them in a single thread.
Retry Quarantine File List	Quarantined files that are in this list will be automatically removed and flagged as unsynchronized and will be retried every second after a delay period (delay is configured by **fc.retryQuarantinesDelay**). Any change event that is detected for the files will trigger a scan of the files where the newest file will win. This list can contain file names (wperms.dat, eperms.dat, requests.dat, deltas.dat, users.dat) or extensions (*.dat,*.abc).
Last Write Override Extensions	Act on every write event performed on these extensions instead of waiting for the last write event prior to the closing of a file.
Expedited Fast Sync File List	Access events and transfer events will be expedited for the list of extension or files in this list.
Expedited Slow Sync File List	Access events received for files or extension in this list will be expedited. Transfers will go through a slow priority queue.
Direct Target Write List	List of files to be updated without the use of a temp file. This list can contain file names (wperms.dat, eperms.dat, requests.dat, deltas.dat, users.dat") or extensions.

5. Click **Apply and Close** or **Apply**.

SNMP Notifications

When you create a job, you can select an existing SMNP notification to apply to the job or you can create a new notification and apply it to the job. You cannot modify or delete an SMNP notification while it is applied to a job. See SMNP Notifications in the Basic Concepts section for more information about SMNP notifications.

To create an SMNP notification:

1. From the **Window** menu, select **Preferences**.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **SNMP Notifications**.

   Any existing SNMP notifications are listed in the **SNMP Notifications** table.

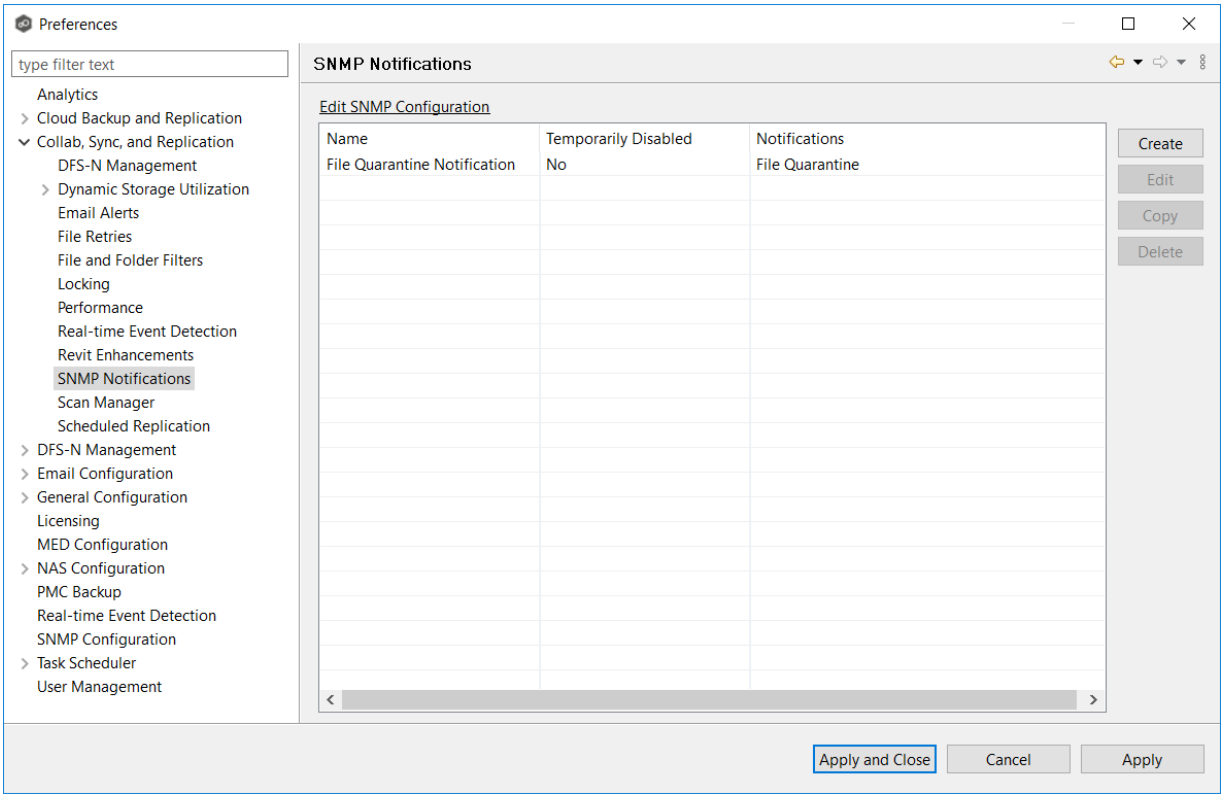

3. Click the **Create** button.

   The **Create SNMP Notification** dialog appears.

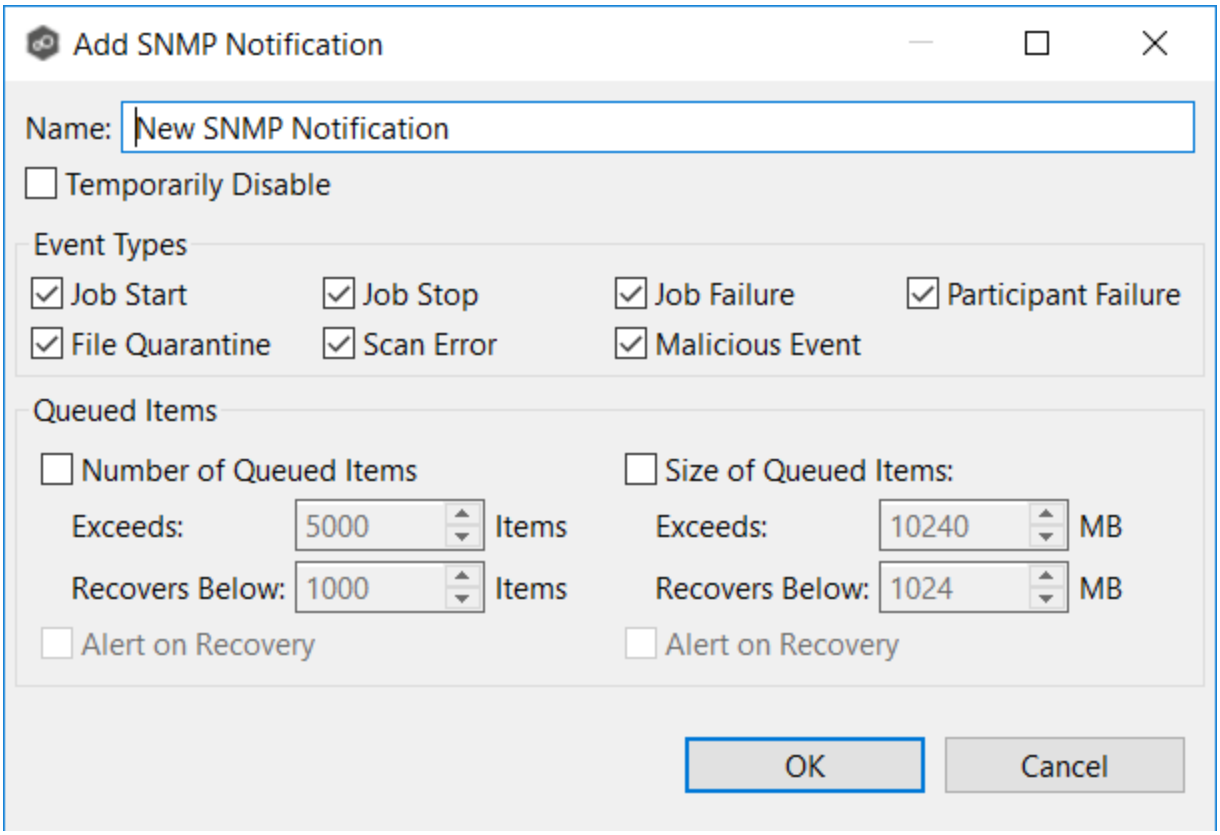

4. Select the types of events that will trigger the generation of an SNMP trap:

Event T	Description
Job Start	Sends a notification when the job starts.
Job Stop	Sends a notification when the job stops.
Job Failure	Sends a notification when the job is aborted because of lack of quorum due to one or more failed participants.
Participant Failure	Sends a notification when a participant timeout occurs, and the participant is taken out of session.
Participant Reconnect	Sends a notification when a participant reconnects to the job and the job resumes with the reconnected participant.
File Quarantine	Sends a notification when a file is marked as quarantined because a file conflict was not able to be resolved.
Scan Error	Sends a notification when an error occurs during the initial synchronization process.
Malicious Event	Sends a notification when Peer MED detects potentially malicious activity. For more information, see MED Configuration.

5. Select options in the **Queued Items** section.

Option	Description
Number of Queued Items section	Select this checkbox if you want alerts to be sent regarding number of items in the queue. This is useful to notify you about when there is a queue backlog potentially due to latency issues. If you select this checkbox, you need to enter two values that work in tandem: • **Exceeds X Items** - Enter the highest number of queued items before an email alert is sent. The default value is 5000. • **Recovers Below X Items** - Enter a value. The default value is 1000. An notification is sent the first time that the **Queued Items** counter has items is greater than the value set in **Exceeds X Items**. The counter's value is displayed in the **Queued Items** column in the in the **Collab, Sync, and Repl Summary** view. The counter's value is a combination of the **Real-time** and **File Sync** queues. Another notification will not be sent until the **Queued Items** counter has dropped below the **Recovers Below x Items** value and then exceeds the **Exceeds X Items** value again. This prevents multiple or redundant alerts from being sent.
Alert on Recovery	Select this option if you want a notification to be sent when the number of queued items has fallen below the **Recovers Below** value.
Size of Queued Items section	Select this if you want notifications to be sent based on the total data size of queued items for a job. This is useful to notify you about when there is a queue backlog potentially due to bandwidth issues. If you select this checkbox, you need to enter two values that work in tandem. • **Exceeds X MB** - Enter the highest number of queued items before an email alert is sent. The default value is 10240 MB. • **Recovers Below X MB** - Enter a value. The default value is 1024 MB . An alert is sent the first time that the **Pending Bytes** for a job has items is greater than the value set in **Exceeds X MB**. The counter's value is displayed in the **Pending Bytes** column in the **Collab, Sync, and Repl Summary** view.
	Another alert will not be sent until the **Pending Bytes** counter has dropped below the **Recovers Below x MB** value and then exceeds the **Exceeds X MB** value again. This prevents multiple or redundant alerts from being sent.

6. Click **Apply and Close** or **Apply**.

   The new notification is listed in the **SNMP Notifications** table and can now be applied to jobs.

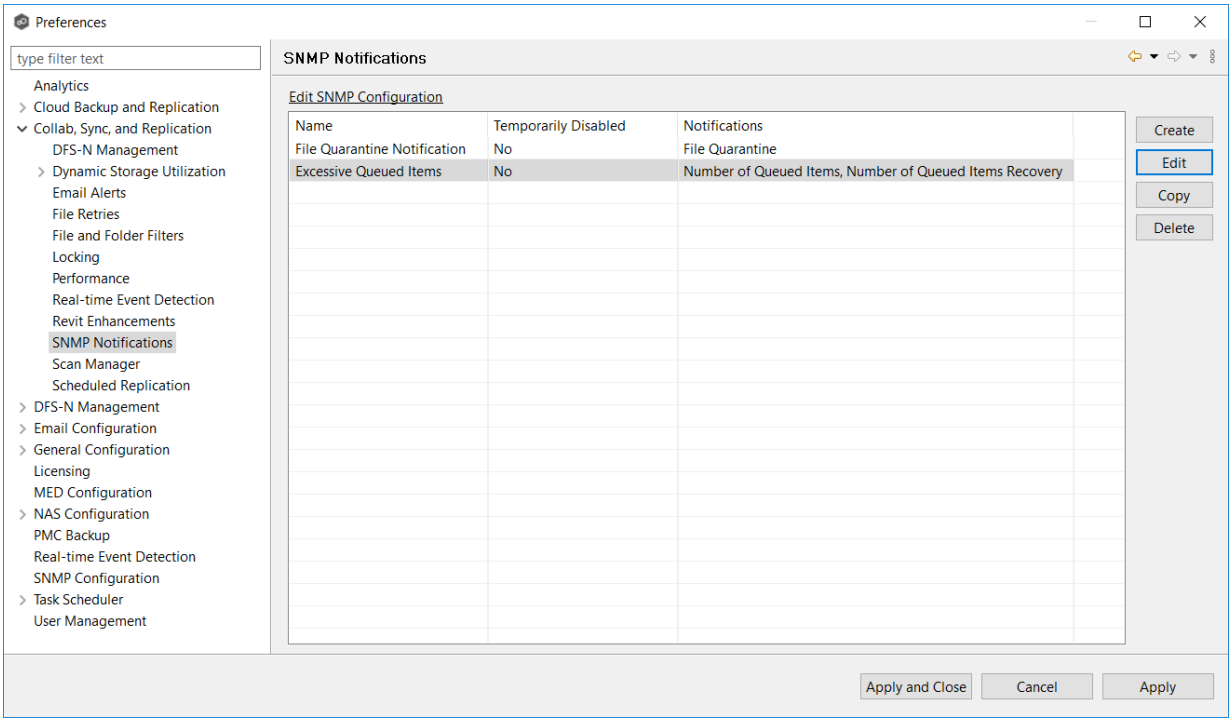

Scan Manager

Several options are available to tune the way scans are performed for File Collaboration, File Synchronization, and File Replication jobs.

To modify the Scan Manager settings for File Collaboration, File Synchronization, and File Replication jobs:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Scan Manager**.

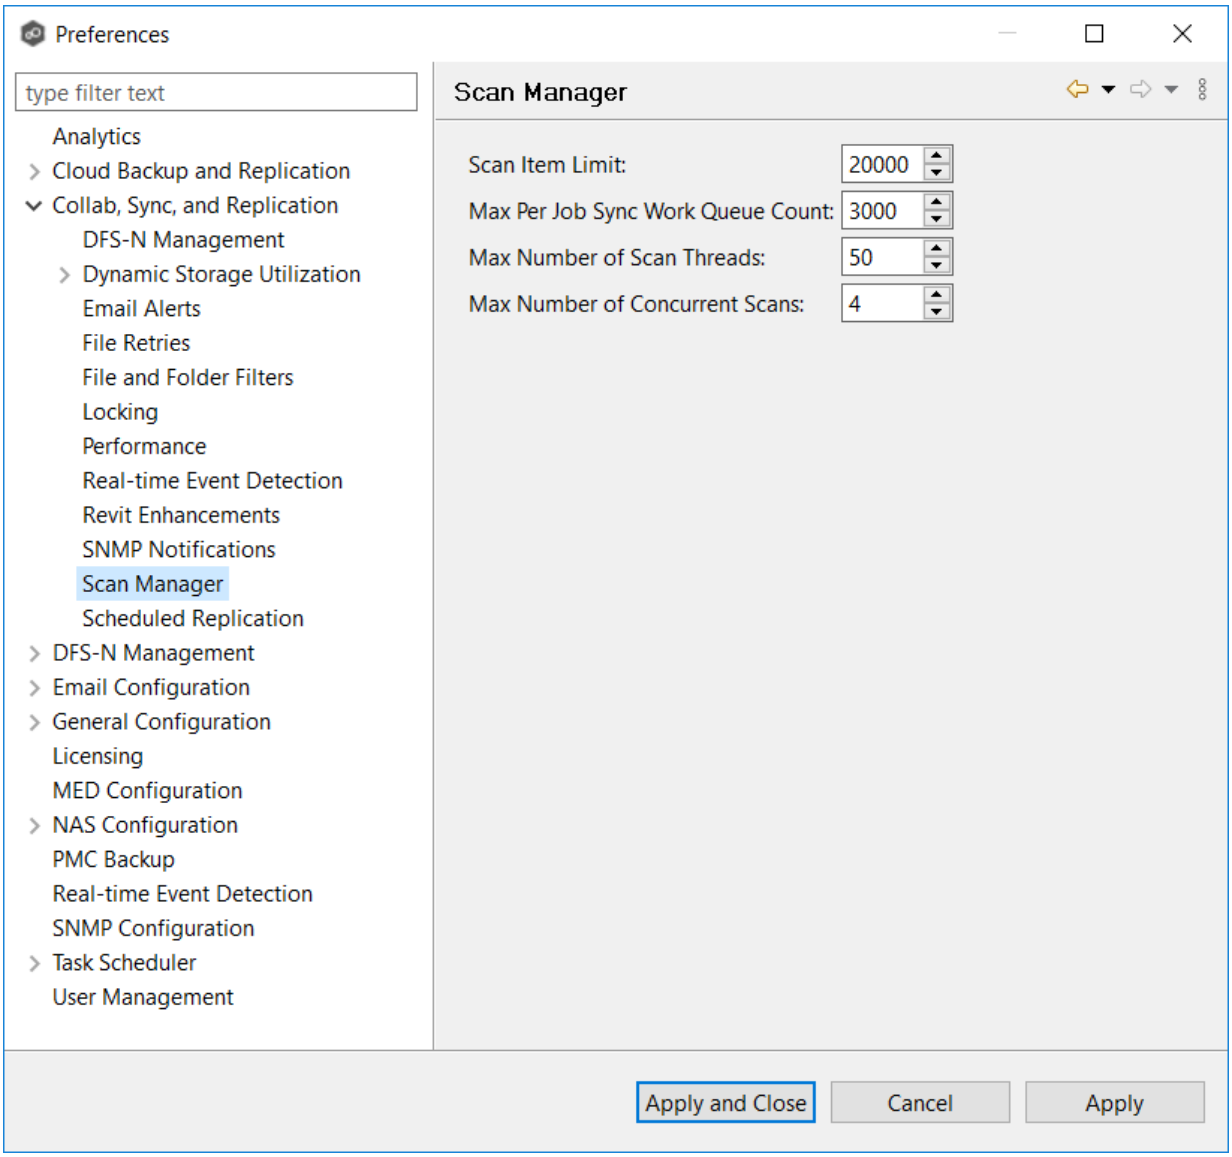

3. Modify the options as needed.

Option	Description
Scan Item Limit	The maximum number of file and folder scan results that are returned in one scan iteration during a job's initial scan. This value is used to constrain the amount of memory used when performing initial scans with a large number of jobs.
Max Sync Work Queue Count	The per job maximum number of pending file synchronization tasks that are queued in memory before pausing the current scan. This value only has an effect on jobs with large numbers of files that must be synchronized during initial synchronization.
Max Number of Scan Threads	The maximum number of threads that can be created to scan folders and files. This number should be set to at least the number of jobs that you are running.
Max Number of Concurrent Scans	The maximum number of scan threads that can be actively working at the same time. This differs from the Max Number of Scan Threads in that not all created scan threads can be simultaneously doing work. For example, if 20 scan threads are configured but only 10 can run concurrently, 10 of the 20 threads will be paused at any one time, waiting for a time slot to continue working. Each of the 20 scan threads will get a chance to work in a round-robin fashion.

4. Click **Apply and Close** or **Apply**.

Scheduled Replication

When you create a job, you can select existing scheduled replication filters to apply to the job or you can create new scheduled replication filters and apply them to the job. This Preferences page lists the existing scheduled replication filters. From this page, you can view, create, edit, and delete scheduled replication patterns. However, you cannot edit or delete a scheduled replication filter while it is applied to a job. See Scheduled Replication in the Advanced Topics section for more information about scheduled replication.

To create a scheduled replication filter:

1. Select **Preferences** from the **Window** menu.

2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Scheduled Replication**.

   Any existing scheduled replication filters are listed in the **Scheduled Replication** table.

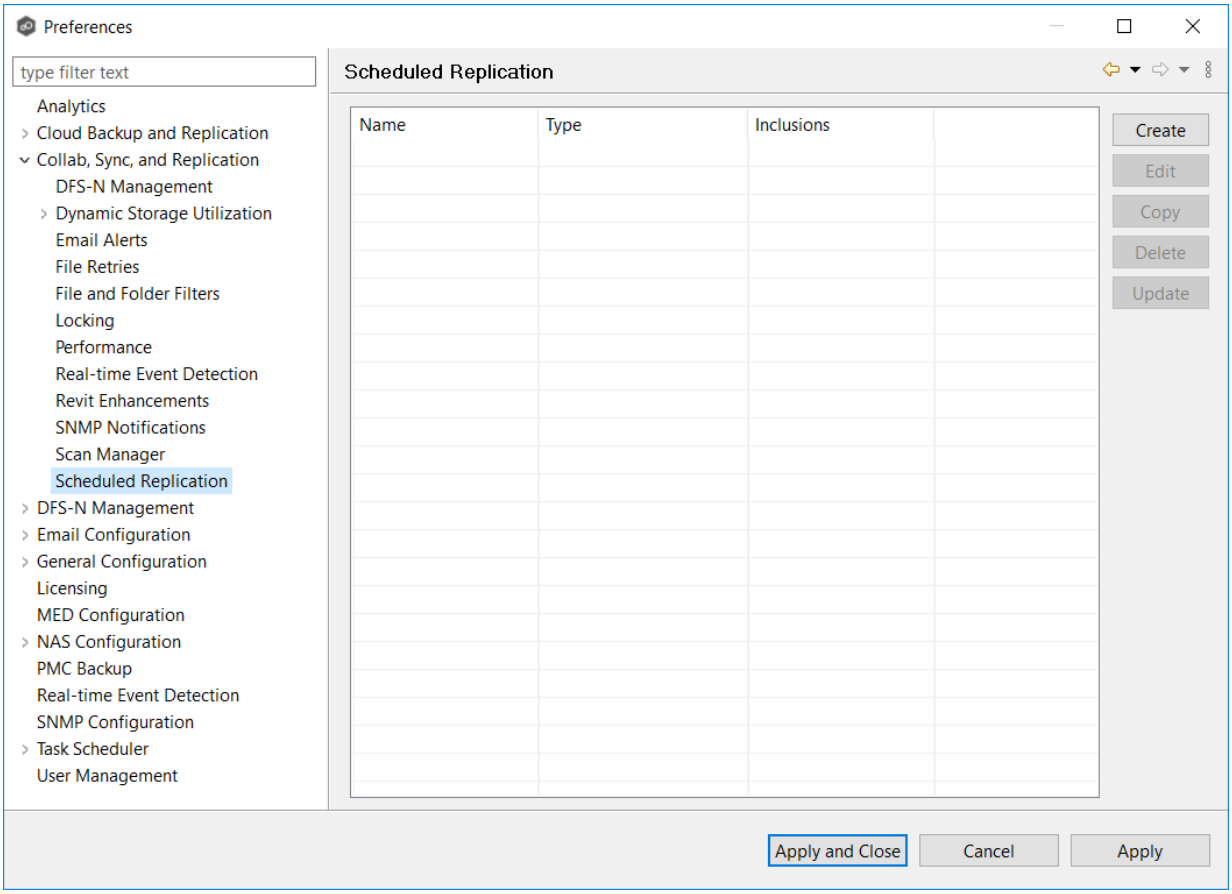

3. Click **Create**.

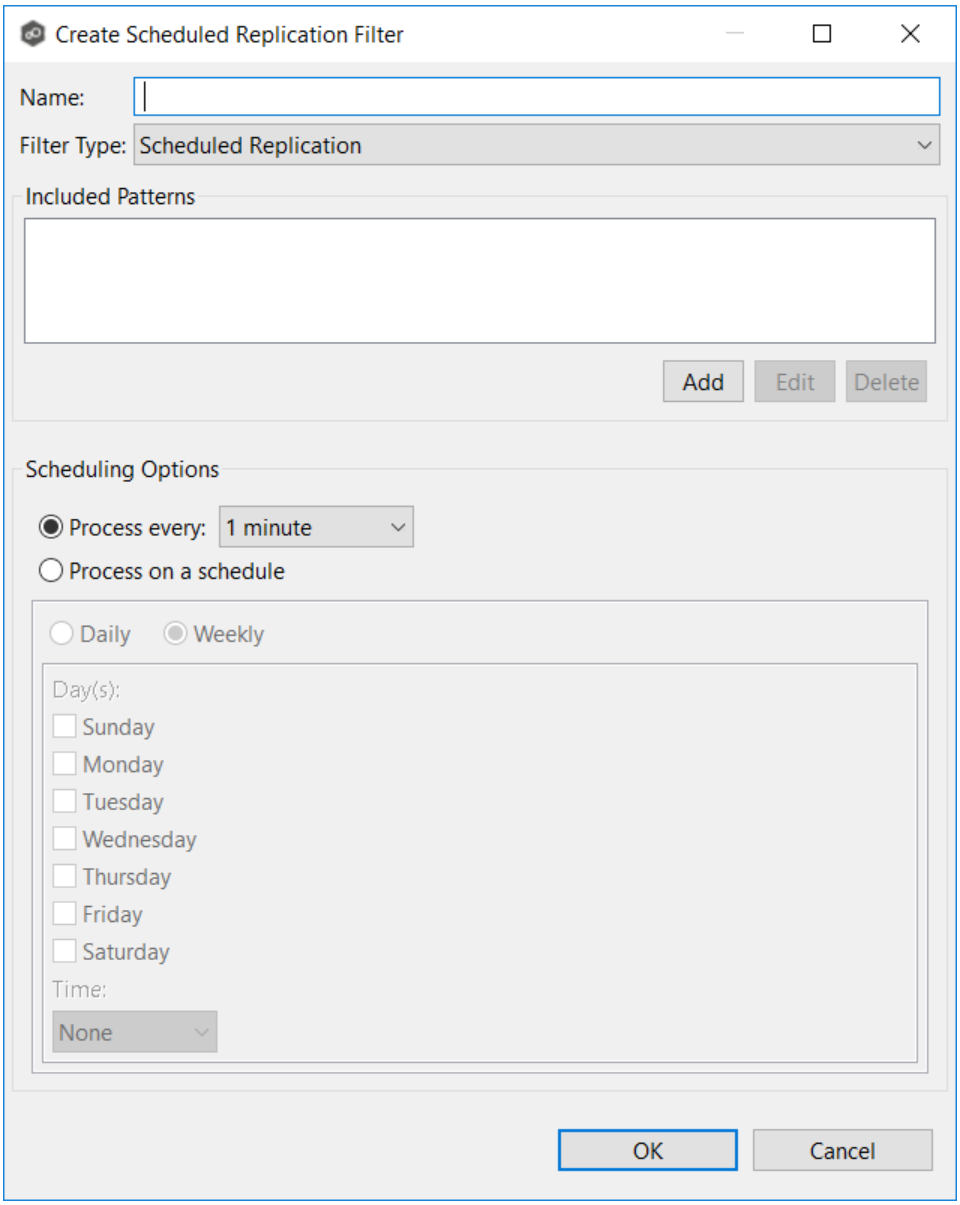

4. Enter a unique name for the filter.

5. Click **Add** under **Included Patterns** to enter a filter pattern for files that you want to delay replication. Repeat to add more filter patterns.

   A **filter pattern** is a character string that defines a logical expression that is evaluated to determine which files and folders match the pattern. A filter pattern can contain complex regular expressions and wildcards.

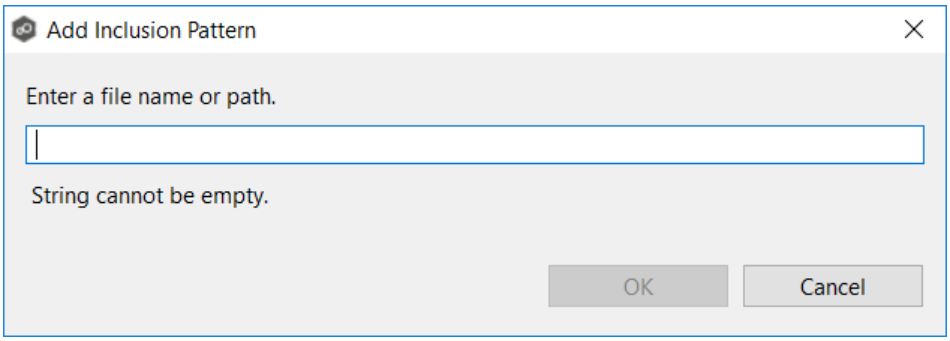

6. Click **OK**.

   The pattern appears in the **Included Patterns** field.

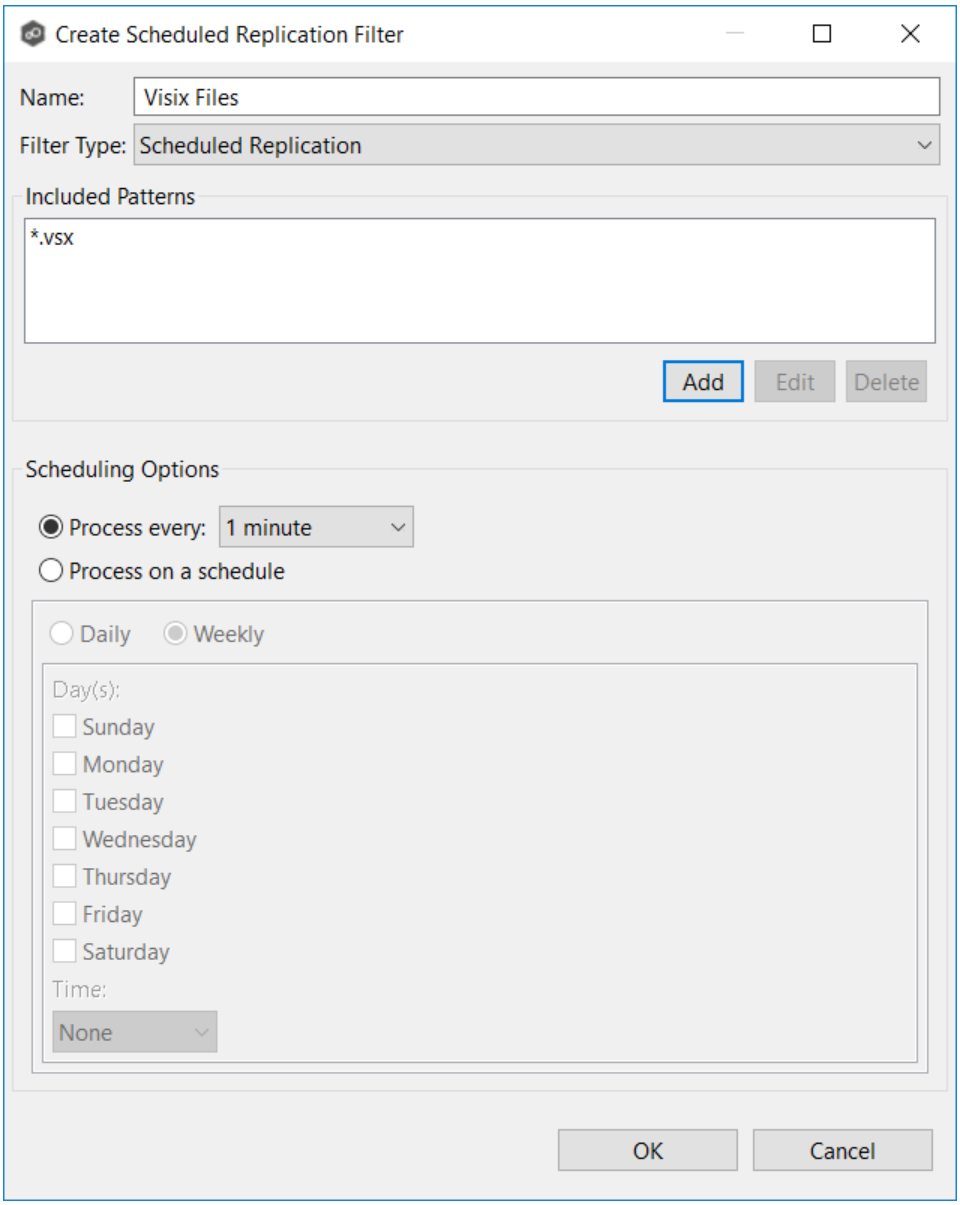

7. Select a scheduling option:

   - **Interval** - Process at a specified interval.
   - **Schedule** - Process at a scheduled time.

8. After selecting a scheduling option, click **OK**.

   The new filter is listed in the **Scheduled Replication** table and can now be applied to jobs.

9. Click **Apply and Close** or **Apply**.

DFS-N Management Job Preferences

To modify settings for DFS-N Management jobs:

1. Select **Preferences** from the **Window** menu.
2. Select **DFS-N Management** in the navigation tree.

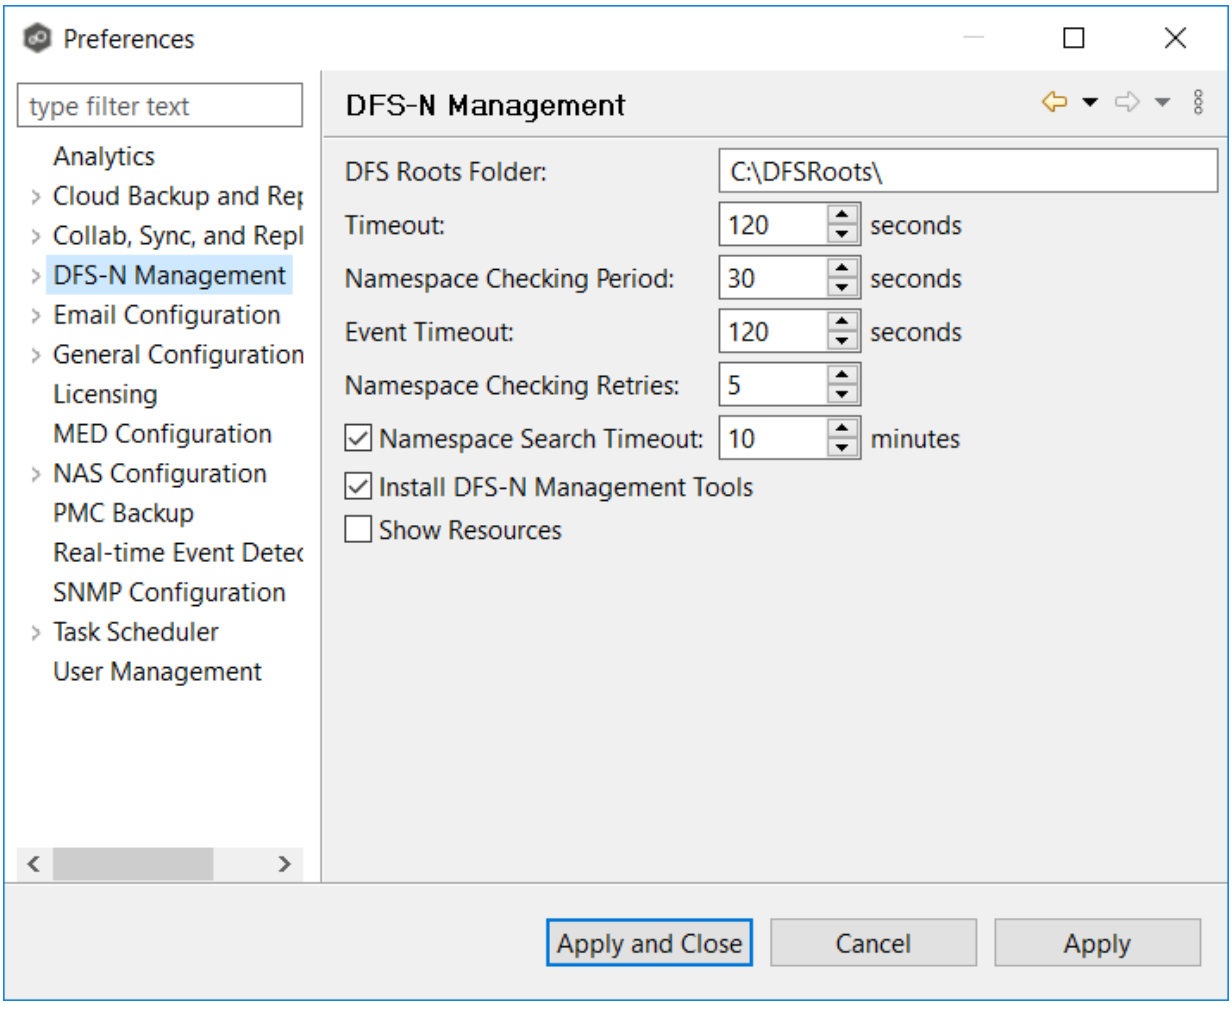

3. Modify settings as needed.

Setting	Description
DFS Roots Folder	The default folder for namespace roots is C:\DFSRoots\. If you want a different folder to be used, enter the local path to the folder.
Timeout	Enter the number of seconds to wait for a response from any Agent. The default value is 120 seconds.
Namespace Checking Period	Enter the number of seconds to delay between checking namespace information calls. This check catches any changes made to a namespace using the Microsoft DFS Management tool. Selecting a low value will negatively affect performance but will reflect changes to the user interface more quickly. The default value is 120 seconds.
Event Timeout	Enter the number of seconds to wait before marking an event containing DFS namespace information from the Agent as timed out. The default value is 120 seconds.
Namespace Checking Retries	Enter the maximum number of times for checking namespace information if the namespace is not found. Once the maximum number is exceeded, the job is stopped. The default value is 5 retries.
Namespace Search Timeout	When a user tries to import a namespace, PeerGFS searches for the namespace. This may take some time, depending on the environment. Enter the number of minutes to before timing out. The default value is 10 minutes.
Install DFS-N Management Tools	Select this option if you want Microsoft's DFS-N Management tools installed when creating or importing a namespace.
Show Resources	Select this option if you want to display individual namespace folders under each namespace in the **Jobs** view.

4. Click **Apply and Close** or **Apply**.

Email Alerts

When you create a job, you can select existing email alerts to apply to the job or you can create new email alerts and apply them to the job. This Preferences page lists the existing email alerts. From this page, you can view, create, edit, and delete email alerts. However, you cannot edit or delete an email alert while it is applied to a job. See Email Alerts in the Basic Concepts section for more information about email alerts.

Note: An SMTP email connection must be configured before email alerts can be sent. See Email Configuration for information about configuring SMTP email settings.

To create an email alert:

1. Select **Preferences** from the **Window** menu.
2. Expand **DFS-N Management** in the navigation tree, and then select **Email Alerts**.

   Any existing DFS-N Management email alerts are listed in the **Email Alerts** table.

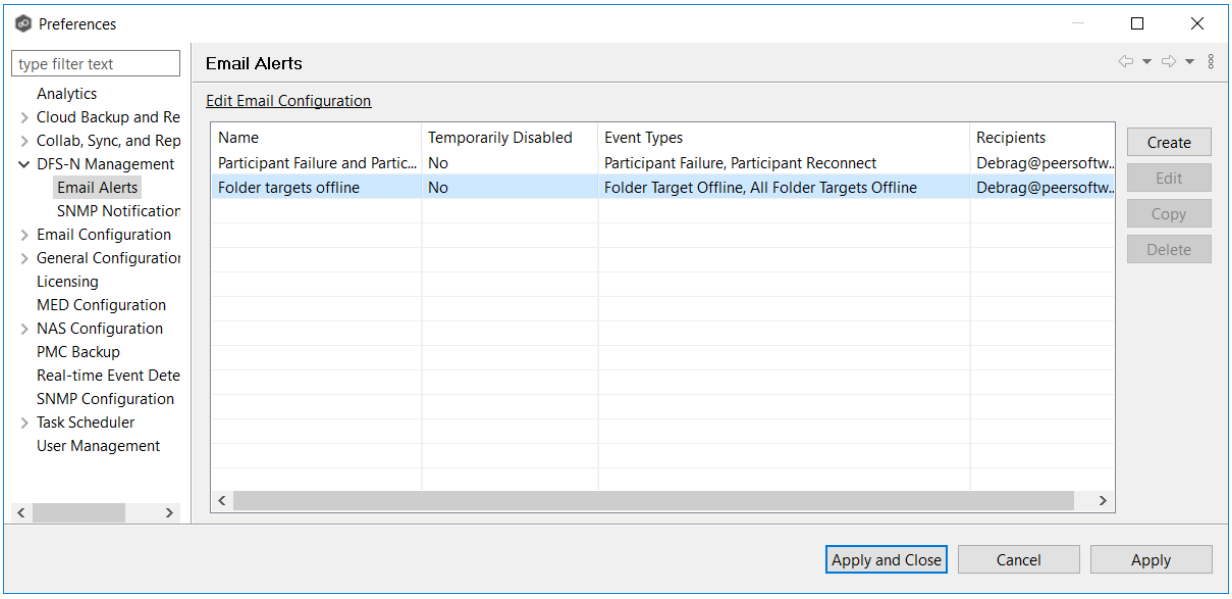

3. Click the **Create** button.

   The **Create Email Alert** dialog appears.

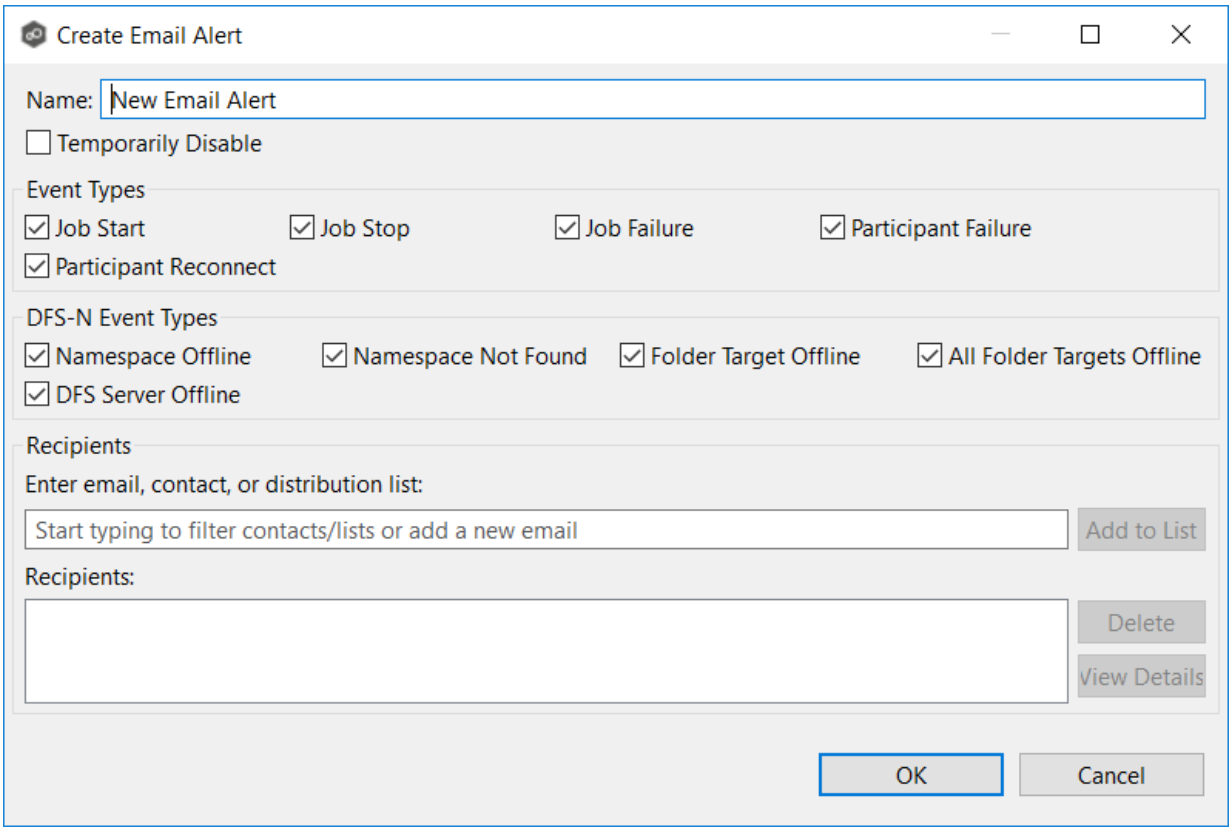

4. Enter a name for the alert.
5. Select the event types to be alerted.

   The event type determines what will trigger the email alert to be sent.

Event Type	Description
Job Start	Sends an alert when the job starts.
Job Stop	Sends an alert when the job stops.
Job Failure	Sends an alert when the job stops unexpectedly.
Participant Failure	Sends an alert when the Management Agent job disconnects or stops responding.

Event Type	Description
Participant Reconnect	Sends an alert when the Management Agent reconnects.

6. Select the DFS-N event types.

Event Type	Description
Namespace Offline	Sends an alert when a namespace goes offline.
Namespace Not Found	Sends an alert when a namespace is not found.
Folder Target Offline	Sends an alert when a folder target goes offline.
All Folder Targets Offline	Sends an alert when all folder targets go offline
DFS Server Offline	Sends an alert when a DFS server goes offline.

7. Enter the alert recipients, and then click **Add to List**.

   The recipients are listed in the **Recipients** field.

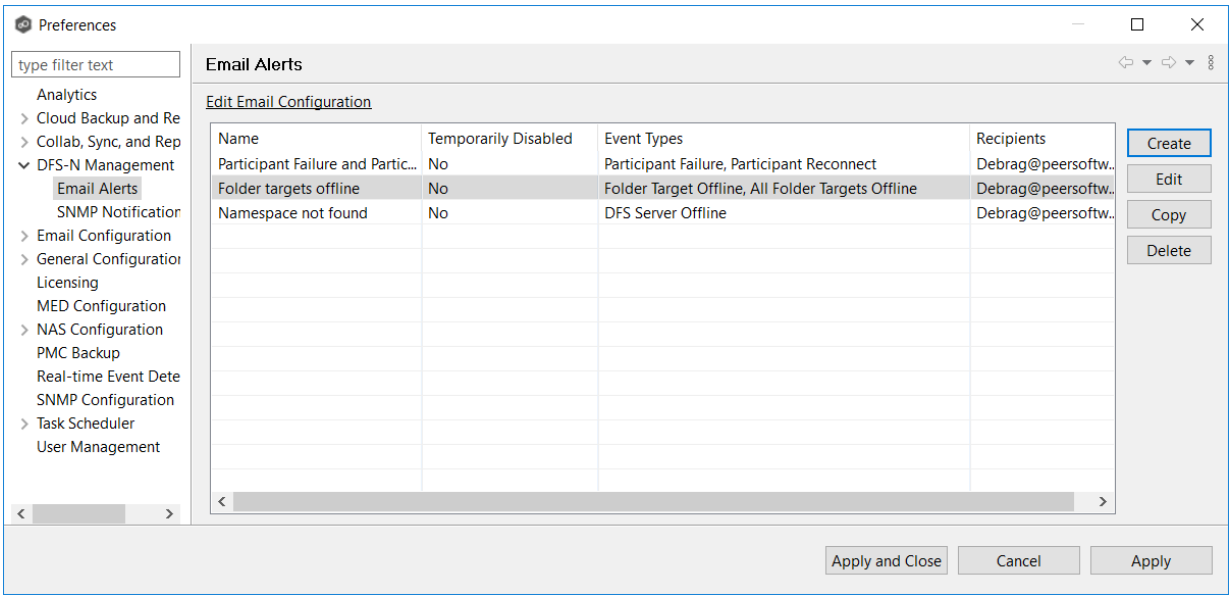

8. Click **OK**.

   The new alert is listed in the **Email Alerts** table and can now be applied to jobs.

9. Click **Apply and Close** or **Apply**.

SNMP Notifications

When you create a job, you can select an existing SMNP notification to apply to the job or you can create a new notification and apply it to the job. You cannot edit or delete an SMNP notification while it is applied to a job. See SMNP Notifications in the Basic Concepts section for more information about SMNP notifications.

Note that the SNMP source address, destination, and prefix must be configured before notifications can be set.

To create an SMNP notification:

1. From the **Window** menu, select **Preferences**.

2. Expand **DFS-N Management** in the navigation tree, and then select **SNMP Notifications**.

   The existing SNMP notifications are listed in the **SNMP Notifications** table.

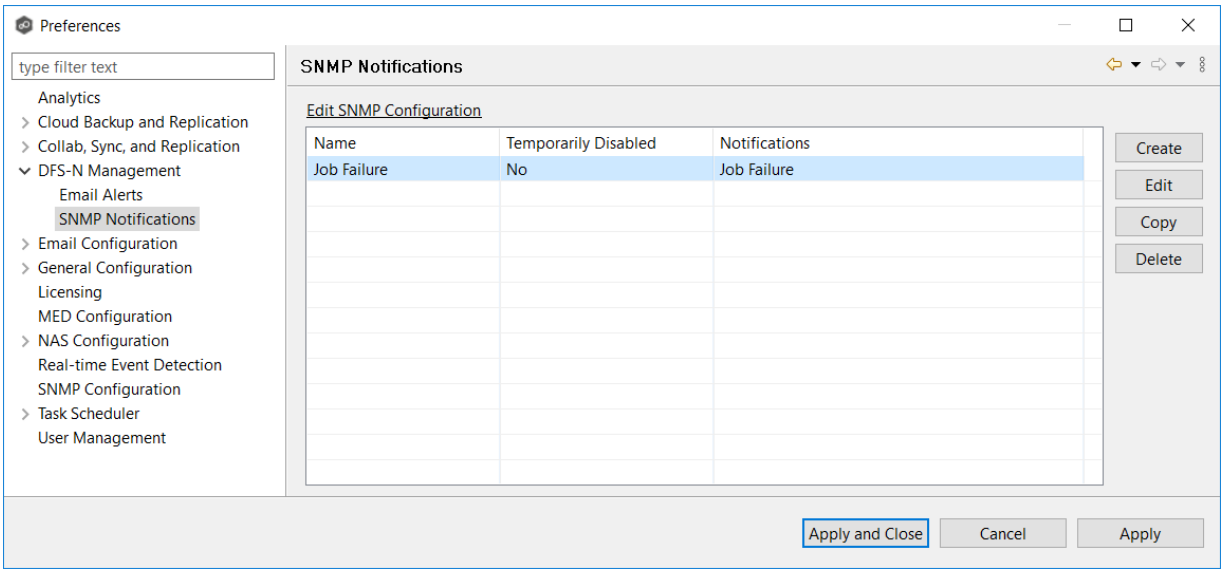

3. Click the **Create** button.

   The **Create SNMP Notification** dialog appears.

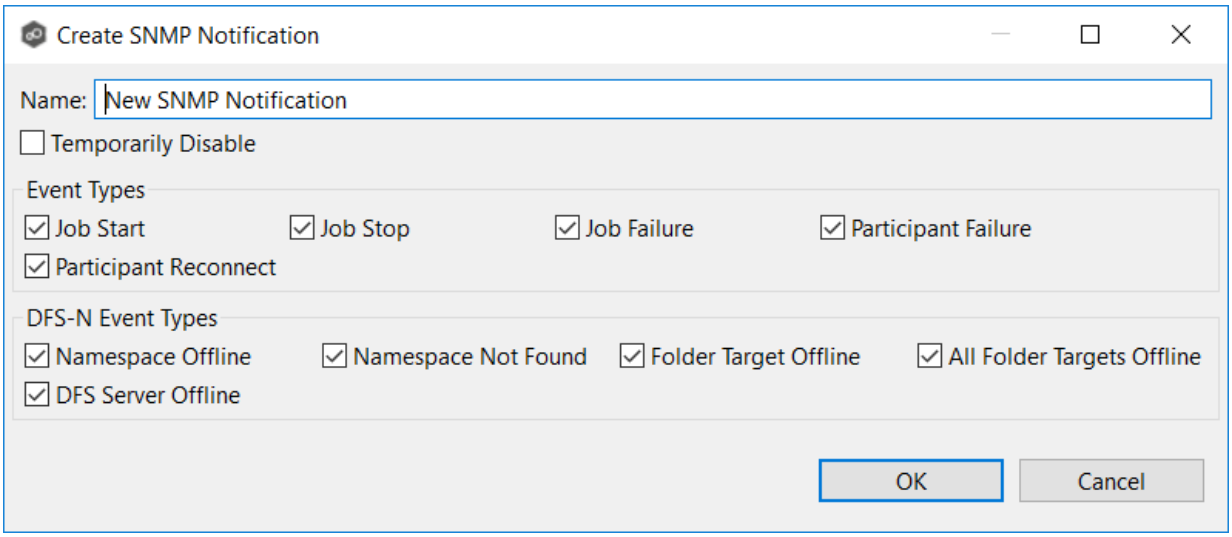

4. Select the types of events that will trigger the generation of an SNMP trap.

Event Type	Description
Job Start	Sends a notification when the DFS-N Management job starts.

Event Type	Description
Job Stop	Sends a notification when the DFS-N Management job stops.
Job Failure	Sends a notification when the DFS-N Management job stops unexpectedly.
Participant Failure	Sends a notification when the Management Agent of the DFS-N Management job disconnects or stops responding.
Participant Reconnect	Sends a notification when the Management Agent of the DFS-N Management job reconnects.

5. Select the DFS-N event types that will trigger the generation of an SNMP trap.

Event Type	Description
Namespace Offline	Sends a notification when a namespace goes offline.
Namespace Not Found	Sends a notification when a namespace is not found.
Folder Target Offline	Sends a notification when a folder target goes offline.
All Folder Target Offline	Sends a notification when all folder targets go offline

Event Type	Description
DFS Server Offline	Sends a notification when a DFS server goes offline.

6. Click **OK**.

   The new notification is listed in the **SNMP Notifications** table and can now be applied to jobs.

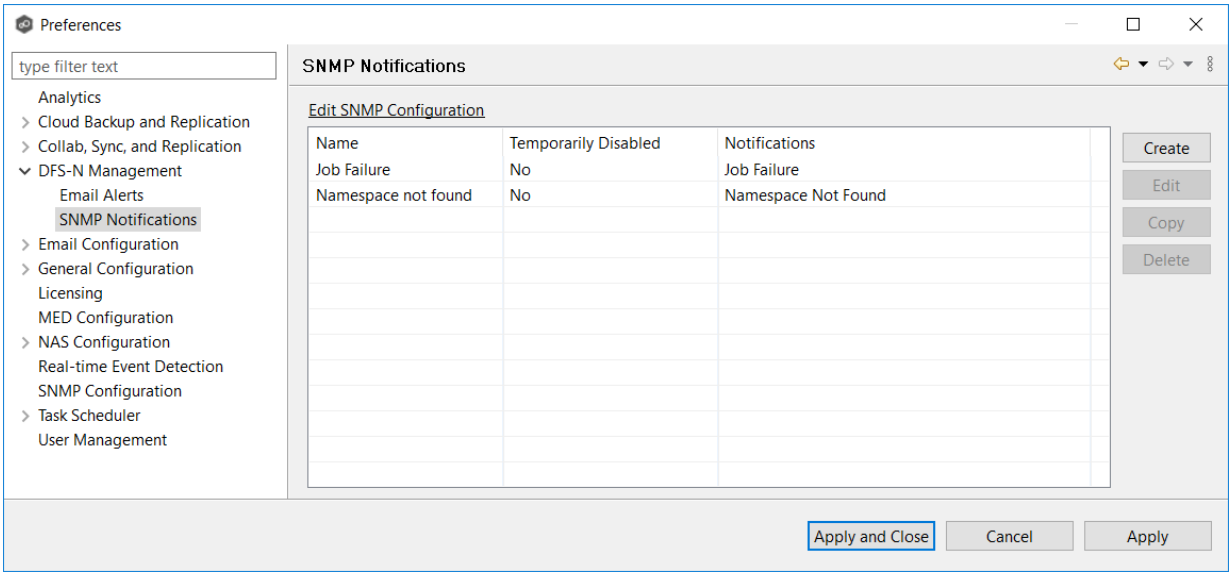

7. Click **Apply and Close** or **Apply**.

Email Configuration

Before Peer Management Center can send emails on behalf of any job, a few key SMTP email settings must be configured. In addition, you can define contacts and distribution lists.

To configure email settings:

1. Select **Preferences** from the **Window** menu.

2. Select **Email Configuration** in the navigation tree.

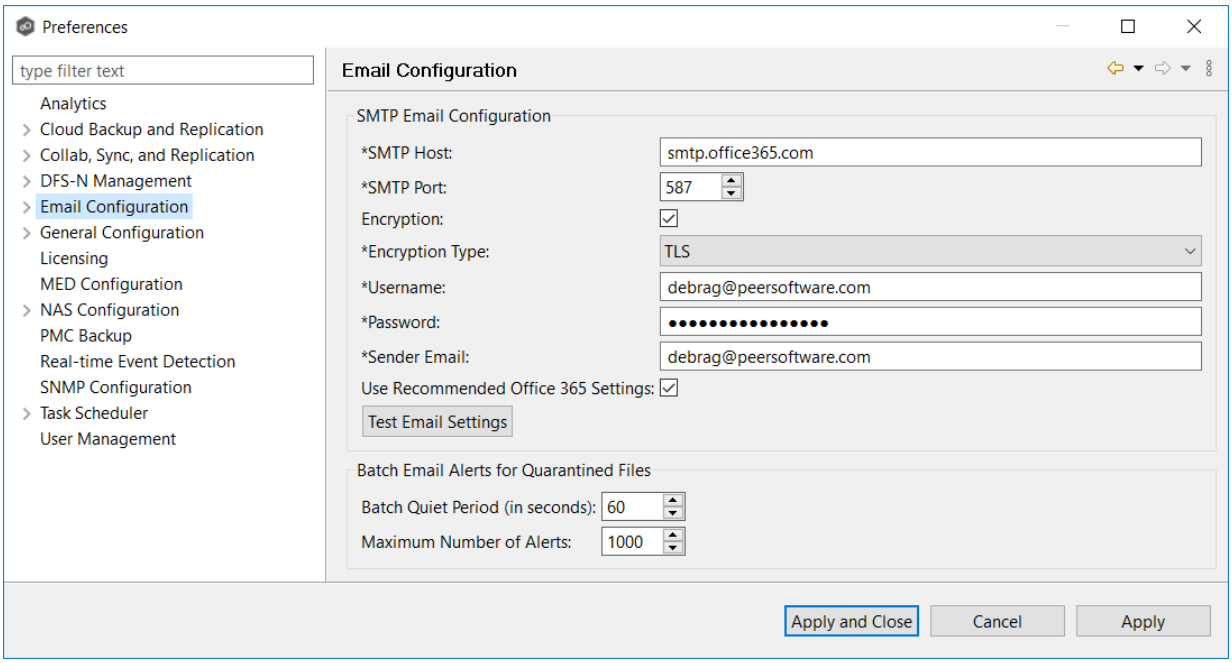

3. Enter values for the following fields:

Field	Description
SMTP Host	Enter the host name or IP address of the SMTP mail server through which Peer Management Center will send emails.
SMTP Port	Enter the TCP/IP connection port (default is 25 and 465 for encryption) on which the mail server is hosting the SMTP service. We recommend that you leave the default setting unless your email provider specifies otherwise.
Encryption	Select this checkbox if the SMTP mail server requires an encrypted connection.
Encryption Type	If encryption is enabled, an encryption method must be selected. TLS and SSL are the available options. If you do not know which one your mail server requires, try one, and then the other.
Username	Enter the user name to authenticate as on the SMTP mail server.
Password	Enter the password for the user specified above.
Sender Email	Enter the email address to appear in the **From** field of any sent emails. This email address sometimes needs to have a valid account on the SMTP mail server.
Use Recommended Office 365 Settings	Select this checkbox if you are connecting to an Office 365 SMTP server. Follow Microsoft's Direct Send recommendations to set up email configuration with an Office 365 SMTP server.

4. (Recommended) Click **Test Email Settings**, enter an email address, and then click **OK**.

   It is highly recommended that you test your SMTP settings before saving them. You will be prompted for an email address to send the test message to. Upon submission, Peer Management Center will attempt to send a test message using the specified settings.

5. Enter values for the fields in the **Batch Email Alerts for Quarantined Files** section:

Field	Description
Batch Quiet Period (in seconds)	Enter the number of seconds to wait before releasing a batch of alerts.
Maximum Number of Alerts	Enter the maximum number of alerts that should be sent in a single email.

6. Click **OK** or **Apply**.

General Configuration

The **General Configuration** settings affect the overall operation of Peer Management Center, Peer Agents, the Peer Broker, and other general operations. They are not specific to jobs or job types.

You can modify the following settings:

General Configuration

Agent Connectivity

Broker Configuration

Email Alerts

Software Updates

Tags Configuration

Web and API Configuration

General Configuration

To modify General Configuration settings:

1. Select **Preferences** from the **Window** menu.
2. Select **General Configuration** in the navigation tree.

   The first page of the General Configuration options is displayed.

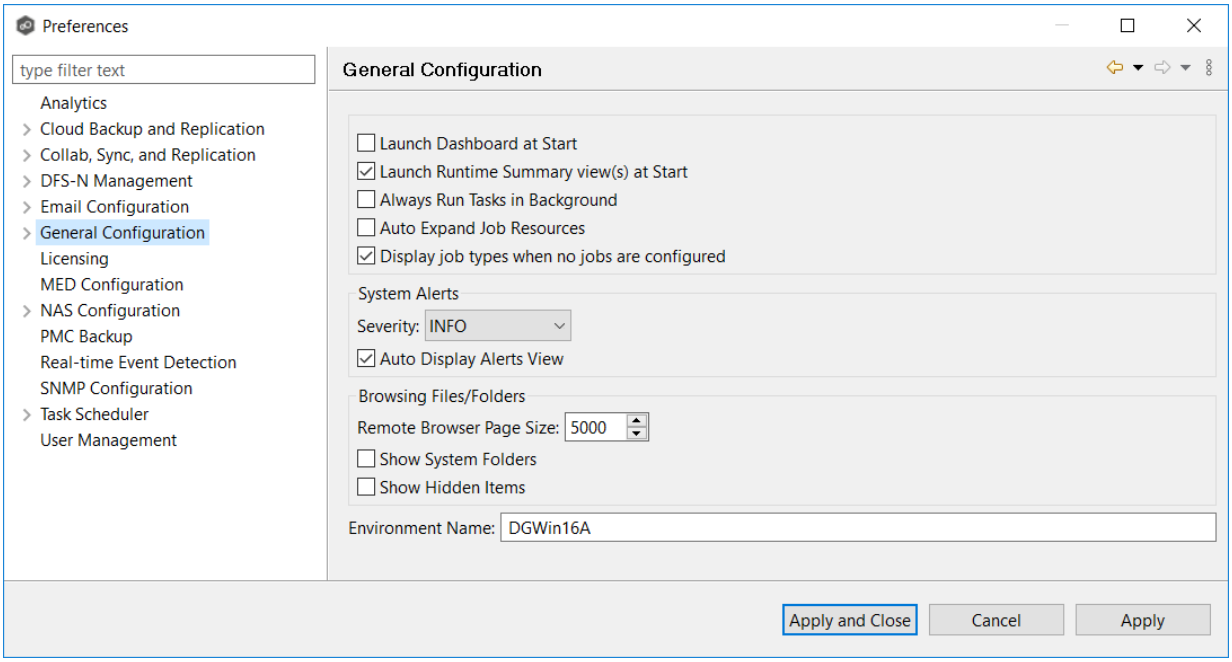

3. Modify the first four settings as needed:

Setting	Description
Launch Dashboard at start	Select this option if you want the Dashboard to be automatically displayed when Peer Management Center is started.
Launch Runtime Summary view(s) at start	Select this option if you want the Summary views to be automatically displayed when Peer Management Center is started. Summary views will be displayed for all job types, even for job types without currently running jobs.

Setting	Description
Always run tasks in background	Select this option to run tasks like log gathering and Agent updates in the background, preventing these tasks from blocking the use of the Peer Management Center client while they run.
Auto expand job resources	Select this option if you want all jobs with associated resources to start expanded in the **Jobs** view. Currently only available for Cloud Backup and Replication jobs and DFS-N Management jobs.
Display job types when no jobs are configured	Select this option if you want to display a job type in the **Jobs** view, even when no jobs of that type have been configured.

4. Select options for alerts regarding the operation of Peer Management Center in the Alerts view:

Option	Description
Severity	Select one of these options: • INFO • DEBUG • TRACE
Auto Display Alerts View	Select this option if you want the alerts to be automatically displayed in the Alerts view.

5. Select options for managing browsing files and folders on remote file systems in the **Browsing Files/Folders** section:

Option	Description
Remote browser page size	Enter the maximum page size for the remote file system browser. This browser is used for selecting paths during the creation of most new jobs.
Show system folders	Select this checkbox to show system folders in the remote file system browser.
Show hidden folders	Select this checkbox to show hidden folders in the remote file system browser.

6. (Optional) Enter the name of your PMC server or environment in the **Environment Name** field; if left blank, reports and dashboards will use the name of the PMC server.

7. Click **OK** or **Apply**.

Agent Connectivity

To modify Agent Connectivity settings:

1. Select **Preferences** from the **Window** menu.

2. Expand **General Configuration** in the navigation tree, and then select **Agent Connectivity**.

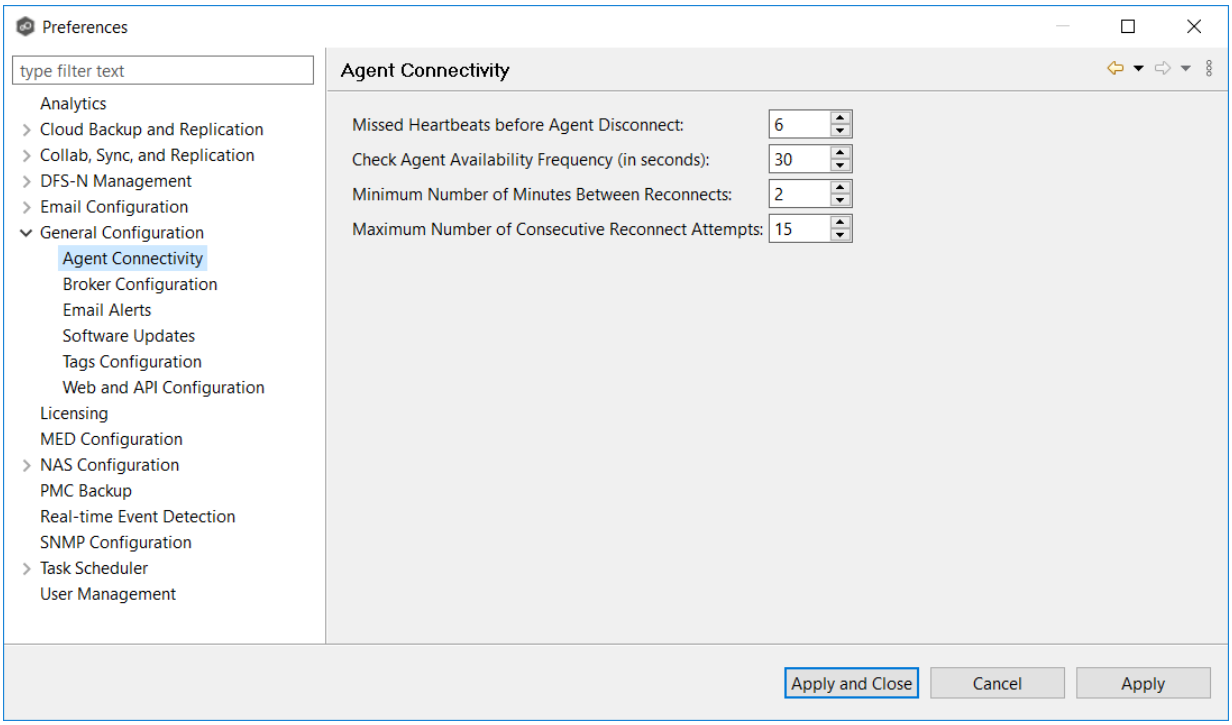

3. Modify the settings as needed:

Option	Description
Missed Heartbeats before Agent Disconnect	Enter the maximum number of heartbeats that can be missed on a host before Peer Management Center labels the Agent as disconnected. If a running job hits a timeout when communicating with a specific Agent, Peer Management Center will check this status to decide if the Agent should be dropped from the job.
Check Agent Availability Frequency (in seconds)	Enter the frequency (in seconds) that Peer Management Center should check whether an Agent is back online.
Minimum Number of Minutes Between Reconnects	Enter the minimum number of minutes that must elapse before Peer Management Center attempts to retry reconnecting to the Agent.
Maximum Number of Consecutive Reconnect Attempts	Enter the maximum number of attempts that Peer Management Center tries to reintegrate a previously connected agent into one or more jobs. Once the maximum number of attempts has been reached, you must manually reintegrate the Agent into affected jobs, typically by restarting the affected jobs.

4. Click **OK** or **Apply**.

Broker Configuration

The **Broker Configuration** page displays a non-editable field that shows the URL used by the Peer Management Center service to connect to the Broker service.

To view the Broker Configuration URL:

1. Select **Preferences** from the **Window** menu.

2. Expand **General Configuration** in the navigation tree, and then select **Broker Configuration**.

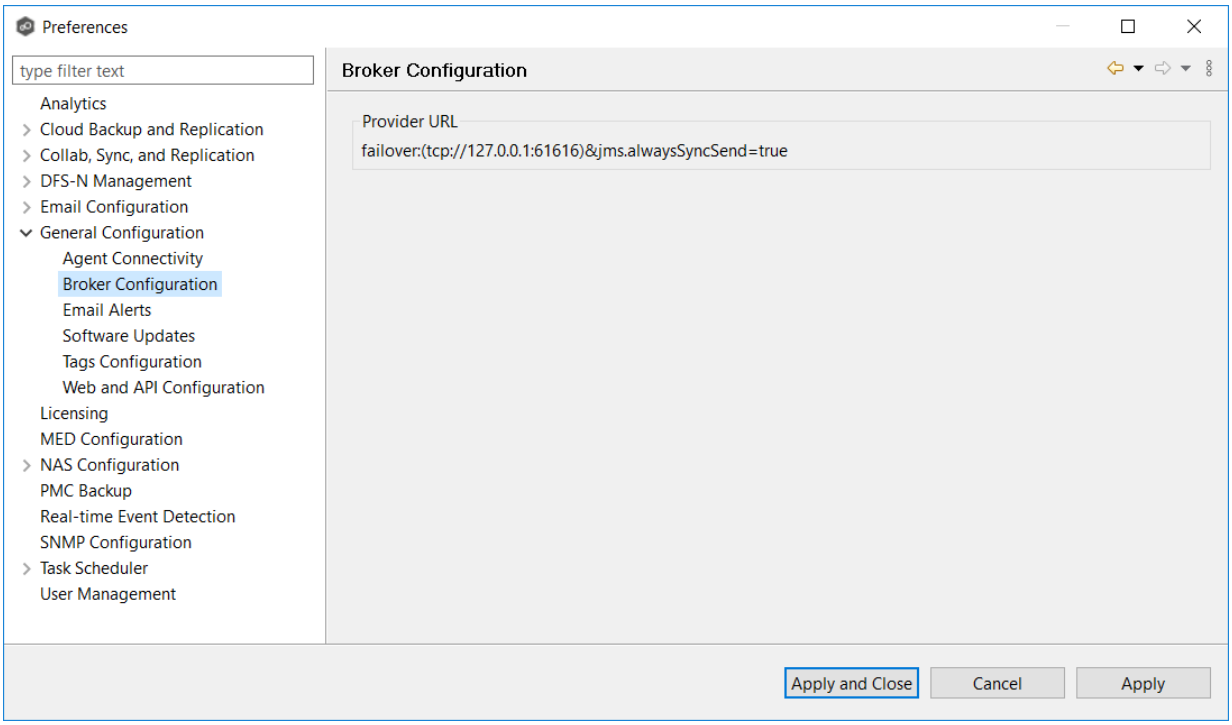

3. Click **OK** or **Apply**.

Email Alerts

System email alerts notify recipients when certain types of system events occur, for example, low memory, low disk space, disconnected agents. This Preferences page lists the existing system email alerts. From this page, you can create, edit, and delete system email alerts. You can also disable and enable alerts. See Email Alerts in the Basic Concepts section for more information about email alerts.

Note: An SMTP email connection must be configured before email alerts can be sent. See Email Configuration for information about configuring SMTP email settings.

To create a system email alert:

1. Select **Preferences** from the **Window** menu.

2. Expand **General Configuration** in the navigation tree, and then select **Email Alerts**.

   Any existing system email alerts are listed in the **Email Alerts** table.

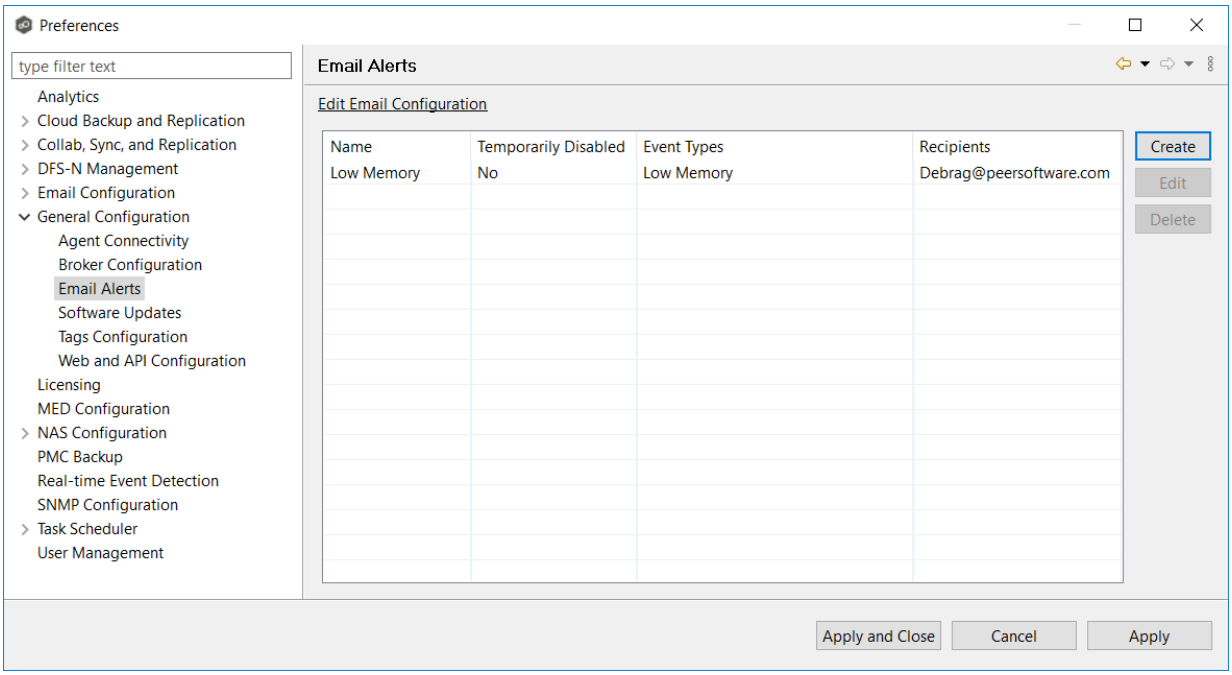

3. Click **Create**.

   The **Create Email Alert** dialog appears.

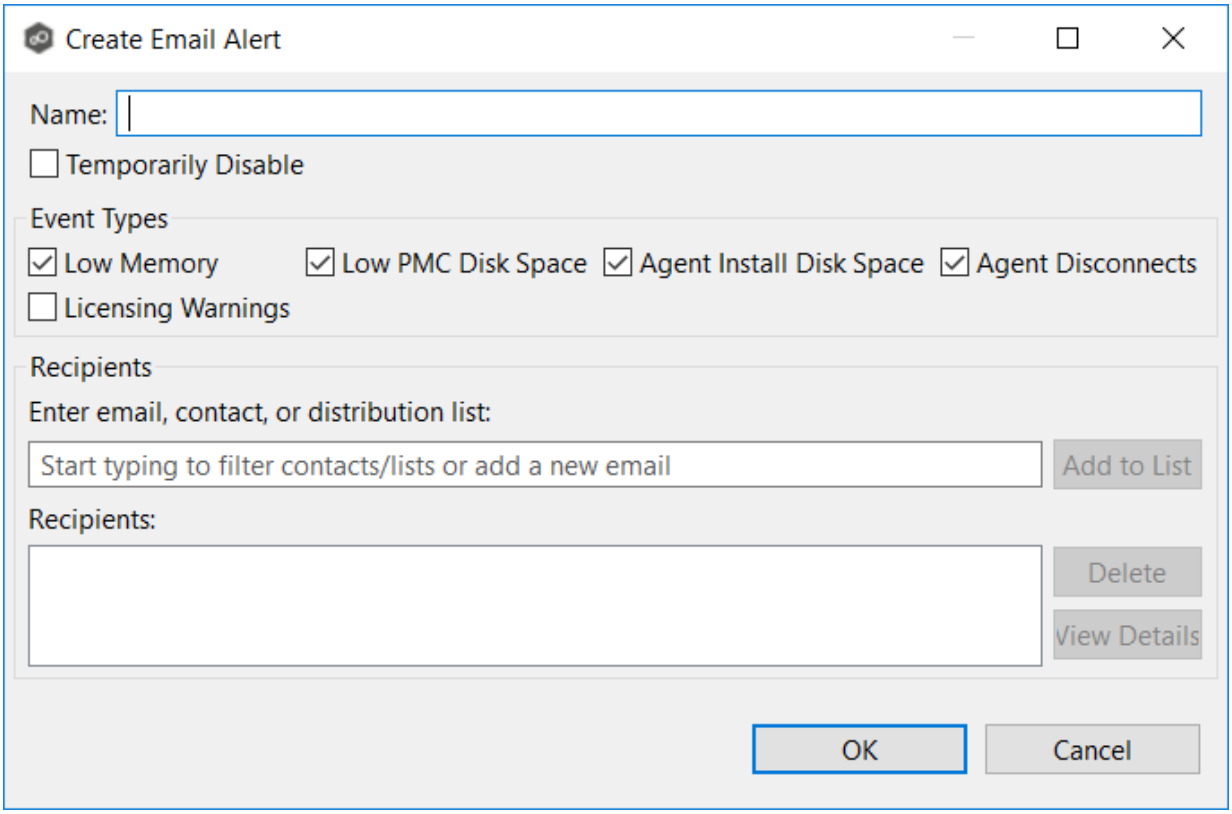

4. Enter a name for the alert.
5. Select the **Enable** checkbox if you want to enable the alert.

   If you choose not to enable the alert, you can enable it later.
6. Select the type of events for which you want alerts sent:

Event Type	Description
Low Memory	Sends an alert when Peer Management Center or connected Agent services are low on memory.
Low PMC Disk Space	Sends an alert when the space on the disk where Peer Management Center software is installed running low.
Agent Install Disk Space	Sends an alert when the space on the disk where the Peer Agent software is installed is running low.
Agent Disconnects	Sends an alert whenever an Agent is disconnected.
License Warnings	Sends an alert when a license has expired or when a license violation is about to occur (for example, when storage usage reaches 95% of maximum storage and when storage usage exceed license limits).

7. Enter alert recipients, and then click **Add to List**.
8. Click **OK** or **Apply**.

   The new alert is listed in the **Email Alerts** table.

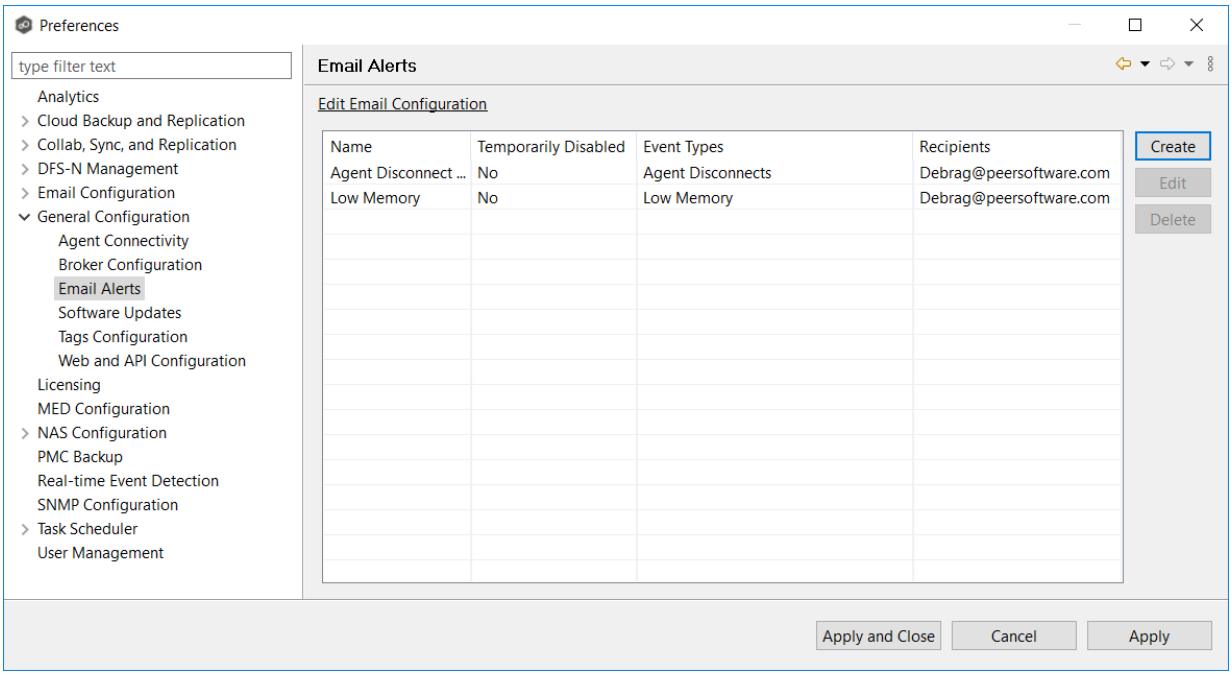

Software Updates

You can configure Peer Management Center to automatically check for updates and download the updates. Peer Management Center to checks for updates every evening at 11 p.m. local time. Only minor updates are automatically downloaded; if a major update is available, a notification appears. Major releases require a new license key and must be requested from Peer Software Support.

You can also manually check for updates. See Updating Peer Management Center for information about manually checking for updates.

To configure Peer Management Center to automatically check for updates:

1. Select **Preferences** from the **Window** menu.

2. Expand **General Configuration** in the navigation tree, and then select **Software Updates**.

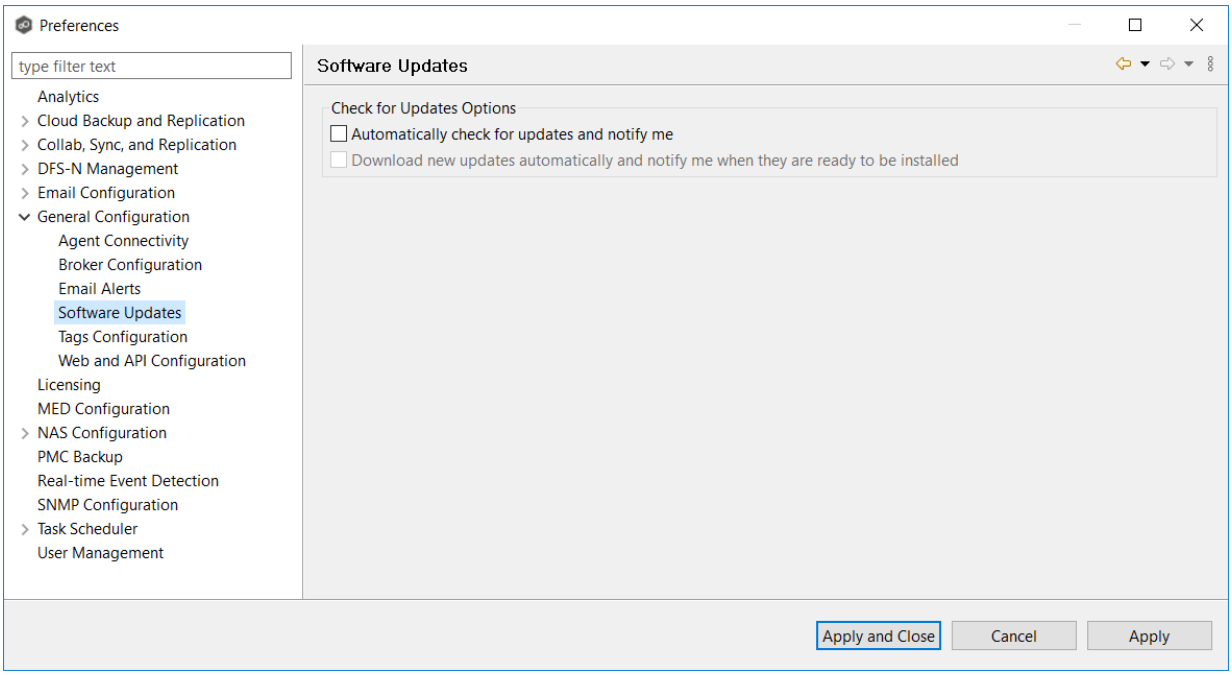

3. Select update options:

   - **Automatically check for updates and notify me** - Select this option if you want to automatically check for updates.

   - **Download new updates automatically and notify me when they are ready to be installed** - Select this option if you want to automatically check for and download available updates.

4. Click **OK** or **Apply**.

   Whenever updates are available, a notification appears in the lower right corner of Peer Management Center.

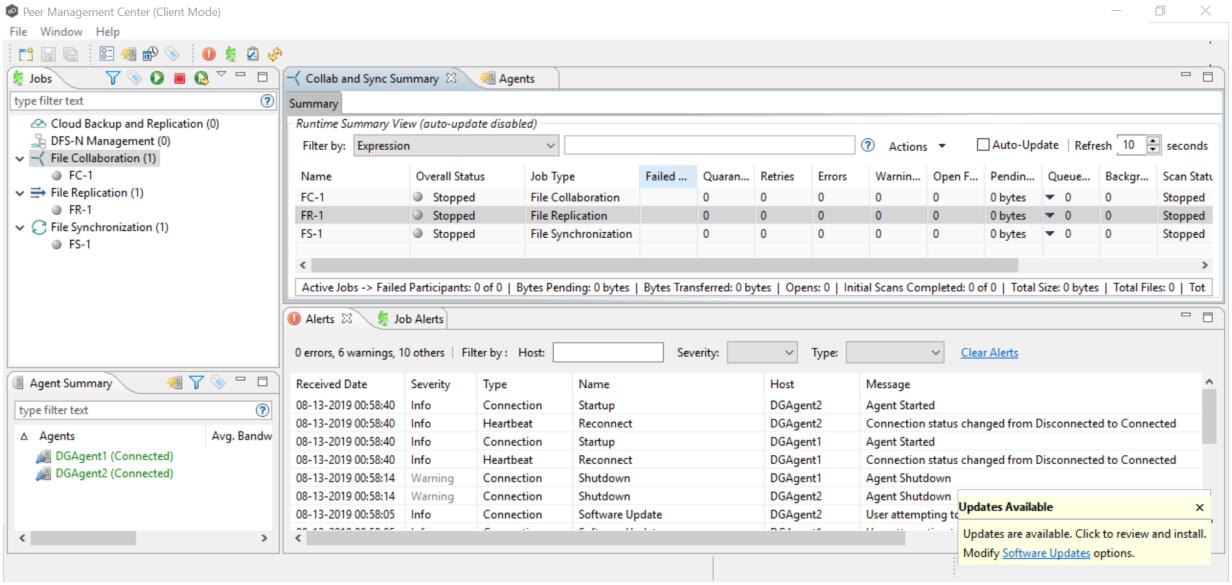

5. Click the notification to review and proceed with the update. See Updating Peer Management Center for details.

Tags Configuration

The **Tags Configuration** page in Preferences is the starting place for creating tags and categories that can later be assigned to resources. See Assigning Tags for more information about assigning to resources.

To create a tag:

1. Select **Preferences** from the **Window** menu.

2. Expand **General Configuration** in the navigation tree, and then select **Tags Configuration**.

   Any existing tags are listed in the **Tags** table.

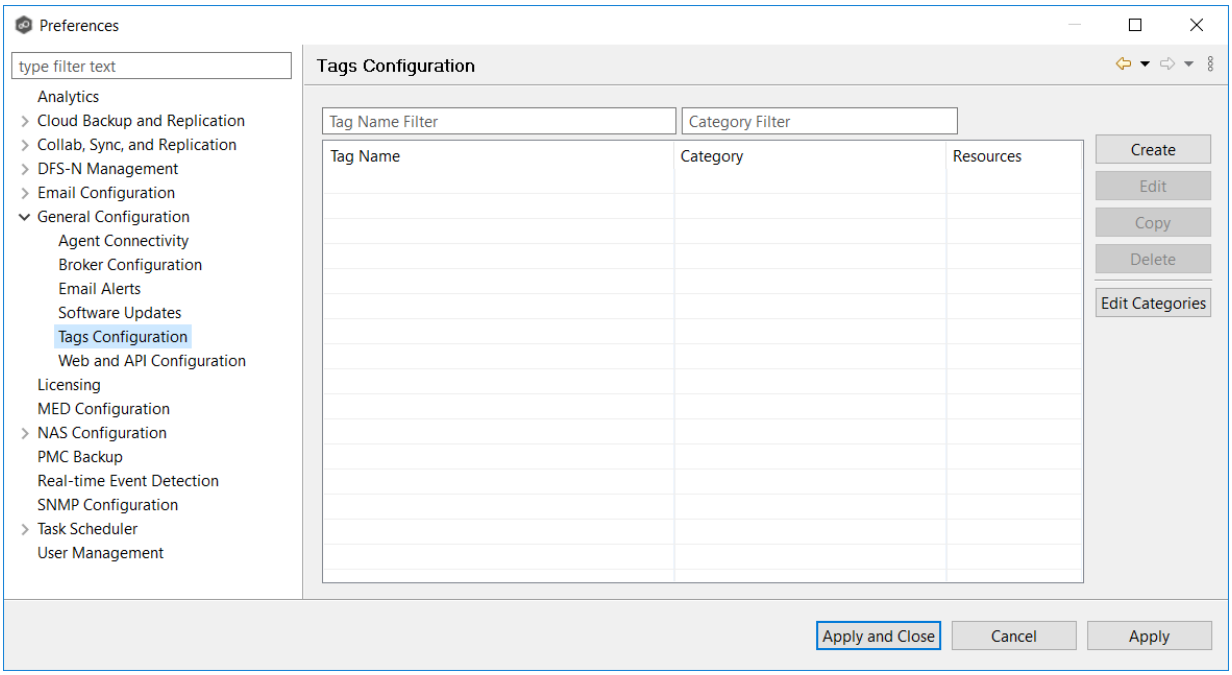

3. Click the **Create** button.

   The **New Tag** dialog appears.

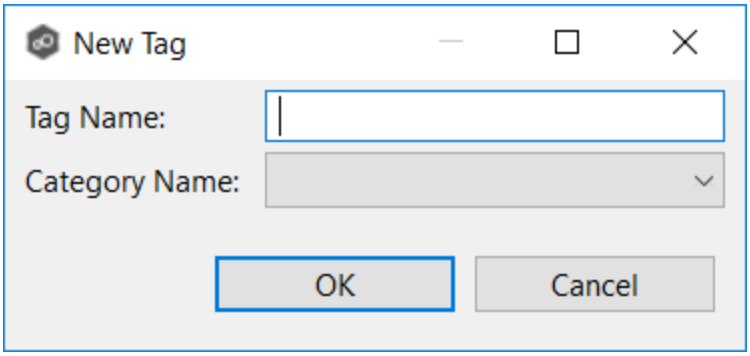

4. Enter a name for a tag.
5. Select a category or create a new category.
6. Click **OK**.

   The tag appears in the **Tags** table.

7. Click **OK** or **Apply**.

Web and API Configuration

As part of the Peer Management Center installation process, you are prompted to configure access to the web and API services. If you do not enable access during the initial installation or want to modify settings at a later date, you can modify them in Web and API Configuration in Preferences.

To modify web and API settings:

1. Select **Preferences** from the **Window** menu.
2. Expand **General Configuration** in the navigation tree, and then select **Web and API Configuration**.

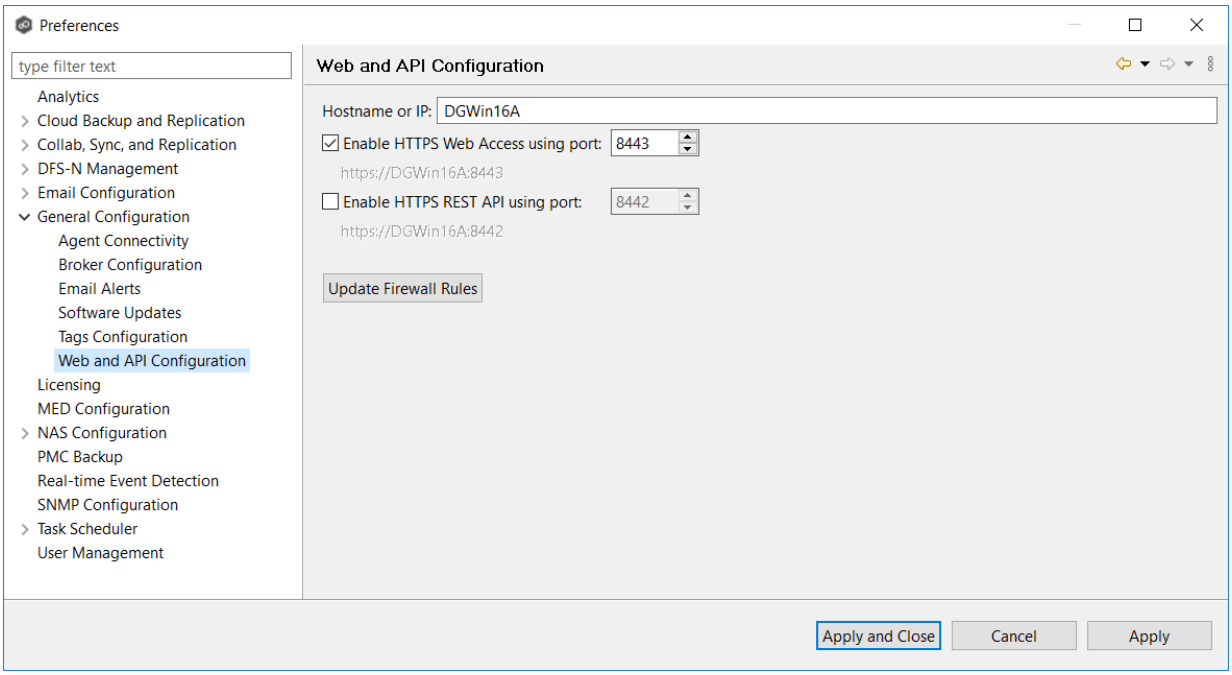

3. Modify the configuration options.

Option	Discription
Hostname or IP	Enter the hostname or IP address via which the services can be accessed: • Enter **localhost** or **127.0.0.1** if you want the services to be accessible only to users of the local server via the loopback interface.

Option	Discription
	• Enter **0.0.0.0** to make the services accessible via all network interfaces. • Enter a specific IP address to restrict access to a specific network interface.
Enable HTTPS Web Access	Select this checkbox to enable HTTP access to the web service using the specified port.
Enable HTTPS REST API	Select this checkbox to enable HTTPS access to the REST API service using the specified port.

4. Click **OK** or **Apply**.

Licensing

Peer Global File System is licensed by the number of unique participants and by the number of terabytes in the watch set.

Installing or Upgrading a License File

After purchasing or requesting a trial download of Peer Management Center, you will receive a license file representing your purchase or trial.

To install a new license file or upgrade an existing license:

1. From the **Window** menu, select **Preferences**.
2. Select **Licensing** in the navigation tree.

   Existing valid licenses are listed in the **Peer Software Licensing Configuration** table.

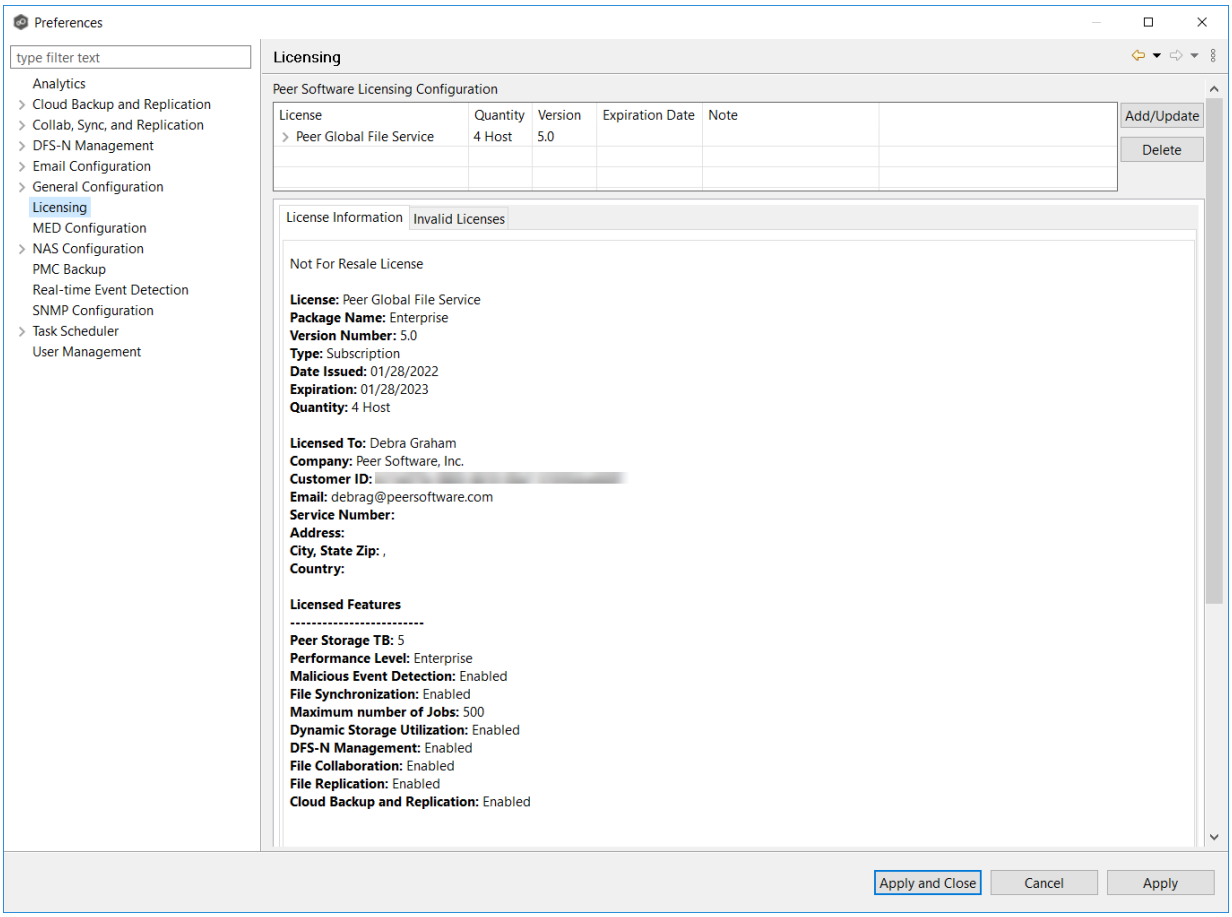

3. Click the **Add/Update** button to browse for a license file.

4. Select the license file, and then click **Open**.

   If you are prompted with a message that an existing license already exists, click **Yes** to overwrite the existing license.

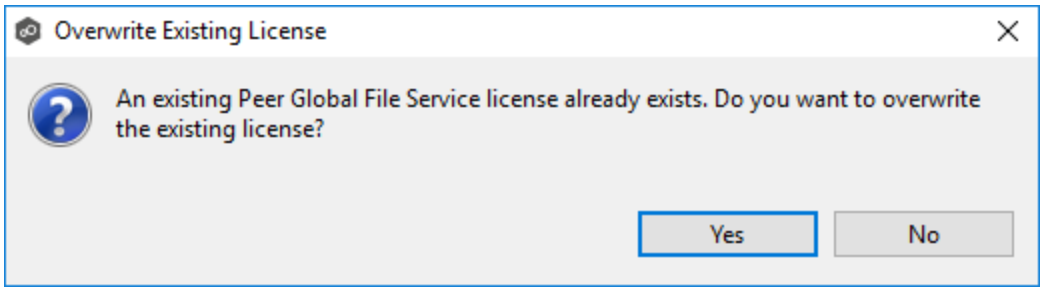

After successful installation of the license, it is listed in the table, along with the license quantity, version, and an expiration date (if applicable). You can now create, configure, and run jobs using the new license.

Note: You will need to restart existing jobs if any of the following applies:

- Software version is different (typically when upgrading to a new version).
- Software package level is different.
- New license is insufficient for the number of existing hosts.

5. Click the license in the table to view details about the license.
6. Click **OK** or **Apply**.

Deleting a License File

To delete a license.

1. From the **Windows** menu, select **Preferences**.
2. Select **Licensing** in the navigation tree.
3. Select the license you want to delete.
4. Click the **Delete** button

   Any job types enabled by that license will be hidden from Peer Management Center.

MED Configuration

Peer's Malicious Event Detection (MED) real-time engine can spot unwanted activity being executed on storage platforms by ransomware, viruses, malware, hackers, or rogue users. MED technology provides alerting capabilities, as well as the ability to minimize the amount of encrypted or deleted content from being replicated to remote locations. Once MED is enabled and jobs are restarted, these capabilities apply to all jobs. For more information, see our knowledge base article Introduction to Peer MED.

Peer MED deploys three different mechanisms for spotting malicious activity, each of which can be enabled and tuned independently. These settings are configured on a global level.

To modify MED settings:

1. From the **Window** menu, select **Preferences**.
2. Select **MED Configuration** in the navigation tree.

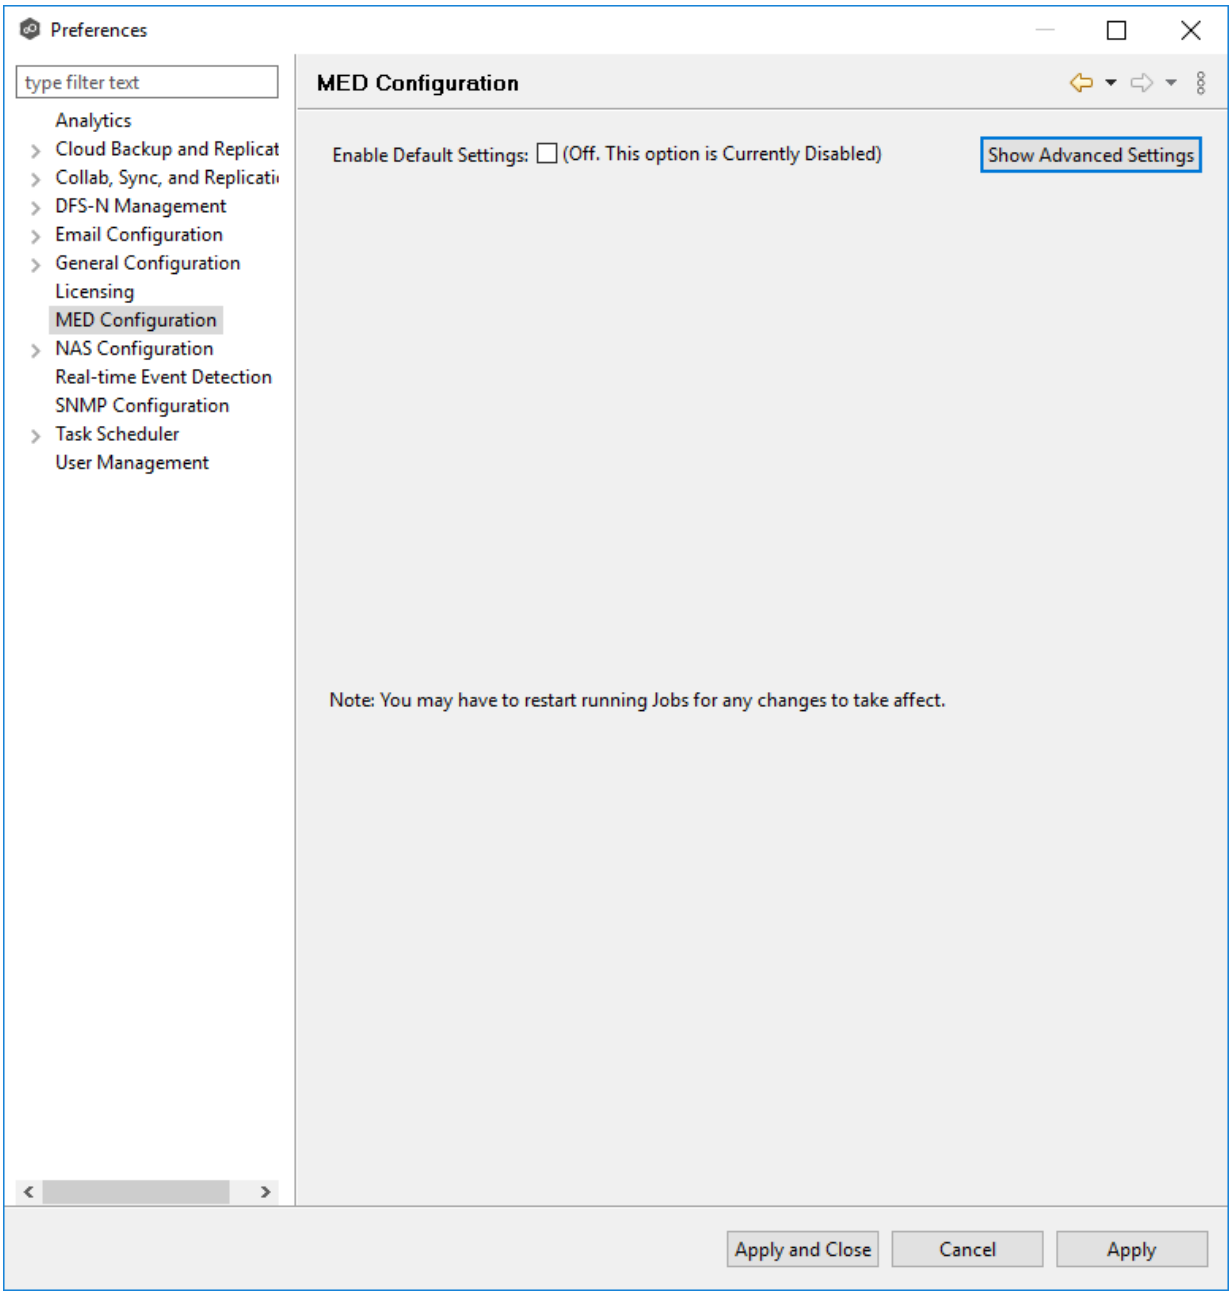

3. Select the **Enable Default Settings** or click **Show Advanced Settings**.

   If you selected **Show Advanced Settings**, the following is displayed.

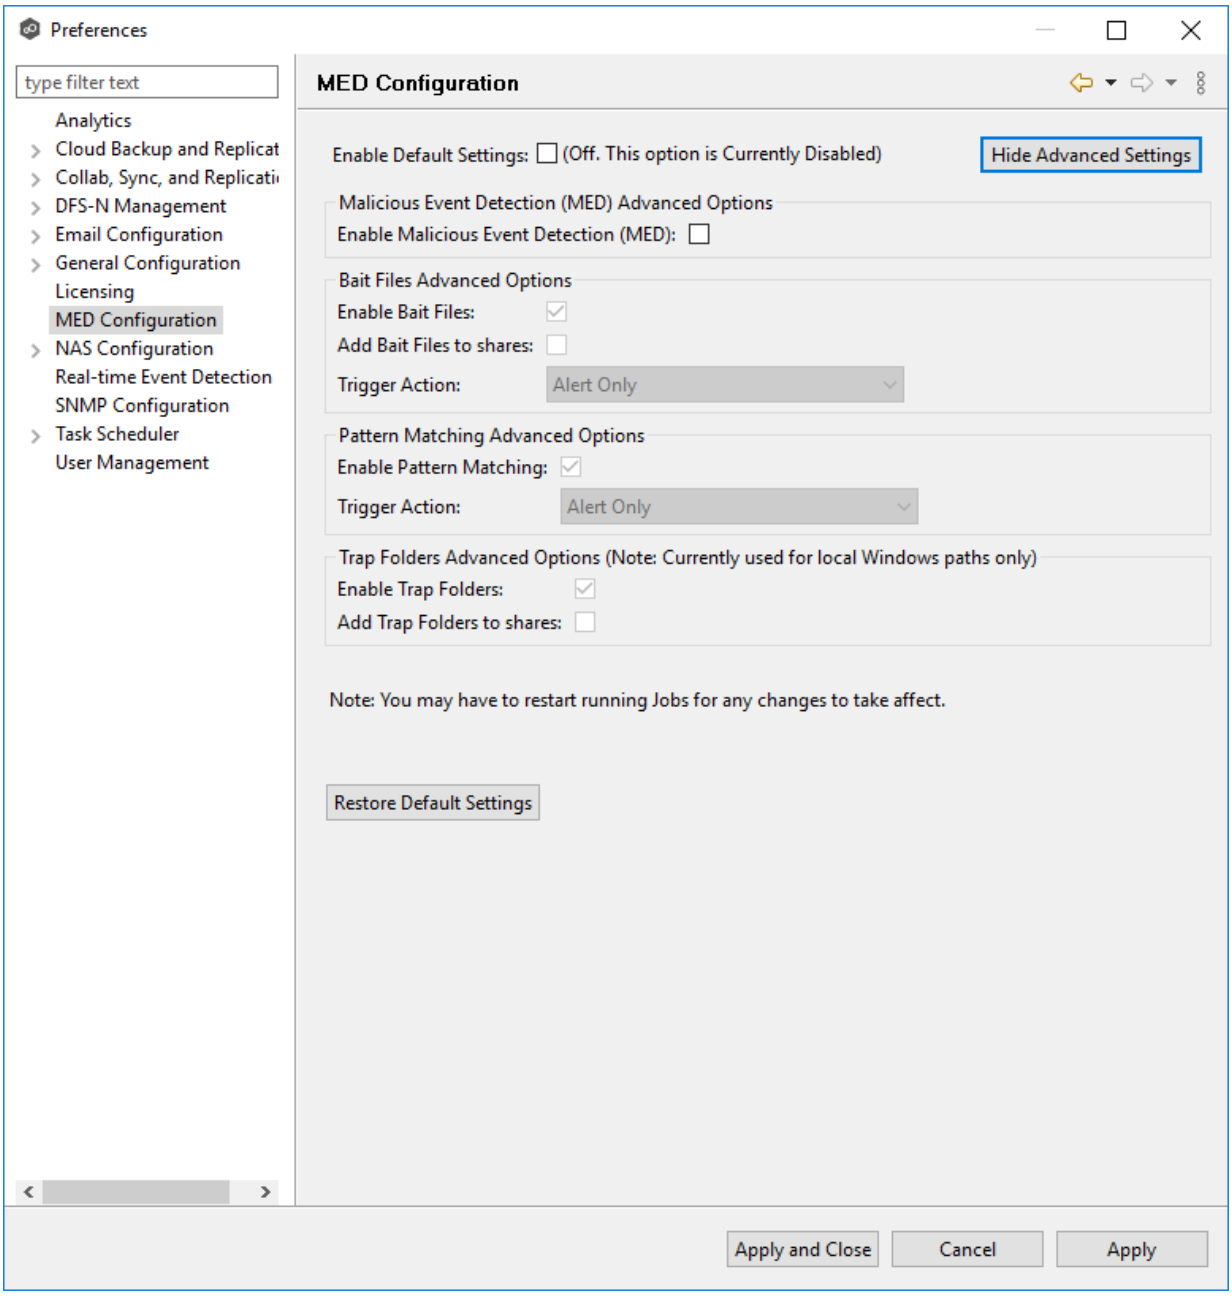

4. Modify the options as needed:
   - Primary MED Options
   - Bait File Advanced Options
   - Trap Folders Advanced Options
5. Click **OK**.

Primary MED Options

The main options are as follows:

Option	Description
Enable Default Settings	Enables/disables Peer MED using default settings. By default, all three MED mechanisms are enabled.
Show/Hide Advanced Settings	Shows/hides options for each of the three MED mechanisms.
Enable Malicious Event Detection (MED)	The master on/off switch for MED. If unchecked, all MED mechanisms will be disabled.
Restore Default Settings	Restores all defaults across the three MED mechanisms.

Bait File Advanced Options

Bait files are files of common types, inserted into the file system in a way that hides them from users. Though hidden, these bait files are likely to be accessed by automated processes (like ransomware) or by mass deletions of entire folder structures. As soon as these files are touched, an action is triggered.

The options for bait files are:

Option	Description
Enable Bait Files	Enables/disables bait file creation and monitoring.
Add Bait Files to shares	At the start of each job, creates bait files under the root of each participant's watch directory. To see the watch directory for a job, review Host Participants and Directories.

Option	Description
Trigger Action	Defines the action to take when MED detects malicious activity on a bait file. See Action Types for more details on available actions.

Action Types

For each MED mechanism, one of four actions can be configured on the detection of malicious activity. These actions are:

Action	Description
Alert Only	Triggers an alert in Peer Management Center. If email alerts are configured for MED Alerts and enabled for a job, an email will also be sent. See Email Alerts in the Basic Concepts section for more information about email alerts. If SNMP traps are configured for MED Alerts and enabled for a job, an SNMP trap will also be sent. See SNMP Notifications in the Basic Concepts section for more information about SNMP notifications.
Alert and Disable Host	Triggers an alert while also removing the afflicted Agent from the job in which the malicious activity was detected. Once disabled, Agents will need to be manually re-enabled for collaboration to resume. See Re-enabling a Disabled Agent Within a Job for details.
Alert and Stop Job	Triggers an alert while also stopping the job where the malicious activity was detected. Jobs will need to be restarted in order for collaboration to resume.
Alert, Disable Host and Stop Job	Triggers an alert, removes the afflicted Agent from the job where the malicious activity was detected, and stops the job. This option is the most aggressive and will require administrators to re-enable Agents as well as restart jobs. See Re-enabling a Disabled Agent Within a Job for details.

An example of an alert as displayed in Peer Management Center is as follows:

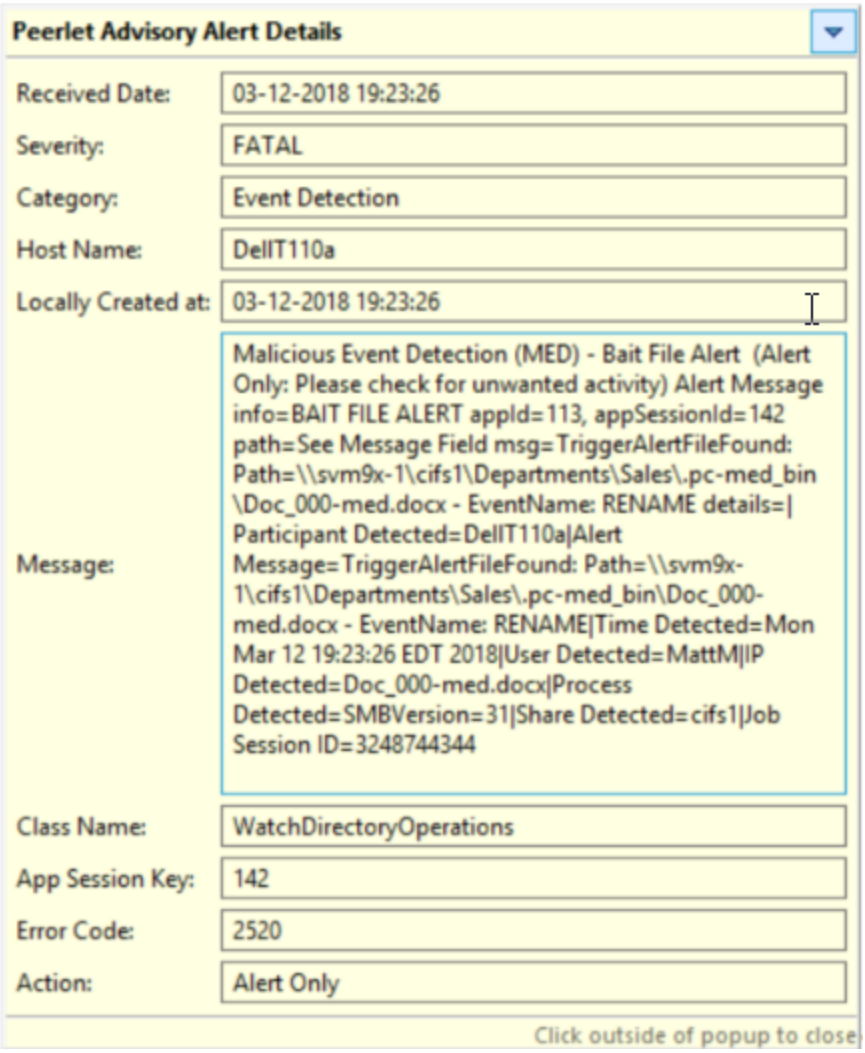

Trap Folders Advanced Options

On Windows file servers, Peer MED can be configured to create hidden, recursive folders that attempt to trap or slowdown ransomware as it enumerates a folder structure. As with the bait files, these folders cannot be seen by users but will be accessible by automated processes. If bait files (above) are enabled, a bait file will be placed within each trap folder, and an action will be triggered as soon as these files are touched.

Options for trap folders are:

Option	Description
Enable Trap Folders	Enables/disables the creation and monitoring of trap folders.

Option	Description
Add Trap Folders to shares	At the start of each job, create trap folders under the root of each participant's configured watch directory. To see the watch directory for a job, review Host Participants and Directories. Note: Trap Folders will only be used with participants that are Windows file servers. As such, these settings will not apply to any other enterprise NAS device.

NAS Configuration

This section contains information about configuring your NAS for use with Peer Global File System:

- Dell EMC Configurations
- NetApp 7-Mode Configurations
- NetApp cDOT Configurations
- Nutanix Configurations

Dell EMC Configurations

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more Dell EMC storage devices within most available job types. These Dell EMC devices can be running Isilon, Unity, or VNX. For detailed information about Dell EMC prerequisites, see Dell EMC Prerequisites.

To create a new Dell EMC configuration:

1. Select **Preferences** from the **Window** menu.
2. Select **NAS Configuration** in the navigation tree.
3. Select **Dell EMC Configurations**.

   The **Dell EMC Configurations** page is displayed. It lists existing configurations.

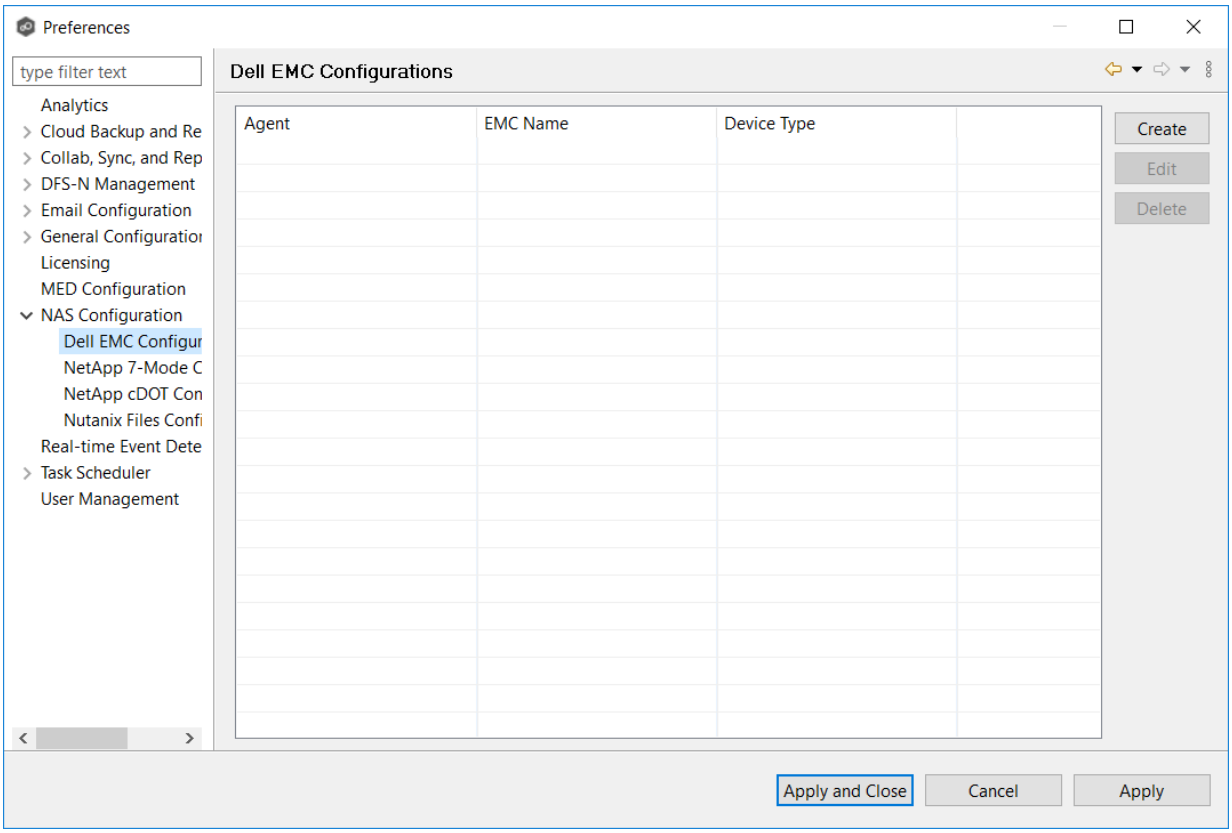

4. Click the **Create** button.

   The **Management Agent** page appears.

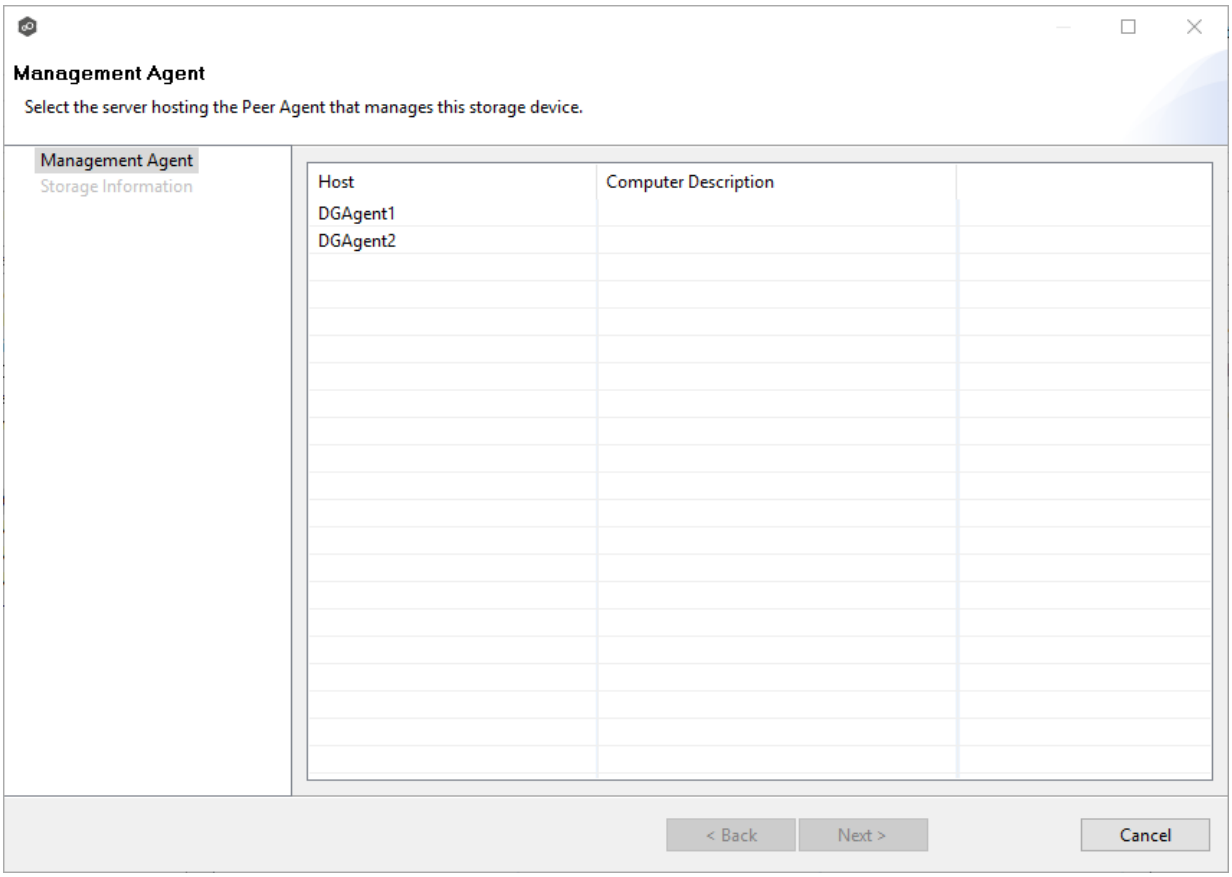

5. Select a Management Agent, and then click **Next**.

   The **Storage Information** page appears. The fields in the **Credentials** section vary, depending on the selected EMC platform type; VNX is selected by default.

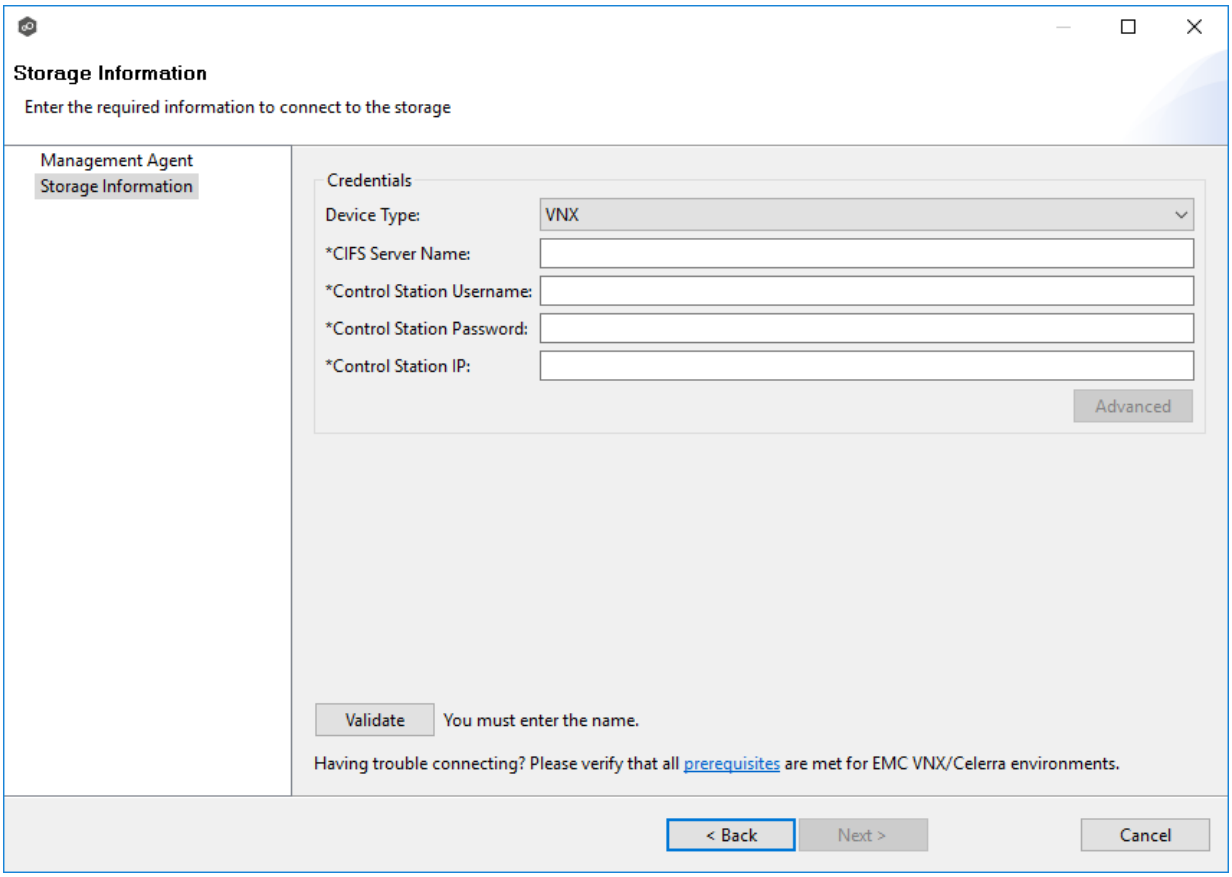

6. Select the device type, and then enter the required values in **Credentials**:

   Dell EMC Isilon Credentials

   Dell EMC Unity Credentials

   Dell EMC VNX Credentials

7. (Optional) Click the **Advanced** button if you want to specify advanced options, and then enter the required values:

   Dell EMC Isilon Advanced Options

   Dell EMC Unity Advanced Options

   Dell EMC VNX Advanced Options

8. Click **Validate**.

9. Click **Next**.

10. Click **OK**.

1. Enter the required values.

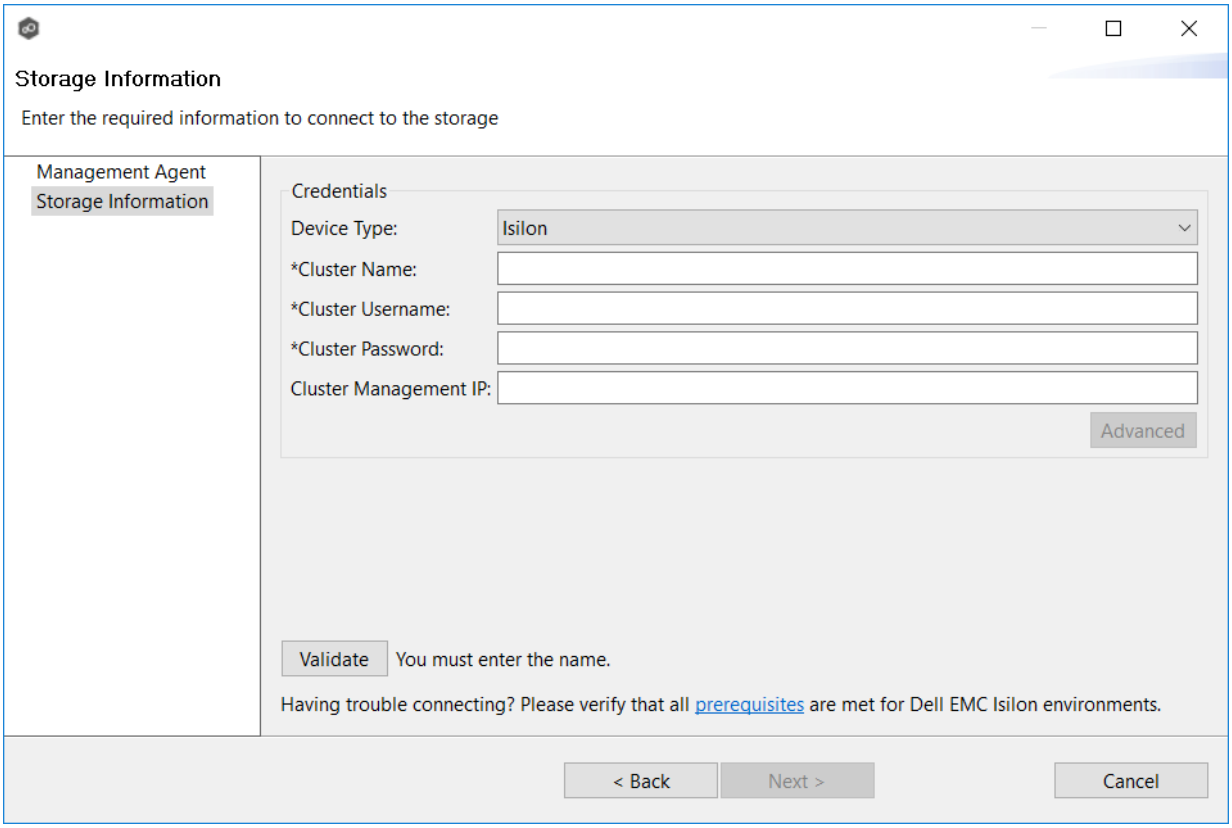

Field	Description
Cluster Name	Enter the name of the EMC Isilon cluster hosting the data to be replicated.
Cluster Username	Enter the user name for the account managing the EMC Isilon cluster.

Field	Description
Cluster Password	Enter the password for account managing the EMC Isilon cluster.
Cluster Management IP	Enter the IP address of the system used to manage the EMC Isilon cluster. Required only if multiple Access Zones are in use on the cluster.

2. (Optional) Click Advanced and enter the required values.
3. Click **Validate**.
4. Click **Finish**.

Dell EMC Isilon Advanced Options

The options are divided into two groups:

- Dell EMC Isilon Options for this job
- Dell EMC Isilon Advanced Settings

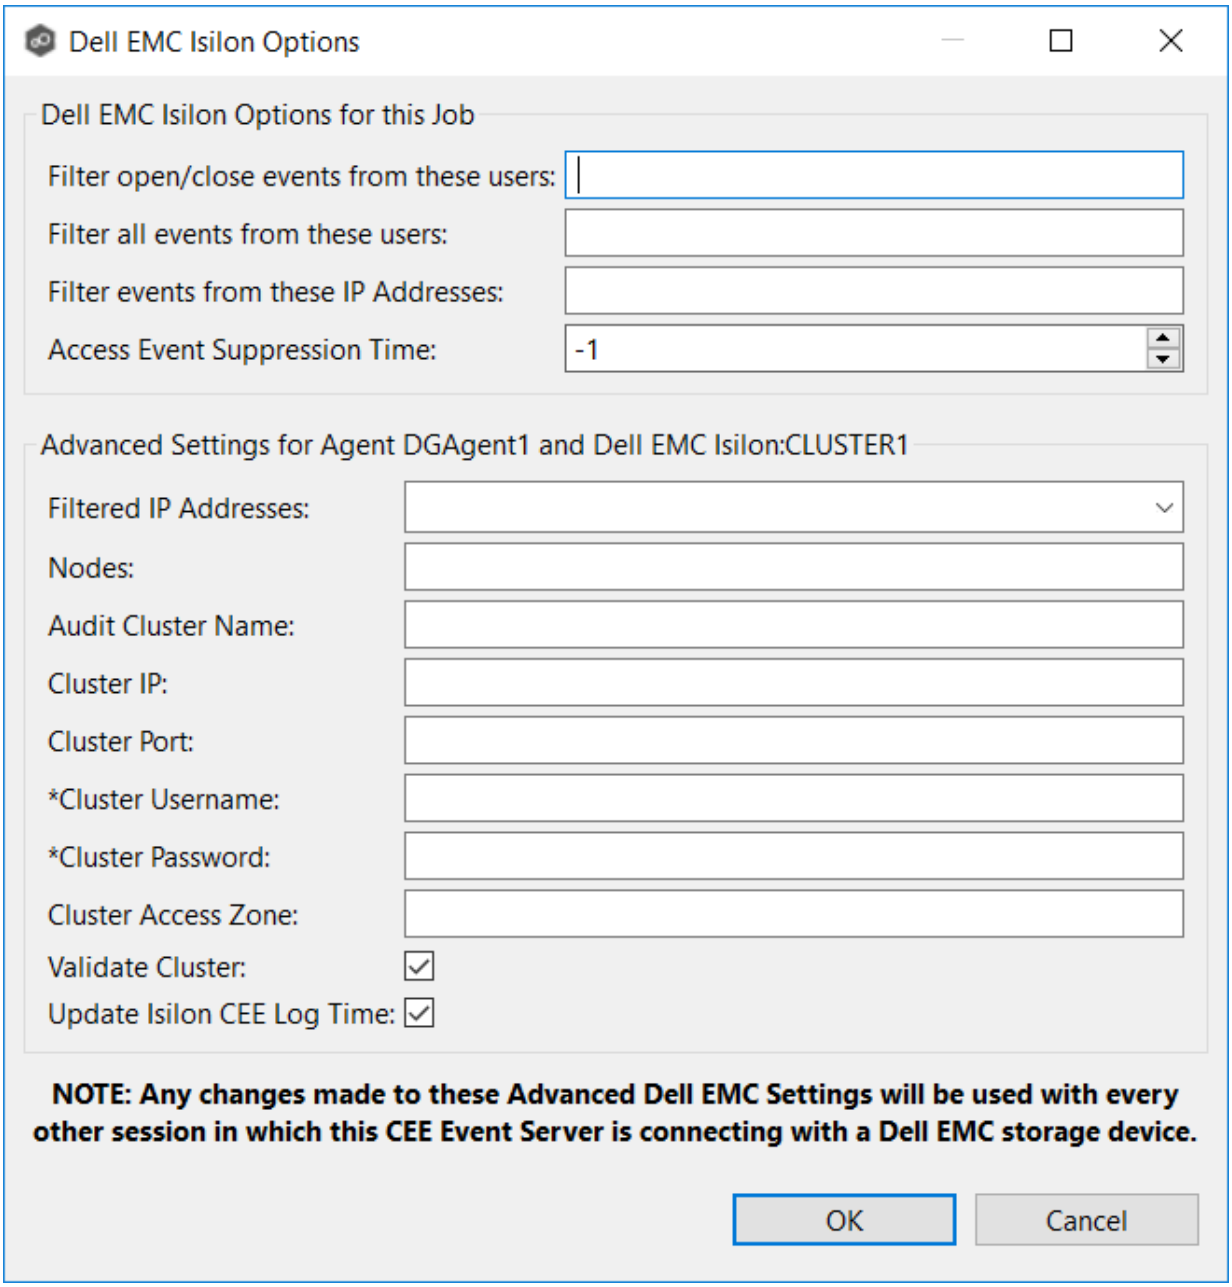

Dell EMC Isilon Options for this Job

The following configuration options are available for Dell EMC Isilon devices jobs:

Option	Description
Filter open/close events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Filter all events from these users	A comma-separated list of user names to exclude from all event detection. For example, if "USER"1 is excluded, any activity generated by USER1 will be ignored, e.g., file is open and modified.
Filter events from these IP Addresses	A comma-separated list of IP addresses to exclude from all event detection. For example, if "192.168.0.100" is excluded, any activity generated by 192.168.0.100 will be ignored, e.g., file is open and modified.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will be adjusted based on the selected NAS platform. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

Dell EMC Isilon Advanced Settings

The following advanced settings are available for Dell EMC Isilon devices:

Option	Description
Filtered IP Addresses	Optional. Events generated from these IP addresses will be filtered. It is recommended that the IP address of the CEE Server is added to this list.
Nodes	Optional. Comma-delimited listed of additional node IP address to query for open files. These addresses must be accessible from the CEE Server where the Agent is running.
Audit Cluster Name	Optional. The hostname that is set in the Isilon audit system configuration.
Cluster IP	Optional. The cluster IP address of the Isilon system.
Custer Port	Optional. The cluster port number of the Isilon system. Default value is 8080.
Cluster Username	The user name used to sign into the Isilon cluster.
Cluster Password	The password used to sign into the Isilon cluster.
Cluster Access Zone	Optional. The name of the access zone that is being monitored.
Validate Cluster	If enabled, the Isilon cluster will be validated both on registration and periodically by a maintenance thread.
Update Isilon CEE Log Time	If enabled, the audit log time on the Isilon cluster will be set to the start time of the first job to communicate with this Isilon cluster.

1. Enter the required values.

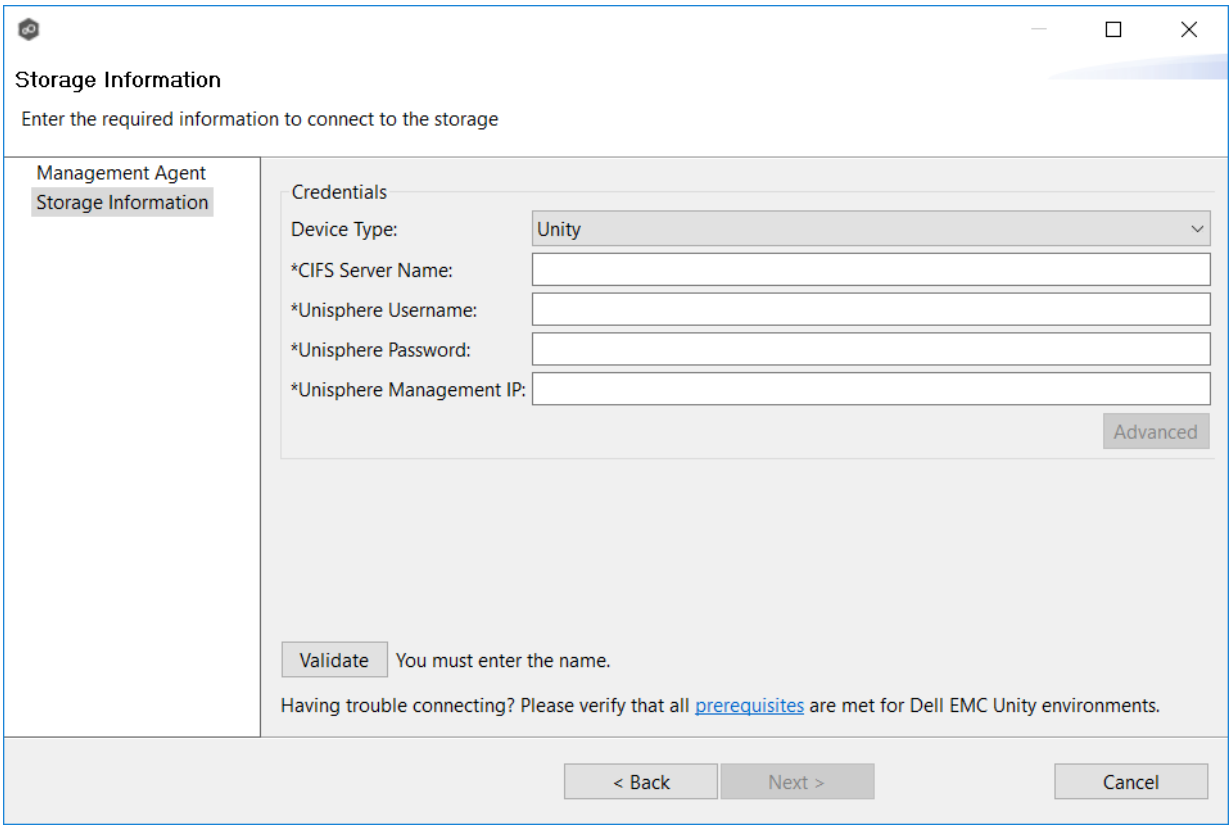

Field Description	Description
CIFS Server Name	Enter the name of the CIFS server hosting the data to be replicated.
Unisphere Username	Enter the user name for the Unisphere account managing the Unity storage device.
Unisphere Password	Enter the password for the Unisphere account managing the Unity storage device.

Field Description	Description
Unisphere Management IP	Enter the IP address of the Unisphere system used to manage the Unity storage device. This should not point to the NAS server.

2. (Optional) Click Advanced and enter the required values.
3. Click **Validate**.
4. Click **Finish**.

Dell EMC Unity Advanced Options

The options are divided into two groups:

- Dell EMC Unity Options for this Job
- Dell EMC Unity Advanced Settings

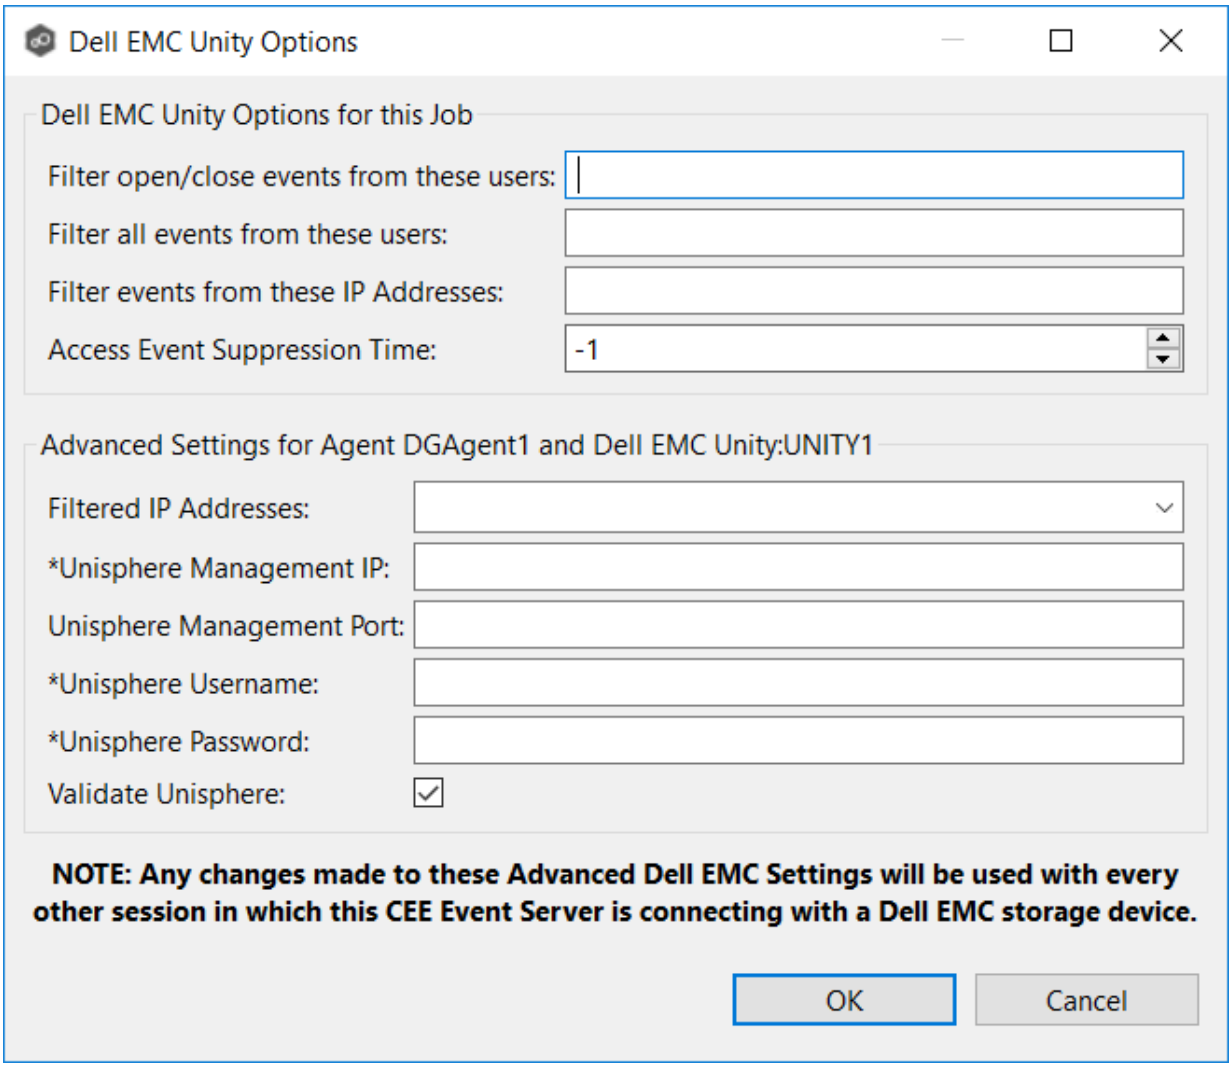

Dell EMC Unity Options for this Job

The following configuration options are available for EMC Unity devices:

Option	Description
Filter open/close events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Filter all events from these users	A comma-separated list of user names to exclude from all event detection. For example, if "USER1" is excluded, any activity generated by USER1 will be ignored, e.g., file is open and modified.
Filter events from these IP Addresses	A comma-separated list of IP addresses to exclude from all event detection. For example, if "192.168.0.100" is excluded, any activity generated by 192.168.0.100 will be ignored, e.g., file is open and modified.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will be adjusted based on the selected NAS platform. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

Dell EMC Unity Advanced Settings

The following advanced settings are available for Dell EMC Unity devices:

Option	Description
Filtered IP Addresses	Optional. Events generated from these IP addresses will be filtered. It is recommended that the IP address of the CEE Server is added to this list.
Unisphere Management IP	The Unisphere Management IP address of the Unity system. This address is used for making API calls to validate configuration.
Unisphere Management Port	Optional. The Unisphere Management port number of the Unity system. Default value is 443.
Unisphere Username	The user name used to sign into Unisphere.
Unisphere Password	The password used to sign into Unisphere.
Validate Unisphere	Optional. If enabled, Unisphere settings will be validated both on registration and periodically by a maintenance thread.

1. Enter the required values.

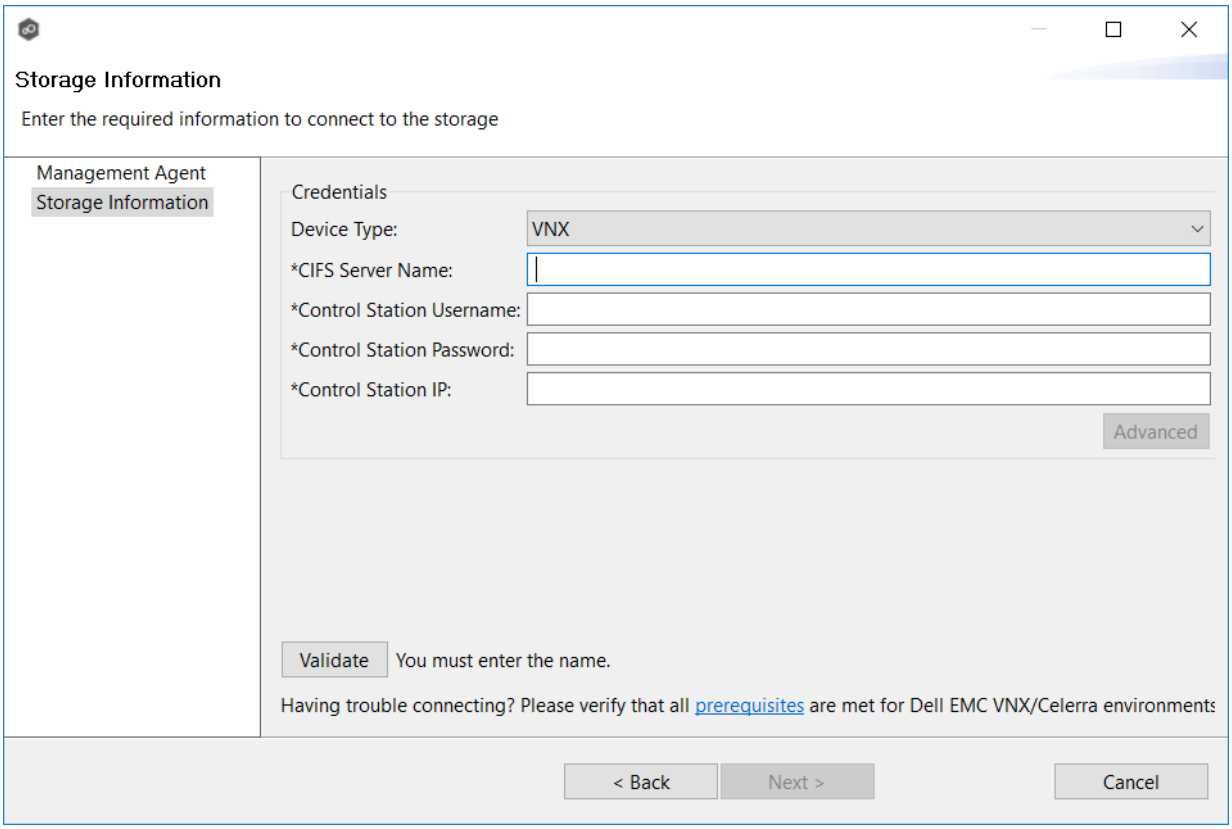

Field	Description
CIFS Server Name	Enter the name of the CIFS Server hosting the data to be replicated.
Control Station Username	Enter the user name for the Control Station account managing the VNX/Celerra storage device.
Control Station Password	Enter the password for the Control Station account managing the VNX/Celerra storage device.

Field	Description
Control Station IP	Enter the IP address of the Control Station system used to manage the VNX/Celerra storage device. This should not point to the CIFS Server.

2. (Optional) Click Advanced and enter the required values.
3. Click **Validate**.
4. Click **Finish**.

Dell EMC VNX/Celerra Advanced Options

The options are divided into two groups:

- Dell EMC VNX Options for this Job
- Dell EMC VNX Advanced Settings

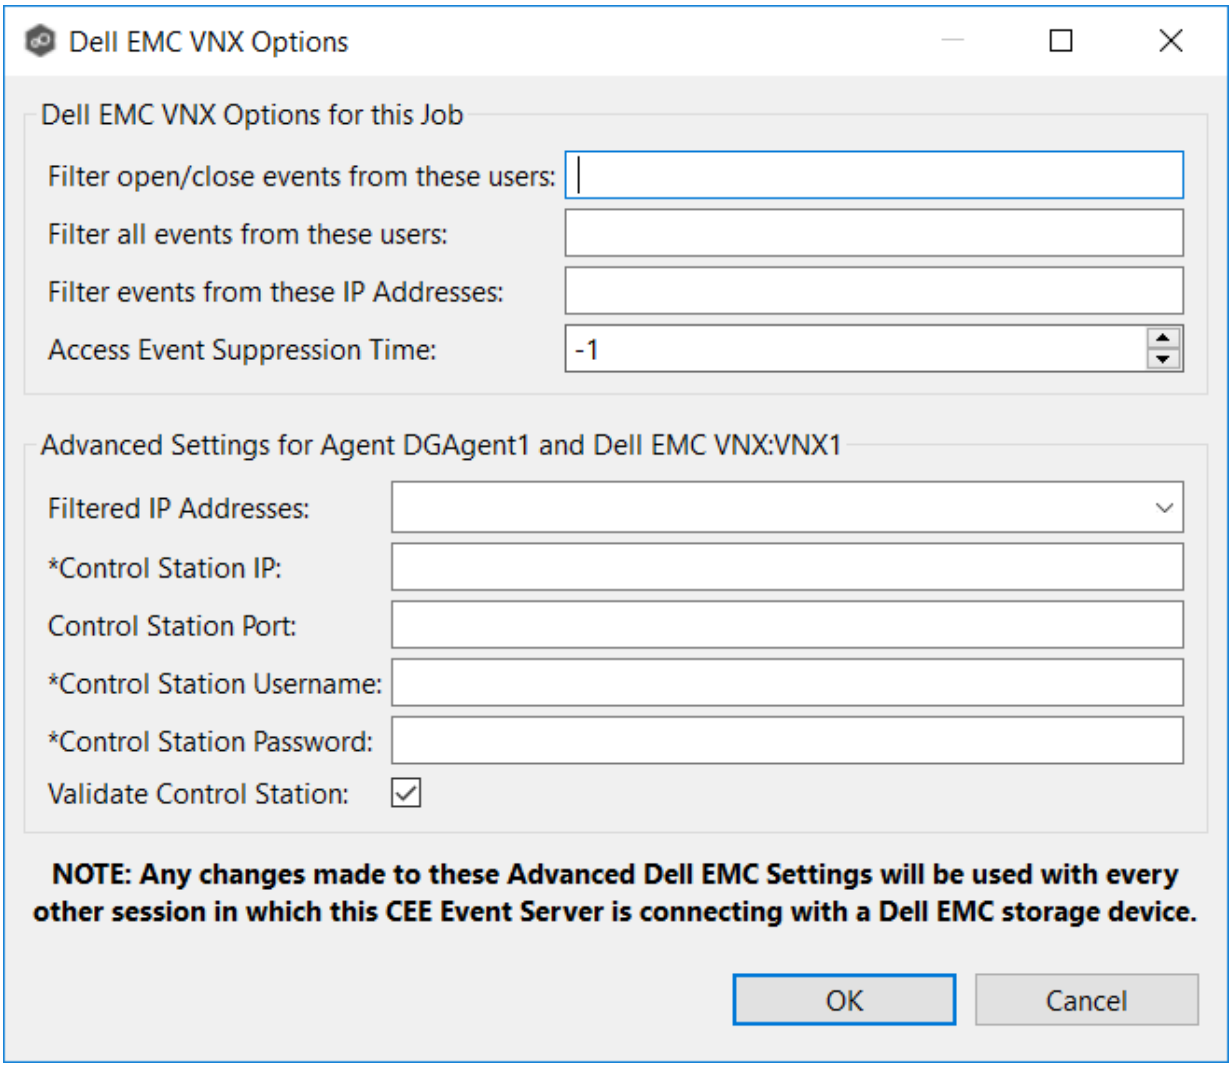

Dell EMC VNX/Celerra Options for this Job

The following configuration options are available for Dell EMC VNX/Celerra devices:

Option	Description
Filter open/close events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Filter all events from these users	A comma-separated list of user names to exclude from all event detection. For example, if "USER1" is excluded, any activity generated by USER1 will be ignored, e.g., file is open and modified.
Filter events from these IP Addresses	A comma-separated list of IP addresses to exclude from all event detection. For example, if "192.168.0.100" is excluded, any activity generated by 192.168.0.100 will be ignored, e.g., file is open and modified.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will be adjusted based on the selected NAS platform. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

Dell EMC VNX/Celerra Advanced Settings

The following advanced settings are available for Dell EMC VNX/Celerra devices:

Option	Description
Filtered IP Addresses	Optional. Events generated from these IP addresses will be filtered. It is recommended that the IP address of the CEE Server is added to this list.
Control Station IP	The Control Station IP address of the VNX/Celerra system.
Control Station Port	Optional. The Control Station Port number of the VNX/Celerra system. The default value is 443.
Control Station Username	The user name used to sign into the VNX/Celerra Control Station.
Control Station Password	The password used to sign into the VNX/Celerra Control Station.
Validate Control Station	Optional. If enabled, the VNX/Celerra Control Station will be validated both on registration and periodically by a maintenance thread.

NetApp 7-Mode Configurations

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more NetApp storage devices within most available job types. These NetApp devices can be running Data ONTAP 7-Mode or clustered Data ONTAP. In order to work with NetApp devices, Peer Management Center leverages the FPolicy API built into the NetApp device. For detailed information about NetApp prerequisites and configuration, see NetApp Prerequisites.

To create a new NetApp 7-Mode configuration:

1. Select **Preferences** from the **Window** menu.
2. Select **NAS Configuration** in the navigation tree.
3. Select **NetApp 7-Mode Configurations**.

The **NetApp 7-Mode Configurations** page is displayed. It lists existing configurations.

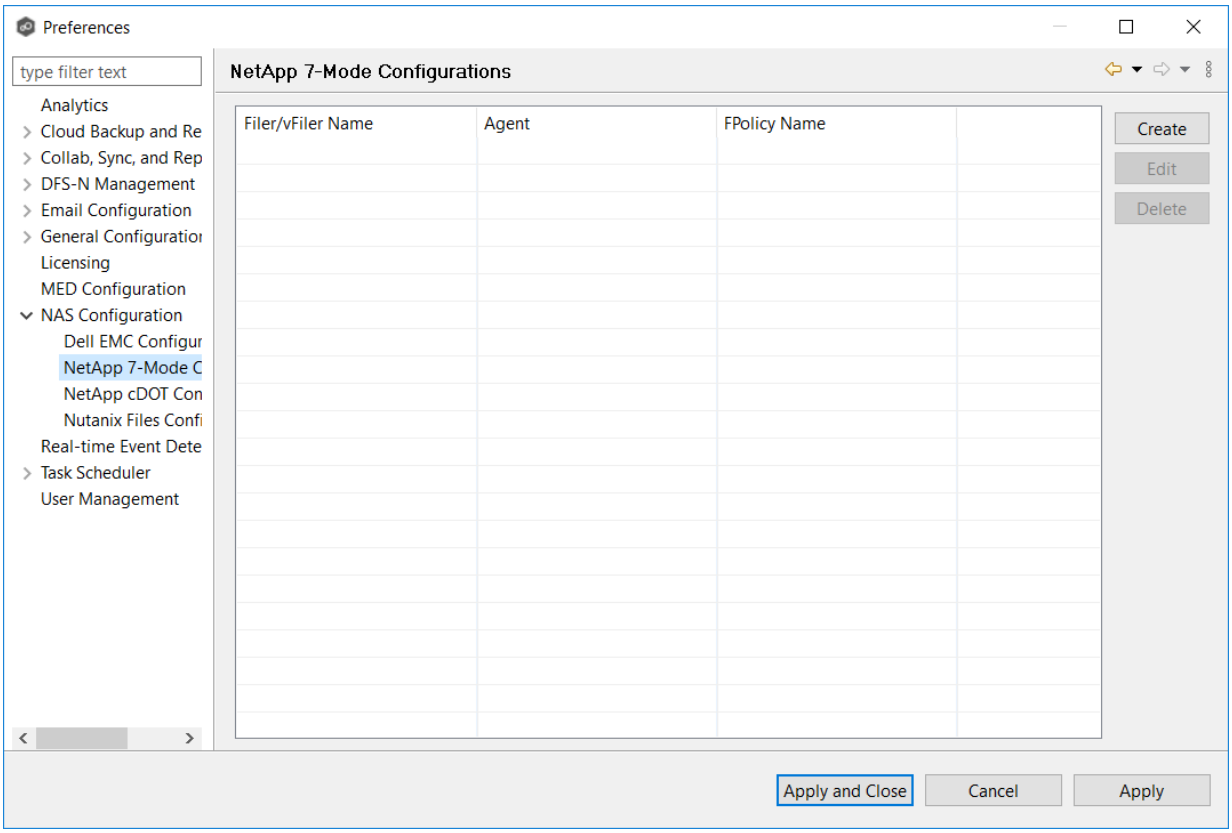

4. Click the **Create** button.

   The **Management Agent** page appears.

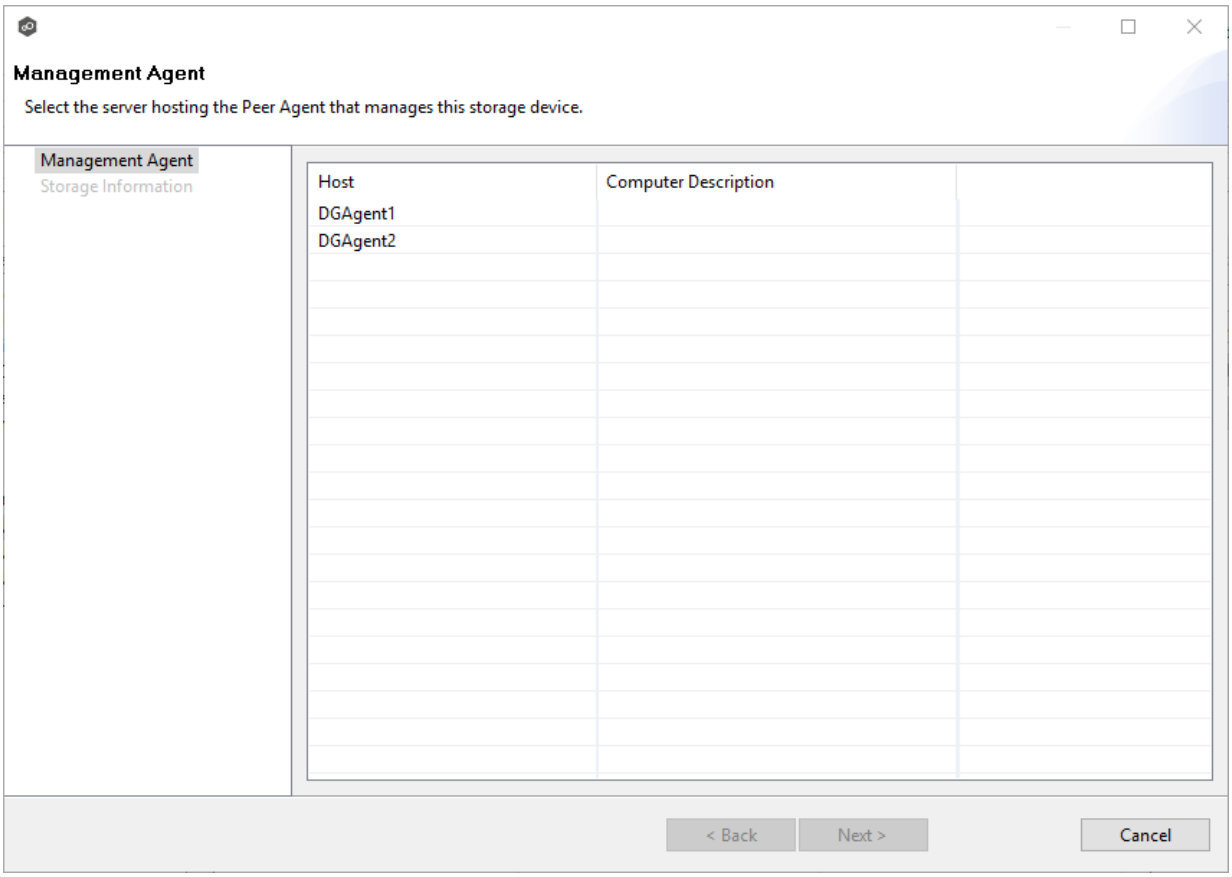

5. Select a Management Agent, and then click **Next**.

   The **Storage Information** page appears.

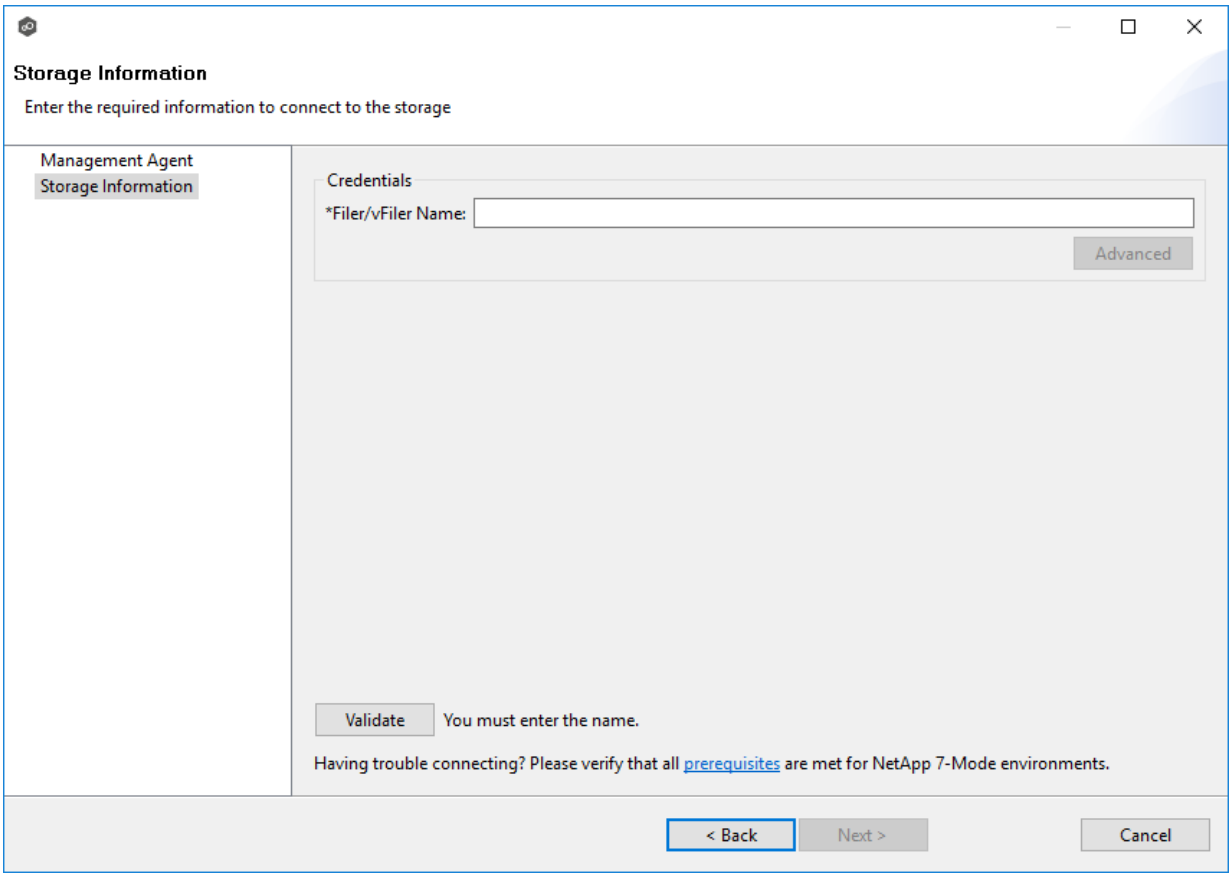

6. Enter the required values in **Credentials**.

Field	Description
Filer Name	Enter the name of the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

7. (Optional) Click Advanced and enter the required values.
8. Click **Validate**.
9. Click **Next**.
10. Click **OK**.

The options are divided into two groups:

- NetApp 7-Mode Options for this Job
- NetApp FPolicy 7-Mode Advanced Settings

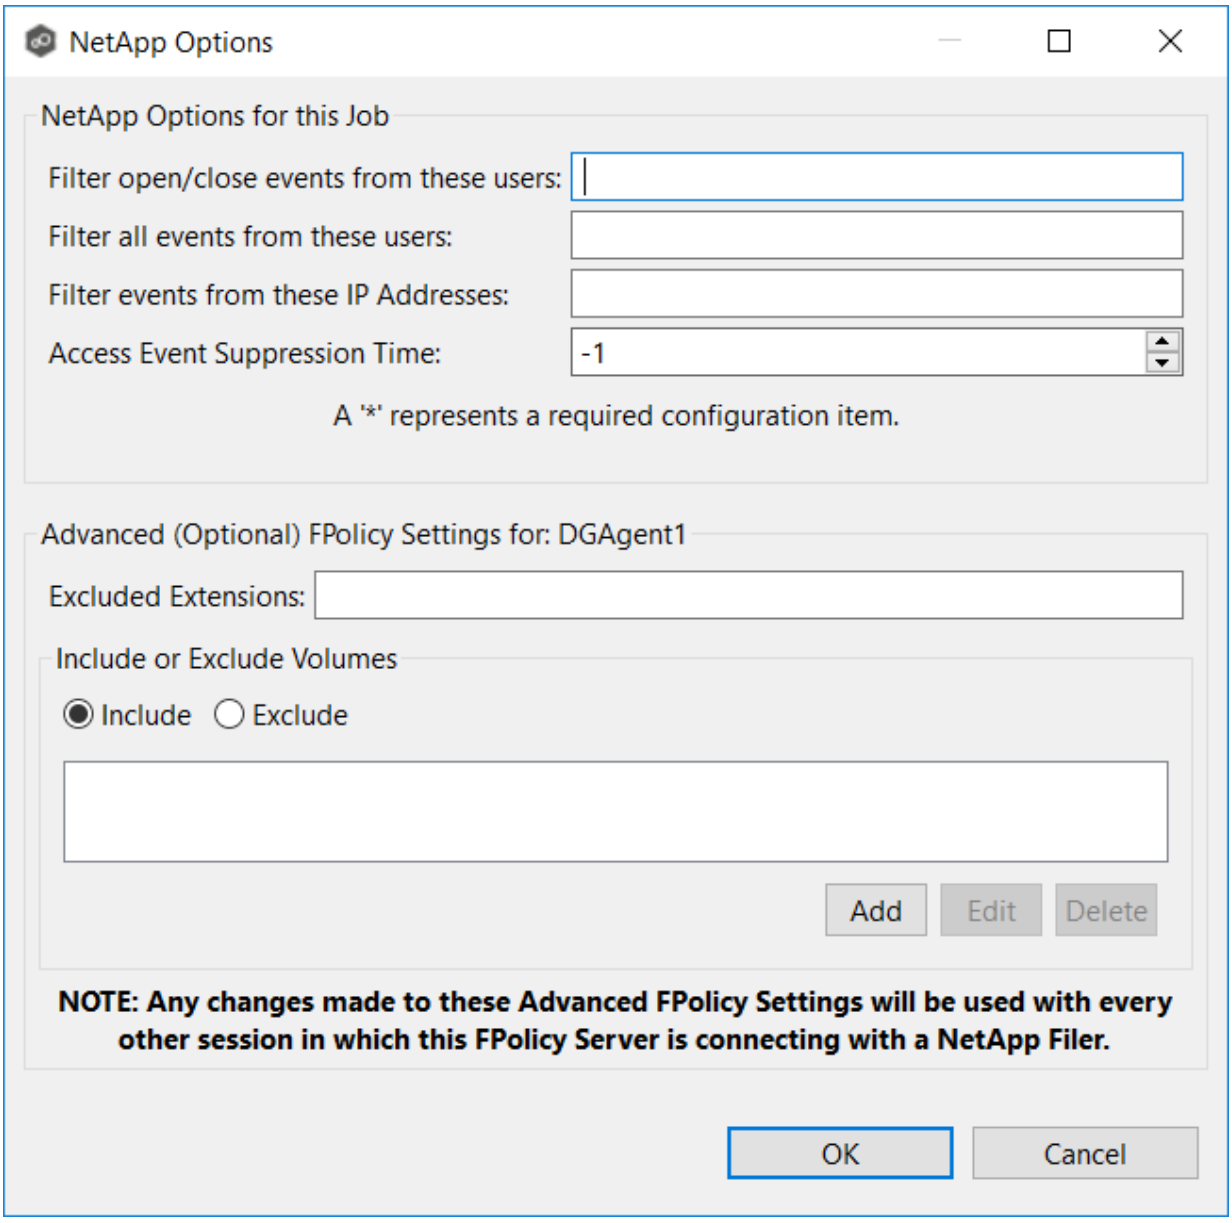

NetApp 7-Mode Options for this Job

The following configuration options are available for NetApp 7-Mode devices:

Option	Description
Filter open/close events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Filter all events from these users	A comma-separated list of user names to exclude from all event detection. For example, if "USER1" is excluded, any activity generated by USER1 will be ignored, e.g., file is open and modified.
Filter events from these IP Addresses	A comma-separated list of IP addresses to exclude from all event detection. For example, if "192.168.0.100" is excluded, any activity generated by 192.168.0.100 will be ignored, e.g., file is open and modified.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will be adjusted based on the selected NAS platform. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

NetApp FPolicy 7-Mode Advanced Settings

Option	Description
Excluded Extensions	Extensions entered here are excluded from event detection on the NetApp Filer. Values are comma-separated and must not contain any periods. FPolicy enables you to restrict a policy to a certain list of file extensions by excluding extensions that need to be screened. Note: The maximum length of a file name extension supported for screening is 260 characters. Screening by extensions is based only on the characters after the last period (.) in the file name. For example, for a file named fle1.txt.name.jpg, file access notification takes place only if a file policy is configured for the jpg extension.
Include or Exclude Volumes	List all volumes on the NetApp Filer to exclude or include based on selected choice. FPolicy enables you to restrict a policy to a certain list of volumes by including or excluding volumes that need to be screened. Using the include list, you can request notifications for the specified volume list. Using the exclude list, you can request notifications for all volumes except the specified volume list. However, by default, both the include and exclude list are empty. You can use the question mark (?) or asterisk (*) wildcard characters to specify the volume. The question mark (?) wildcard character stands for a single character. For example, entering vol? in a list of volumes that contain vol1, vol2, vol23, voll4, will result in only vol1 and vol2 being matched. The asterisk (*) wildcard character stands for any number of characters that contain the specified string. Entering *test* in a list of volumes to exclude from file screening excludes all volumes that contain the string such as test_vol and vol_test.

NetApp cDOT Configurations

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more NetApp storage devices within most available job types. These NetApp devices can be running Data ONTAP 7-Mode or clustered Data ONTAP. In order to work with NetApp devices, Peer Management Center leverages the FPolicy API built into the NetApp device. For detailed information about NetApp prerequisites and configuration, see NetApp Prerequisites.

To create a new NetApp cDOT configuration:

1. Select **Preferences** from the **Window** menu.
2. Select **NAS Configuration** in the navigation tree.
3. Select **NetApp cDot Configurations**.

   The **NetApp cDOT Configurations** page is displayed. It lists existing configurations.

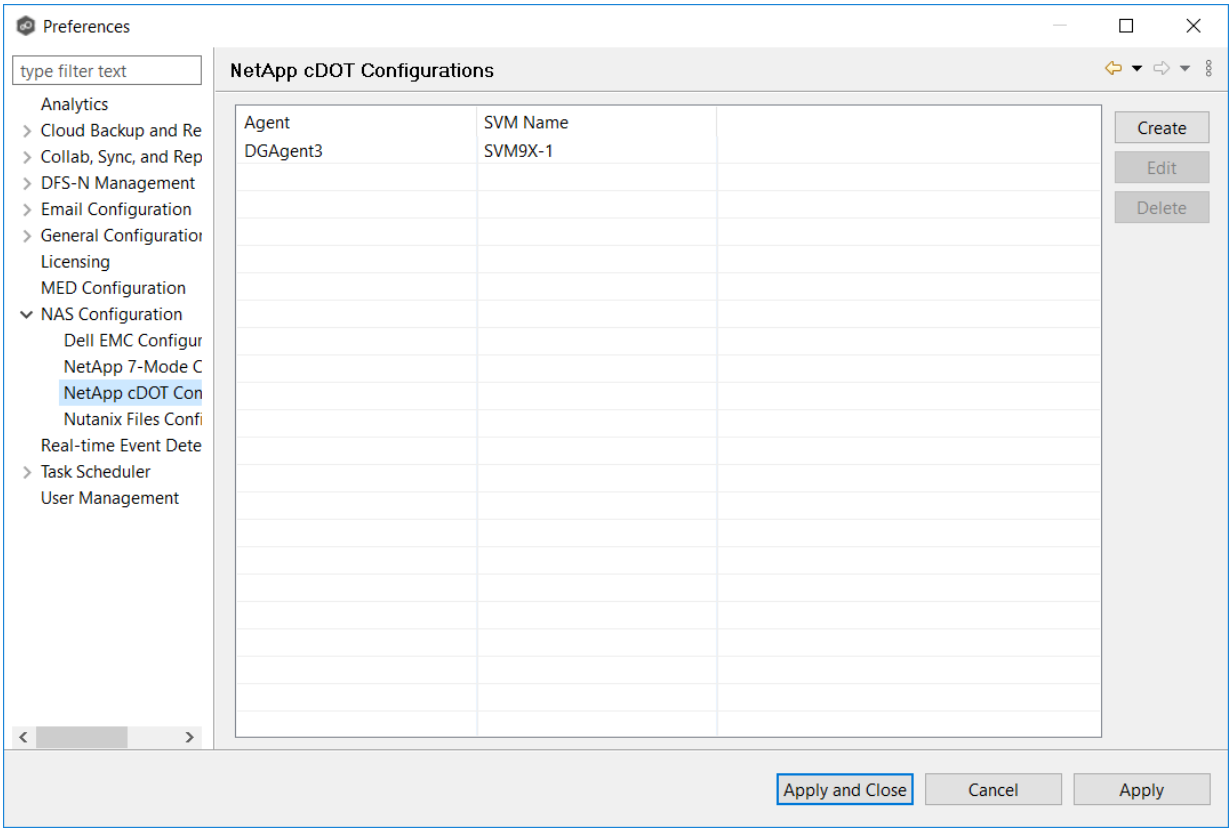

4. Click the **Create** button.

   The **Management Agent** page appears.

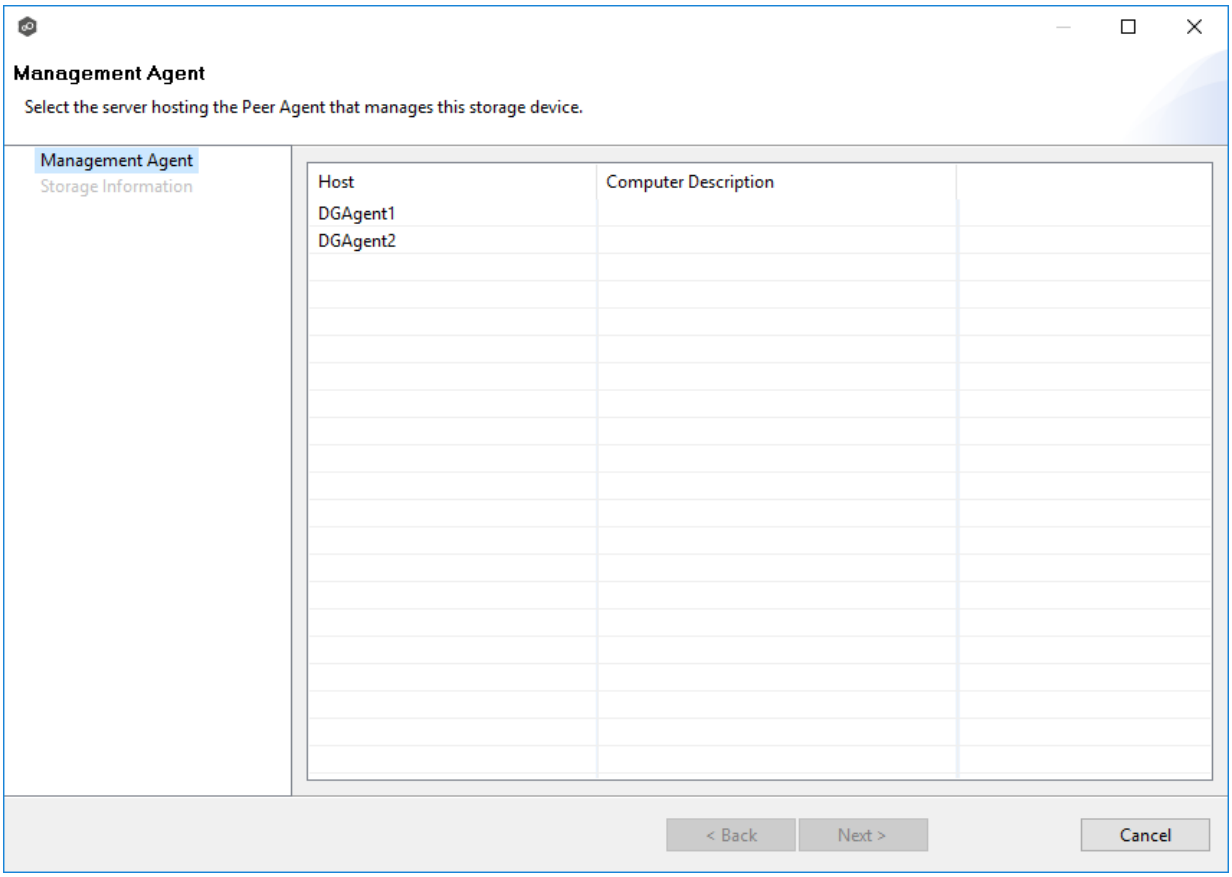

5. Select a Management Agent, and then click **Next**.

   The **Storage Information** page appears.

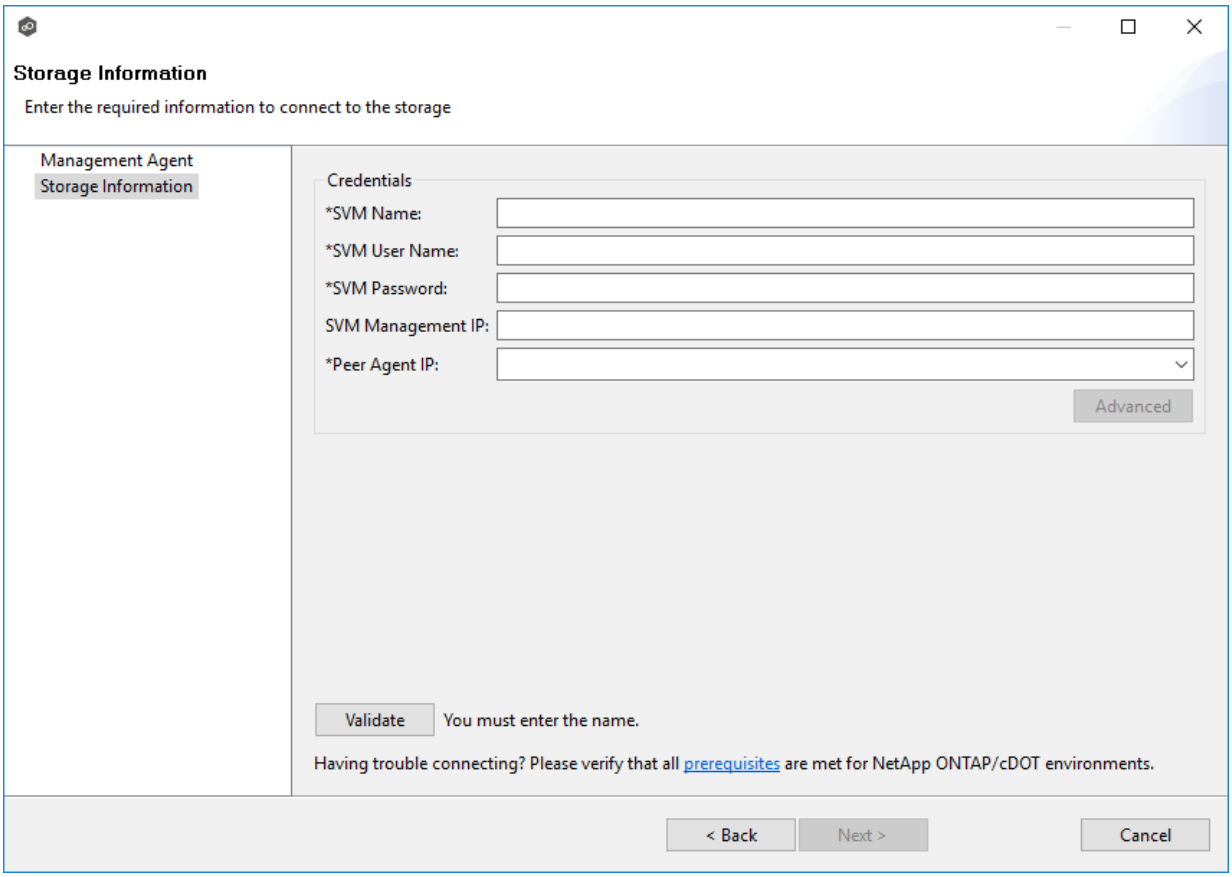

6. Enter the required values in **Credentials**.

Field	Description
SVM Name	Enter the name of the Storage Virtual Machine hosting the data to be replicated.
SVM Username	Enter the user name for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	Enter the IP address used to access the management API of the NetApp Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already allow management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address.

7. (Optional) Click Advanced and enter the required values.
8. Click **Validate**.
9. Click **Next**.
10. Click **OK**.

The options are divided into two groups:

- NetApp cDOT Options for this Job
- NetApp Advanced FPolicy cDOT Settings

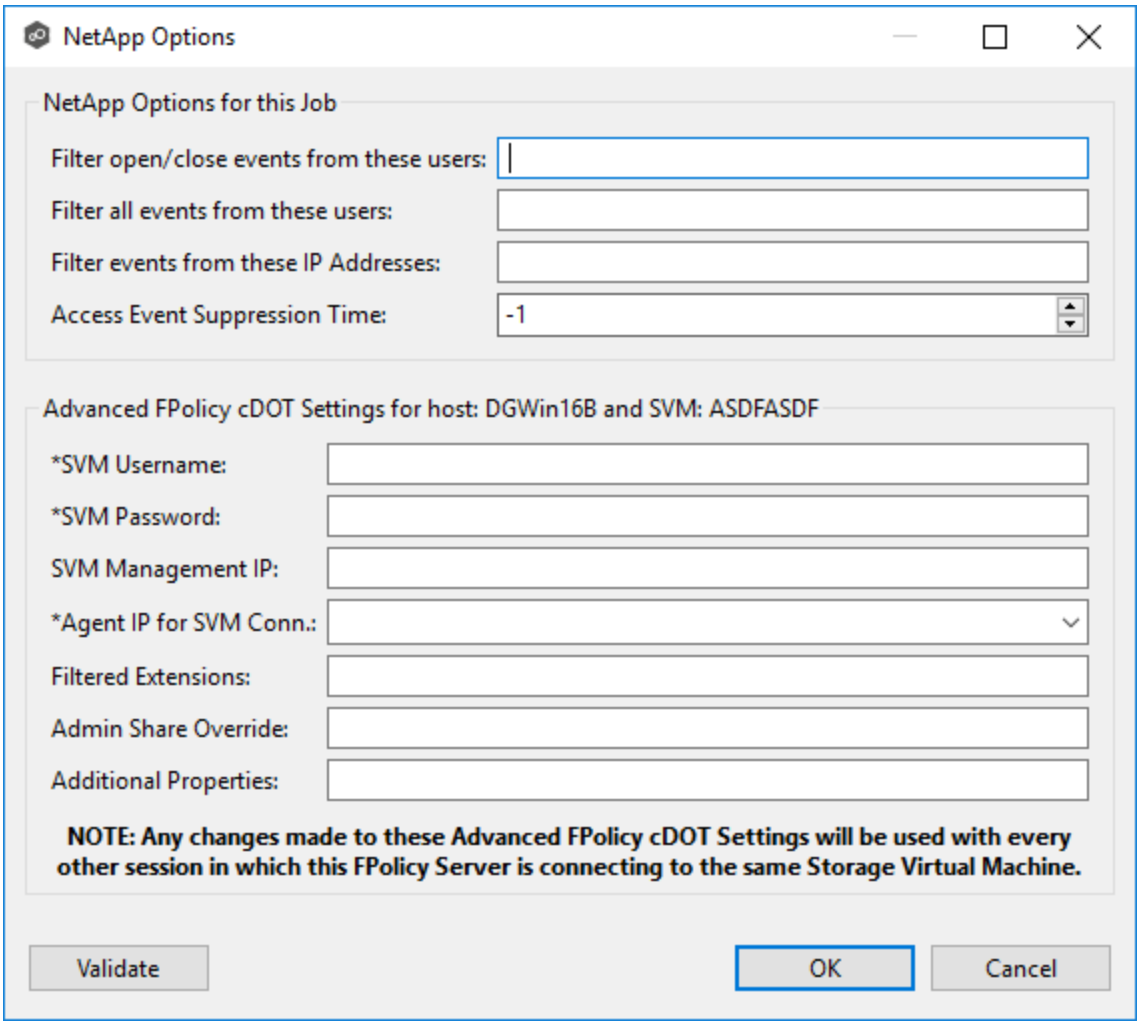

NetApp cDOT Options for this Job

The following configuration options are available for NetApp cDOT devices:

Option	Description
Filter open/close events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Filter all events from these users	A comma-separated list of user names to exclude from all event detection. For example, if "USER1" is excluded, any activity generated by USER1 will be ignored, e.g., file is open and modified.
Filter events from these IP Addresses	A comma-separated list of IP addresses to exclude from all event detection. For example, if "192.168.0.100" is excluded, any activity generated by 192.168.0.100 will be ignored, e.g., file is open and modified.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will be adjusted based on the selected NAS platform. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

NetApp Advanced FPolicy cDOT Settings

Option	Description
SVM Username	The account name of the VSAdmin or similar account on the SVM that has the appropriate access to ONTAPI.
SVM Password	The password of the VSAdmin or similar account on the SVM that has the appropriate access to ONTAPI. This value will be encrypted.
SVM Management IP	Optional. If the primary data LIF for the SVM (whose IP address is registered in DNS) does not support management calls, enter the management IP address of SVM.
Agent IP for SVM Conn.	The IP address over which this Peer Agent will connect to the configured SVM. This MUST be an IP address.
Filtered Extensions	Optional. A comma-separated list of file extensions to exclude (without a leading asterisk (*).
Admin Share Override	Optional. Enter the administrative-type share that you created on the cDOT SVM. To take advantage of performance improvements when using this option, the share must be created at the root of the SVM's namespace (/). Ideally it should be named to something similar to PMCShare$ to prevent users from being able to see it.

Nutanix Files Configurations

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more Nutanix Files (formerly Acropolis File Services or AFS) clusters within most available job types. For detailed information about Nutanix prerequisites, see Nutanix Prerequisites.

To create a new Nutanix configuration:

1. Select **Preferences** from the **Window** menu.

   The **Preferences** dialog appears.

2. Select **NAS Configuration** in the navigation tree.

3. Select **Nutanix Configurations**.

   The **Nutanix Files Configurations** page is displayed. It lists existing configurations.

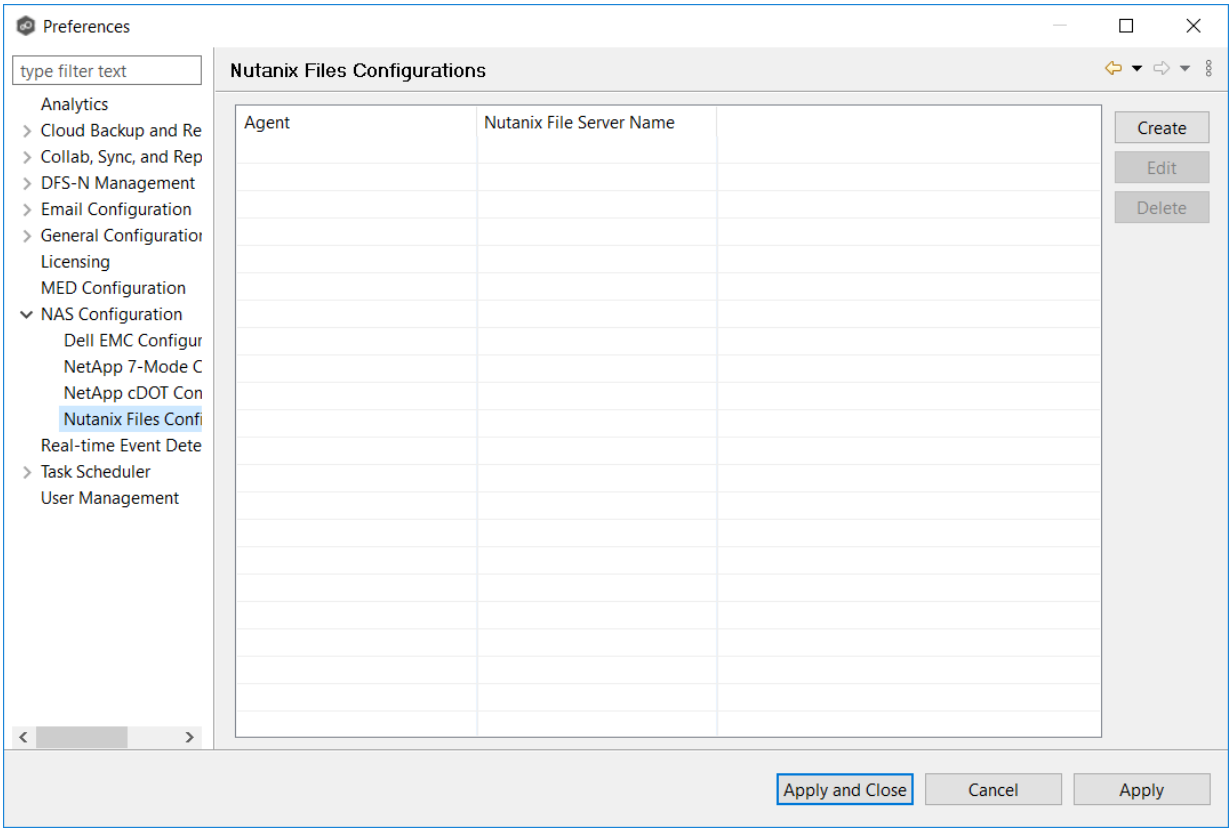

4. Click the **Create** button.

   The **Management Agent** page appears.

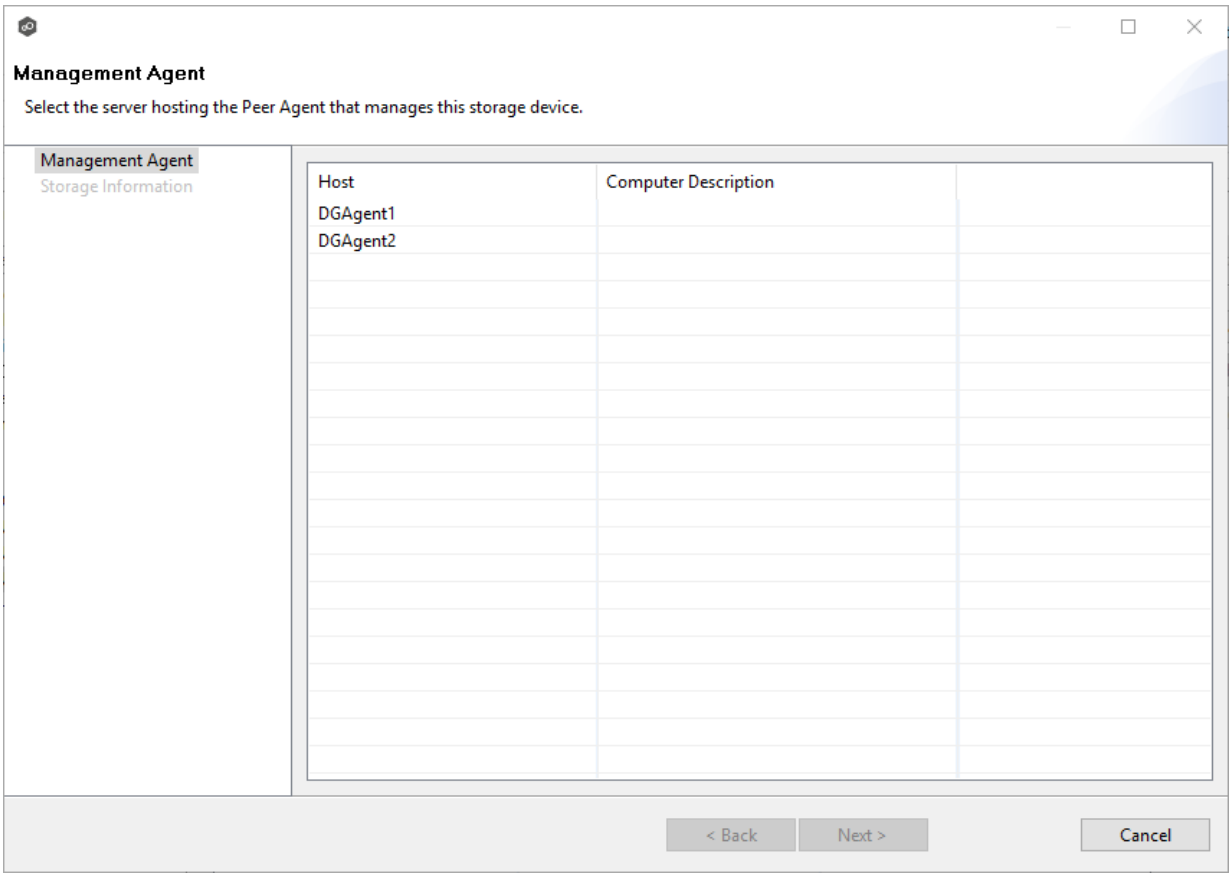

5. Select a Management Agent, and then click **Next**.

   The **Storage Information** page appears.

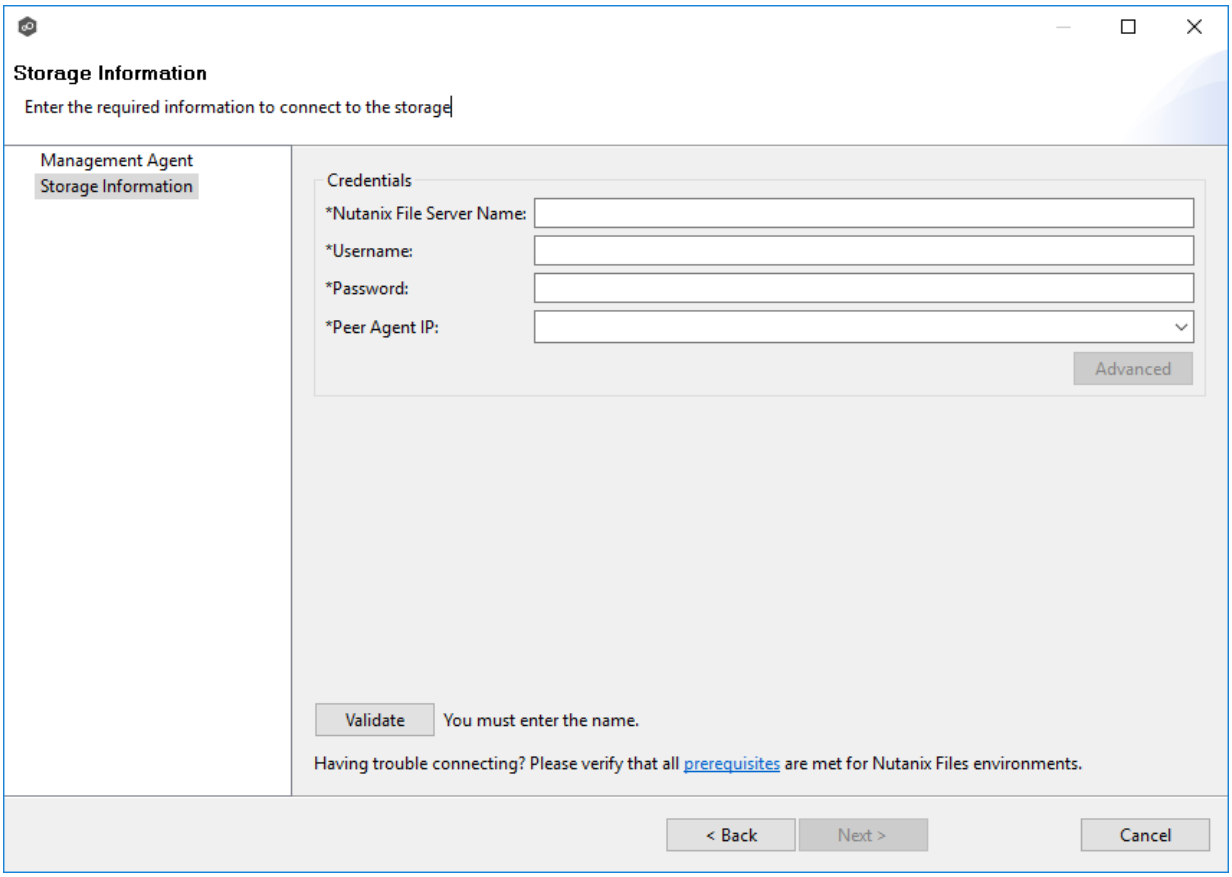

6. Enter the required values in **Credentials**.

Field	Description
Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address. This should not point to the Files cluster itself.

7. (Optional) Click Advanced button and then enter the required values.
8. Click **Validate**.
9. Click **Next**.
10. Click **OK**.

The options are divided into two groups:

- Nutanix Files Options for this Job
- Advanced Settings

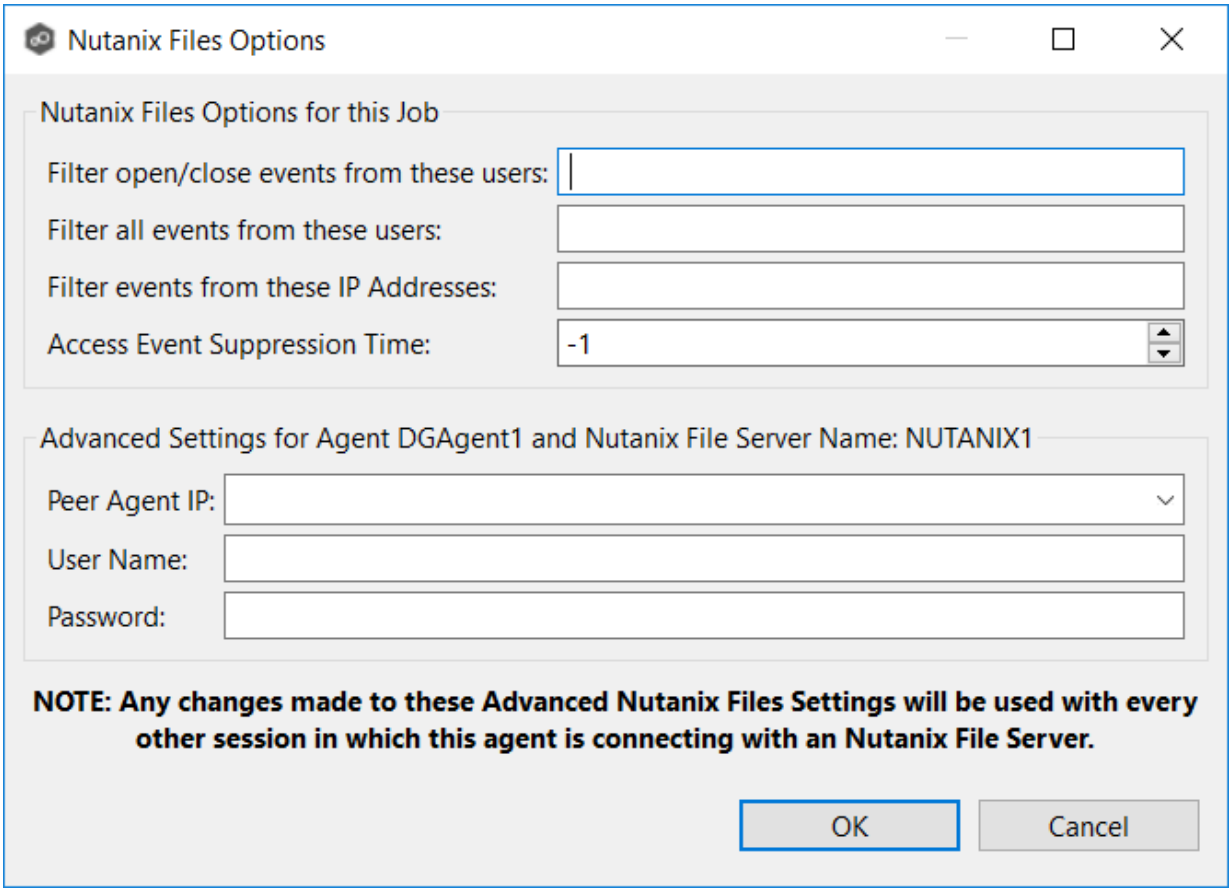

Nutanix Files Options for this Job

The following configuration options are available for Nutanix Files devices:

Option	Description
Filter open/close events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Filter all events from these users	A comma-separated list of user names to exclude from all event detection. For example, if "USER1" is excluded, any activity generated by USER1 will be ignored, e.g., file is open and modified.
Filter events from these IP Addresses	A comma-separated list of IP addresses to exclude from all event detection. For example, if "192.168.0.100" is excluded, any activity generated by 192.168.0.100 will be ignored, e.g., file is open and modified.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will be adjusted based on the selected NAS platform. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

Advanced Settings

The following advanced settings are available for Nutanix Files devices:

Option	Description
Peer Agent IP	The IP address over which the configured Files cluster will send activity to this Peer Agent. This must be an IP address.
User Name	User name used to access the APIs on the Files cluster.
Password	Password used to access the APIs on the Files cluster.

Real-time Event Detection Preferences

Several options are available to tune the way real-time event detection occurs. These options apply to all job types, except for DFS-N Management and PeerSync Management.

Note: There are also real-time event detection settings applicable to most job types in Peer Management Center. See Real-time Event Detection in the File Collab, Sync, and Repl, and Locking Preferences topic for more information.

To view and modify real-time event detection settings for all job types:

1. Select **Preferences** from the **Window** menu.
2. Select **Real-time Detection** in the navigation tree.

   The following page is displayed.

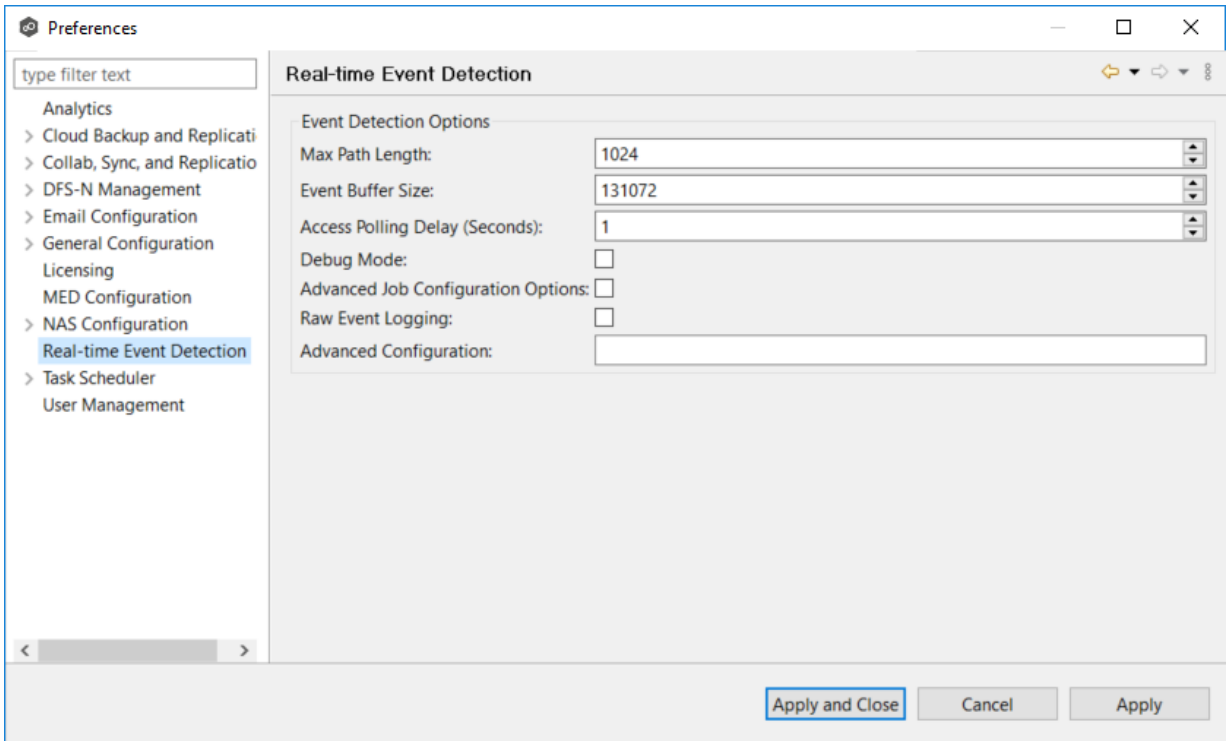

3. Modify values as needed:

Option	Description
Max Path Length	The maximum length in characters of a file or folder path that can be detected and worked with. In rare cases, this can be increased to 2048 or even 4096 but doing so will impact memory usage of the Peer Agents.
Event Buffer Size	The buffer size used by the Peer Agents to communicate with various Windows and enterprise NAS platform APIs.
Access Polling Delay (Seconds)	Controls how often a Peer Agent will poll a Windows File Server for its open files list.
Debug Mode	Turns on debug logging for real-time detection. This logs additional information that is often useful in troubleshooting issues but can increase overhead.
Advanced Job Configuration Options	When selected, enables advanced job-level options tied to real-time event detection.
Raw Event Logging	When selected, turns on raw logging. This logs every single event that we receive from a storage platform, even ones that we may be able to consolidate and coalesce. This additional information is often useful in troubleshooting issues but will increase overhead.
Advanced Configuration	A list of strings to enable advanced real-time detection options not found in the GUI. This should only be used when instructed by Peer Software support.

4. Click **OK** or **Apply**.

SNMP Configuration

Before Peer Management Center can send SMNP notifications on behalf of any job, a few key SNMP settings must be configured.

To configure SMNP settings:

1. Select **Preferences** from the **Window** menu.

2. Select **SNMP Configuration** in the navigation tree.

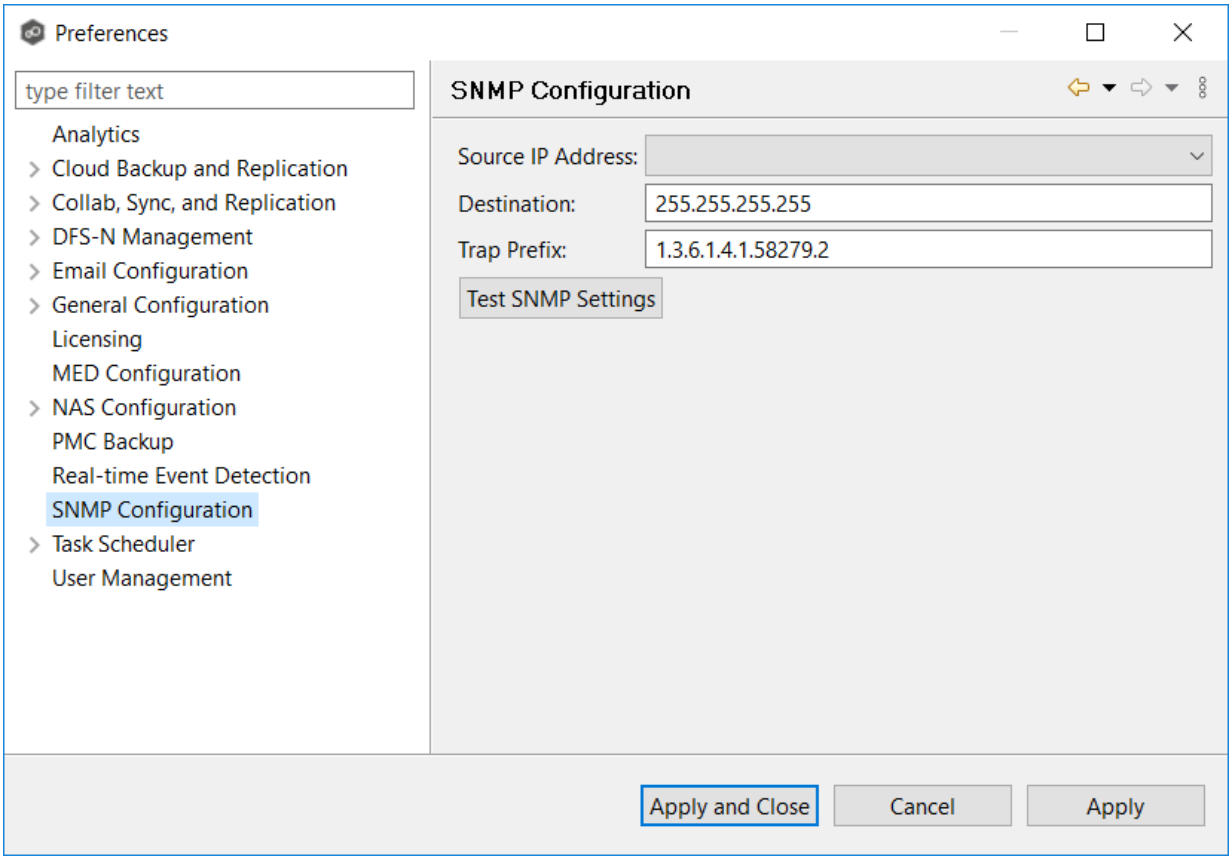

3. In the **Source IP Address** field, select or manually enter the IP address over which the trap will be sent.

4. In the **Destination** field, enter the destination host name, IP address, or broadcast address.

5. For **Trap Prefix**, enter a prefix that will help to identify whether the message is coming from different instances of Peer Management Center or from different jobs. In the default prefix, "1.3.5.1.4.1" represents IANA-registered private enterprises, "58279" is reserved for Peer Software, and the trailing ".2" represents the Peer Management Center.

6. Click **Test SNMP Settings**, and then click **OK** in the **Test Connection** dialog.

   You can verify the result by checking in your SNMP management tool.

7. Click **Apply and Close** or **Apply**.

User Management

Peer Management Center has both a rich client interface and a web client interface. The **User Management** page allows you to manage users of the web client interface. From this page, you can manage web client users, manage web roles, and configure Active Directory authentication.

Note: The User Management page can be accessed by any rich client user but only by web client users that have an **Administrator** role.

To access the User Management page:

1. Select **Preferences** from the **Window** menu.

2. Select **User Management** from the navigation tree.

   The **User Management** page is displayed:

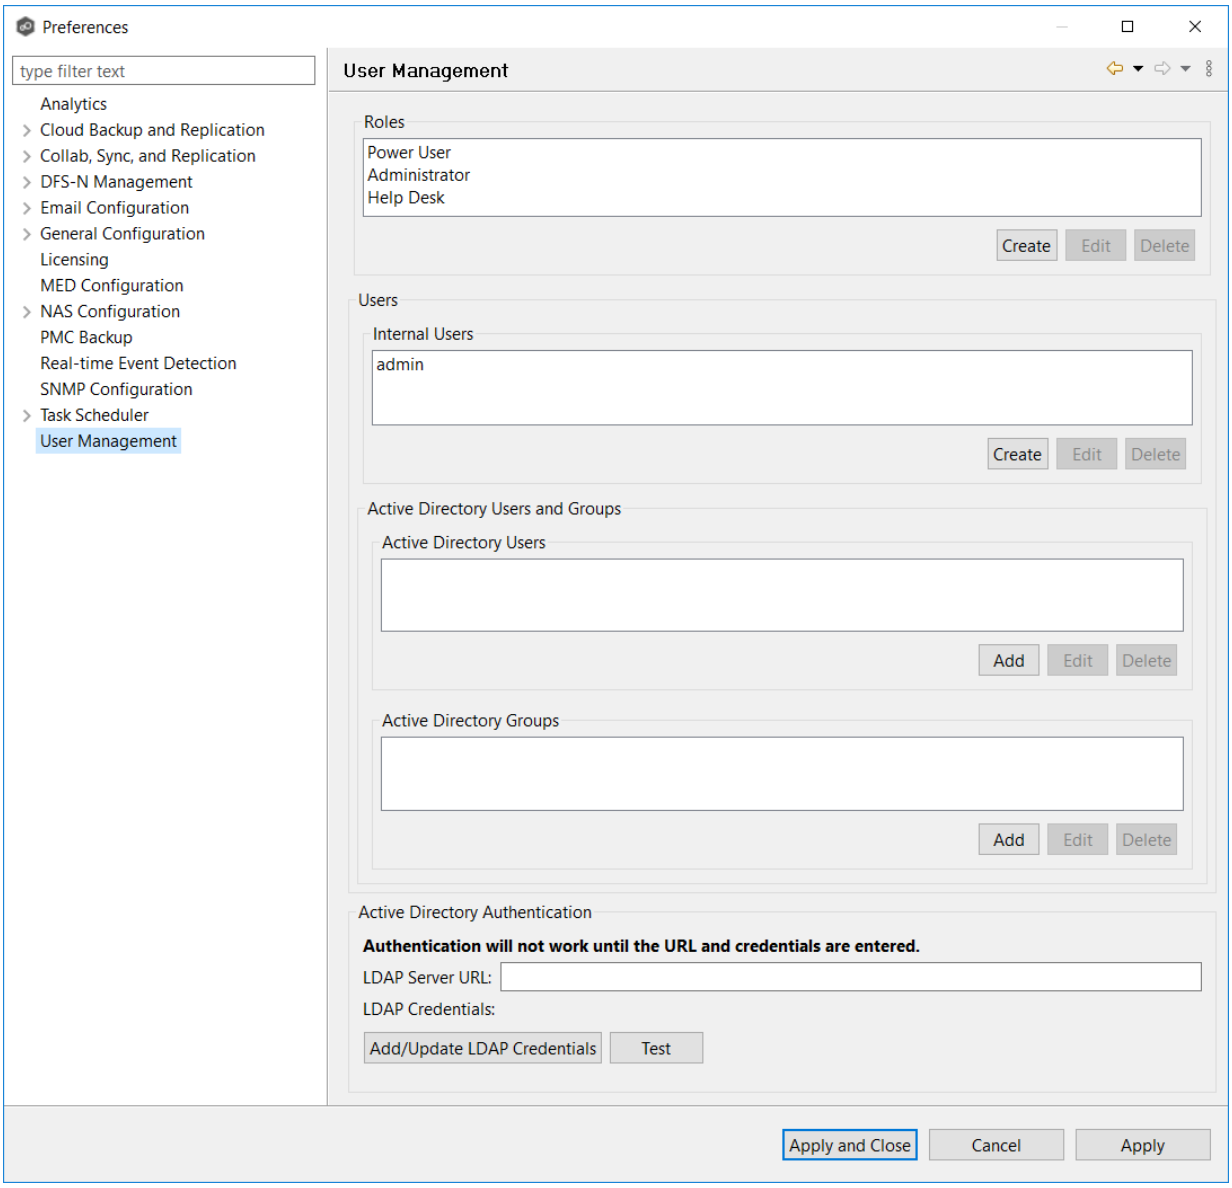

3. From this page, you can add, edit, and remove web roles, internal users, active directory users and groups, and configure Active Directory authentication.

Managing Web Client Users

Web client users are users that access Peer Management Center through the web client.

Web users can be divided into two types based on how their access to the web client is authenticated:

- Internal users - Users whose access to the web client is authenticated through the internal PMC database.
- Active Directory (AD) users and groups - Users whose access to the web client is authenticated through Active Directory.

You add, modify, and delete web users through the User Management page in Preferences. The **User Management** page is also where you specify the Active Directory account that will be used when Peer Management Center queries Active Directory for authentication.

Management of web users can be performed through the rich client or through the web client by a user with an **Administrator** role. For more information, see:

- Managing Internal Users
- Managing Active Directory Users and Groups
- Configuring Active Directory Authentication

Managing internal users involves:

- Creating internal users
- Editing internal users
- Deleting internal users

Creating an Internal User

To add an internal user, follow these steps:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.

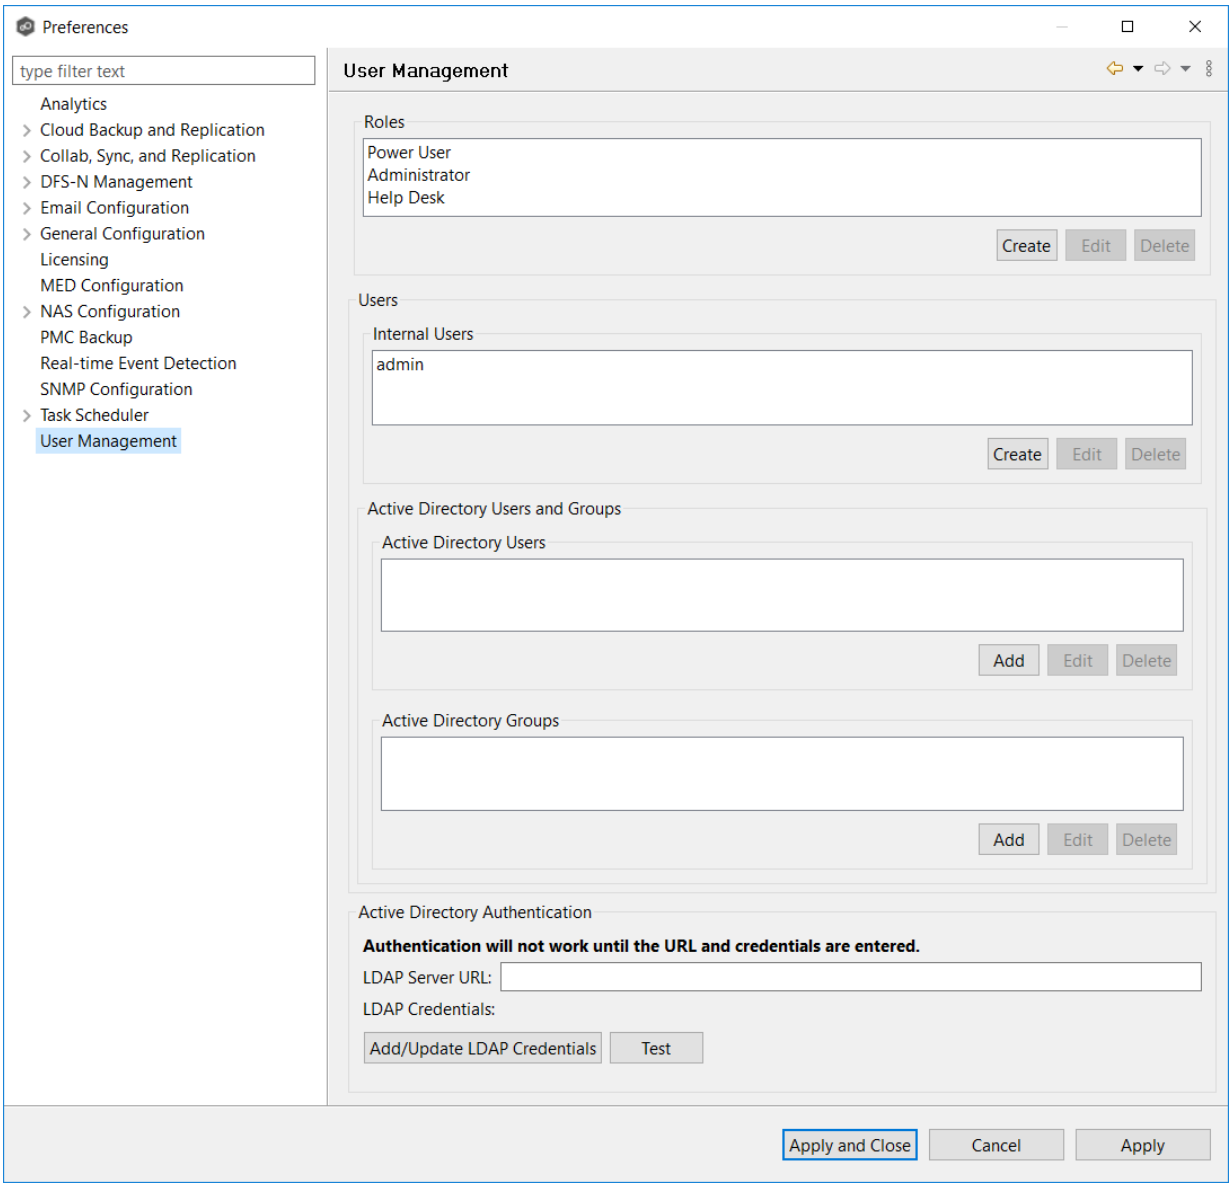

3. Select the **Create** button for Internal Users.

   The **New Internal User** dialog appears.

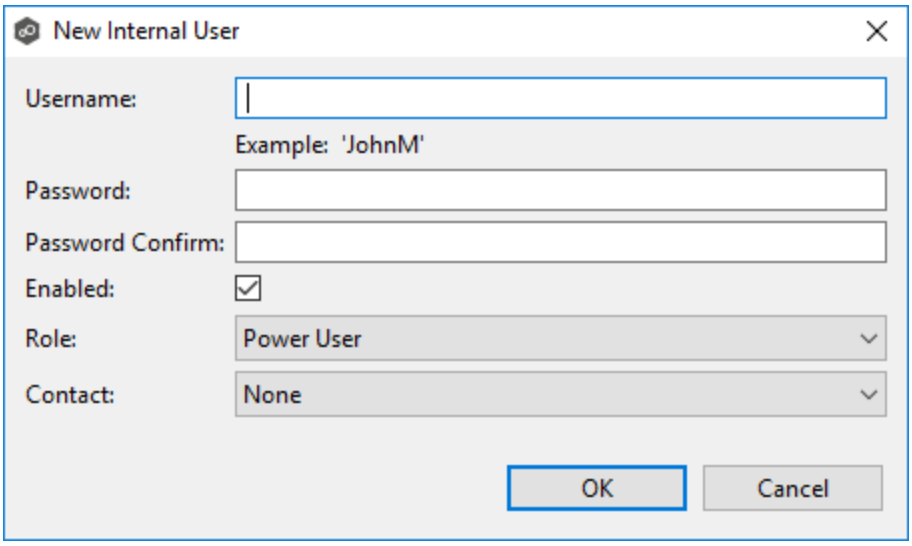

4. Enter the following information.

   - **Username**: The username can contain letters, numbers, and spaces; it cannot contain special characters. The minimum number of characters is 6; the maximum number of characters is 20.

   - **Password**: The minimum number of characters is 6; the maximum number of characters is 20. The password cannot be the same as the username.

   - **Password Confirm**: Re-enter the password you entered.

   - **Enabled**: Select this checkbox if you want to enable this user to access Peer Management Center. You can enable or disable the user at a later date by editing the user.

   - **Role:** Select the web role you want to assign to the user. It can be a standard role or a custom role. For more details on the available roles, see Web Roles

   - **Contact**: Select the user's email address from the drop-down list. If the user's email address does not appear in the list, you can add it to **Contacts** in the Email Configuration in Preferences.

5. Click **OK**.

   The new user appears in the list of internal users on the User Management page.

6. Click **Apply and Close** or **Apply**.

Editing an Internal User

Once an internal user has been created, its user name, password, email address, and web role can all be changed.

Note: The default admin user cannot be renamed, nor can its role be changed. However, you should change the default password for the default admin user.

To edit an internal user from Peer Management Center:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.
3. Select the user from the list of internal users.
4. Click **Edit**.
5. Make the changes in the **Edit User Information** dialog.
6. Click **OK**.
7. Click **Apply and Close** or **Apply**.

Deleting an Internal User

Once the account of an internal user is deleted, that user can no longer access Peer Management Center through the web client.

Note: The default admin user cannot be deleted.

To delete an Active Directory user or group from Peer Management Center:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.
3. Select the user from the list of internal users.
4. Click **Delete**.
5. Click **OK** in the **Remove User** dialog.
6. Click **Apply and Close** or **Apply**.

Managing Active Directory users involves:

- Adding an Active Directory User or Group
- Editing an Active Directory User or Group
- Deleting an Active Directory User or Group

Adding an Active Directory User or Group

To add Active Directory users and groups to Peer Management Center, follow these steps:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.

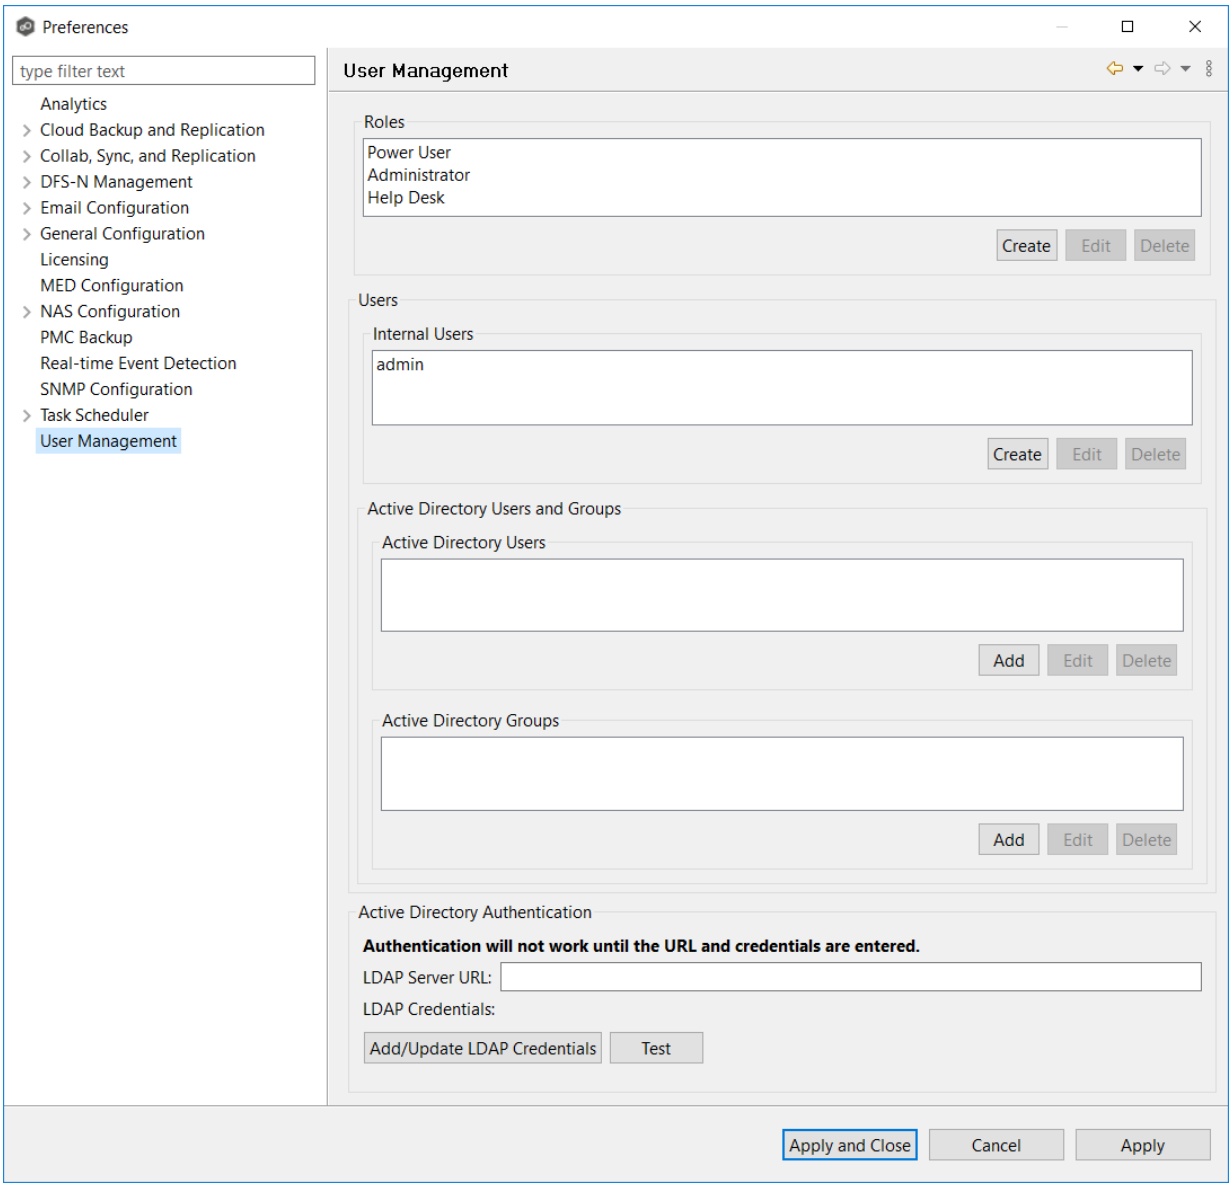

3. Add an Active Directory user or group by clicking the appropriate **Add** button.
4. Enter the information required in the dialog that appears:
   - For an individual user, enter the domain name, user name, and select a role.

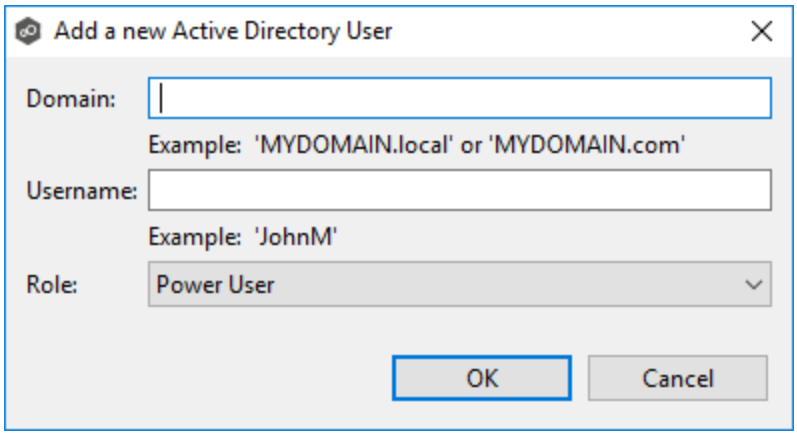

- For a user group, enter the domain name, group name, and select a role.

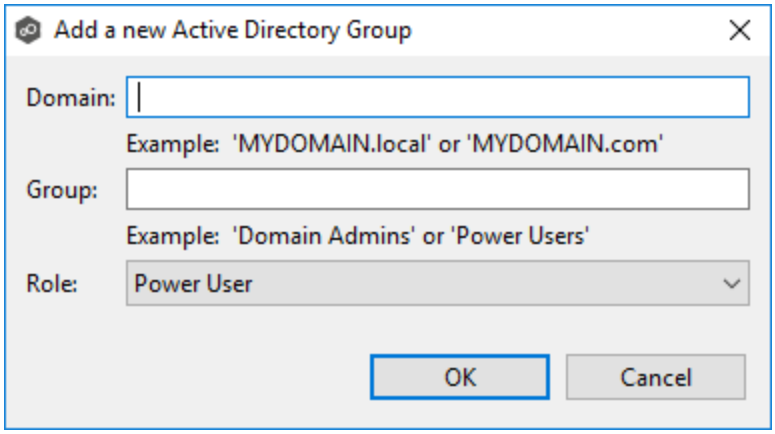

Directory users and groups are saved in the following format:
username@mydomain.local

5. Click **OK**.

   The added user or group appears in the list of Active Directory users or groups.

6. Click **OK**.

Editing an Active Directory User or Group

To edit an Active Directory user or group:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.
3. Select the AD user or group from the list of AD users or groups.

4. Click **Edit**.
5. Make the changes.
6. Click **OK**.

Deleting an Active Directory User or Group

If you delete an Active Directory user or group from Peer Management Center, that user or group will no longer have access to Peer Management Center through the web client. However, deleting the AD user or group from Peer Management Center does not delete that user or group from the Active Directory.

To delete an Active Directory user or group from Peer Management Center:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.
3. Select the AD user or group from the list of AD users or groups.
4. Click **Delete**.
5. Confirm that you want to delete the user or group.
6. Click **OK**.

To configure Active Directory authentication, you need:

- The URL of the LDAP server
- The LDAP administrator credentials

Active Directory users won't be able to access Peer Management Center until the authentication is configured.

To configure Active Directory authentication:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.

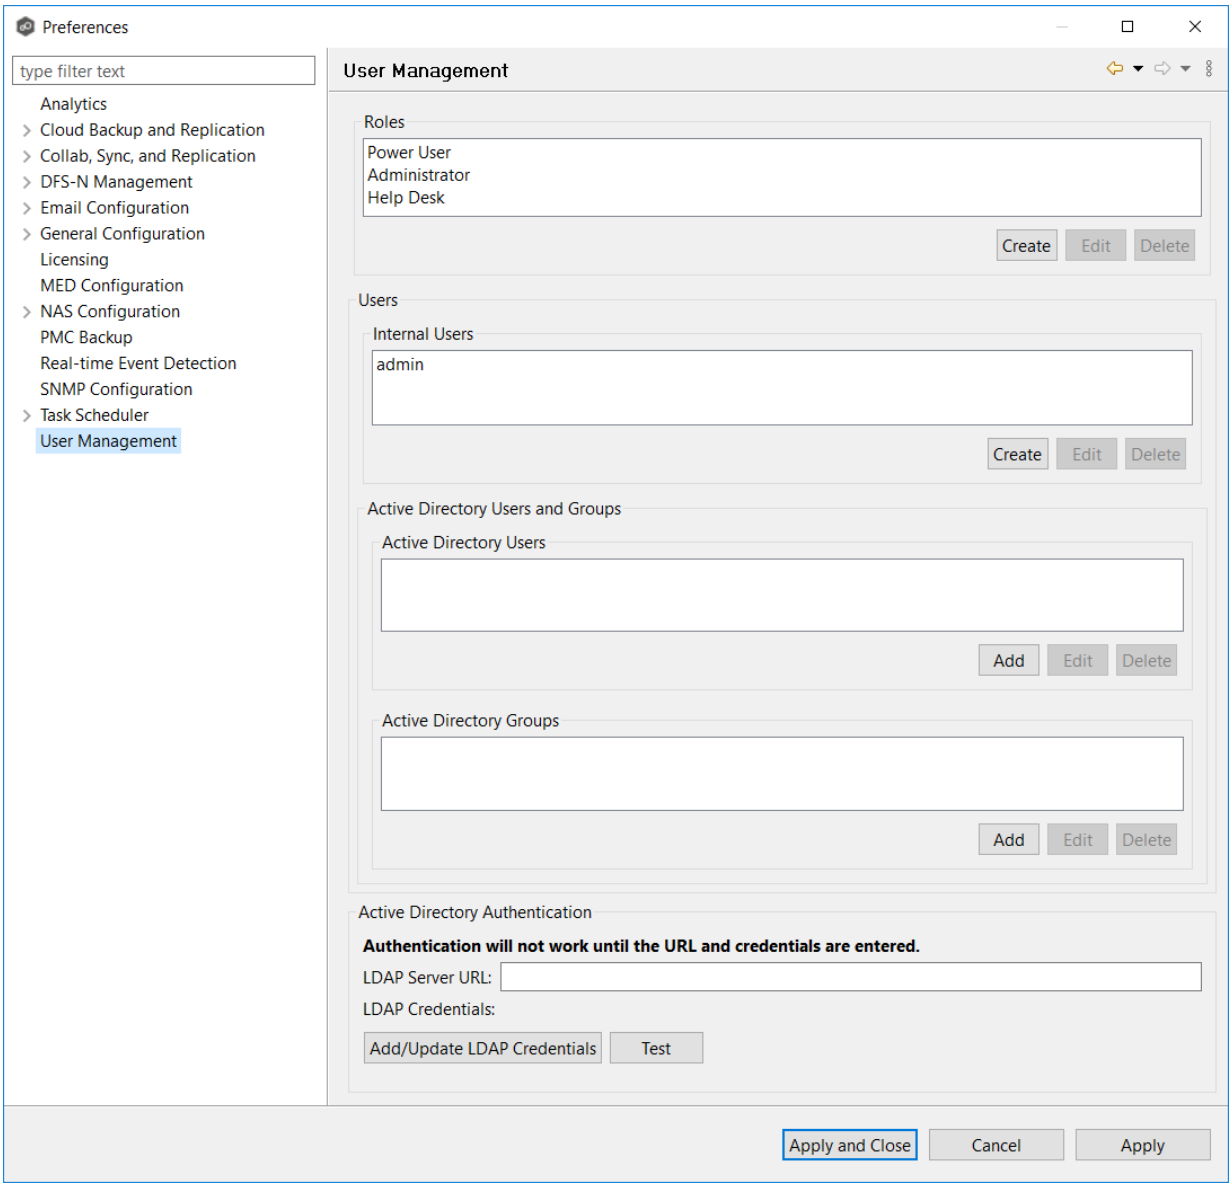

3. In the **URL** field in the **Active Directory Authentication** section, enter the URL of the LDAP server on the network using one of the following formats:

   - ldap://MYDOMAIN.LOCAL
   - ldaps://MYDOMAIN.LOCAL

4. Click **Add/Update LDAP Admin User**.

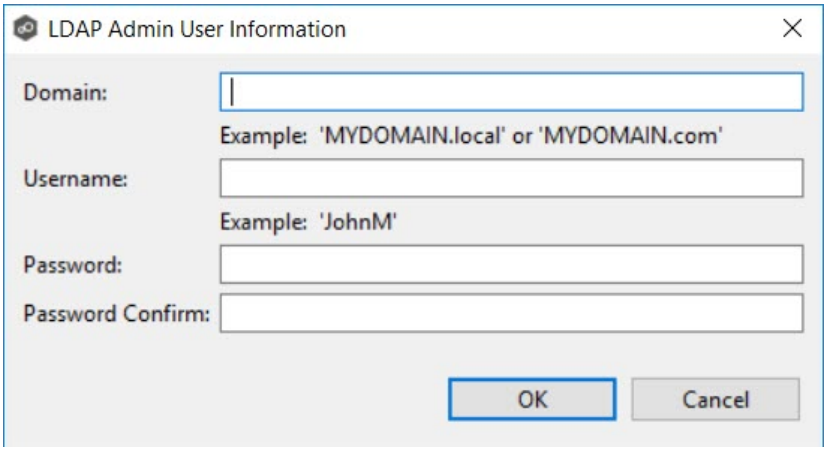

5. Enter the domain name, user name, and password.
6. Confirm the password.
7. Click **OK**.

   The LDAP user's information appears below the **Add/Update LDAP Admin User** button.
8. Click **Test** to verify the connection to the LDAP server.
9. Click **OK**.

Managing Web Roles

Managing web roles involves:

- Creating custom web roles
- Editing and deleting web roles
- Assigning tags to web roles

To create a custom role:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.

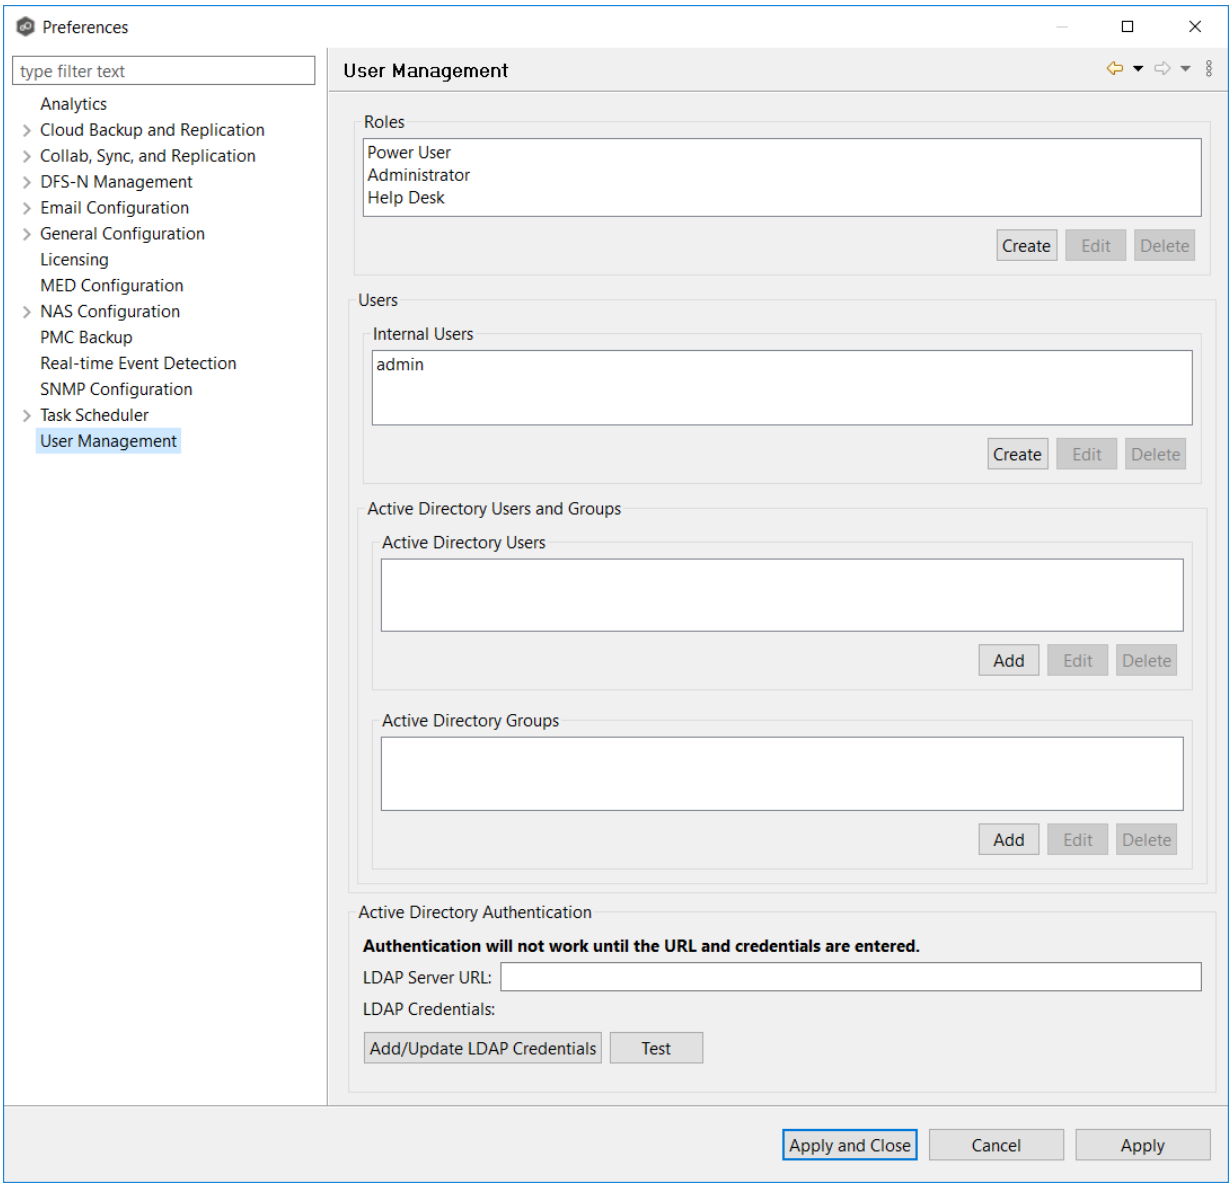

3. Click the **Create** button in the **Roles** section.

   The **General** tab of **New Role** dialog is displayed.

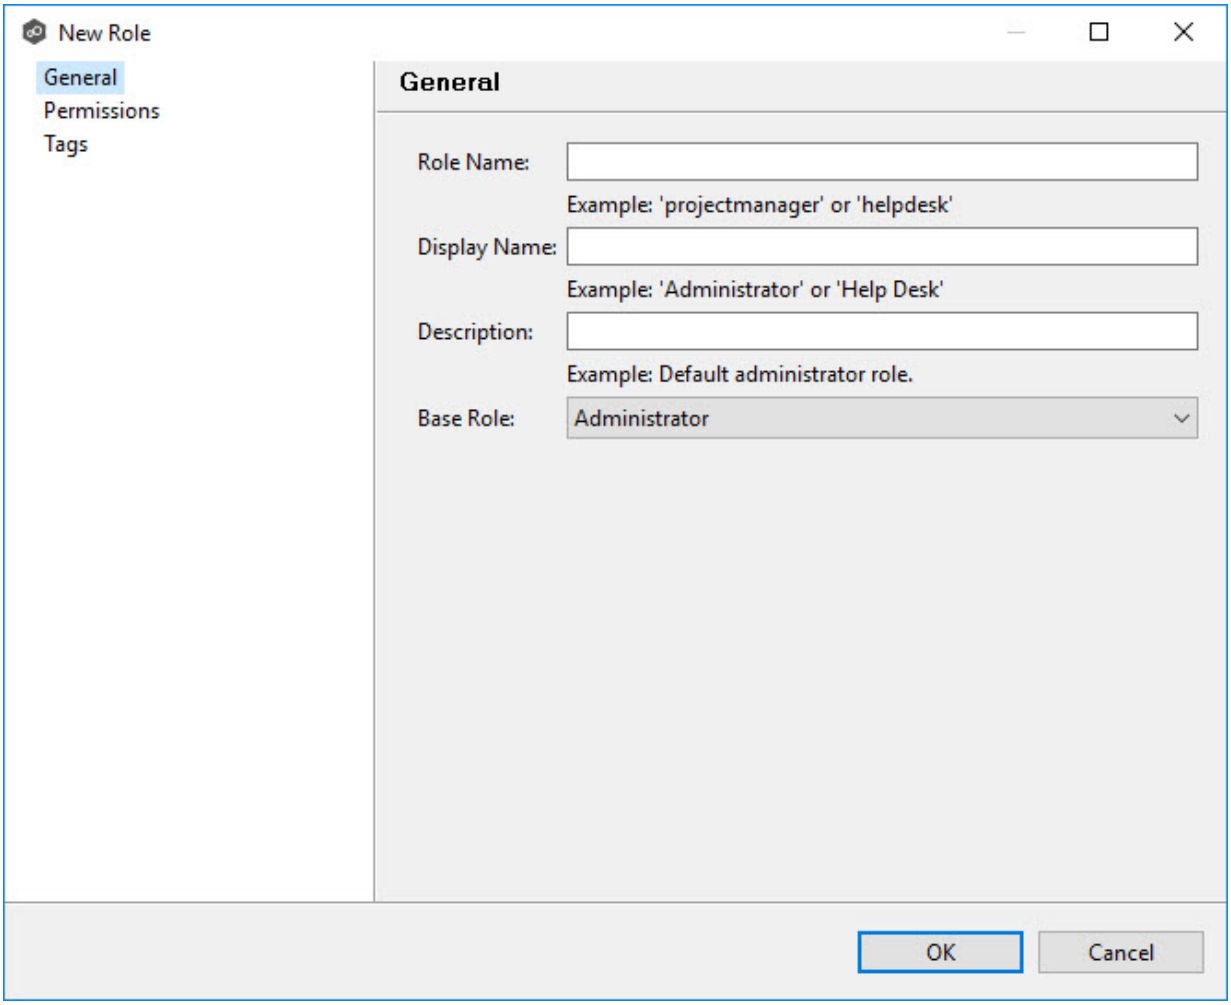

4. Enter the following information:

   - **Role Name**: A Role Name can contain only letters and numbers; it cannot contain any spaces or special characters. The maximum number of characters is 20. The Role Name is used in the internal Peer Management Center database.

   - **Display Name**: The Display Name can contain spaces and special characters, in addition to letter and numbers. The Display Name is displayed in the PMC user interface and reports.

   - **Description**: (Optional) Use the Description to provide a brief summary of the intended use of the role.

5. Select a standard web role on which to base the custom role.

6. Select the **Permissions** tab.

The **Permissions** tab displays a table of the permissions that are available to be modified for the new role. The **Access** column displays the current level of access that the role has to the resource. There are three levels of access: **Full access**, **View-only access**, and **No access**.

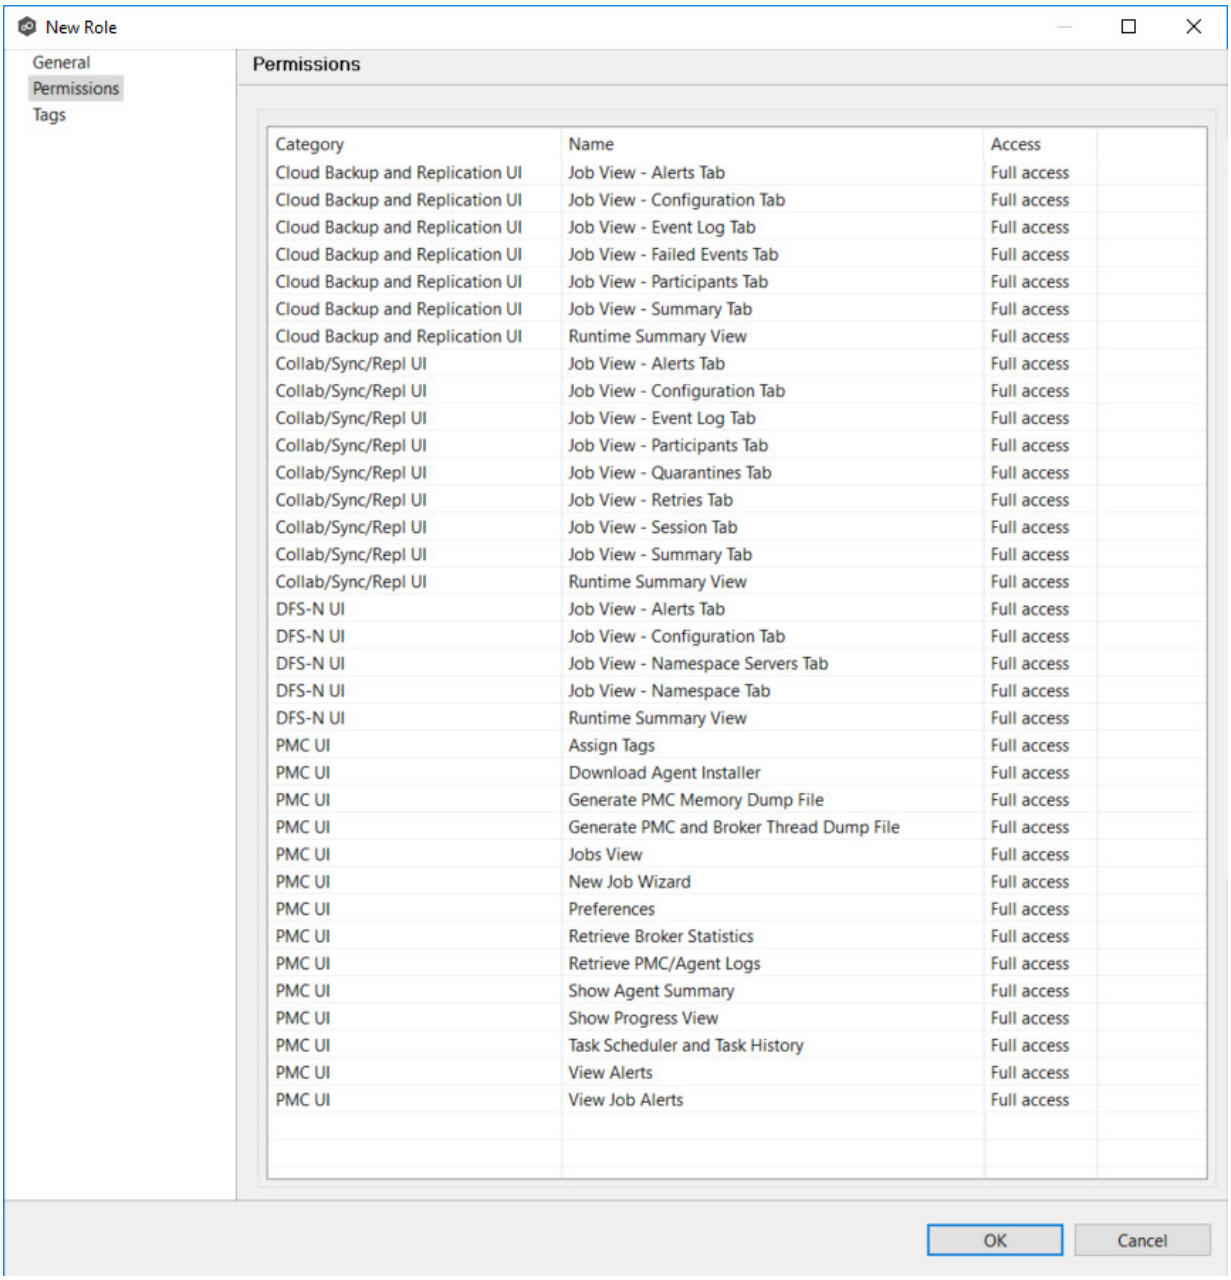

7. For each permission that you want to modify, click in the **Access** column, and then select the access level that you want for the new role.

8. (Optional) Click **Tags** to assign tags to the role.

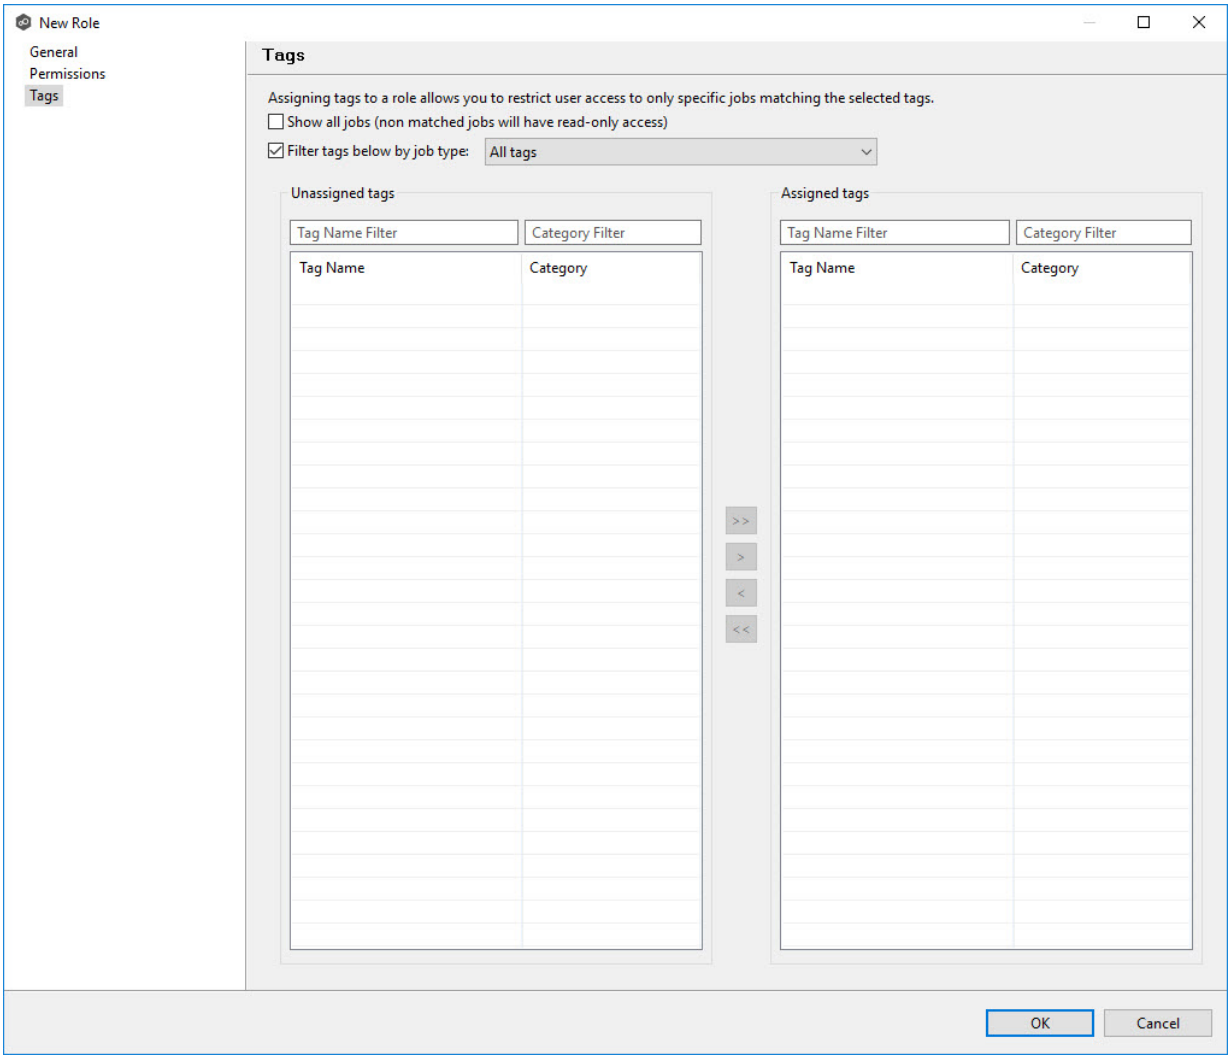

See Assigning Tags for more information.

9. Click **OK**.

   The new role appears in the **Roles** section.

Editing a Web Role

You can edit a custom web role, changing its Role Name and Display Name, its base role, associated permissions, and tags assigned to the role.

Editing of a standard role is much more restricted. It is limited to modifying the tags assigned to the role. You cannot edit its names or associated permissions.

To edit a web role, select the role in the **Roles** section in the **User Management** page, and then click **Edit**.

Deleting a Web Role

You cannot delete a standard web role.

To delete a custom web role, select the role in the **Roles** section in the **User Management** page, and then click **Delete**.

For information about assigning a tag to a web role, see Assigning Tags.

Cloud Backup and Replication Jobs

This section provides information about creating, running, and managing a Cloud Backup and Replication job:

- Overview
- Before You Create Your First Cloud Backup and Replication Job
- Creating a Cloud Backup and Replication Job
- Running a Cloud Backup and Replication Job
- Monitoring Your Cloud Backup and Replication Jobs
- Deleting a Cloud Backup and Replication Job
- Recovering Data from the Cloud

Overview

Cloud Backup and Replication brings file to object replication into Peer Software's capabilities for enterprise NAS environments. Leveraging the same real-time engine that powers Peer Software's multi-site, multi-vendor replication, Cloud Backup and Replication efficiently pushes data into Microsoft Azure or Amazon S3 storage in an open format that is immediately consumable by other applications and services.

Use cases for Cloud Backup and Replication include: (1) pushing exact replicas of on-premises data sets into object storage for use with burstable compute and cloud-borne services and (2) tape replacement-style backup to object with point-in-time recovery capability.

Before You Create Your First Cloud Backup and Replication Job

We strongly recommend that you configure the Cloud Backup and Replication settings (including proxy configurations), as well as other global settings such as SMTP configuration, email alerts, and before configuring your first Cloud Backup and Replication job. See Preferences for details on what and how to configure these settings.

In addition, we recommend that you set up your destination storage account before creating the job.

Creating a Cloud Backup and Replication Job

The **Create Job Wizard** walks you through the process of creating a Cloud Backup and Replication job. The process consists of the following steps:

Step 1: Job Type and Name

Step 2: Source Storage Platform

Step 3: Management Agent

Step 4: Proxy Configuration

Step 5: Storage Information

Step 6: Source Paths

Step 7: File Filters

Step 8: Destination

Step 1: Job Type and Name

1. Open Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **Create New Job** wizard displays a list of job types you can create.

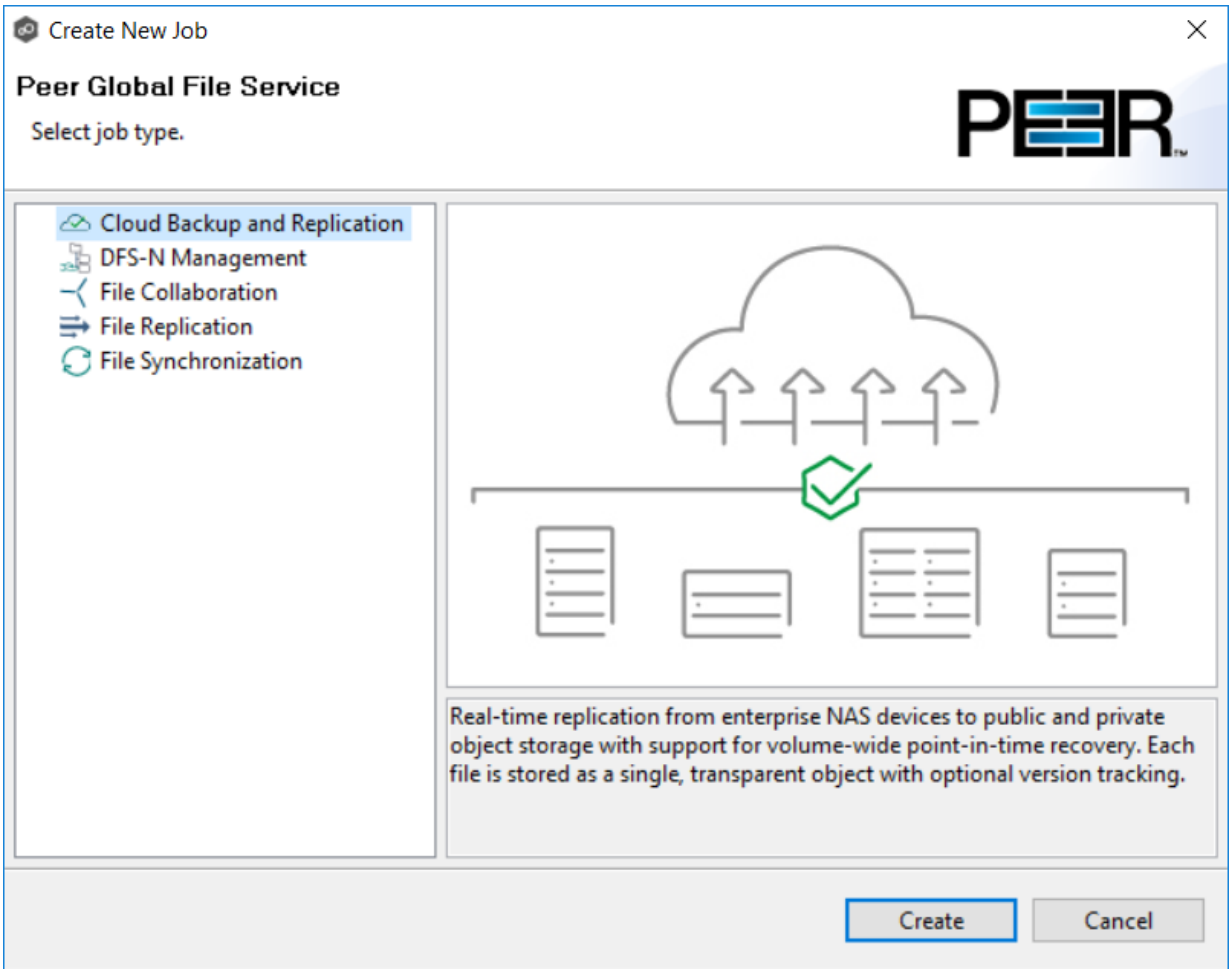

3. Click **Cloud Backup and Replication**, and then click **Create**.
4. Enter a name for the job in the dialog that appears.

   The job name must be unique.

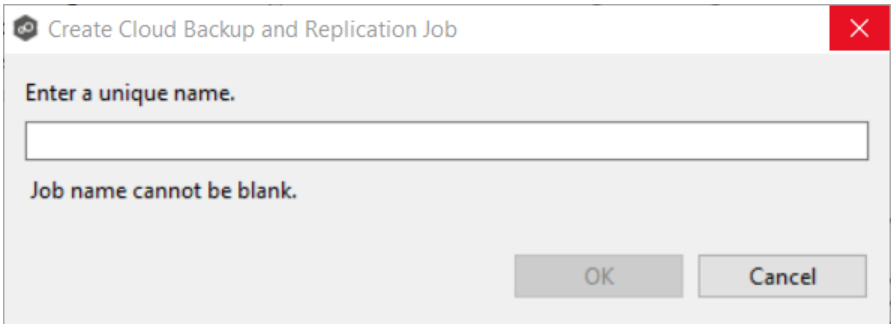

5. Click **OK**.

   The Source Storage Platform page appears.

Step 2: Source Storage Platform

The **Source Storage Platform** page lists the types of source storage platforms that Cloud Backup and Replication supports. The source storage device hosts the data you want to replicate.

1. Select the type of storage platform you want to replicate.

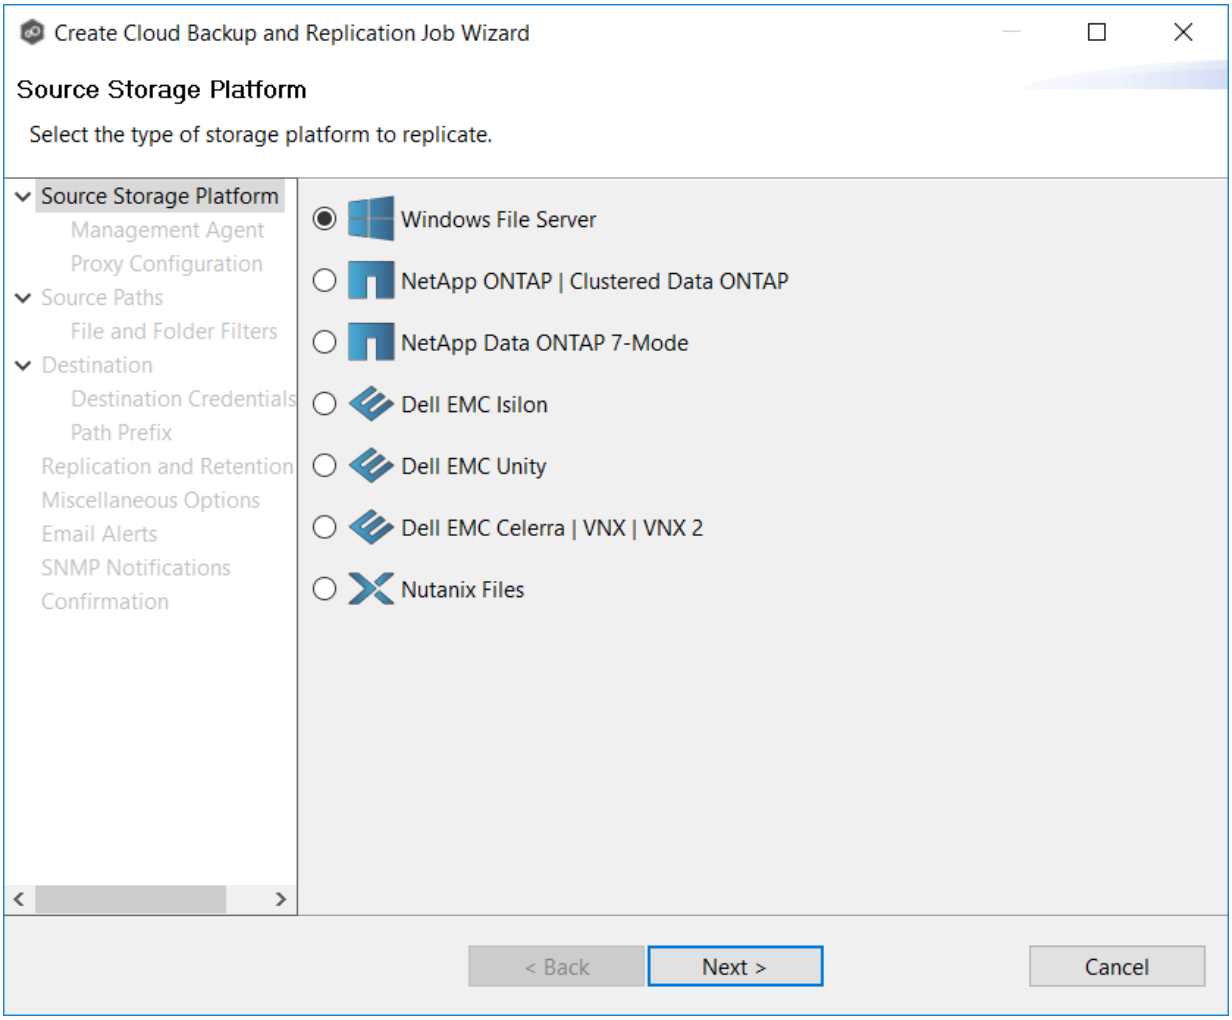

2. Click **Next**.

   The Management Agent page appears.

Step 3: Management Agent

Each storage device that you want to replicate must have a Peer Agent that manages that device. The Peer Agent that manages a device is known as its Management Agent. You can have more than one Agent managing a storage device—however, the Agents must be managing different volumes/shares/folders on that storage device.

The **Management Agent** page lists the available Agents. In this step, you should select the Agent that manages the volumes/shares/folders you want to replicate in this job.

1. Select the Management Agent for the volume/share/folder you want replicated.

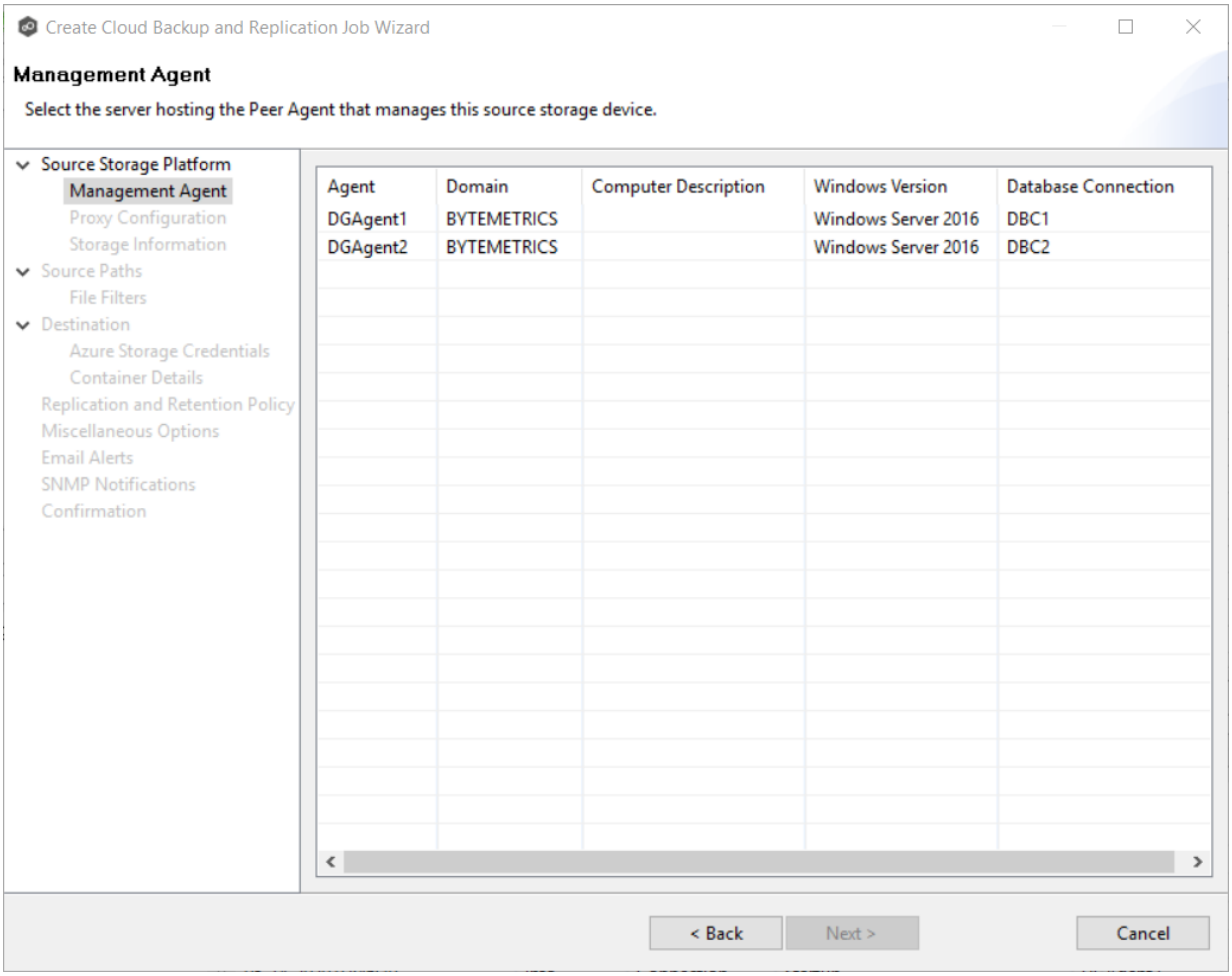

Note: If you select an Agent that does not have a database connection listed in the **Database Connection** column, a message prompts you to create the connection:

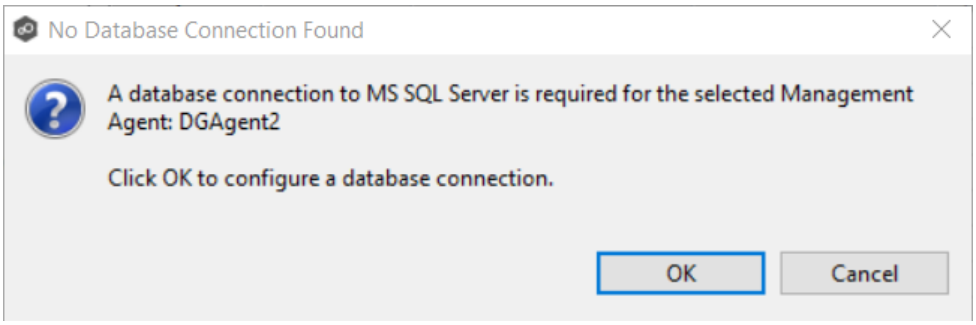

Click **OK** and then configure the database connection for the selected Management Agent. See Database Connections for instructions about creating a database connection.

2. Click **Next**.

   The Proxy Configuration page appears.

Step 4: Proxy Configuration

If you do not need a proxy server to connect to outside networks, skip this step and proceed to Step 5.

If you do need a proxy server to connect to outside networks, you have three options:

- Create a new proxy configuration.
- Use the existing proxy configuration. If there is an existing proxy configuration, details about the configuration will be displayed on the page.

  You may have created one in advance through Cloud Back and Replication Preferences or when you created another Cloud Backup and Replication job. Once a proxy configuration is created for a source storage platform, that proxy configuration is used for all Cloud Backup and Replication jobs using that agent.
- Edit an existing proxy configuration. Click the **Edit Proxy Configuration** link to edit the existing proxy.

  If you edit the proxy configuration, it affects other jobs using the same agent. Editing an existing proxy configuration has the potential to create problems with the other jobs.

If there is not an existing proxy configuration for the selected management agent, follow these steps to create a new proxy configuration:

1. Click **Create Proxy Configuration**.

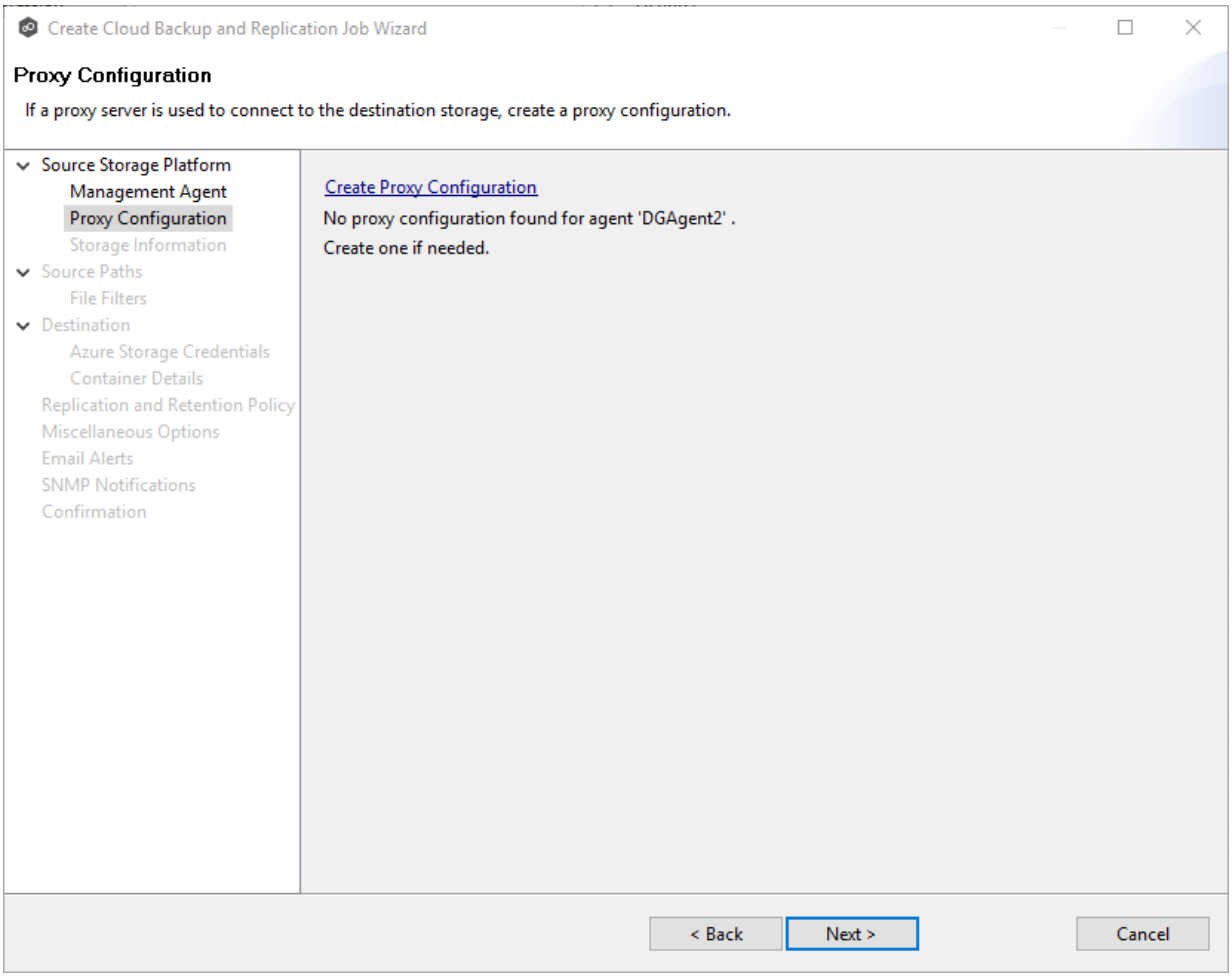

The **Proxy Configuration** page is displayed. Existing proxies are listed in the Proxy Configuration table.

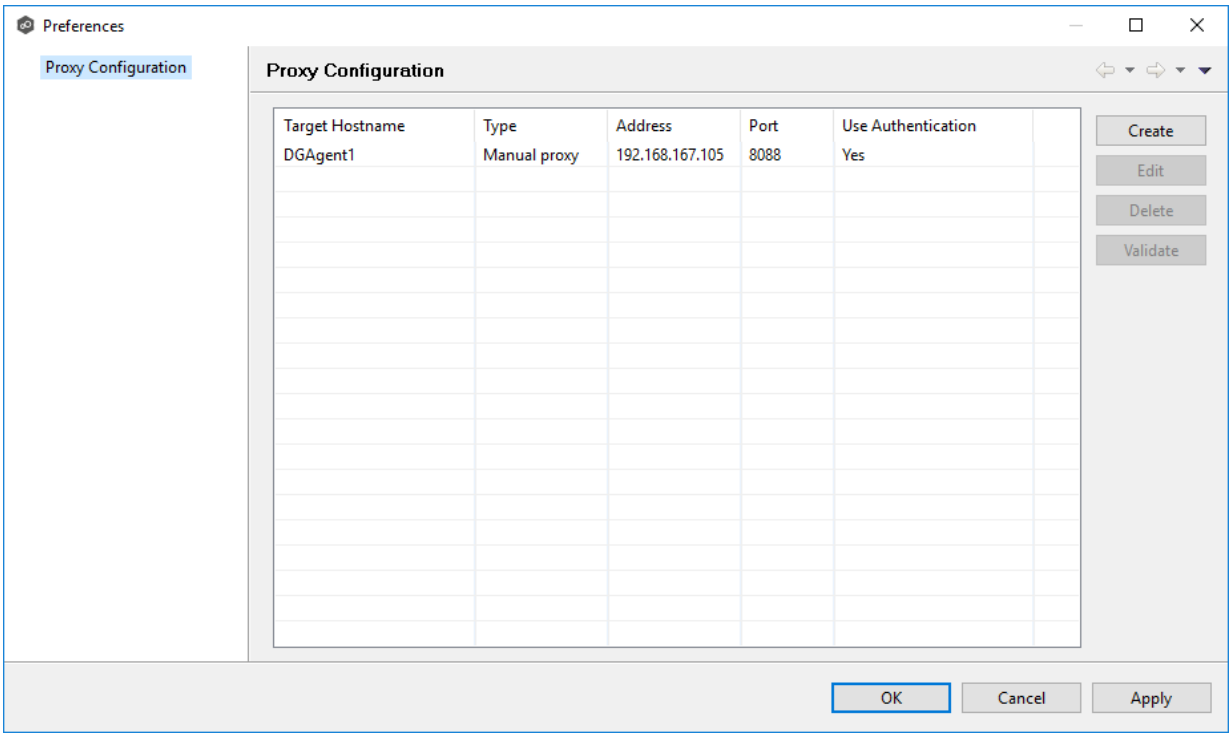

2. Click the **Create** button.

   The **Create Proxy Configuration** dialog is displayed. The Agent you selected in **Step 3: Management Agent** is preselected.

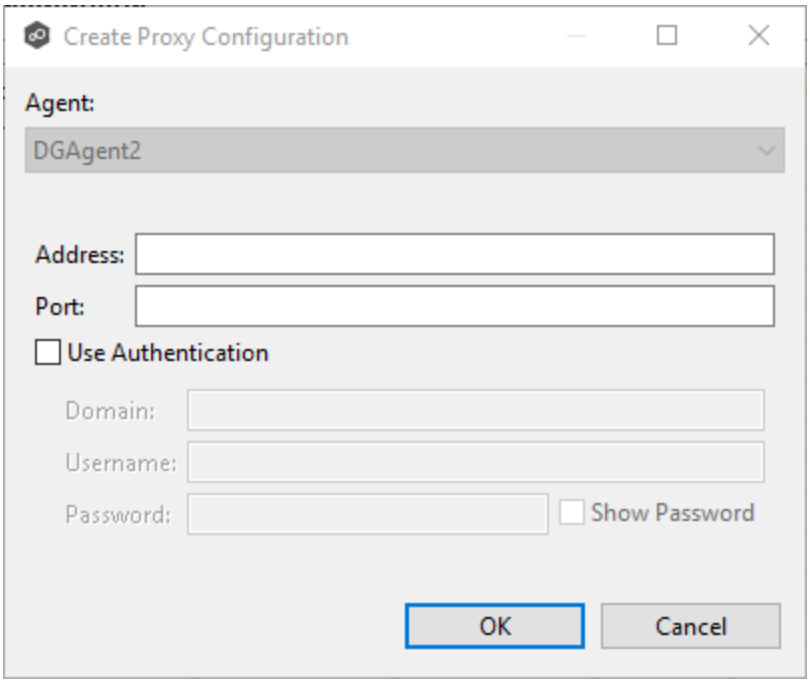

3. Enter values for the following fields:

Address	Enter the IP address or fully qualified domain name of the proxy server.
Port	Enter the port number.
User Authentication	Select this checkbox if the proxy server requires authentication.

4. If your proxy server requires authentication, click the **User Authentication** checkbox, and then supply the necessary values.

Domain	Enter the domain name on the proxy server.
Username	Enter the user name for the proxy server.
Password	Enter the password for the proxy server.

5. Click **OK**.

   If you already have jobs managed by this agent, a message appears and identifies those jobs. They will now use the proxy as well.

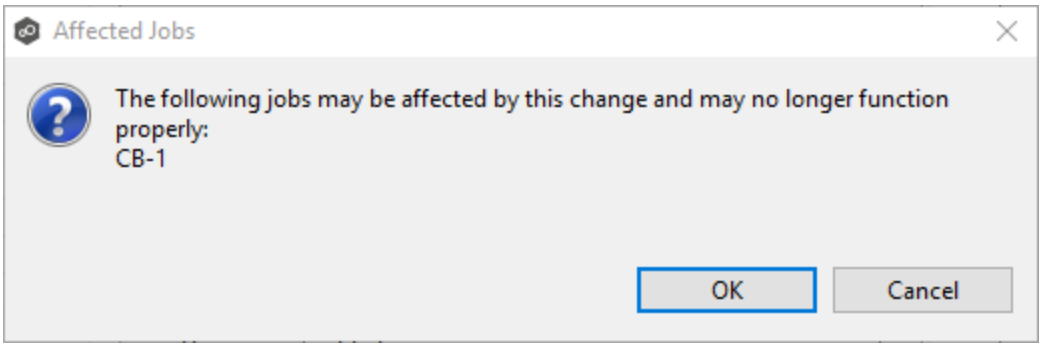

   After you click OK, the **Proxy Configuration** page is redisplayed. The proxy you just created now appears in the table.

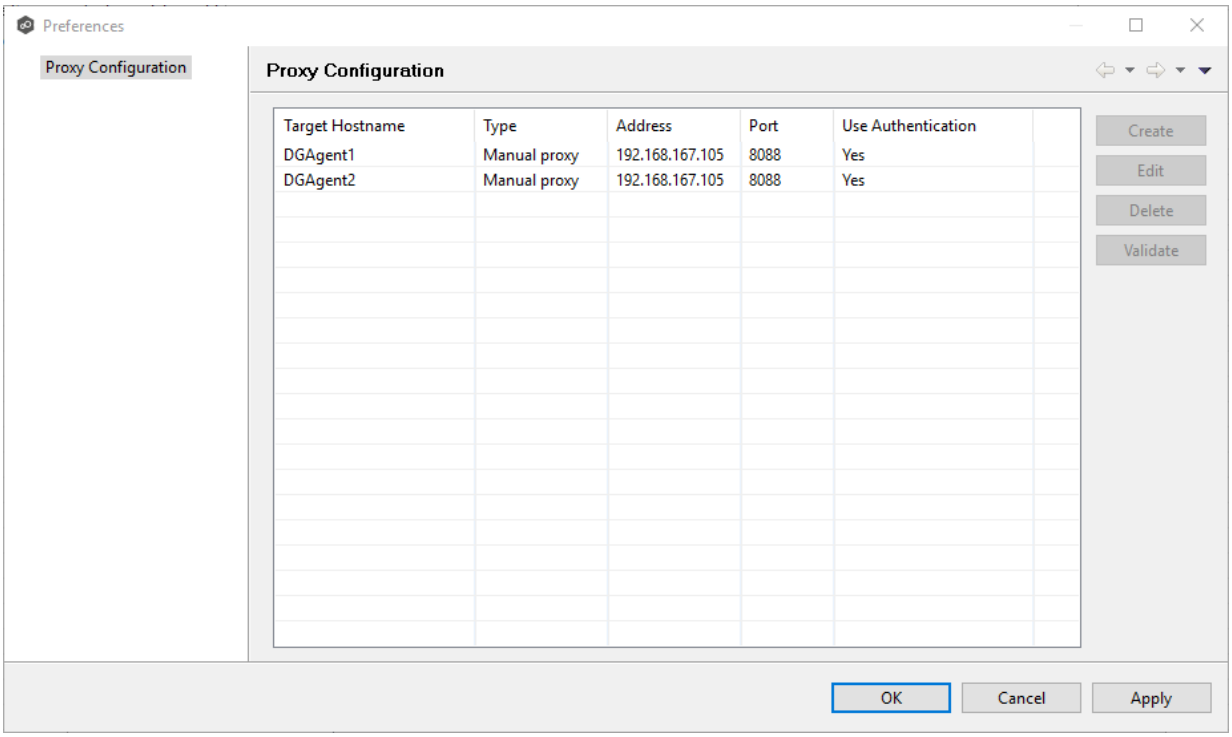

6. (Optional) Select the proxy you created and click **Validate**.

   The **Validate Proxy Configuration** dialog appears.

7. Select the target storage account, and then click **OK**.

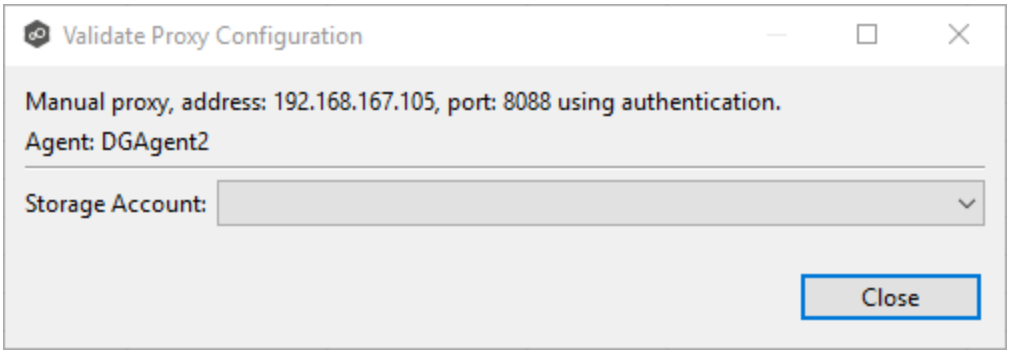

   Once you click OK, PeerGFS tests the connection to the target storage account using the proxy.

8. Click **OK** in the **Validation Result** dialog.

9. Click **Close** in the **Validate Proxy Configuration** dialog.

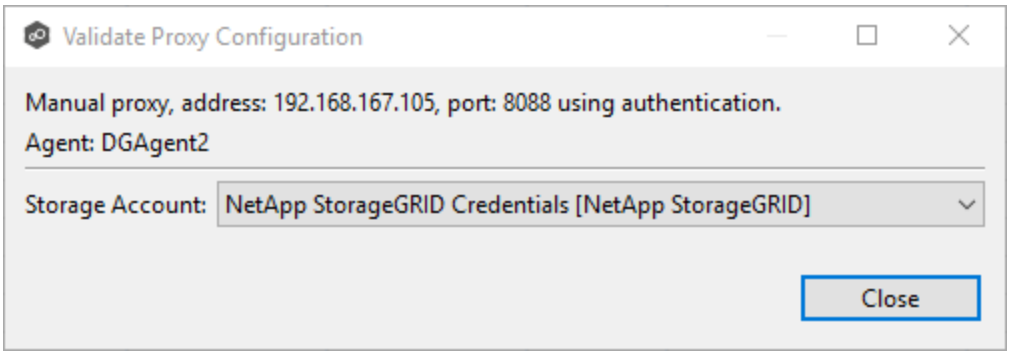

10. Click **OK** in the Proxy Configuration page.

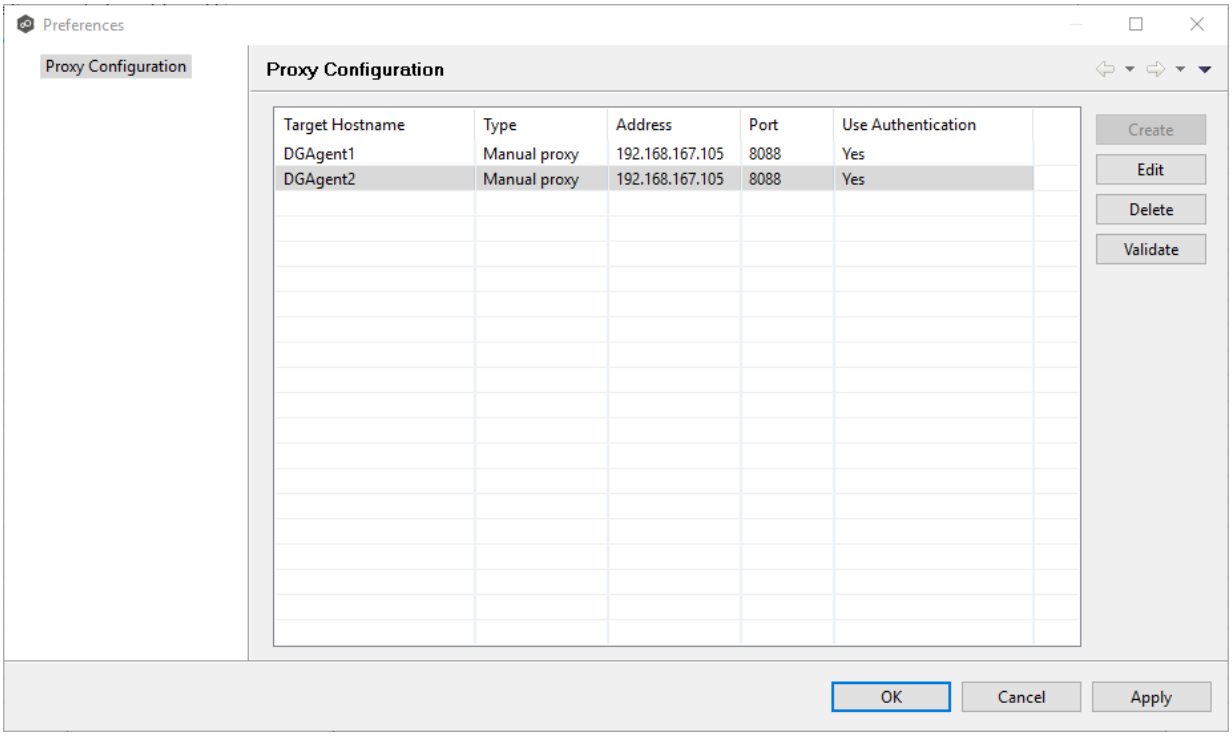

The **Proxy Configuration** page now displays the details about the proxy configuration.

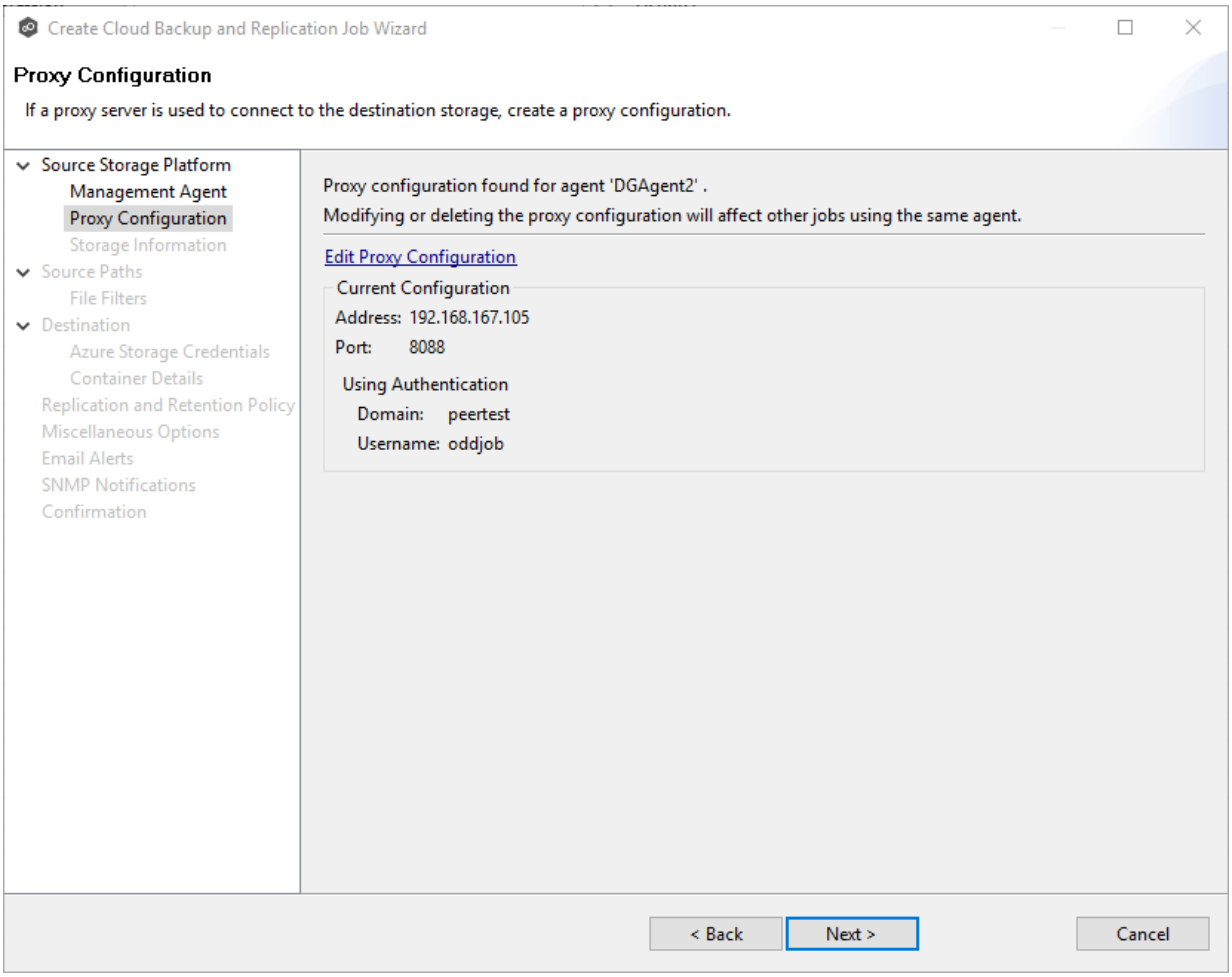

11. Click **Next**.

    The Storage Information page appears.

Step 5: Storage Information

The **Storage Information** page requests the credentials necessary to connect to the storage device you want to replicate.

Note: If you selected **Windows File Server** in Step 2, this page doesn't appear; skip to Step 5: Source Paths.

1. Select **New Credentials** to enter a new set of credentials for the source storage platform or select **Existing Credentials**.

2. If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue with Step 5: Source Paths.

   If you selected **New Credentials**, enter the credentials for connecting to the source storage device. The information you are prompted to enter varies, depending on the type of storage platform:

   NetApp ONTAP | Clustered Data ONTAP

   NetApp Data ONTAP 7-Mode

   Dell EMC Isilon

   Dell EMC Unity

   Dell EMC Celerra | VNX |VNX 2

   Nutanix Files

3. Click **Validate** to test the credentials, and then click **OK** in the confirmation message that appears.

4. Click **Next**.

   The Source Paths page appears.

1. Enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated or select existing credentials.

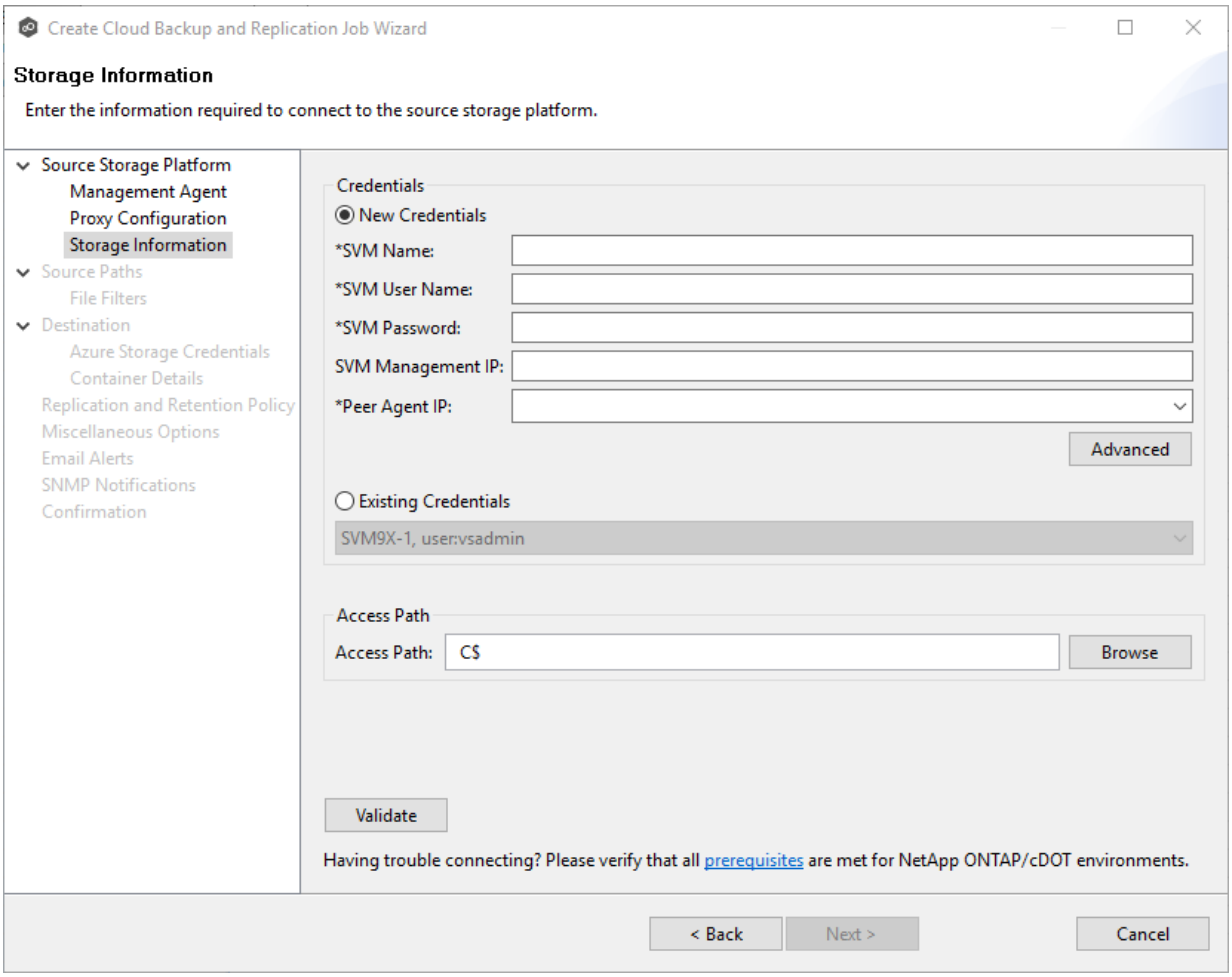

SVM Name	Enter the name of the Storage Virtual Machine hosting the data to be replicated.
SVM Username	Enter the user name for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	Enter the IP address used to access the management API of the NetApp Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already allow management access, this field is not required.

Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address.
Access Path	Use only when experiencing access issues. Contact Peer Software support for more information.

2. Click **Validate**.
3. Click **Next**.

   The Source Paths page is displayed.

1. Enter the credentials to connect to the NetApp 7-Mode filer or vFiler hosting the data to be replicated or select existing credentials.

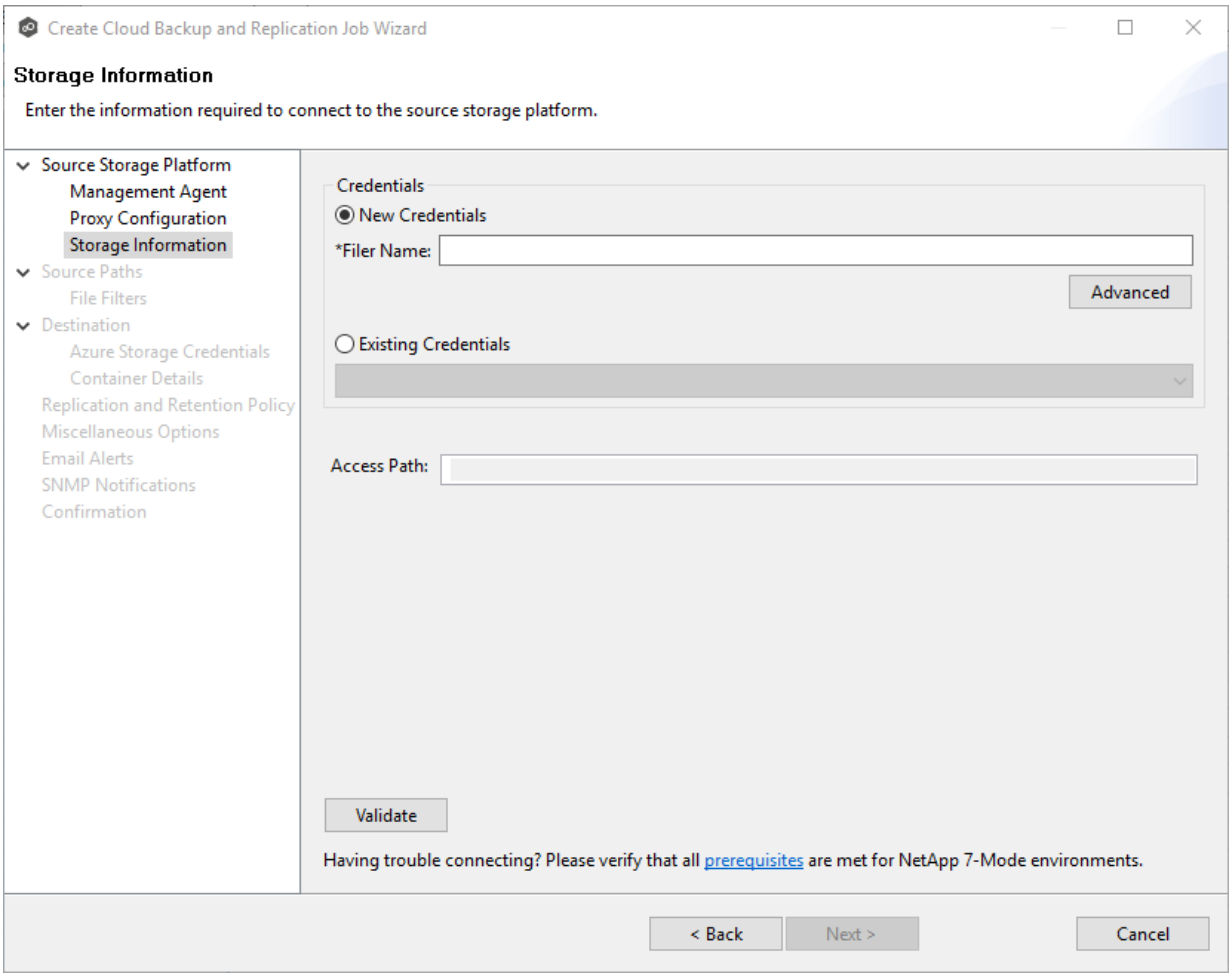

Filer Name	Enter the name of the NetApp 7-Mode filer or vFiler hosting the data to be replicated.
Access Path	Use only when experiencing access issues. Contact Peer Software support for more information.

2. Click **Validate**.
3. Click **Next**.

   The Source Paths page is displayed.

1. Enter the credentials to connect to the EMC Isilon cluster hosting the data to be replicated or select existing credentials.

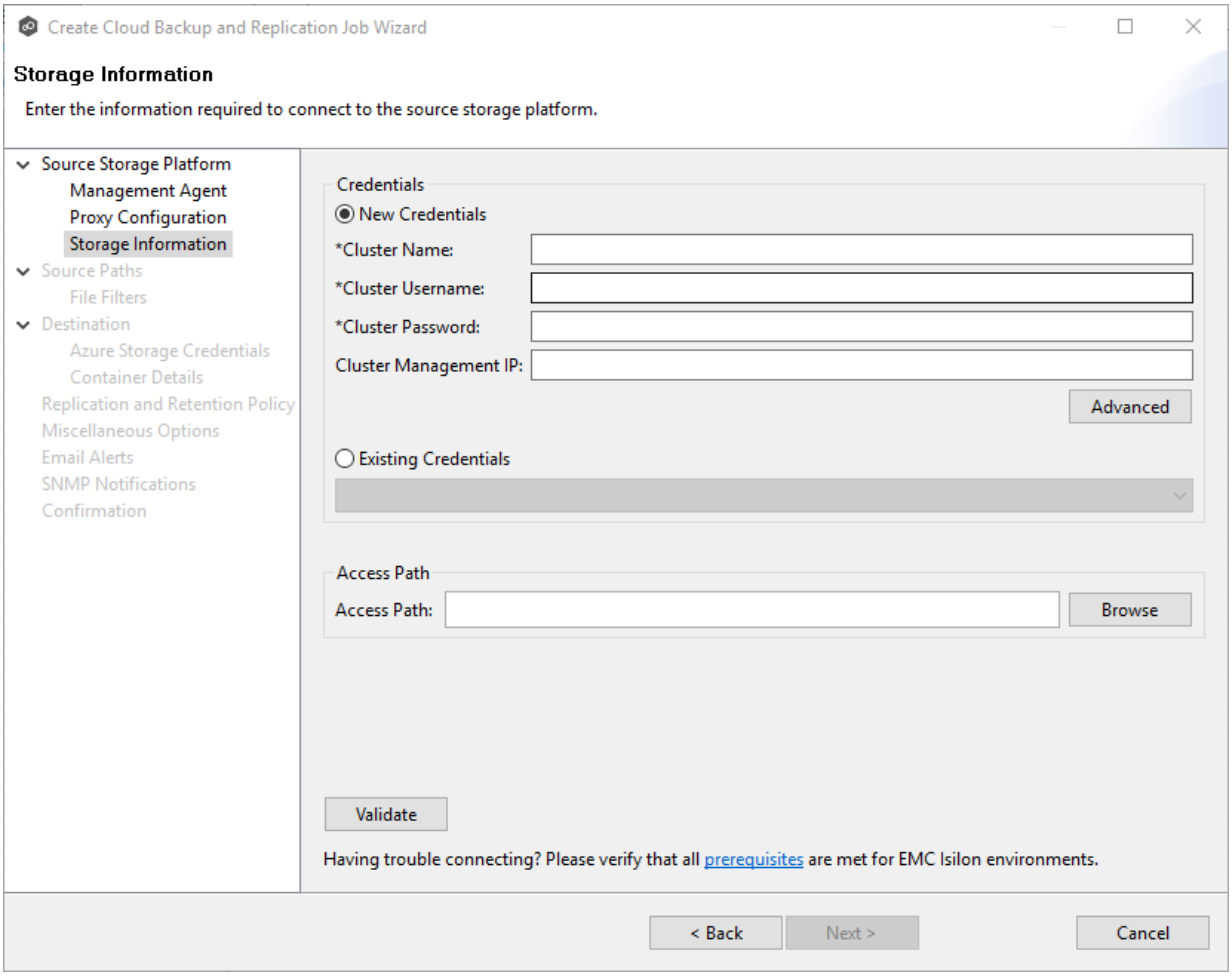

Cluster Name	Enter the name of the EMC Isilon cluster hosting the data to be replicated.
Cluster Username	Enter the user name for the account managing the EMC Isilon cluster.

Cluster Password	Enter the password for account managing the EMC Isilon cluster.
Cluster Management IP	Enter the IP address of the system used to manage the EMC Isilon cluster. Required only if multiple Access Zones are in use on the cluster.
Access Path	Use only when experiencing access issues. Contact Peer Software support for more information.

2. Click **Validate**.
3. Click **Next**.

   The Source Paths page is displayed.

1. Enter the credentials to connect to the CIFS Server hosting the data to be replicated or select existing credentials.

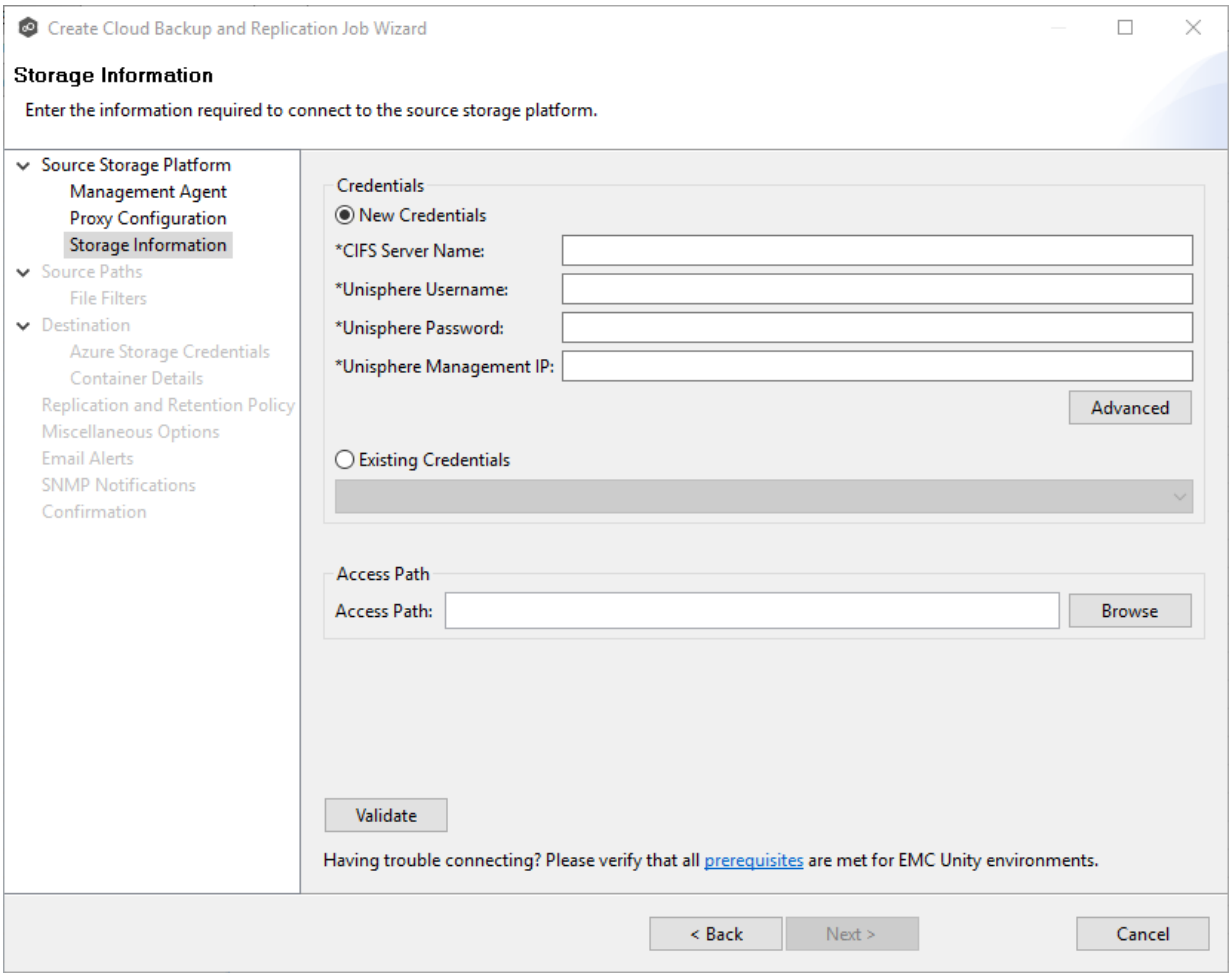

CIFS Server Name	Enter the name of the CIFS server hosting the data to be replicated.
Unisphere Username	Enter the user name for the Unisphere account managing the Unity storage device.
Unisphere Password	Enter the password for the Unisphere account managing the Unity storage device.
Unisphere	Enter the IP address of the Unisphere system used to manage the Unity storage device. This should not point to the CIFS server.

Management IP	
Access Path	Use only when experiencing access issues. Contact Peer Software support for more information.

2. Click **Validate**.
3. Click **Next**.

   The Source Paths page is displayed.

1. Enter the credentials to connect to the CIFS Server hosting the data to be replicated or select existing credentials.

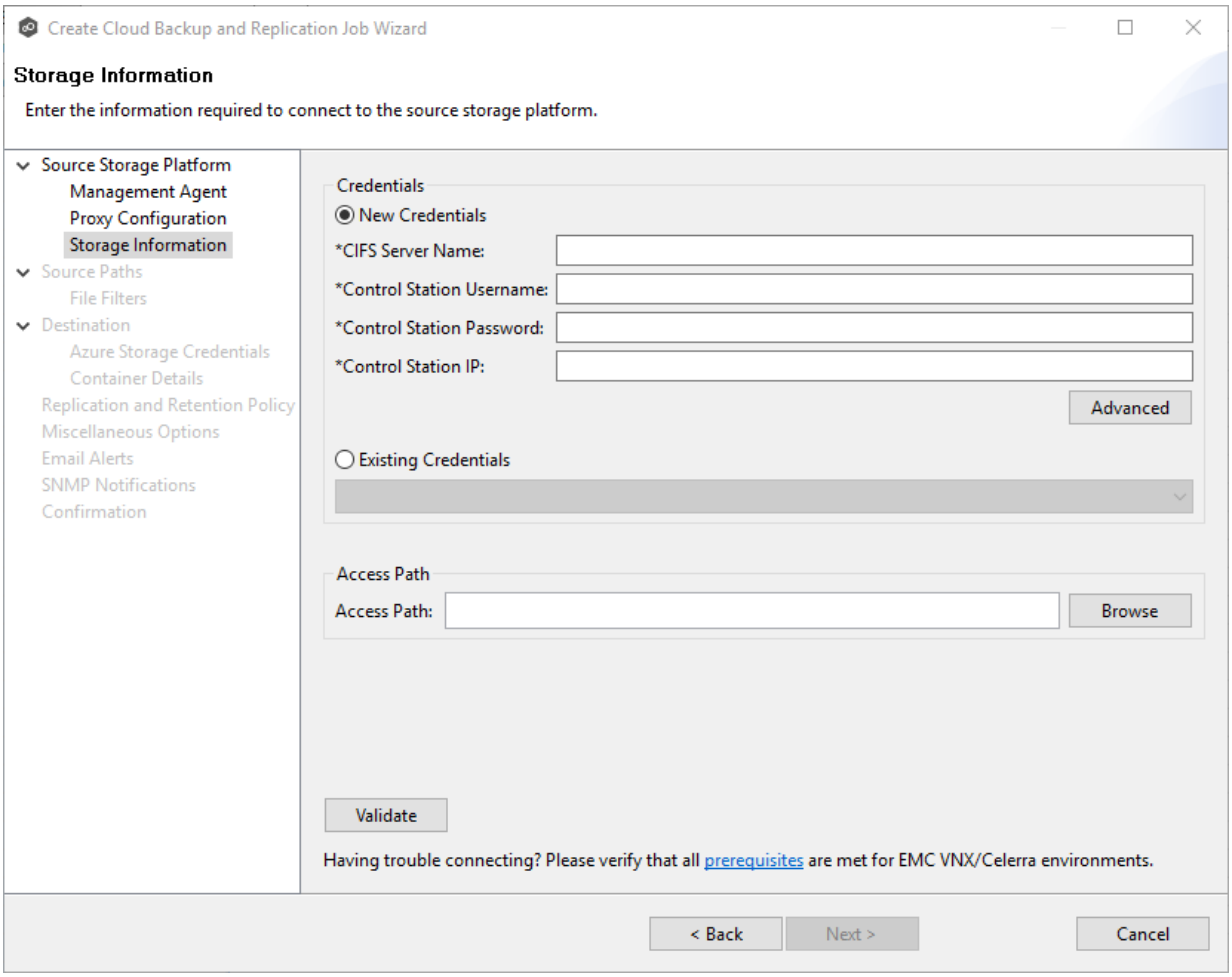

CIFS Server Name	Enter the name of the CIFS Server hosting the data to be replicated.
Control Station Username	Enter the user name for the Control Station account managing the Celerra/VNX storage device.
Control Station Password	Enter the password for the Control Station account managing the Celerra/VNX storage device.
Control Station	Enter the IP address of the Control Station system used to manage the Celerra/VNX storage device. This should not point to the CIFS Server.

IP	
Access Path	Use only when experiencing access issues. Contact Peer Software support for more information.

2. Click **Validate**.
3. Click **Next**.

   The Source Paths page is displayed.

1. Enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated or select existing credentials.

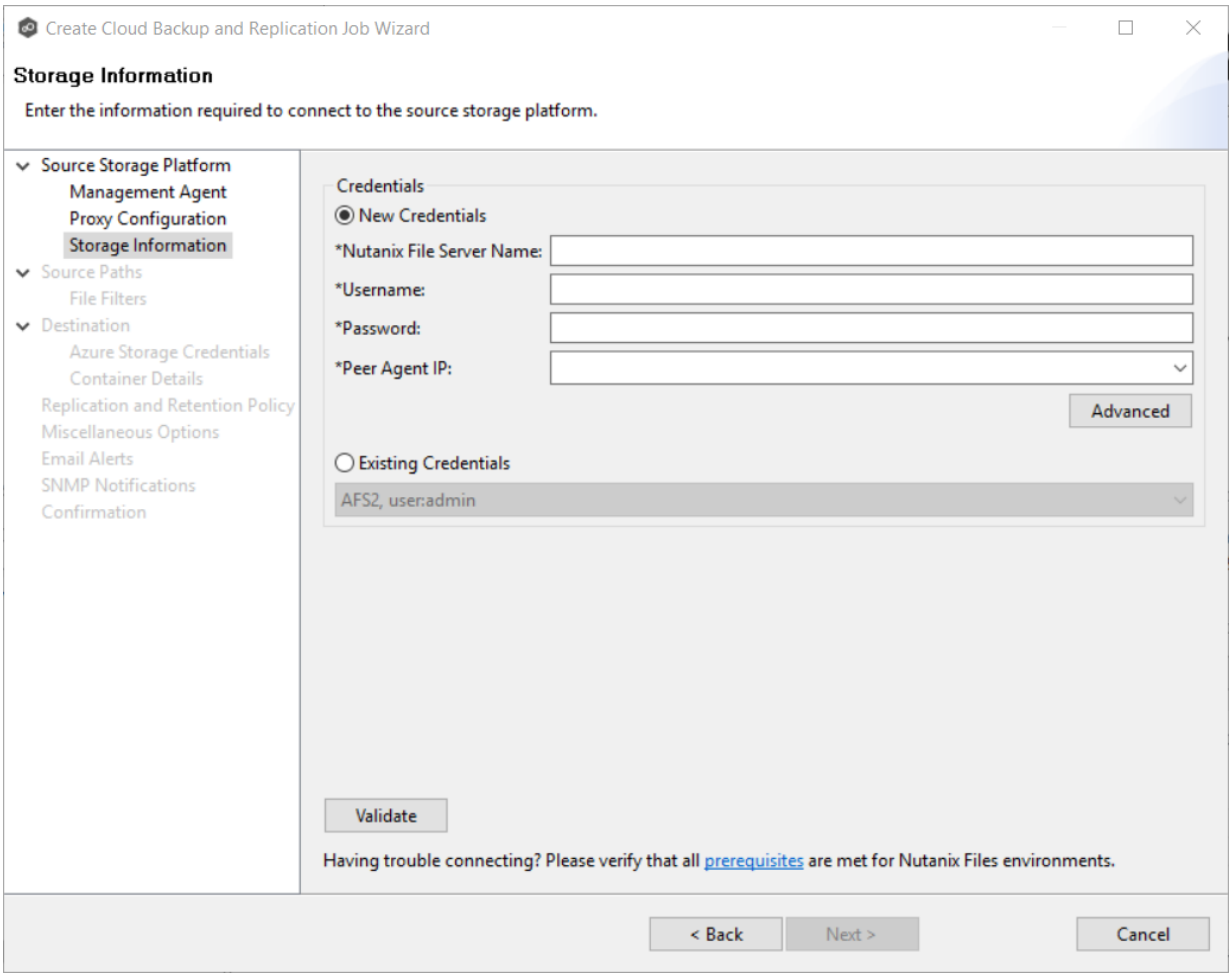

Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files server must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address. This should not point to the Files cluster itself.

2. Click **Validate**.

3. Click **Next**.

   The Source Paths page is displayed.

Step 6: Source Paths

The **Source Paths** page displays a list of available volumes to replicate. You can choose to replicate an entire volume or selectively replicate files and folders. The files/folders/volumes selected for replication are referred to as the watch set.

1. Select the paths to the files/folders/volumes you want to replicate.

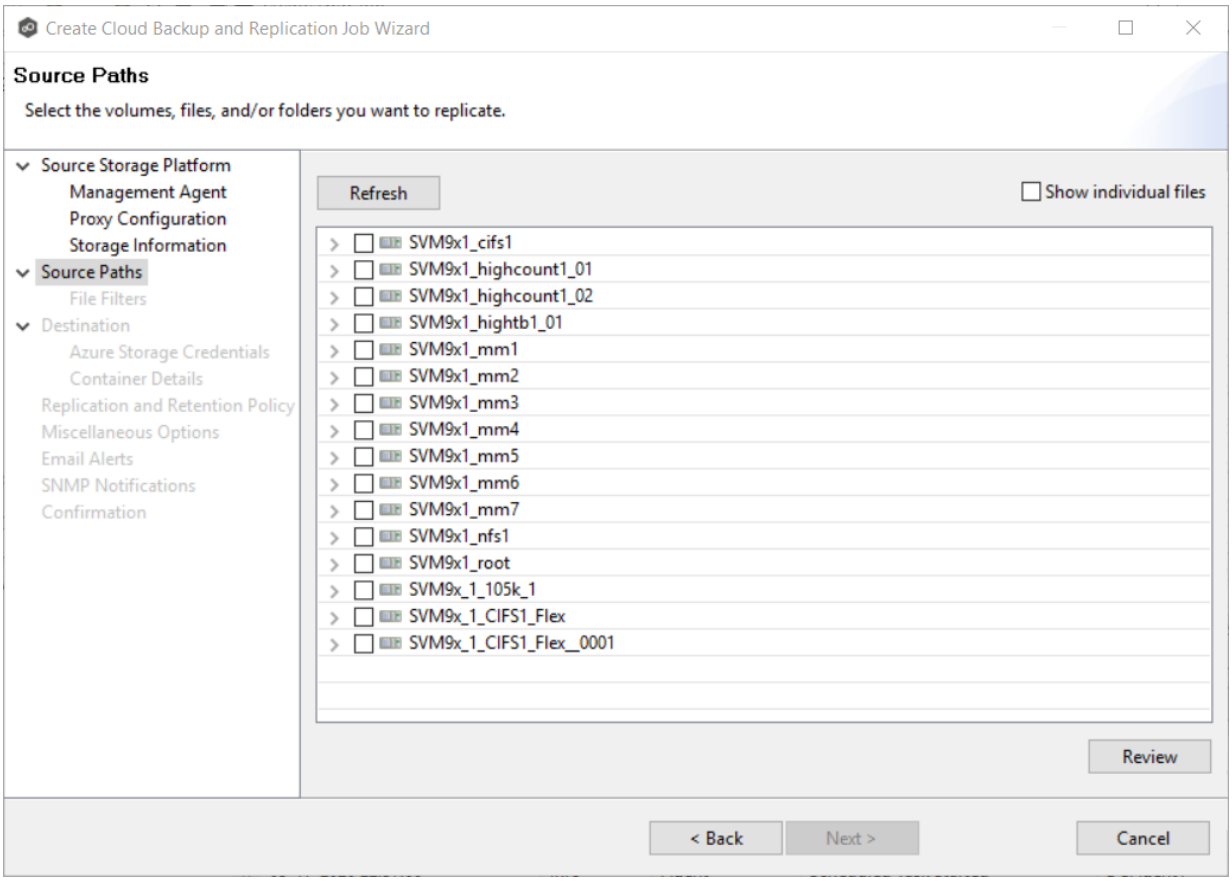

   To replicate:

The entire volume (all files and folders, including subfolders and their files)	Select the volume checkbox.
All files at the root level of the volume (but no folders)	Expand the volume, scroll to the bottom of the expanded list, and select **All Files**.
A specific folder and its content (including subfolders and their files)	Expand the volume, find the desired folder, and select its checkbox.
All files within a specific folder (but not the folder)	Expand the folder and select **All Files**.
Specific files and folders	Select the **Show individual files** checkbox, expand the folders, and select the files and folders you want to replicate.

2. (Optional) Click the **Review** button to see your selections.

3. Click **Next**.

   The File and Folder Filters page appears.

Step 7: File and Folder Filters

The **File and Folder Filters** page displays a list of file and folder filters. By default, all files and folders selected in the **Source Paths** page will be replicated. A file or folder filter enables you to exclude and/or include files and folders from the job based on file type, extension, name, or directory path. Any file or folder that matches the filter is excluded or included from replication, depending on the filter's definition.

1. Select the file and folder filters you want to apply to the job.

   If you want to create a new file or folder filter or modify an existing one, click **Edit File Filters**. See File and Folder Filters in the Preferences section for information about creating or modifying a file filter.

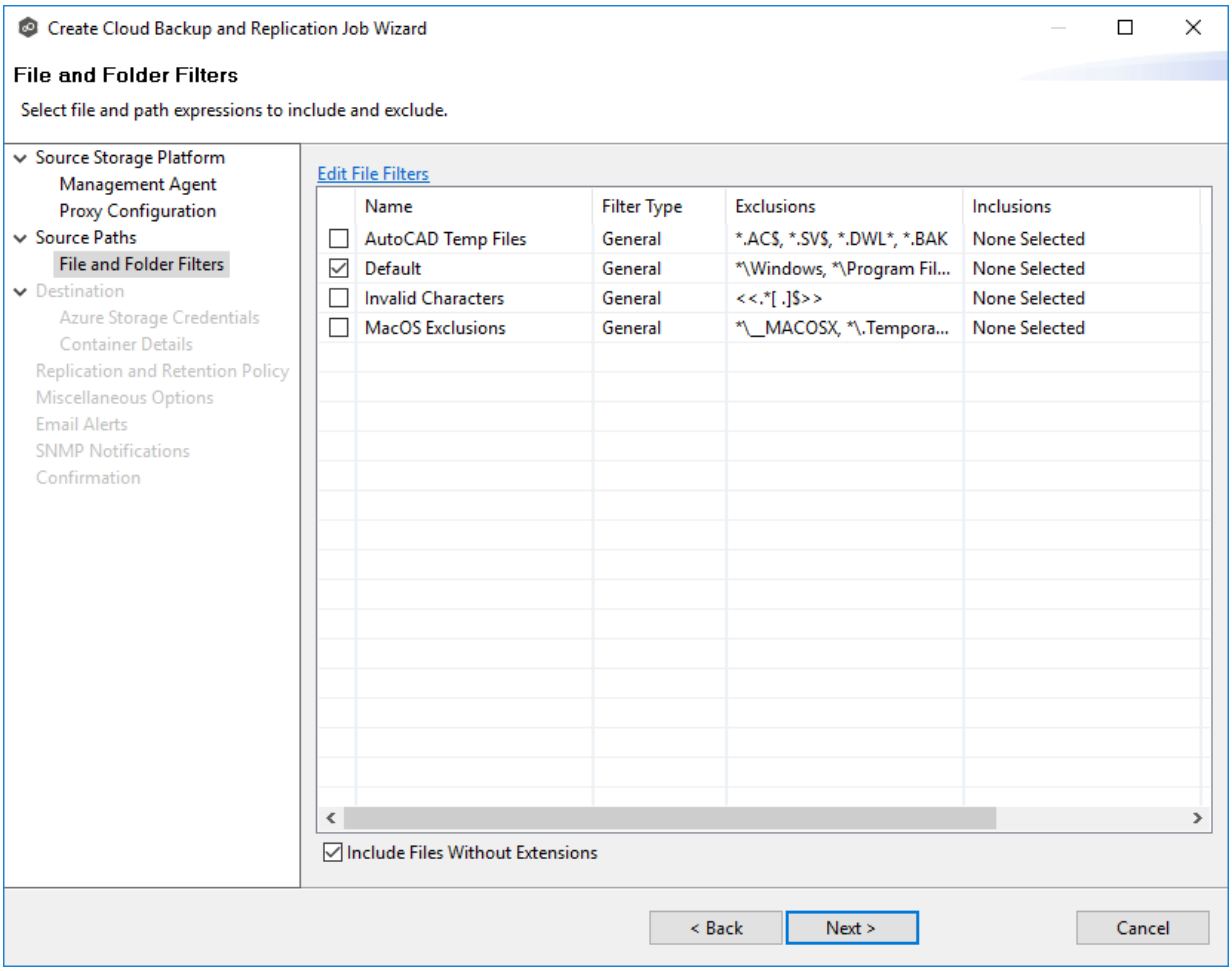

2. Select the **Include Files Without Extensions** checkbox if you want to replicate files that do not have extensions.

   Note: Files without extensions are ignored during replication unless you select this checkbox.

3. Click **Next**.

   The Destination page appears.

Step 8: Destination

The **Destination** page displays a list of the available storage platforms to which Cloud Backup and Replication can replicate. Currently, the following platforms are supported:

- Microsoft Azure

- Amazon S3
- NetApp StorageGRID
- Nutanix Objects

In addition, some S3-compatible platforms are also supported. Contact your Peer Software Sales representative to see if the S3 compatible platform you want to use is supported.

Important: You should create the storage account before creating the Cloud Backup and Replication job.

1. Select the type of destination storage platform.

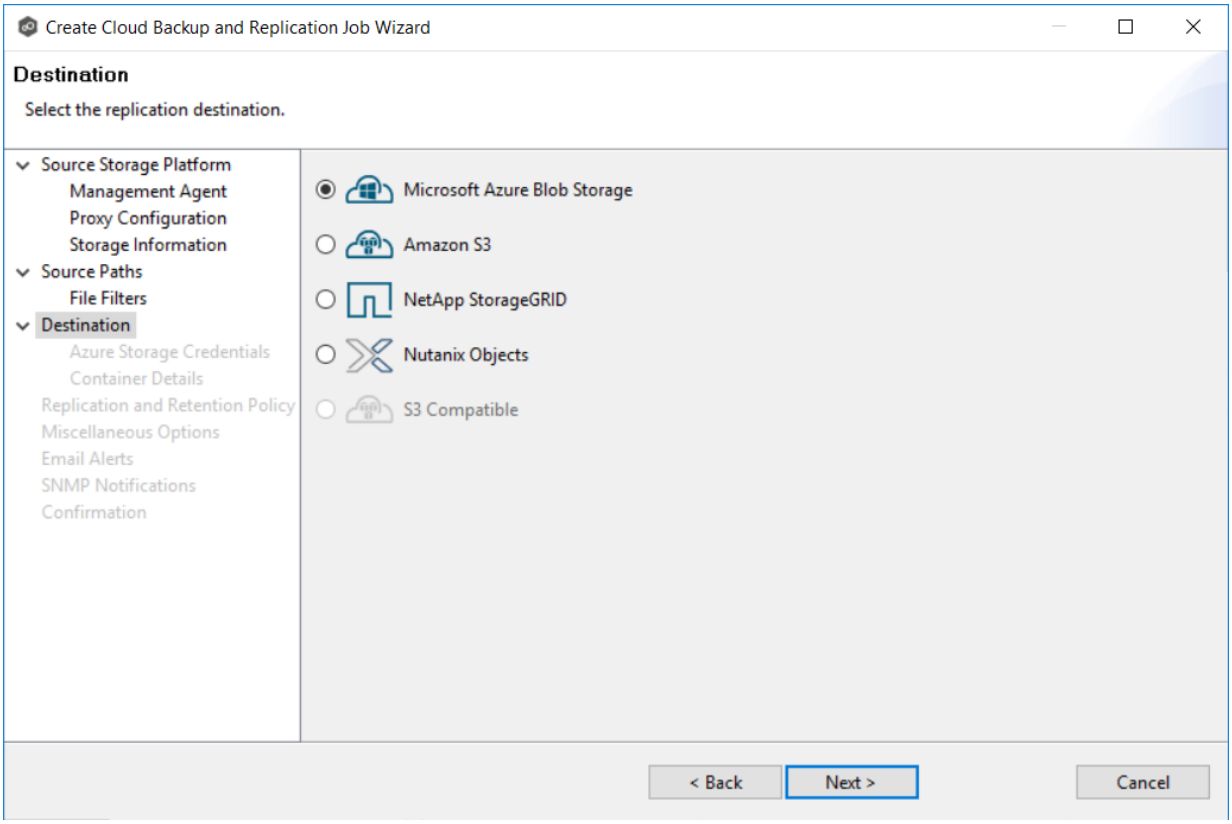

2. Click **Next**.

   The Destination Credentials page appears.

Step 9: Destination Credentials

The **Credentials** page requests the credentials necessary to connect to the destination storage account.

1. Select **New Credentials** to enter a new set of credentials for the destination storage device or select **Existing Credentials**.

2. If you selected **Existing Credentials**, select a credential from the drop-down list.

   If you selected **New Credentials**, enter the credentials for connecting to the destination storage account. The information you are prompted to enter varies, depending on the type of storage platform:

   Azure Blob Storage Credentials

   Amazon S3 Credentials

   NetApp StorageGRID

   Nutanix Objects

3. Click **Next**.

   The **Details** page for the selected destination storage account.

1. Enter the credentials to connect to a Microsoft Azure storage account. General Purpose and Blob storage accounts are supported.

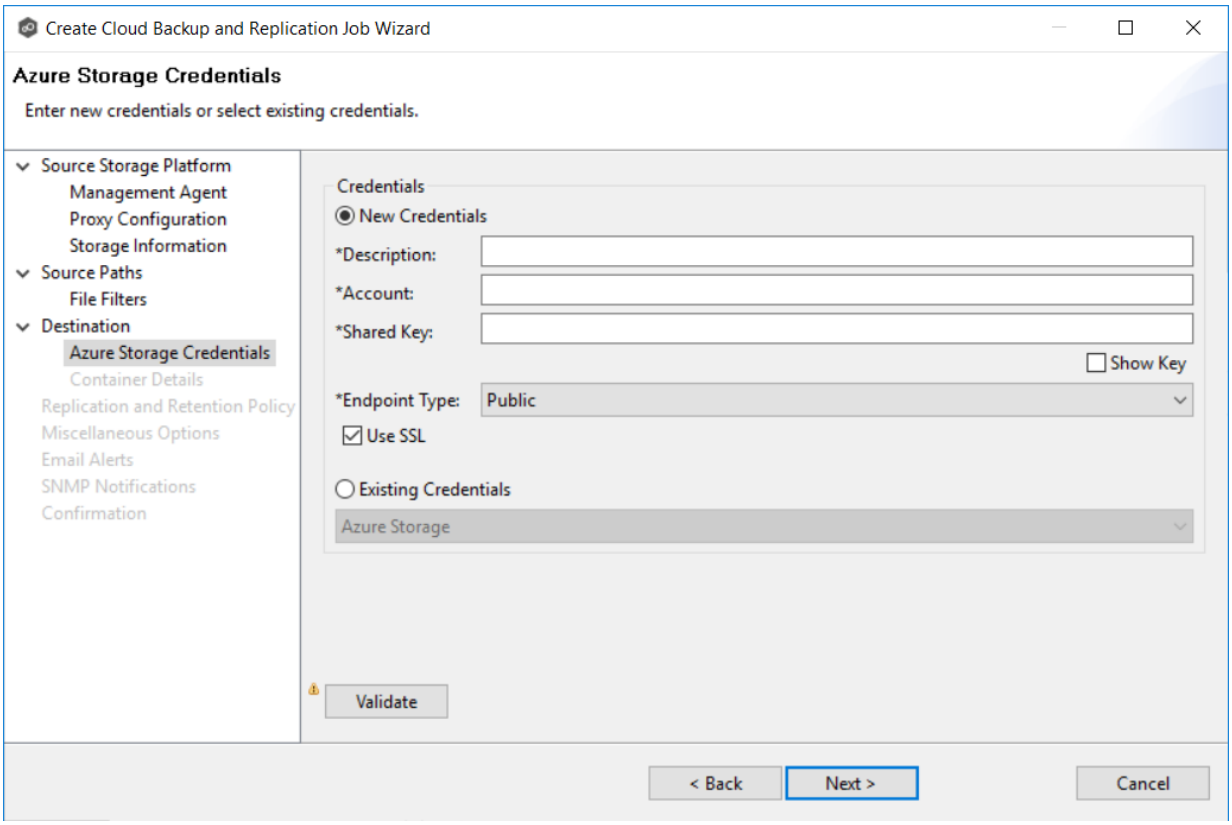

Description	Enter a name for the credentials.
Account	Enter the name of the Azure storage account, which can be found in the Azure Portal.
Shared Key	Enter one of the shared keys for the Azure Storage account. The shared keys can be found in the Azure Portal.
Endpoint Type	Select the type of data center endpoint. The options are: **Public**, **Germany**, **China**, **US Government**, and **Custom**.
Endpoint	If you selected **Custom** for **Endpoint Type**, the **Endpoint** field appears. Enter the IP address of the endpoint.
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

2. Click **Validate** to test the connection.

   If you are using a proxy in your environment and you get an error while trying to validate, you may want to check the proxy configuration in Preferences.

3. Click **Next**.

   The Container Details page appears.

1. Enter the credentials to connect to an Amazon S3 storage account.

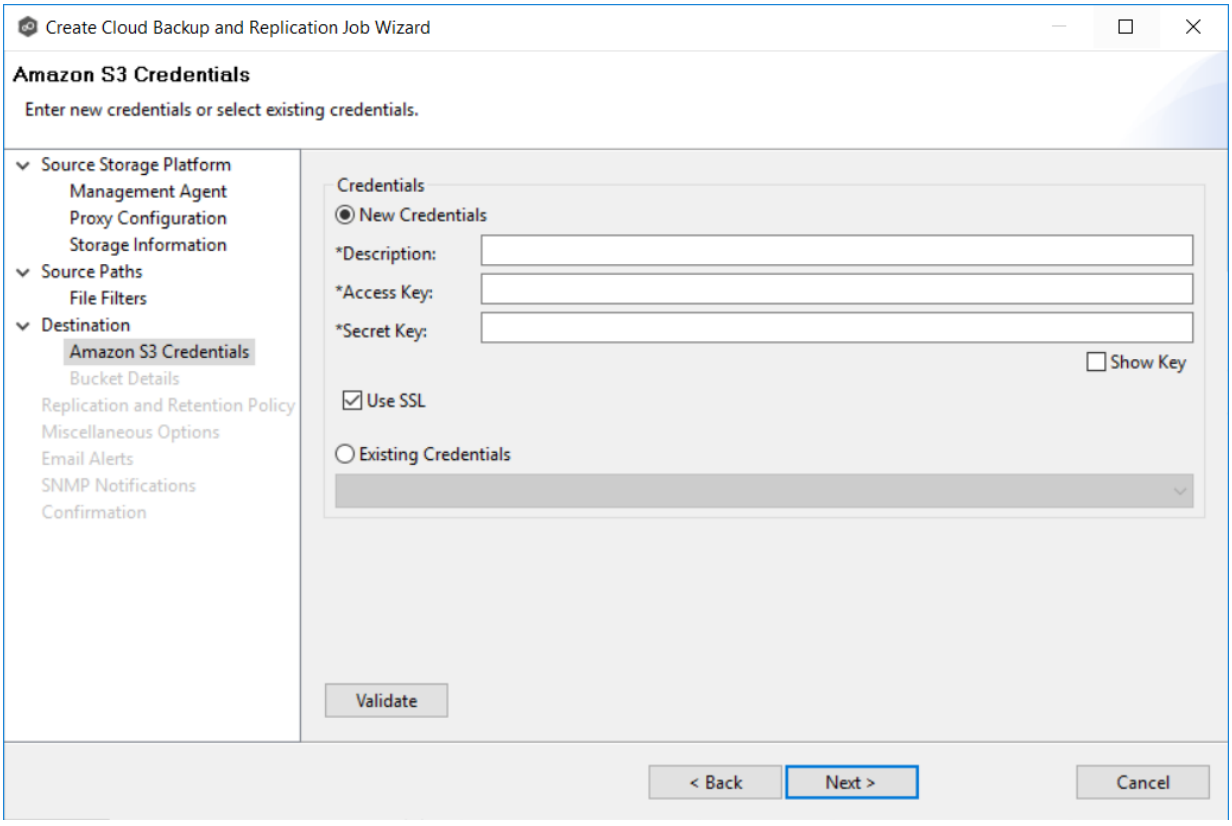

Description	Enter a name for the credentials.

Access Key	Enter one of the shared keys of the Amazon S3 Storage account, which can be found in the Amazon AWS portal.
Secret Key	Enter the secret key of the Amazon S3 Storage account, which can be found in Amazon AWS portal.
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

2. Click **Validate** to test the connection.

   If you are using a proxy in your environment and you get an error while trying to validate, you may want to check the check the proxy configuration in Preferences.

3. Click **Next**.

   The Bucket Details page appears.

1. Enter the credentials to connect to a NetApp StorageGRID storage account.

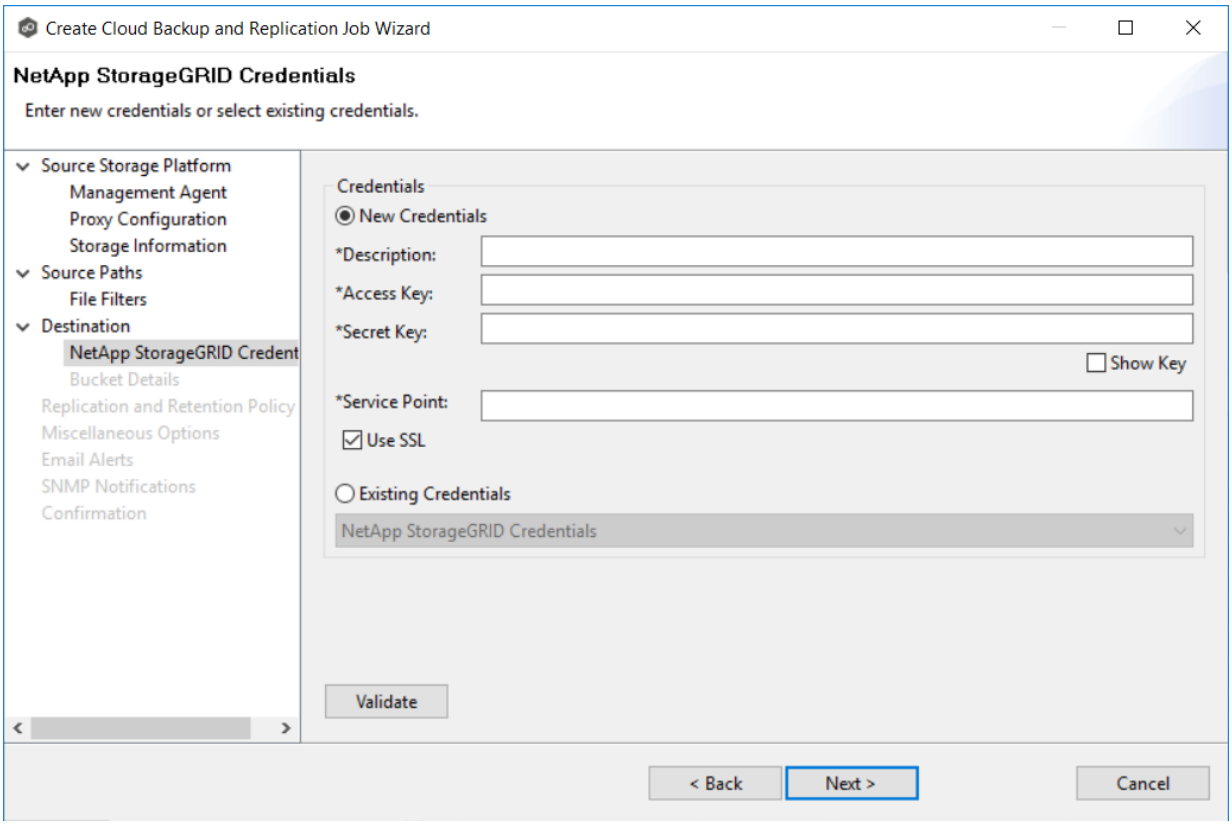

Description	Enter a name for the credentials.
Access Key	Enter one of the shared keys of the NetApp StorageGRID account, which can be found in the Tenant Manager.
Secret Key	Enter the secret key of the NetApp StorageGRID account, which can be found in the Tenant Manager.
Service	Enter the IP or name of the object store.

Point	
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

2. Click **Validate** to test the connection.
3. Click **Next**.

   The Container Details page appears.

1. Enter the credentials to connect to a Nutanix Objects storage account.

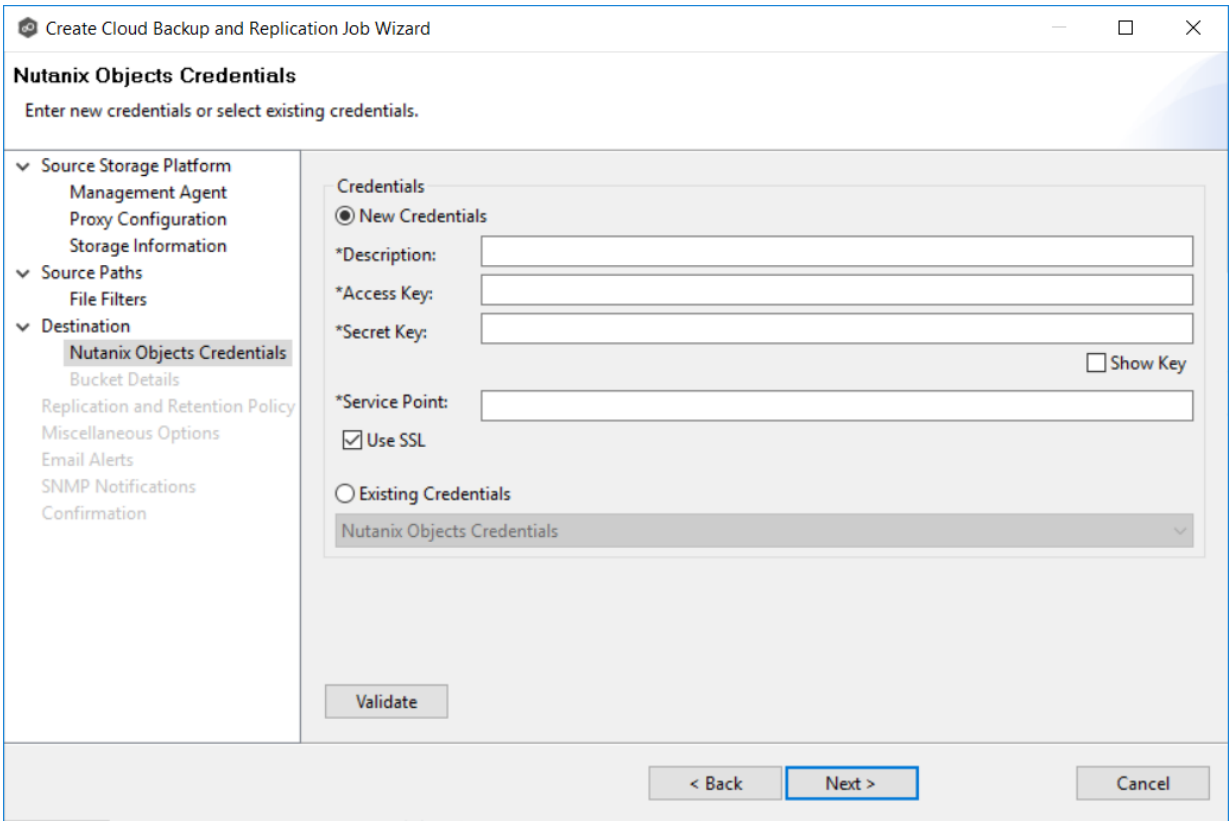

Description	Enter a name for the credentials.
Access Key	Enter one of the shared keys of the Nutanix Objects account, which can be found in Prism Central.
Secret Key	Enter the secret key of the Nutanix Objects account, which can be found in Prism Central.
Service Point	Enter the IP or name of the object store.
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

2. Click **Validate** to test the connection.

3. Click **Next**.

   The Container Details page appears.

Step 10: Container or Bucket Details

The **Container Details** or **Bucket Details** page allow you to create a new storage container or bucket or choose an existing one.

1. Select **New Container**/**New Bucket** to create a new storage container/bucket; otherwise, select **Existing Container**/**Existing Bucket** to choose an existing one.

2. If you selected **Existing Container** or **Existing Bucket**, select a container or bucket from the drop-down list.

   If you selected **New Container** or **New Bucket**, enter the requested information. The information you are prompted to enter varies, depending on the type of storage platform:

   Azure Blob Storage Container Details

Amazon S3 Bucket Details

NetApp StorageGRID Bucket Details

Nutanix Objects Bucket Details

3. Click **Next**.

   The Replication and Retention Policy page appears.

1. Select **New Container** to create a new container or select **Existing Container**.

   Choose **Existing Container** if:

   - You (or someone else) already created a container you want to use.
   - You want to use a container that was created outside Peer Management Center.
   - You don't have the permissions required to create a new container and want to use one that someone else will create.

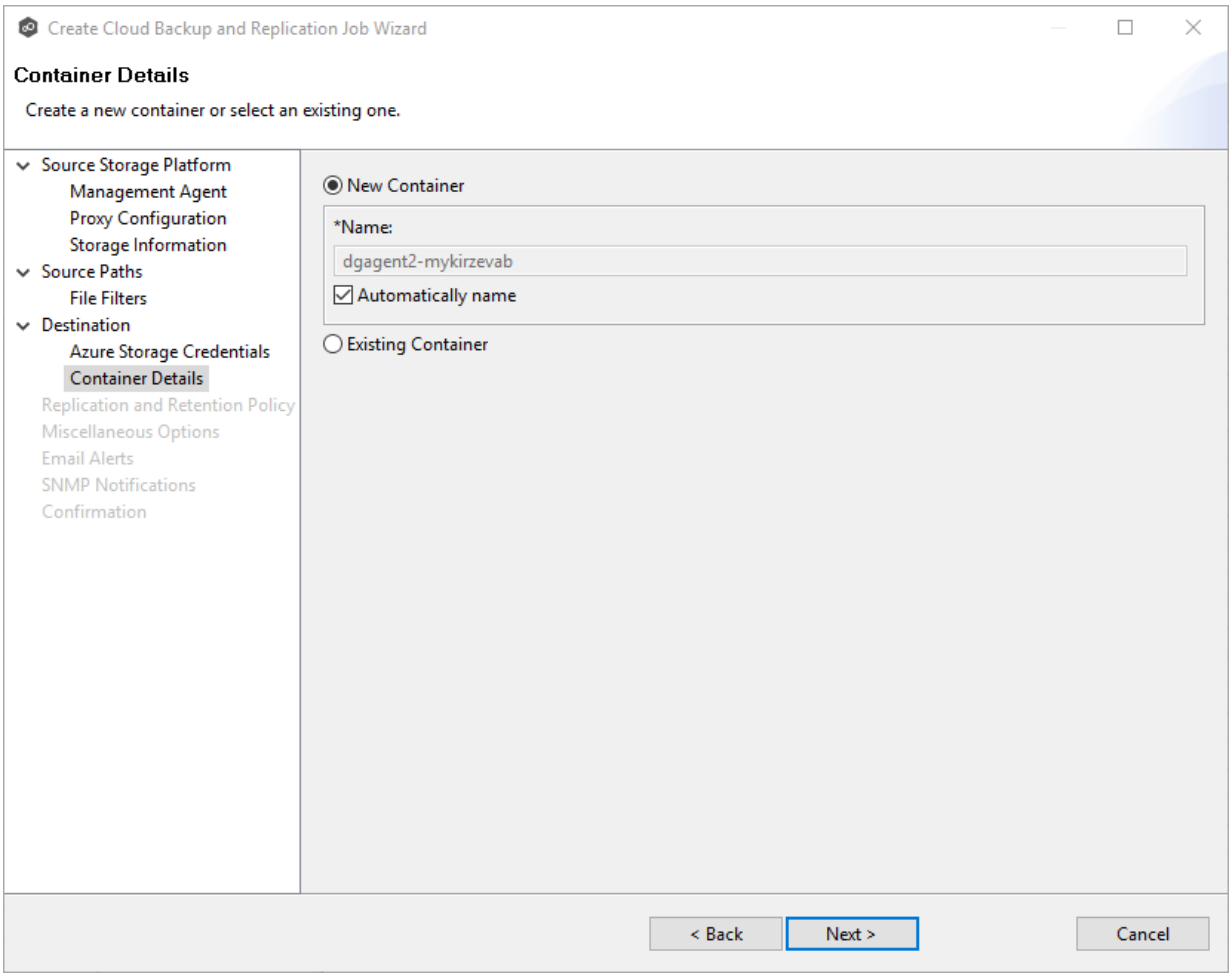

2. If you selected **Existing Container**, select a container from the drop-down list. If the container does not appear in the list because the person who has the permissions to create a container has not yet created the bucket, click the **Reload** button after the container is created. The container will appear in the updated list.

   If you selected **New Container**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the container; the container name must conform to the following naming rules:

   - A container name must be unique.
   - A container name must be a valid DNS name.
   - A container name must start with a letter or number, and can contain only letters, numbers, and the dash (-) character.
   - Every dash (-) character must be immediately preceded and followed by a letter or number; consecutive dashes are not permitted in container names.

- All letters in a container name must be lowercase.
- A container name must be from 3 through 63 characters long.

For more information about container names, see Naming and referencing containers, blobs, and metadata.

3. Click **Next**.

   The Replication and Retention Policy page appears.

1. Select **New Bucket** to create a new bucket or select **Existing Bucket**.

   Choose **Existing Bucket** if:

   - You (or someone else) already created a bucket you want to use.
   - You want to use a bucket that was created outside Peer Management Center.
   - You don't have the permissions required to create new buckets and want to use one that someone else will create.

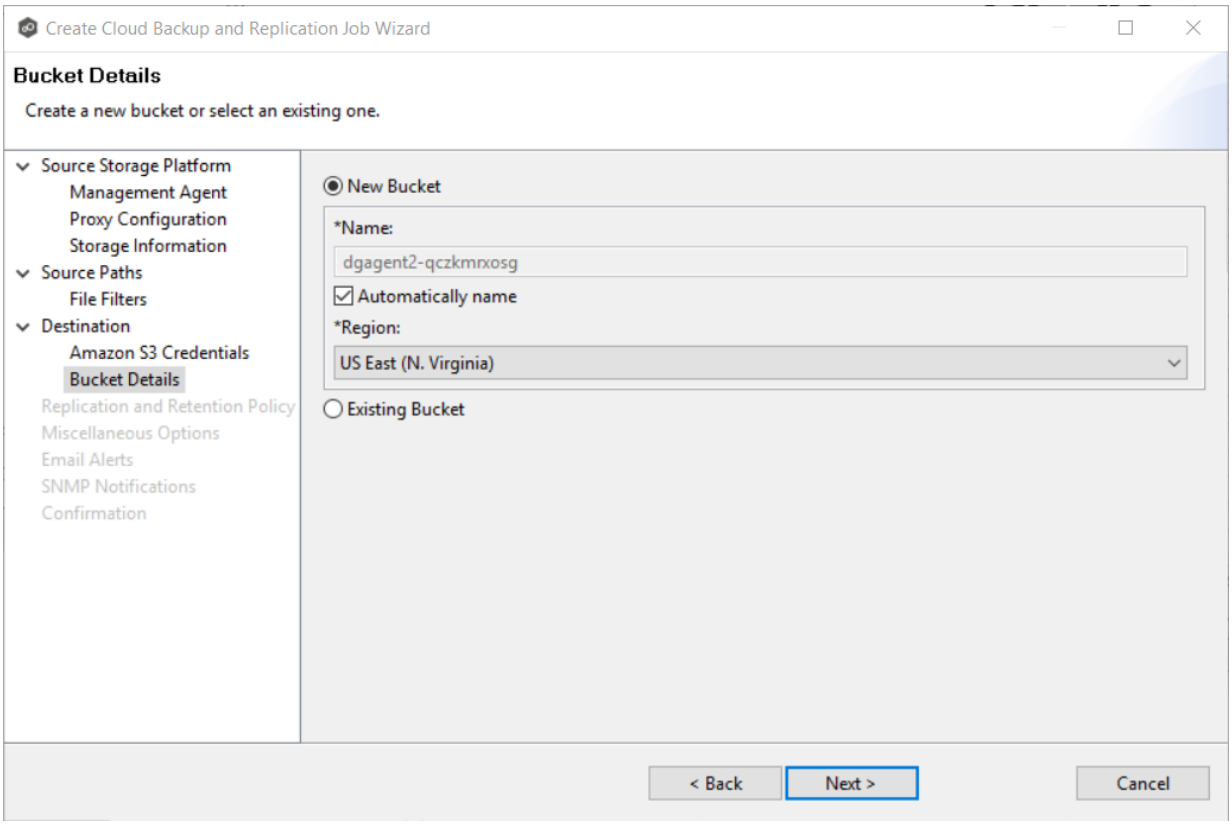

-

2. If you selected **Existing Bucket**, select a bucket from the drop-down list. If the bucket does not appear in the list because the person who has the permissions to create a bucket has not yet created the bucket, click **Reload** after the bucket is created. The bucket will appear in the updated list.

   If you selected **New Bucket**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the bucket; the bucket name must conform to the following naming rules:

   - A bucket name must be unique across all existing bucket names in Amazon S3 (that is, across all AWS customers). For more information, see Bucket Restrictions and Limitations.
   - Bucket names must comply with DNS naming conventions. For information about legacy non-DNS-compliant bucket names, see Bucket Restrictions and Limitations.
   - A bucket name must start with a lowercase letter or number.
   - A bucket name must not contain uppercase characters or underscores.
   - A bucket name must be from 3 through 63 characters long.

- A bucket name must be a series of one or more labels. Adjacent labels are separated by a single period (.). Bucket names can contain lowercase letters, numbers, and hyphens. Each label must start and end with a lowercase letter or a number.
- A bucket name must not be formatted as an IP address (for example, 192.168.5.4).
- When you use virtual hosted–style buckets with Secure Sockets Layer (SSL), the SSL wildcard certificate only matches buckets that don't contain periods. To work around this, use HTTP or write your own certificate verification logic. We recommend that you do not use periods (.) in bucket names when using virtual hosted–style buckets.
- Choose a bucket name that reflects the objects in the bucket because the bucket name is visible in the URL that points to the objects that you're going to put in your bucket. After you create the bucket, you cannot change the name, so choose wisely.

For information about naming buckets, see Rules for Bucket Naming in the Amazon Simple Storage Service Developer Guide.

3. Select the region where you want the bucket to reside.

   Important: After you have created a bucket, you cannot change its region.

4. Click **Next**.

   The Replication and Retention Policy page appears.

1. Select **New Bucket** to create a new bucket or select **Existing Bucket**.

   Choose **Existing Bucket** if:

   - You (or someone else) already created a bucket you want to use.
   - You want to use a bucket that was created outside Peer Management Center.
   - You don't have the permissions required to create new buckets and want to use one that someone else will create.

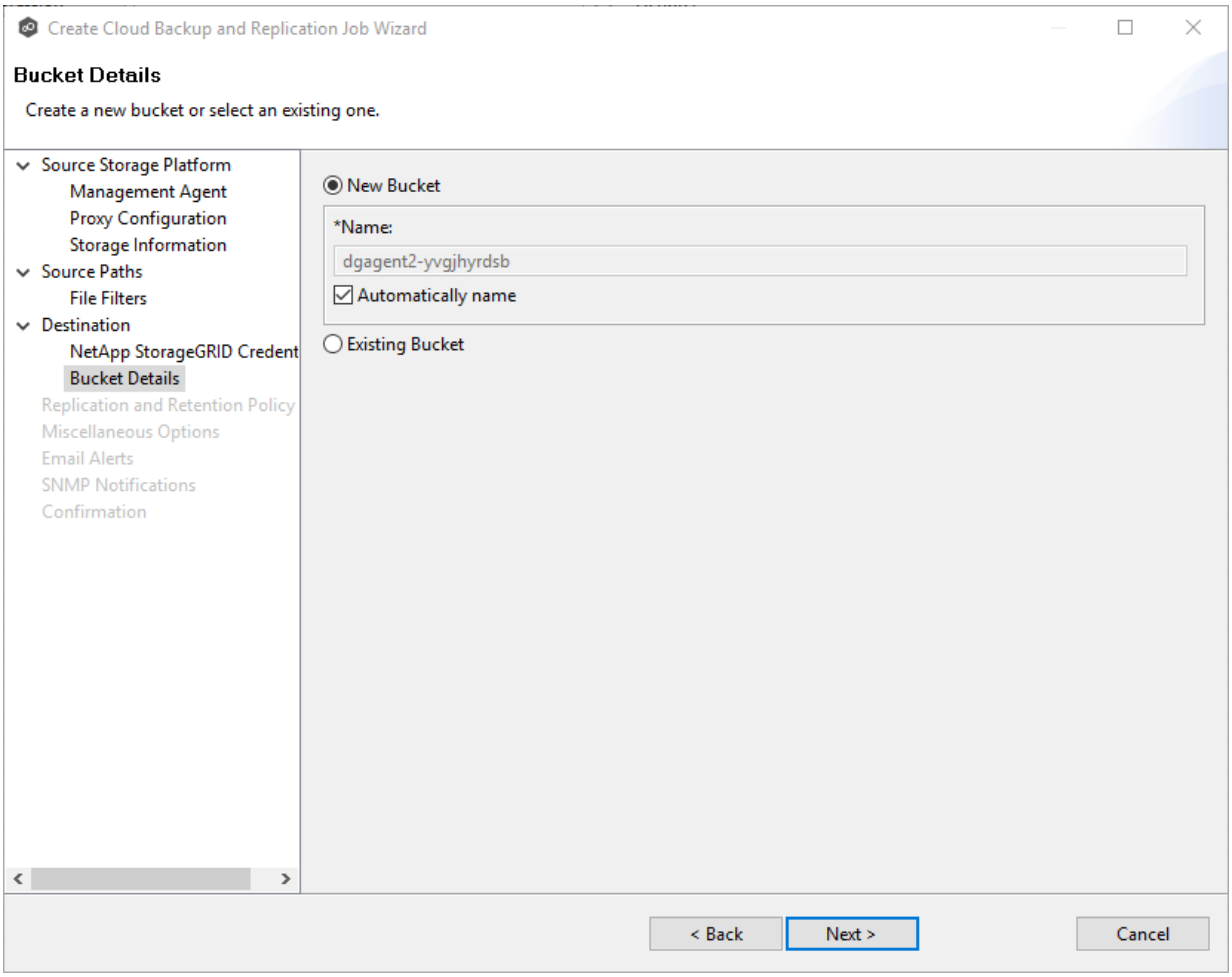

1. If you selected **Existing Bucket**, select a bucket from the drop-down list. If the bucket does not appear in the list because the person who has the permissions to create a bucket has not yet created the bucket, click **Reload** after the bucket is created. The bucket will appear in the updated list

   If you selected **New Bucket**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the bucket; the bucket name must comply with the following rules:

   - Must be unique across each StorageGRID Webscale system (not just unique within the tenant account).
   - Must be DNS compliant.
   - Must contain between 3 and 63 characters.
   - Can be a series of one or more labels, with adjacent labels separated by a period. Each label must start and end with a lowercase letter or a number and can only use lowercase letters, numbers, and hyphens.

- Must not look like a text-formatted IP address.
- Should not use periods in virtual hosted-style requests because periods will cause problems with server wildcard certificate verification.

For information about naming buckets, see Rules for Bucket Naming in the Amazon Simple Storage Service Developer Guide.

2. Click **Next**.

   The Replication and Retention Policy page appears.

1. Select **New Bucket** to create a new bucket or select **Existing Bucket**.

   Choose **Existing Bucket** if:

   - You (or someone else) already created a bucket you want to use.
   - You want to use a bucket that was created outside Peer Management Center.
   - You don't have the permissions required to create new buckets and want to use one that someone else will create.

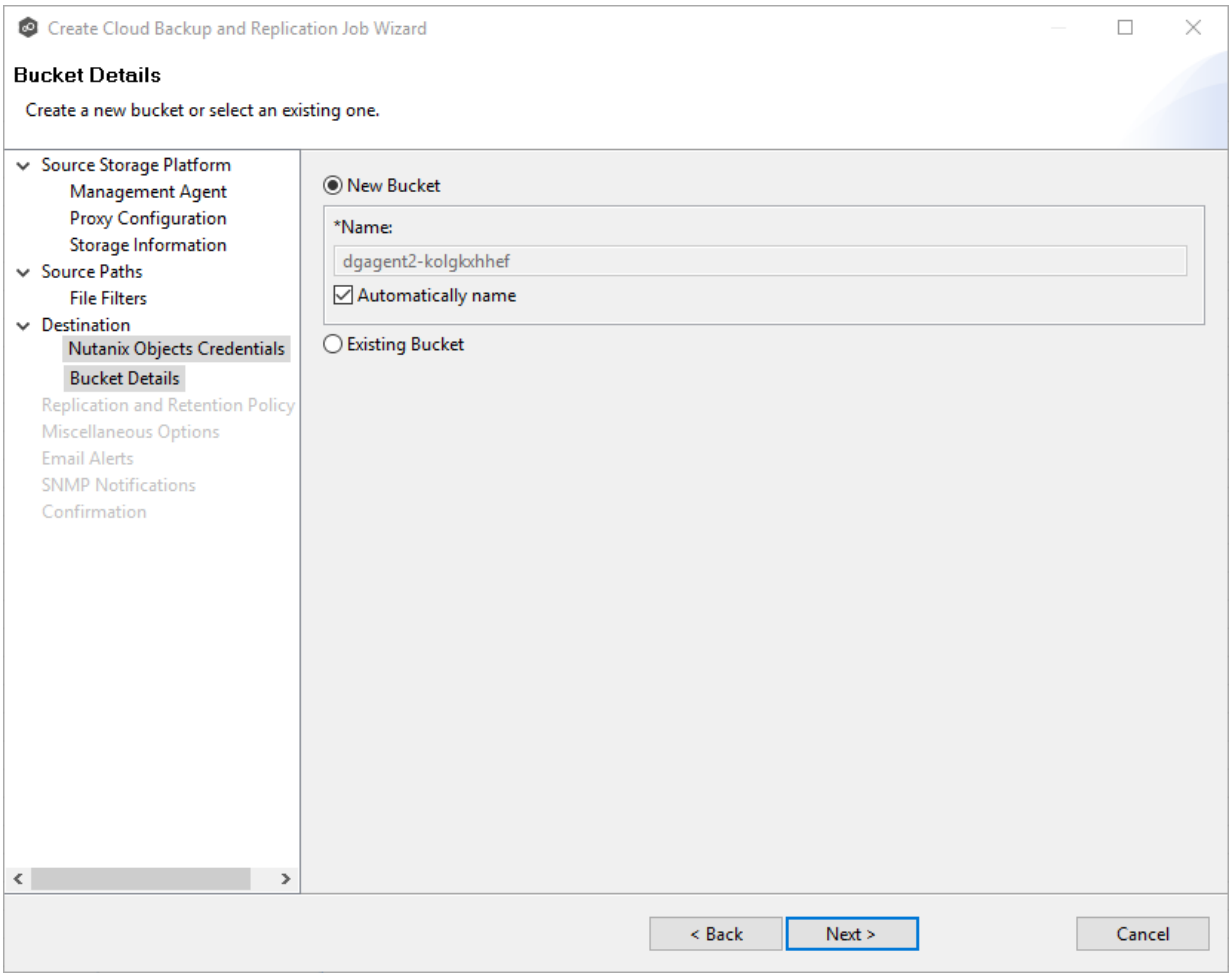

2. If you selected **Existing Bucket**, select a bucket from the drop-down list. If the bucket does not appear in the list because the person who has the permissions to create a bucket has not yet created the bucket, click **Reload** after the bucket is created. The bucket will appear in the updated list

   If you selected **New Bucket**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the bucket; the bucket name must comply with the following rules:

   - Must start with a number or a letter.
   - Must be 3 - 255 characters long.
   - Can contain lowercase letters, numbers, underscores (_), and dashes (-).
   - There may be additional restrictions on bucket names in some AWS regions. We recommend that you create bucket names that are DNS-compliant, if you want to access objects using URL. For more information, see Amazon Simple Storage Service Console user's guide.

3. Click **Next**.

   The Replication and Retention Policy page appears.

Step 11: Replication and Retention Policy

Each Cloud Backup and Replication job must have a Replication and Retention policy. A Replication and Retention policy specifies:

- How often you want to scan the storage device for replication or if you want to replicate in real-time.
- Whether you want to take snapshots of the data. A **snapshot** captures the state of a file system at a point in time. There are two types of snapshots:
  - A **destination snapshot** captures an image of the data on the destination storage device immediately after replication. Destination snapshots are useful for recovering data from different period of times. Destination snapshots track versions of changed files and file system structure that can be used for data recovery. For more information about recovering data, see Recovering Data.
  - A **source snapshot** captures an image of the data on the source storage device immediately before replication. Sources snapshots are useful for replicating open and locked files, which otherwise may not be able to be replicated. A source snapshot also ensures that the replicated data is coming from a static version of the source file system. For details about using source snapshots, see Step 13: Source Snapshots.
- How long you want to retain destination snapshots.

The **Replication and Retention Policy** page enables you to create a new Replication and Retention policy or choose an existing policy.

1. Select **New Policy** or **Existing Policy**.

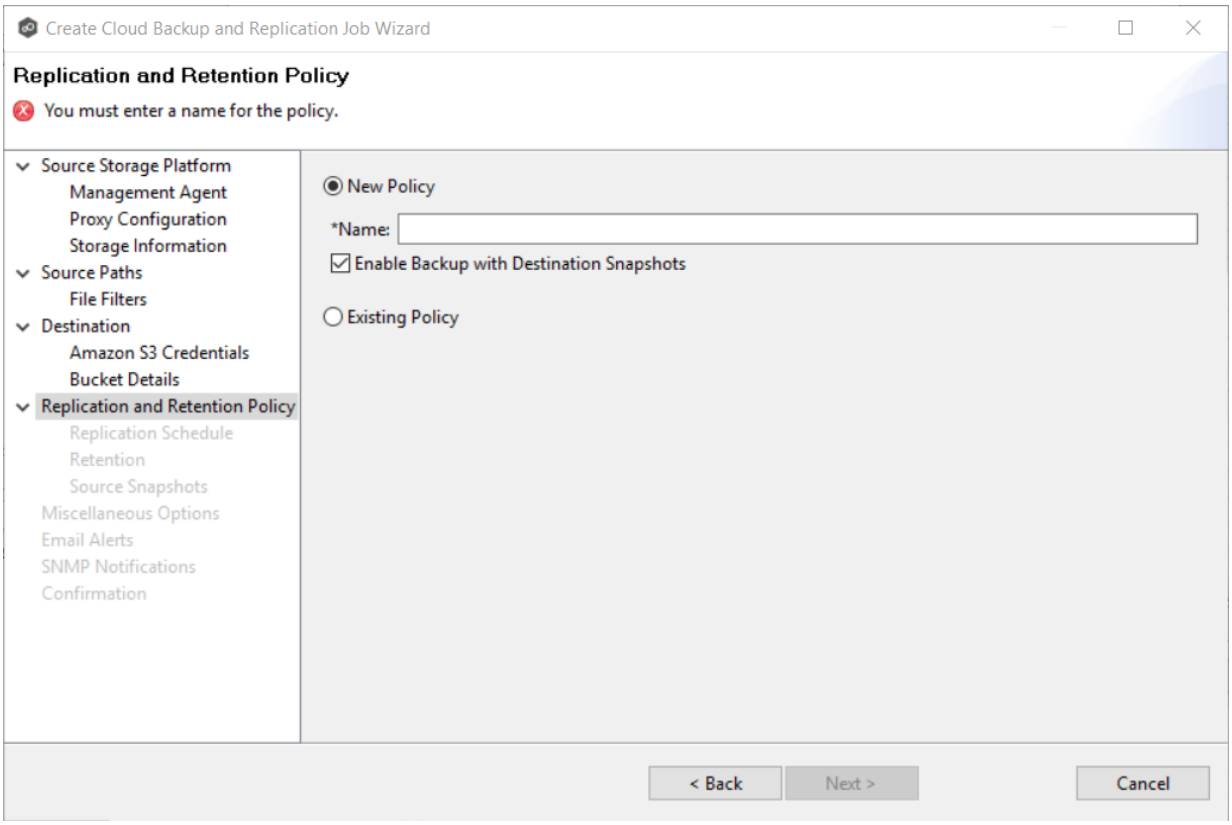

2. If you selected **Existing Policy**, select a policy from the drop-down list, and then click **Next**. Continue with Step 14. Miscellaneous Options.

   If you selected **New Policy**, enter a name for the policy in the **Name** field.

3. Select **Enable Backup with Destination Snapshots** if you want to replicate what is on premises to the destination storage device, while taking destination snapshots at specified points in times.

4. Click **Next**.

   The Replication Schedule page appears.

Step 12: Replication Schedule

The **Replication Schedule** page enables you to select the frequency of the replication and when snapshots should be taken. Replication can be performed on a scheduled, batched real-time, or a continuous real-time basis.

1. Select the frequency of the replication:

- Scheduled Scans – Select this option if you want to replicate files on a scheduled basis. A scan of changes to the file system occurs on a scheduled basis, either daily or weekly, and replication of changes occurs as the scan progresses.

- Batched Real-time – Select this option if you want to continuously monitor changes to the file system but replicate changes on scheduled basis. Changes are monitored in real-time and only the latest version of changed file is replicated at scheduled times. An initial scan can be performed to establish a baseline.

- Continuous Data Protection – Select this option if you want continuously monitor changes and replicate changes in real-time. Whenever a file changes, the change is replicated in real-time.

2. Click **Next**.

   The Retention page appears.

If you selected **Scheduled Scans** for the replication frequency:

1. Select the **Scan at Start** checkbox if you want a baseline replication to be performed.

2. Select **Daily** or **Weekly** for the frequency of the scans:

   - Select **Daily** if you want replications performed every day. You can schedule one to four scans per day

   - Select **Weekly** if you want to select specific days for replication. You can select one scan per day.

3. Select the day(s) and time(s) when you want the replication performed:

   - If you selected **Daily**, select the times you want the scans performed. Then, if you selected **Enable Backup with Destination Snapshots** in Step 10, choose when snapshots are taken (you must take at least one snapshot). If you did not select the backup option, the **Take Destination Snapshot** options will not appear.

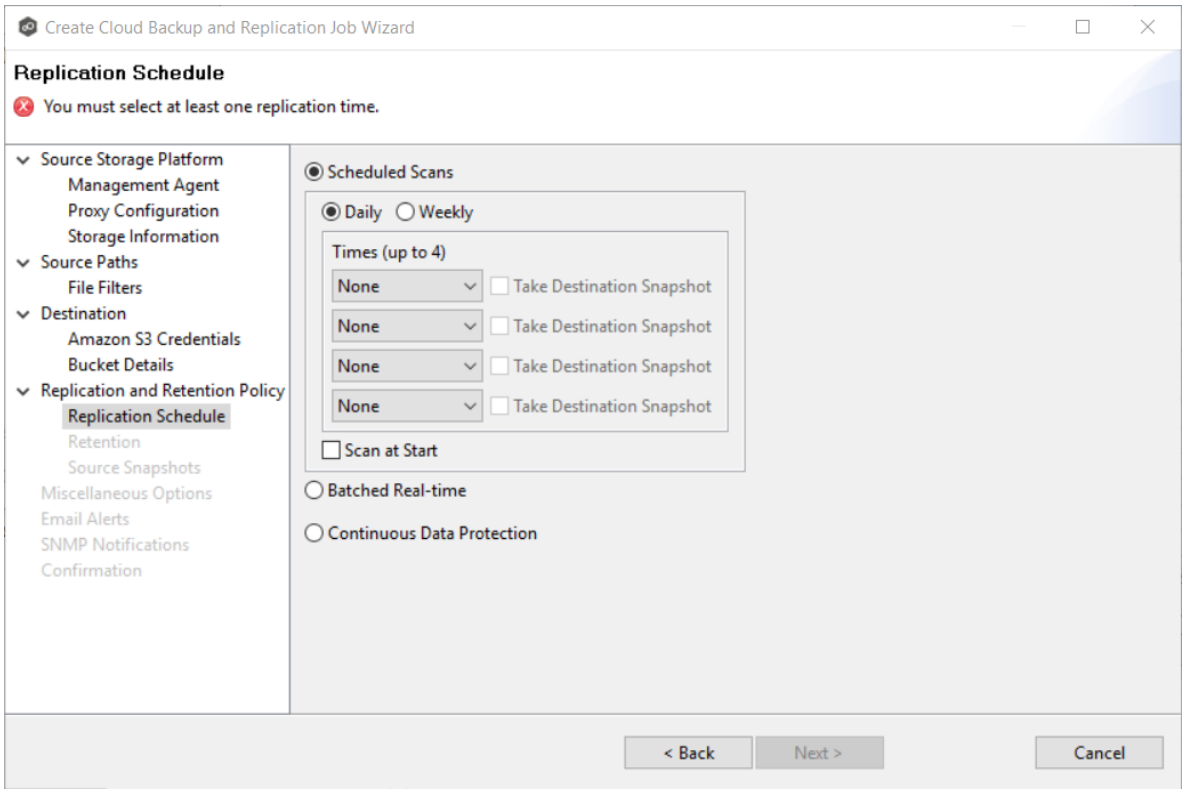

- If you selected **Weekly**, select the day(s) and time you want the replication performed. Then, if you selected **Enable Backup with Destination Snapshots** in Step 10, choose when snapshots are taken. You must take at least one snapshot. If you did not select the backup option, the **Destination Snapshot** option will not appear.

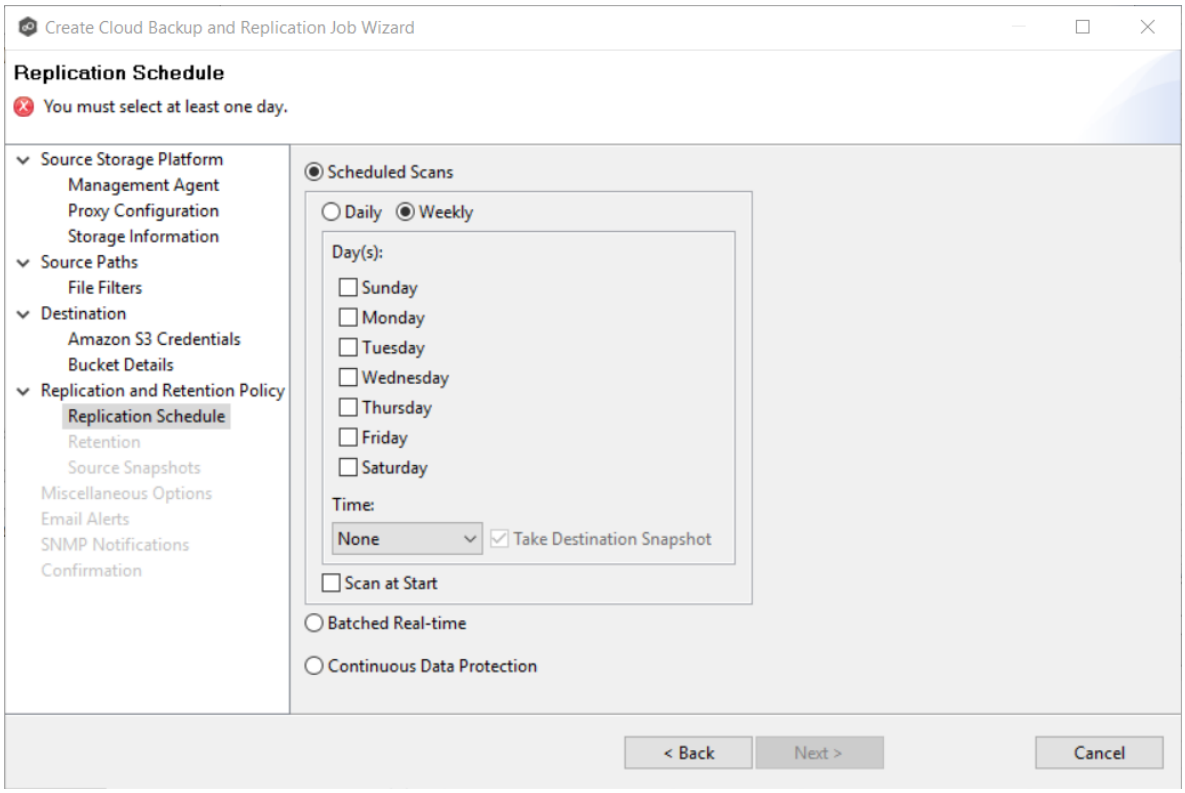

4. Click **Next**.

   The Retention page appears.

If you selected **Batched Real-time** for the replication frequency:

1. Select **Scan at Start** if you want a baseline replication to be performed.

2. Select the frequency of the replications; you can schedule one to four replications per day.

3. If you selected **Enable Backup with Destination Snapshots** in Step 10, choose when destination snapshots are taken (you must take at least one snapshot). The destination snapshot will be taken after the files have been replicated. If you did not select the backup option, the **Take Destination Snapshot** option will not appear.

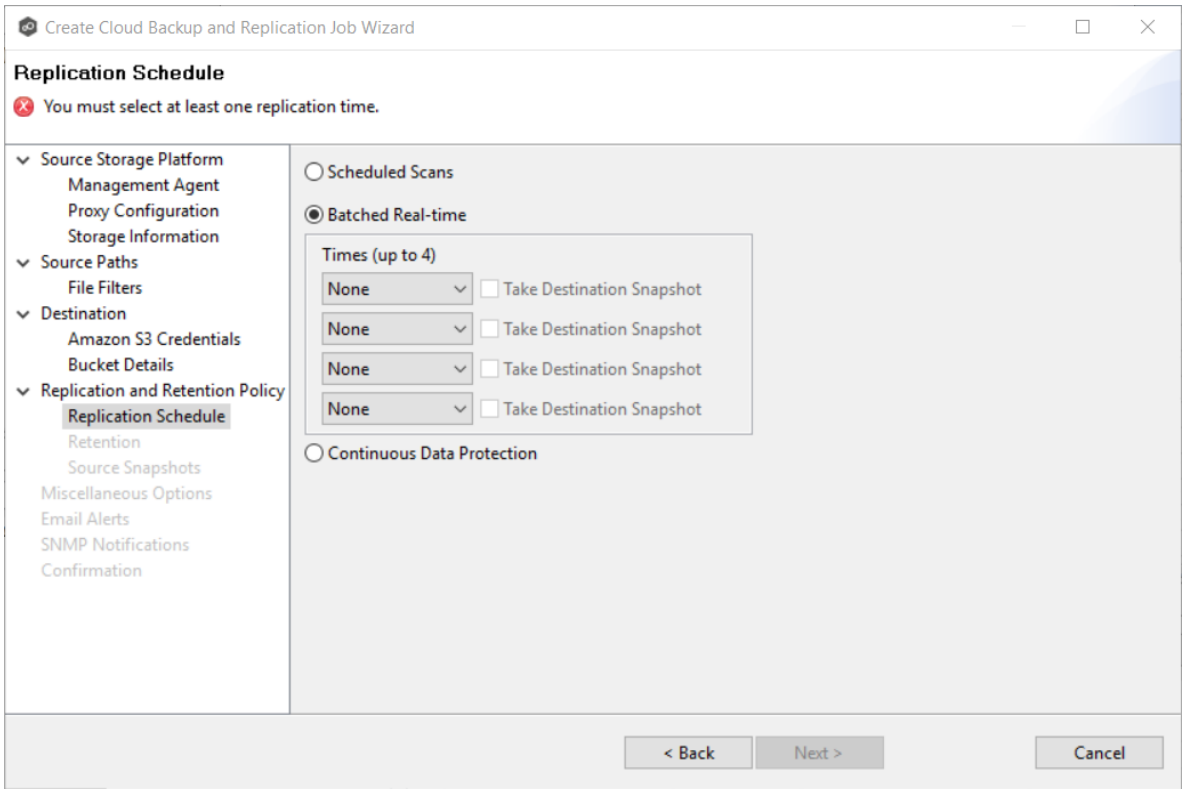

4. Click **Next**.

   The Retention page appears.

If you selected **Continuous Data Protection** for the replication frequency:

1. Enter a value for **Processing Delay** if you want the replication to occur after a slight delay. A delay is useful to ensure that when a file or folder is created and quickly renamed, only the latest copy of the file or folder is replicated. This reduces WAN usage.

2. If you selected **Enable Backup with Destination Snapshots** in Step 10, choose when snapshots are taken (you must take at least one snapshot). If you did not select the backup option, the **Take Destination Snapshot at** options will not appear.

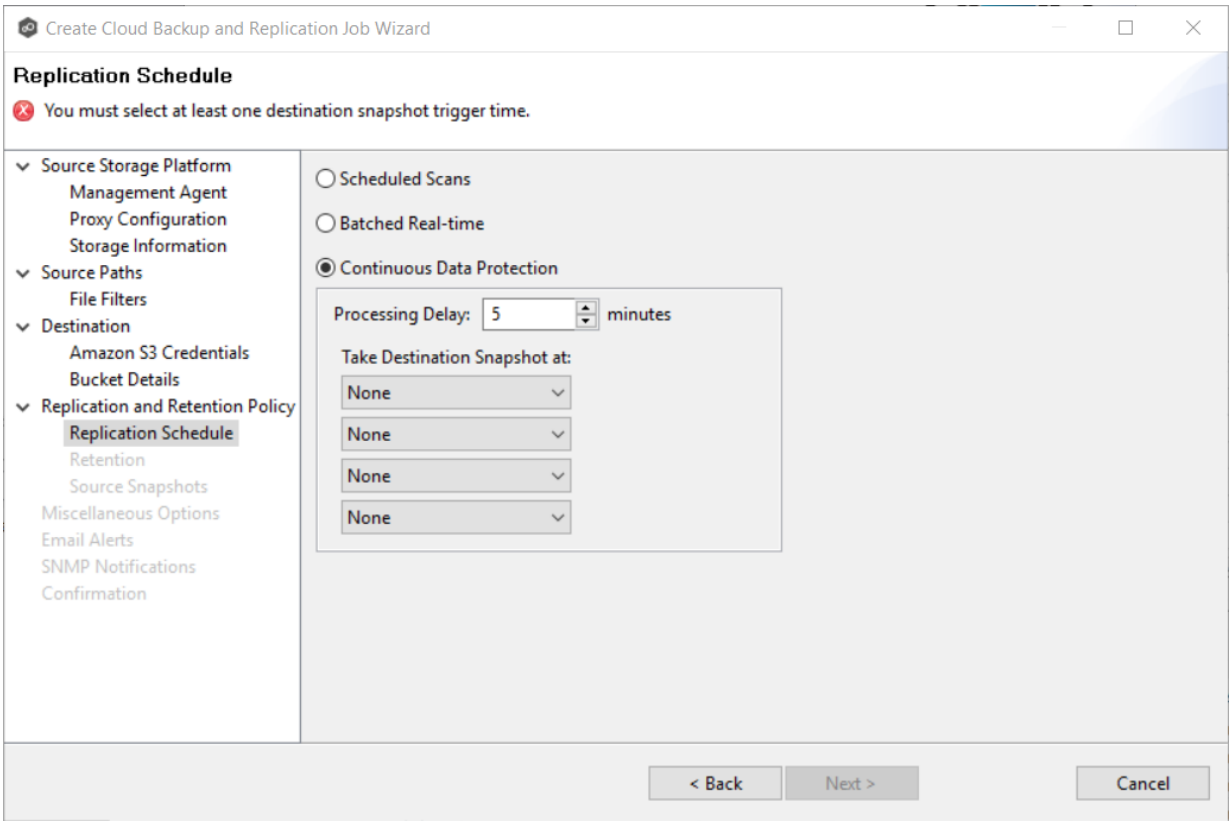

3. Click **Next**.

   The Retention page appears.

Step 13: Retention

The **Retention** page enables you to define how long you want to retain destination snapshots. You have the option to retain destination snapshots on a daily, weekly, monthly, and yearly basis. If you did not select the **Enable Backup with Destination Snapshots** in Step 10, the **Retention** page will not appear.

1. Select the **Purge all versions between snapshots** checkbox if you do not want to indefinitely retain all versions.

2. Select the retention options. The options vary according to the replication schedule you selected.

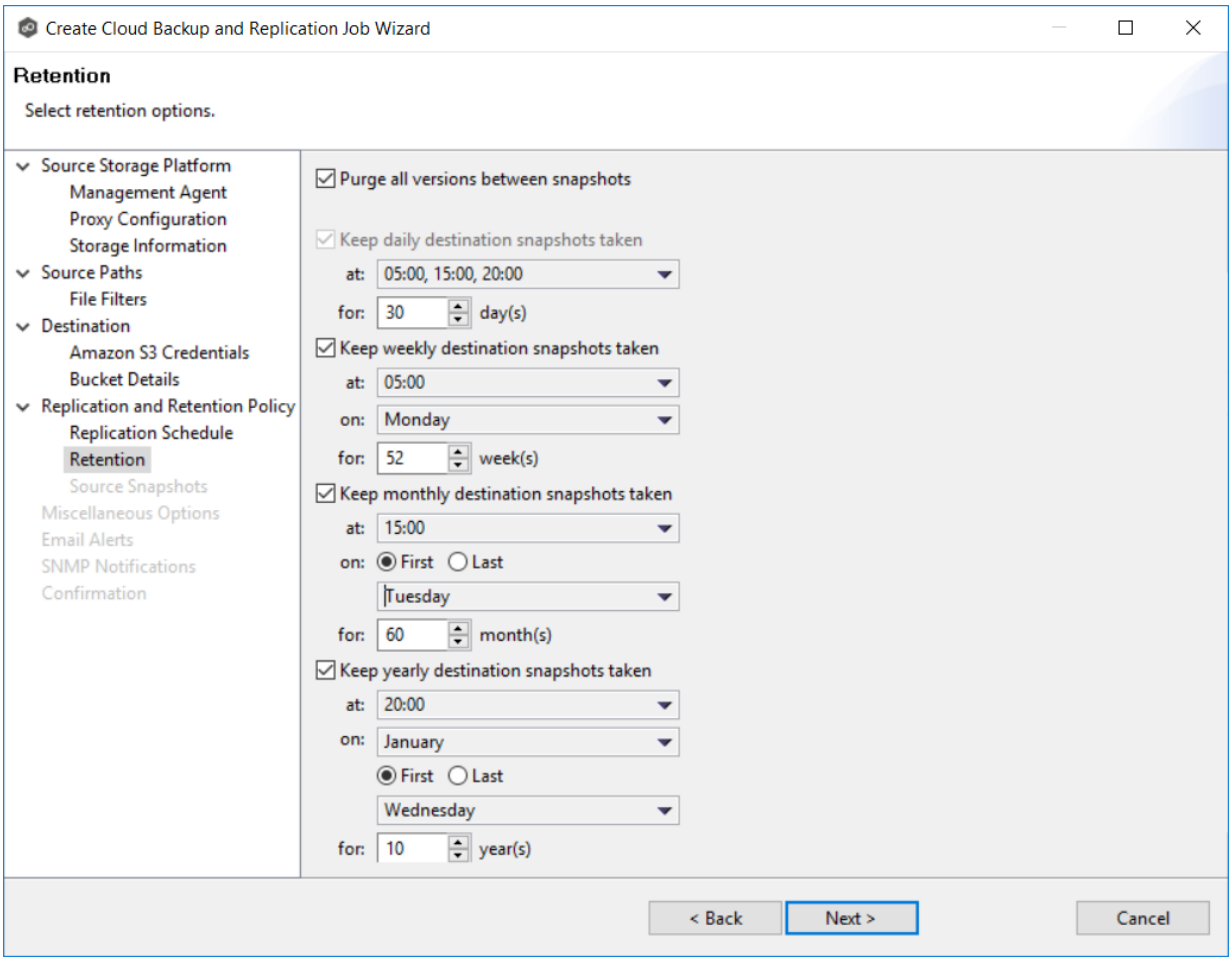

3. Click **Next**.

   The Source Snapshots page appears.

Step 14. Source Snapshots

The **Source Snapshots** page enables you to choose whether to take snapshots of the source storage before the items are replicated. A source snapshot is a read-only point-in-time version of the volume. A source snapshot allows the creation of consistent backups of a volume, ensuring that the contents do not change and are not locked while the backup is being made. It can be used to provide a consistent state of a managed file, e.g., pst files, and help with errors accessing files that are currently open.

1. Select a source snapshot option:

   - Select the **Disabled** option if you do not want to take source snapshots.

- Select the **Use only for Destination Snapshots** option when you want the source snapshot to be stored on the destination storage as the destination snapshot rather than an actual destination snapshot. To use this option, you must have selected the **Enable Backup with Destination Snapshot** in Step 10.

- Select **Always use when replicating** when you want to replicate always using source snapshots.

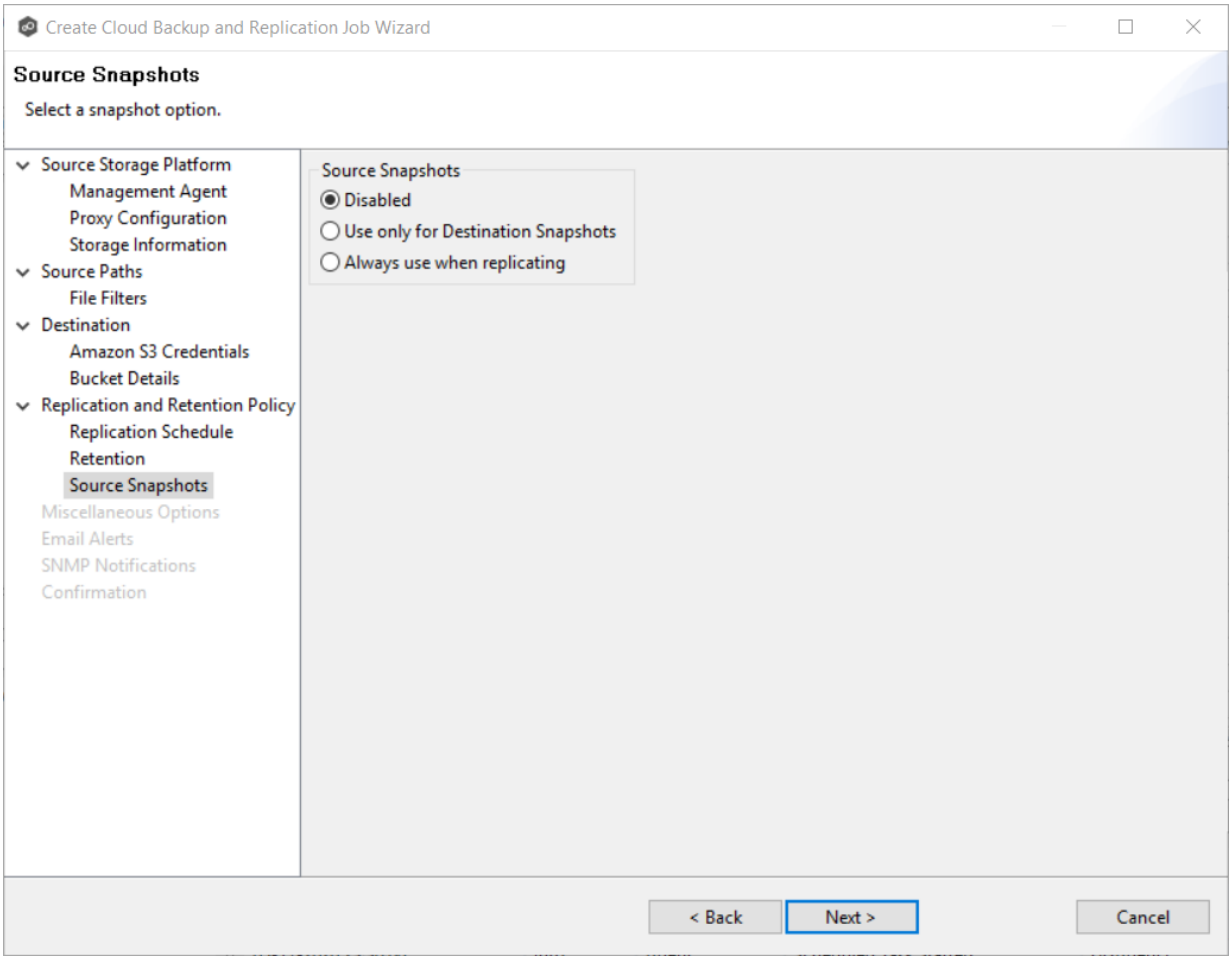

2. Click **Next**.

   The Miscellaneous Options page appears.

Step 15: Miscellaneous Options

The **Miscellaneous Options** page displays various options; the options available depend on the destination storage platform selected.

1. Select the options to apply to this job.

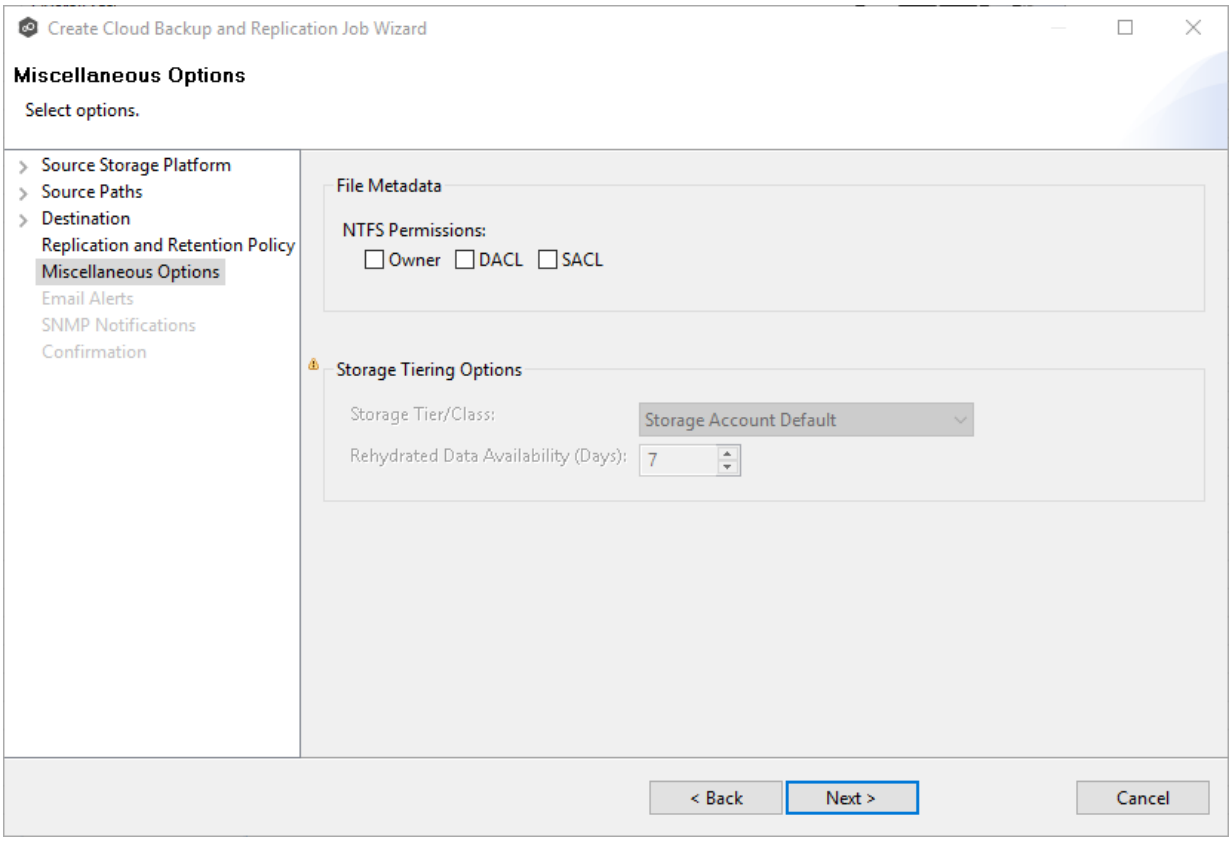

Option	Description
NTFS Permissions	If you want NTFS permissions metadata included in the replication, select the elements to include: • **Owner** – The NTFS Creator-Owner who owns the object (which is, by default, whomever created it). • **DACL** – A Discretionary Access Control List identifies the users and groups that are assigned or denied access permissions on a file or folder. • **SACL** - A System Access Control List enables administrators to log attempts to access a secured file or folder. It is used for auditing. See File Metadata Synchronization for more information about NTFS permissions metadata.

Option	Description
Storage Tier/Class	Only available with Azure Blob Storage with either a General Purpose v2 (GPv2) account or Blob Storage Account. Select a storage tier. If you do not select a tier, it will default to the tier you configured on your Azure Storage account. Azure Storage offers three storage tiers for blob object storage so that you can store your data most cost-effectively depending on how you use it: • **Azure Hot Storage Tier** is optimized for storing data that is accessed frequently. • **Azure Cool Storage Tier** is optimized for storing data that is infrequently accessed and stored for at least 30 days. • **Azure Archive Storage Tier** is optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements (on the order of hours). The archive storage tier is only available at the blob level and not at the storage account level. To read data in archive storage, Cloud Backup and Replication must first change the tier of the blob to hot or cool. This process is known as rehydration and can take up to 15 hours to complete. **Rehydrated data** remains in hot or cool storage for a specified number of days before Cloud Backup and Replication automatically returns it to archive storage.
Rehydrated Data Availability (Days)	Only available with Azure Blob Storage with either a General Purpose v2 (GPv2) account or Blob Storage Account. Rehydrated data is automatically returned to archive storage after a specified period. Enter the number of days for rehydrated data to remain in hot or cool storage before returning to archive storage. The default is seven days.

2. Click **Next**.

   The Email Alerts page appears.

Step 16: Email Alerts

This step is optional.

An email alert notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in Preferences and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply.

To create a new alert, see Email Alerts in the Preferences section.

To apply an existing email alert to the job:

1. Click the **Select** button.

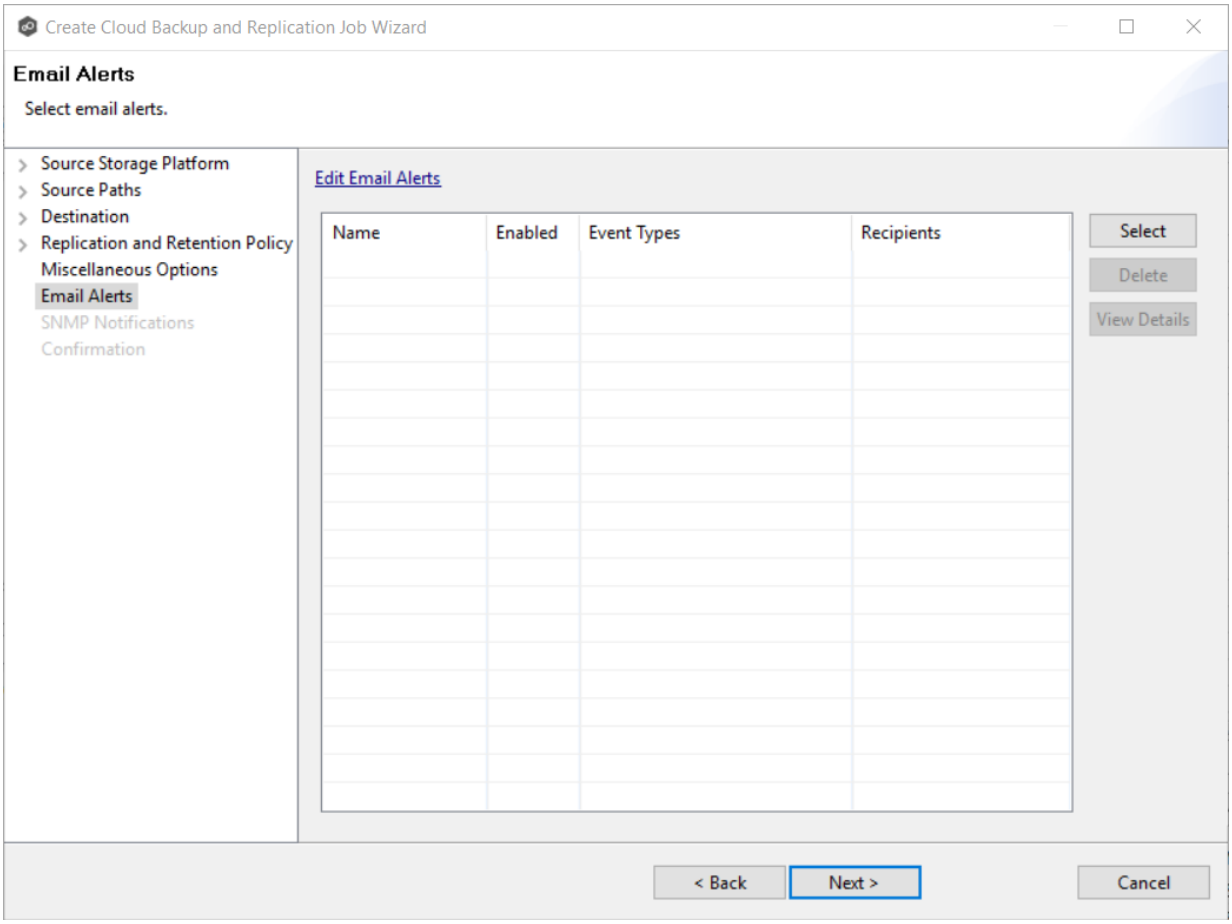

The **Select Email Alert** dialog appears.

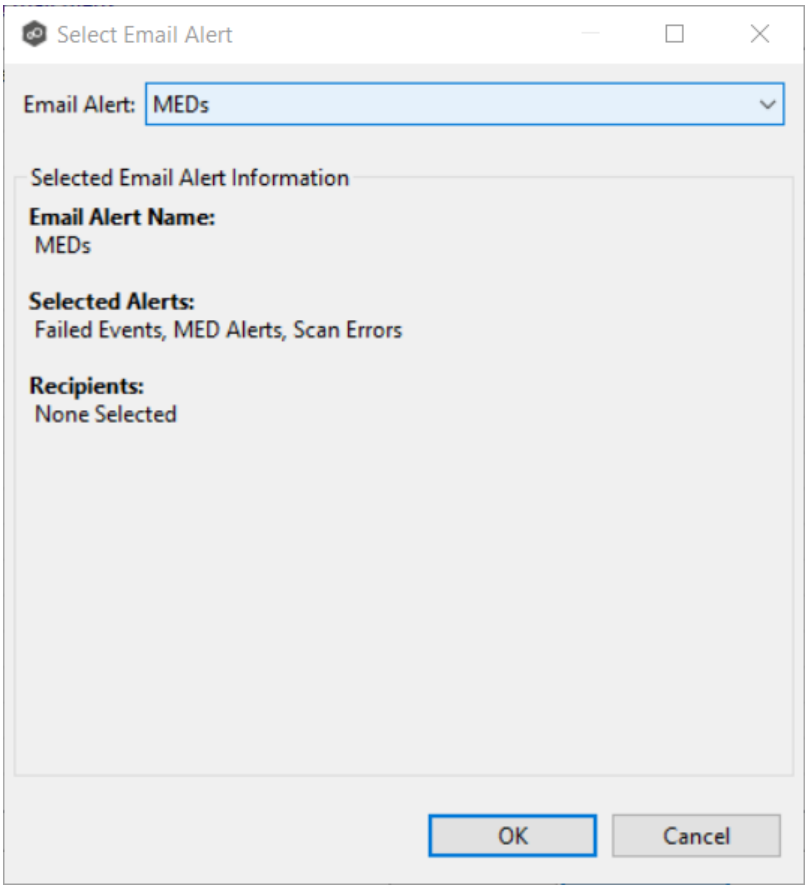

2. Select an alert from the **Email Alert** drop-down list, and then click **OK**.

   The alert is listed in the **Email Alerts** page.

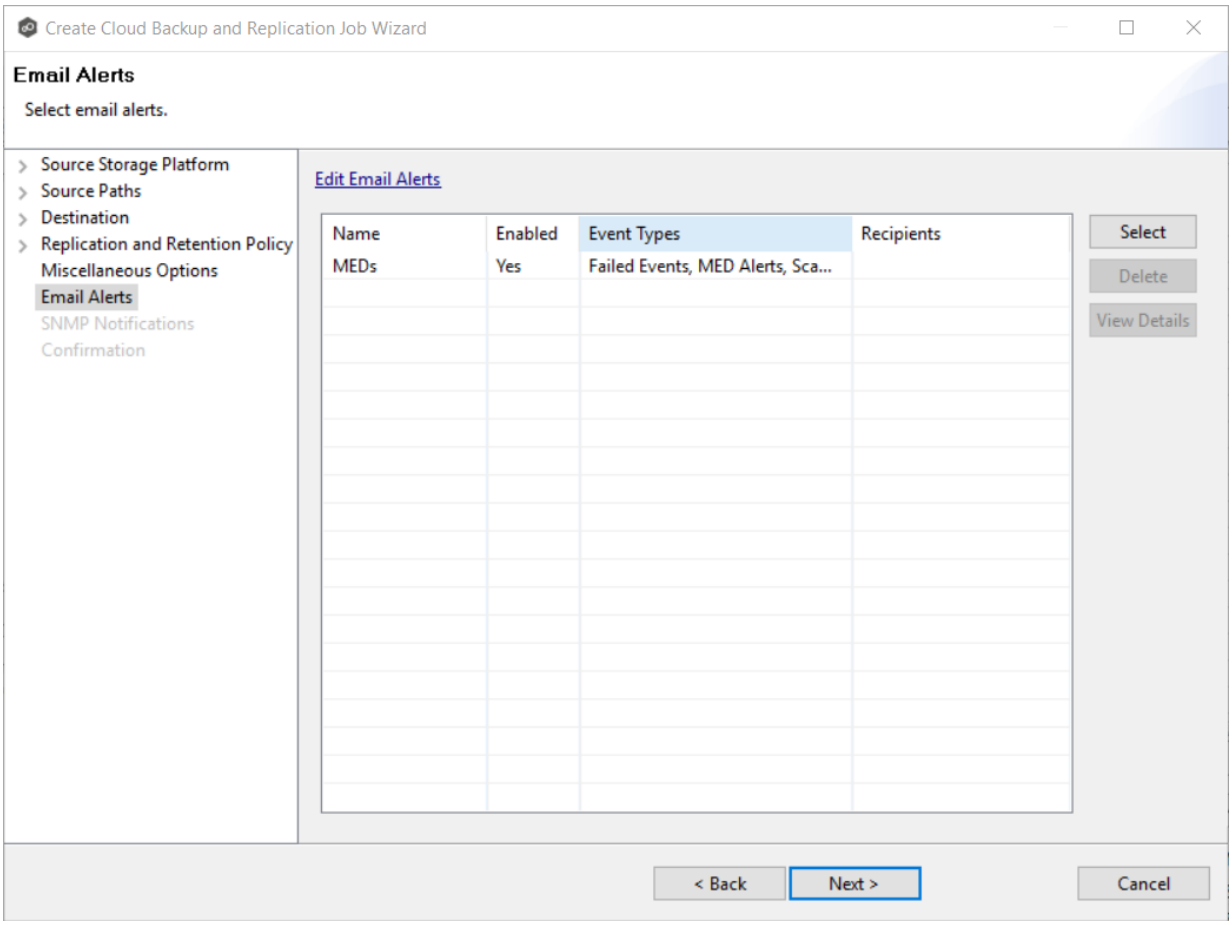

3. (Optional) Repeat steps 1-3 to apply additional alerts.

4. Click **Next**.

   The SNMP Notifications page appears.

Step 17: SNMP Notifications

This step is optional.

An SNMP notification notifies recipients when certain type of event occurs, for example, session abort, host failure, system alert. The **SNMP Notifications** page displays a list of notifications that have been applied to the job. When you first create a job, this list is empty. Like email alerts and file filters, an SNMP notification is defined in Preferences and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create SNMP notifications in advance. However, from this wizard page, you can select an existing SNMP notification to apply to the job or create new SNMP notifications.

To apply an existing SNMP notification to the job or disable notifications:

1. Select an SNMP notification from the drop-down list.

   To disable, select **None - Disabled**.

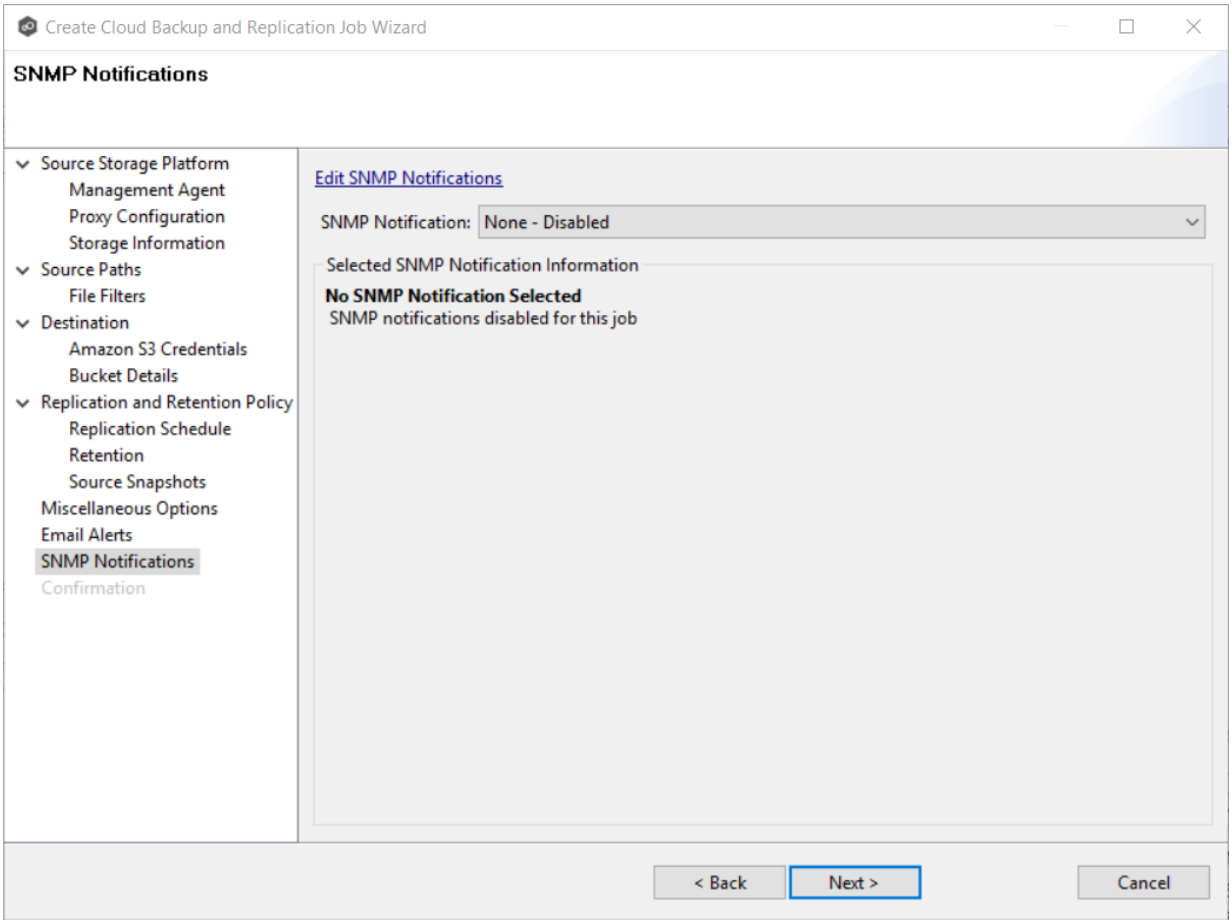

2. Click **Next**.

   The Confirmation page appears.

Step 18: Confirmation

The **Confirmation** page displays the job configuration.

1. Review the job configuration.

2. If you need to modify the job configuration after reviewing it, click **Back** until you reach the appropriate page and make your changes.

   Note: You cannot change the job name.

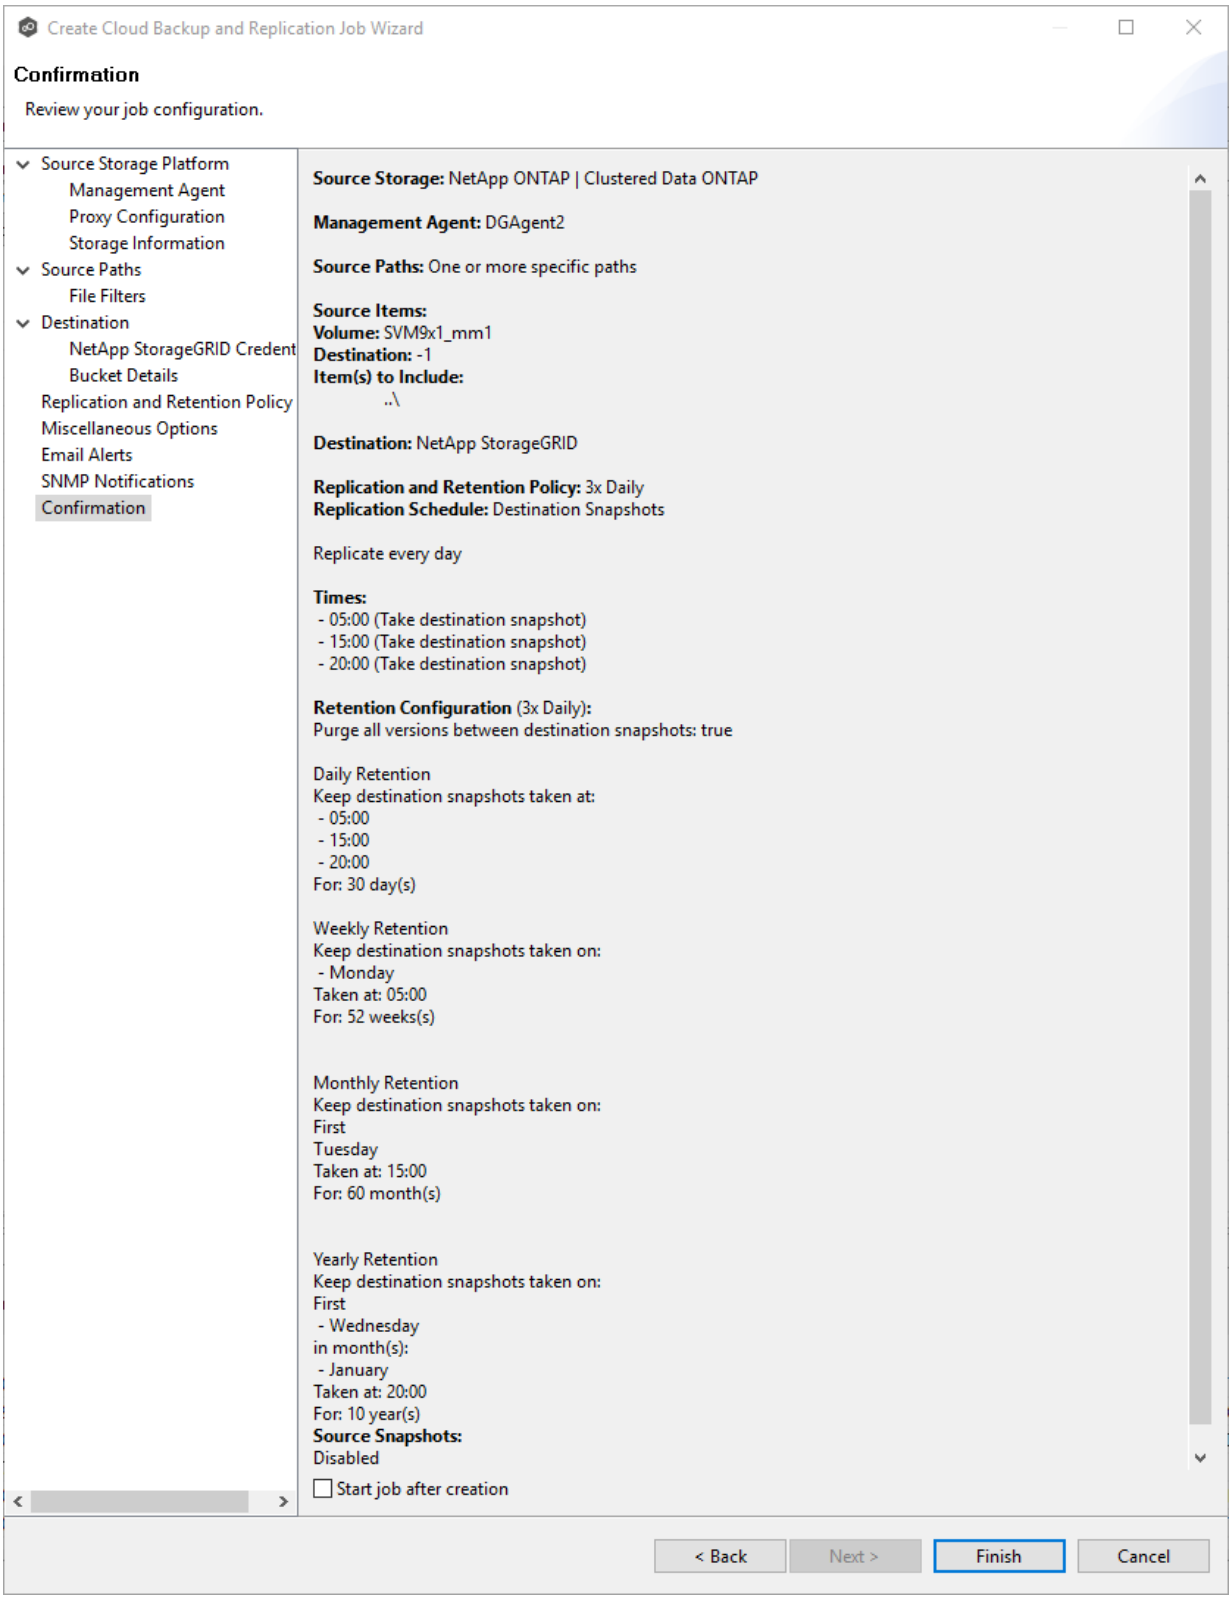

3. Select the **Start job after creation** checkbox if you want the job to start immediately after clicking **Finish**.

4. Click **Finish**.

   The **Summary** tab in the **Cloud Backup and Replication Job** runtime view is displayed.

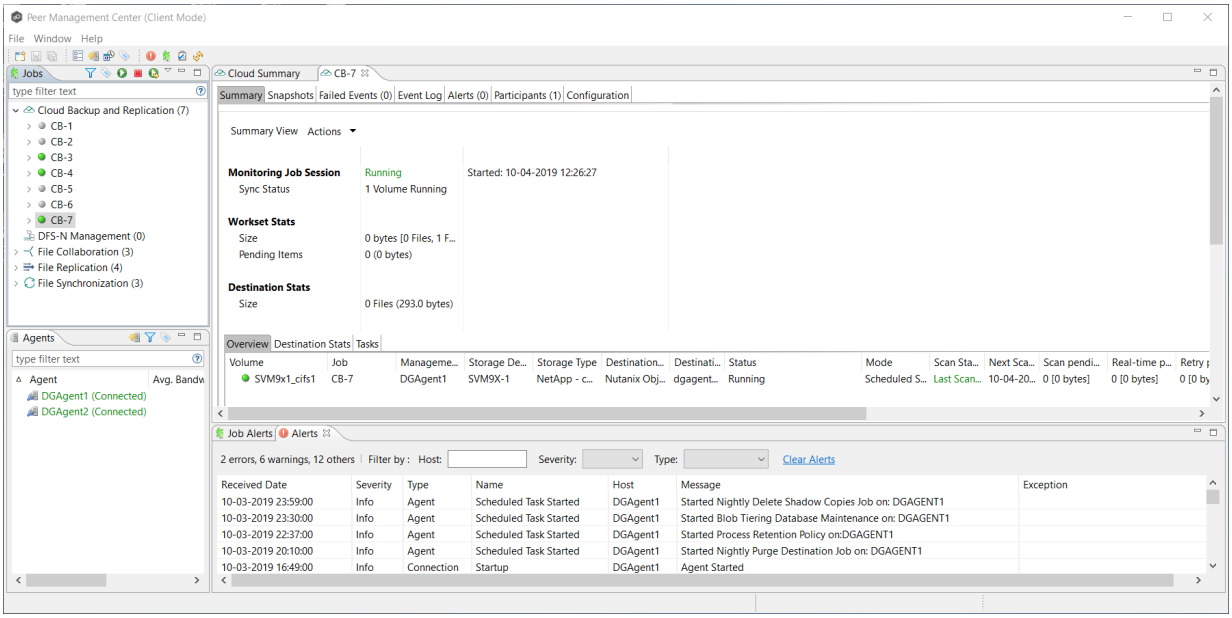

Running a Cloud Backup and Replication Job

This section describes:

- Starting a Cloud Backup and Replication Job
- Stopping a Cloud Backup and Replication Job

Starting a Cloud Backup and Replication Job

When running a Cloud Backup and Replication job for the first time, you must manually start it. After the initial run, a job will automatically start, even when the Peer Management Center server is rebooted.

Note: You cannot run two jobs concurrently on the same volume if the watch sets contain an overlapping set of files and folders.

To manually start a job:

1. Choose one of these options:

   - Right-click the job name in the **Jobs** view.

   - Open a job and then click the **Start/Stop** button to the left of the **Status** field in the bottom left corner of the job's view (shown below).

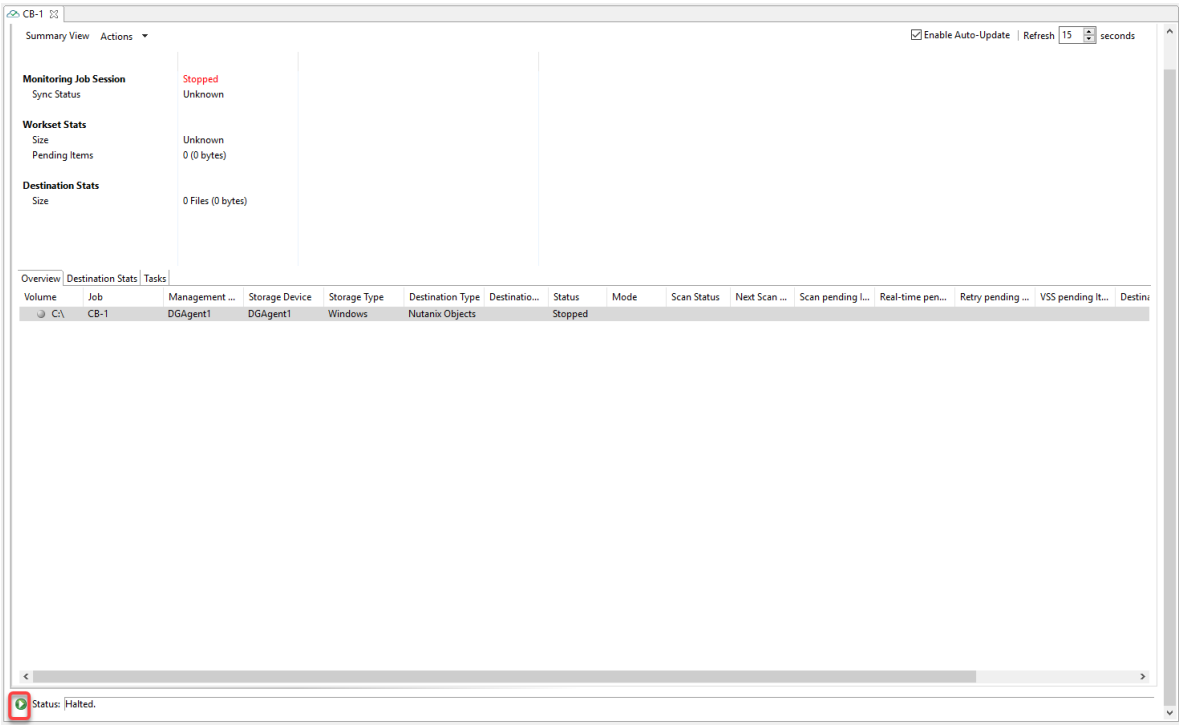

2. Click **Yes** in the confirmation dialog.

   After the job initialization has completed, the job will run. Once the job starts, the icon next to the job name in the **Jobs** view changes from gray to green.

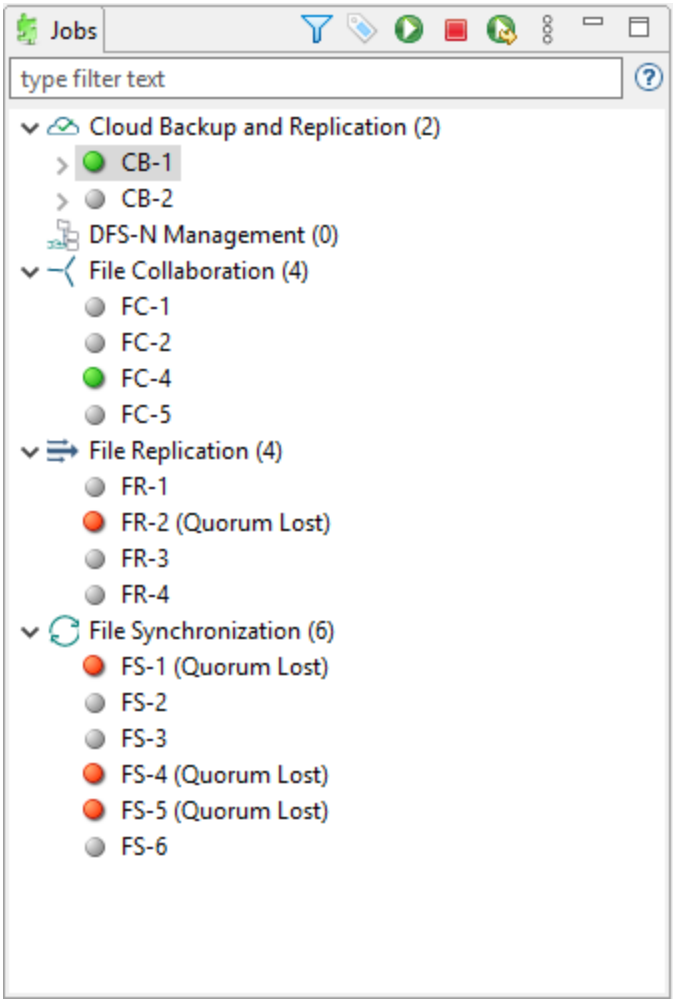

Stopping a Cloud Backup and Replication Job

You can stop a Cloud Backup and Replication job at any time.

To stop a Cloud Backup and Replication job:

1. Right-click the job name in the **Jobs** view or in the **Cloud Backup and Replication Job Summary** view, and then choose **Stop** from the pop-up menu.

   Or open the job and then click the **Start/Stop** button to the left of the **Status** field in the bottom left corner of the job's runtime view (shown below)

2. Click **Yes** in the confirmation dialog.

   The icon next to the job name in the **Jobs** view changes from green to red.

Monitoring Cloud Backup and Replication Jobs

Monitoring your Cloud Backup and Replication jobs is an important aspect of successfully replicating to the cloud. Monitoring involves checking the execution of a running job, checking the status of a job, reviewing performance statistics, making sure snapshots are created correctly, identifying problems such as a server outage, seeing how much data has been uploaded, and so forth. Cloud Backup and Replication provides several views to help you monitor the health and performance of your Cloud Backup and Replication jobs.

Many of the views are customizable tables. You can sort the columns in the view, filter by columns, add and subtract columns from the default display, and so forth.

To display a view:

- Double-click **Cloud Backup and Replication** in the **Jobs** view to display the summary view for all Cloud Backup and Replication jobs. The **Volume Summary** tab of the **Cloud Summary** view is displayed.

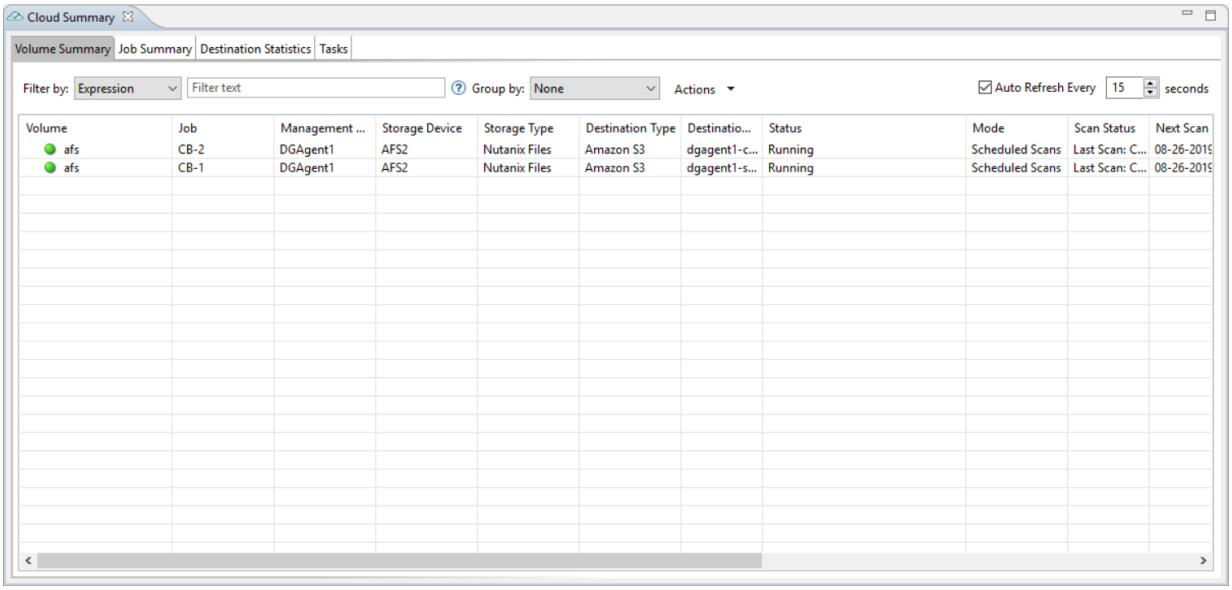

- Double-click the name of a Cloud Backup and Replication job in the **Jobs** view to display the runtime view associated with that job. The **Summary** tab of the runtime view is displayed.

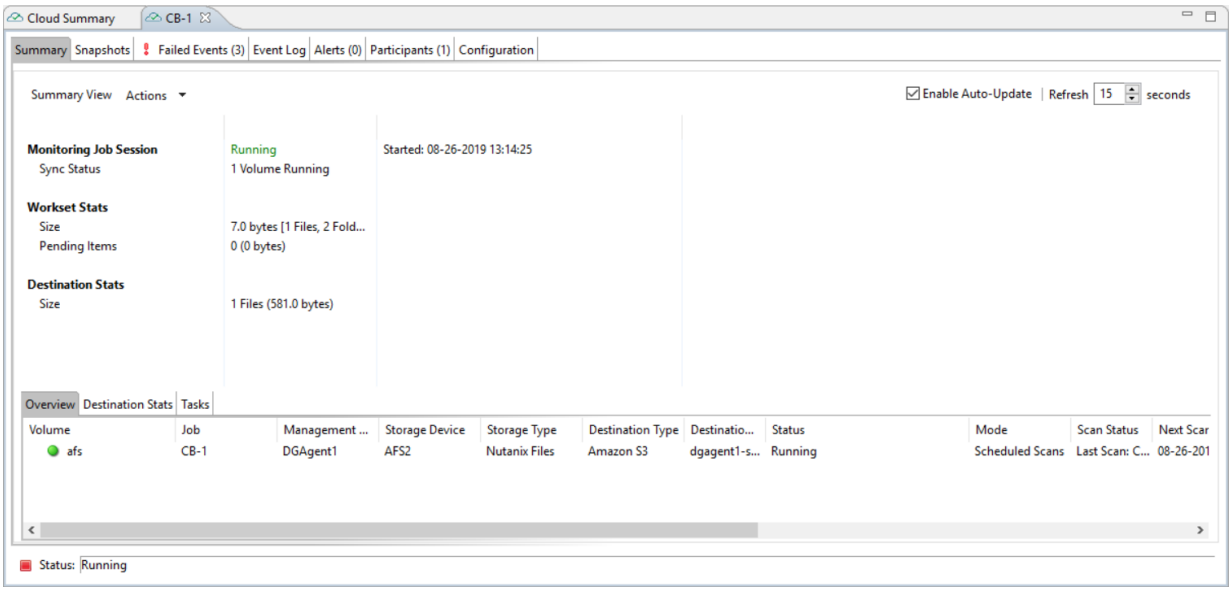

Deleting a Cloud Backup and Replication Job

To delete a Cloud Backup and Replication job:

1. Right-click on the job name in the **Jobs** view, and then choose **Delete** from the menu. A confirmation dialog appears.

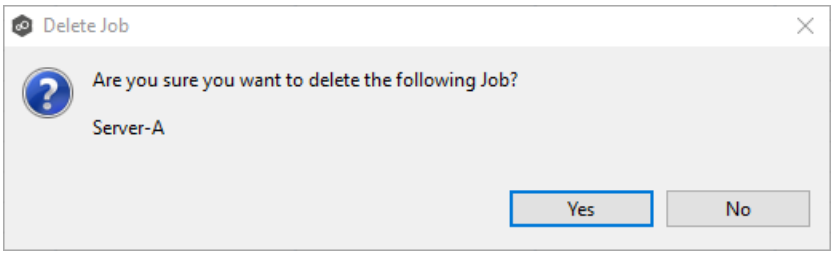

2. Click **OK** in the confirmation dialog.

   Another dialog appears, prompting you to choose whether to delete data associated with the job.

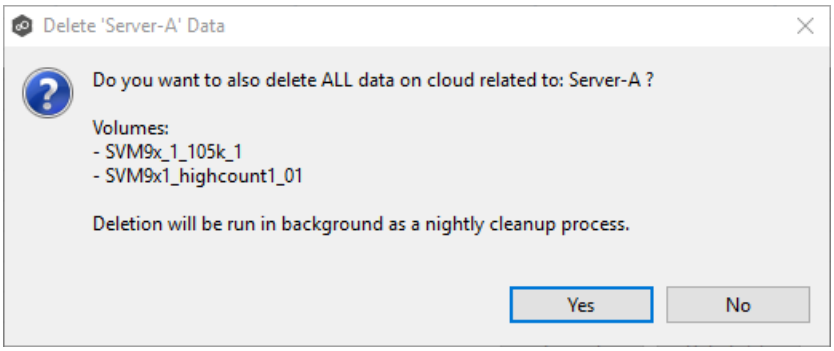

3. Click **Yes** or **No**.

   If you click **Yes**, the data associated with this job will be deleted as part of a nightly clean-up process in addition to the job itself. If you click **No**, the data will not be deleted but the job will be deleted.

Recovering Data

When you need to recover data from the cloud to on-premises, you can use the **Data Recovery** wizard. To restore data, you must have an existing Cloud Backup and Replication job that has been replicating that data.

Note: You can recover data from a running job. However, if you plan to restore the data to the original location, you should stop the job first.

To recover data:

1. Open Peer Management Center.

2. In the **Jobs** view, identify the Cloud Backup and Replication job that replicated the data you want to restore.

3. Right-click the job name, and then select **Recover Volume/File(s)** from the menu.

   The **Recovery Wizard** opens and displays the **Volume to Recover** page. The **Storage Device** field on the page is a read-only field that displays the name of the source storage device.

4. Select the volume that was the source of the replicated data from the **Volume** drop-down list.

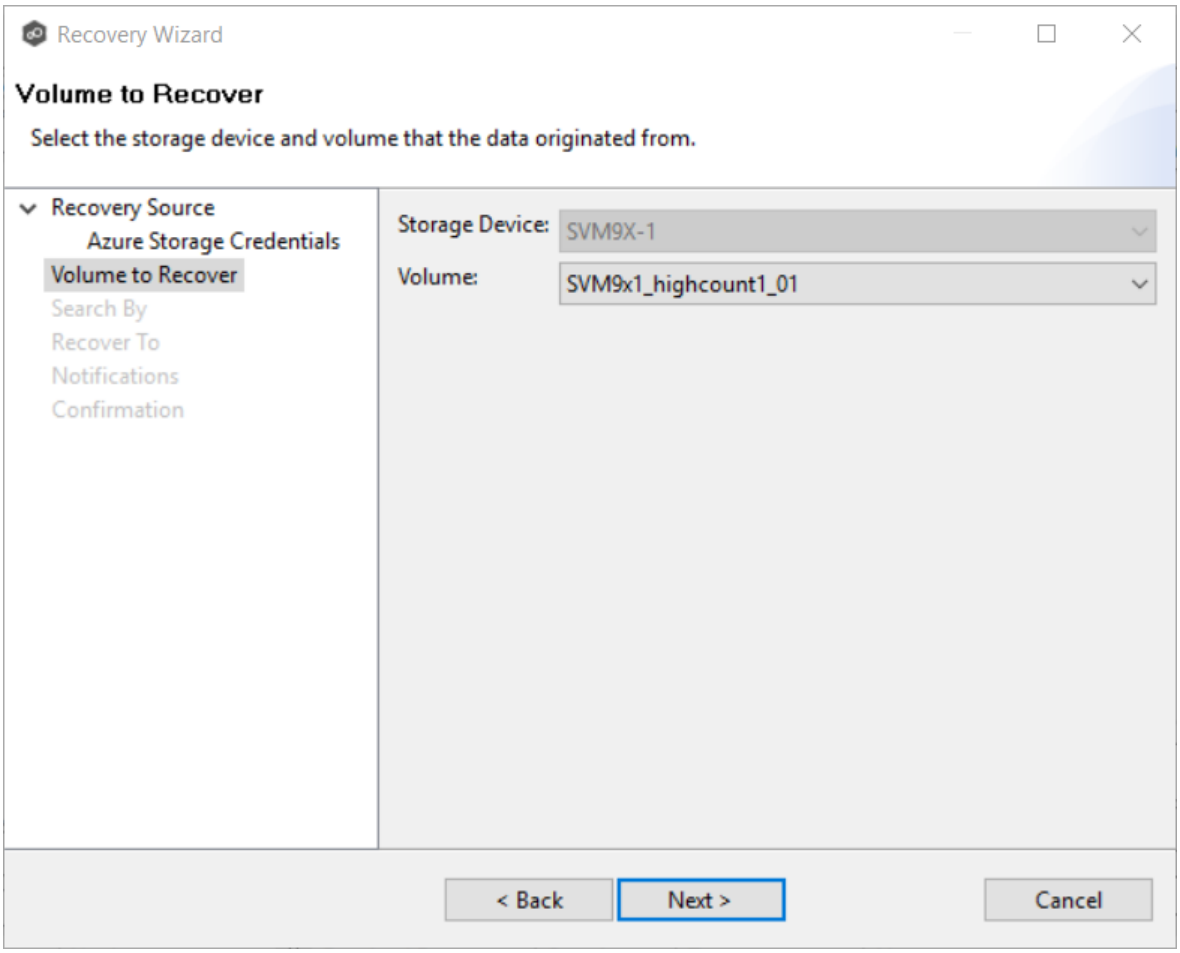

5. Click **Next**.

   The **Search By** page is displayed.

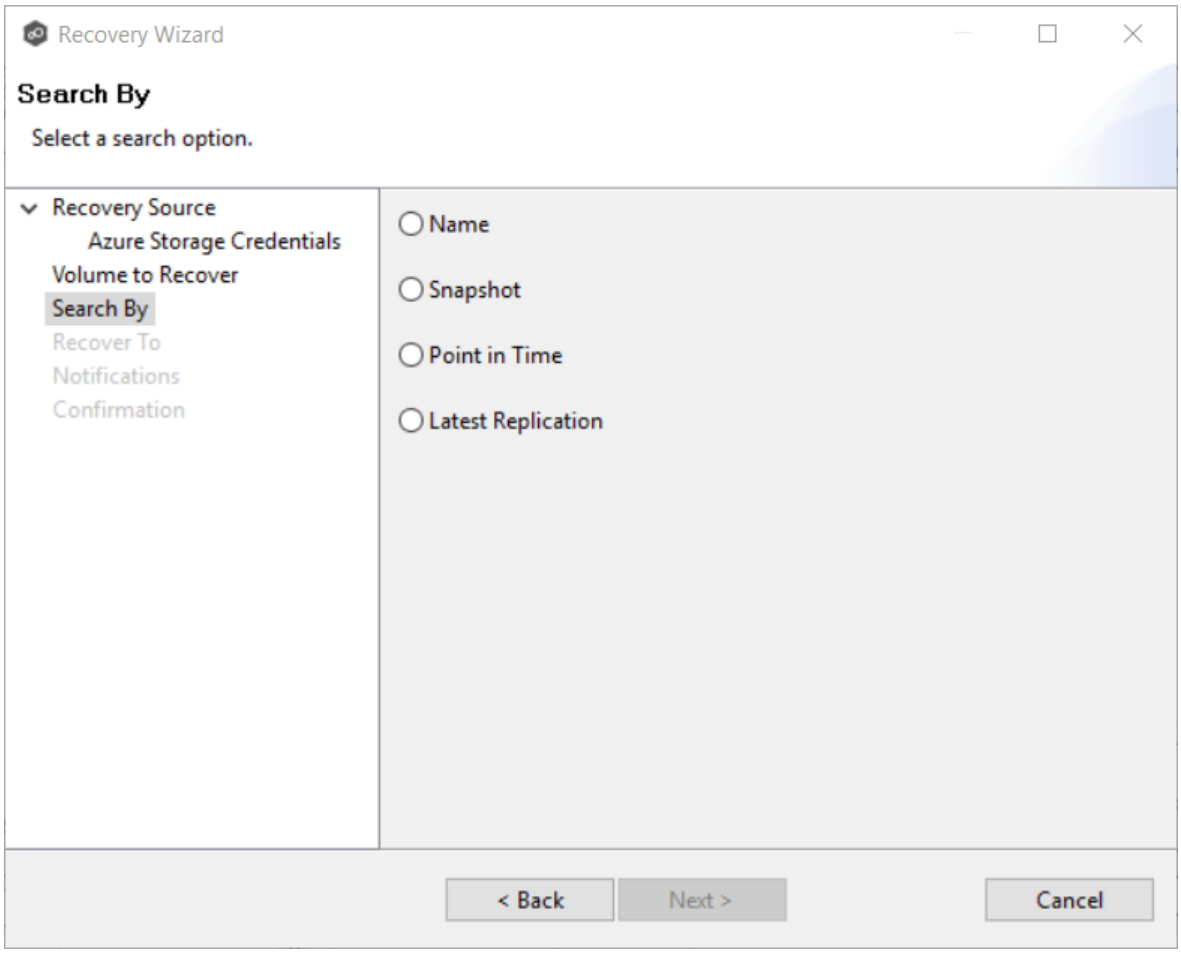

6. Select one of the search options.
   - Name
   - Snapshot
   - Point in Time
   - Latest Replication
7. Click **Next** and continue with Recovery Options.

   The search pages vary according to the search option you selected.

Search Options

1. The search options are:

   - Name
   - Snapshot
   - Point in Time
   - Latest Replication

Use the **Search by Name** option if you know any part of the name of a file or folder but don't know which folder contained it on the original volume on premises.

To search by name:

1. Enter a search string in the **Name** field.

   The search string can be a full or partial name and can include wildcards. If you do not enter a search string, all files and folders will be listed in the search results.

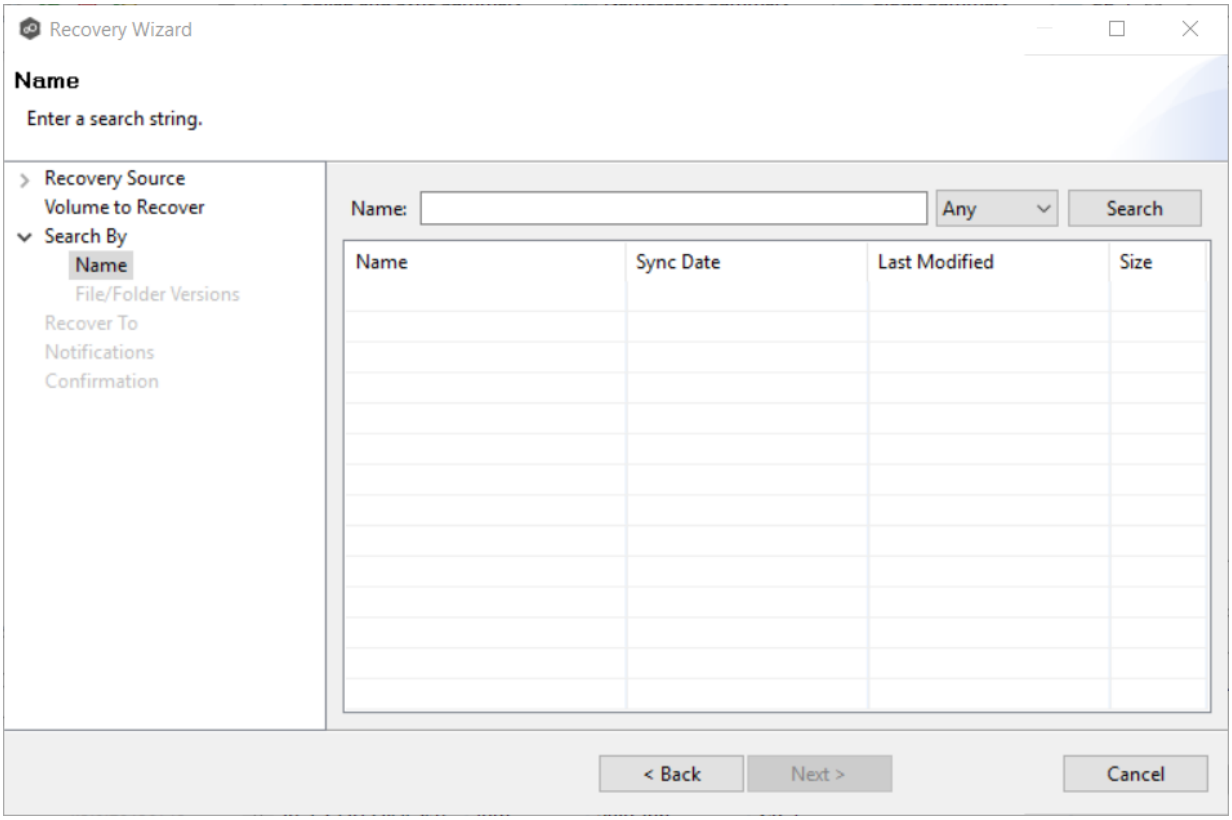

2. Select **File** or **Folder** from the **Any** drop-down list; if you want to search for both files and folders, select **Any**.

3. Click **Search**.

   A list of matching files and/or folders appears. The **Sync Date** column shows the date the file was replicated; the **Last Modified Date** column shows the last known date and time that the file was changed on premises.

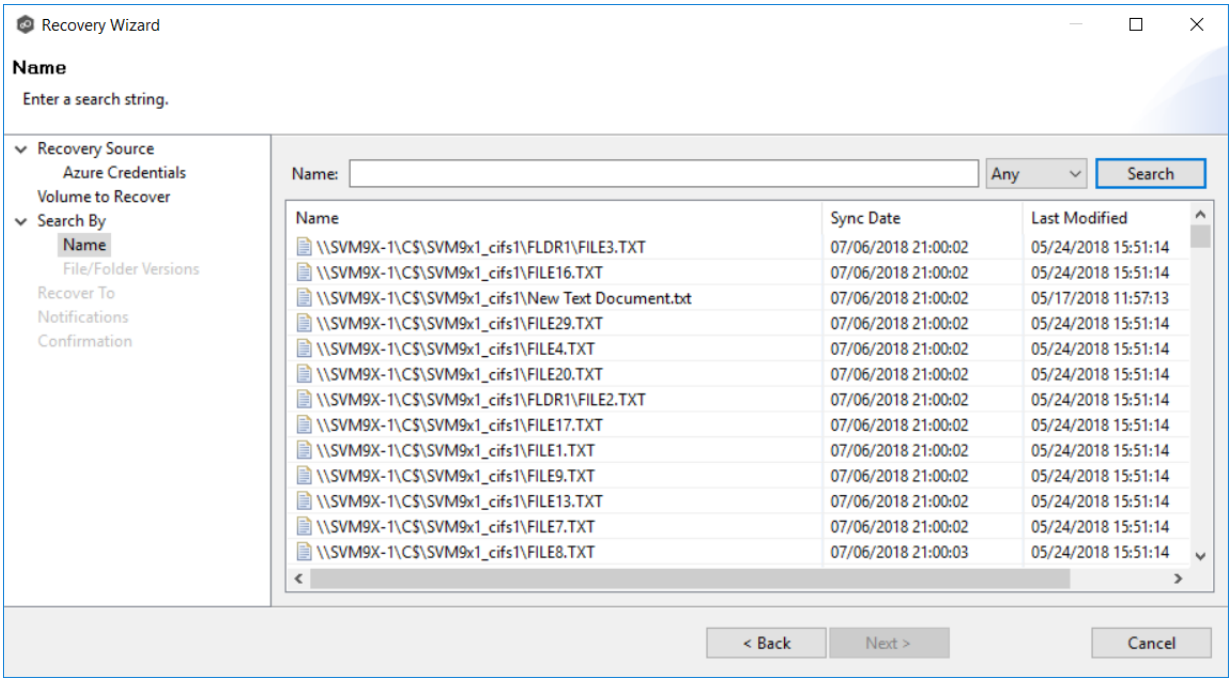

4. Select the file or folder to recover.

5. Click **Next**.

   The **File/Folder Versions** page appears. Your options will vary, depending on whether you are recovering a file or folder.

6. If you selected a file to recover, all available versions of that file are presented below the calendar. Select the time of the desired version and then click elsewhere in the page.

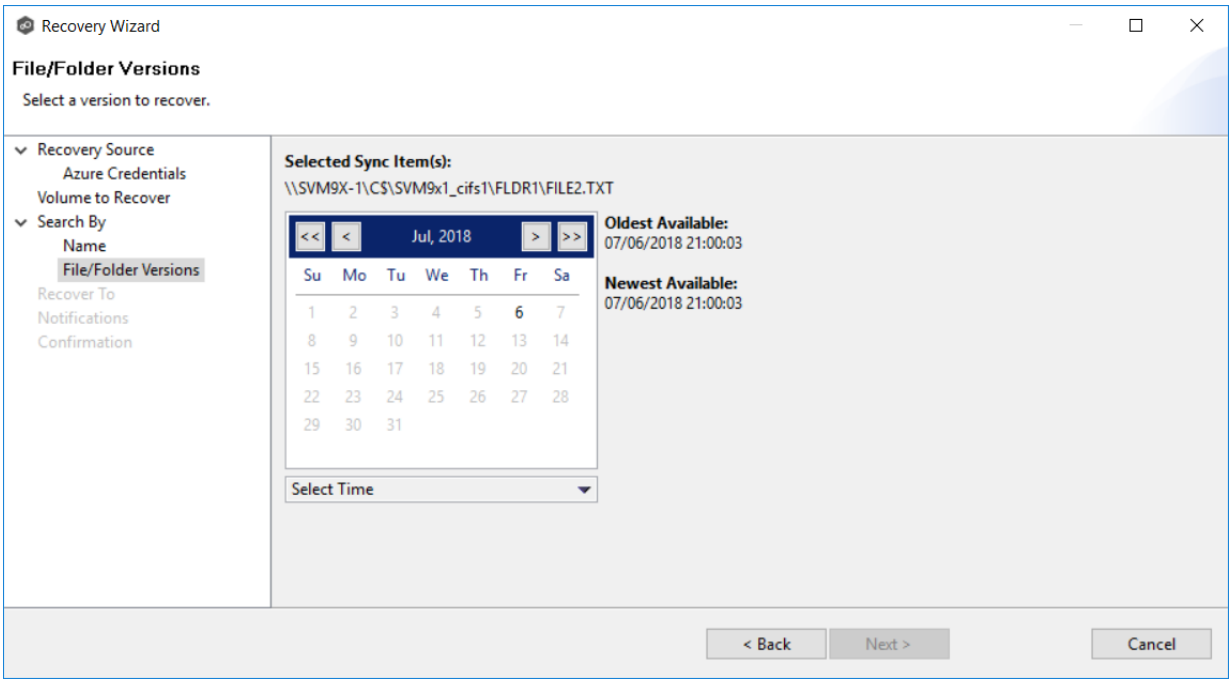

If you selected a folder to recover, you have two options. You can recover the contents of the folder based on a snapshot that was previously taken, or you can recover the contents of the folder as it existed at a specific point in time. Select one of the options, select a time, and then click elsewhere in the page.

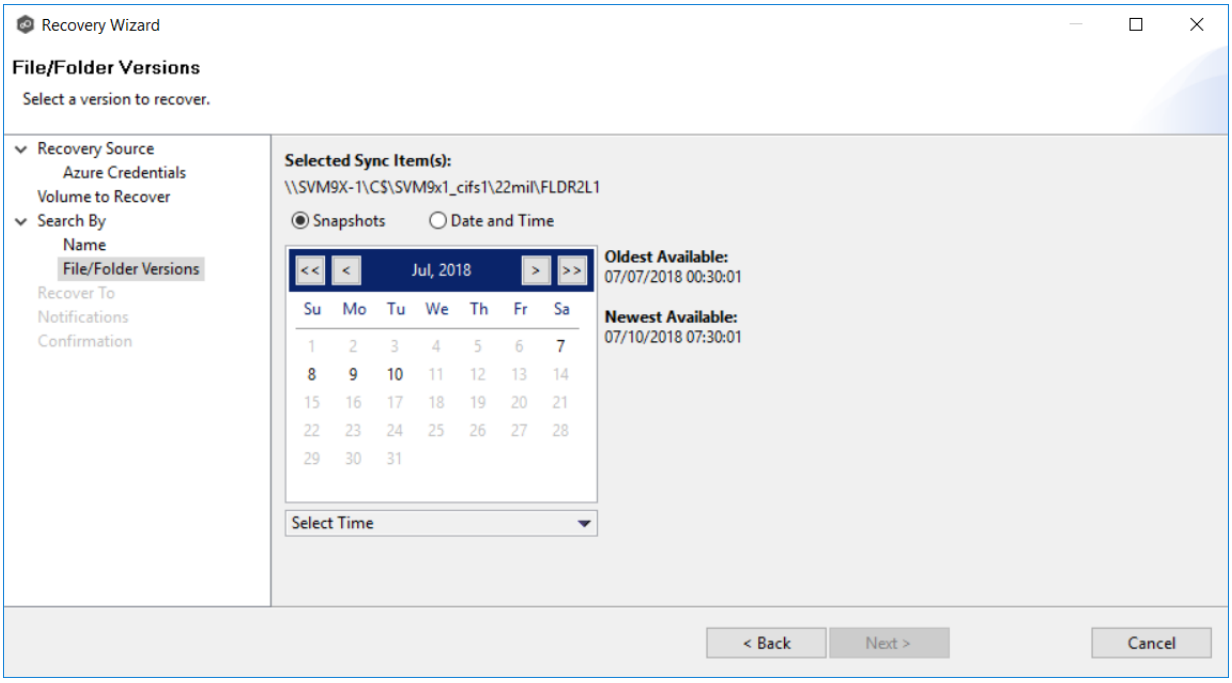

7. Click **Next** and continue with Recovery Options.

Use the **Search by Snapshot** option if you want to recover data by browsing a previously taken destination snapshot. All available snapshots will be represented in the calendar widget below.

To search by snapshot:

1. Select the date of the snapshot.

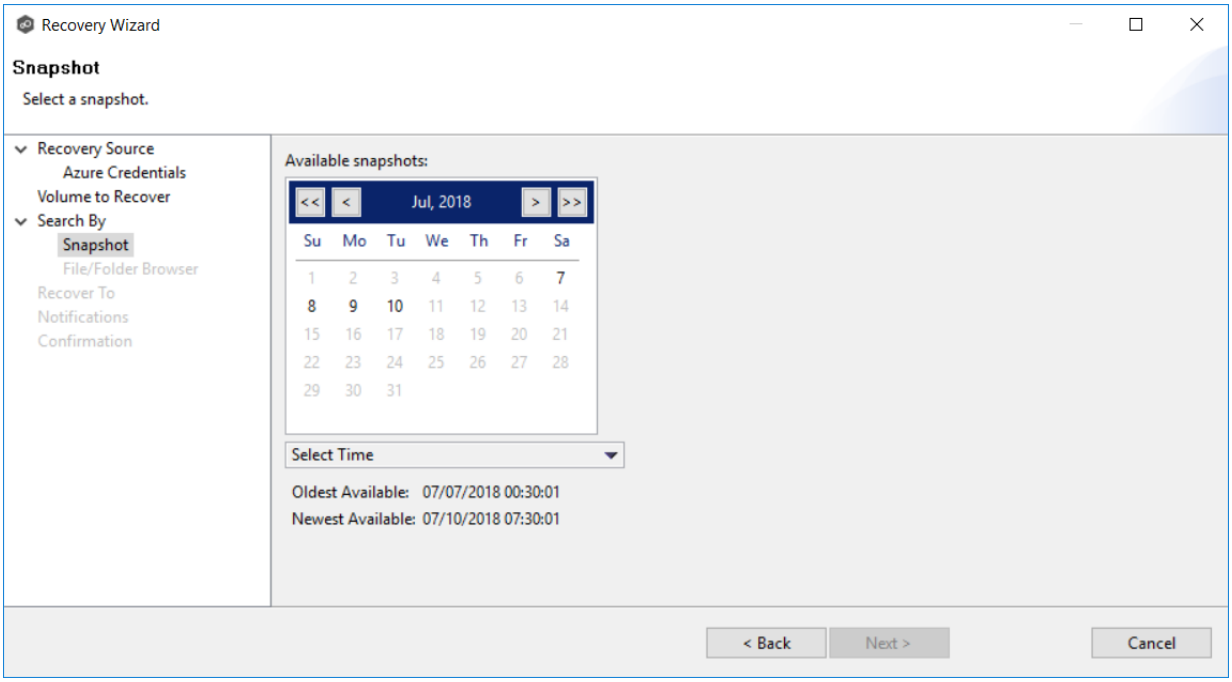

2. Select the time of the snapshot, and then click elsewhere in the page.

3. Click **Next**.

   The **File/Folder Browser** page appears.

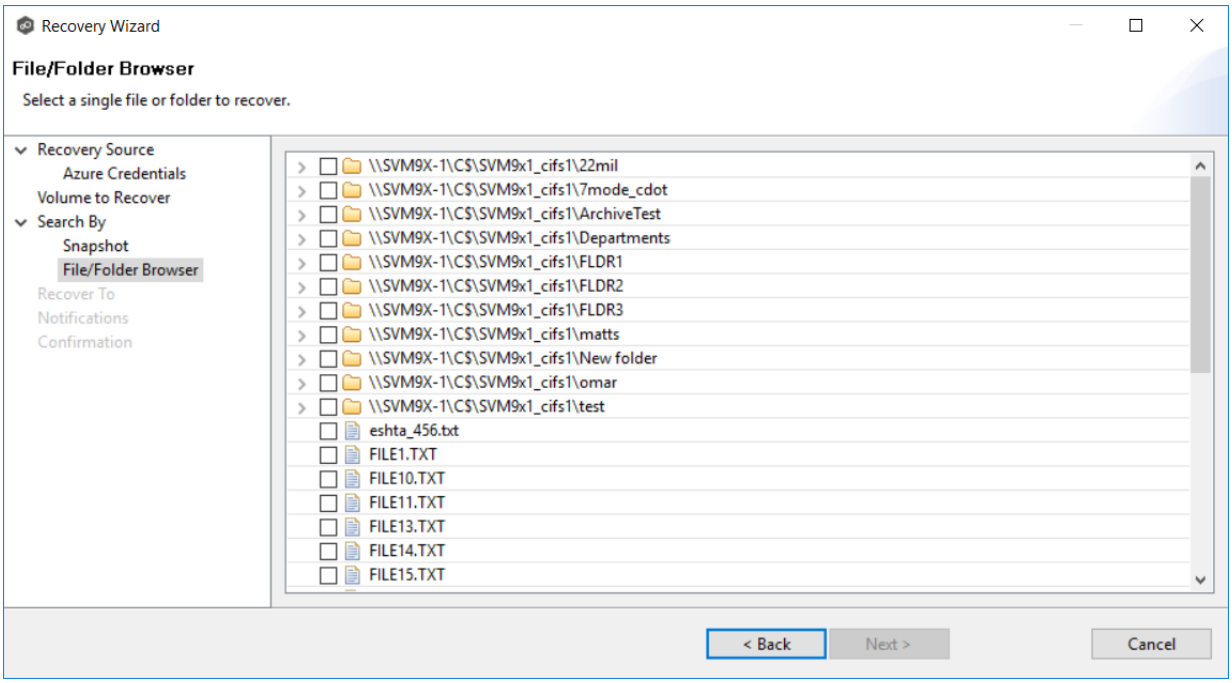

4. Select the file or folder to restore. If no snapshots are available, click **Back** and select a different search option.

5. Click **Next** and continue with Recovery Options.

Use the **Search by Point in Time** option if you want to restore a data from a specific point in time. This option does not require that a snapshot was taken and is very useful if you selected Continuous Data Protection, where replication is performed on an on-going basis

To search by a point in time:

1. Select a date.

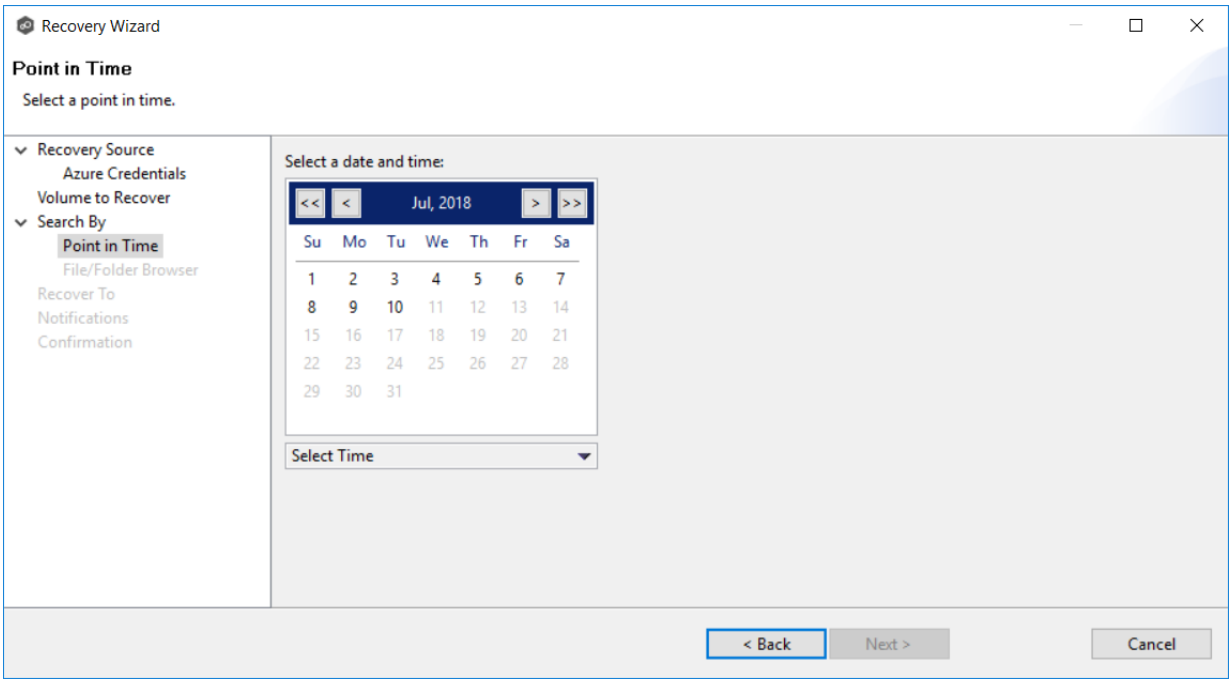

2. Select a date and time, and then click elsewhere in the page.
3. Click **Next**.

   The **File/Folder Browser** page appears.

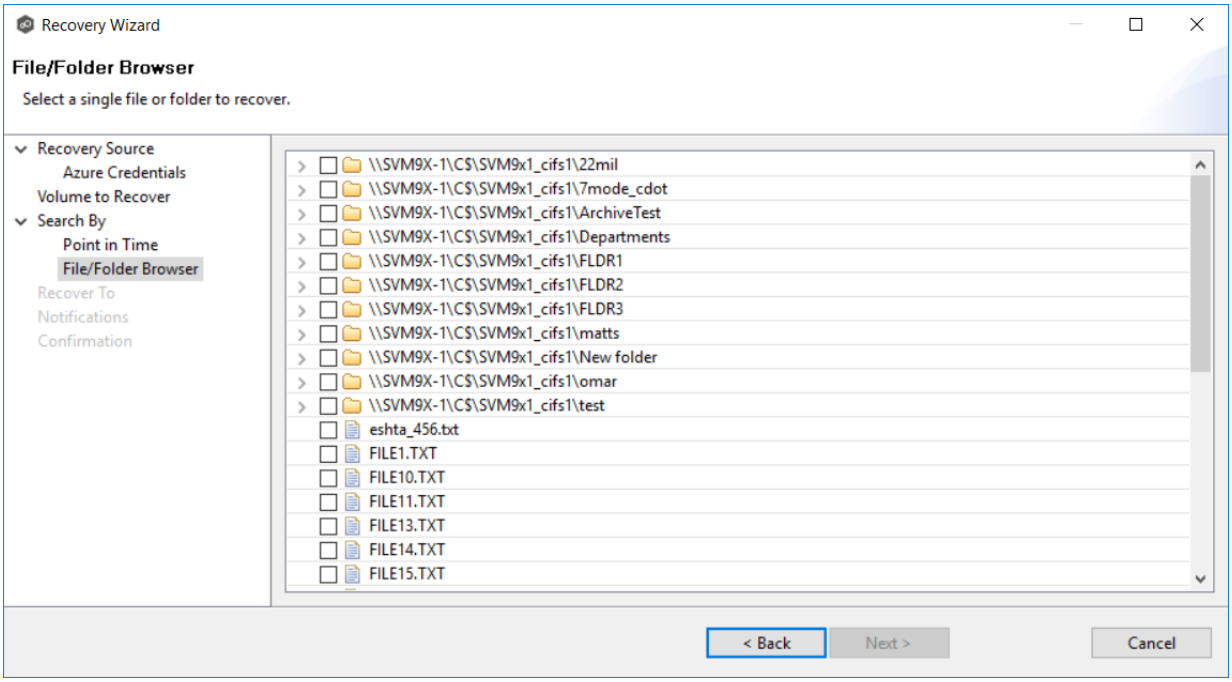

4. Select the file or folder to restore.

5. Click **Next** and continue with Recovery Options.

Use the **Search by Latest Replication** option if you want to restore from the latest replication. For example, you may want to restore data from the last time that replication occurred rather than a snapshot or a point in time. This option is very useful if you selected Continuous Data Protection, where replication is performed on an on-going basis.

To search by latest replication:

1. Select the file or folder to restore.

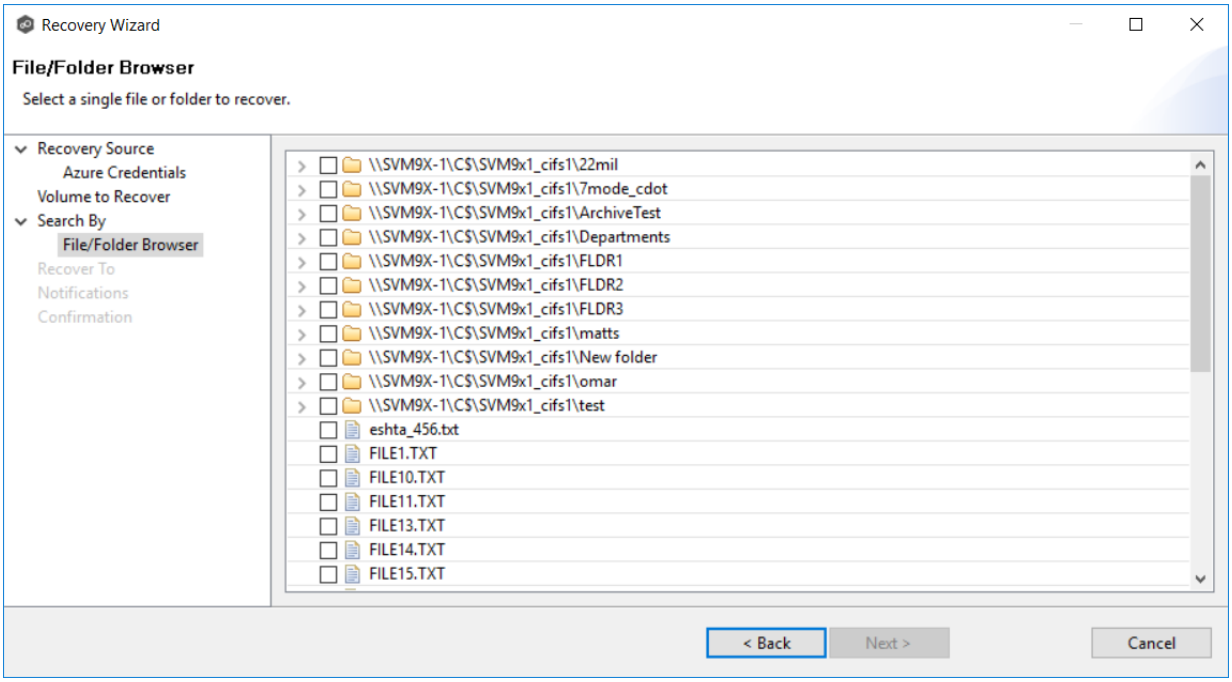

2. Click **Next** and continue with Recovery Options.

Recovery Options

After you select the data to recover, the **Recover To** page appears.

1. Select the recovery location. You have two options:

- **Another Location** - Enter the UNC path to a location on another storage device.

- **Original Location** - Browse to a location on the device hosting the management agent. However, we recommend not restoring directly to the original location, especially if the job is currently running. If the version that is restored is older than the latest version in the destination storage, the restored version will not be backed up until the next scan.

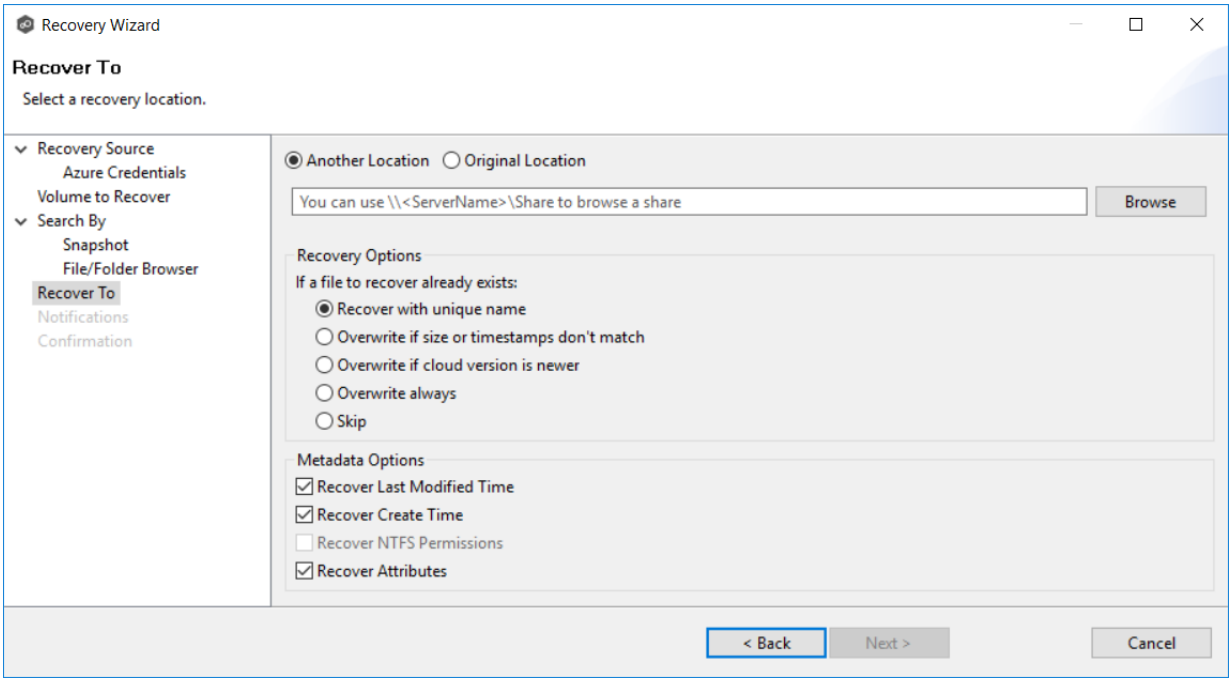

2. Select the recovery options for when the file to recover already exists in the recovery location:

Recovery Option	Select this option if you want to:
Recover with unique name	Ensure that the existing file is not overwritten with the cloud version.
Overwrite if sizes or timestamps don't match	Overwrite the existing file with the cloud version if the sizes or timestamps the existing file do not match the cloud version.
Overwrite if cloud version is newer	Overwrite the existing file if the cloud version has a more recent modification date.

Overwrite always	Always overwrite the existing file with the cloud version.
Skip	Skip recovering a file if the file already exists.

3. Select the recovery metadata options:

Metadata Option	Select this option if you want to:
Recover Last Modified Time	Set the last modification time of a recovered file to match the last modification time stored at upload rather than the time at which it was recovered.
Recover Create Time	Set the creation time of a recovered file to match the creation time stored at upload rather than the time at which it was recovered.
Recover NTFS Permissions	Set the NTFS permissions of any recovered files and folders to match the original permissions when those files and folders were uploaded.
Recover	Set the attributes of any recovered files and folders to match the original attributes when those files and folders were uploaded.

4. (Optional) Click the **Review** button to see your selections.

5. Click **Next**.

   The **Notifications** page appears.

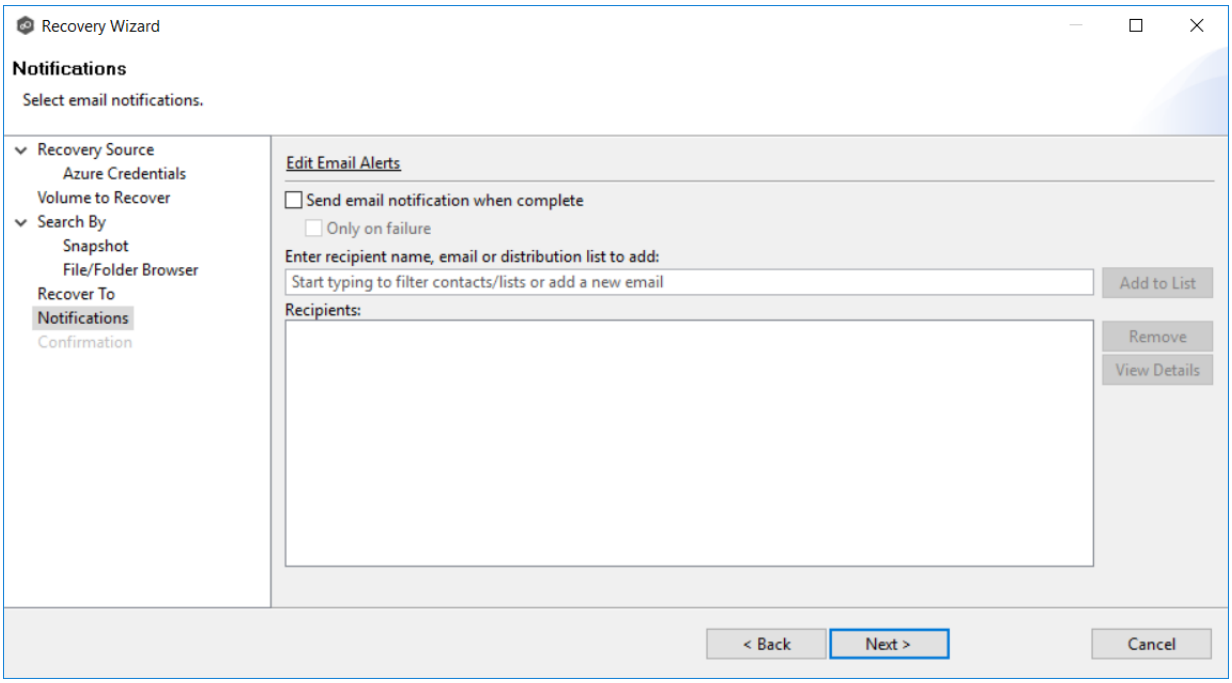

6. (Optional) Select the **Send email notification when complete** checkbox if you want notifications sent when the recovery process is complete. Select **Only on failure** if you want notifications sent only if the recovery does not successfully complete.

7. If sending notifications, enter recipients and add them to the list.

8. Click **Next**.

   The **Confirmation** page is displayed.

9. Review your recovery settings.

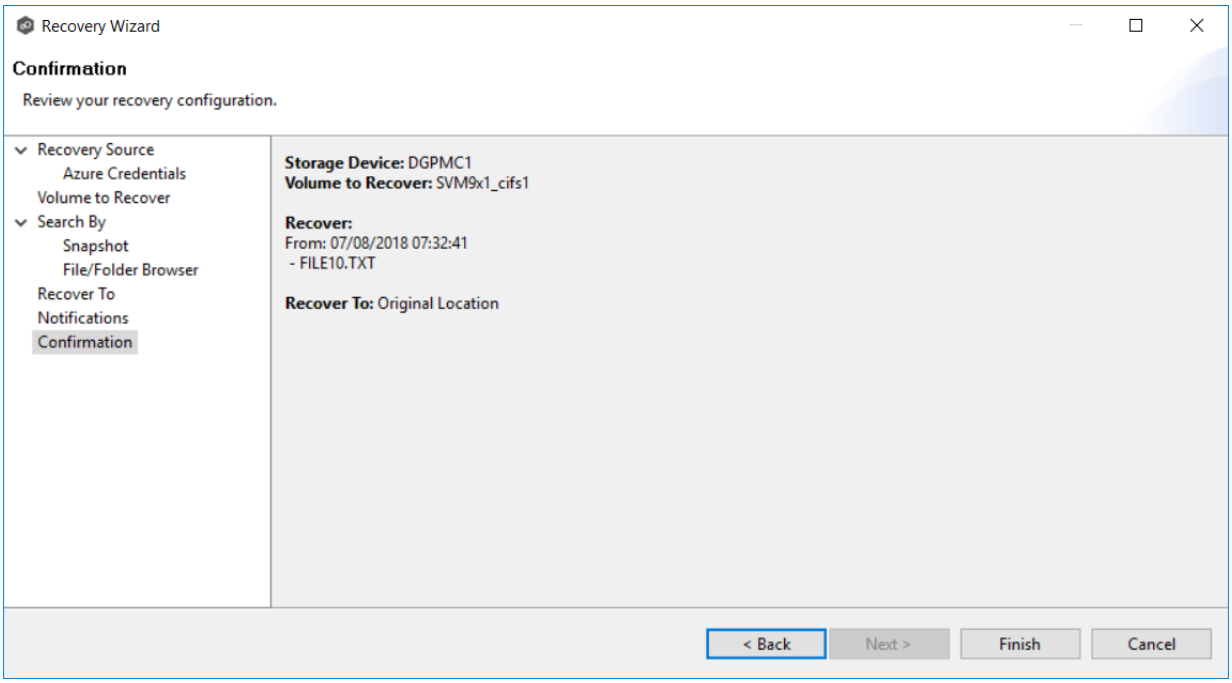

10. Click **Finish**.

DFS-N Management Jobs

A DFS namespace enables you to group shared folders that are located on different servers into one or more logically structured namespaces. Each namespace appears to users as a single shared folder with a series of subfolders. However, the underlying structure of the namespace can consist of numerous file shares that are located on different servers and in multiple sites. See DFS Namespaces for more information about the benefits of using DFS namespaces in PeerGFS.

PeerGFS enables you to create a namespace and manage various activities related to it, such as creating namespace folders, adding folder targets, and linking the namespace to a File Collaboration or File Synchronization job. You could manage DFS namespace using Microsoft tools; however, you can manage DFS namespaces through a dedicated job type in Peer Management Center, the DFS-N Management job.

This section provides information about creating, editing, running, and managing a DFS-N Management job:

- Creating a DFS-N Management Job

- Running a DFS-N Management Job
- Managing DFS Namespaces
  - Adding an Existing Namespace
  - Adding a Namespace Server
  - Adding a Namespace Folder
  - Adding a Namespace Folder Target
- Linking a DFS Namespace to File Collaboration and File Synchronization Jobs

Creating a DFS-N Management Job

The **Create Job** Wizard walks you through the process of creating a DFS-N Management job. The process consists of the following steps:

Step 1: Job Type

Step 2: Management Agent

Step 3: Agent Verification

Step 4: Namespace Name

Step 5: Namespace Servers

Step 6: Namespace Settings

Step 7: Namespace Folders

Step 8: Email Alerts

Step 9: SNMP Notifications

Step 10: Review

Step 11: Results

Step 1: Job Type

1. Open Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **Create New Job** wizard displays a list of job types you can create.
3. Click **DFS-N Management**, and then click **Create**.

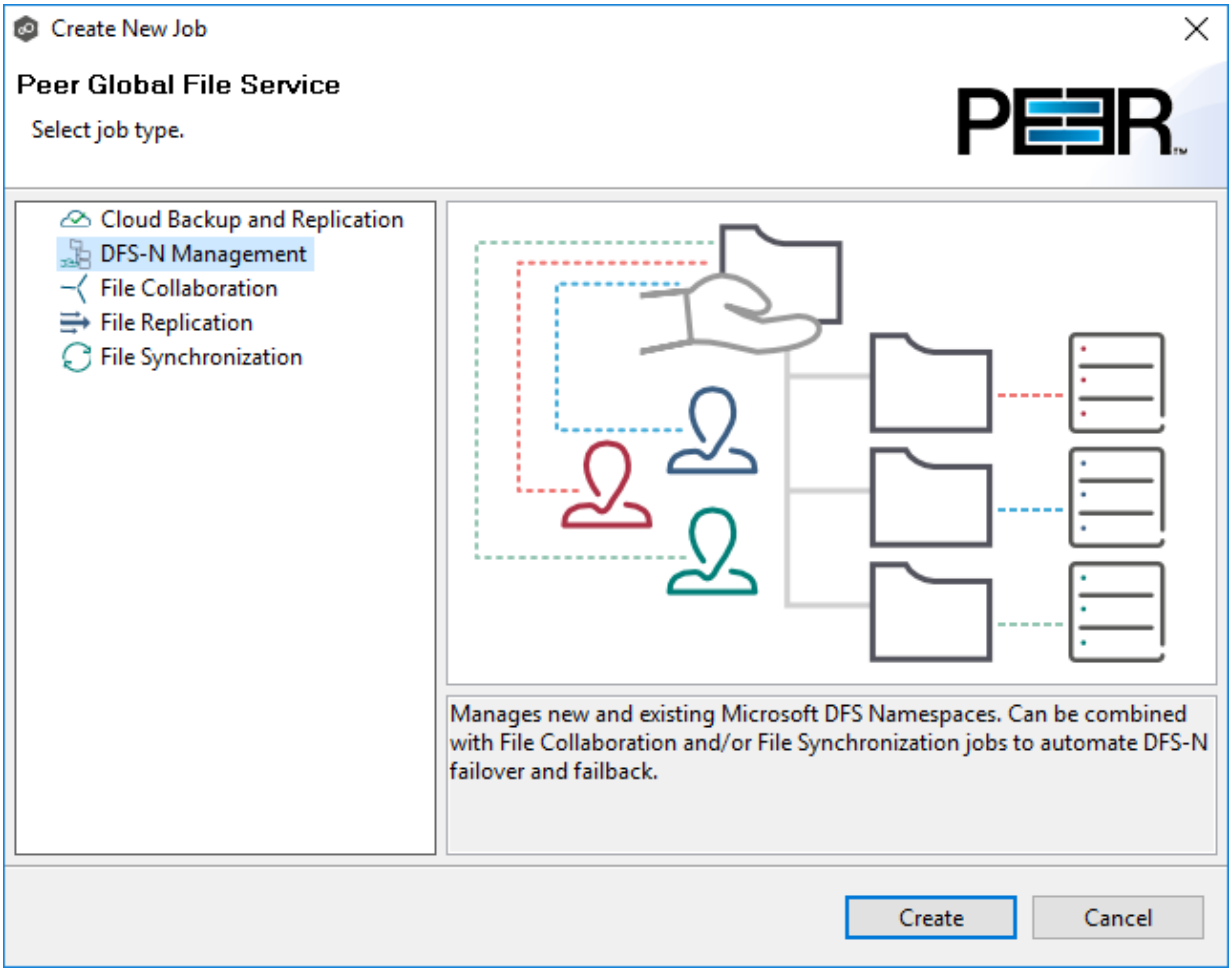

The following dialog appears.

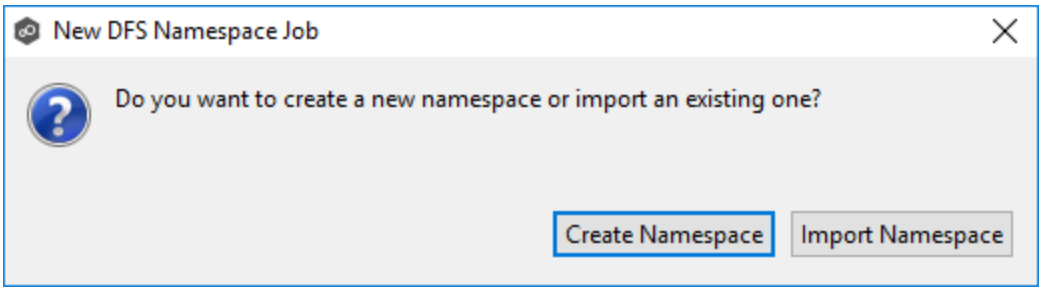

4. If you have an existing namespace you want to import, click **Import Namespace**, and then follow the steps for importing an existing namepace.

   Otherwise, click **Create Namespace**.

   The Management Agent page appears.

Step 2: Management Agent

On the **Management Agent** page, you select a Management Agent for this DFS-N Management job from the list of servers that have a Peer Agent installed.

Recommendation: Select an Agent that will be dedicated to managing DFS-N Management jobs. This enables the Agent to continue managing the namespace even if other Agent servers go down. If you use a dedicated Agent for DFS-N Management jobs, this Agent will not count against the number of licensed servers.

To reduce the number of Windows servers in your environment, you can install an Agent that runs on the same server as Peer Management Center and use this Agent to manage DFS-N Management jobs. This Agent will also not count against the number of licensed servers.

1. Select an Agent that is in the domain of the DFS namespace of where you want to create the new DFS namespace.

   Note: If you select an Agent that has **No** in the **DFS Mgmt. Enabled** column, the Microsoft DFS PowerShell Management toolkit will be installed in Step 3: Agent Verification.

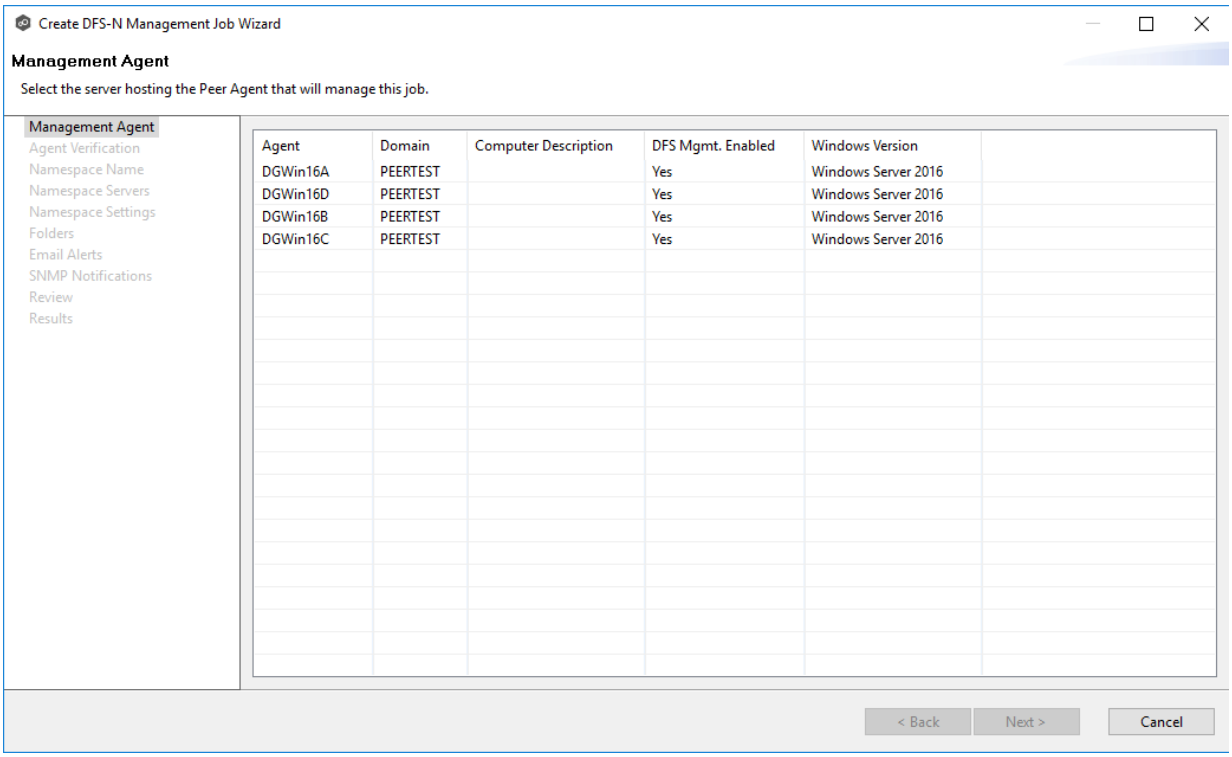

2. Click **Next**.

   The Agent Verification page appears.

Step 3: Agent Verification

The purpose of the **Agent Verification** page is to verify that environment of the selected Management Agent is set up properly to communicate with DFS Namespaces. For example, the Microsoft DFS PowerShell Management toolkit must be installed on the same system as the Management Agent and configured correctly. If the toolkit is not already installed, you will be able to install it during the verification process.

Note: The verification process does not include checking whether DFS Services is running because DFS Services doesn't have to run on the Agent server itself; it typically runs on a domain controller.

1. Click **Start Verification**.

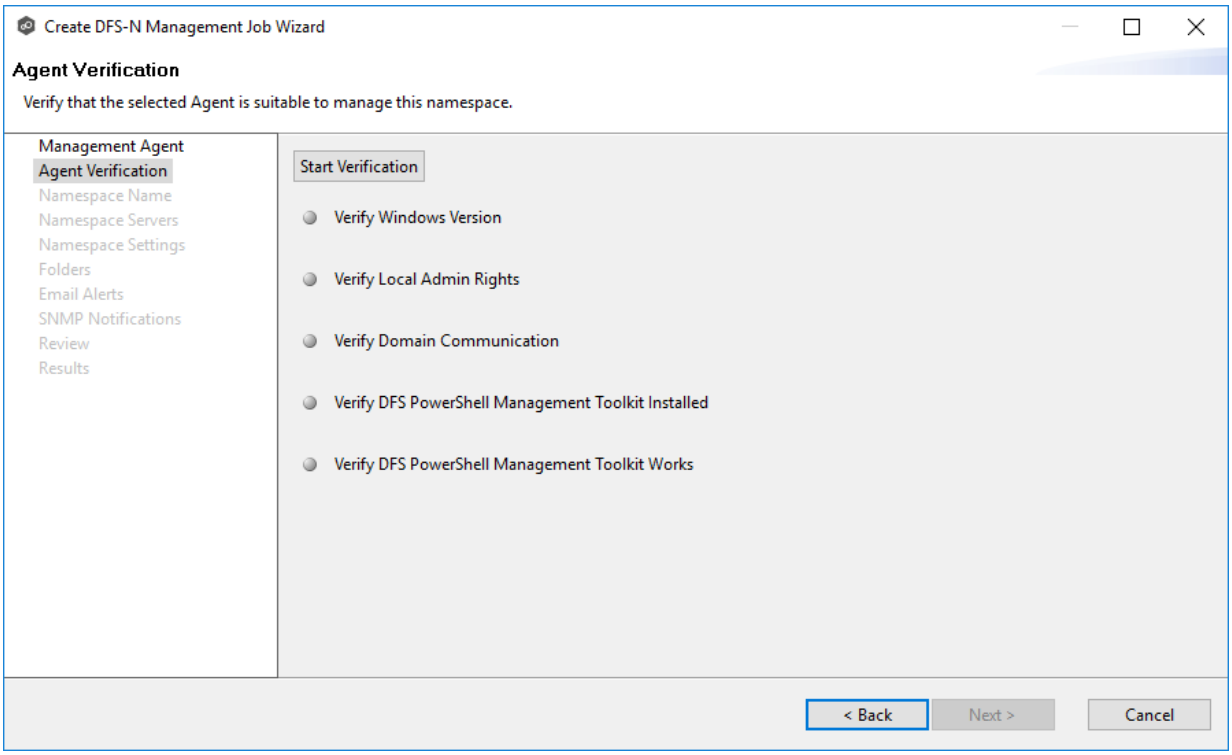

2. If the DFS PowerShell Management toolkit is not installed, click the **Install** button that will appear next to **Verify DFS PowerShell Management Toolkit Installed**.

   After the toolkit is installed, the verification continues. A green dot signifies that the verification of that element was successful.

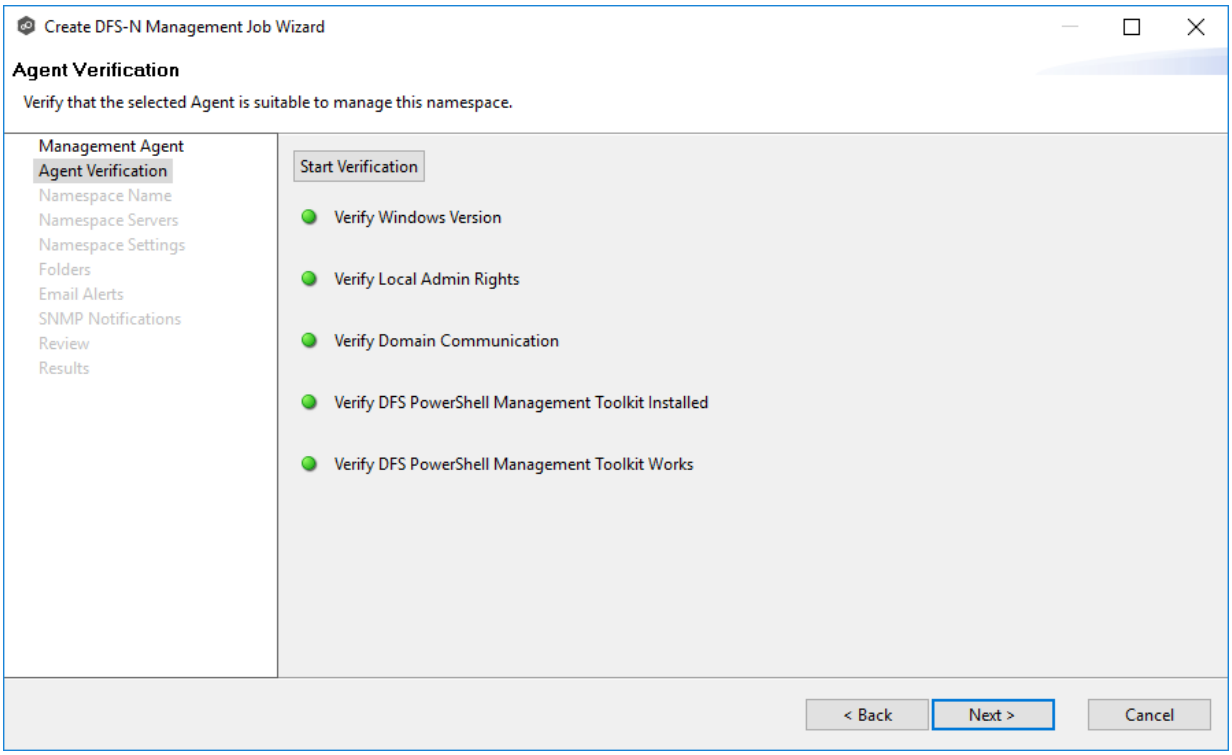

3. After the verification has successfully completed, click **Next**.

   The Namespace Name page appears.

Step 4: Namespace Name

On the **Namespace Name** page, you enter a name for the new namespace. The name of the namespace will also be the name of this DFS-N Management job.

1. Enter a name for the namespace.

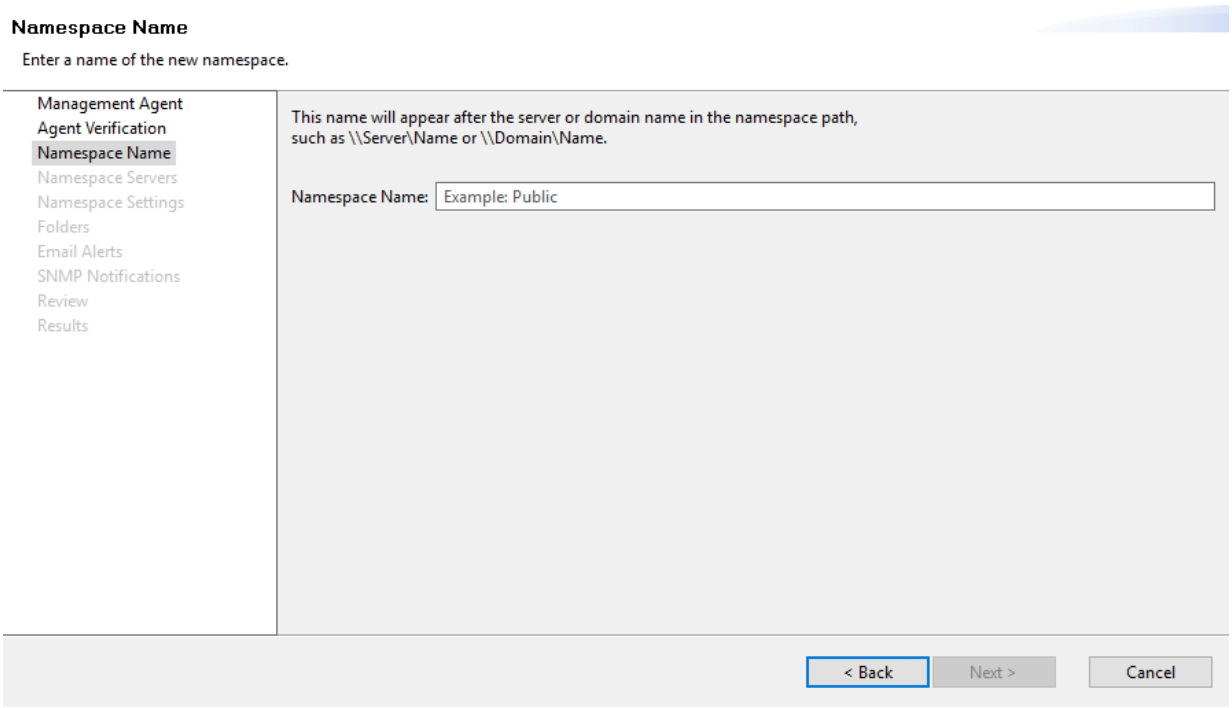

2. Click **Next**.

   The Namespace Servers page appears.

Step 5: Namespace Servers

On the **Namespace Servers** page, you select one or more servers to host the namespace. A namespace server must be a member server or domain controller in the domain in which the namespace is configured. The Microsoft DFS Namespaces service must also be running on the namespace server.

1. Enter the fully qualified domain of a namespace server in the **Server Name** field, and then click **Add**.

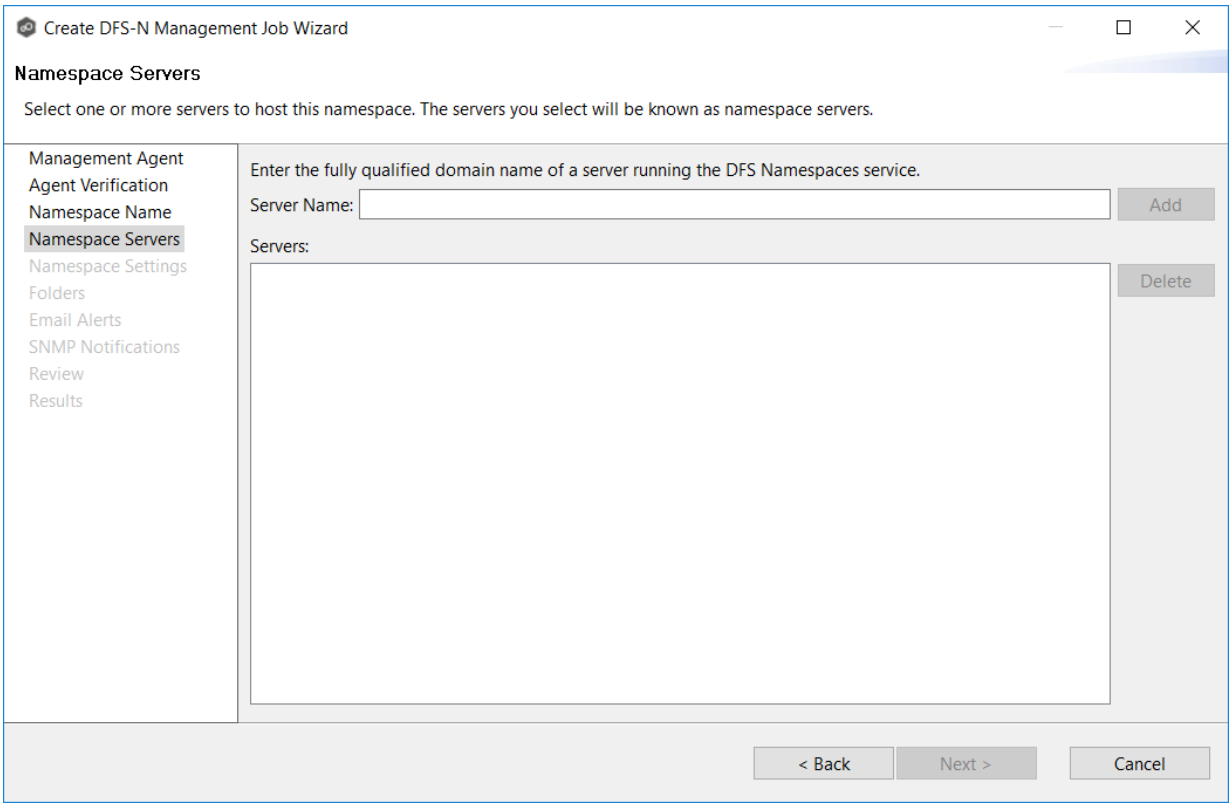

The server path is listed in the **Servers** area below.

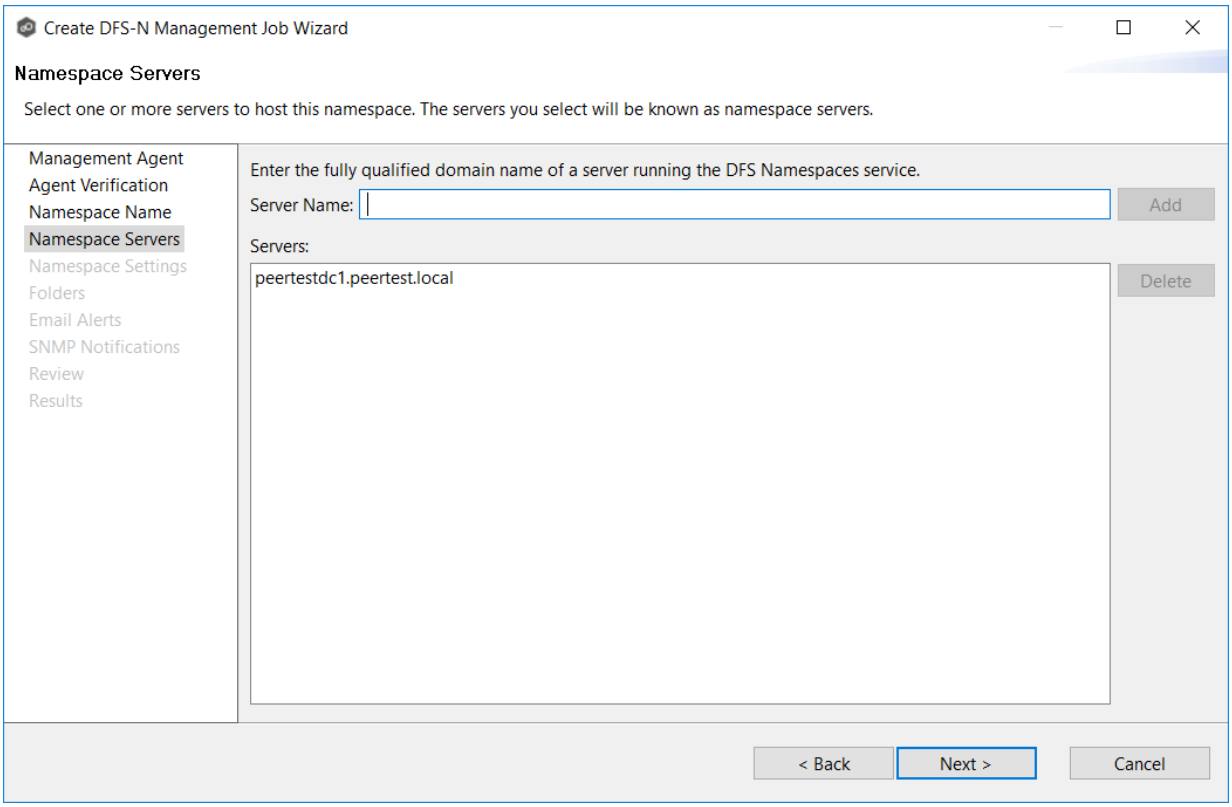

2. Add additional servers if desired.

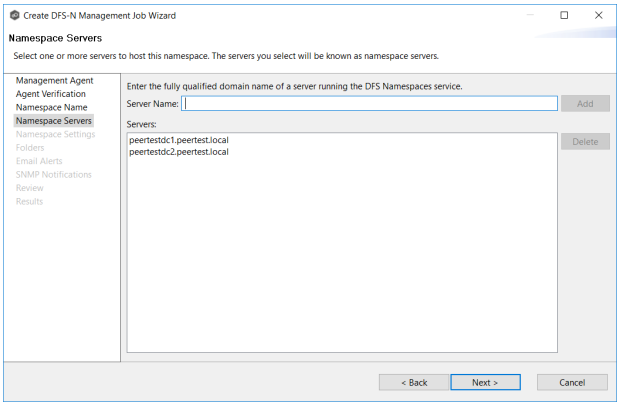

3. Click **Next**.

   The Namespace Settings page appears.

Step 6: Namespace Settings

The **Namespace Settings** page displays the namespace servers selected for the job. You can modify the following namespace server's settings if necessary:

- The local path to the DFS root share for the namespace. The default location of the DFS root share is under C:\DFSRoots\ and is specified in DFS-N Management Job Preferences.
- The access permissions to the namespace server. By default, all users have full access.

To edit a server's settings:

1. In the **DFS Root Local Path** column for the server, the prepopulated path is recommended. If your DFS root share location is different than the default (C:\DFSRoots\), you can modify the first part of **DFS Root Local Path** to match. The second part of the path should be the name of the namespace.

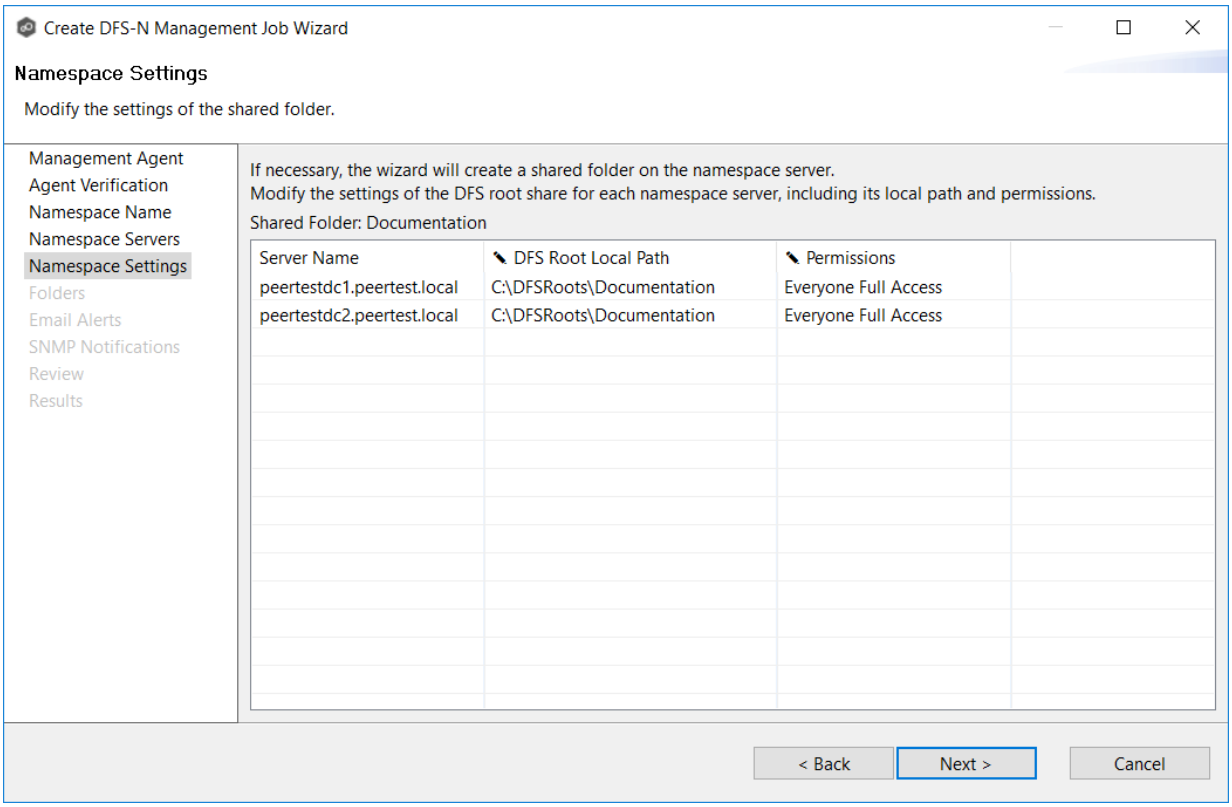

2. In the **Permissions** column for the server, select the desired access level from the drop-down menu.

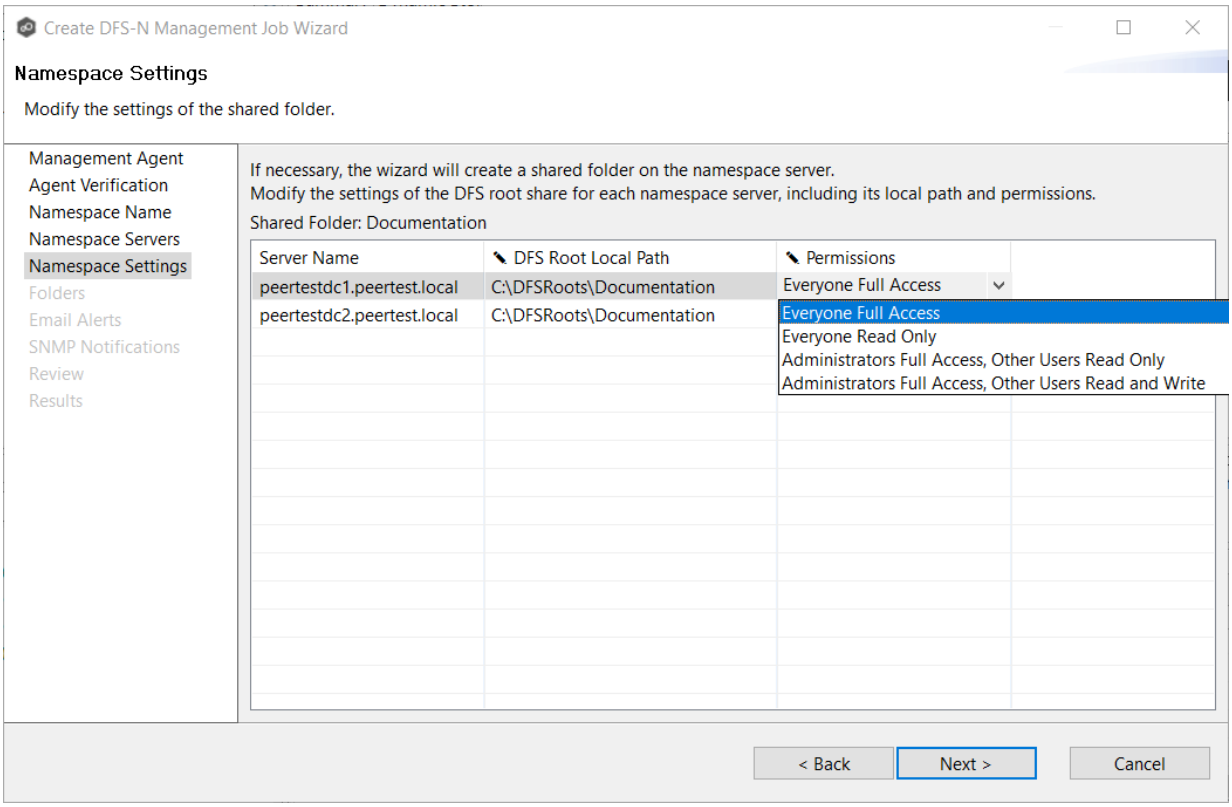

3. (Optional) Modify the path and permissions for other servers.

4. Click **Next**.

   The Namespace Folders page appears.

Step 7: Folders

Initially, the **Folders** page displays the only the namespace path. After namespace folders and folder targets are added, they are displayed in a tree structure.

1. Select the namespace path, and then click the **Create** button.

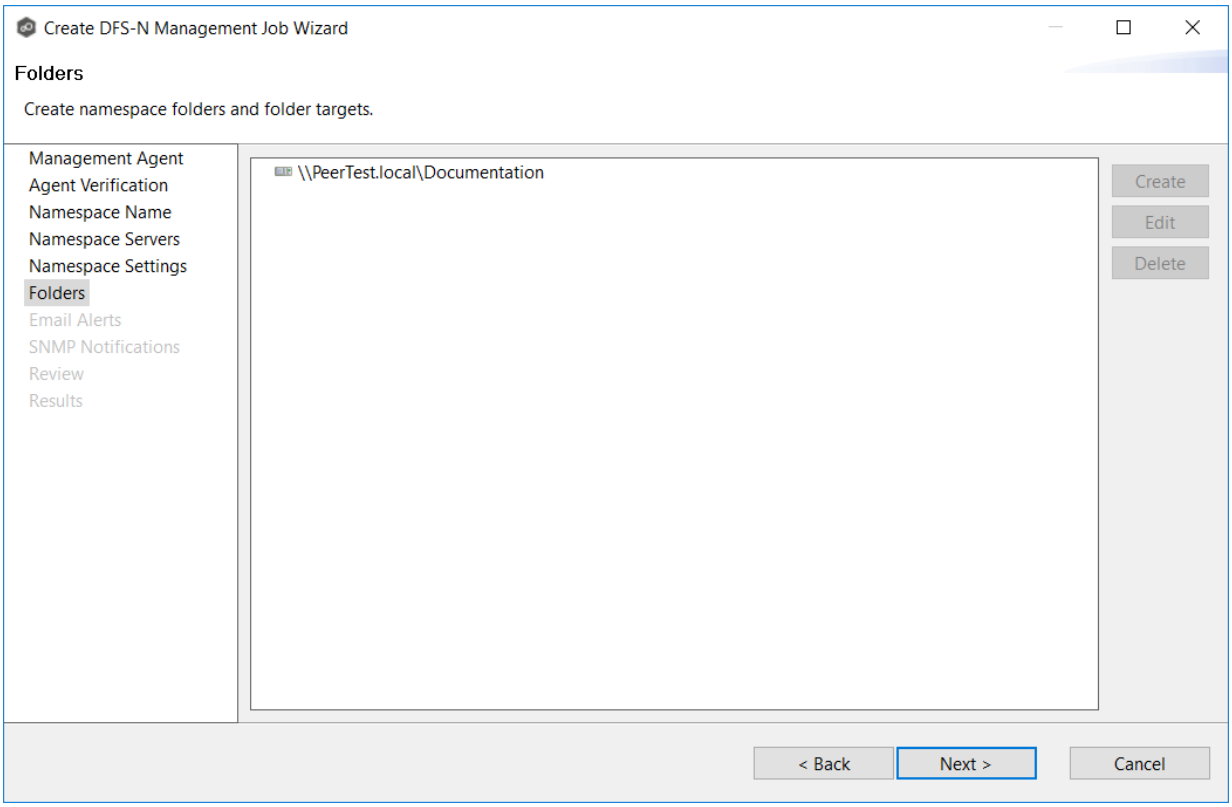

The **Folder Name** dialog appears.

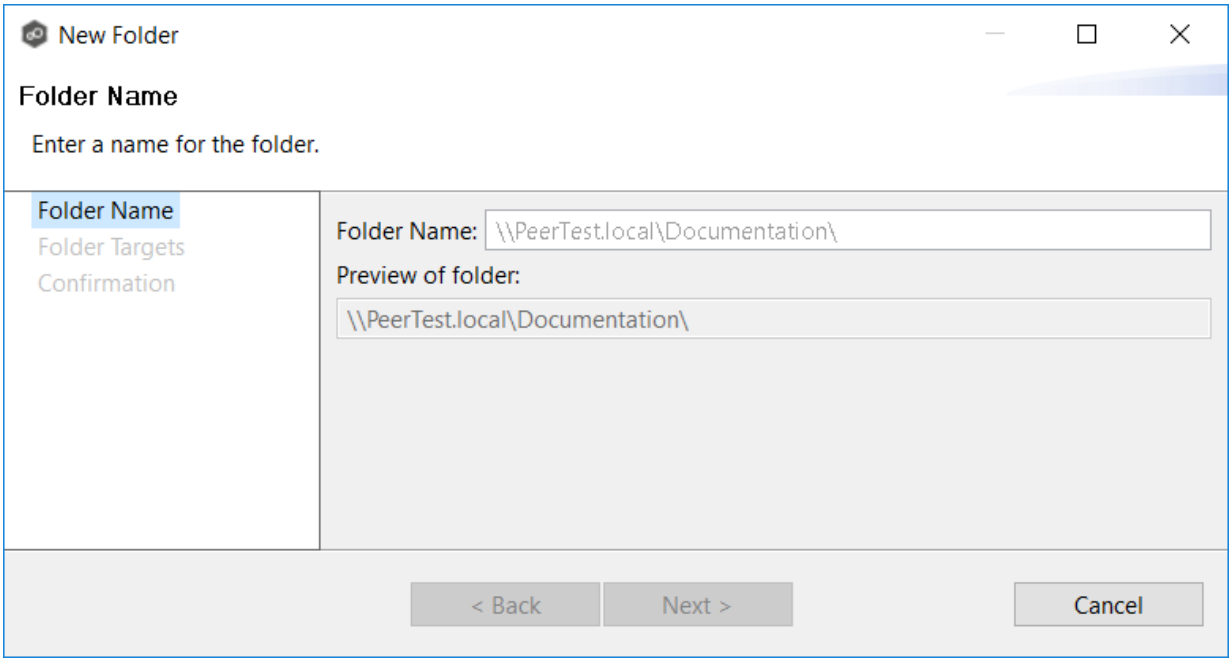

2. Enter a name for the namespace folder in the **Folder Name** field.

After you enter the folder name, a preview of the folder and path name appears below the **Folder Name** field.

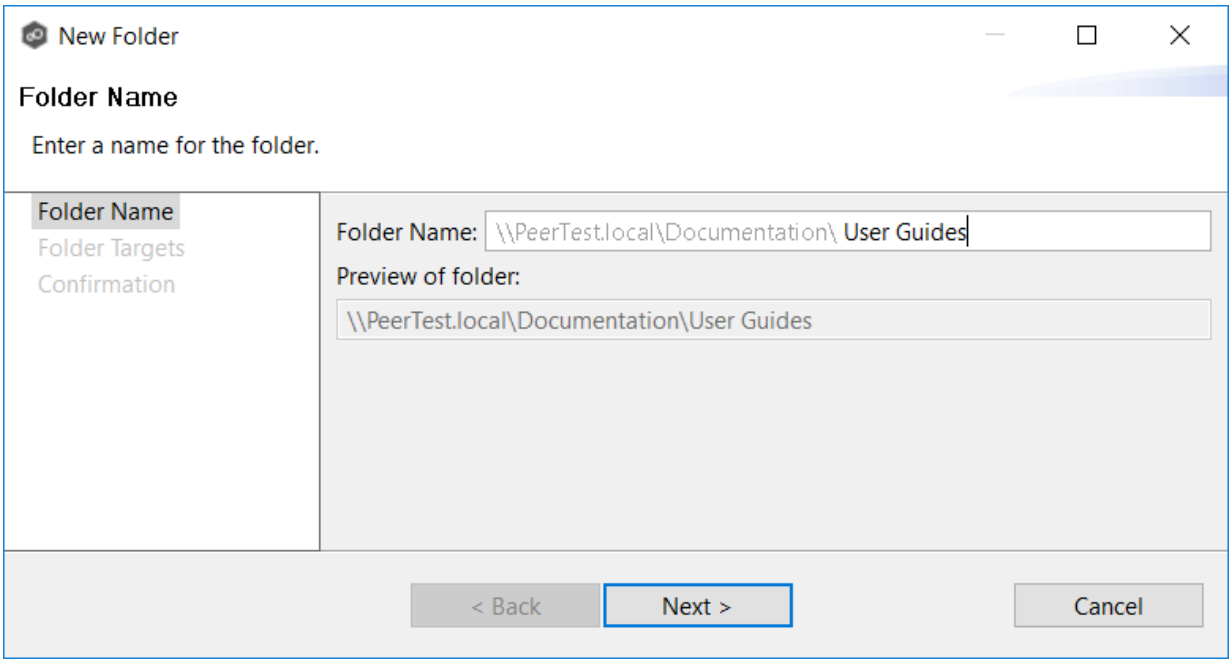

3. Click **Next**.

The **Folder Targets** dialog appears.

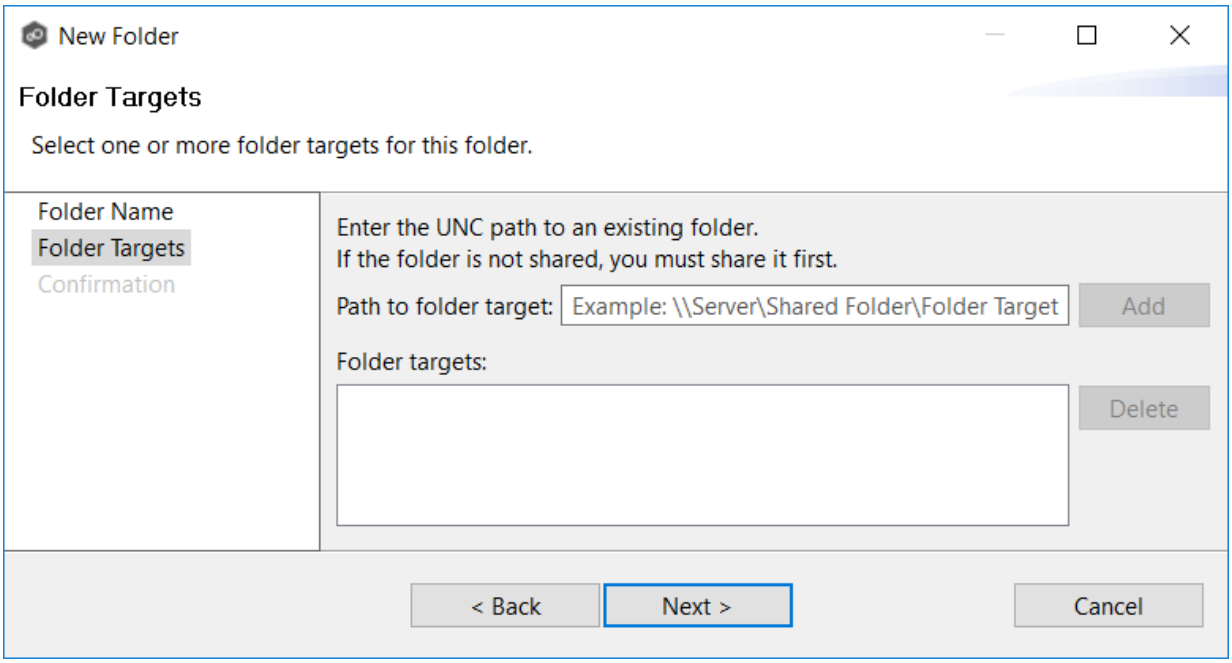

4. (Optional) To add a folder target, enter the UNC path to a shared folder, and then click **Add**.

   If you haven't created your folders target yet, you can skip to Step 6 and add folder targets to the job later.

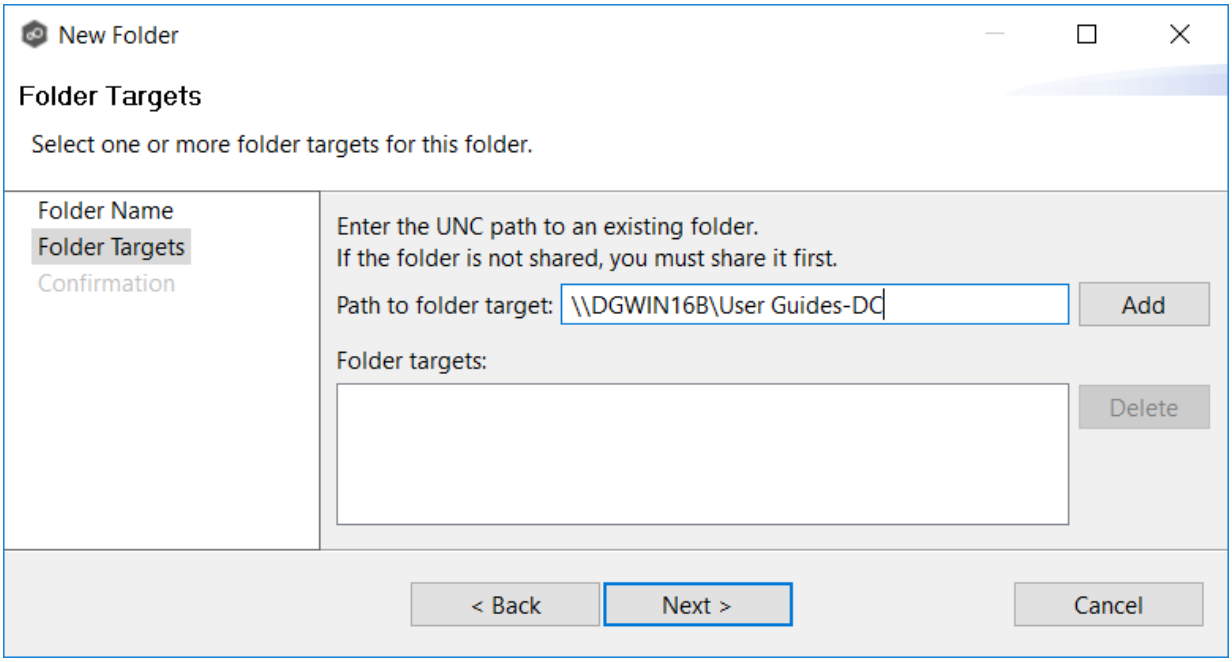

   After the share is validated, it appears in the **Folder targets** area:

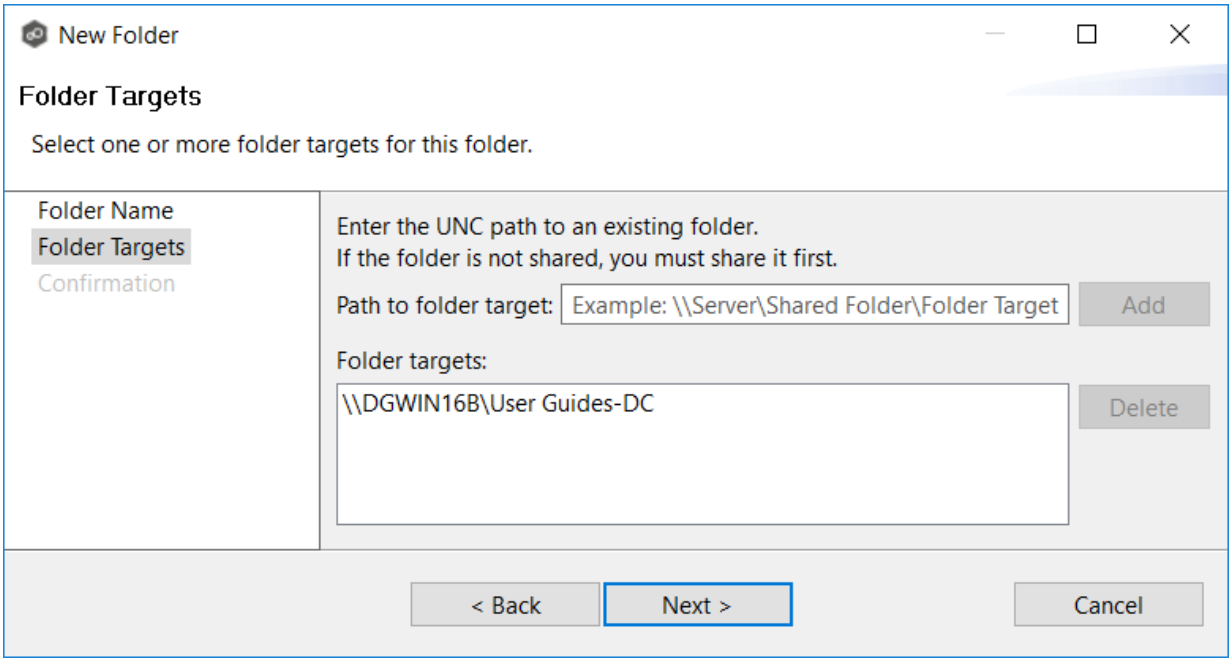

5. (Optional) Add additional folder targets.

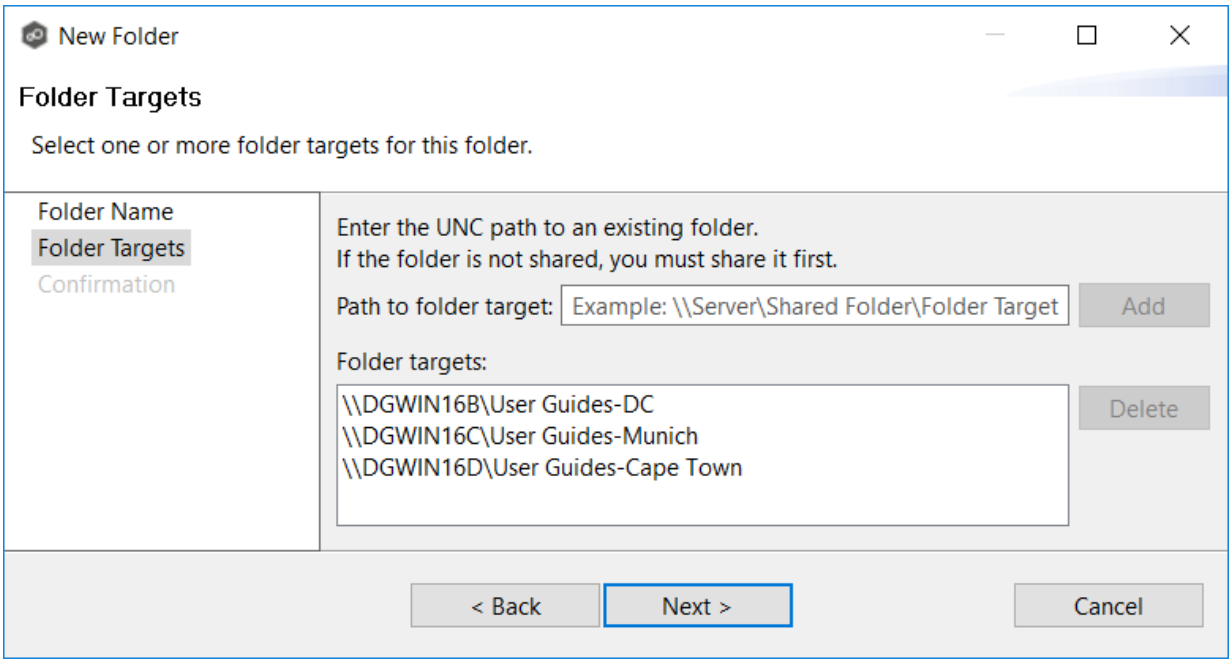

6. Click **Next**.

   The **Confirmation** dialog appears.

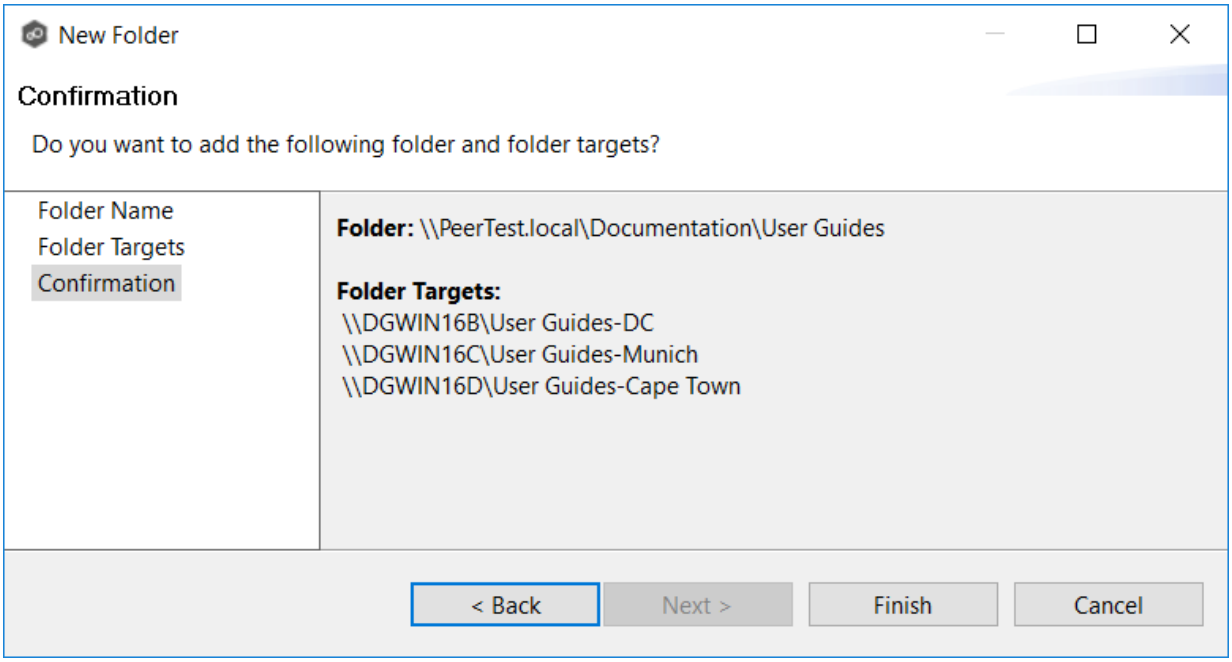

7. Review the folders and folder targets.

8. Click **Back** to add more folder and folder targets; otherwise, click **Finish**.

   The **Folders** page reappears.

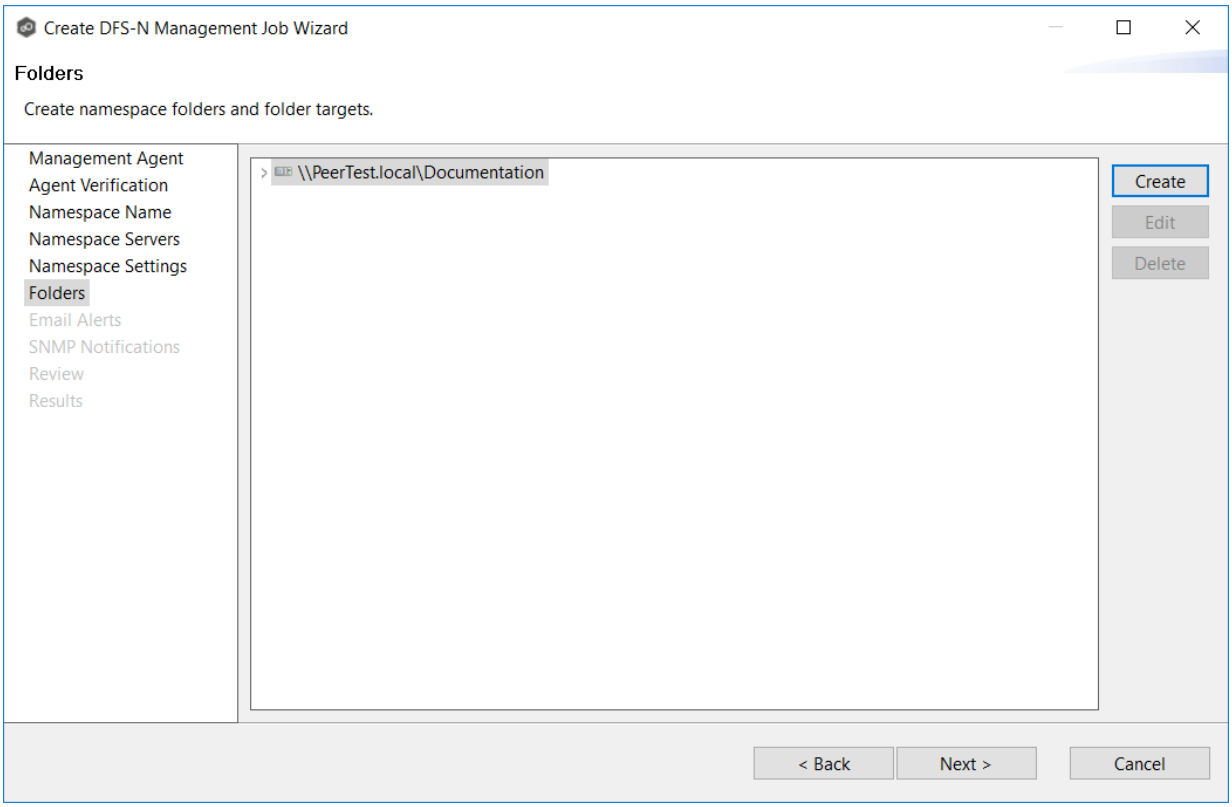

9. Expand the tree to view the folders and folder targets you added.

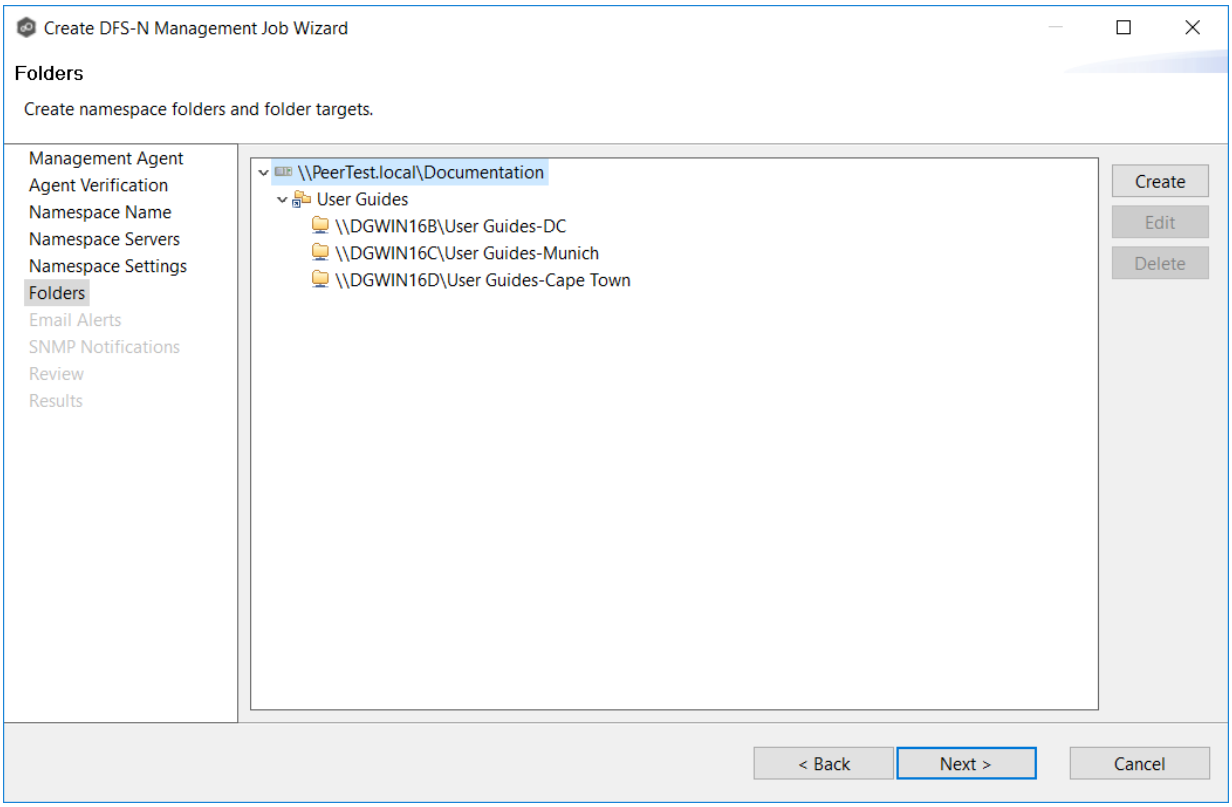

10. Click **Next**.

    The Email Alerts page appears.

Step 8: Email Alerts

This step is optional.

An email alert notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in Preferences and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply. To create a new alert, see Email Alerts in the Preferences section.

To apply an existing email alert to the job:

1. Click the **Add** button.

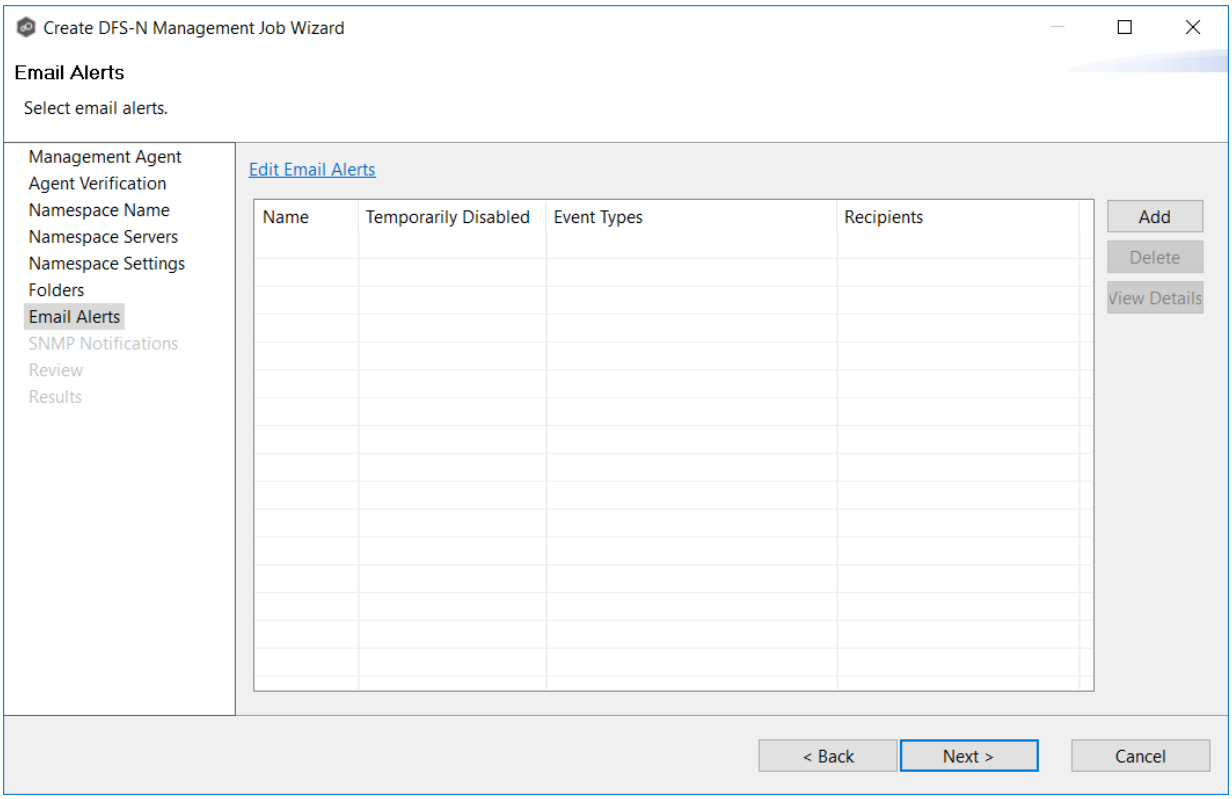

The **Select Email Alert** dialog appears.

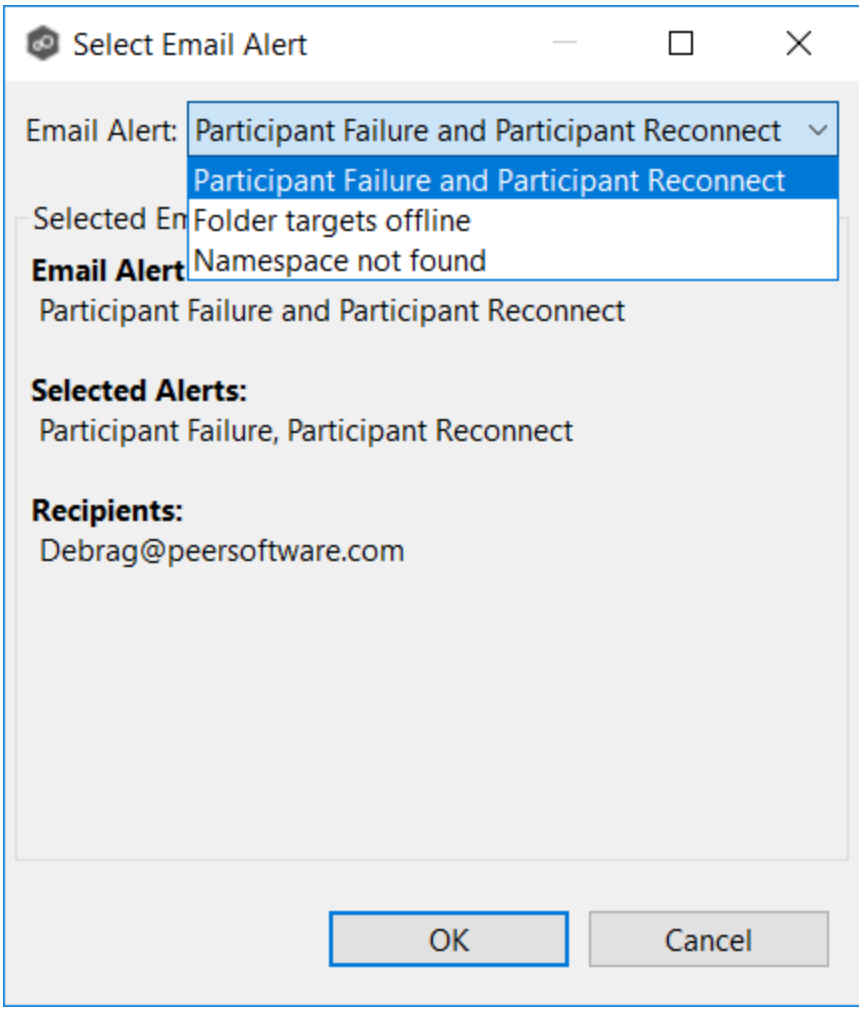

2. Select an alert from the **Email Alert** drop-down list, and then click **OK**.

   The alert is listed in the **Email Alerts** page.

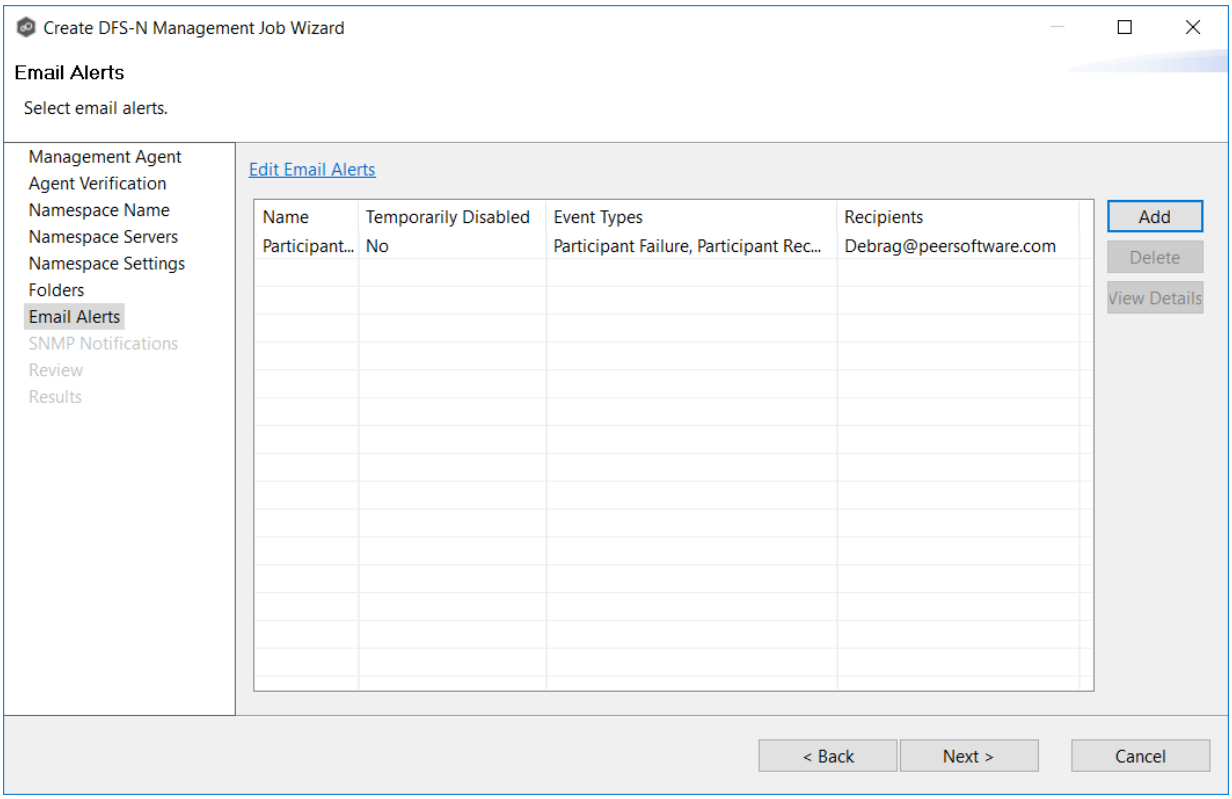

3. (Optional) Repeat steps 1-3 to apply additional alerts.

4. Click **Next**.

   The SNMP Notifications page appears.

Step 9: SNMP Notifications

This step is optional.

An SNMP notification notifies recipients when certain type of event occurs, for example, session abort, host failure, system alert. The **SNMP Notifications** page displays a list of notifications that have been applied to the job. When you first create a job, this list is empty. Like email alerts and file filters, an SNMP notification is defined in Preferences and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create SNMP notifications in advance. However, from this wizard page, you can select an existing SNMP notification to apply to the job or create new SNMP notifications. To create a new alert, see SNMP Notifications in the Preferences section.

To apply an existing SNMP notification to the job or disable notifications:

1. Select an SNMP notification from the drop-down list.

   If you don't want to send a SNMP notification, select **None - Disabled**.

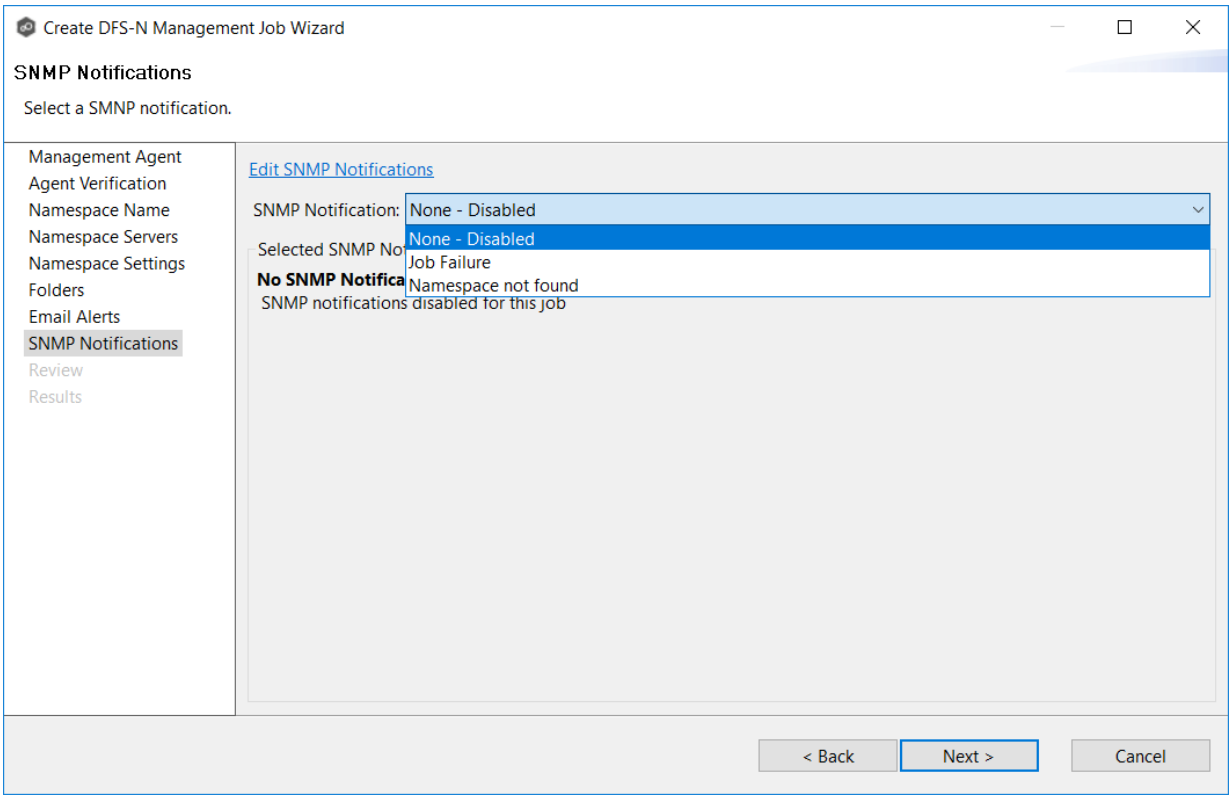

   If you select a notification, details about the notification appear in the Selected SNMP Notification Information section.

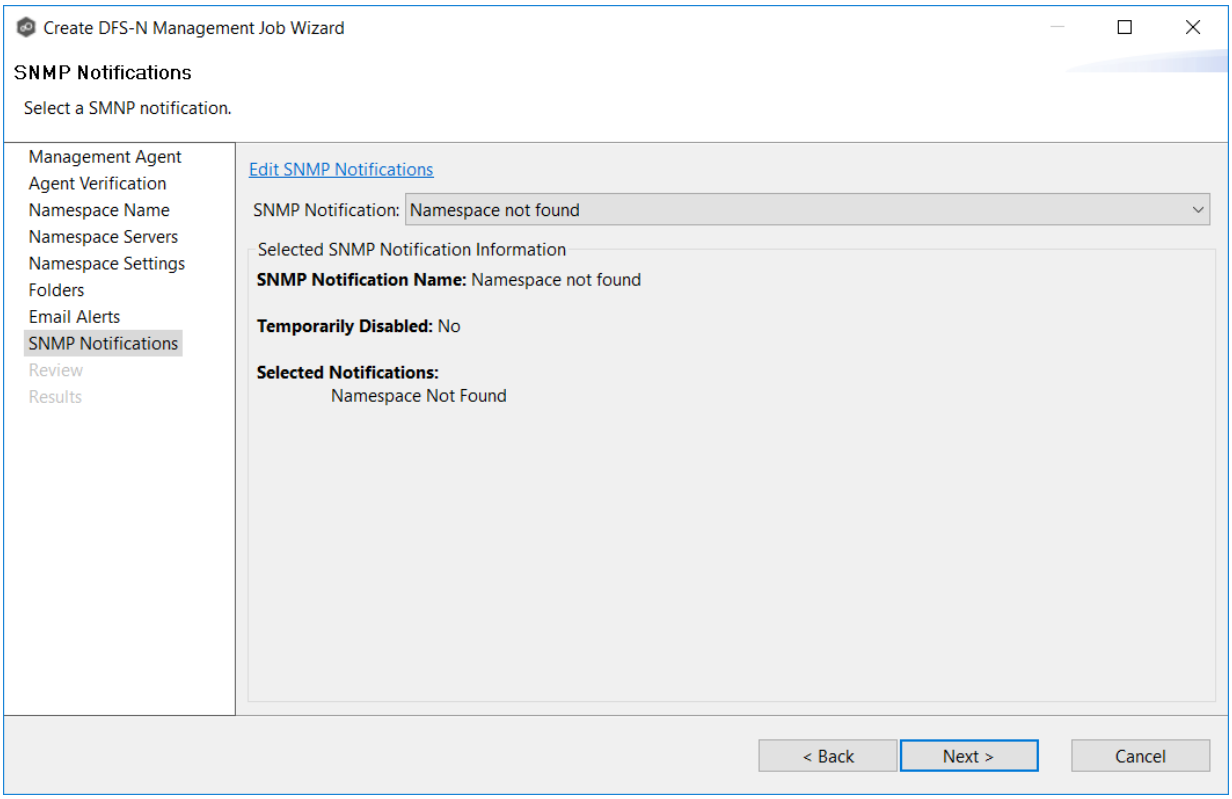

2. Click **Next**.

   The Review page appears.

Step 10: Review

The **Review** page allows you to review the configuration before it is actually created.

1. Review the namespace configuration.

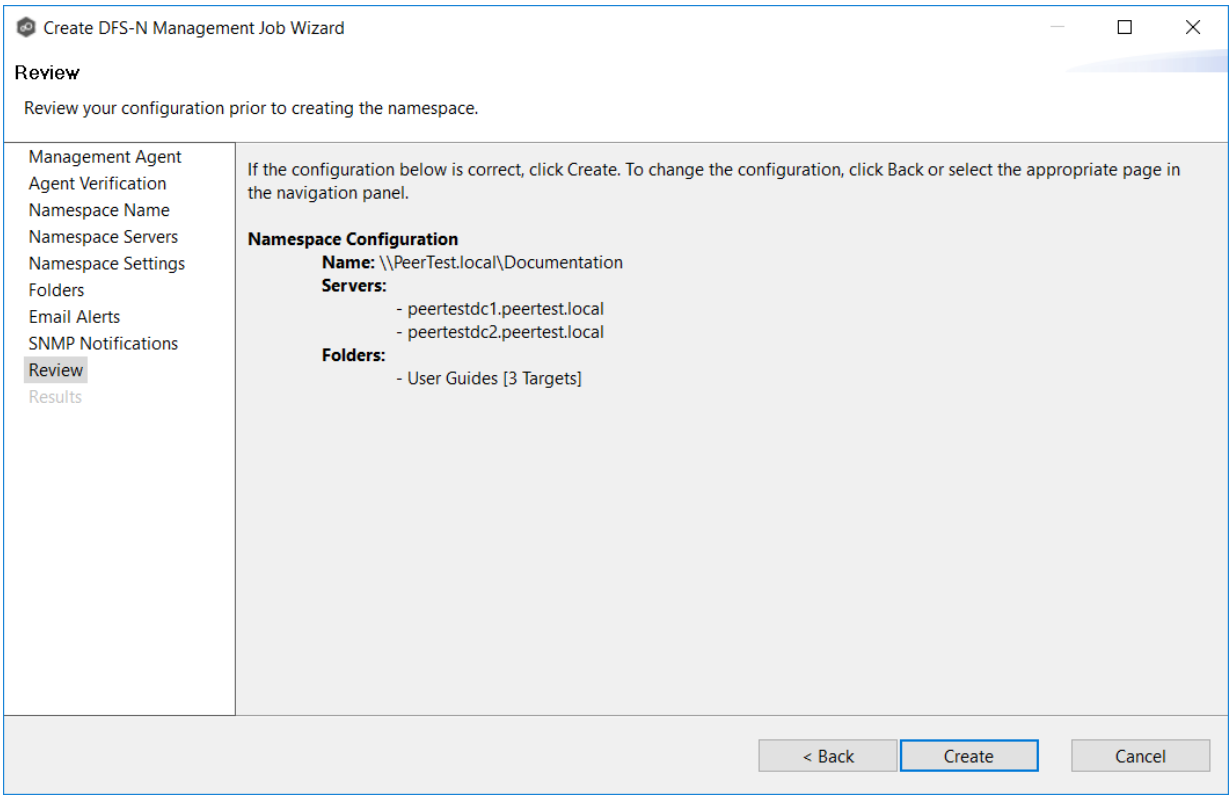

2. Click **Create** if the configuration is correct; otherwise, click **Back** and correct the configuration.

   After you click **Create**, the Results page appears.

Step 11: Results

The **Results** page has two tabs: **Tasks** and **Errors**.

1. Review the results in the **Tasks** and **Errors** tabs.

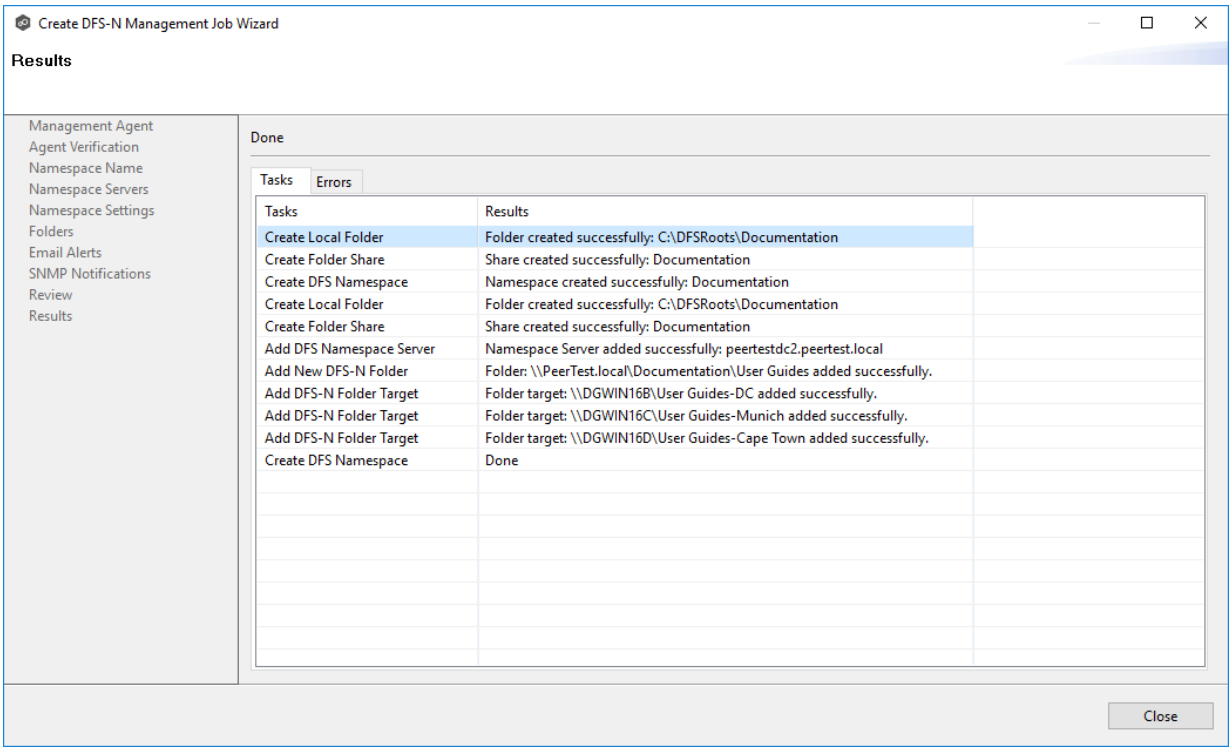

2. Review each task to verify that it was successful; it there were any errors, click the **Errors** tab to view more details about the problem.

3. After reviewing, click **Close**.

   If there were no errors, the job automatically starts and the summary view for the new job is displayed.

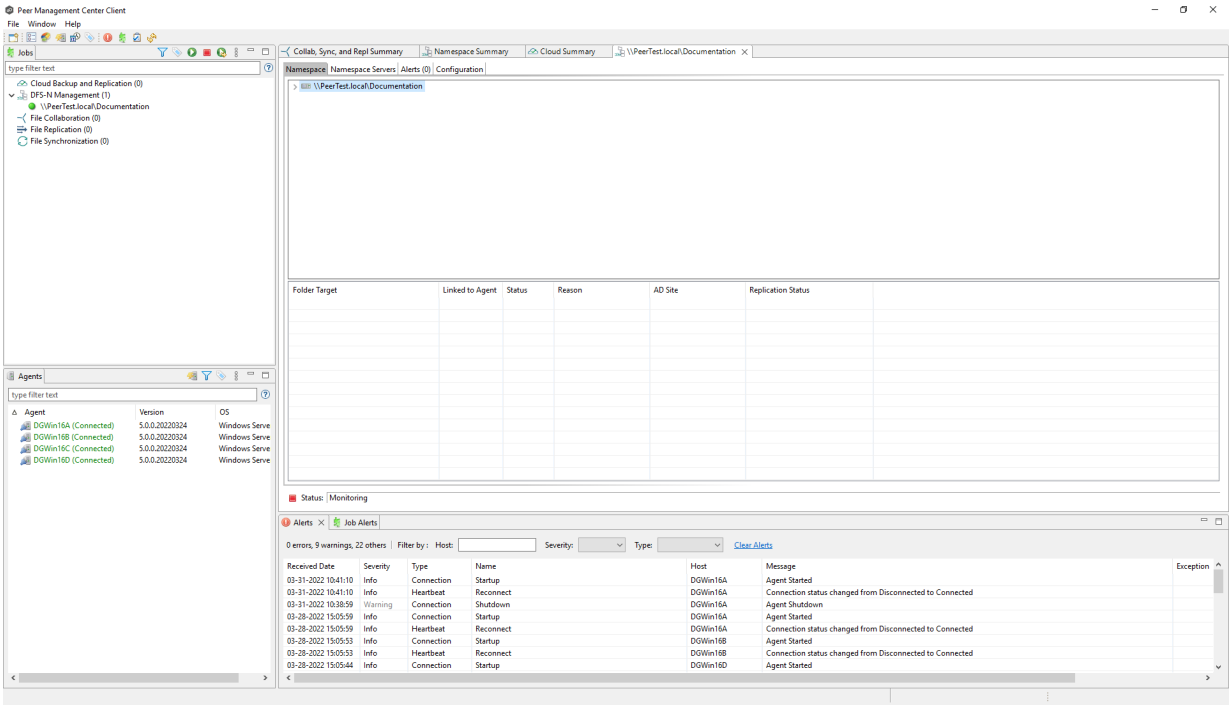

4. Select the namespace in the **Namespace** tab and then select the namespace folder to view its folder targets.

   The folder targets appear in the panel below.

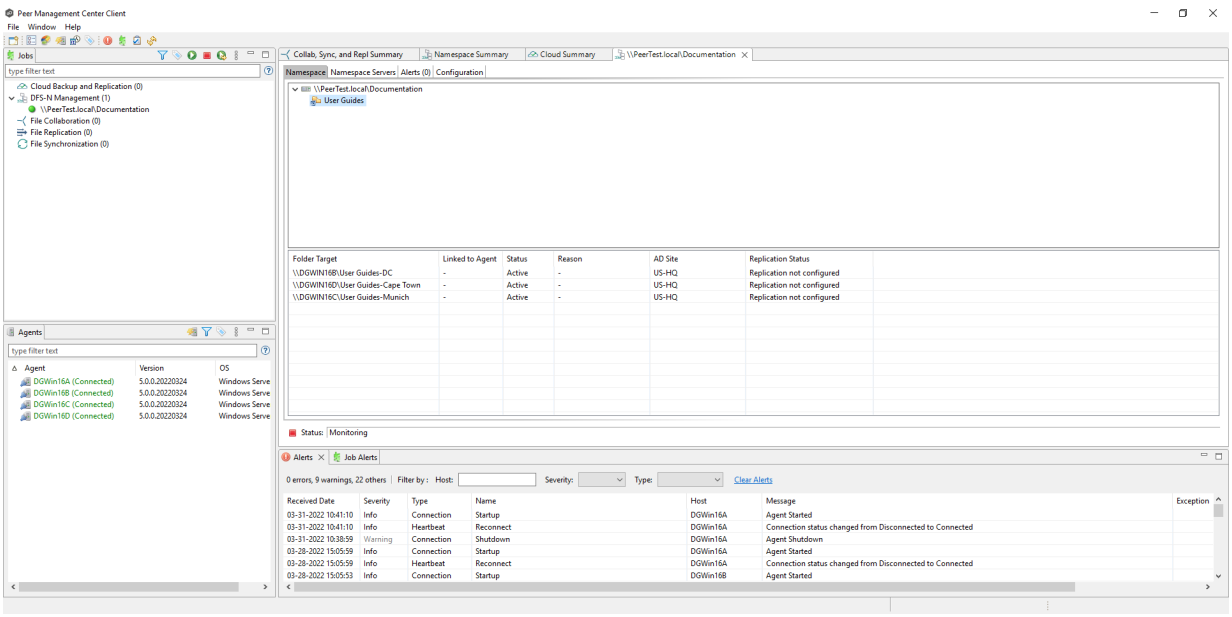

Importing an Existing Namespace

If you have an existing namespace that you want to use in in a File Collaboration or File Synchronization job, you can import the namespace. Importing the namespace automatically creates a new DFS-N Management job with the same name as the imported namespace.

You can then either link the namespace to an existing File Collaboration or File Synchronization job or create a new File Collaboration or File Synchronization job that uses the namespace.

To import an existing namespace:

1. Right-click anywhere in the **Runtime Summary** tab of the **Namespace Summary** view, and then select **Import Existing Namespaces** (or right-click the DFS-N Management job type in the **Jobs** view.

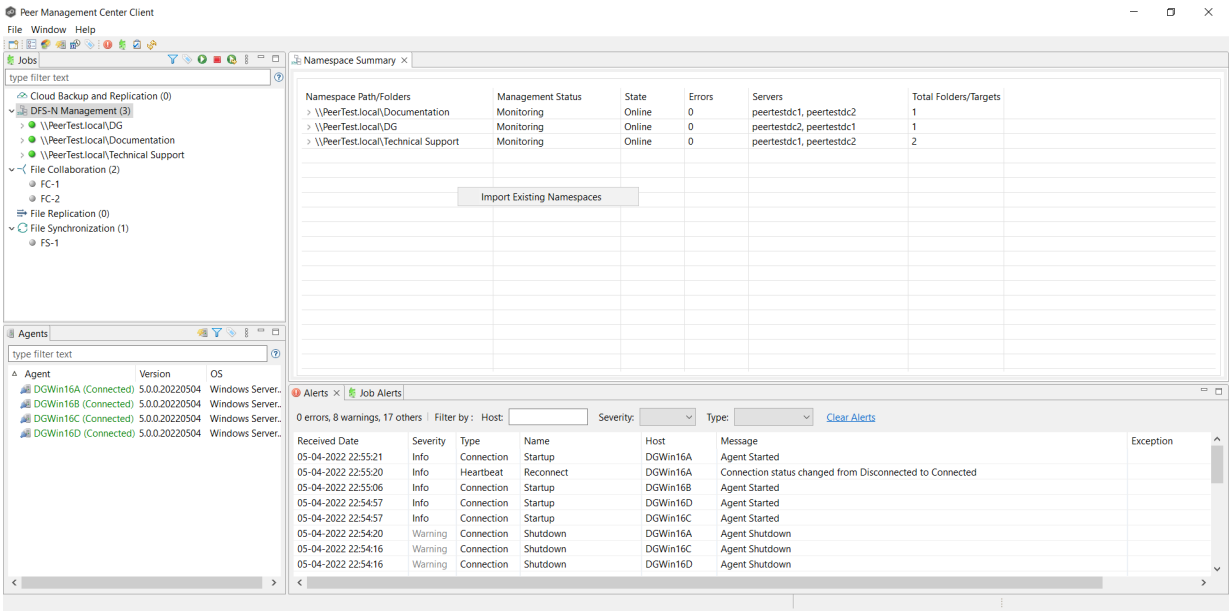

The **Import Existing Namespace** wizard appears.

2. Select a Management Agent, and then click **Next**.

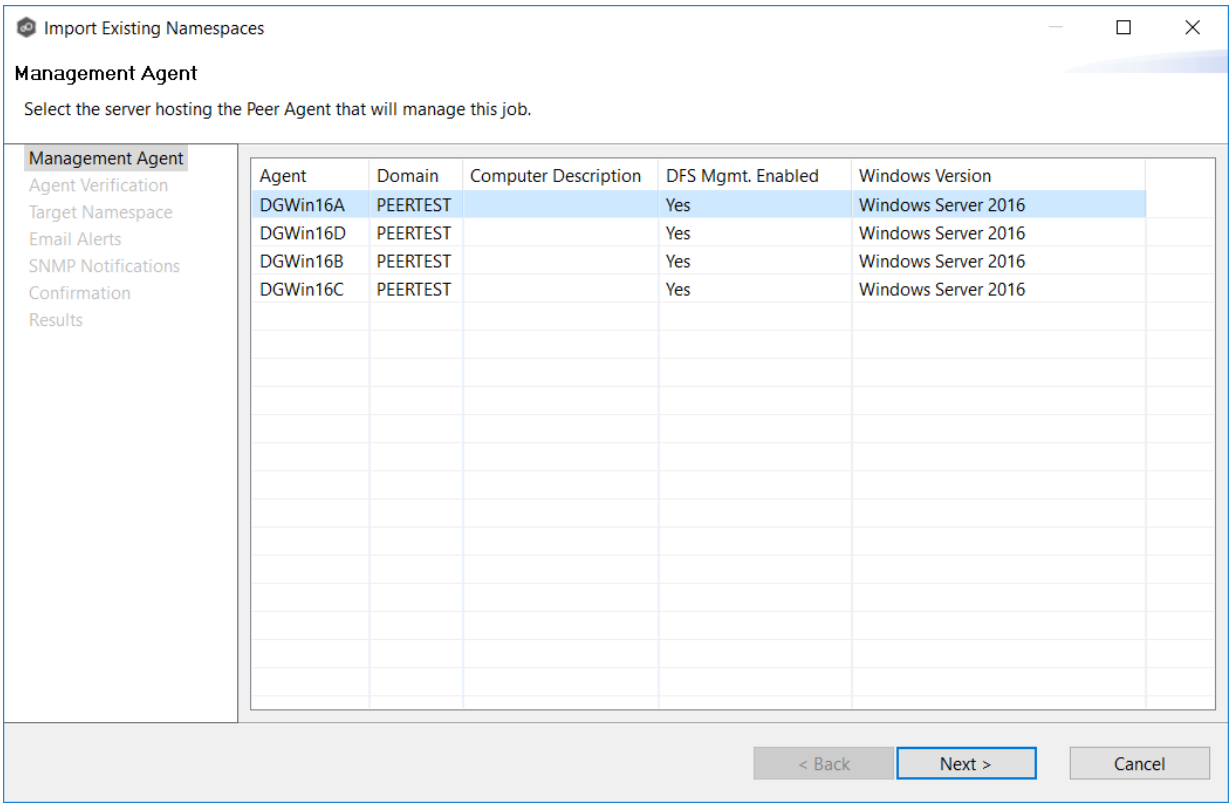

3. Click **Start Verification** to verify the Agent environment is set up to manage DFS namespaces.

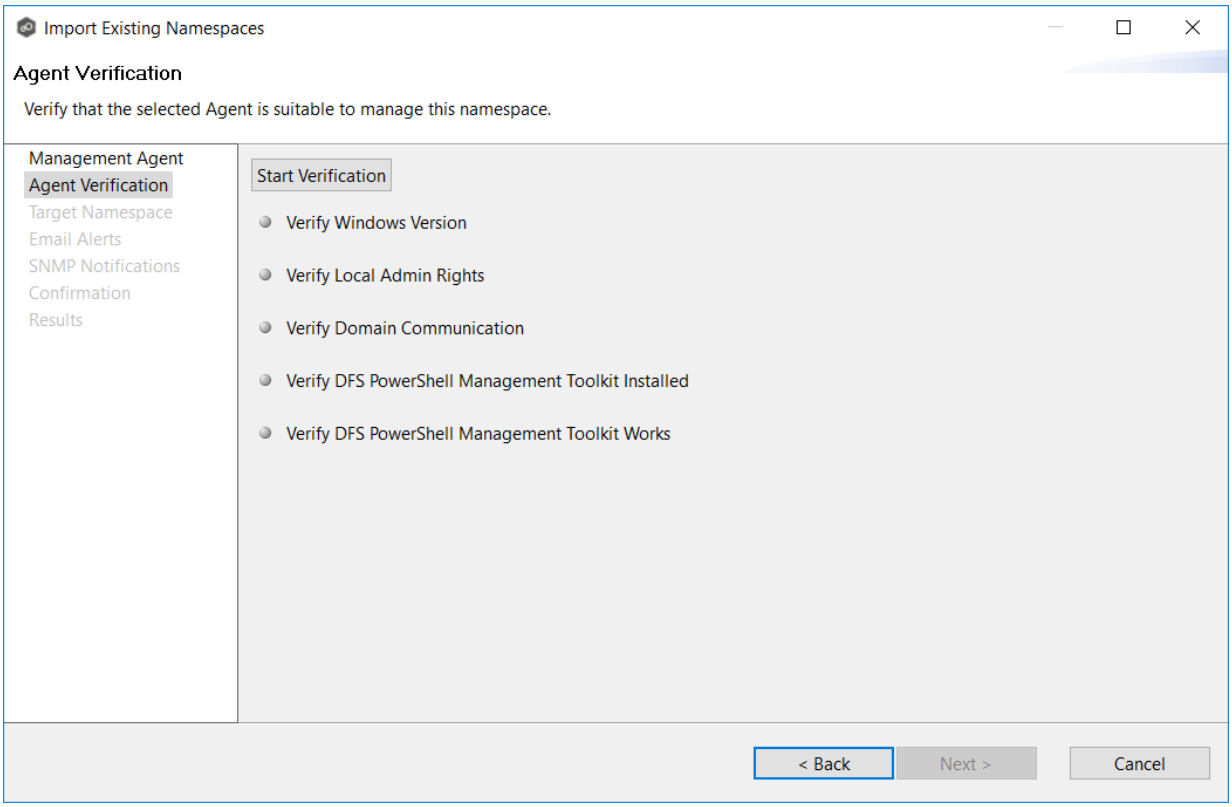

4. If the DFS PowerShell Management toolkit is not installed, click the **Install** button that will appears next to **Verify DFS PowerShell Management Toolkit Installed**.

   After the toolkit is installed, the verification continues. A green dot signifies that the verification of that element was successful.

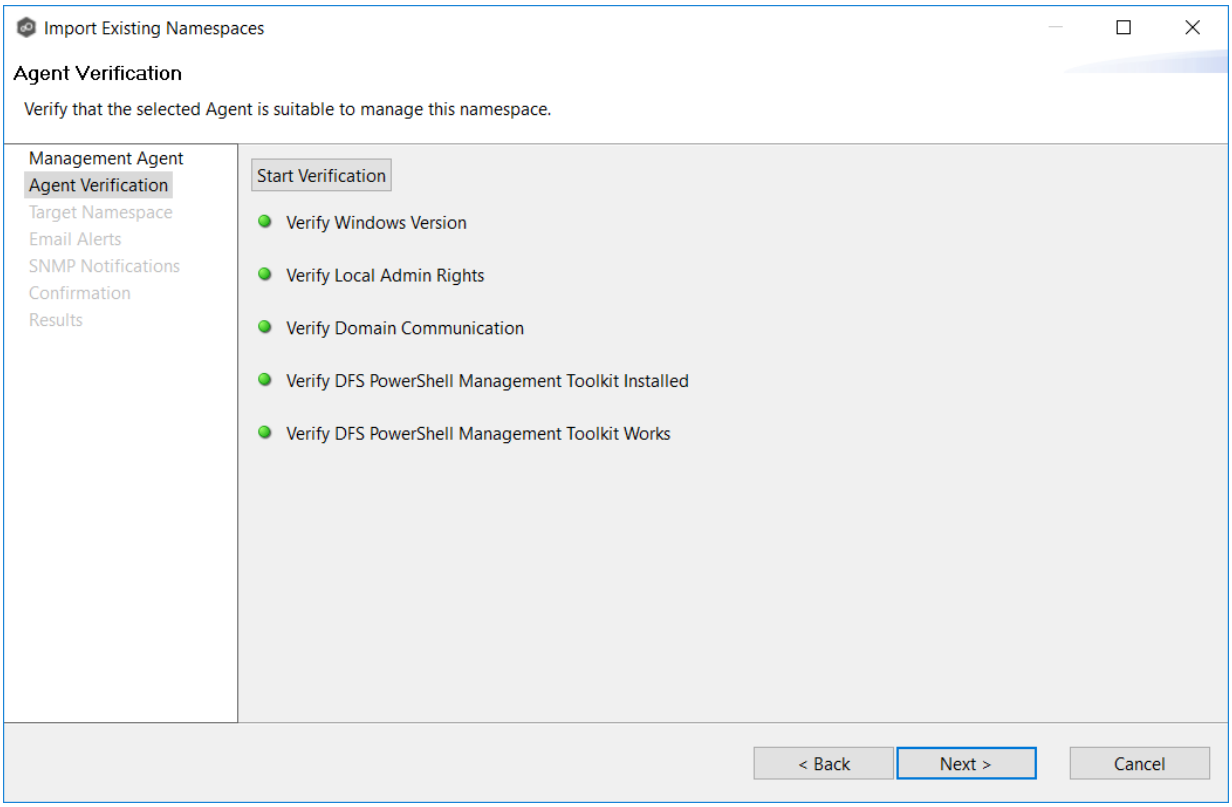

5. After the verification has successfully completed, click **Next**.

   The **Target Namespace** page appears. You have two options for selecting the namespace to import: either by entering its name or by selecting it from a list of namespaces.

6. If you choose **Select By Name**, enter the namespace name, and then click **Validate**. After the namespace is validated, skip to Step 8.

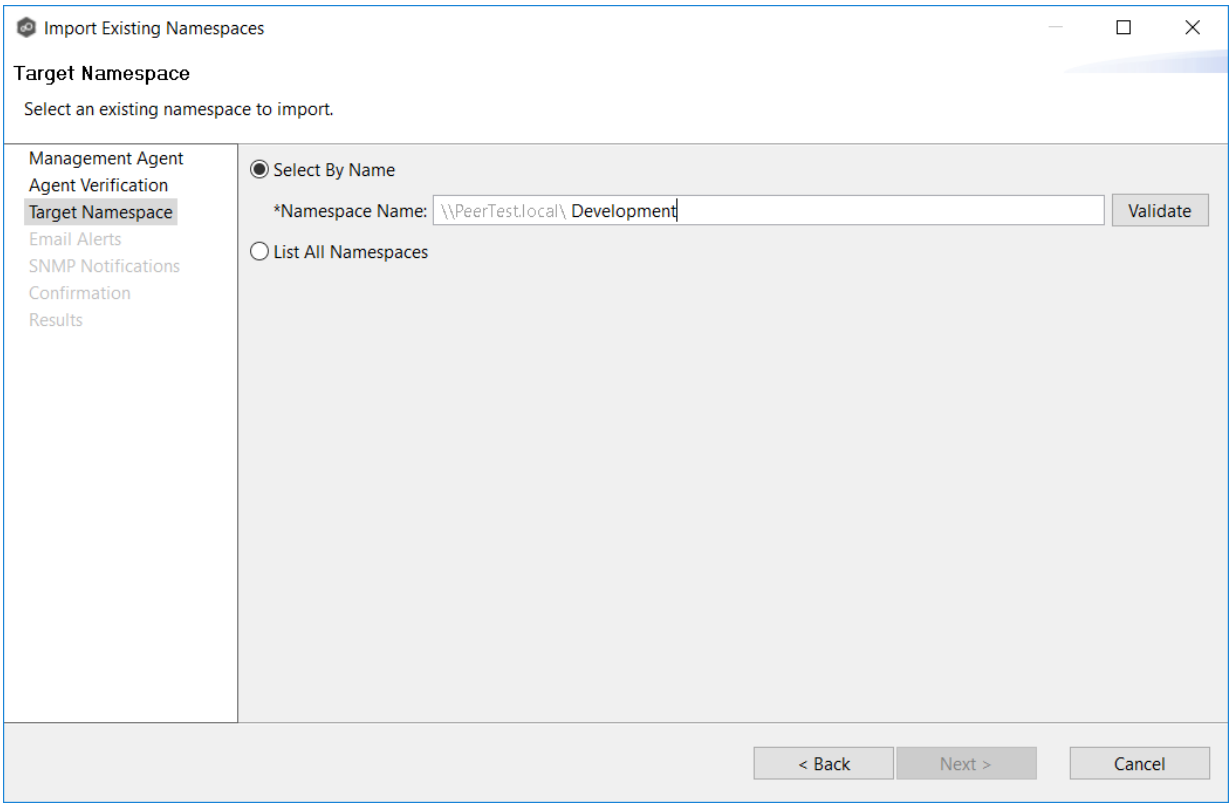

7. If you choose **List All Namespaces**, click **Next**.

   The **Existing Namespace** page appears; it displays a table listing the existing namespaces.

   Note: It may take a few minutes for existing namespaces to appear in the table.

8. Select one or more existing namespaces from the table, and then click **Next**.

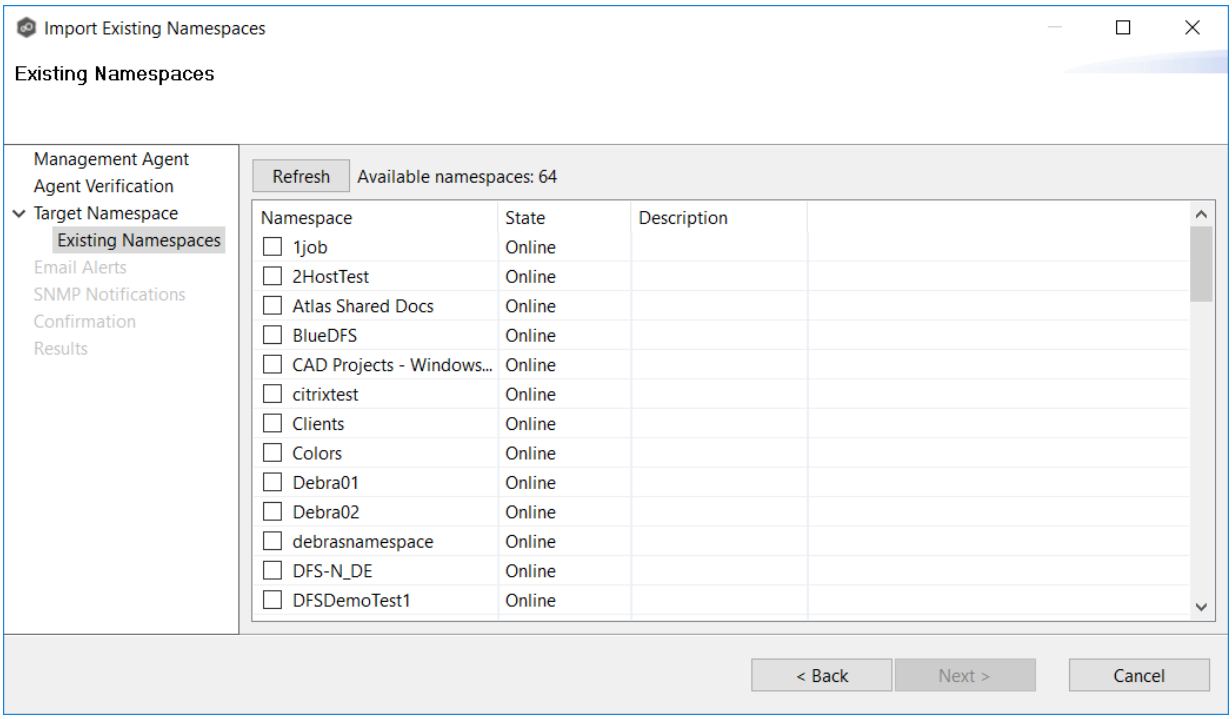

The **Email Alerts** page appears.

9. (Optional) Select or create email alerts to apply to the job, and then click **Next**.

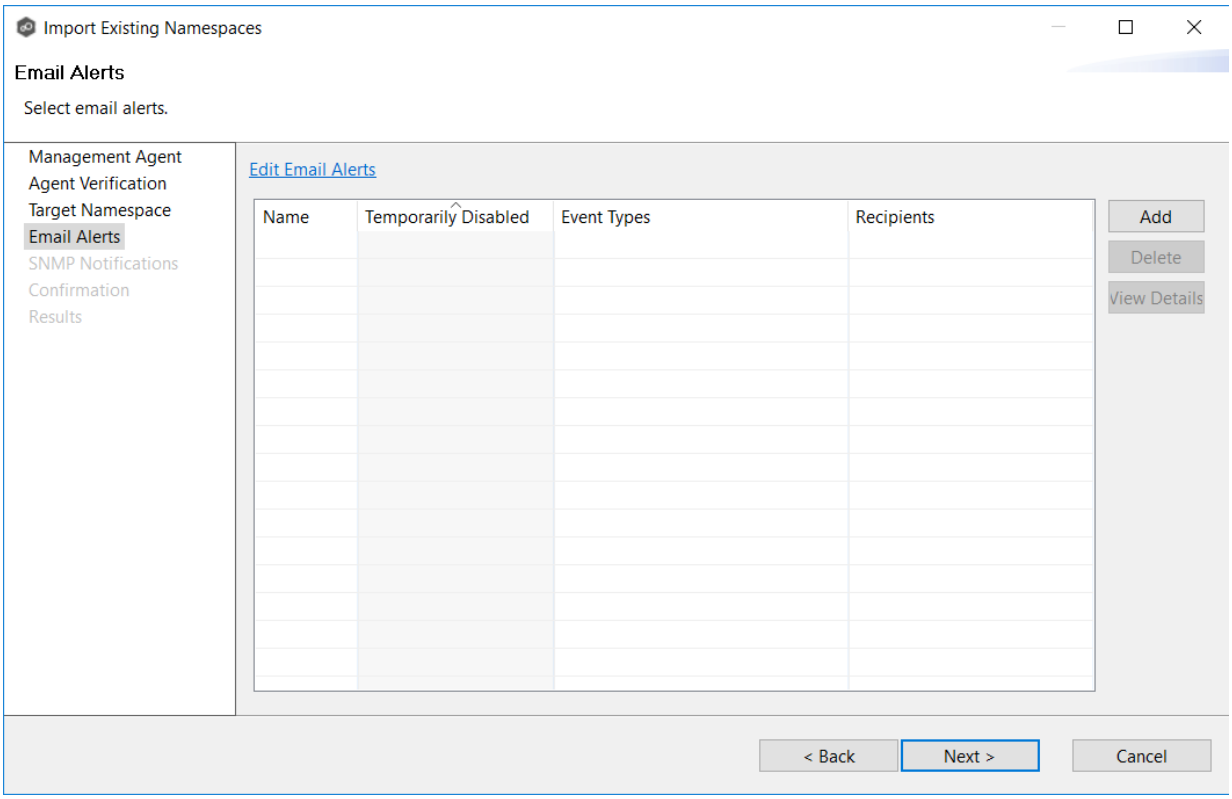

The **SNMP Notifications** page appears.

10. (Optional) Select or create an SNMP notification to apply to the job, and then click **Next**.

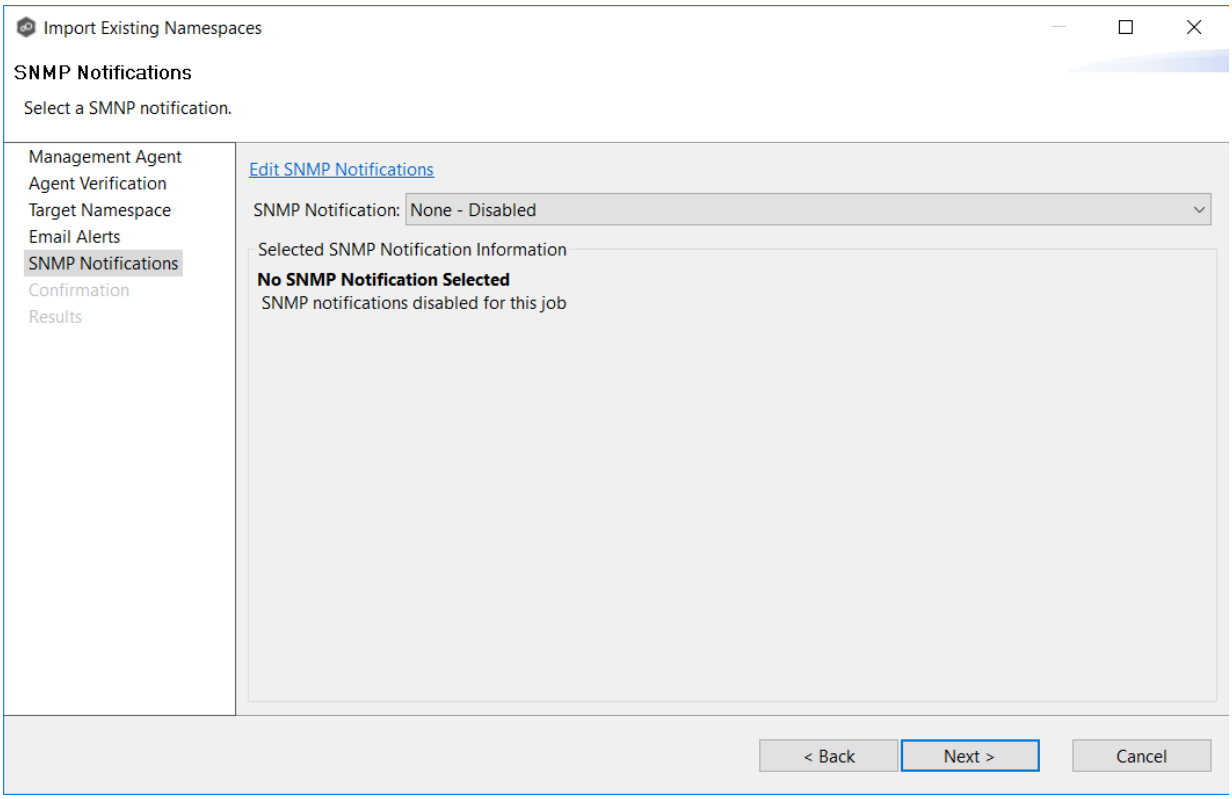

The **Confirmation** page appears.

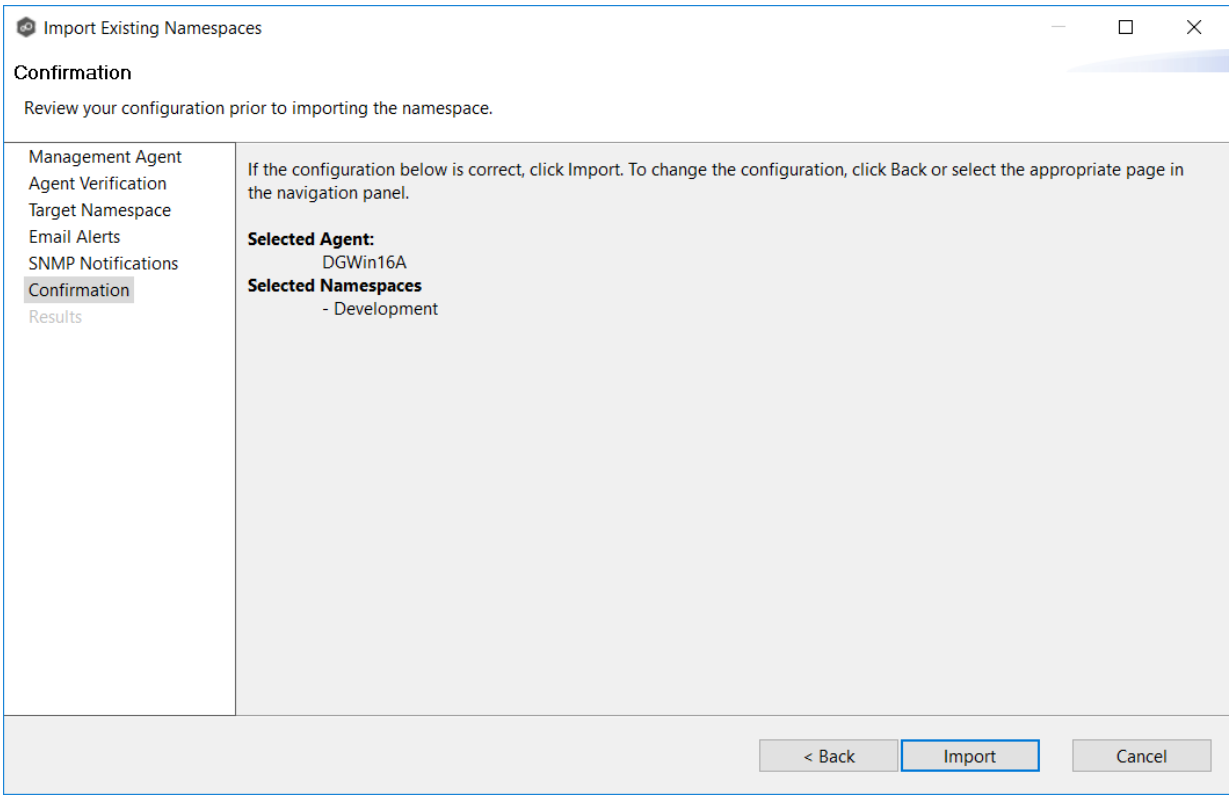

11. Review the configuration.

12. If you need to modify the job configuration after reviewing it, click **Back** until you reach the appropriate page and make your changes.

13. Once you are satisfied with the job configuration, click **Import**.

    The **Results** page appears.

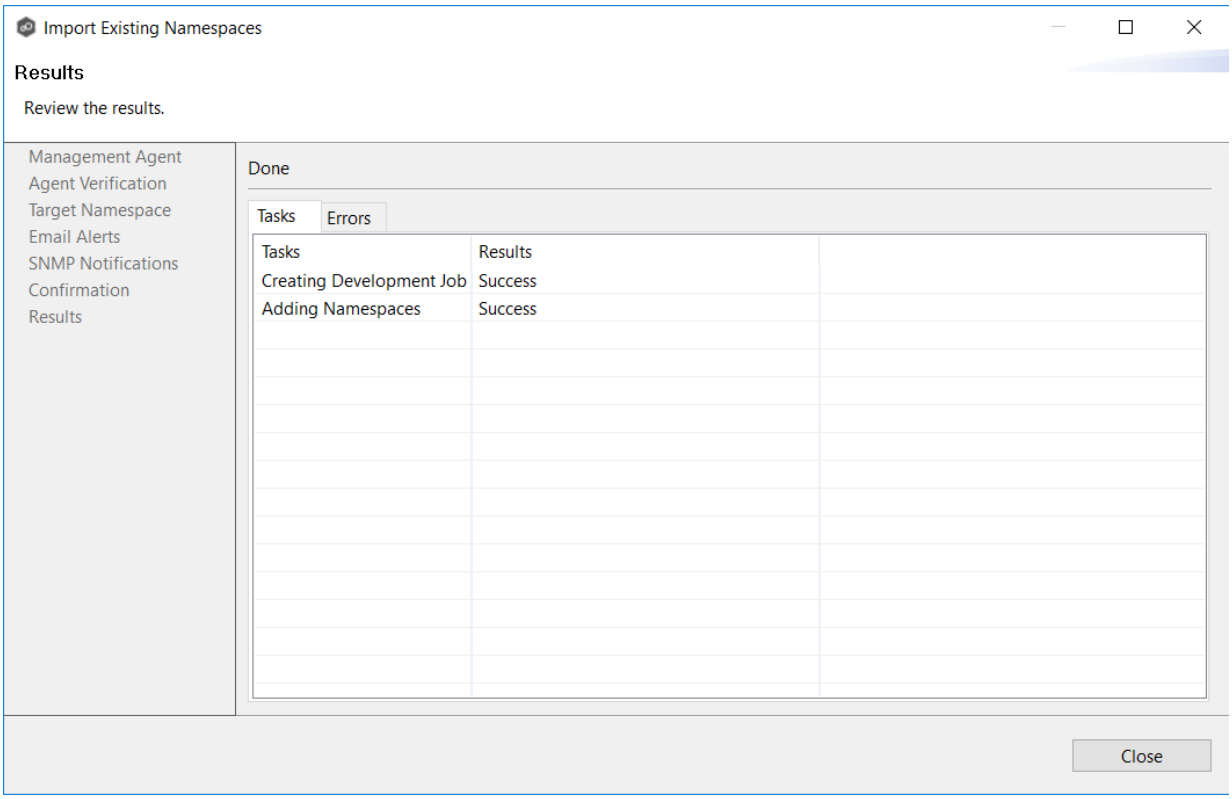

14. Review the results, and then click **Close**.

    A DFS-N Namespace job is created for each namespace you added. The new job(s) are displayed in the **Jobs** view. Runtime views for the jobs are also displayed. The jobs automatically start running. The namespaces can now be linked to File Collaboration and File Synchronization jobs.

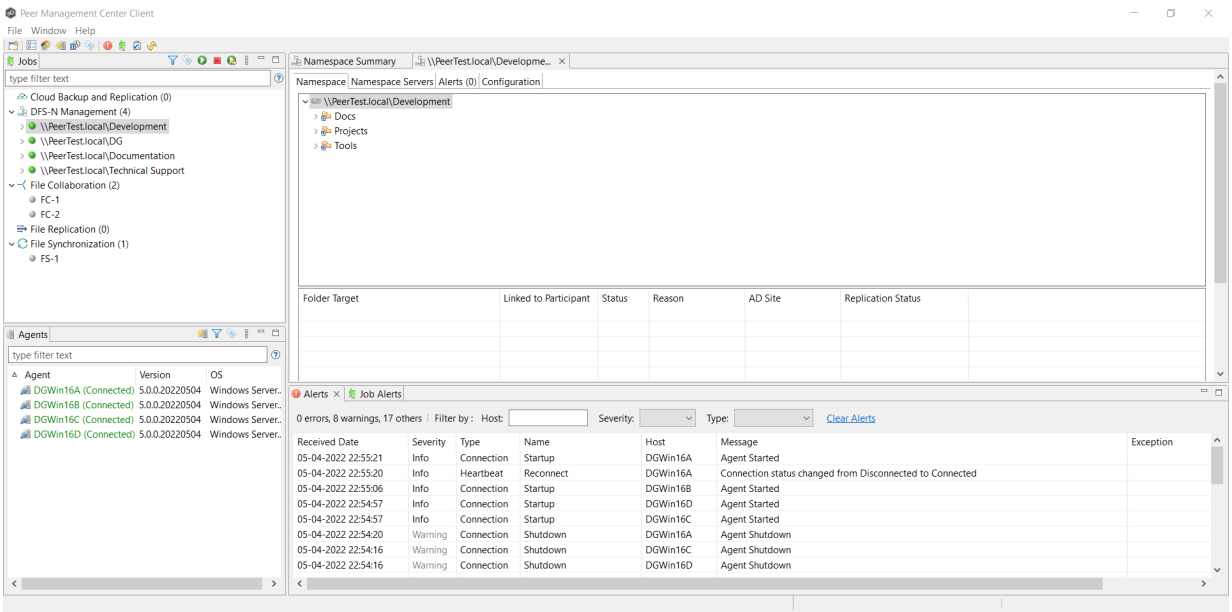

Running a DFS-N Management Job

This section describes:

- Starting a DFS-N Management Job
- Stopping a DFS-N Management Job

Starting a DFS-N Management Job

To manually start a DFS-N Management job:

1. Choose one of these options:
   - Right-click the job name in the **Jobs** view, and then select **Start**.

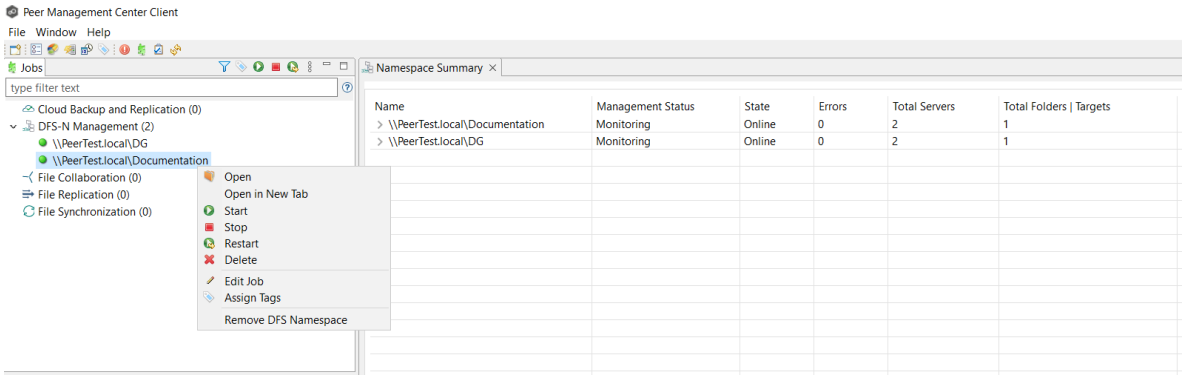

   - Open the job and then click the **Start/Stop** button in the bottom left corner of the job's **Namespace** tab in the DFS-N Management runtime view. You may need to scroll to the bottom of the tab to see the **Start/Stop** button.

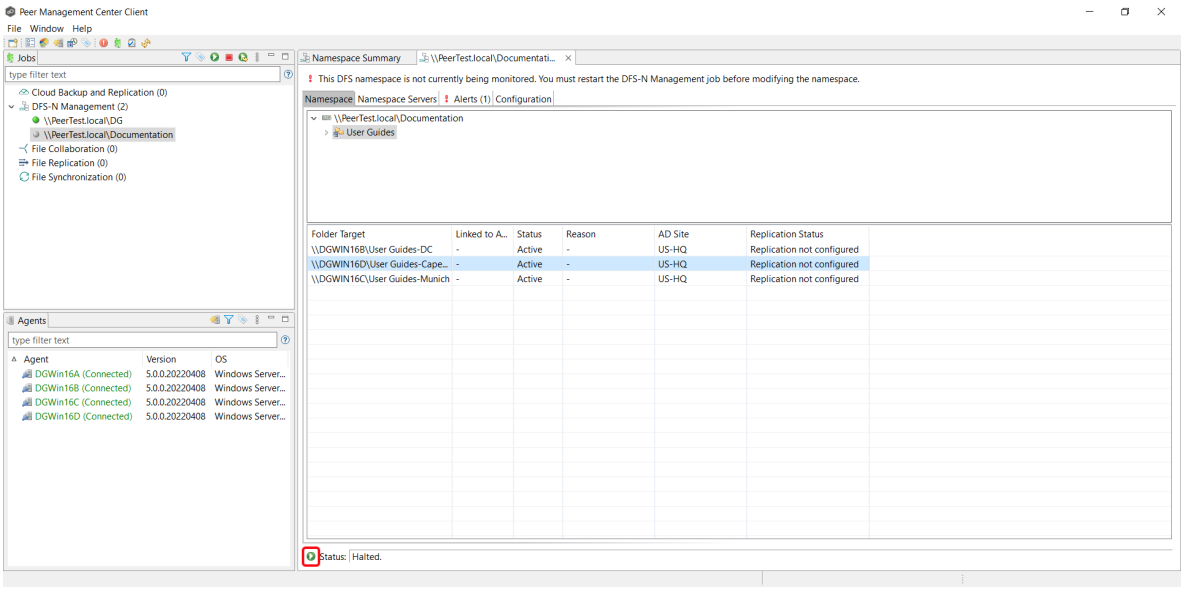

2. Click **Yes** in the confirmation dialog.

   After the job initialization has completed, the job will run. Once the job starts, the icon next to the job name in the **Jobs** view changes from gray to green.

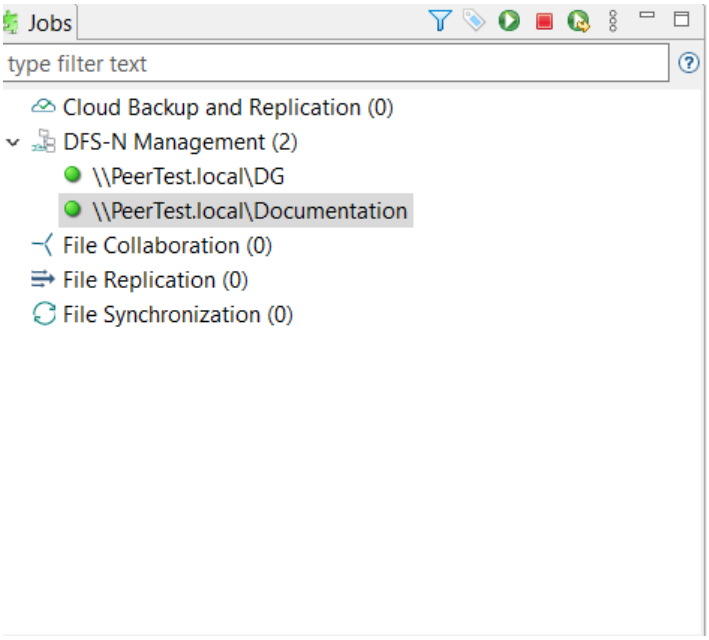

Stopping a DFS-N Management Job

You can stop a DFS-N Management job at any time. Note that you cannot edit a DFS-N Management job while it is stopped.

To stop a DFS-N Management job:

1. Right-click the job name in the **Jobs** view, and then choose **Stop** from the context menu.

   Or open the job and click the **Start/Stop** button in the bottom left corner of the job's **Namespace** tab in the runtime view.

2. Click **Yes** in the confirmation dialog.

   The icon next to the job name in the **Jobs** view changes from green to red.

Managing DFS Namespaces

This section describes:

- Adding a Namespace Server
- Adding a Namespace Folder
- Adding a Namespace Folder Target

Adding a Namespace Server

You can add a namespace server to a namespace.

To add a namespace server:

1. Double-click the name of a DFS-N Management job in the **Jobs** view or in the **Namespace Summary** view to open the runtime view for the job.

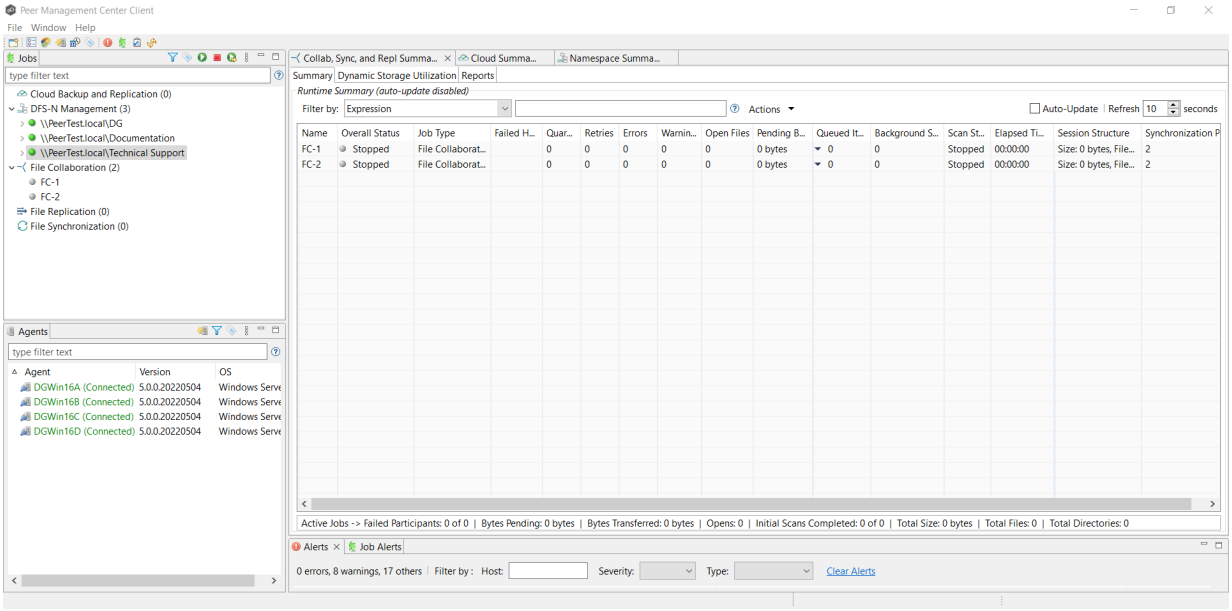

The runtime view for the job is displayed.

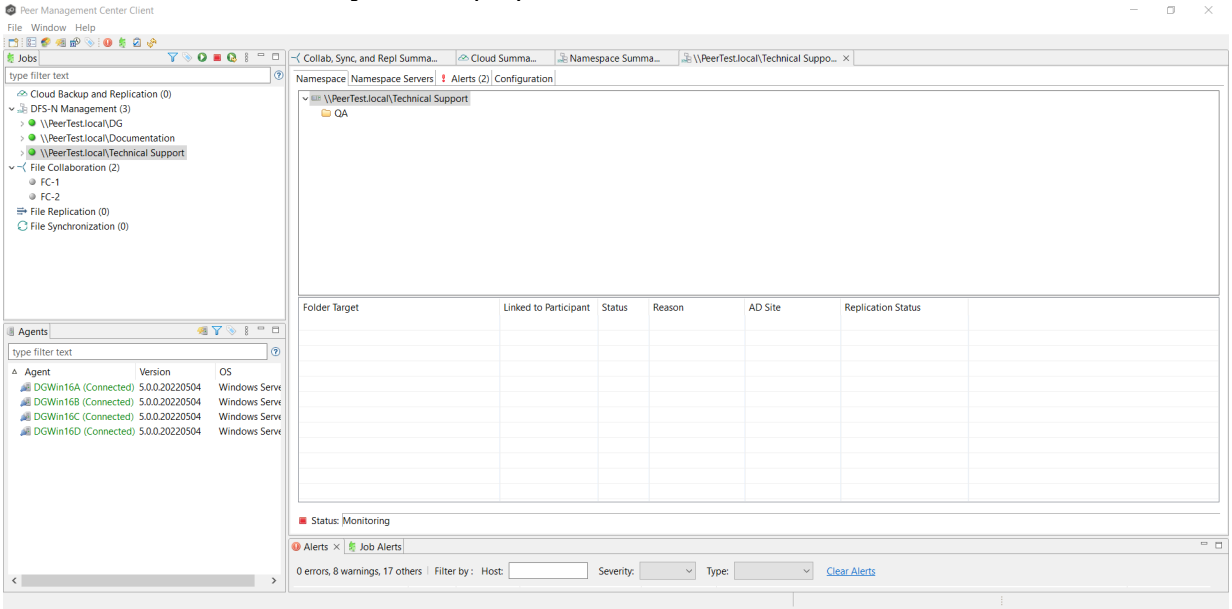

2. Click the **Namespace Servers** tab.

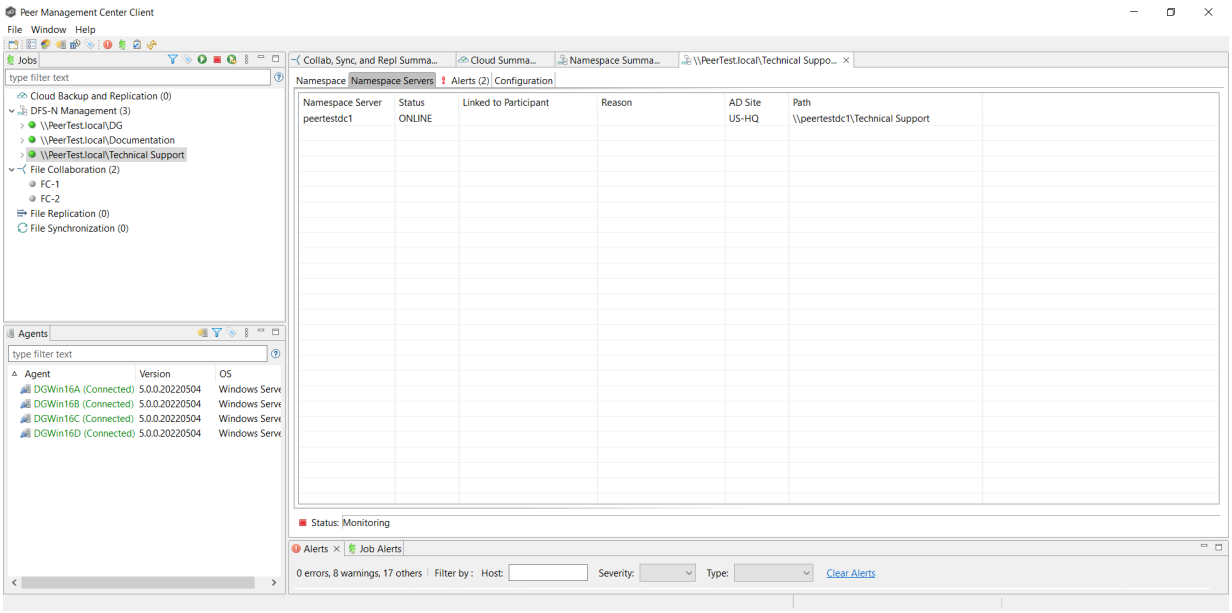

3. Right-click anywhere in the **Namespace Servers** tab, and then select **Add Servers**.

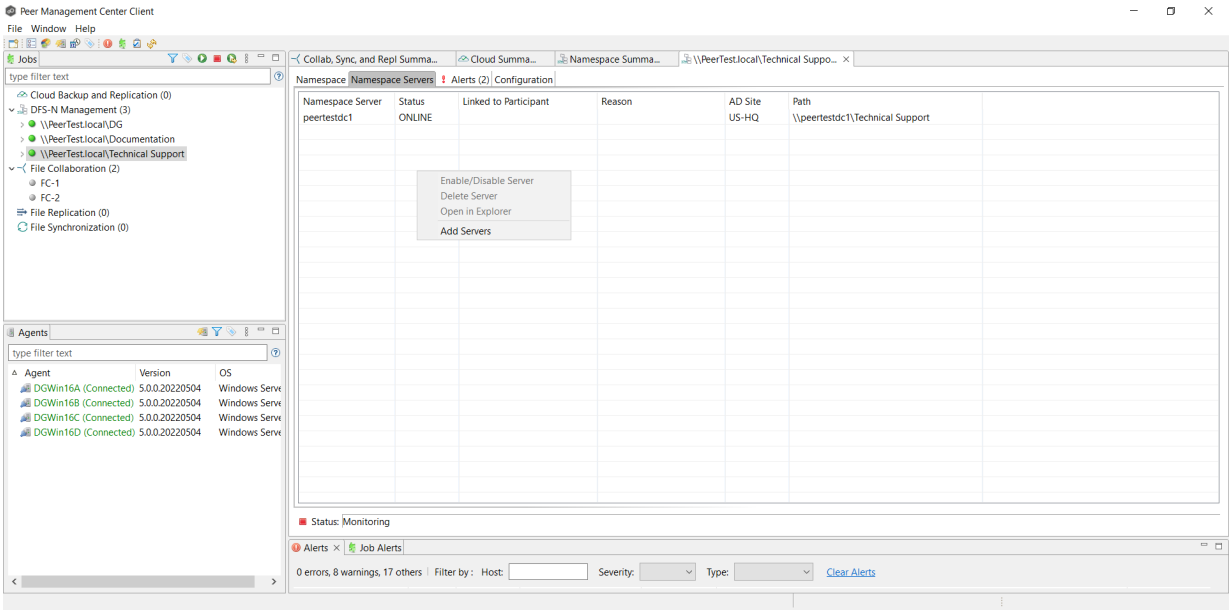

The **Add DFS Namespace Server** wizard appears.

4. Enter the fully qualified domain name (FQDN) of a namespace server in the **Server Name** field, and then click **Add**.

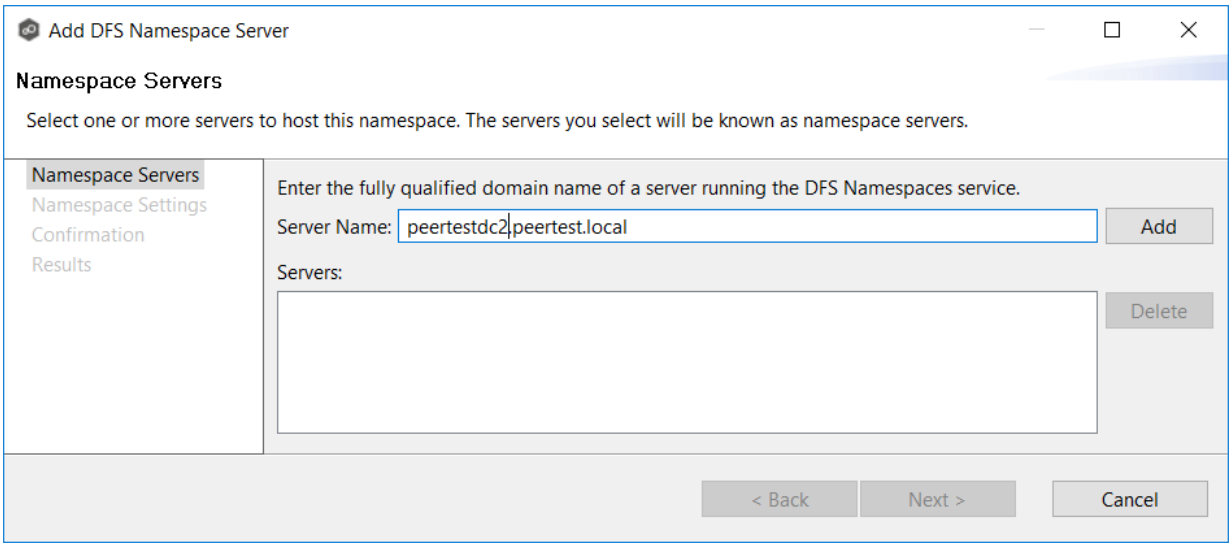

The server FQDN is listed in the area below.

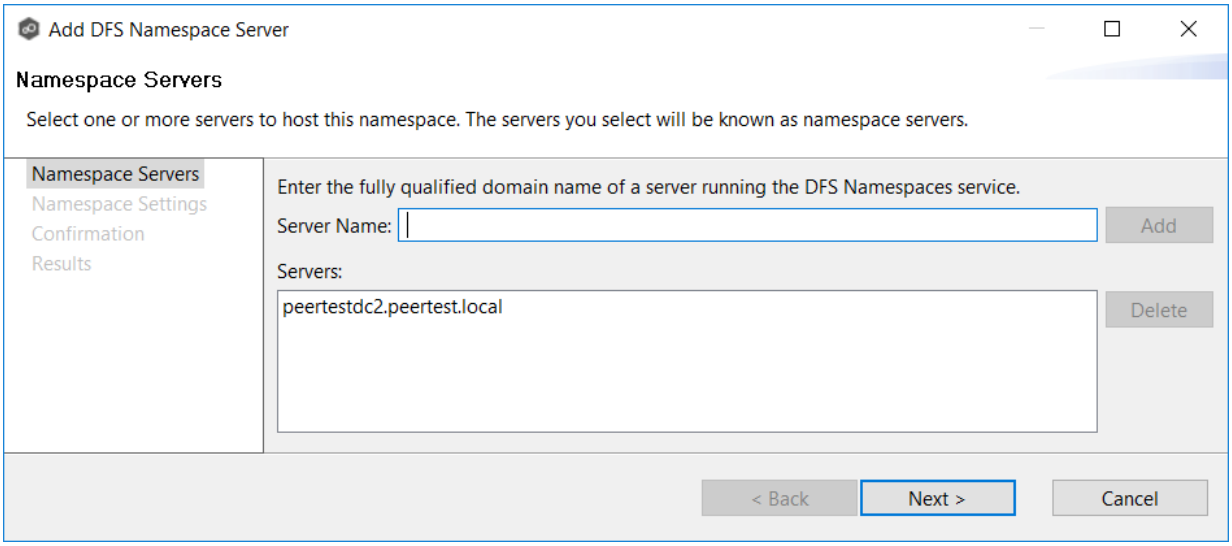

5. Add additional servers if desired.
6. Click **Next**.

   The **Namespace Settings** page is displayed.

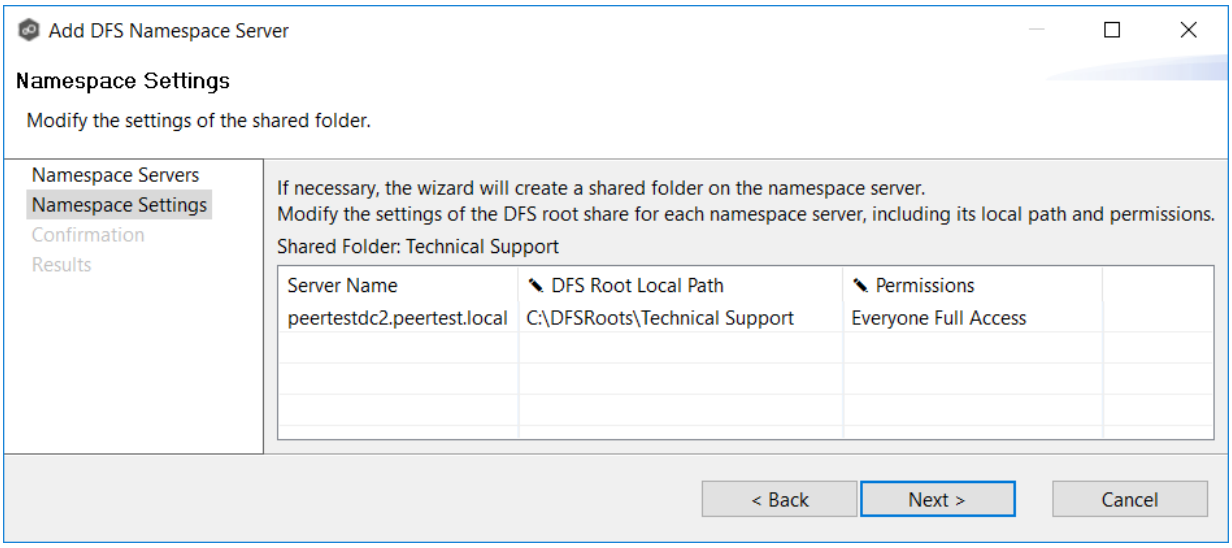

7. (Optional) Edit the namespace server settings: **DFS Root Share Path** and **Permissions**.

   - To modify the local path to the DFS root share for the namespace, type a new path in the **DFS Root Local Path** column. The default location of the DFS root share is under C:\DFSRoots\ and is specified in DFS-N Management Job Preferences.

   - To modify the access permissions, select a new set using the drop-down menu in the **Permissions** column.

8. Click **Next**.

   The **Confirmation** page is displayed.

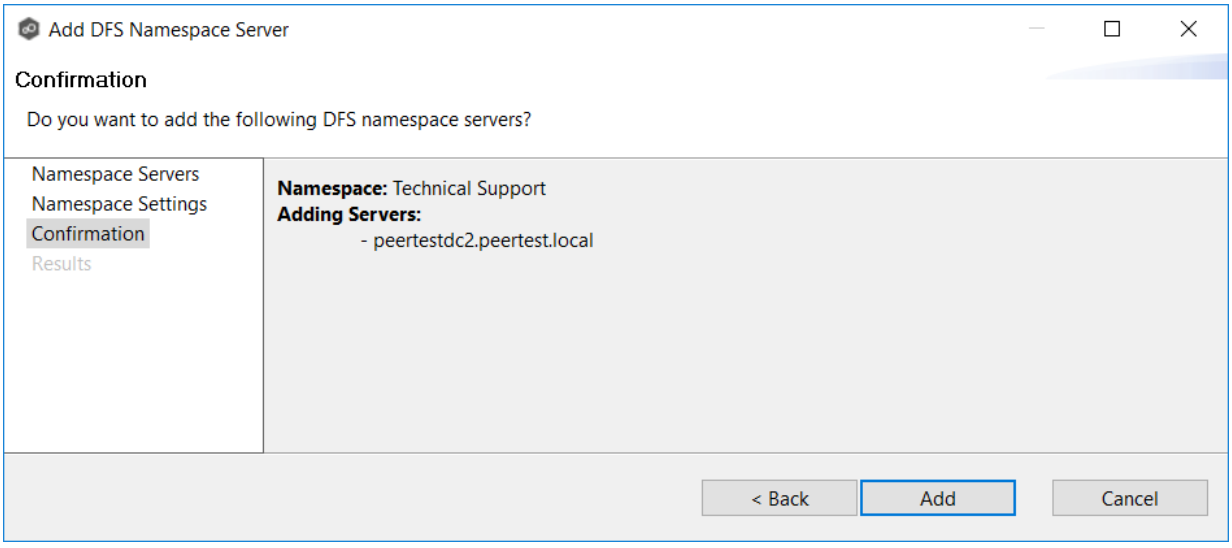

9. Review the namespace server configuration.

10. Click **Add** if the configuration is correct; otherwise, click **Back** and correct the configuration.

    The **Results** page is displayed.

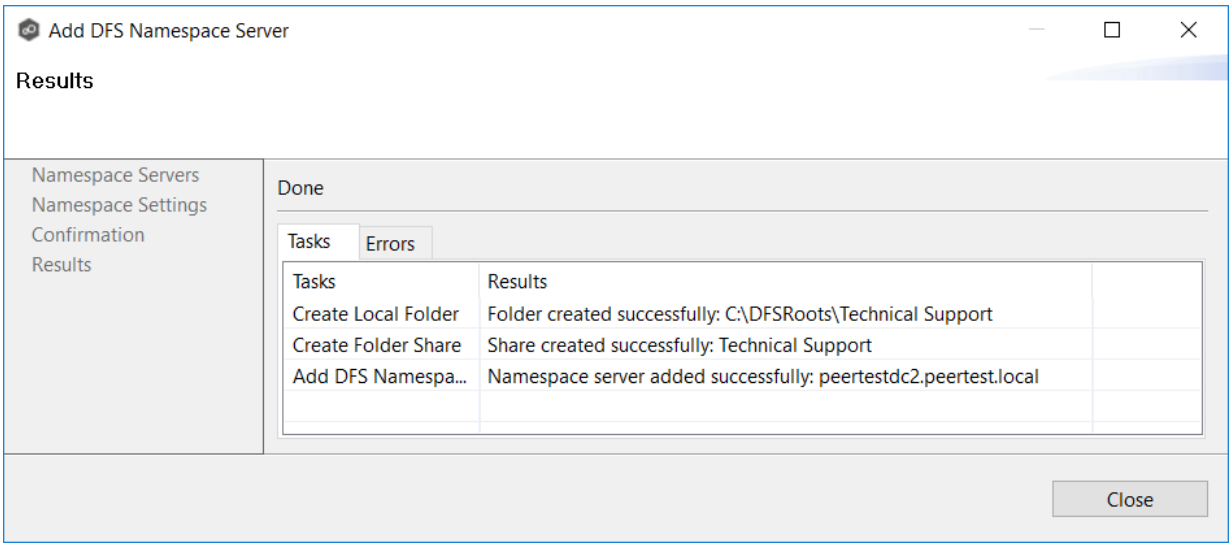

11. Click **OK**.

    The newly added server is listed in the **Namespace Servers** tab.

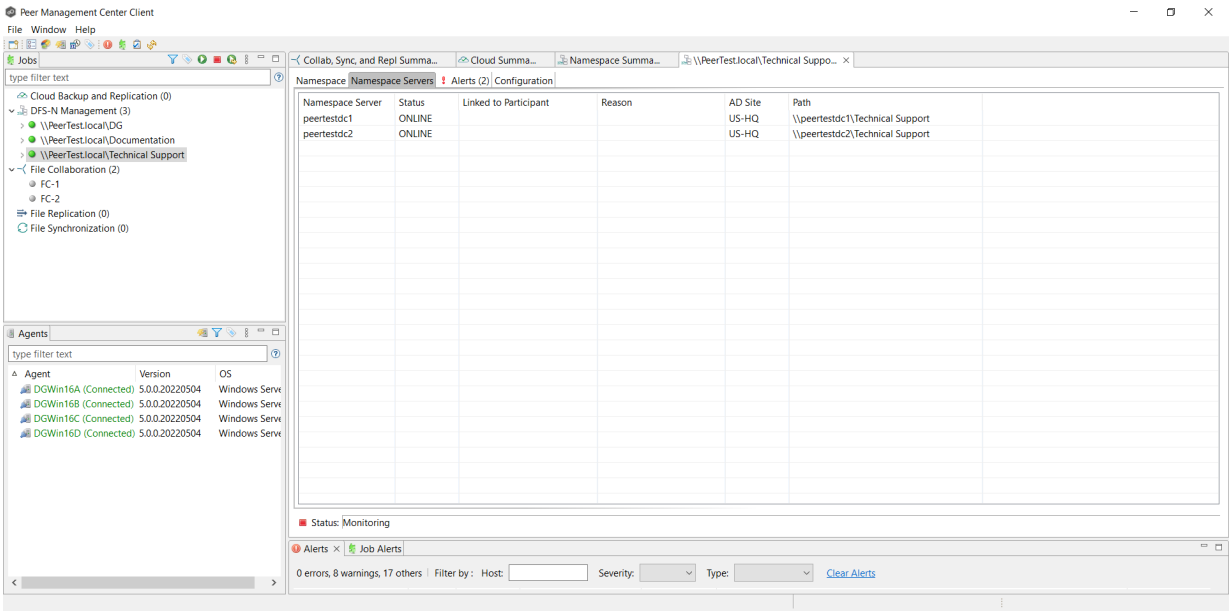

Adding a Namespace Folder

You can add a namespace folder to a namespace. At the same time, you can its folder targets or you can add folder targets later.

Note that PeerGFS does not currently support creation of nested namespace folders. In other words, you cannot add a namespace folder that is a subfolder of a namespace folder, although this can be done when using the Microsoft DFS Management Tool.

A DFS-N Namespace job must be running before you can edit it.

To add a namespace folder to a namespace:

1. Double-click the DFS-N Management job name in the **Jobs** view or in the Namespace Summary view to open the runtime view for the job.

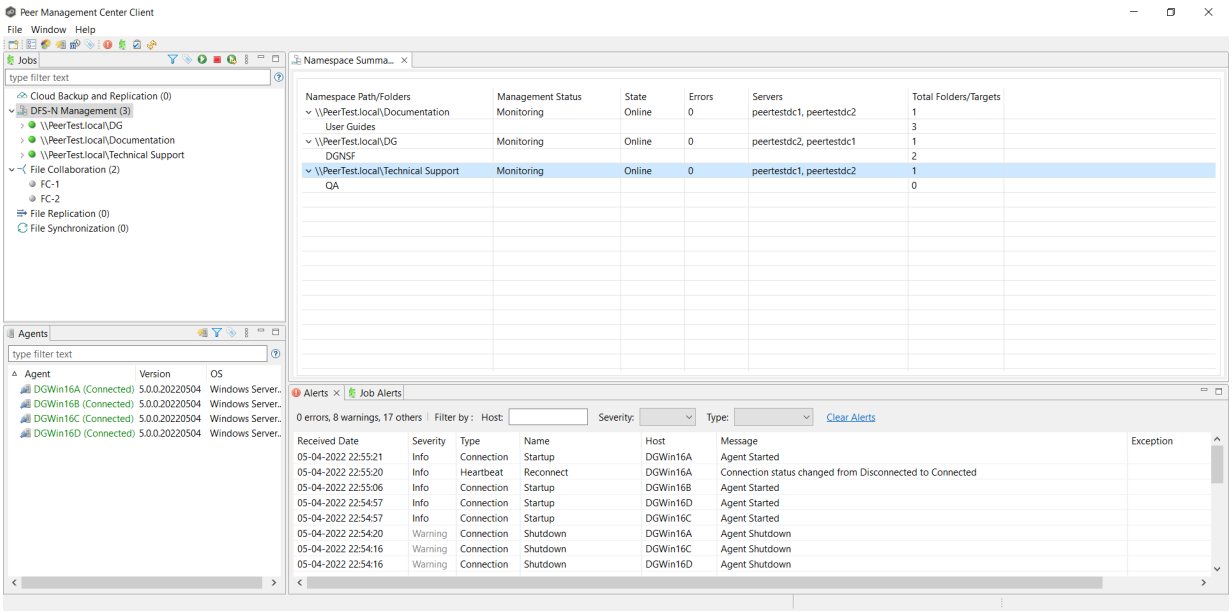

The runtime view for the job is displayed.

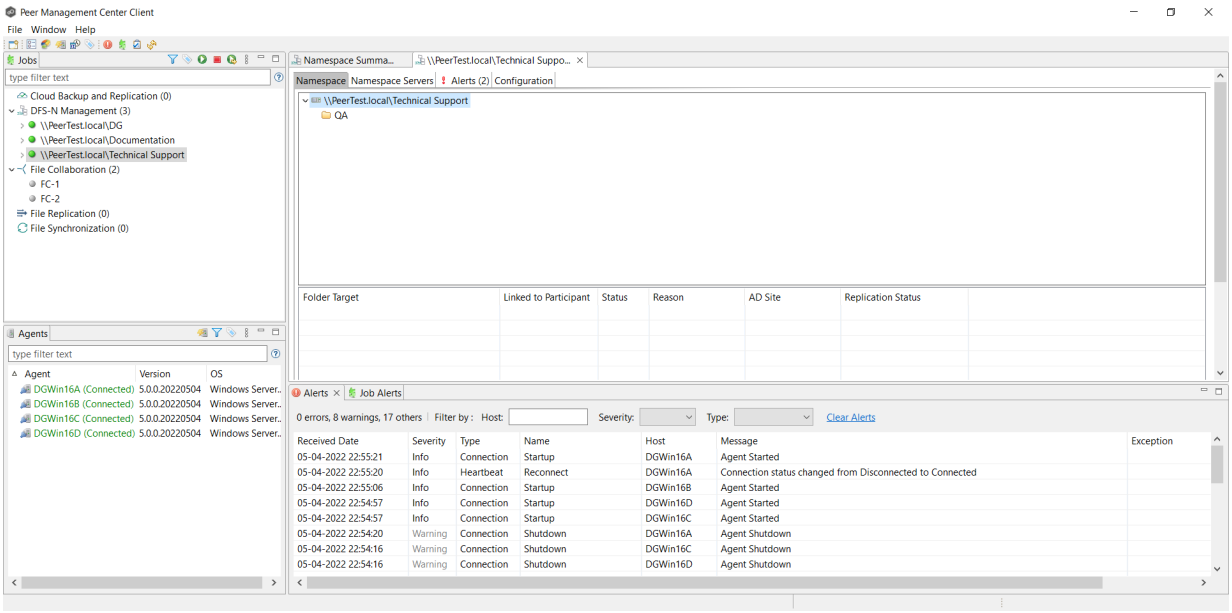

2. Right-click anywhere in the **Namespace** tab, and then select **Add Folder**.

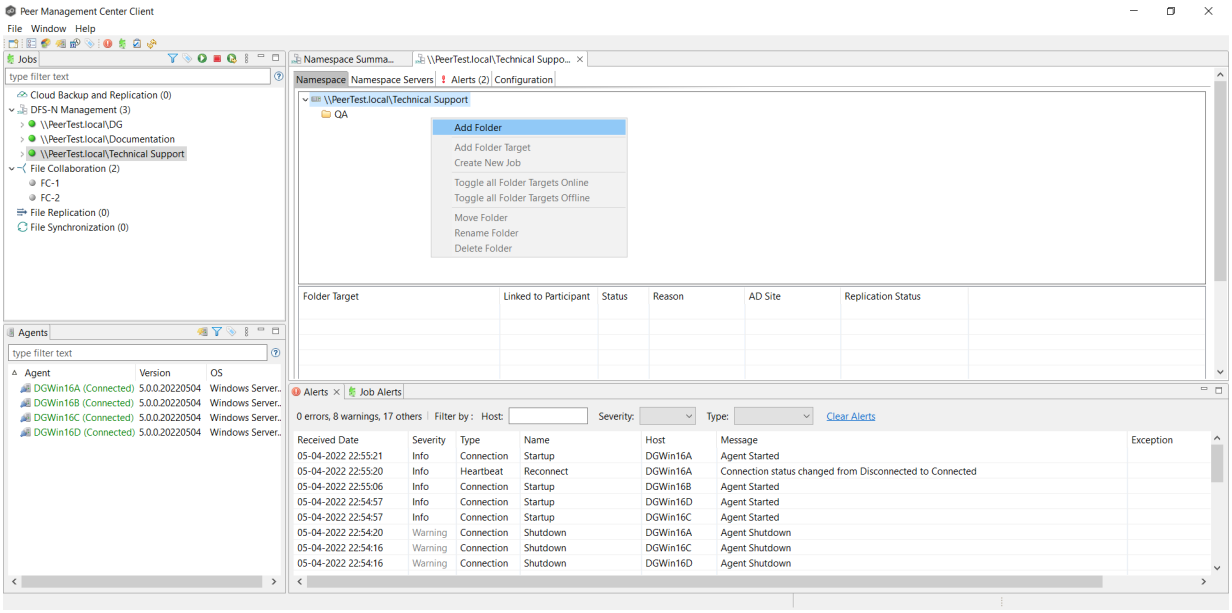

The **New Folder** wizard appears.

3. Enter a name for the namespace folder in the **Folder Name** field.

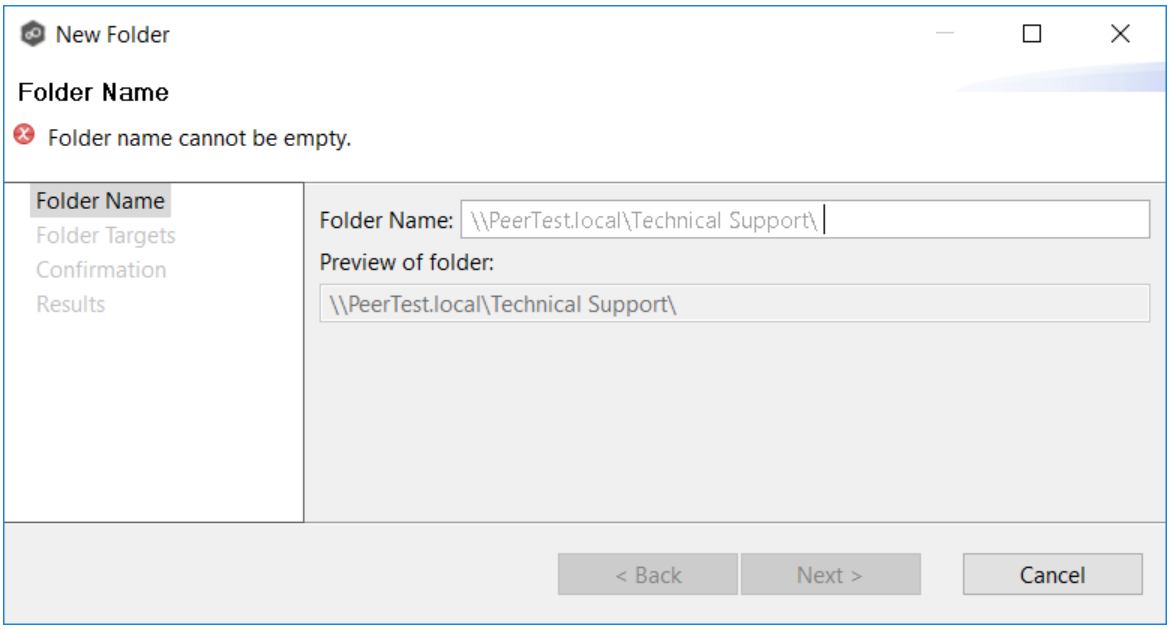

As you enter the folder name, a preview of the folder name and path appear below.

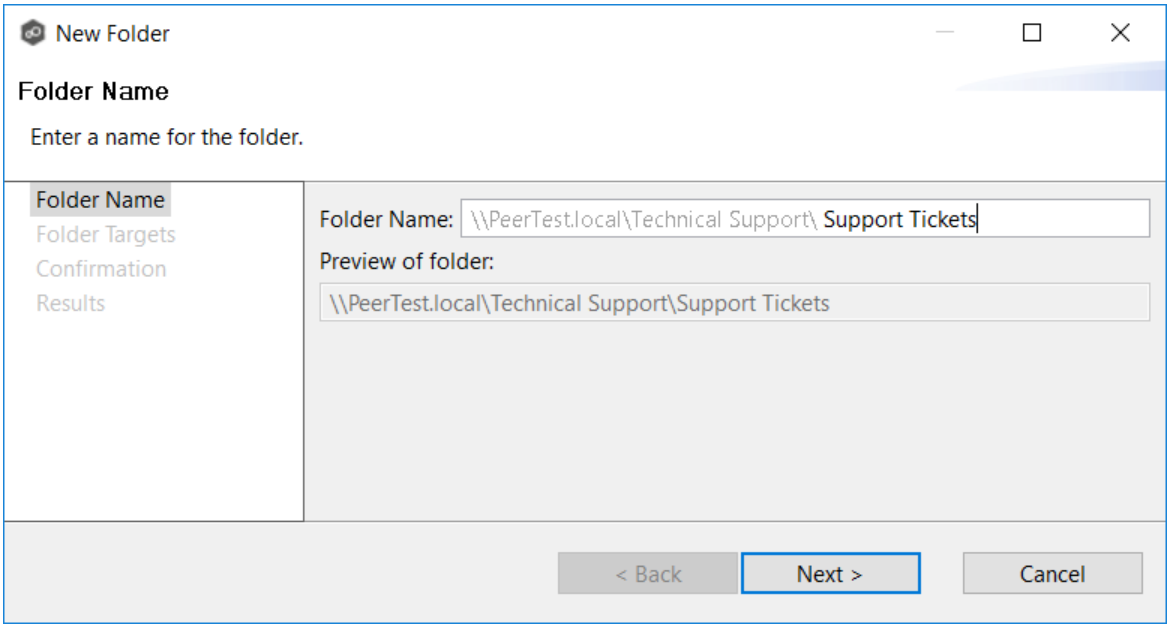

4. Click **Next**.

The **Folder Targets** page is displayed. It is optional to add folder targets for the namespace folder at this point. You can add them later if you wish. If you choose to add the folder targets now, they must already exist and be shared.

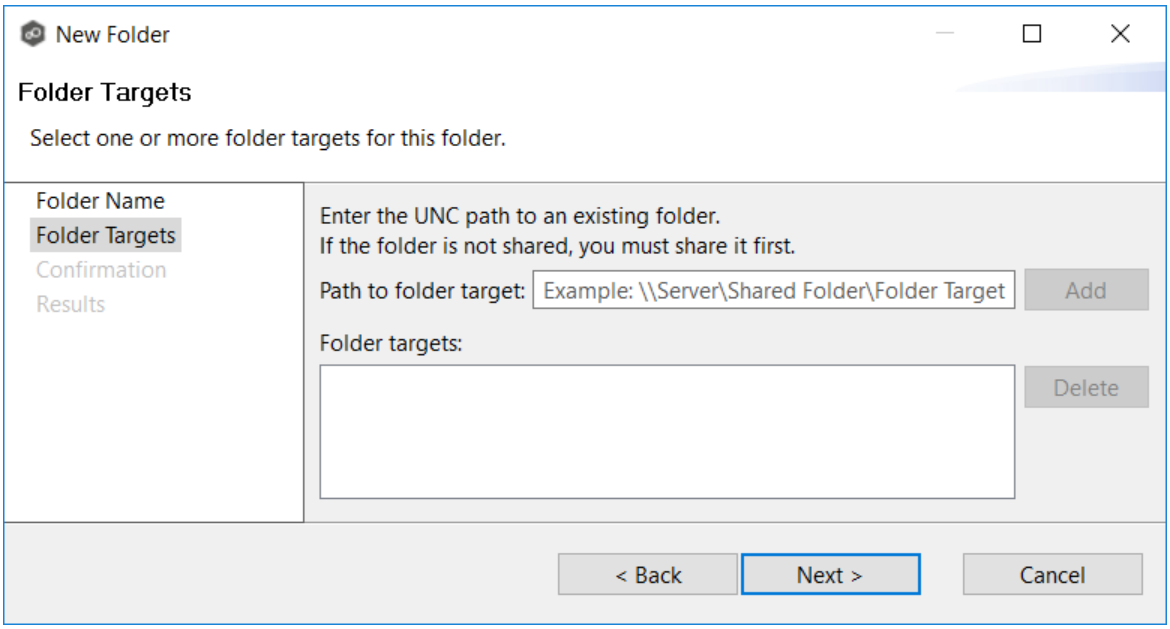

5. Click **Next** if you do not want to add folder targets at this point and continue with Step 9.

6. (Optional) Enter the UNC path to the shared folder you want to be a folder target.

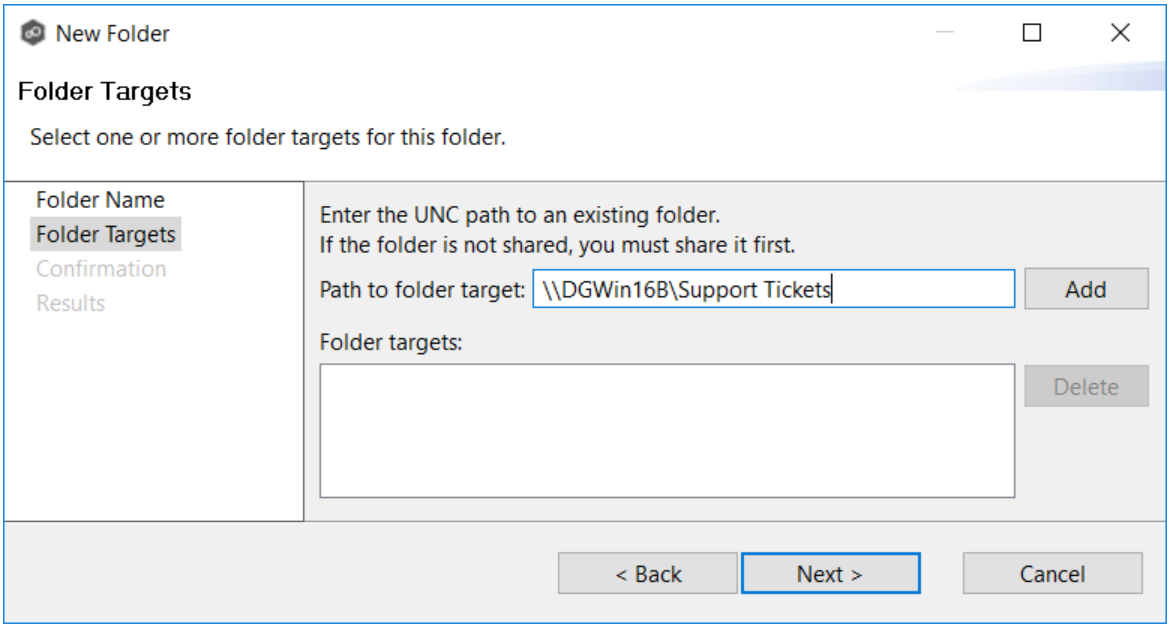

7. Click **Add**.

   The folder target is added to the **Folder targets** section.

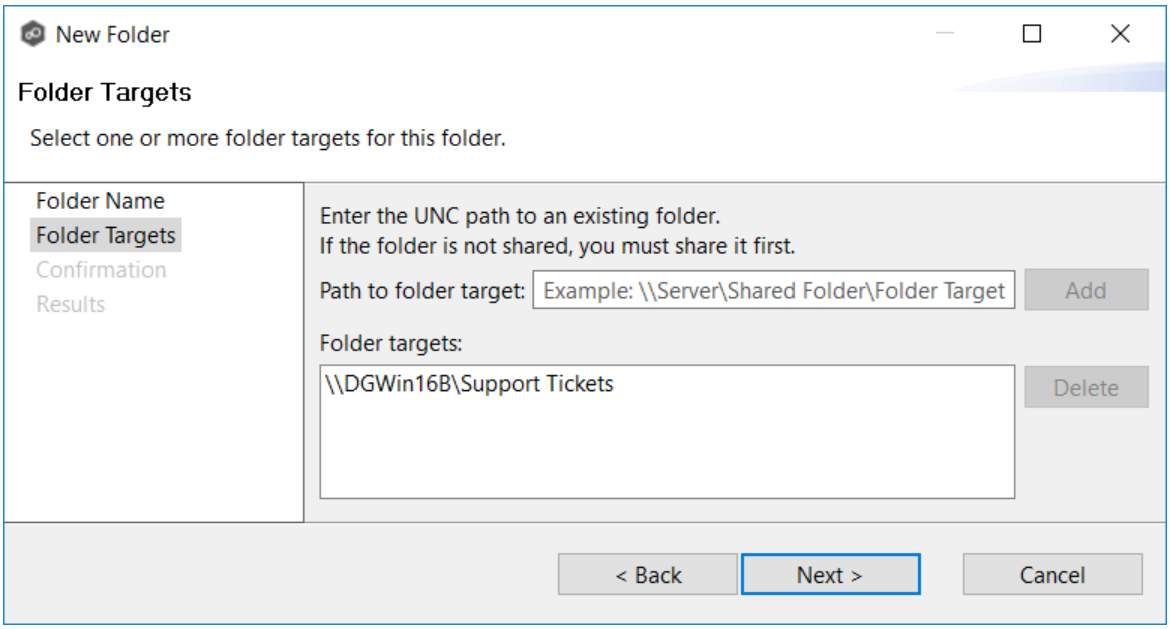

8. Repeat Steps 6-7 to add additional folder targets if desired.

9. Click **Next**.

   The **Confirmation** page is displayed.

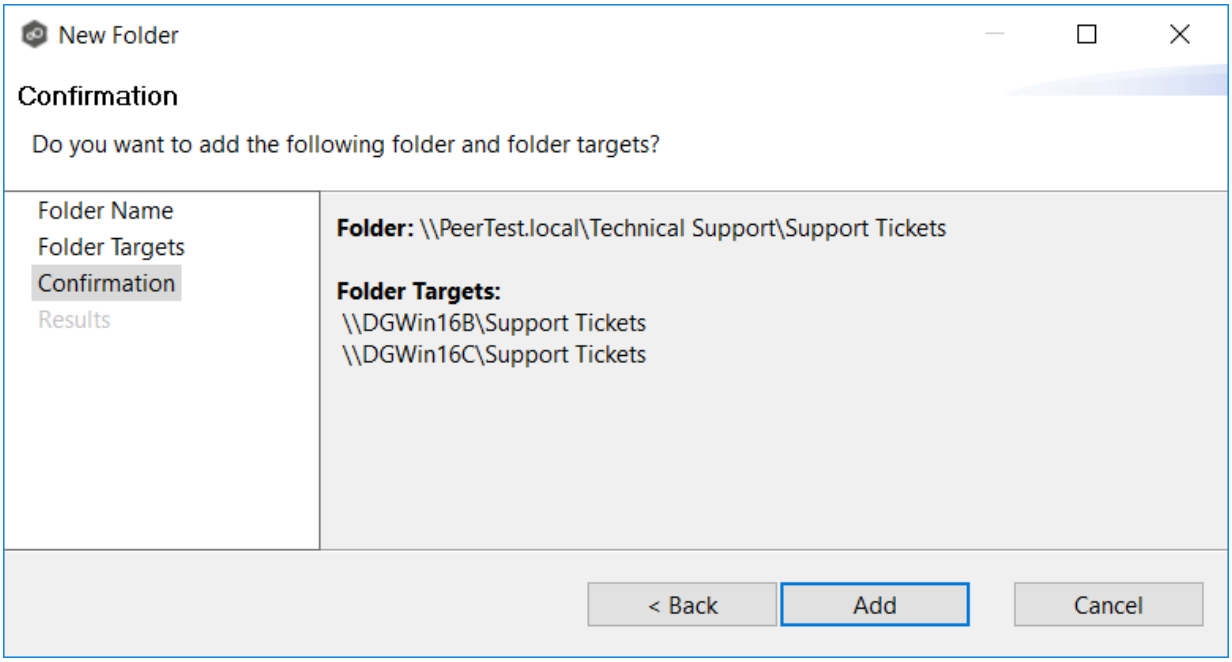

10. Review the folders and folder targets.

11. Click **Add** if the configuration is correct; otherwise, click **Back** and correct the configuration.

    The **Results** page is displayed.

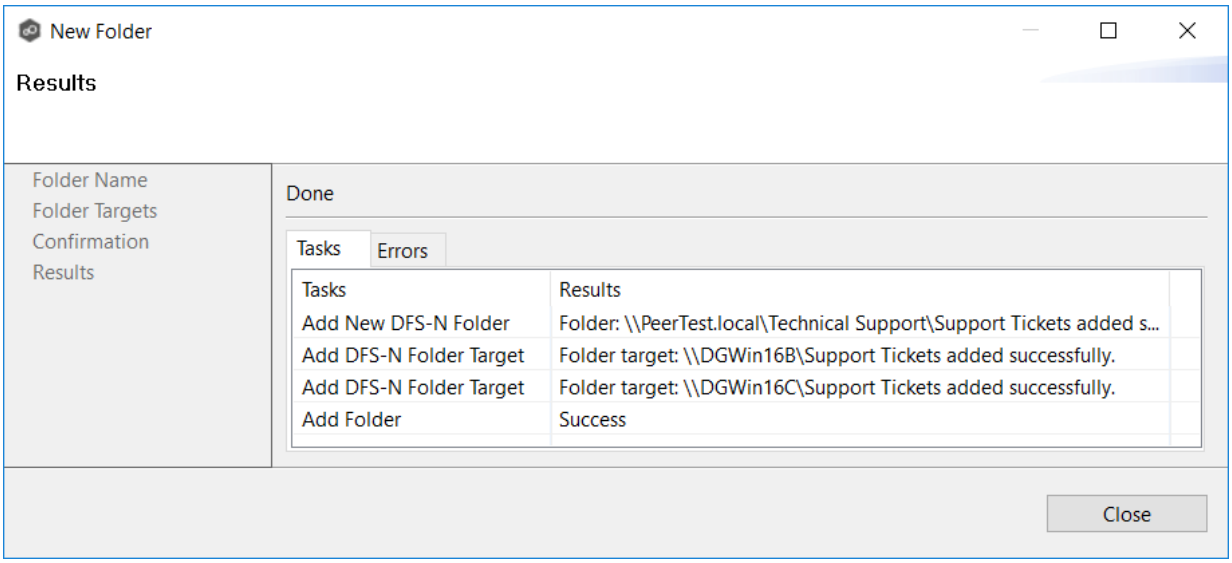

12. Click **Close**.

    The runtime view for the job is displayed. The newly added folders are listed in the job's **Namespace** tab.

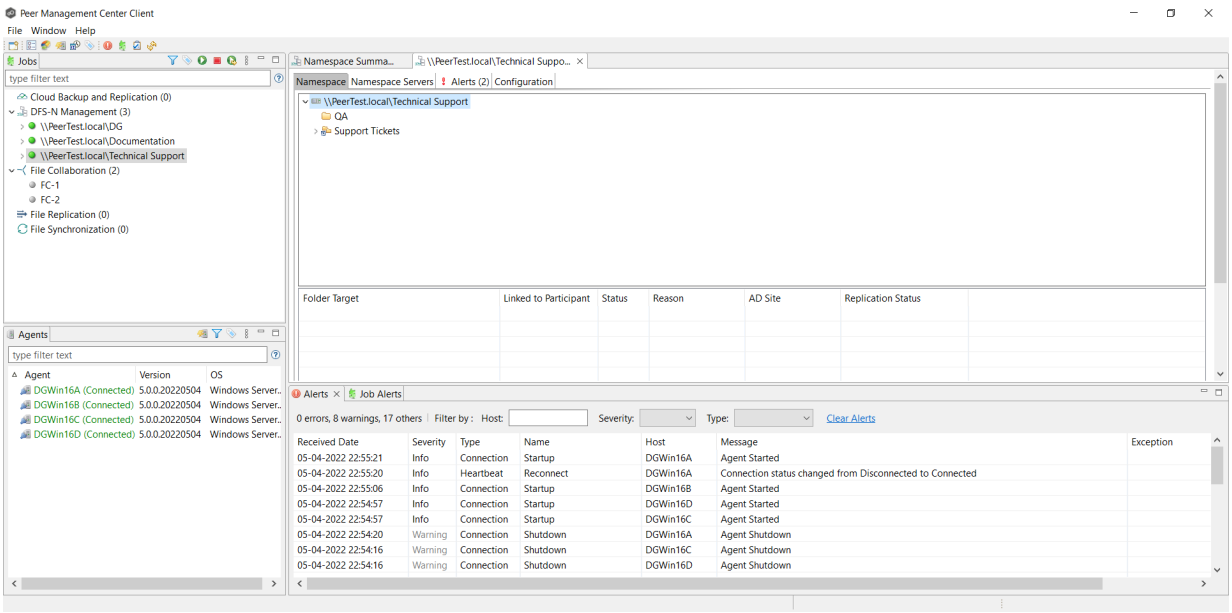

13. Click the folder you just added.

The newly added folder targets are listed in the **Folder Target** section of the tab. (Depending on how many namespace folders you have, you may need to scroll to view the **Folder Target** section.)

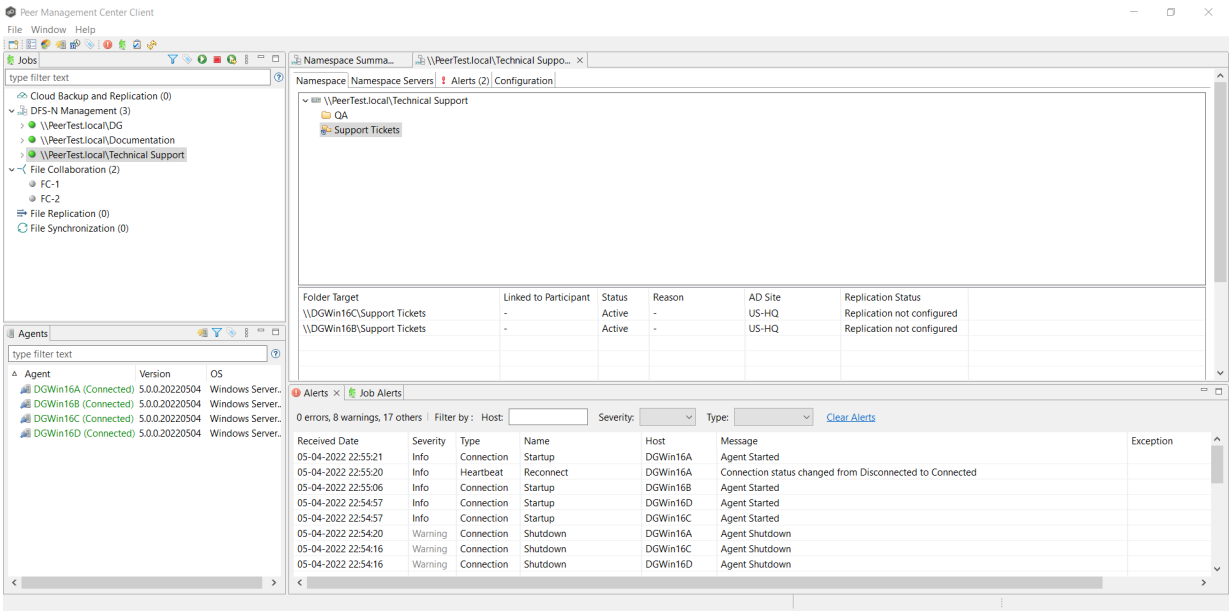

Adding a Namespace Folder Target

You can add a folder target for a namespace folder.

Note: A DFS-N Namespace job must be running before you can edit it.

To add a folder target for a namespace folder:

1. Double-click the job name in the **Jobs** view or the Namespace Summary view to open the runtime view for the job.

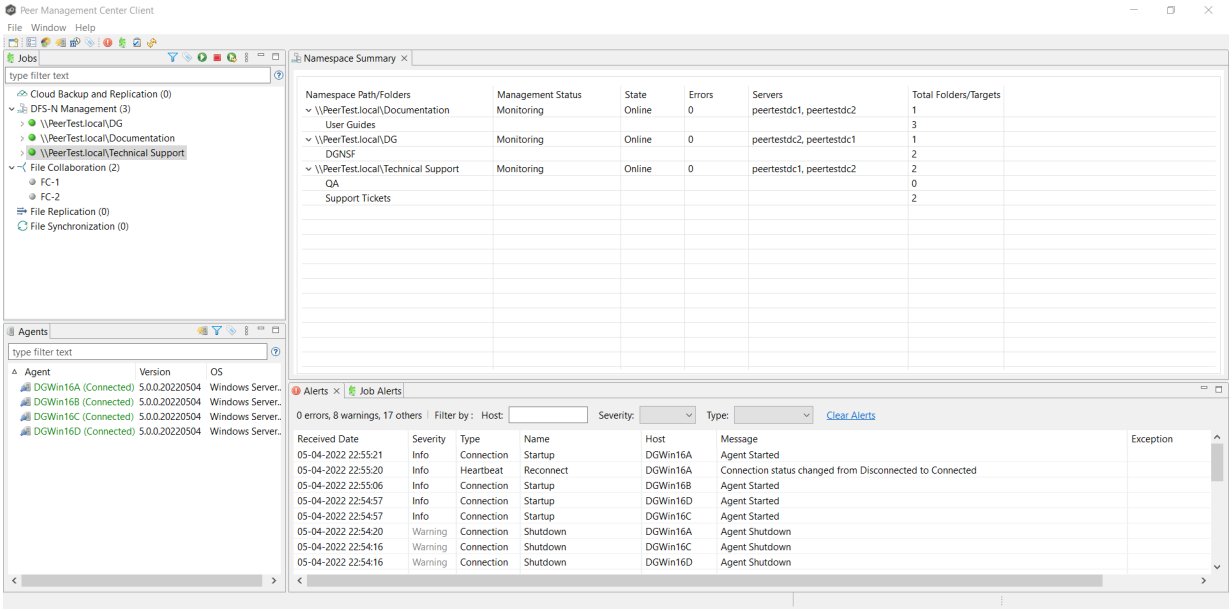

The runtime view for the job is displayed.

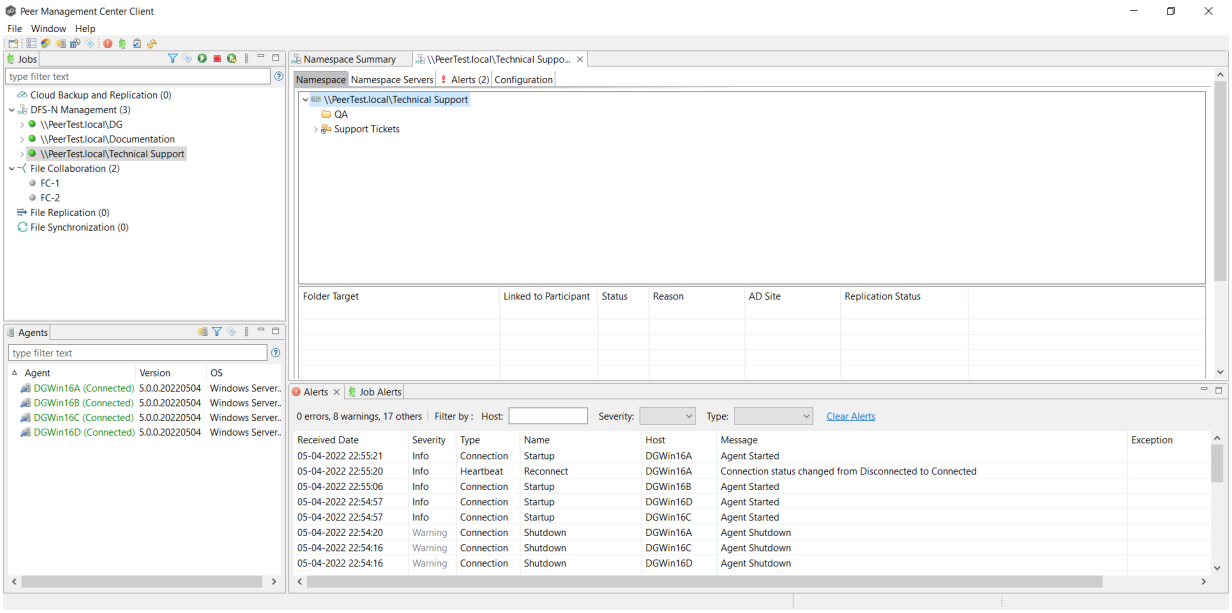

2. Right-click the folder you want to add a folder target to, and then select **Add Folder Target**.

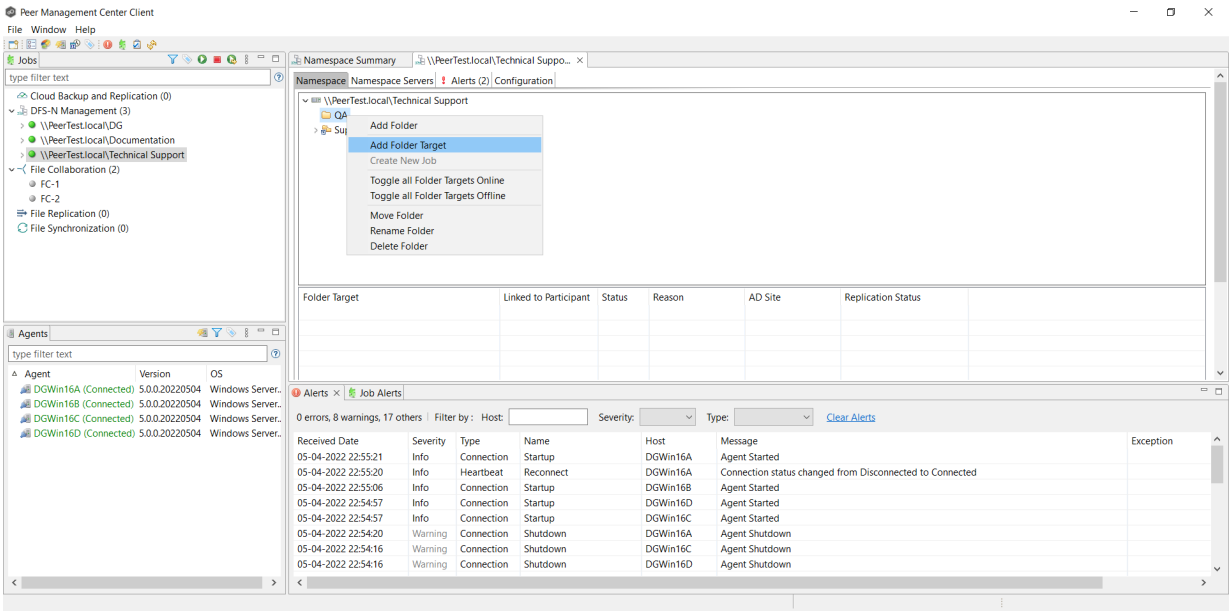

The **New Folder Target** wizard appears.

3. Enter the UNC path to a shared folder.

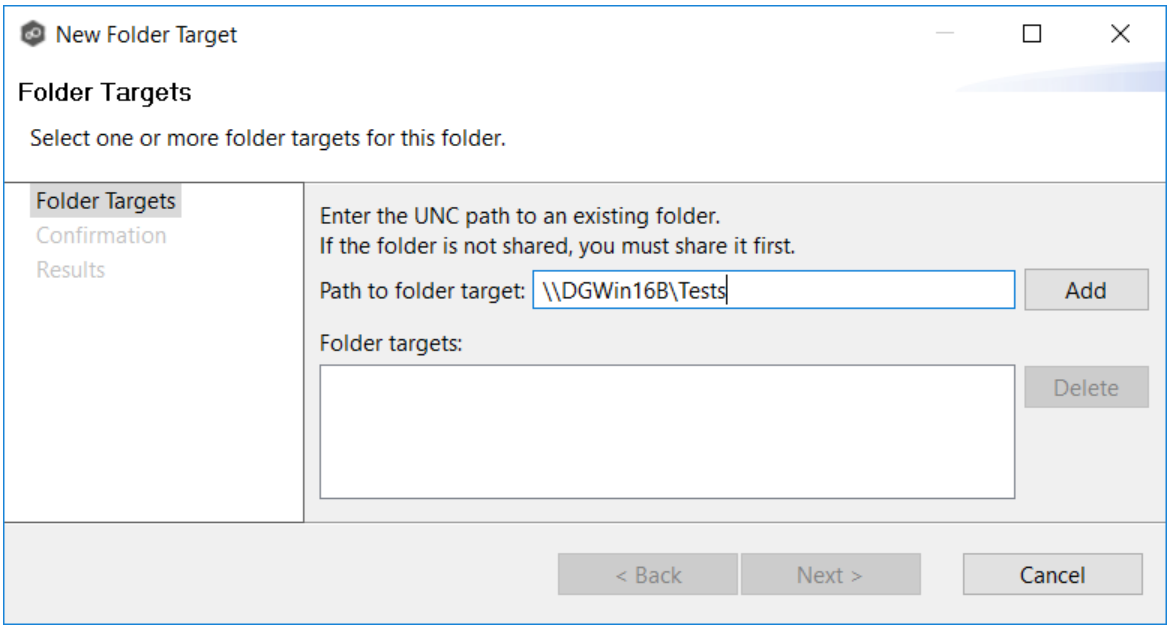

4. Click **Add**.

The folder target is added to the **Folder targets** section

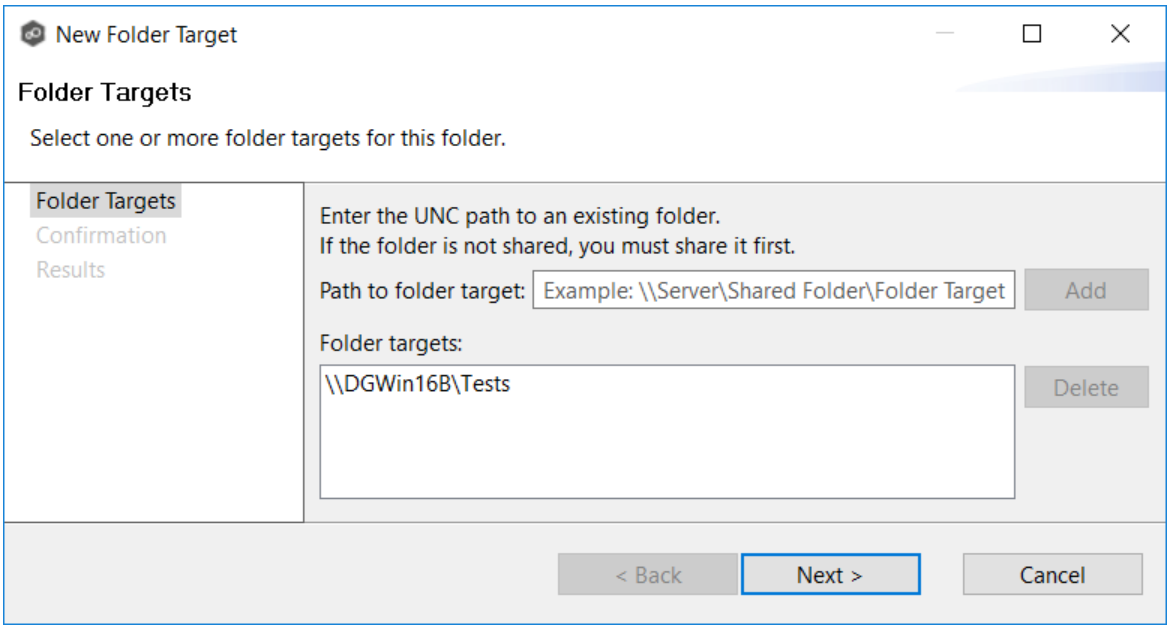

5. Repeat Steps 3-4 to add additional folder targets if desired.

6. Click **Next**.

   The **Confirmation** page is displayed.

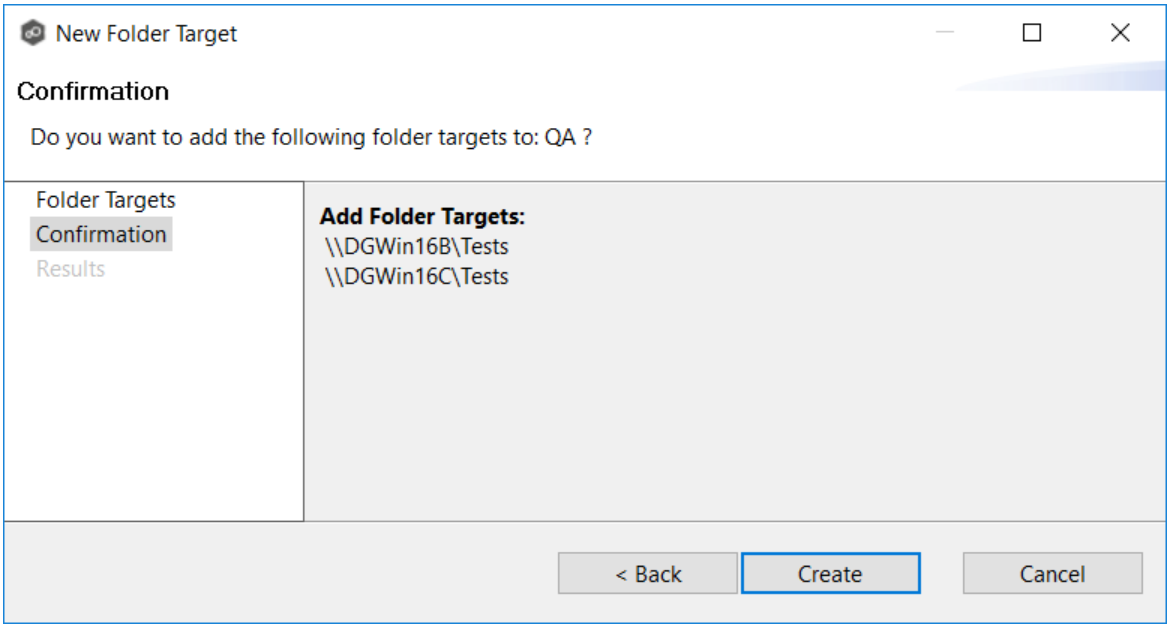

7. Review the folder targets.

8. Click **Create** if the configuration is correct; otherwise, click **Back** and correct the configuration.

The **Results** page is displayed.

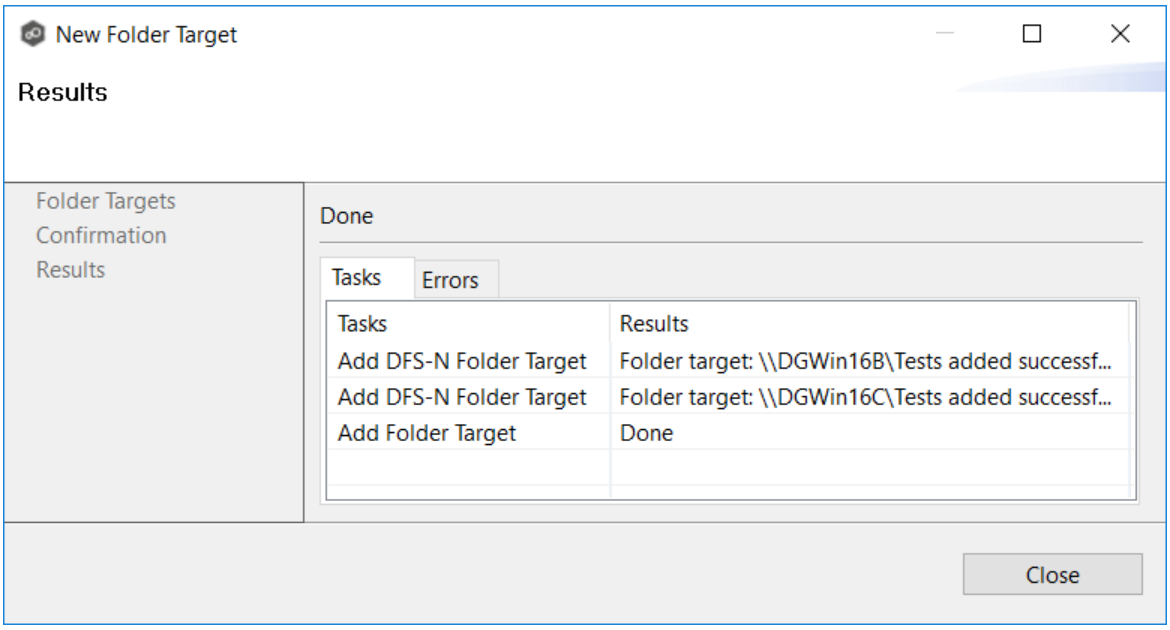

9. Click **Close**.

The runtime view for the job is displayed.

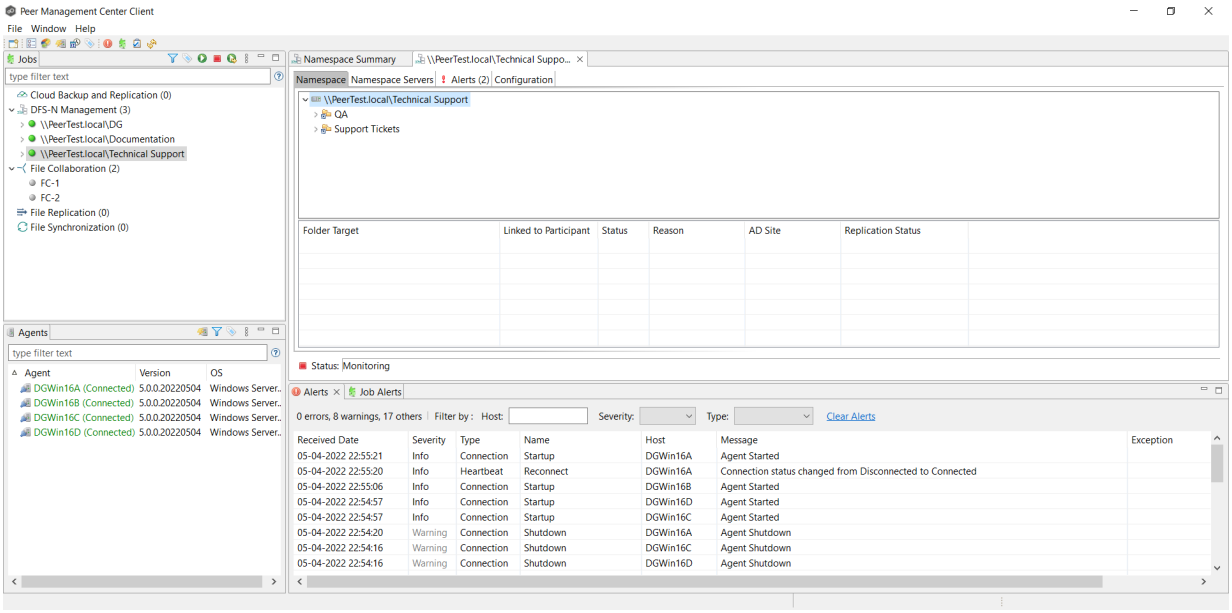

10. Click the folder you just added.

The newly added folder targets are listed in the **Folder Target** section of the job's **Namespace** tab. (Depending on how many namespace folders you have, you may need to scroll to view the **Folder Target** section.)

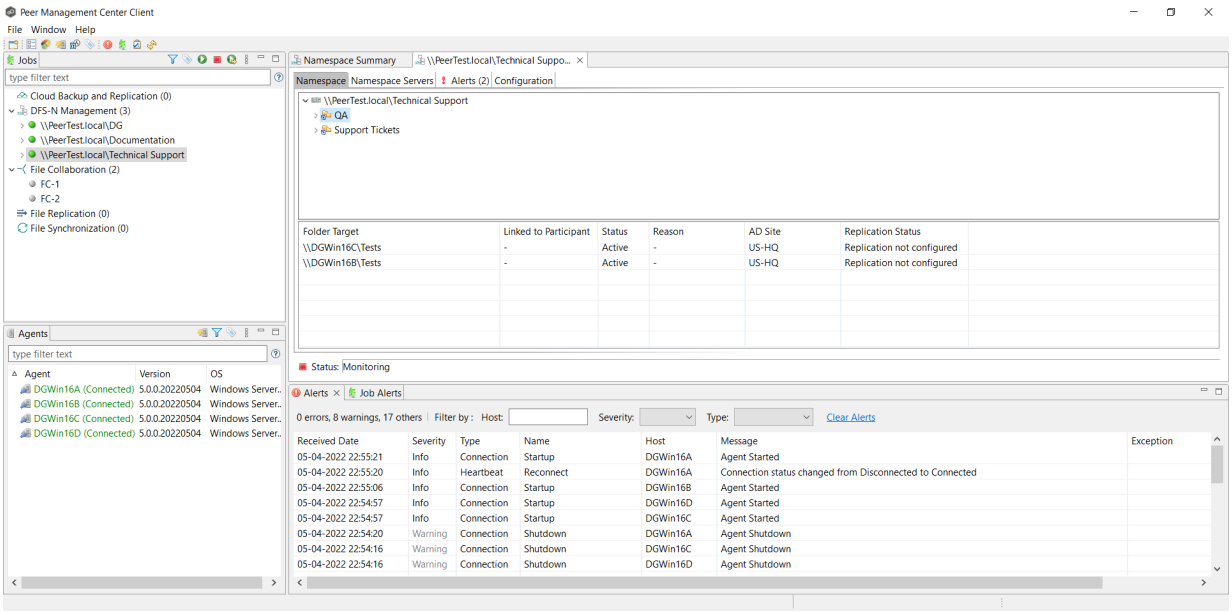

Linking a DFS Namespace to File Collaboration and File Synchronization Jobs

The primary benefit of using PeerGFS's ability to manage DFS namespaces is that PeerGFS can be configured to automatically disable and enable folder targets when they become unavailable, helping to control failover and failback. This is a manual process using Microsoft DFS Namespaces but PeerGFS can automatically disable or enable targets based on the state of a linked collaboration/synchronization job. This ensures a folder target is not available to users until a failed server has come back online and is in sync again.

To take advantage of the failover and failback capabilities, the File Collaboration or File Synchronization job must be linked to the DFS-N Management job that manages the namespace.

There are various ways to link a File Collaboration or File Synchronization job to a DFS-N Management job, including:

- If the File Collaboration or File Synchronization job does not yet exist, you can:
  - Create a File Collaboration or File Synchronization job and link it to a namespace while you are creating the job. When creating the job, you are given the option to create a new namespace, import one, or use an existing one. The DFS-N Management job is automatically created when using this method. See Step 8: DFS Management in Creating a File Collaboration job or in Creating a File Synchronization job for more information.

- Create one from the DFS namespace folder. See Create a File Collaboration or Synchronization Job from a Namespace Folder for step-by-step instructions.

- If the File Collaboration or File Synchronization job already exists, edit the job and link it to the DFS-N Management job. Use the DFS-N settings page in the Edit File Collaboration Job wizard and DFS-N settings page in the Edit File Synchronization Job wizard to link them. See Linking a Namespace with an Existing File Collaboration or Synchronization Job for step-by-step instructions.

Note: Currently, only File Collaboration and File Synchronization jobs can be linked to a DFS-N Management job.

Creating a File Collaboration or File Synchronization Job from a Namespace Folder

You can create a File Collaboration or File Synchronization job from a DFS namespace folder. These steps require that the DFS namespace has been already created and is being managed by Peer Management Center.

To create a File Collaboration or File Synchronization job from a namespace folder:

1. From the **Jobs** view, open the DFS-N Management job managing the namespace.

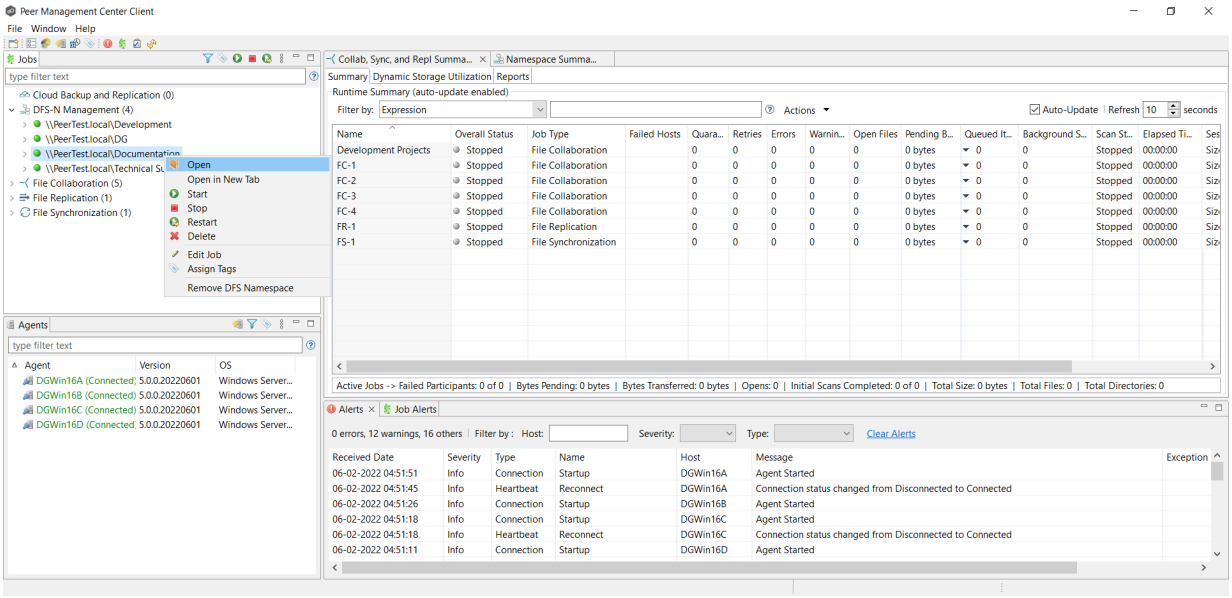

The DFS-N Management Job runtime view is displayed.

2. Open the **Namespace** tab if it is not already displayed.

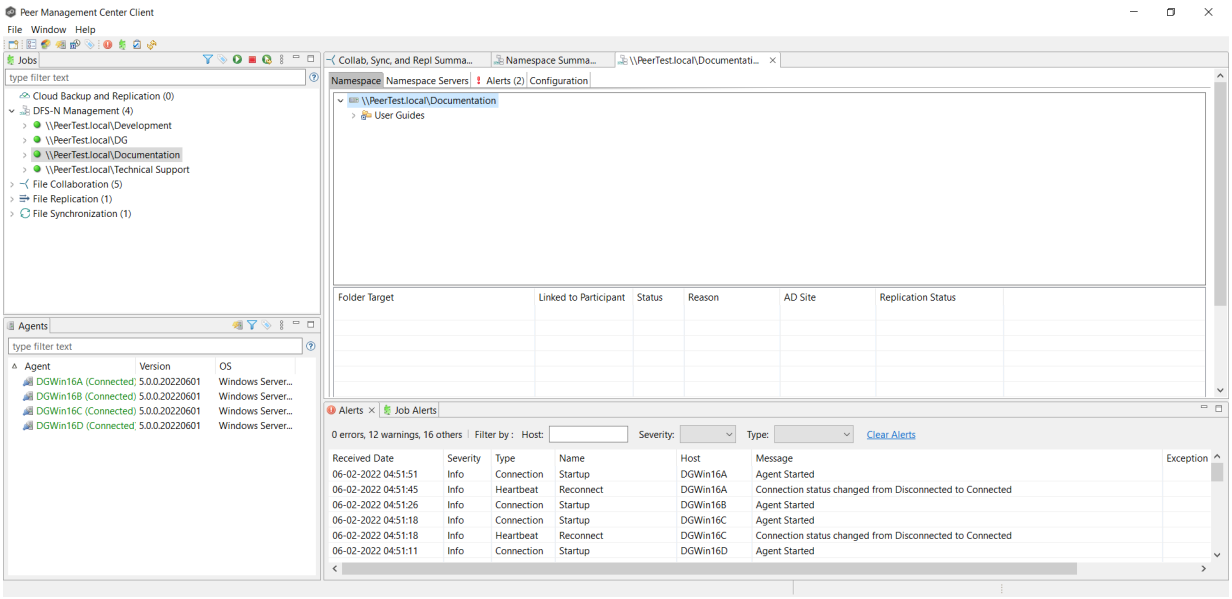

3. In the **Namespace** tab, right-click the desired namespace folder and select **Create New Job**.

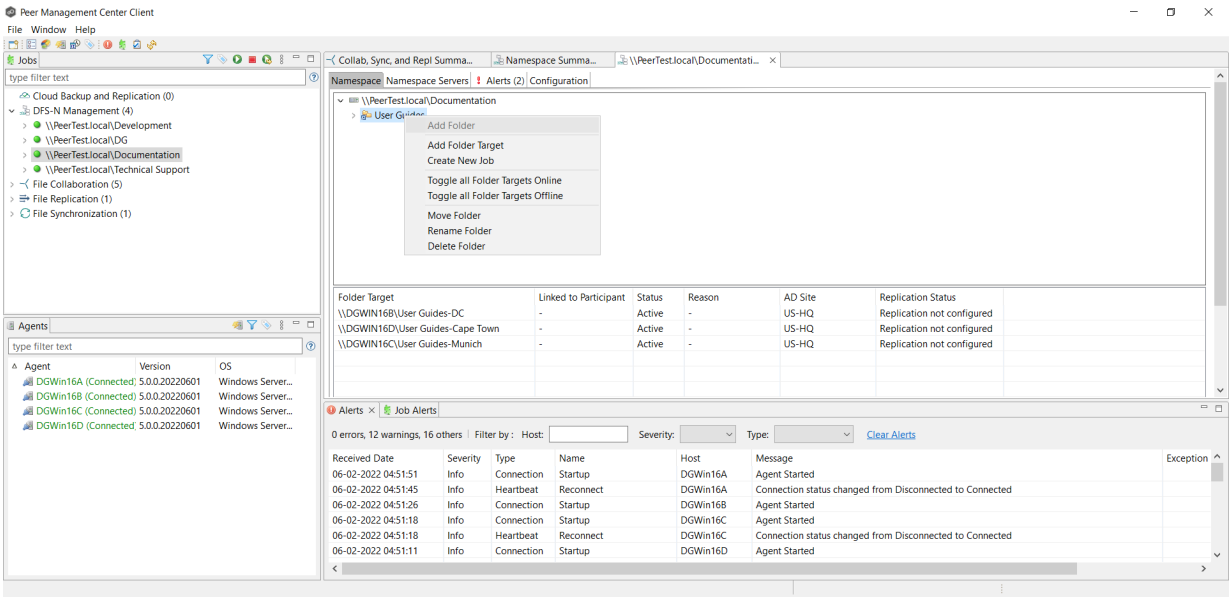

The **Create New Job** wizard displays a list of job types you can create: File Collaboration and File Synchronization. The other job types are not supported for use with DFS namespace management.

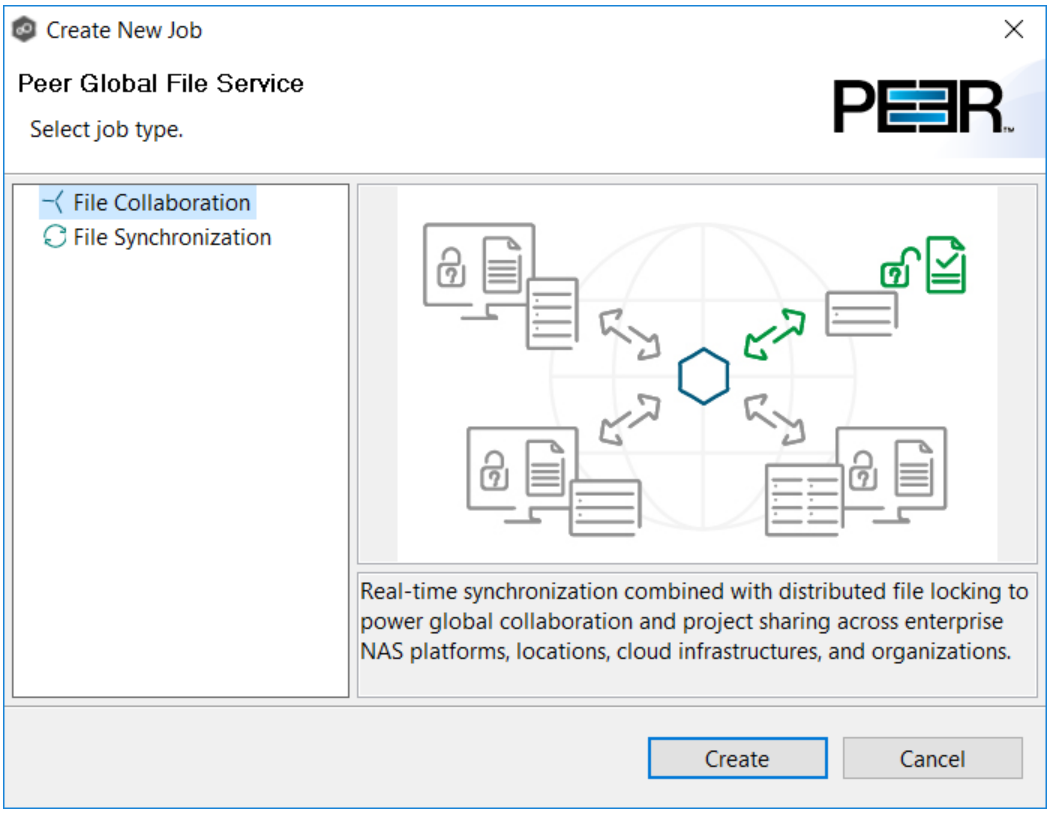

4. Select a job type and click **Create**:

   - Select **File Collaboration** if locking is required in additional to replication (for example, for data sets with shared project files).

   - Select **File Synchronization** if no locking is required (for example, with home directory and user profile datasets).

5. Follow the wizard prompts as it walks you through creating the job:

   - Naming the job

   - Adding the participants

   - Choosing File Metadata options

   - Choosing Application Support options.

   - Selecting email alerts

   The process is the same as described in Creating a File Collaboration Job and Creating a File Synchronization Job.

Once you have selected your email alerts, the **DFS Namespace** page is displayed.

The **Enable DFS Linking** checkbox is preselected and the **Select Existing DFS Namespace** option is selected.

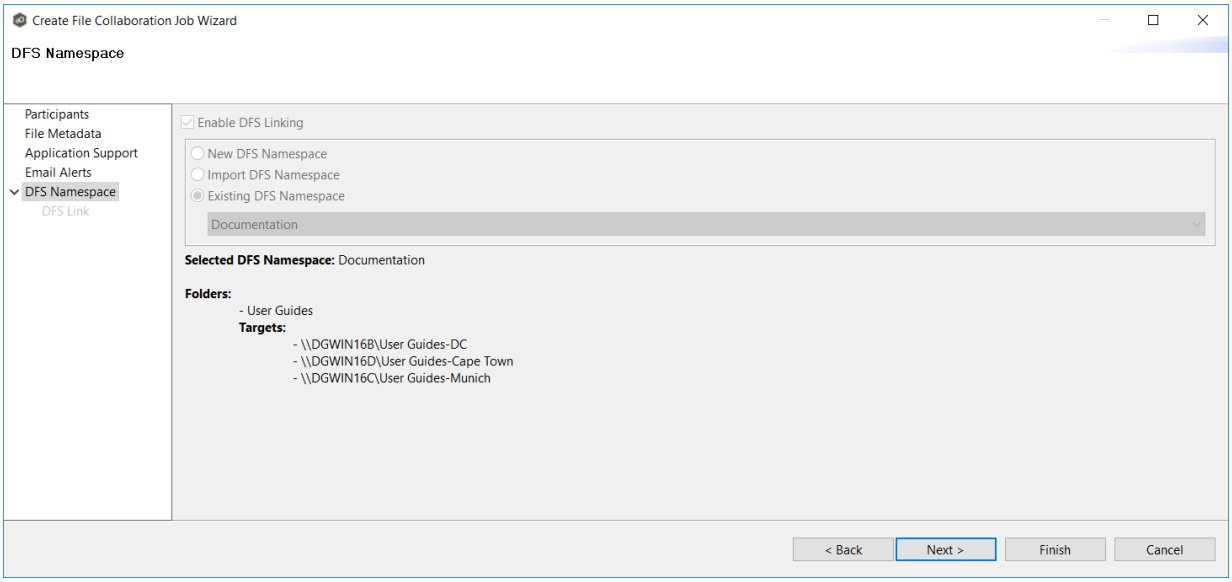

6. Click **Next** if you want to create folder targets; otherwise click **Finish** and continue with Step 9.

   You can create folder targets later if you wish. See Adding a Namespace Folder Target for step-by-step instructions.

   If you clicked **Next**, the **DFS Link** page appears. The purpose of this page is to link the watch path of each participant of a collaboration/synchronization job (specified in the **Path** page) to the appropriate folder target.
   Initially, the **Folder Target** column in the **Linked Participants and Folder Targets** table will not contain any folder targets. After linking, a folder target should be listed in this column for each participant.

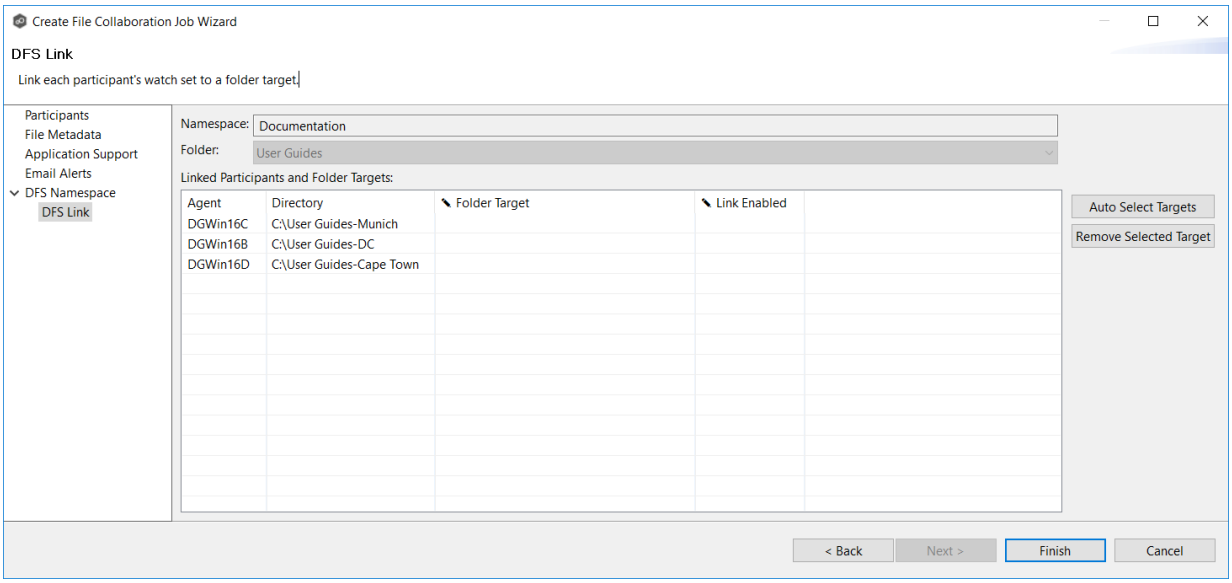

7. Link the participants to folder targets using one or more of the following methods:

   - Click **Auto Select Targets** to have PeerGFS select the targets for you. If you use this method, PeerGFS will try to match the watch path specified for each participant. It will populate the **Folder Target** column with its matches.

   - Manually link the participants by typing a path in the **Folder Target** column.

   - Use a combination of auto select and manually select. For example, you can use auto select, and if the correct targets are not selected, you can manually enter the folder target path in the Folder Target column.

   - If an incorrect target appears in the **Folder Target** column, you can select the incorrect target, and then click **Remove Selected** Target.

   After linking targets, the **Linked Participants and Folder Targets** table should contain a folder target for each participant.

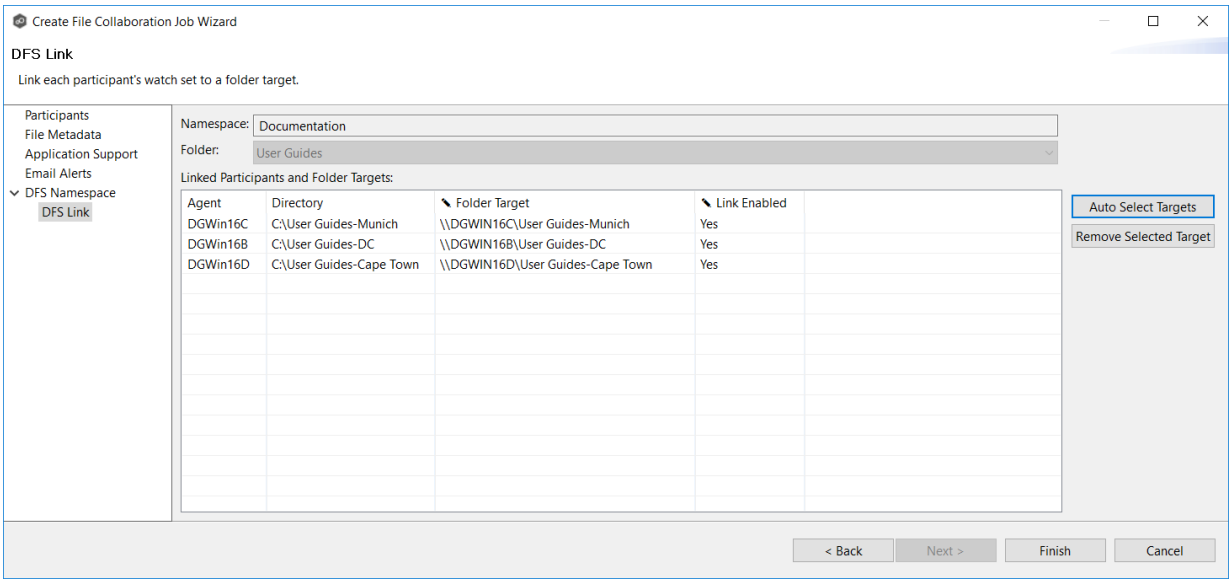

8. If the **Link Enabled** column is blank, select **Yes** for each participant.

9. Click **Finish**.

   The runtime view for the newly created File Collaboration or File Synchronization job is displayed.

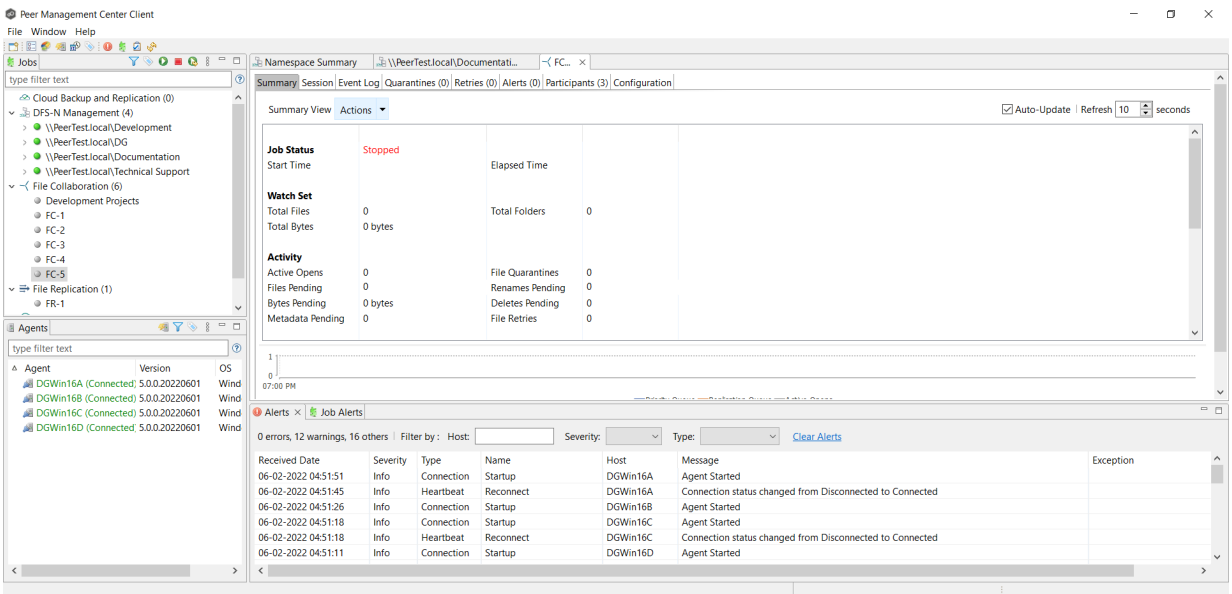

10. Click the DFS-N Management job tab to display its runtime view.

    The **Linked to Participant** column is now populated.

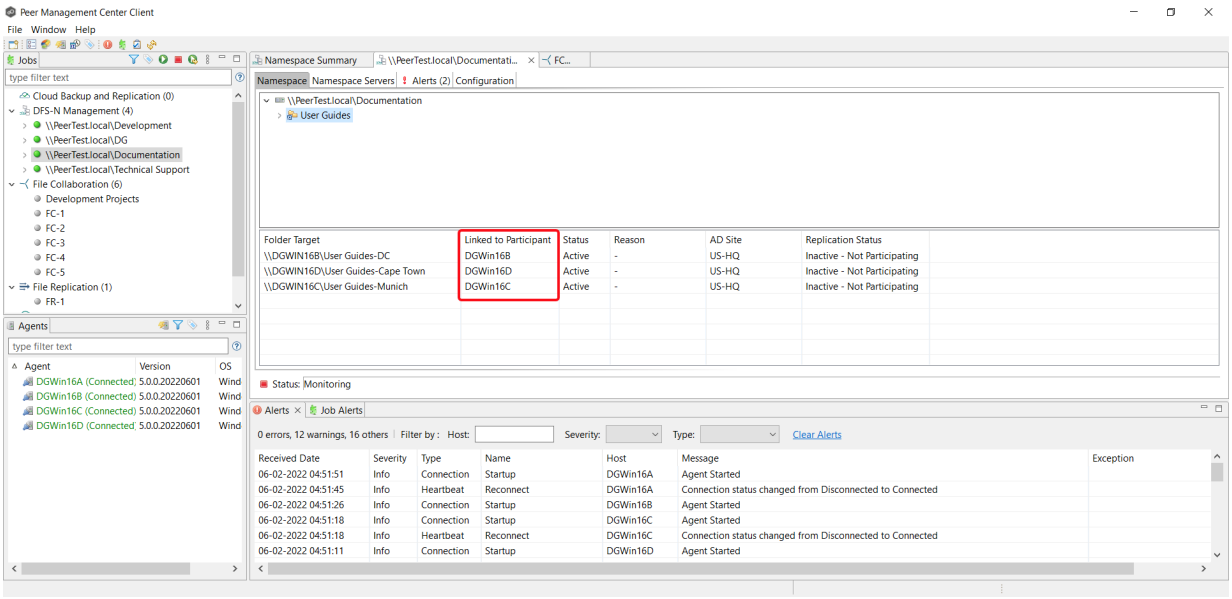

Linking an Existing Namespace Folder to an Existing File Collaboration or File Synchronization Job

You can link an existing DFS namespace folder with an existing File Collaboration or File Synchronization job. These steps require that the DFS namespace has been already created and is being managed by a DFS-N Management job.

To link a namespace folder with an existing File Collaboration or File Synchronization job:

1. Select the File Collaboration or File Synchronization job in the **Jobs** view.

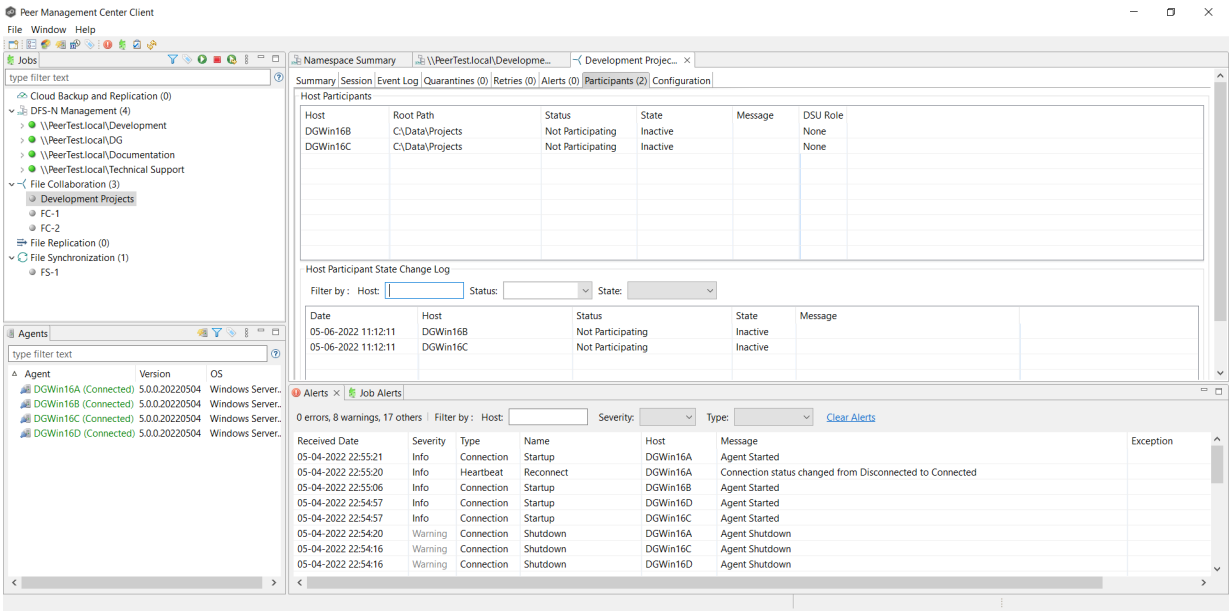

2. Right-click and select **Edit Job**.

   The **Edit Job** wizard appears.

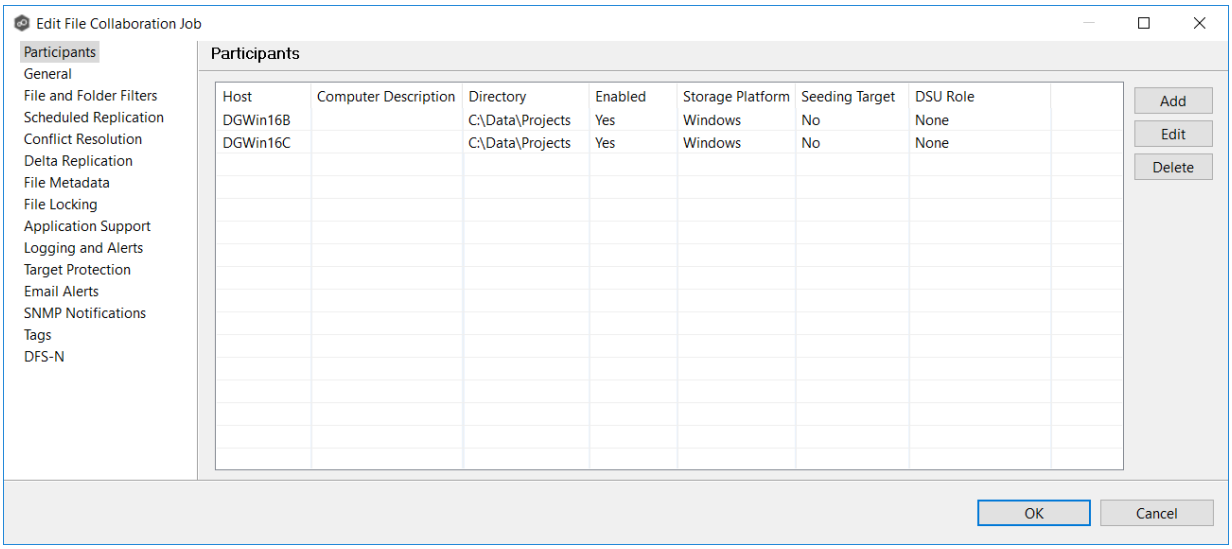

3. Select **DFS-N** in the navigation tree on the left.

   The DFS-N page is displayed.

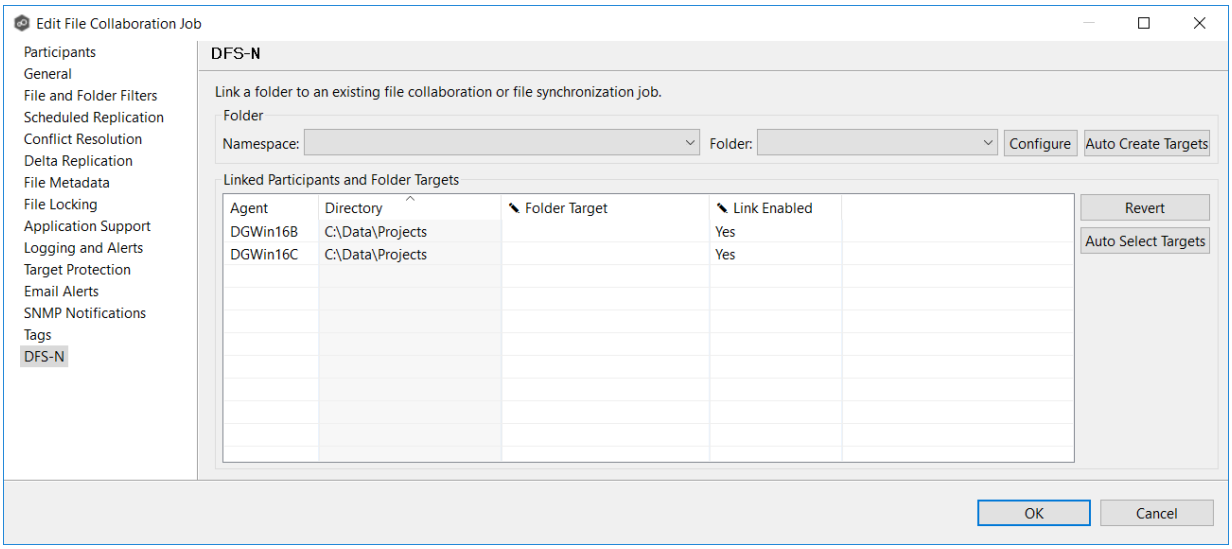

4. In the **DFS Namespace Folder** area, select the namespace you want to link to from the first drop-down list.

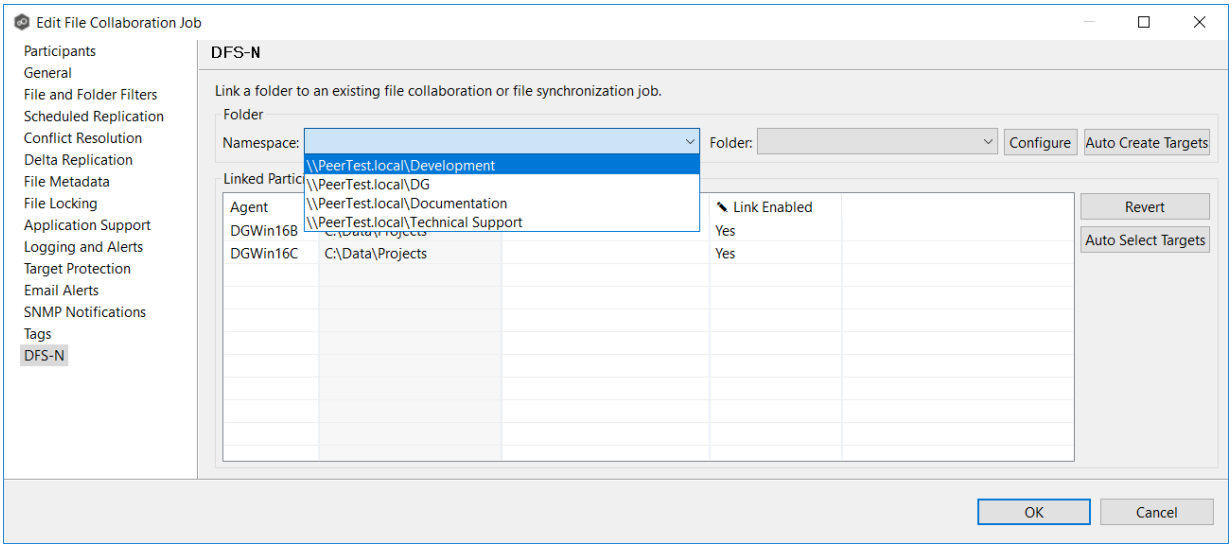

Once you've selected a namespace, a list of available namespace folders appears in the **Folder** drop-down list.

5. Select the namespace folder from **Folder**.

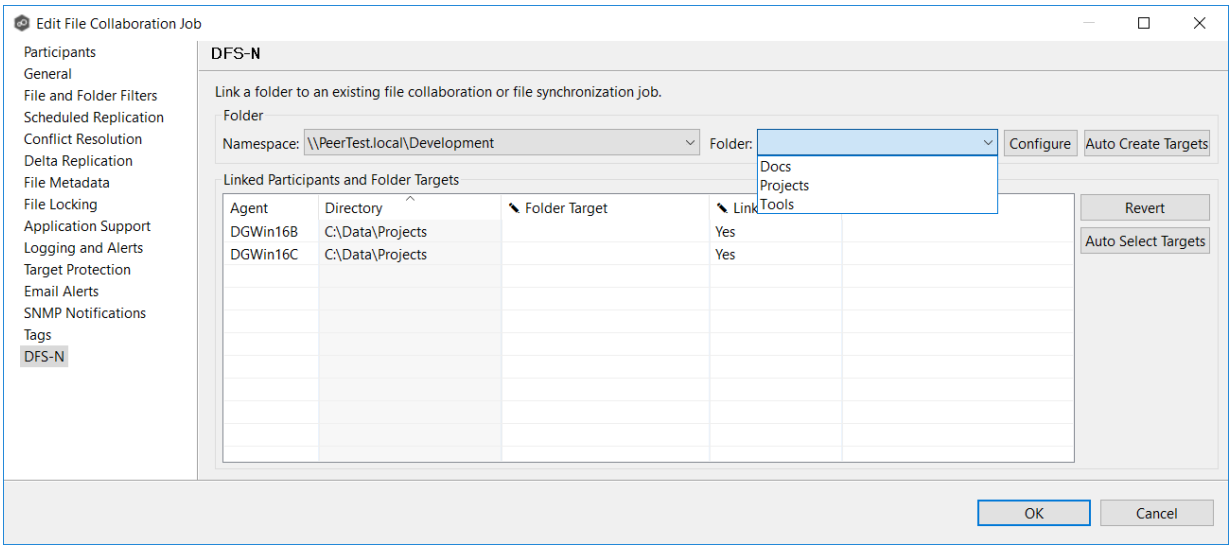

Once you've selected a namespace and a folder, you need to link each participant to a folder target.

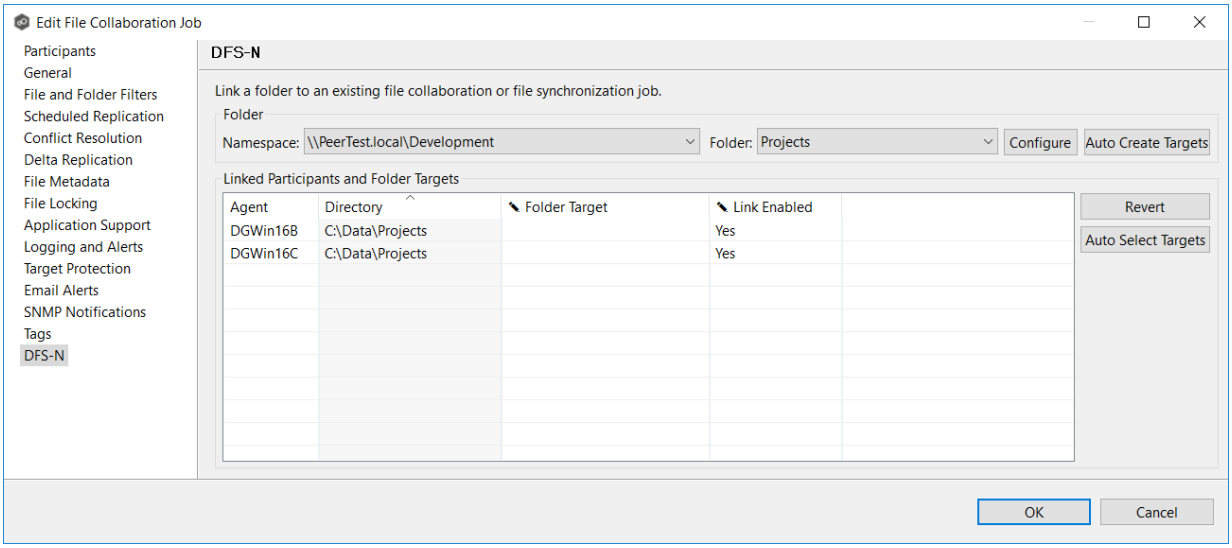

6. Click **Auto Select Targets** to have PeerGFS attempt to automatically map the participants with folder target.

   After auto selection, the linked participants and folder targets are displayed in the **Linked Participants and Folder Targets** table.

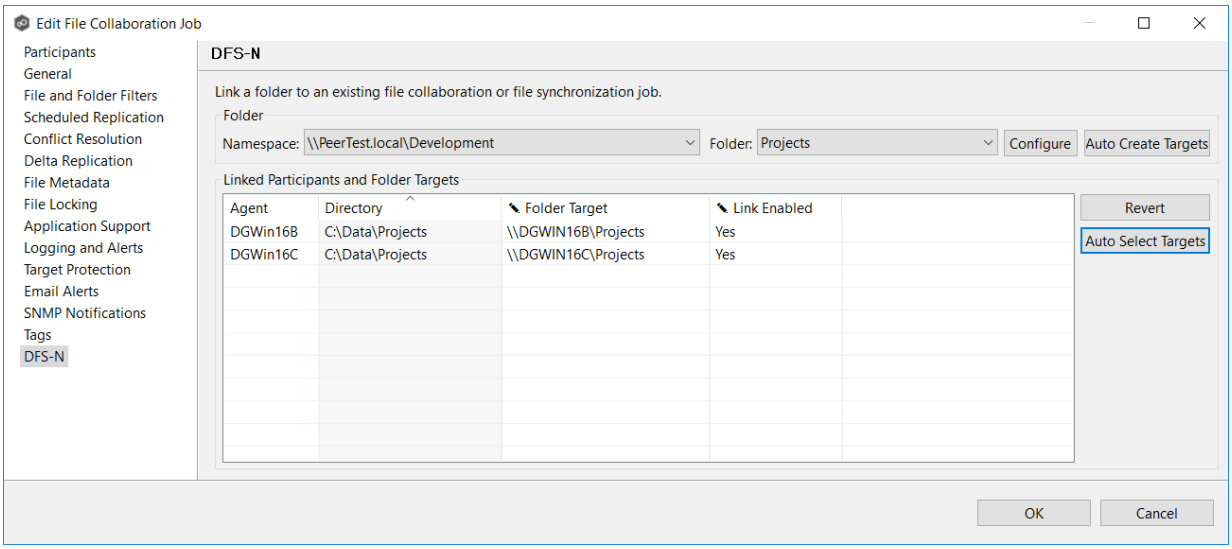

Tip: In most cases, clicking the **Auto Select Targets** button, PeerGFS will be able to automatically link a folder target with the appropriate participant. However, if a folder does not have the appropriate folder targets, click the **Auto Create Targets** button. The wizard that appears will use the paths configured in your File Collaboration or File Synchronization job and try to automatically create folder targets for you.

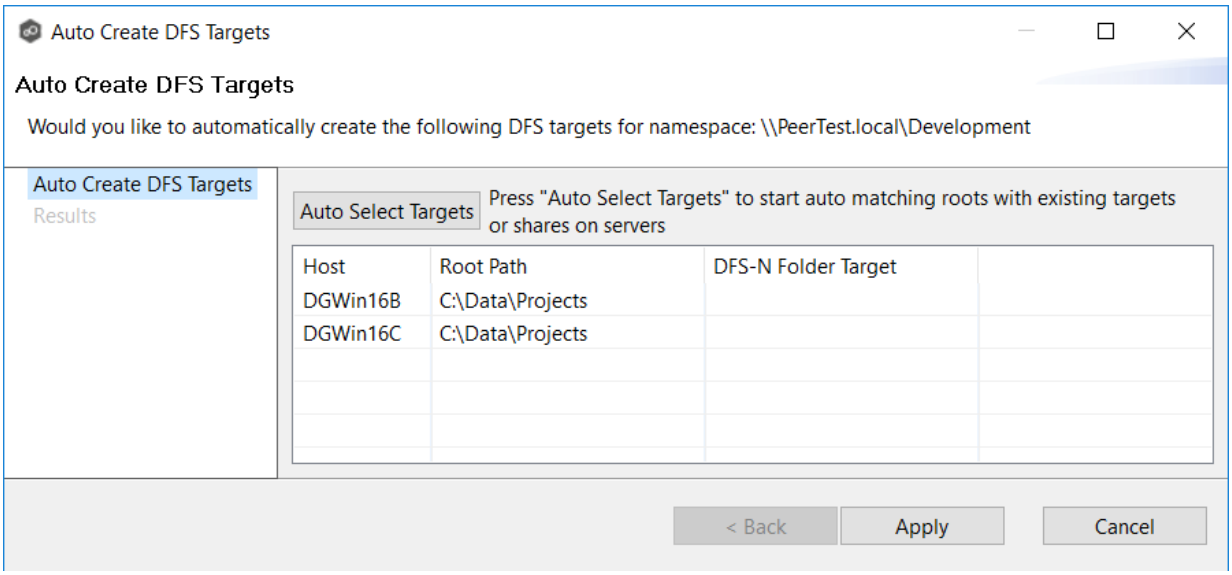

7. Review the values in the **Link Enabled** column; if **No** appears , select **Yes** from the drop-down list.

8. Once all participants are linked to the appropriate folder targets, click **OK** to save your changes.

The runtime window appears. From this point forward, if this collaboration or synchronization job is running along with its paired DFS-N Management job, Peer Management Center will automatically failover and failback folder targets.

1. To confirm that the namespace is linked to the file collaboration or file synchronization job:

   a. Open the namespace job.

   b. In the job runtime view, open the Namespace tab.

   c. In the top panel of the tab, select a folder name to confirm that the appropriate folder targets are linked to participants. The folder targets and linked participants will be displayed in the bottom panel of the tab.

1. Start the FC/FS job.

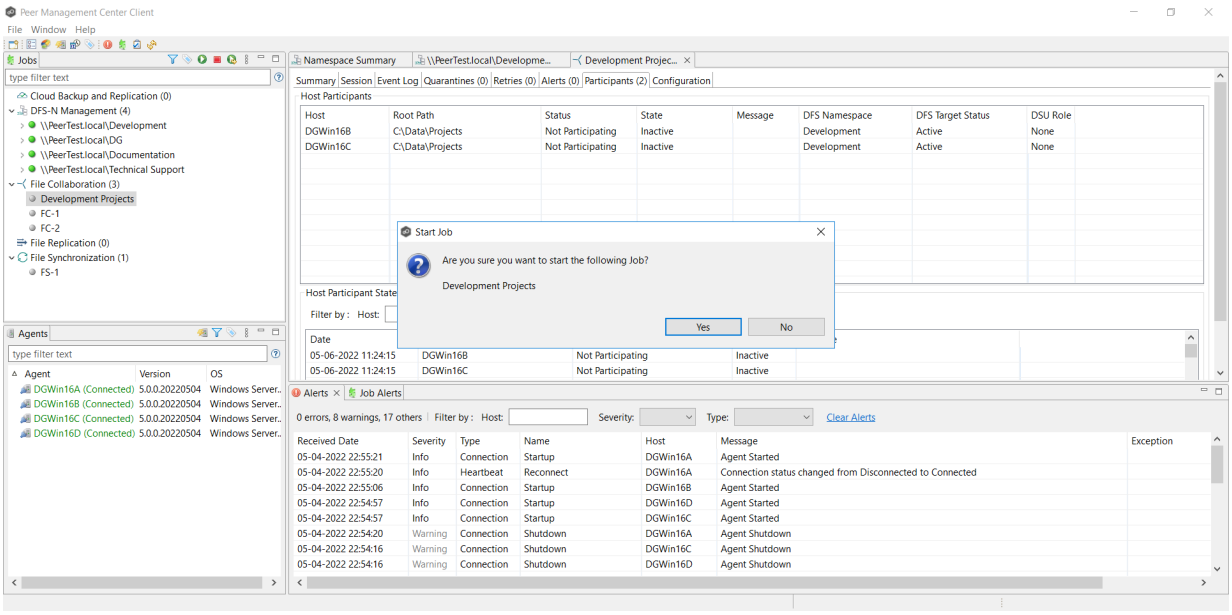

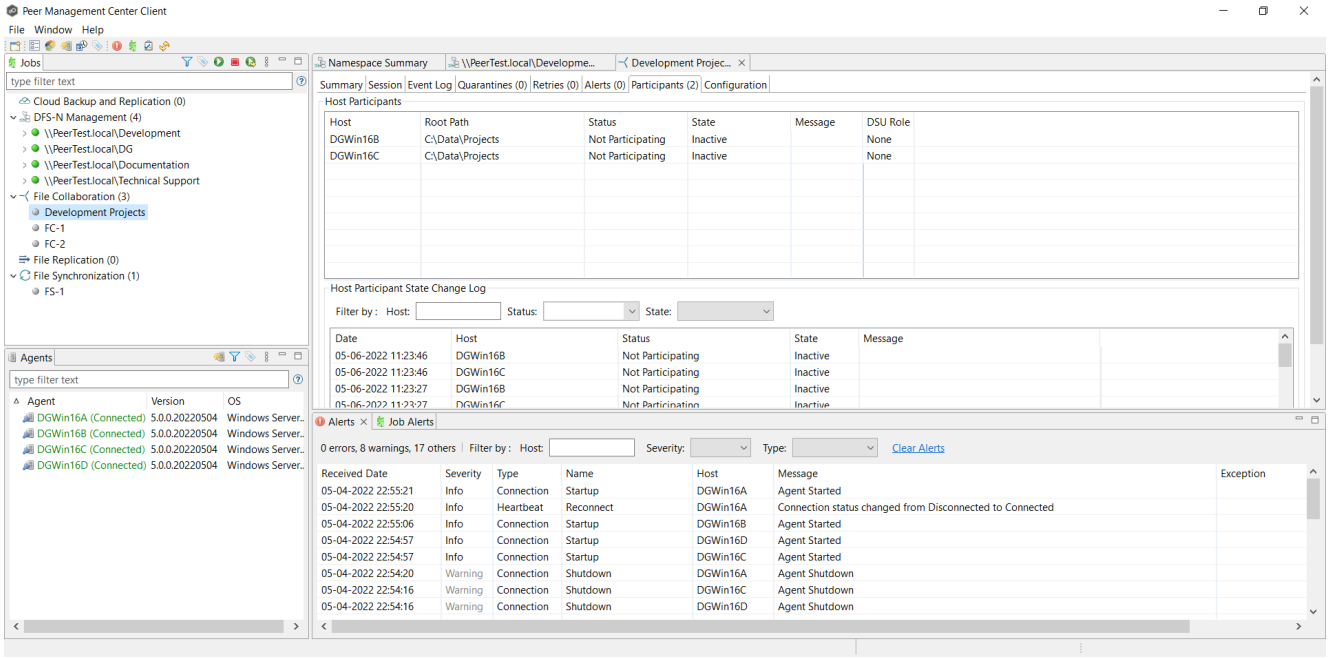

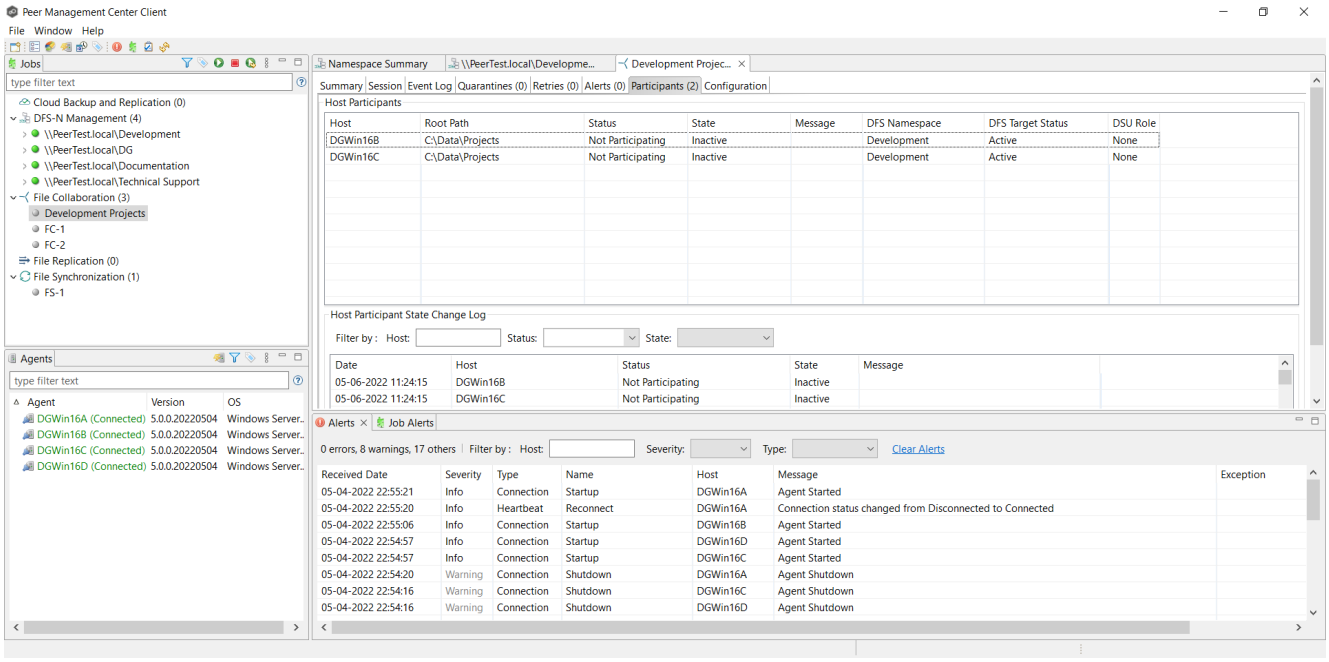

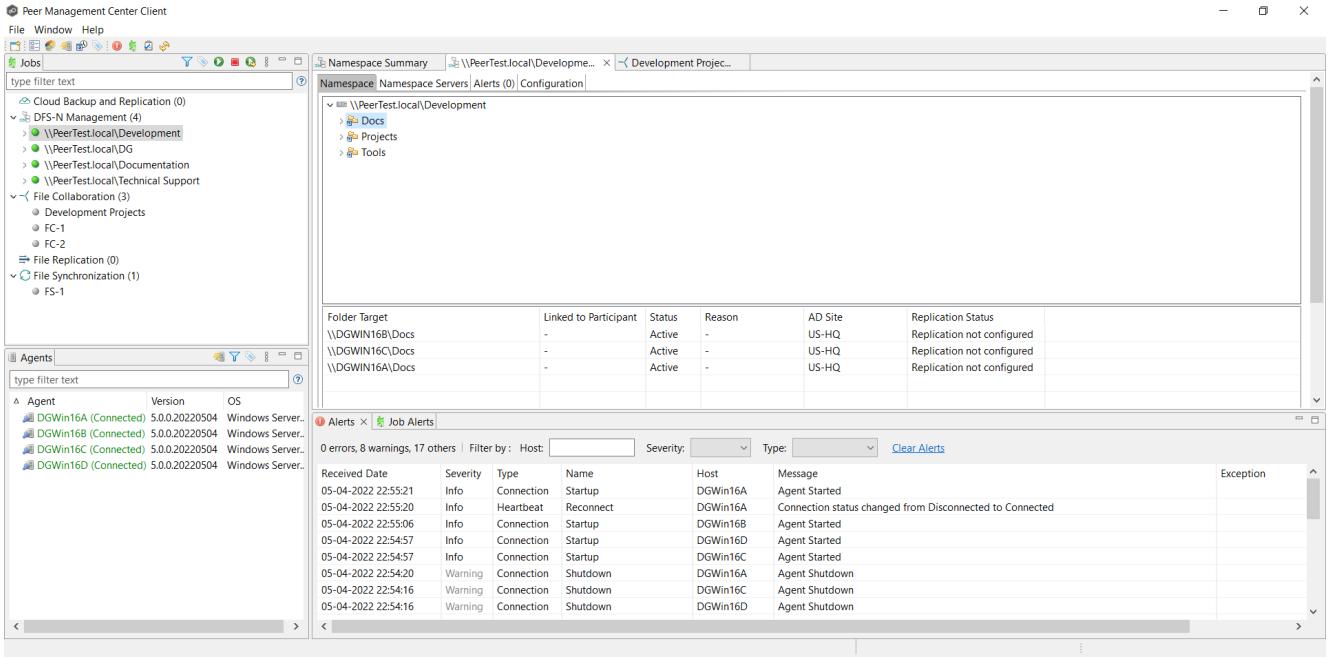

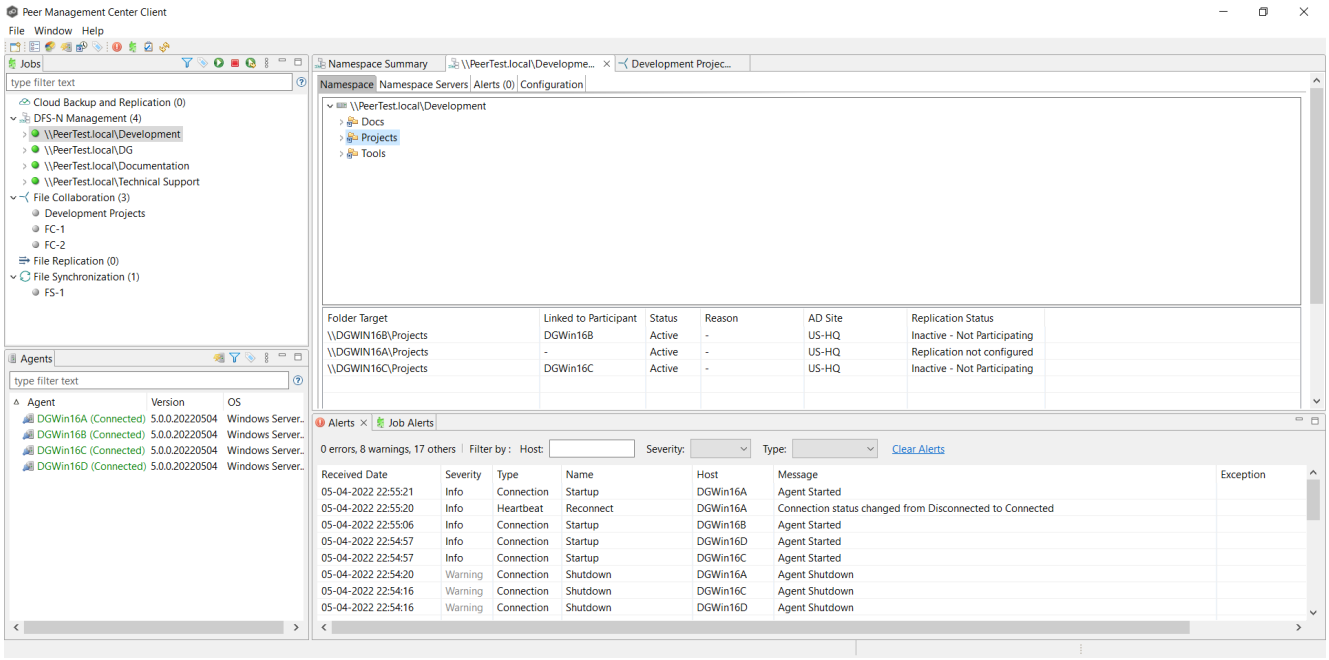

Start the FC/FS job.

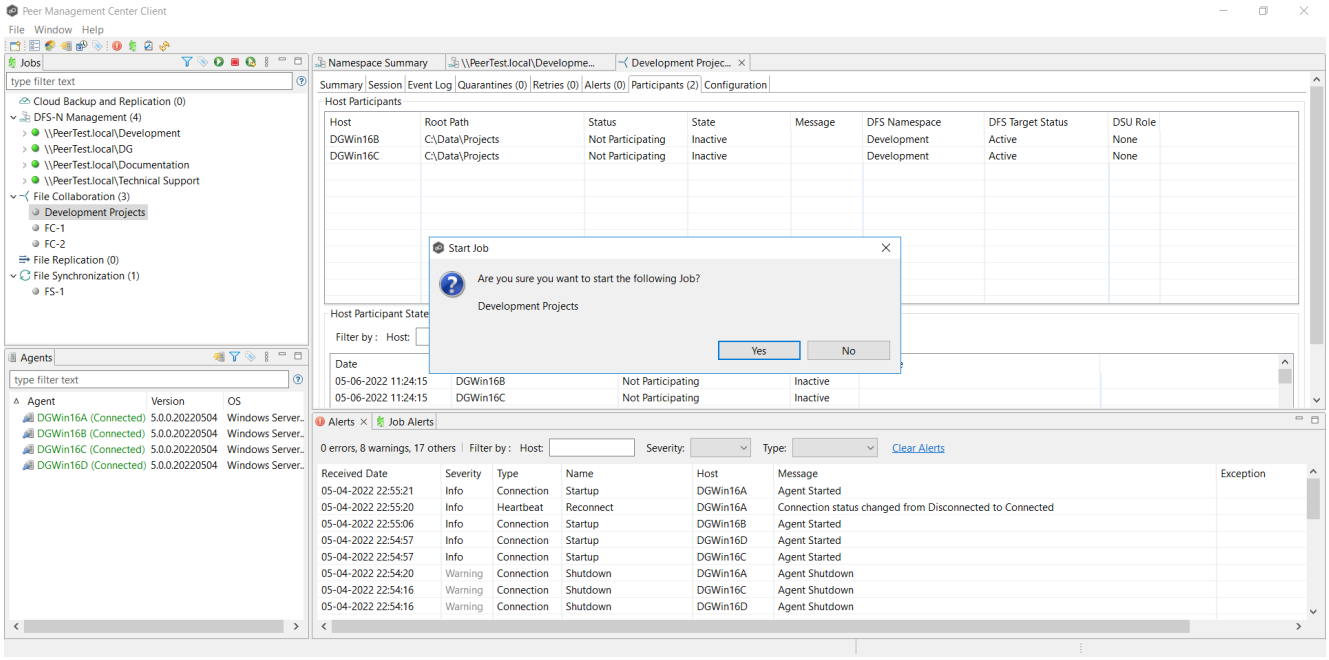

After the job is started, the participants tab shows ???

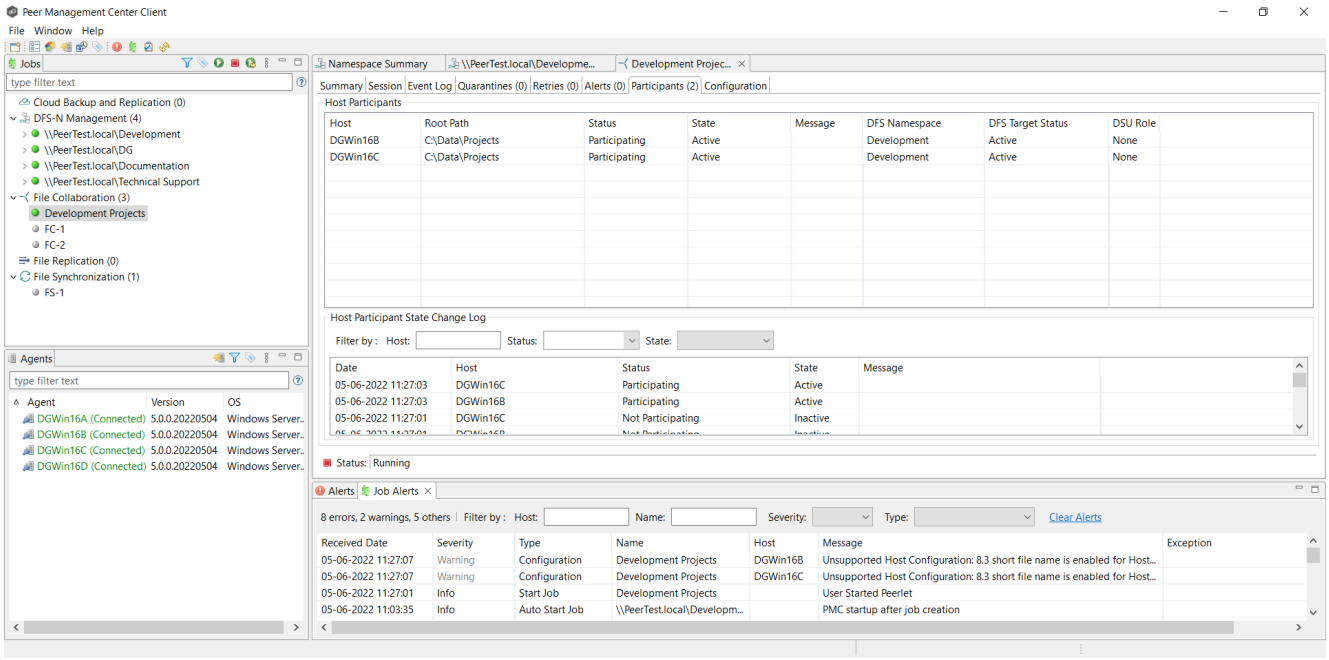

Click the Namespace tab and then click a folder. The folder target panel below shows which folder target is linked to a participant (can multiple folder targets linked to a participant?), the to view the folder targets.

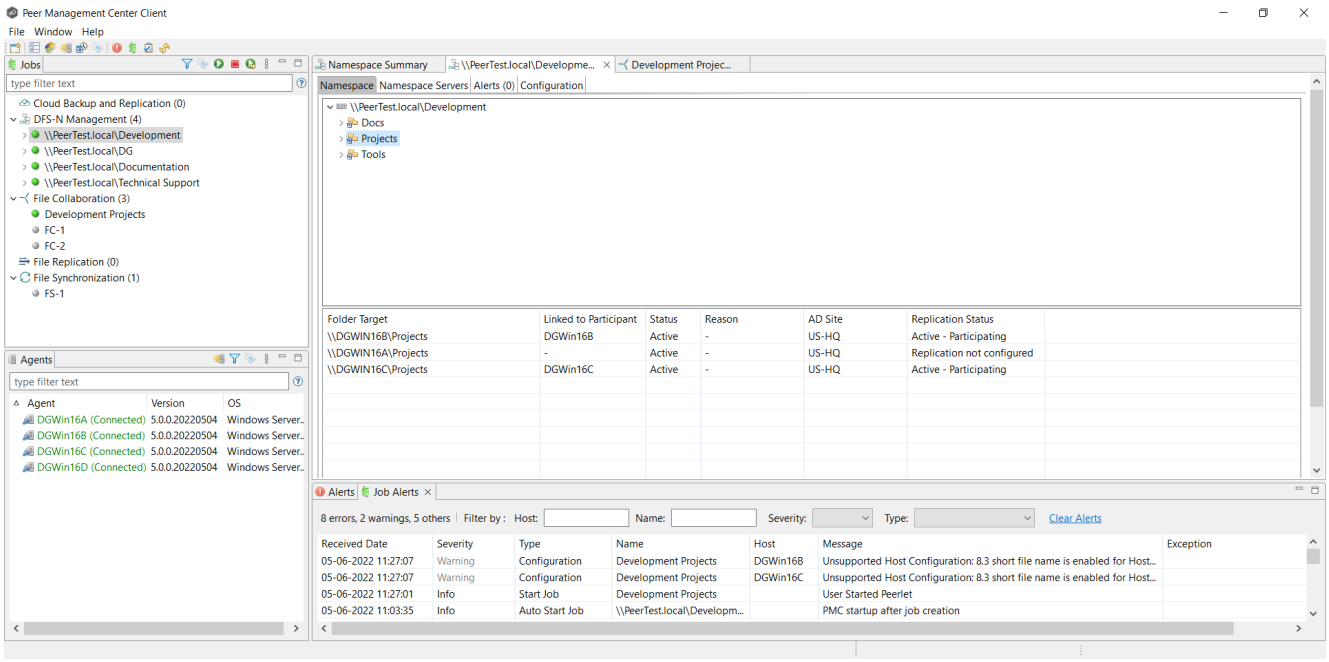

File Collaboration Jobs

This section provides information about creating, editing, running, and managing a File Collaboration job:

- Overview
- Before You Create Your First File Collaboration Job
- Creating a File Collaboration Job
- Editing a File Collaboration Job
- Running and Managing a File Collaboration Job
- Runtime Job Views

Overview

A File Collaboration job provides distributed teams a fast and efficient way to collaborate with shared project files. Unlike other file collaboration solutions that centralize files into a single data repository that cause slow file access across a WAN, a File Collaboration job replicates shared project files to each office site in a distributed environment so that end users are

guaranteed high-speed LAN access to shared files no matter their file size. Version conflicts are prevented through integrated distributed file locking.

By keeping hot data local, File Collaboration maximizes end user productivity. Because files are close to the users, their applications, and their compute resources, the actual performance is as fast as possible from a physical view. At the same time File Collaboration ensures version conflicts are eliminated with file locking.

Before You Create Your First File Collaboration Job

We strongly recommend that you configure the File Collaboration settings (e.g., SMTP notifications), as well as other global settings such as SMTP email settings, email alerts, and file filters before configuring your first File Collaboration job. See Preferences for details on these settings.

Creating a File Collaboration Job

The **Create Job** wizard walks you through the process of creating a File Collaboration job. The process consists of the following steps:

- Step 1: Job Type and Name
- Step 2: Participants
- Step 3: File Metadata
- Step 4: Application Support
- Step 5: Email Alerts
- Step 6: Save Job

Additional configuration options, such as applying file filters and specifying delta level replication, are available when editing a File Collaboration job.

Step 1: Job Type and Name

1. Open Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **Create New Job** wizard displays a list of job types you can create.

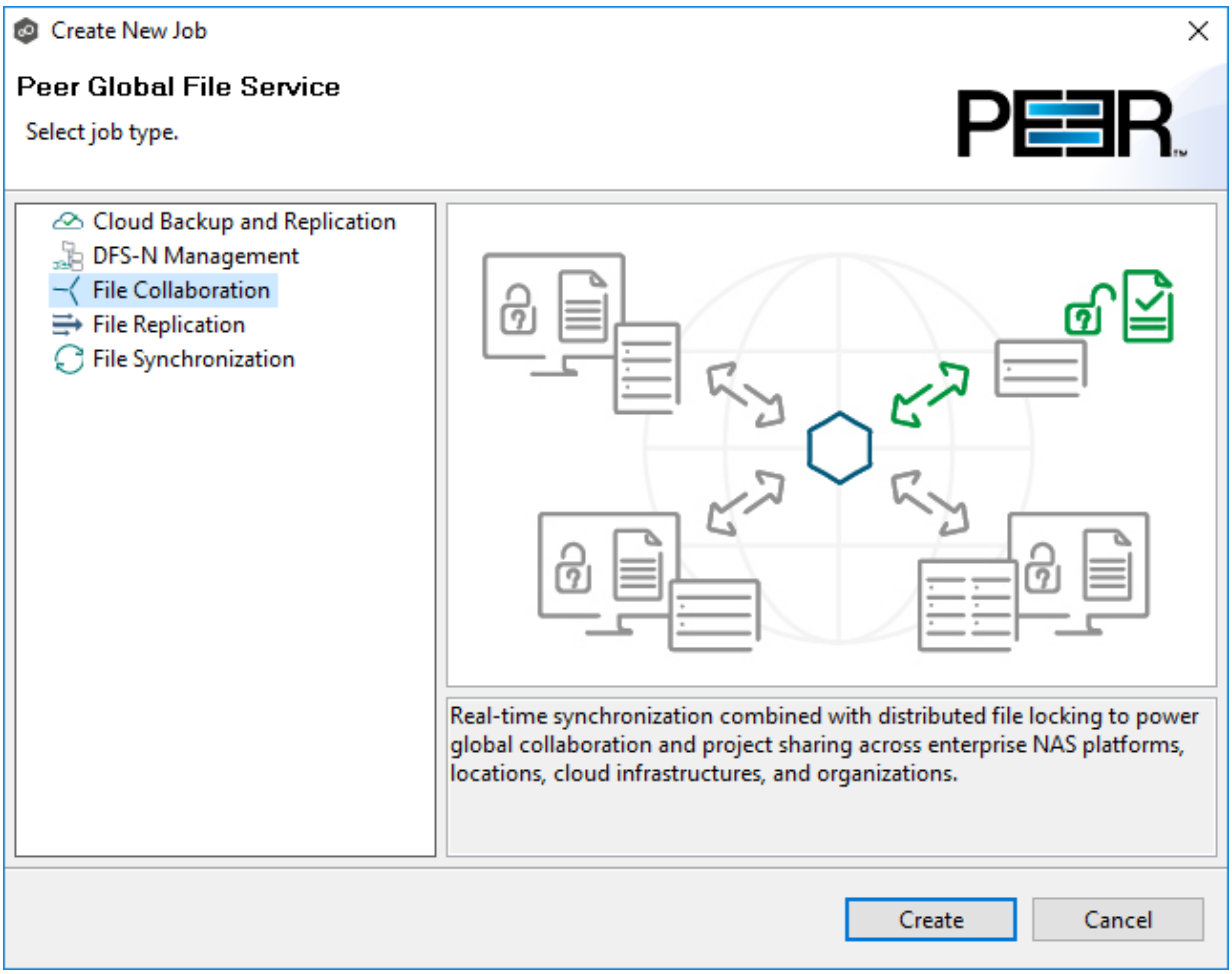

3. Click **File Collaboration**, and then click **Create**.
4. Enter a name for the job in the dialog that appears.

   The job name must be unique.

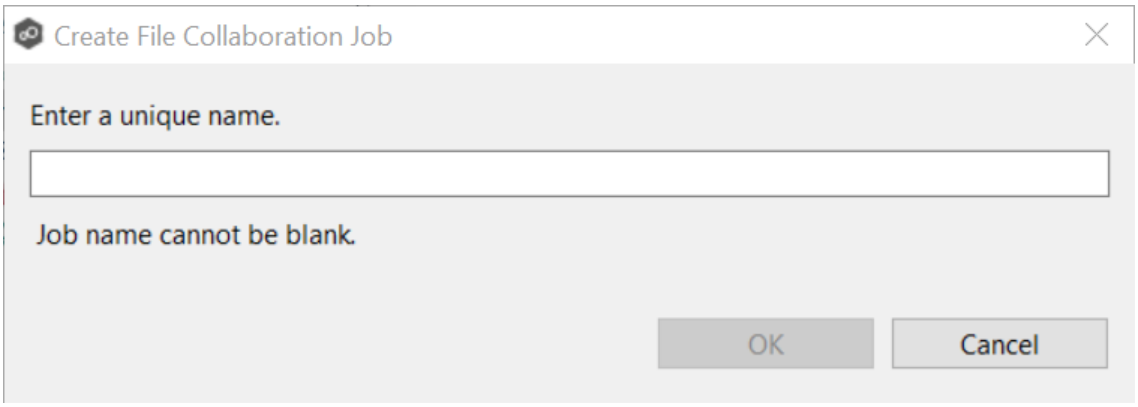

5. Click **OK**.

   The Participants page appears.

Step 2: Participants

After selecting the job type and naming the job, the **Participants** page is displayed. It contains a table that will display the job participants once you have added them. A File Collaboration job must have two or more participants. A **participant** consists of an Agent and the volume/share/folder to be replicated. A File Collaboration job synchronizes the files of participants in real-time and adds distributed locking to avoid version conflicts.

1. Click the **Add** button to start the process of adding a participant.

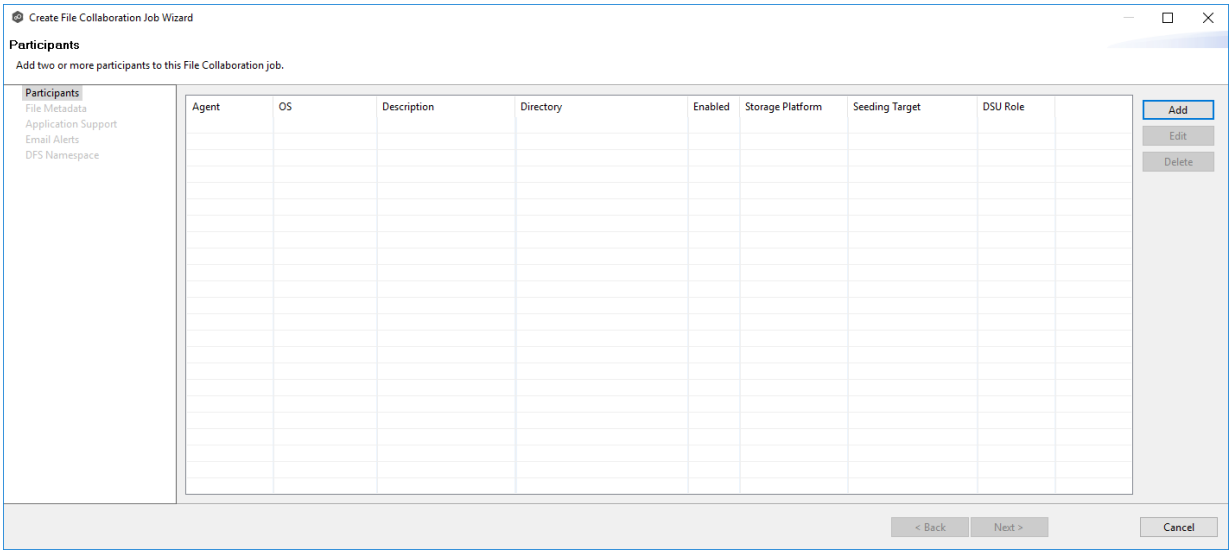

The **Add New Participant** wizard opens; it walks you through the steps for adding a participant:

a. Selecting a Management Agent, which is the Agent that will manage the storage device that hosts data you want to replicate.

b. Selecting the type of storage platform that hosts data you want to replicate.

c. Entering the credentials needed to access a specific storage device and providing other storage information.

d. Entering the path to the watch set (the data that you want to replicate) and selecting whether participant will be a seeding target.

e. (Optional) Enabling Dynamic Storage Utilization for the participant.

Once you have added a participant, it is listed in the **Participants** table.

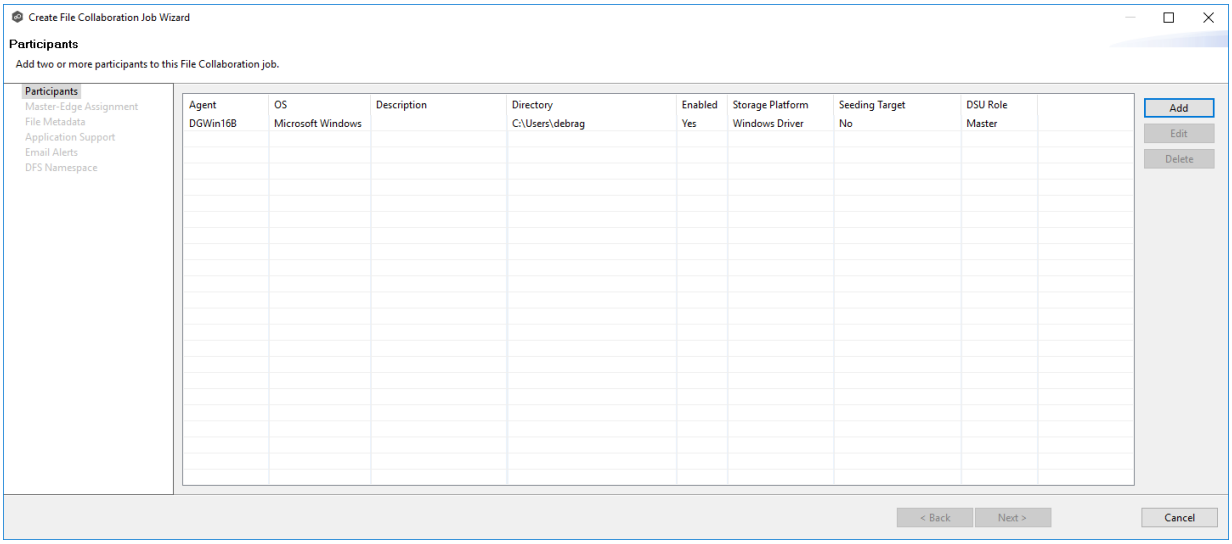

2. To add more participants, click **Add** and repeat the steps for each participant you want to add to the job.

3. Once you have added all the participants, click **Next**.

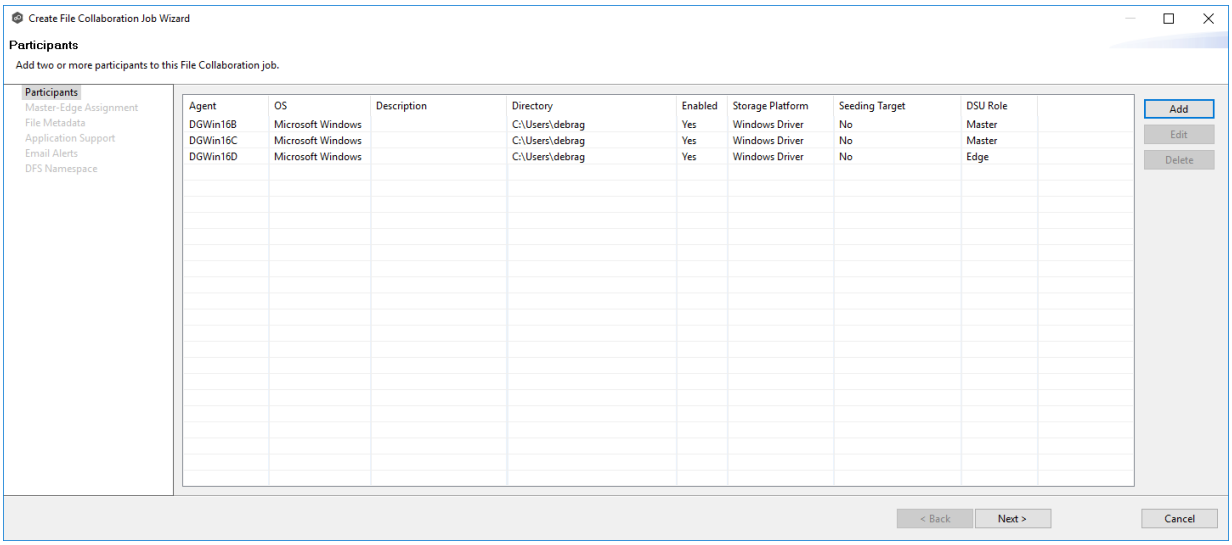

The page that appears next depends on options you selected when adding a participant:

- if you have enabled Dynamic Storage Utilization for a participant, the Master-Edge Assignment page appears.
- Otherwise, the File Metadata page appears.

The **Management Agent** page lists available Agents. You can have more than one Agent managing a storage device—however, the Agents must be managing different volumes/shares/folders on the storage device. For your File Collaboration job, you should select a Management Agent to manage the volumes/shares/folders you want to replicate in this job.

1. Select an Agent to manages the storage device.

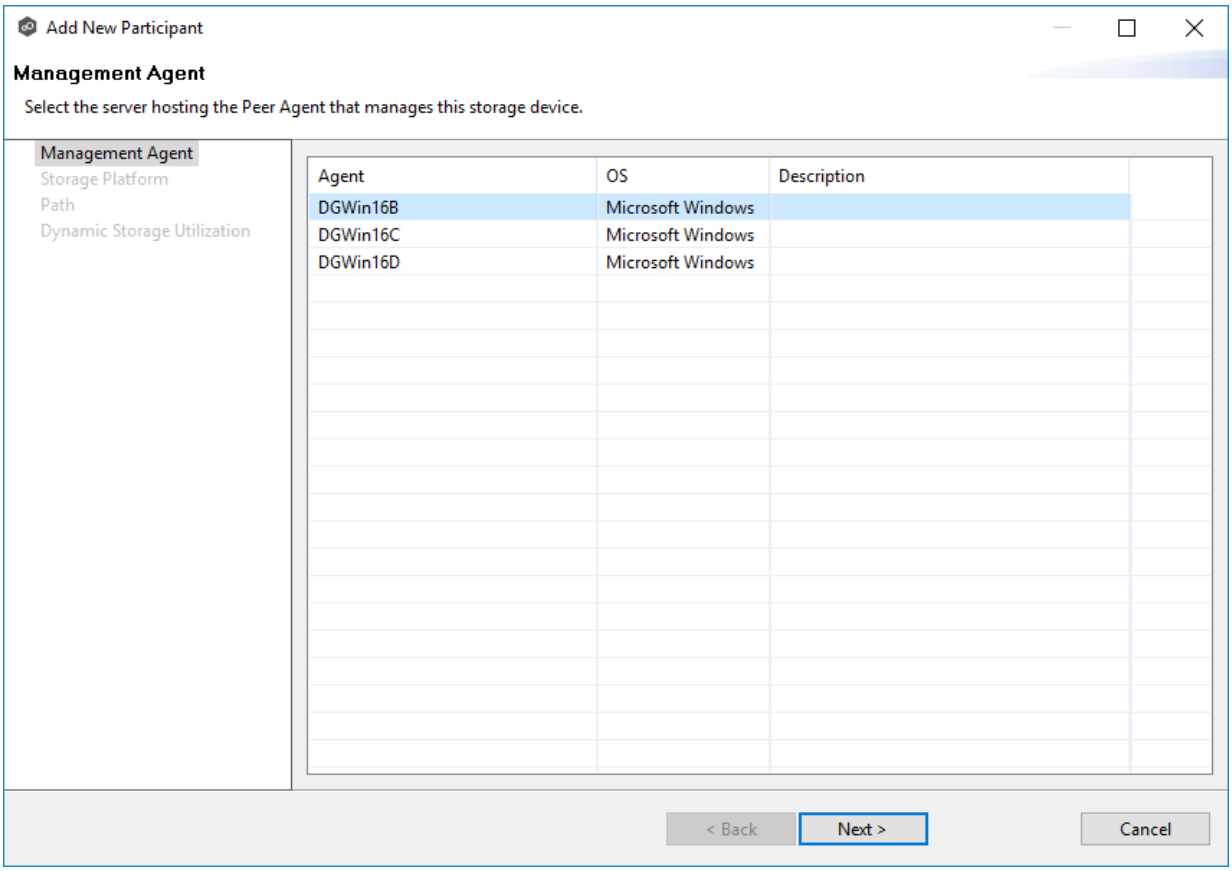

Tip: If the Agent you want is not listed, the Peer Agent Service may not be running on the server hosting the Agent. Try restarting the Peer Agent Service. If the service successfully connects to the Peer Management Broker, then the list is of available agents will be updated with that Agent.

2. Click **Next**.

   The Storage Platform page is displayed.

The **Storage Platform** page lists the types of storage platforms that File Collaboration supports.

1. Select the type of storage platform that hosts the data you want to replicate.

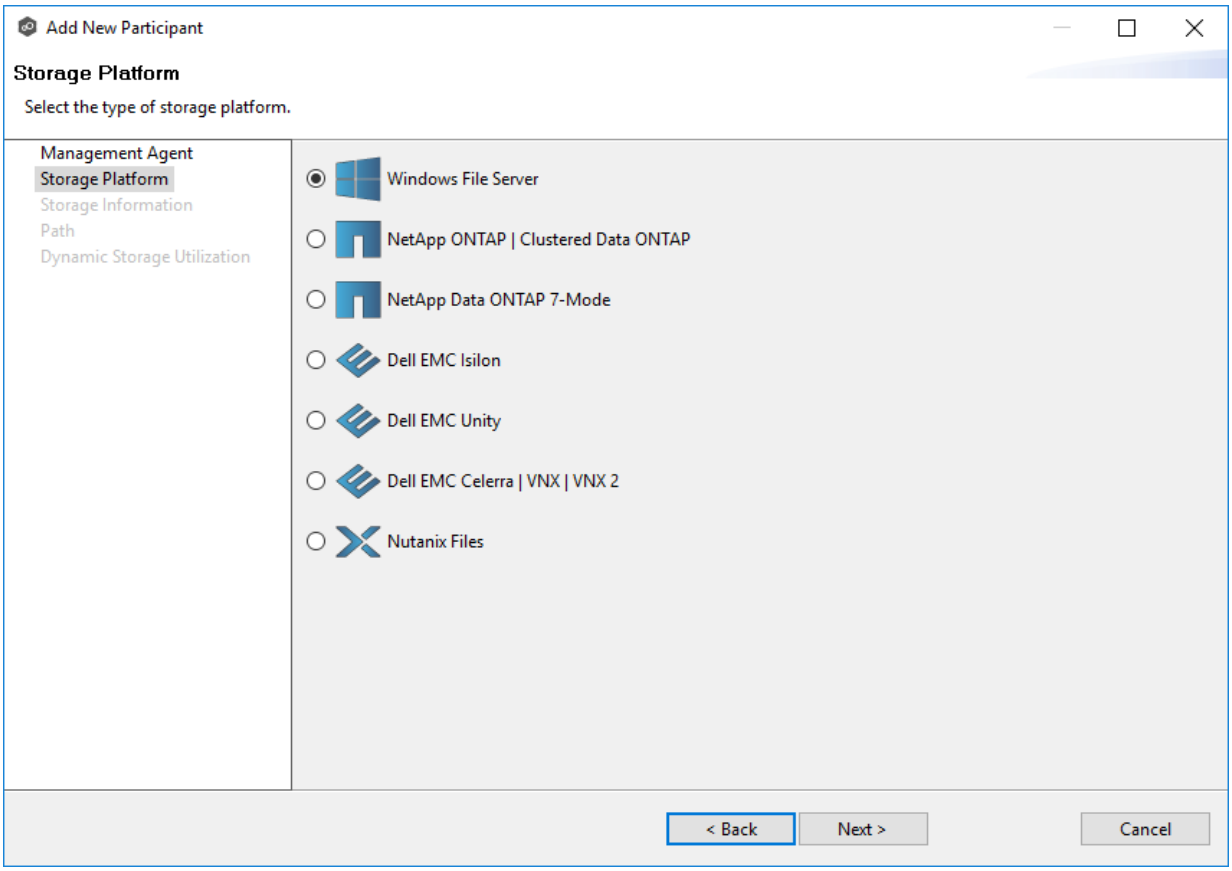

2. Click **Next**.

   The Storage Information page is displayed.

On the **Storage Information** page, you will select a specific storage device containing data that you want to replicate and enter other information about the storage device.

The contents of the **Storage Information** page varies, depending on your selection in the Storage Platform page:

- If you selected **Windows File Server**, you are prompted for configuration information relating to Windows File Server. Continue with the Windows File Server page.
- If you selected any other type of storage device other than Windows File Server, the **Storage Information** page requests the credentials necessary to connect to that storage device. Continue with the steps below.

Storage Platforms Other than Windows File Server

For storage platforms other than Windows File Server:

1. Select **New Credentials** or **Existing Credentials**.

2. If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the Path page.

   If you selected **New Credentials**, enter the credentials for connecting to the storage device. The information you are prompted to enter varies, depending on the type of storage platform:

   NetApp ONTAP | Clustered Data ONTAP

   NetApp Data ONTAP 7-Mode

   Dell EMC Isilon

   Dell EMC Unity

   Dell EMC Celerra | VNX | VNX 2

   Nutanix Files

3. Click **Validate** to test the credentials.

   After the credentials are validated, a success message appears.

4. Click **Next**.

   The Path page is displayed.

Windows File Server

1. Select the **Detector Type**:

   - Select **Windows Driver** for more robust logging and better performance (Recommended).

   - Select **Windows** if suggested by Peer Technical Support.

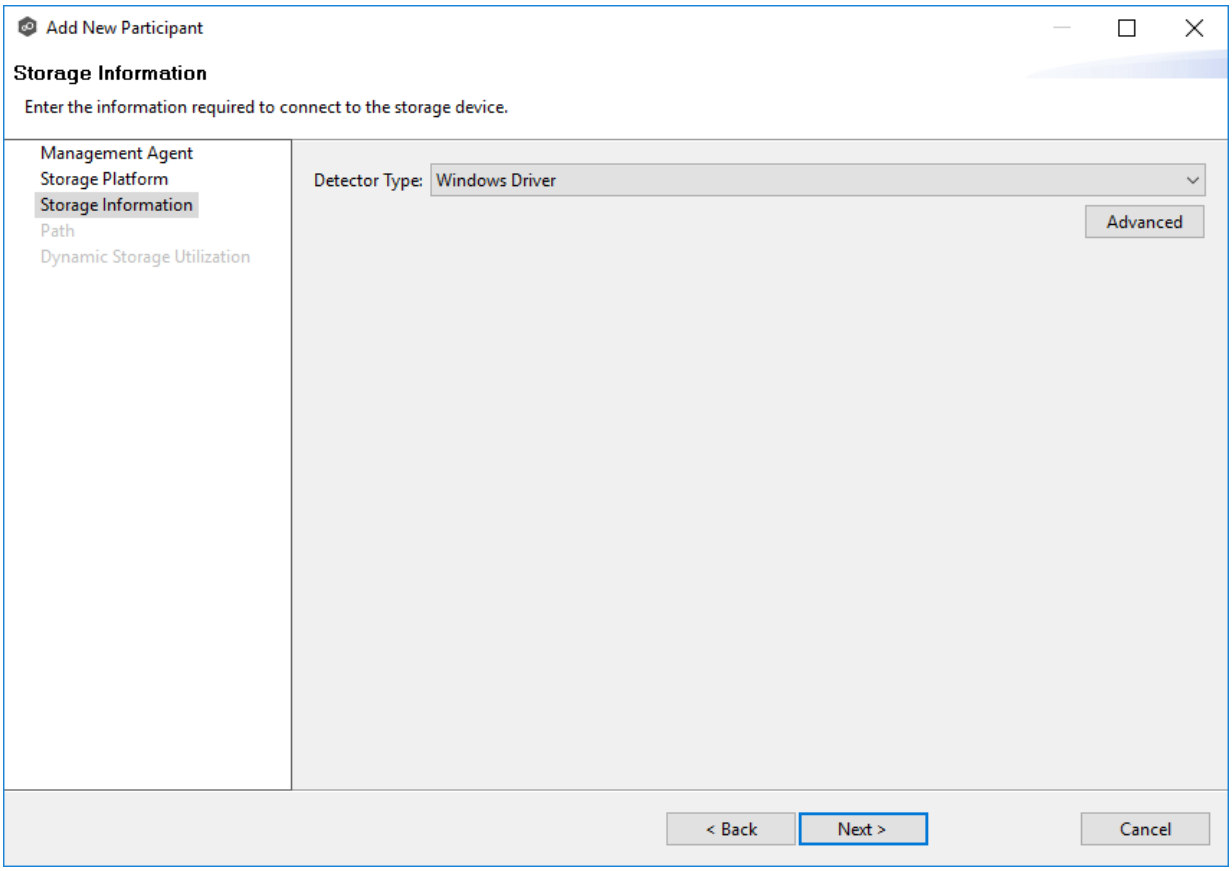

2. Click **Advanced** if you want to set advanced options.

3. Click **Next**.

1. Modify the options as desired.

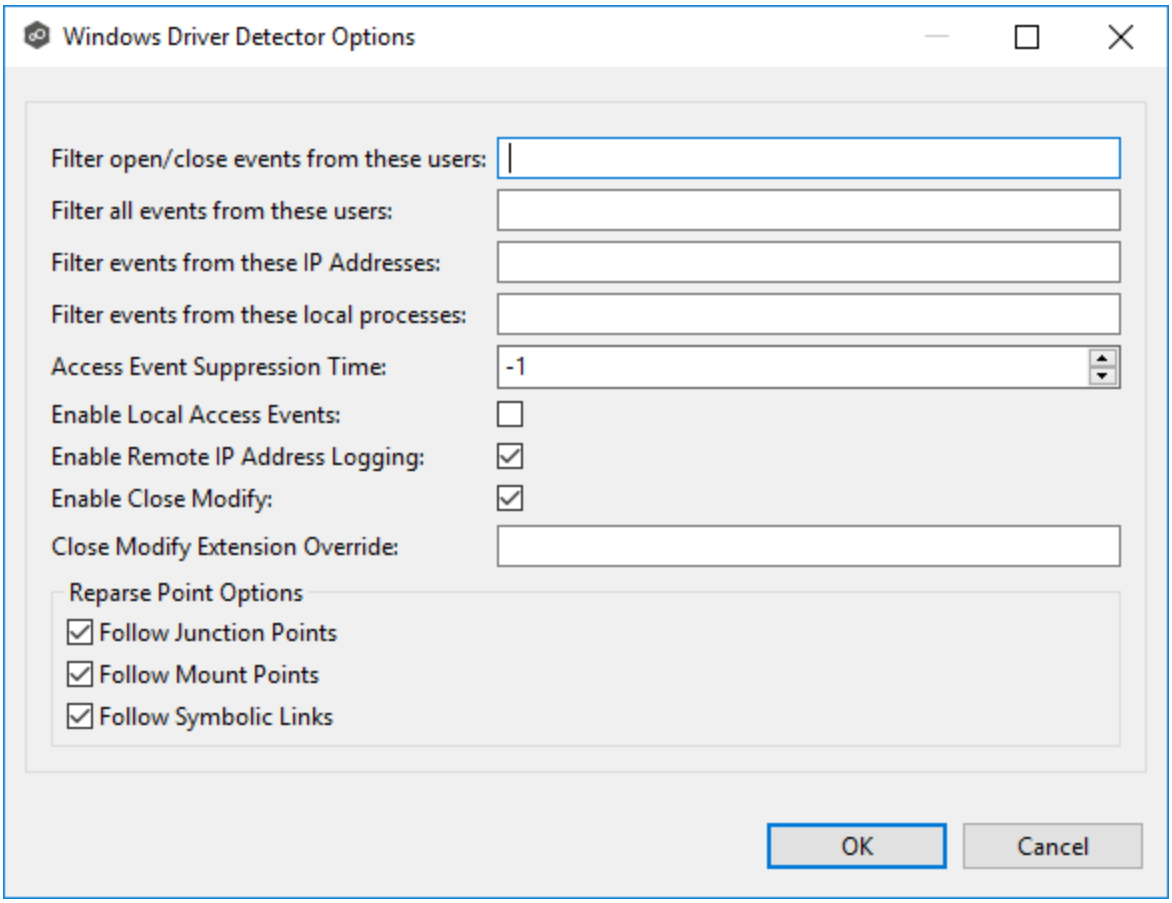

Option	Description
Filter open/close events from these users	Enter a comma-separated list of user account names from which all file opens and closes will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user name under which a backup and/or archival service is running.
Filter all events from these users	Enter a comma-separated list of user account names from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user name under which a backup and/or archival service is running.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the IP addresses on which a backup and/or archival service is running.
Filter events from these local processes	Enter a comma-separated list of local process names on the Agent server from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the specific process names under which a backup and/or archival service is running.
Access Event Suppression Time	Enter the number of seconds to delay an open event before being processed. Use this option to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time an open event will be delayed based on the load of the system.
Enable Local Access Events	Enable tracking of opens and closes that are performed locally on the Agent server.
Enable Remote IP Logging	Enable logging of client IP addresses for all real-time activity.
Enable Close Modify	When enabled, no modify or write events will be detected. Instead, replication of a modified file will be performed when the file is closed.
Close Modify Extension Override	Enter a comma-separated list of exclusions for the Enable Close Modify option. All modify/write events will be detected for these files. This is important for those who rely on sync-on-save functionality.

For more information about junction points or symbolic links, contact <% SUPPORT_EMAIL%

2. Click **OK**.

NetApp ONTAP | Clustered Data ONTAP

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

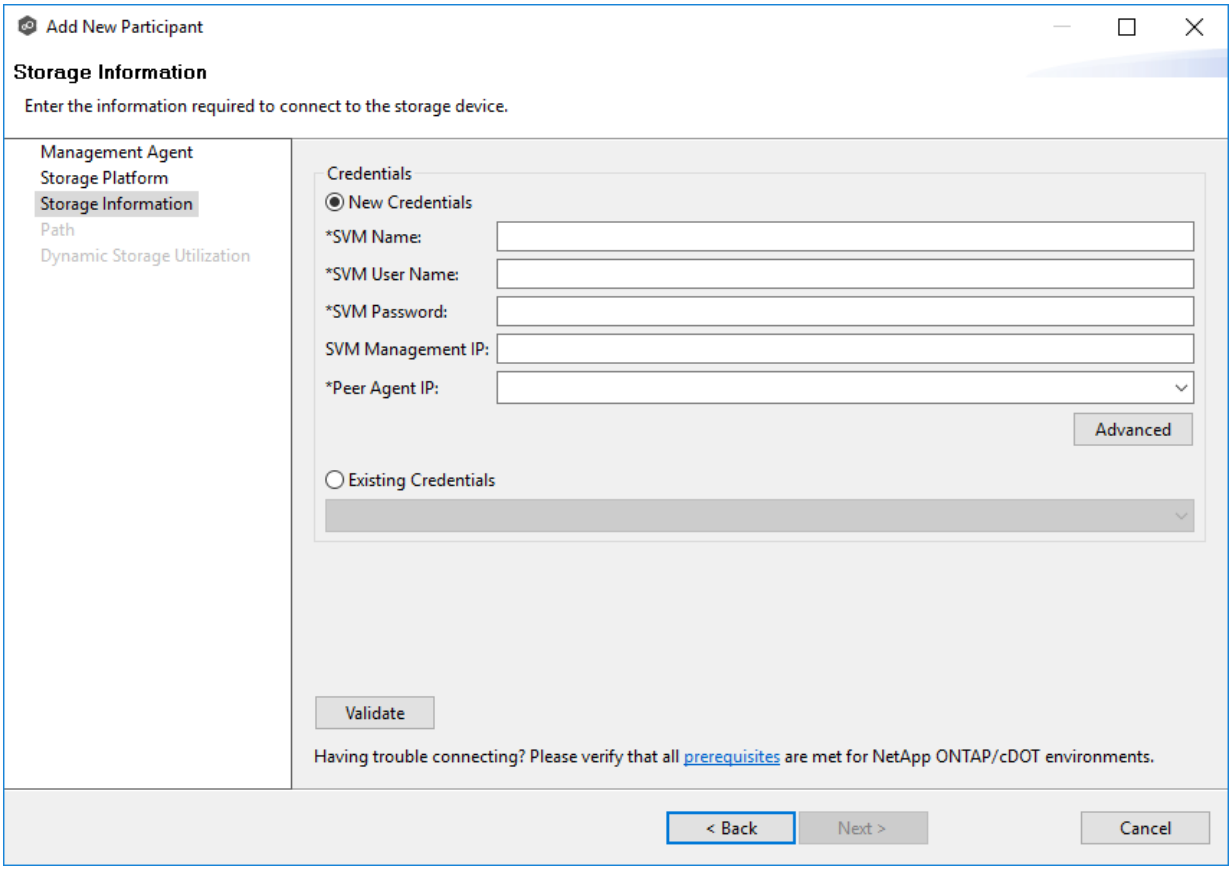

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
SVM Name	Enter the name of the Storage Virtual Machine hosting the data to be replicated.
SVM User Name	Enter the user name for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	(Optional) Enter the IP address used to access the management API of the NetApp Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already allow management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

NetApp Data ONTAP 7-Mode

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

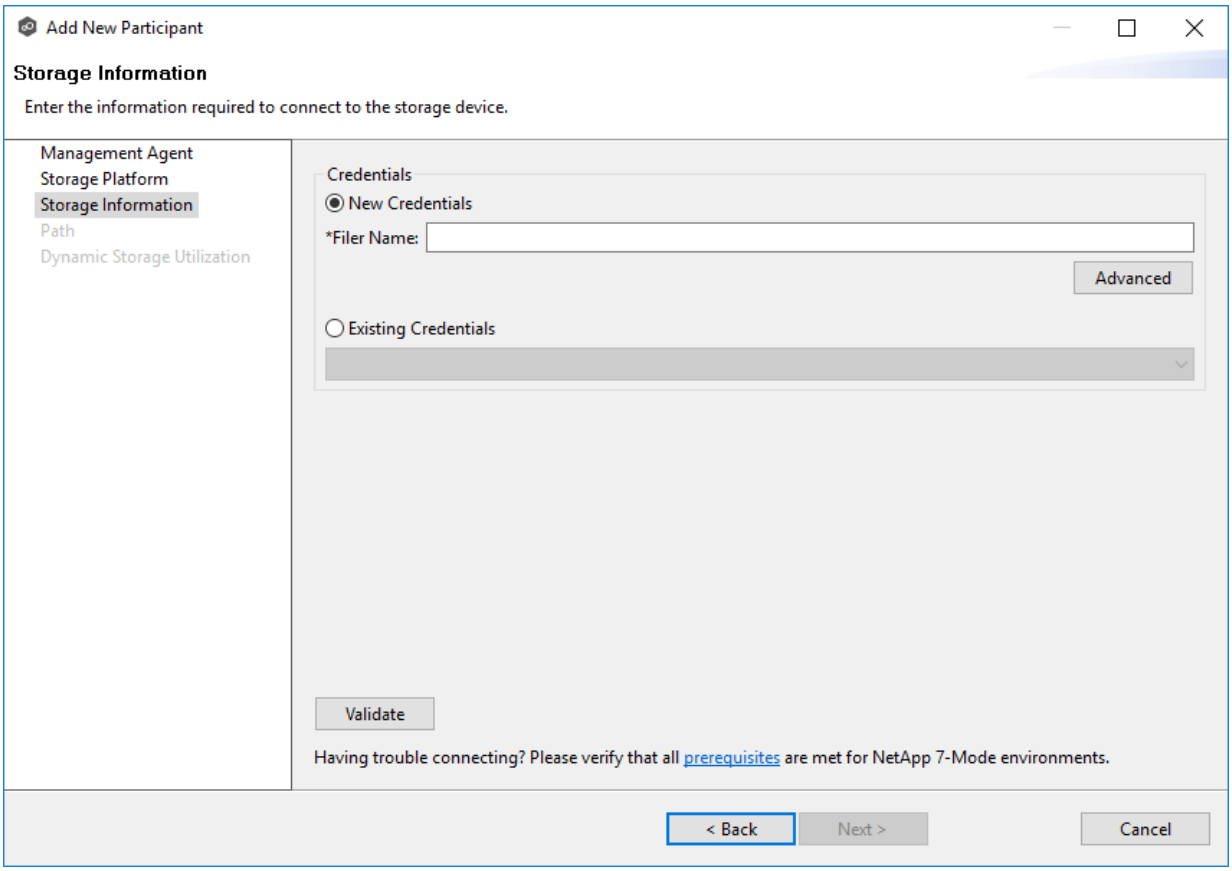

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
Filer Name	Enter the name of the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

Dell EMC Isilon

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect the EMC Isilon cluster hosting the data to be replicated.

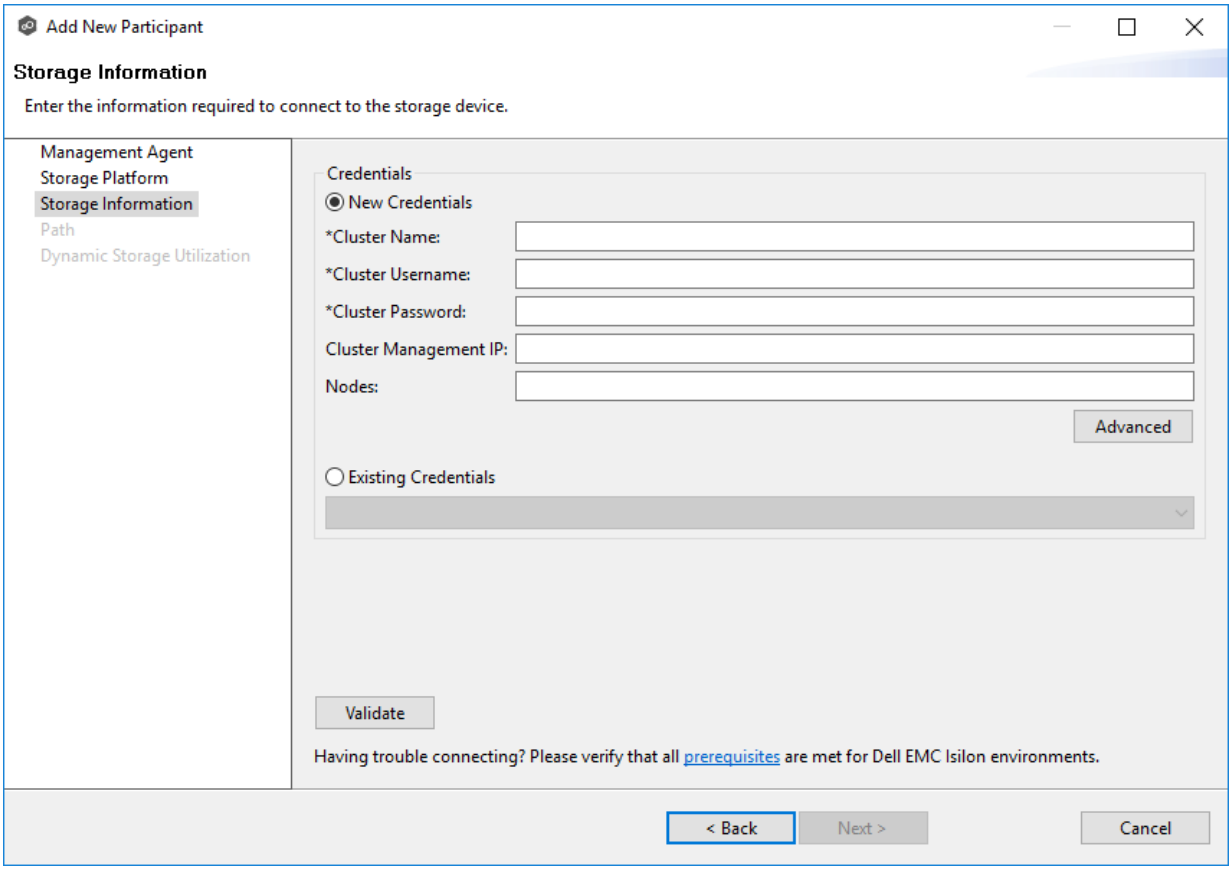

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
Cluster Name	Enter the name of the Dell EMC Isilon cluster hosting the data to be replicated.
Cluster Username	Enter the user name for the account managing the Dell EMC Isilon cluster.
Cluster Password	Enter the password for account managing the Dell EMC Isilon cluster.
Cluster Management IP	(Optional) Enter the IP address of the system used to manage the Dell EMC Isilon cluster. Required only if multiple Access Zones are in use on the cluster.
Nodes	(Optional) Enter one IP from each node in the Isilon cluster that the Agent can access to perform open file lookups. Use commas to separate nodes.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

Dell EMC Unity

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the NAS server hosting the data to be replicated.

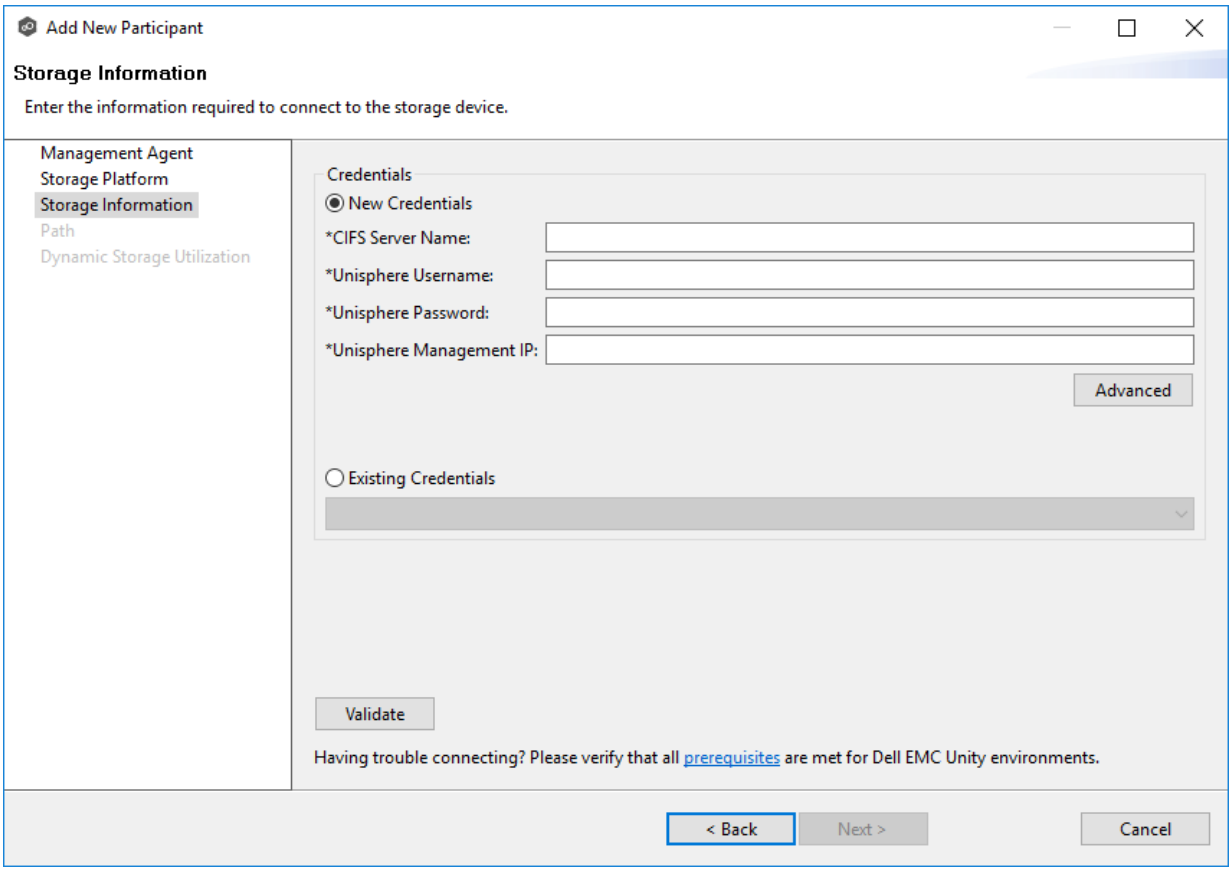

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
CIFS Server Name	Enter the name of the NAS server hosting the data to be replicated.
Unisphere Username	Enter the user name for the Unisphere account managing the Unity storage device.
Unisphere Password	Enter the password for the Unisphere account managing the Unity storage device.
Unisphere Management IP	Enter the IP address of the Unisphere system used to manage the Unity storage device. This should not point to the NAS server.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

Dell EMC Celerra | VNX | VNX 2

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the CIFS Server hosting the data to be replicated or select existing credentials.

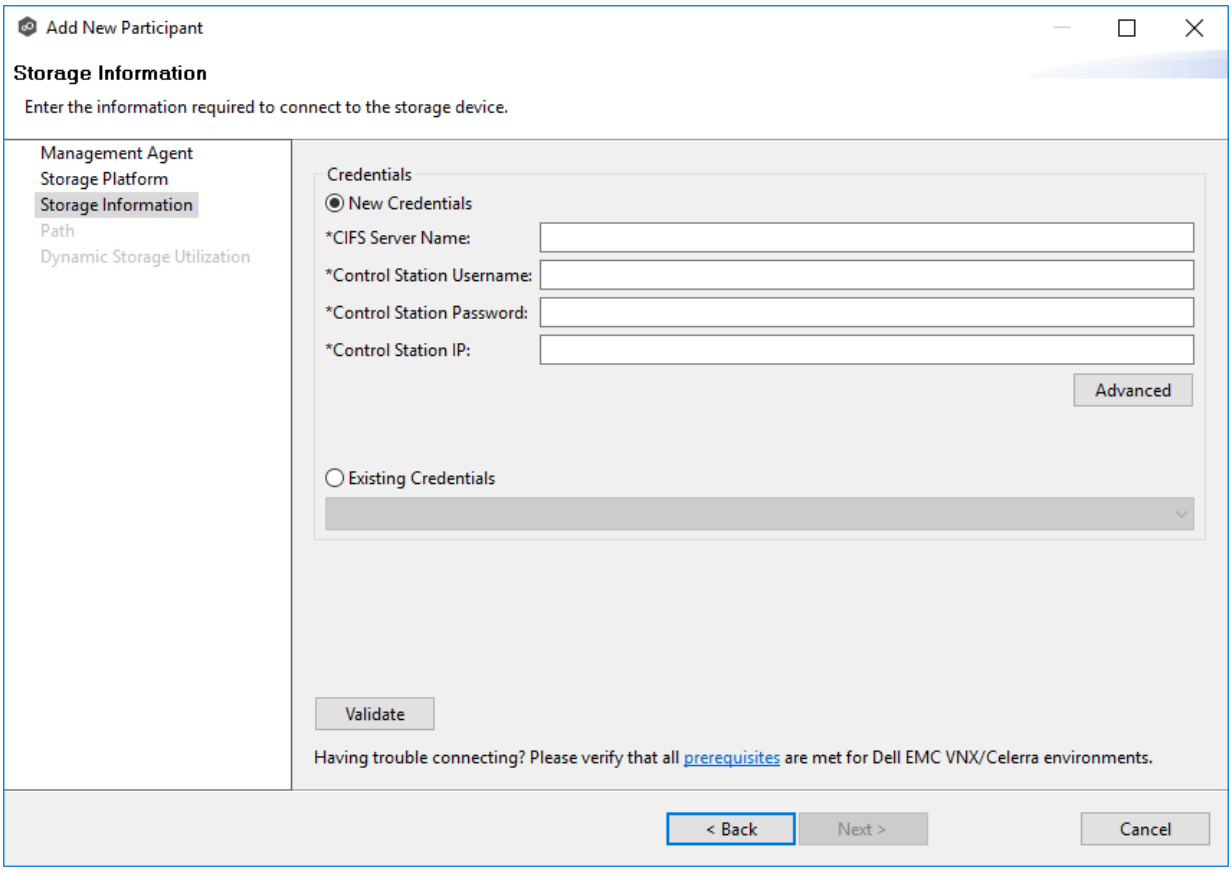

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
CIFS Server Name	Enter the name of the CIFS Server hosting the data to be replicated.
Control Station Username	Enter the user name for the Control Station account managing the Celerra/VNX storage device.
Control Station Password	Enter the password for the Control Station account managing the Celerra/VNX storage device.
Control Station IP	Enter the IP address of the Control Station system used to manage the Celerra/VNX storage device. This should not point to the CIFS Server.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

Nutanix Files

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated.

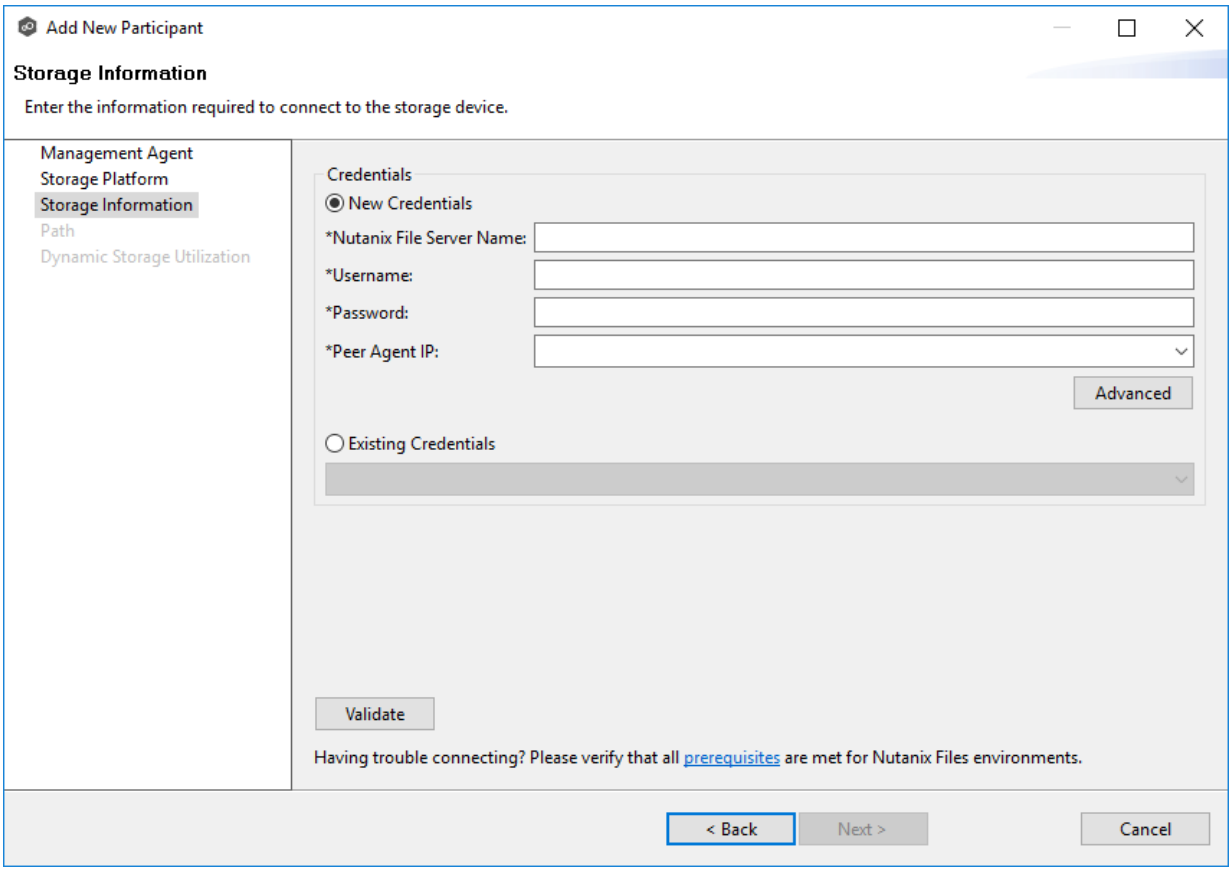

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address. The address should not point to the Files cluster itself.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

The **Path** page is where you specify the path to the volume/share/folder you want to replicate. This volume/share/folder is referred to as the watch set. The watch set can contain only a single volume/share/folder. If you want to replicate multiple volumes/shares/folders, you need to create a separate job for each one.

1. Browse to or enter the path to the watch set.

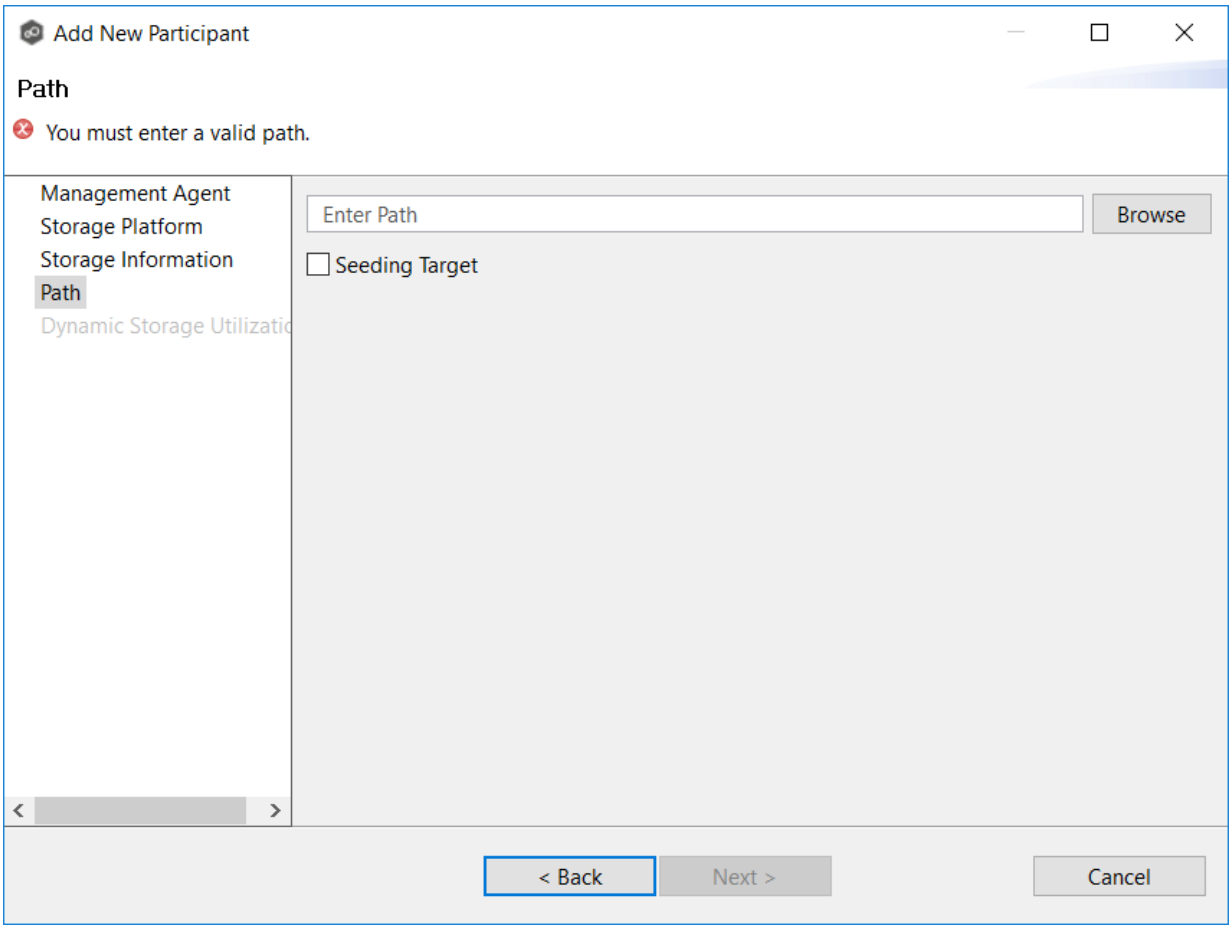

2. (Optional) Select the **Seeding Target** checkbox, and then click **OK** in the dialog that appears.

   If you select this option, a message describing seeding behavior is displayed. Multiple participants in a File Collaboration job can be set as smart data seeding targets; however, at least one participant should not be set as a smart data seeding target. All participants that are not set as seeding targets will become sources for the smart data seeding targets. For more information about smart data seeding, see Smart Data Seeding or contact support@peersoftware.com.

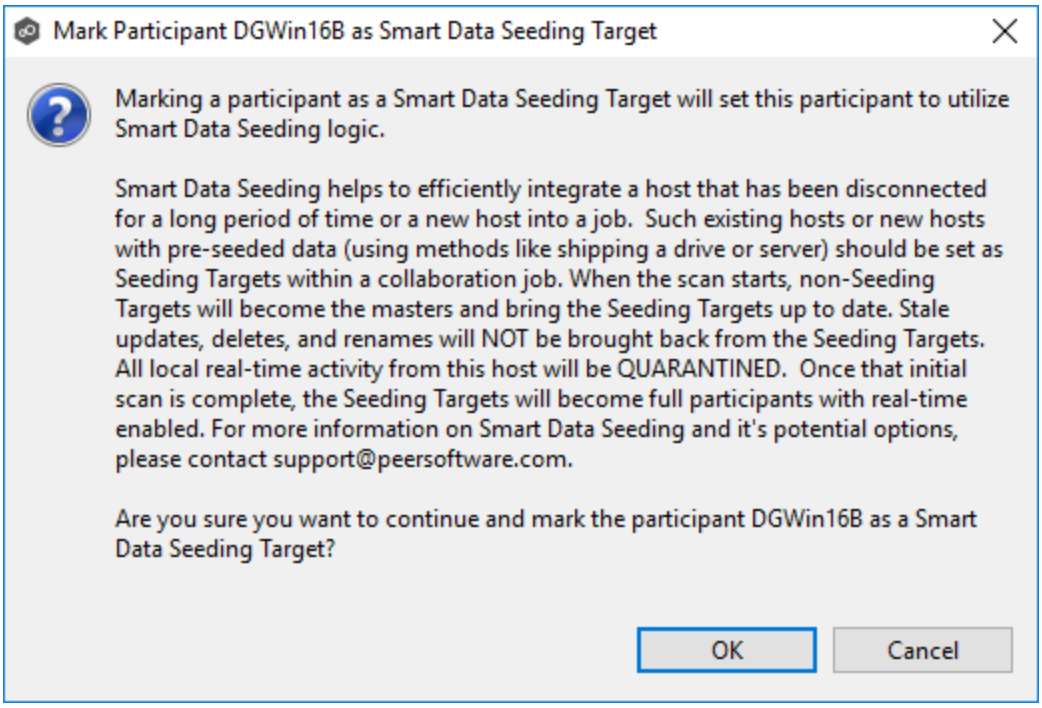

3. Click **Next** if you want this participant to use Dynamic Storage Utilization; otherwise, click **Finish** to complete the wizard for this participant.

   If you click **Finish**, return to Step 2: Participants to add more participants, if applicable. A File Collaboration job must have at least two participants.

Dynamic Storage Utilization (DSU) is a method for conserving space on storage devices by caching files until needed. DSU saves space by stubbing files and rehydrating them as needed. DSU is optional; if you don't need to conserve space on the storage device managed by the Agent, then you do not need to select this option.

If you enable Dynamic Storage Utilization (DSU) for a participant, you must designate the participant as either a **master** or **edge** participant.

- **Master participant** - A master participant always has complete set of files for that job. None of the files are stubbed; they are stored physically on that device.
- **Edge participant** - A subset of the files stored on a master participant are physically stored on an edge participant; the rest of the files on an edge participant are stub files that take up minimal space. DSU allows users to seamlessly retrieve stubbed files directly from a master participant as needed; when retrieved, the local stub file is rehydrated so that the full file is stored locally on the edge participant.

A job can have master and edge participants, as well as participants that don't have either role. If you do not choose to enable DSU for a participant, it will always have a full set of files like a master participant but will not be used to serve file content to any edge participants.

Notes:

- A participant can be a master participant for some jobs and an edge participant for other jobs.
- A job needs at least one master participant that isn't a seeding target. If there is only one master participant for the job, it should not be a seeding target.

For more information about DSU, see Dynamic Storage Utilization in Advanced Topics.

1. Select the **Enable Dynamic Storage Utilization** checkbox if you want this participant to be able to use Dynamic Storage Utilization; otherwise, click **Finish**.

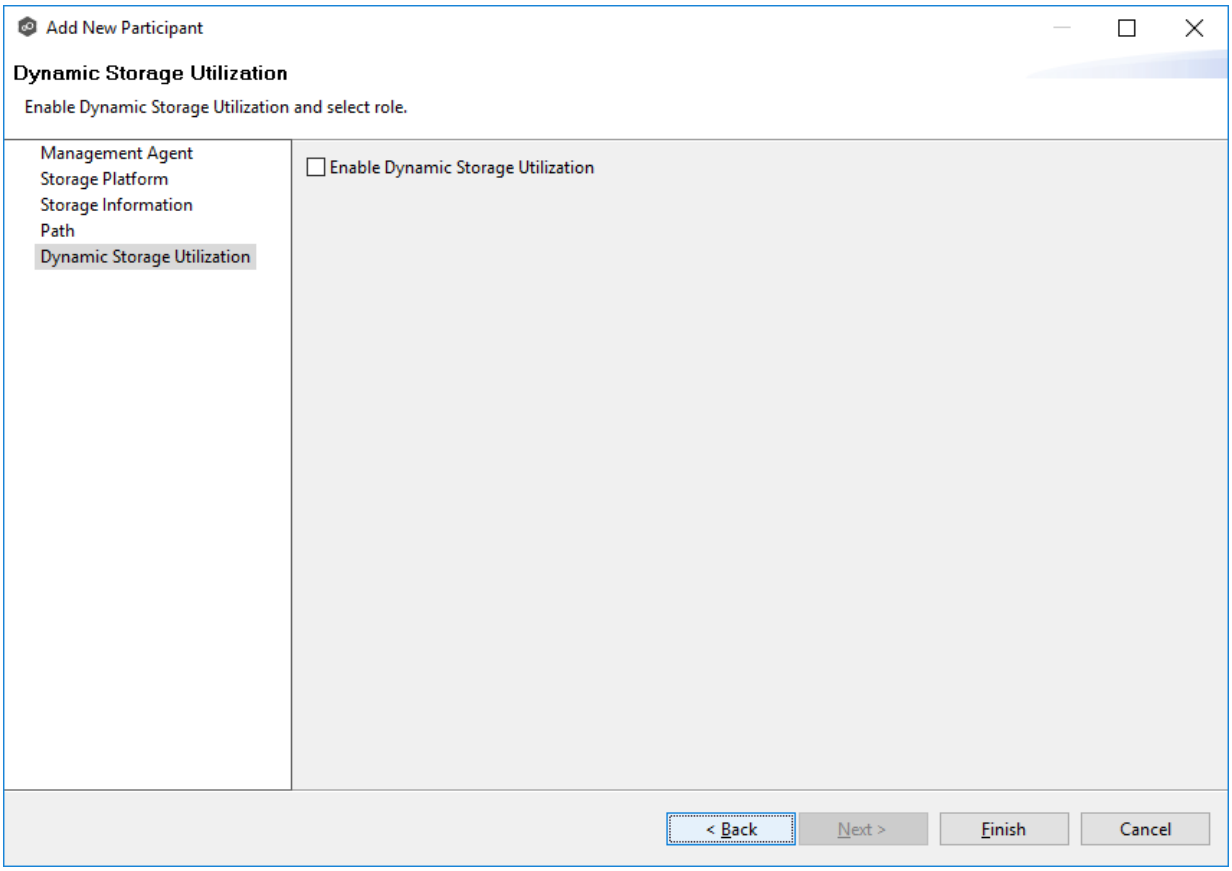

   If you enable DSU, the DSU role options are displayed; the **Master** role is selected by default.

2. Choose a DSU role for the participant:

- Choose **Master** if the storage device managed by the Agent will contain complete copies of all files for this job. Any type of storage platform can be a master participant.

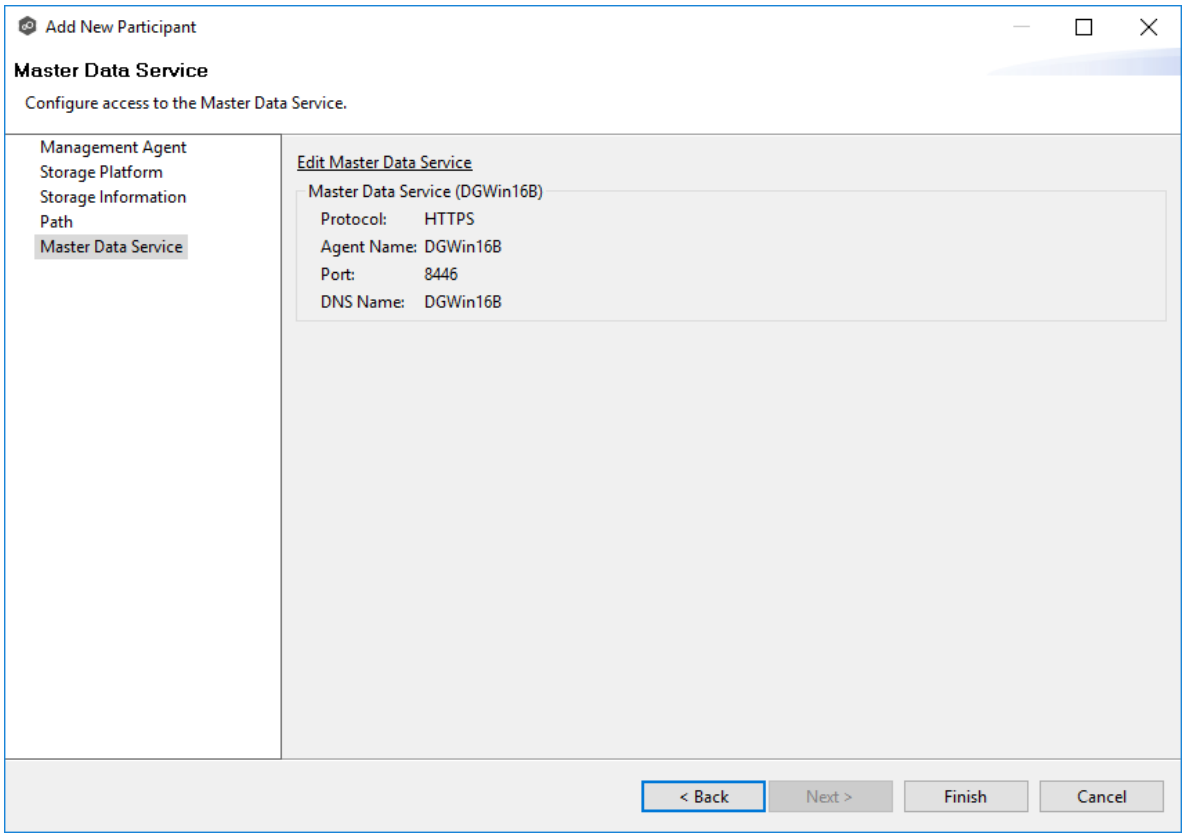

- Choose **Edge** if you want to conserve space on the storage device managed by the Agent. Only Windows File Servers can be an edge participant.

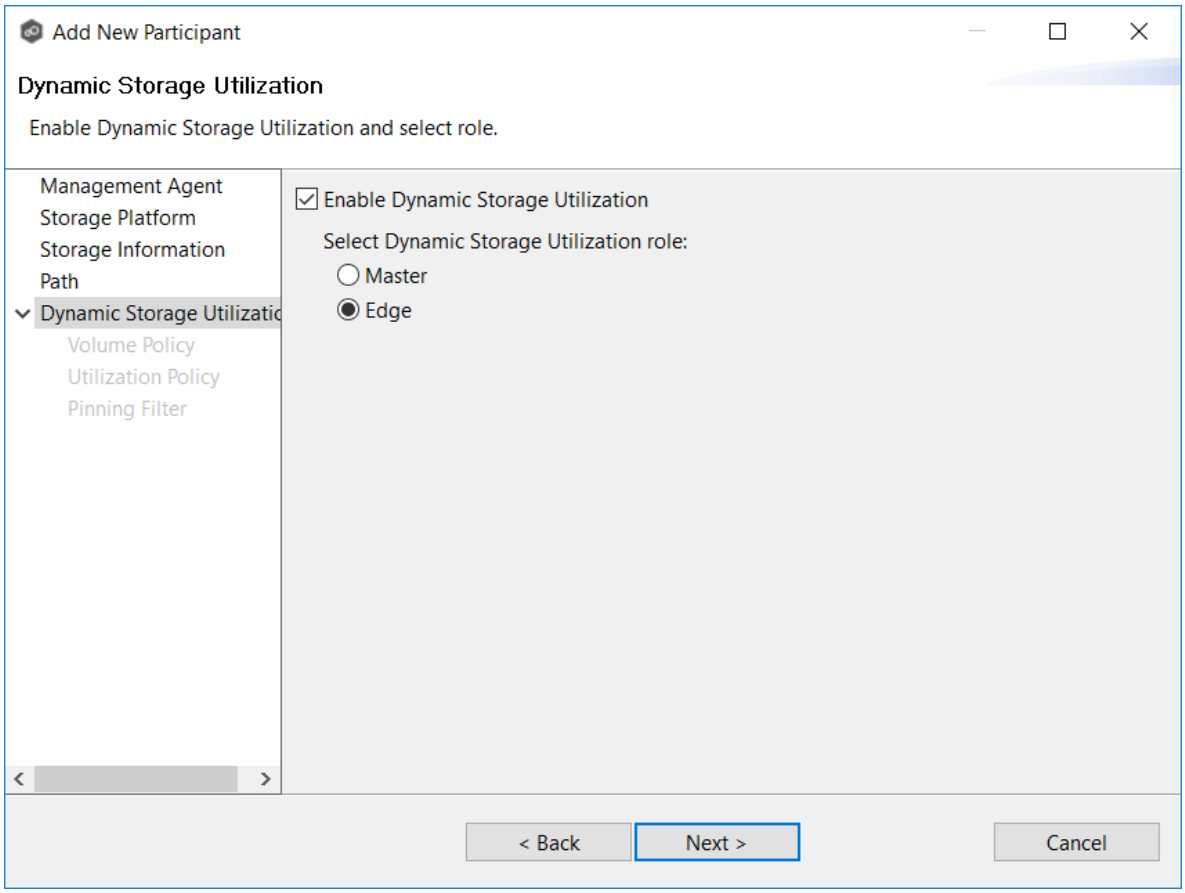

3. Click **Next**.

   - If you selected **Master**, continue with the Master Data Service page.
   - If you selected **Edge**, continue with the Volume Policy page.

Master Data Service

The **Master Data Service** page appears if you chose the master role for the participant. The Master Data Service handles requests from edge participants for files on a master participant. The Master Data Service is installed on the Agent server as part of the Peer Agent installation process.

The first two fields on this page are automatically populated:

- **Protocol**: This field lists the protocol that will be used to transfer file content between master participants and edge participants. HTTPS is currently the only available option as

it requires only a single open firewall port on each master participant and does a better job of handling latency over WAN-based connections.

- **Agent Name**: This field lists the name of the Management Agent that you selected at the beginning of Step 2.

1. (Optional) Enter a value for **Agent Alias**. The value can be a hostname, FDQN, or IP address.

   A value for this field is required only if the name of the Agent cannot be converted to an IP address via DNS. If an alias is entered, it will be used by the Edge Service on edge participants to connect with the Master Data Service. If no alias is entered, the Agent's name will be used.

2. (Optional) Modify the port number that the Master Data Service will listen on for this master participant.

   A default value for the port number, 8446, is set when the Agent is installed. If you modify the port number, the Master Data Service is started with the new port number.

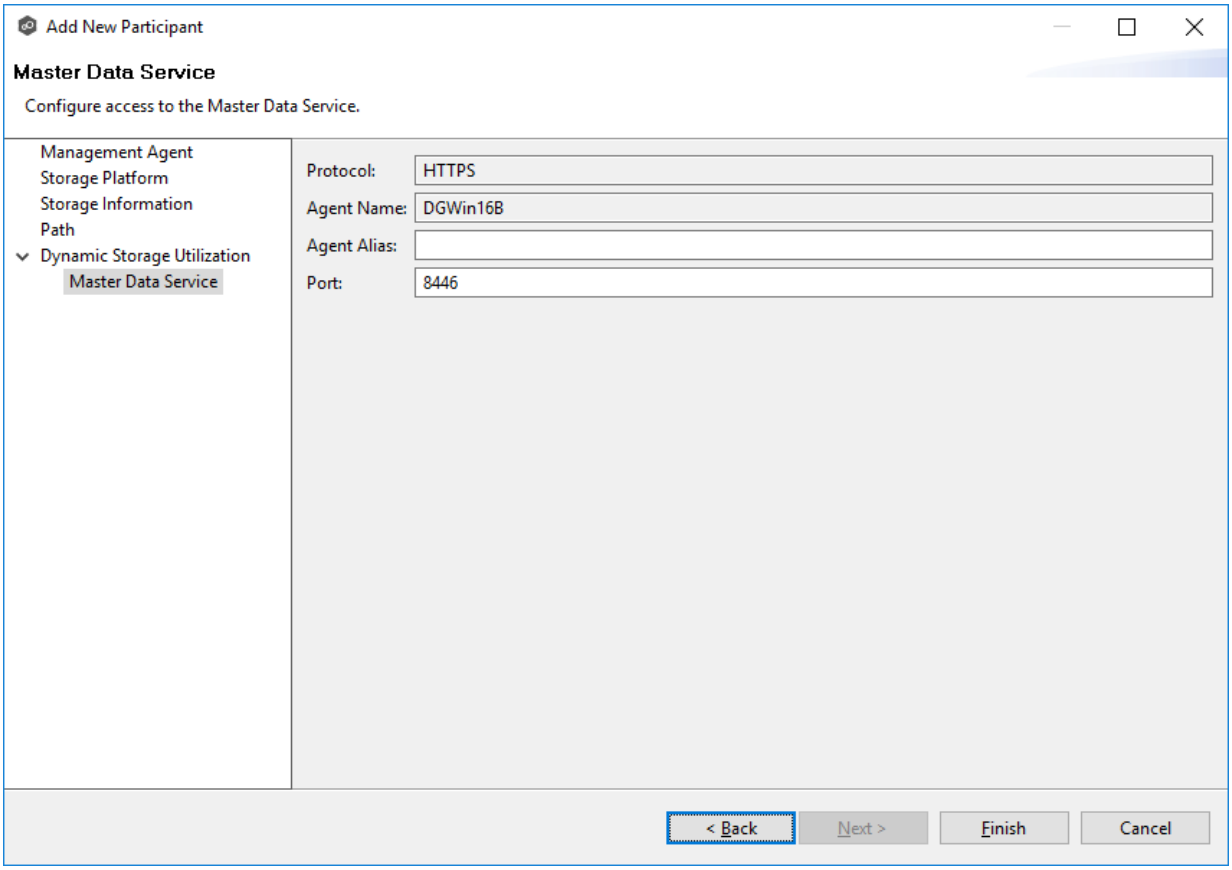

Note: If the Agent you selected is already being used as a master participant in another job utilizing DSU, then the existing Master Data Service parameters will be displayed. You can edit the values by clicking the **Edit Master Data Service** link. If you modify the

port number, the Master Data Service will be restarted and the new port number will take effect immediately. Any modifications you apply will be applied to every other job that use this Agent as a master participant.

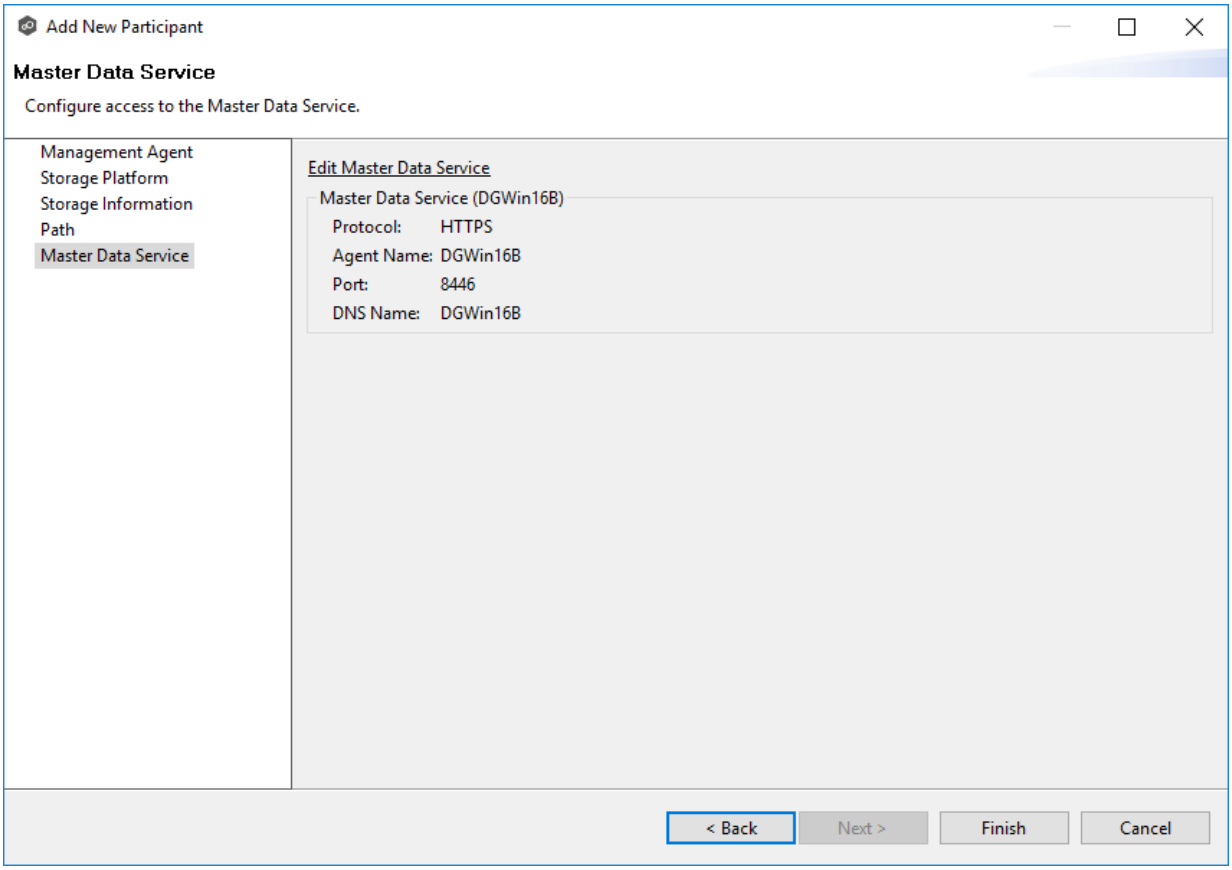

3. Click **Finish**.

The **Participants** page reappears. The participant is listed in the **Participants** table with the **Master** role.

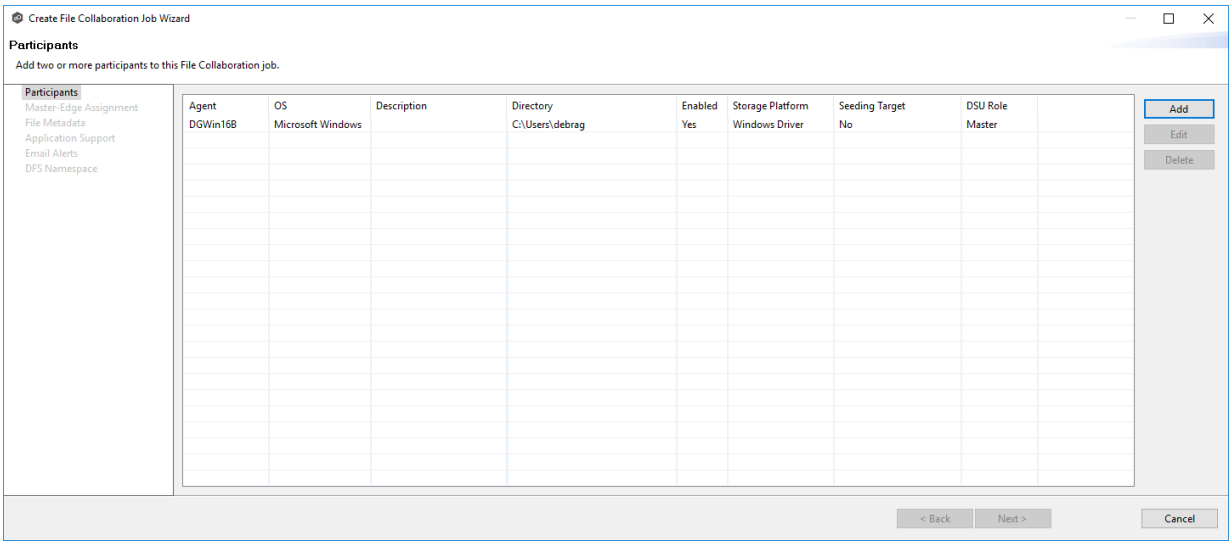

4. Continue adding more participants if applicable or continue with Step 4: Master-Edge Assignment.

Volume Policy

The **Volume Policy** page appears if you chose the edge role for the participant.

A volume policy is applied when a caching scan is run. The primary purpose of a **volume policy** is to specify how much space is available to DSU on a specific volume (or drive letter), i.e, to define the **cache size**. The cache size specifies the maximum amount of disk space you want to allocate to DSU for fully hydrated files on the volume specified by the path on the **Path** page. For example, if the participant is configured to monitor D:\Data, the volume policy for this participant would apply to the D volume.

The cache size can be specified as a percentage of the volume disk space or as a fixed size. For example, if an edge participant is configured to monitor a volume that has 1 TB of disk space, and you tell DSU to use 75% of that volume, then up to 750 GB of files could be locally available on the volume monitored by that edge participant. For optimal performance, we recommend that this cache be dedicated to DSU's use on this volume.

A volume policy applies to each job where the following three elements are true:

- DSU is enabled for the job.
- The participant is an edge participant.
- The paths specified for each job share the same volume.

To create a volume policy:

1. In the **Cache Size** section, choose an option for setting the cache size:

   - Use up to X % of this volume
   - Use up to X size of this volume

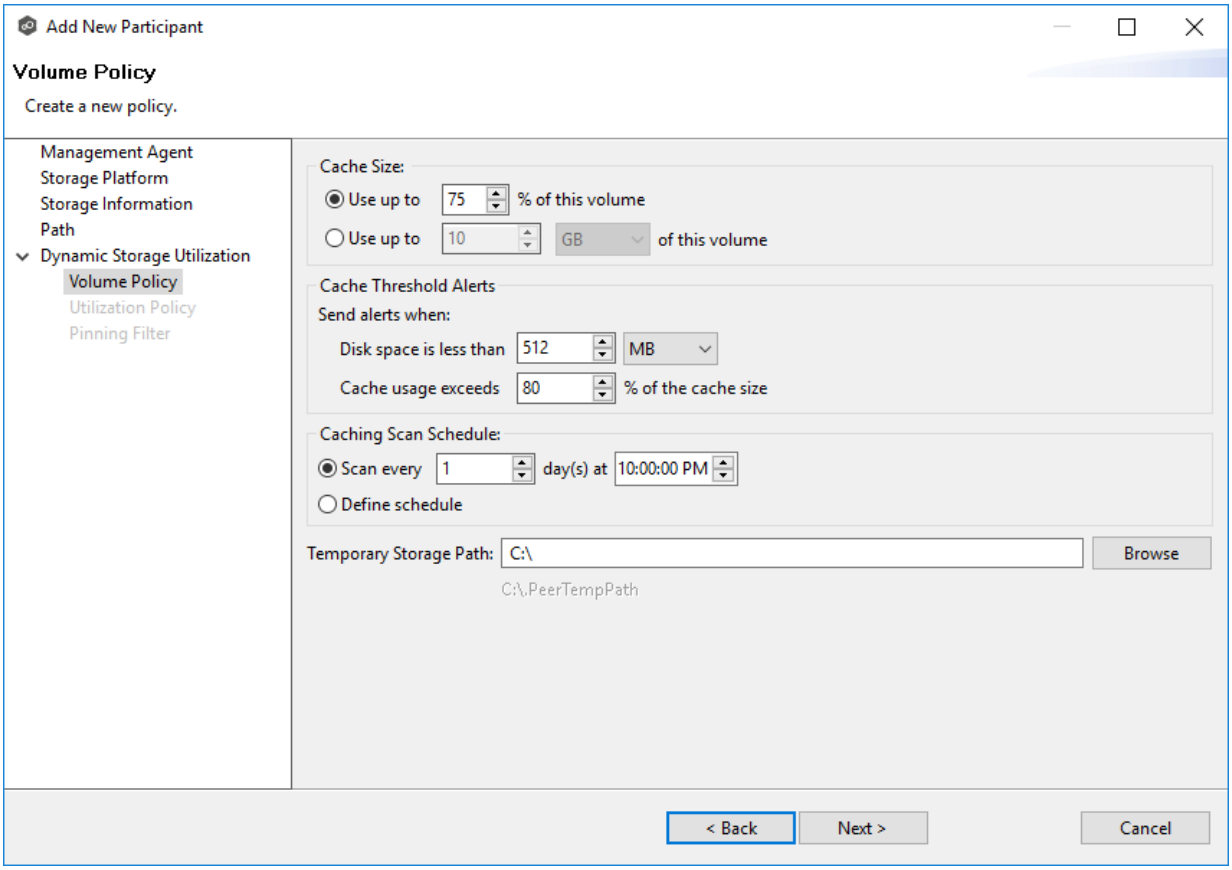

2. In the **Cache Threshold Alerts** section, set threshold values for automatic alerts about free disk space and cache usage.

   Peer Management Center will automatically display alerts in the **Alerts** tab:

   - The amount of free disk space on the volume falls below the specified value. For example, if a 1 TB volume has 500 MB of free space and the threshold is set to 512 MB, an alert will be sent.
   - Cache usage on the volume exceeds the specified percentage of the cache size. For example, if the cache size is set to 80%, equating to 750 GB, DSU will start sending alerts when it has used 600 GB.

You can also send cache threshhold alerts via email alerts and SNMP notifications. You configure these in Dynamic Storage Utilization preferences for File Collaboration and File Synchronization jobs.

3. In the **Caching Scan Schedule**, set the frequency that the volume should be scanned to optimize the balance of fully hydrated vs stubbed files.

   This scan can be run daily at a specified time or you can define a more customized schedule.

4. In the **Temporary Storage Path** field, enter a path or browse to the location you want to be used as temporary storage space.

   The temporary storage space will be used to store the content of stub files as they are are being rehydrated. The content of files undergoing rehydration are referred to as **file blocks**. File blocks are fixed-length chunks of data that are read into memory when requested by an application. DSU will create a subfolder named **.PeerTempPath** under the location that you specify and temporarily store the file blocks in that subfolder.

   For optimal performance, we recommend that the temporary storage space be on the same volume as the watch set. If that is not possible, it should be on a high performance disk.

5. Click **Next**.

   The Utilization Policy page appears.

Note: If the Agent you selected is already being used as an edge participant in another job utilizing DSU, the existing volume policy will be displayed on this page. You can edit the existing volume policy; however, be aware that any modifications to the volume policy will be applied to every other job that uses this Agent as an edge participant and "touches" the same volume.

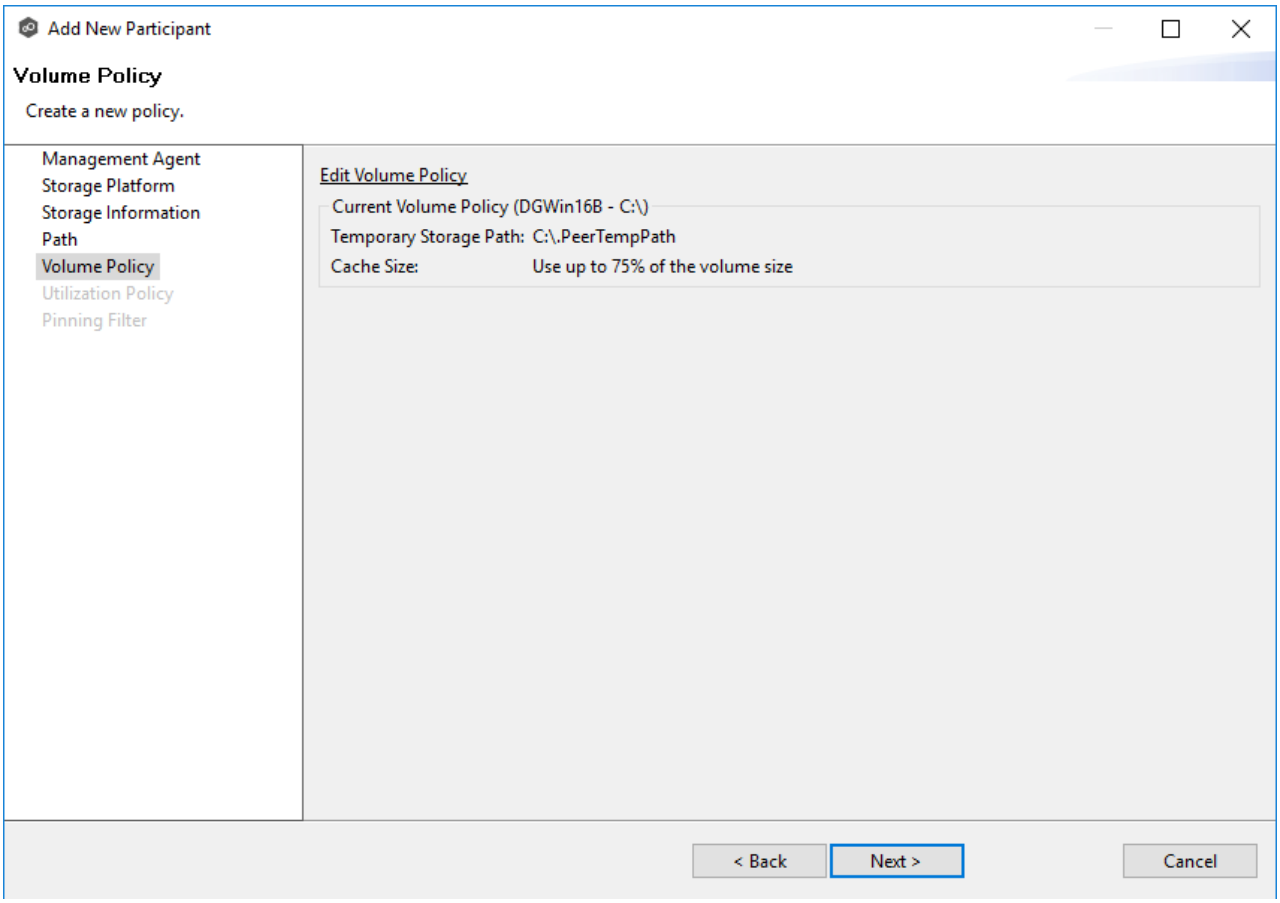

Utilization Policy

The **Utilization Policy** page appears if you chose the edge role for the participant. The primary purpose of a **utilization policy** is to specify the parameters that govern when files on this edge participant should be stubbed or fully hydrated. Whereas the volume policy controls how much space is available to DSU on a specific volume (or drive letter), the utilization policy controls whether to stub or hydrate a file.

Utilization policy parameters are based on the size of the files to be potentially stubbed and when they were last accessed and modified. A utilization policy enables you to balance getting the best performance while keeping the cache as full as possible.

You can select an existing utilization policy to apply to the job or create a new utilization policy. Whereas a volume policy is specific to a volume, a utilization policy can be reused for multiple jobs.

1. Select **New Policy** or **Existing Policy**.

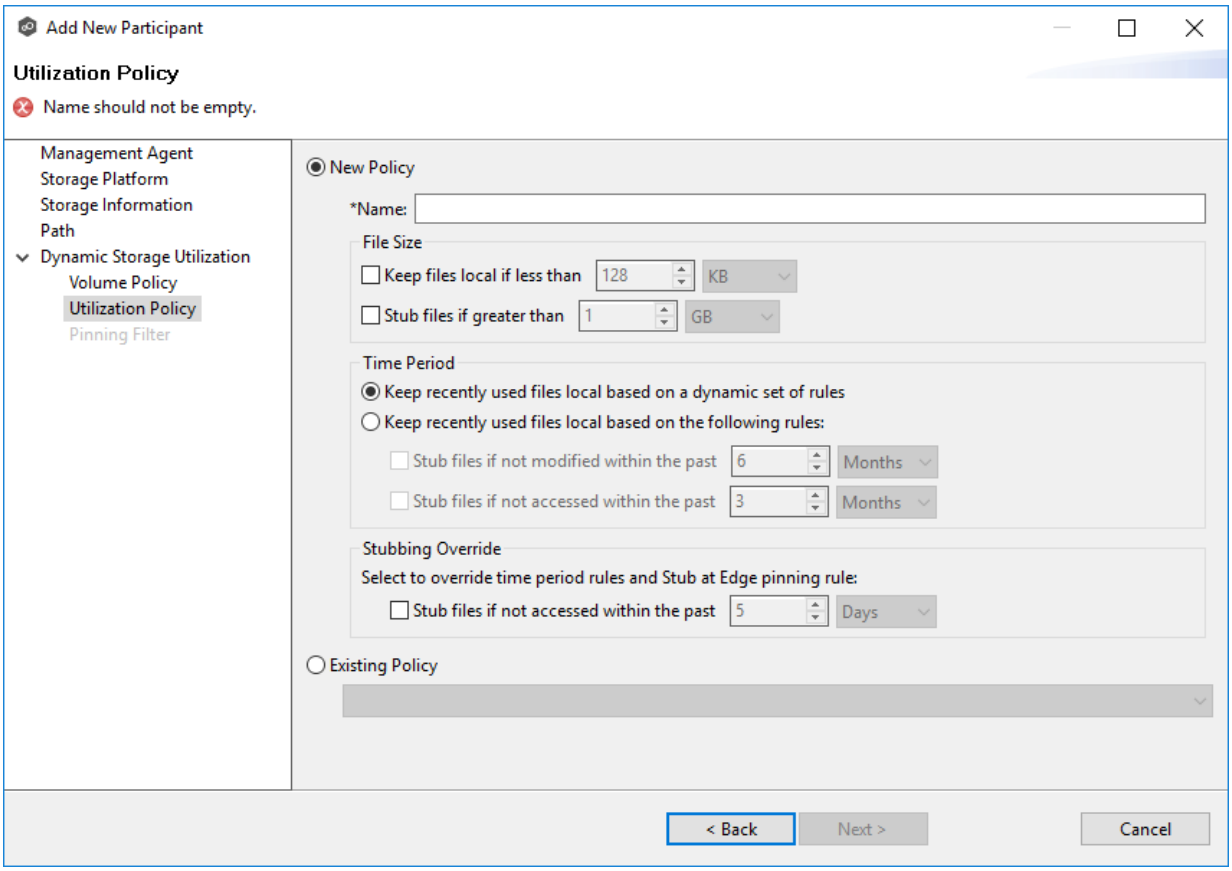

2. If you selected **Existing Policy**, select the policy, and then click **Next**.

   If you selected **New Policy**, enter a name for the policy.

3. (Optional) In the **File Size** section, select one or both options:

Field	Description
Keep files local if less than X size	Select this option if you want files under a specified size to remain local.
Stub files if greater than X size	Select this option if you want files over a specified size to be stubbed.

4. (Optional) In the **Time Period** section, select one of the options:

Field	Description
Keep recently used files local based on a dynamic set of rules	Select this option if you want DSU to control when to stub files based on last accessed and last modified times. DSU dynamically adjusts the accessed and modified time periods it uses based on the total amount of space that DSU is actively using on a volume.
Keep recently used files local based on the following rules	Select this option if you want to specify the dates used when stubbing files based on the last accessed and last modified.

5. If you selected the **Keep recently used files local based on the following rules** option in **Time Period**, two additional options are enabled; select the options to be used and specify the time periods to be used:

Field	Description
Stub files if not modified within the past X time period	Select this option and specify a time range to use the last modified date of a file to determine if it should be kept local or stubbed.
Stub files is not accessed within the past X time period	Select this option and specify a time range to use the last accessed date of a file to determine if it should be kept local or stubbed.

6. (Optional) In the **Stubbing Override** section, select the checkbox and a time period if you want to use the last accessed date of all recently hydrated files to determine if they should be kept local or re-stubbed.

7. Click **Next** or **Finish**.

   If you click **Next**, the Pinning Filter page appears.

Pinning Filter

The **Pinning Filter** page allows you to create a new pinning filter or select an existing pinning filter to apply to the job. A **pinning filter** specifies whether specific files or files in a particular

directory are always stubbed or always local on an edge participant. A pinning filter similar is to a utilization policy—it can be applied to multiple jobs. If there is a conflict between a pinning filter and utilization policy (where, for example, you might have something set to be always stubbed), the pinning filter will take precedence. Pinning filters are optional.

1. Select one of the options: **No Filter**, **New Filter**, or **Existing Filter**.

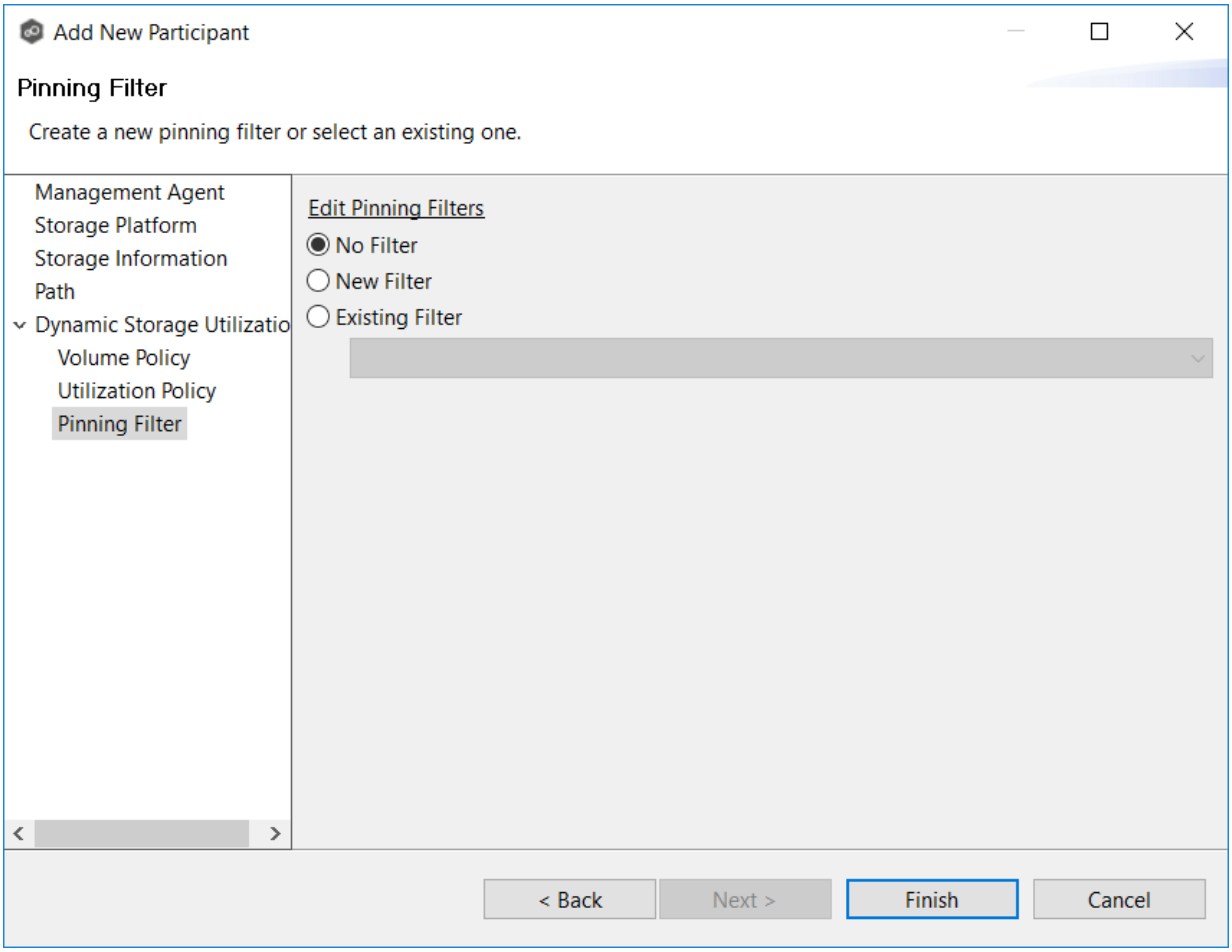

2. If you selected **No Filter**, click **Finish**; if you selected **Existing Filter**, select the filter, and then click **Finish**.

   If you selected **New Filter**, enter a name for the filter.

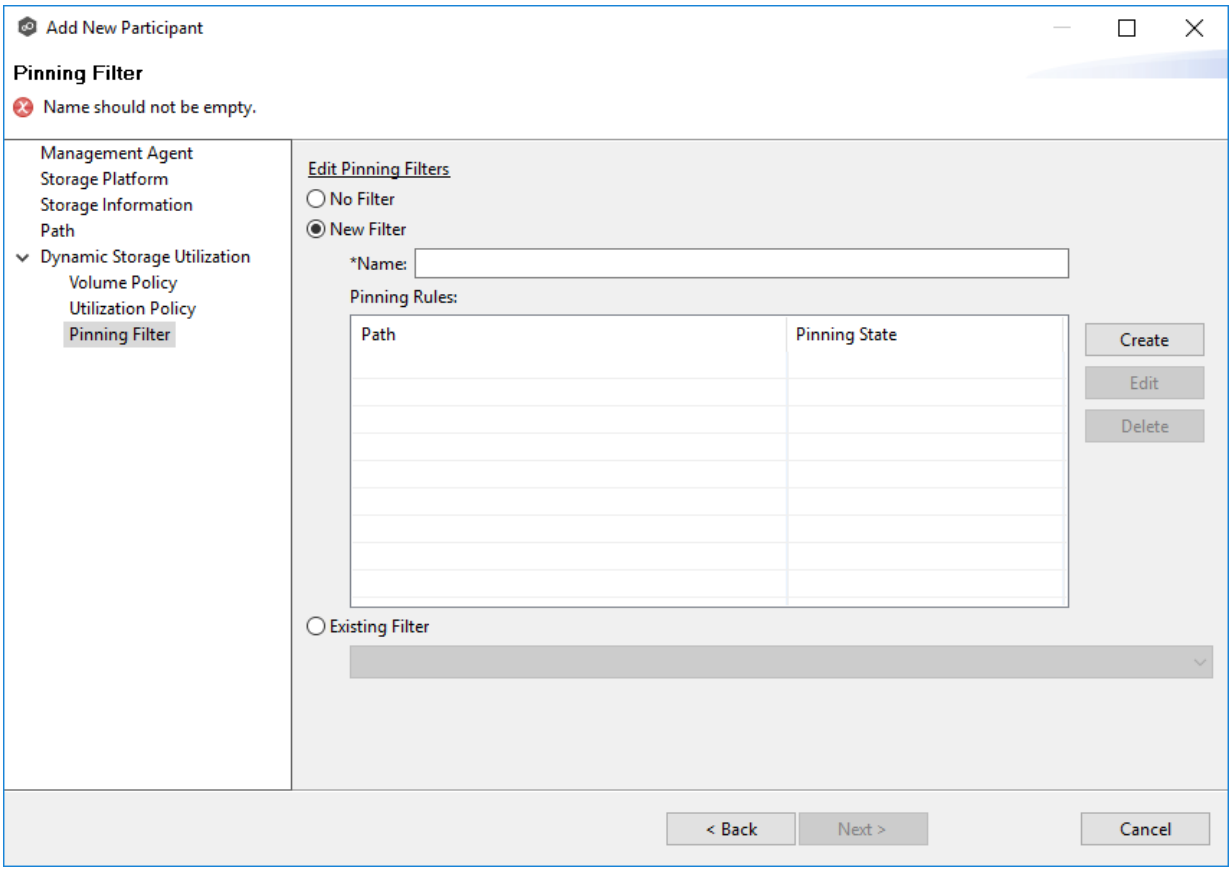

t

3. Enter a name for the filter.
4. Click **Create**.

   The **Create Pinning Rule** dialog appears.

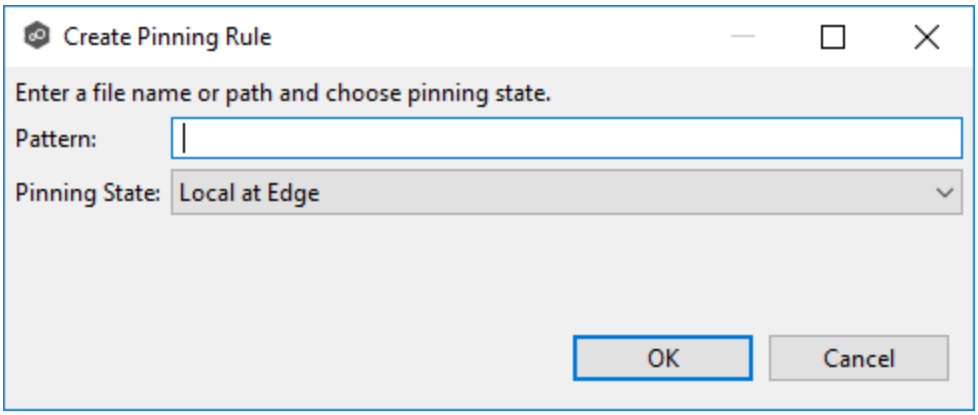

5. Enter a file name or path in the **Pattern** field and then choose a pinning state: **Local at Edge** or **Stubbed at Edge**.

6. Click **OK**.

   The rule appears in the filter table.

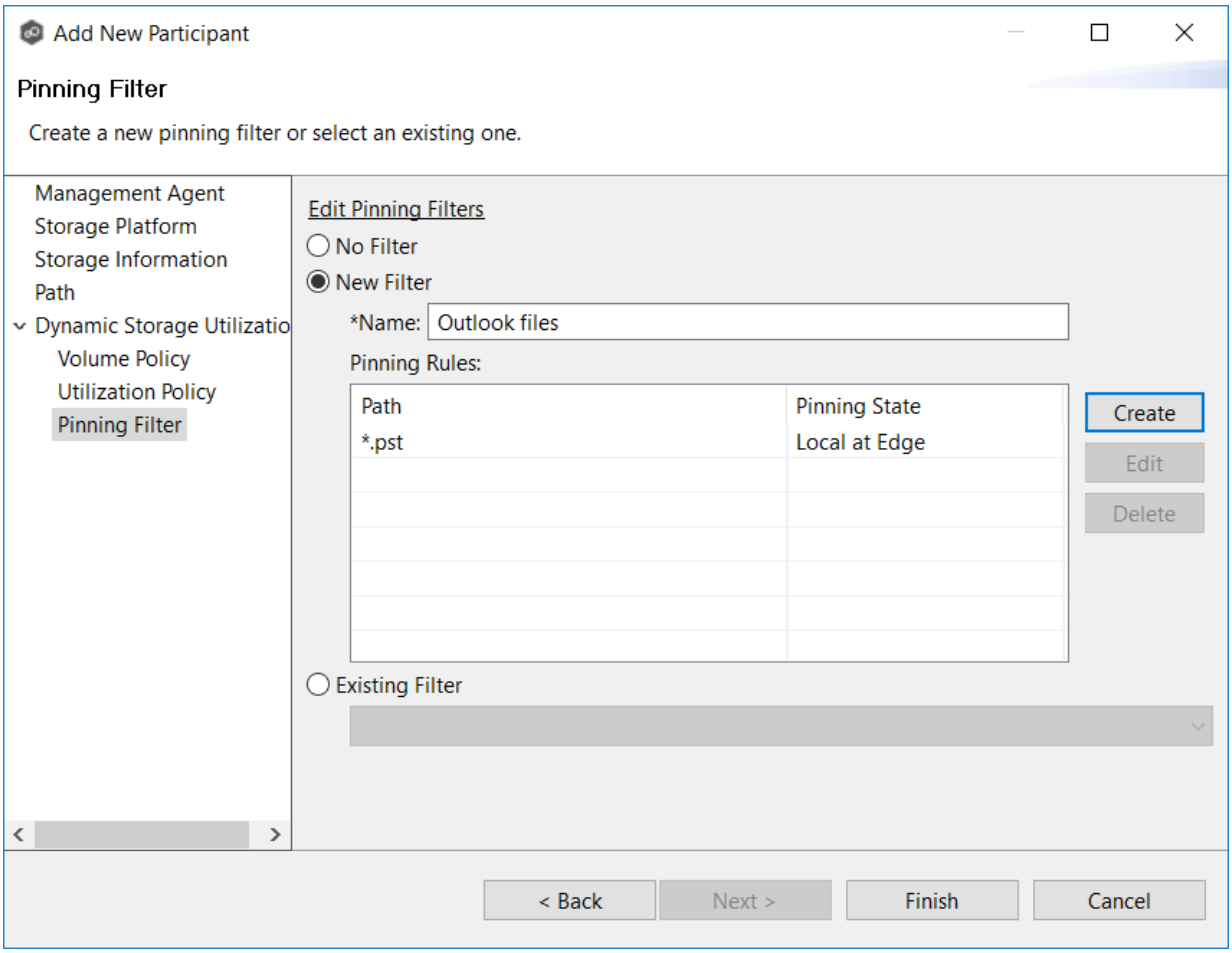

7. (Optional) Create additional pinning rules.

8. Click **Finish**.

   The **Participants** page reappears. The participant is listed in the **Participants** table with the **Edge** role.

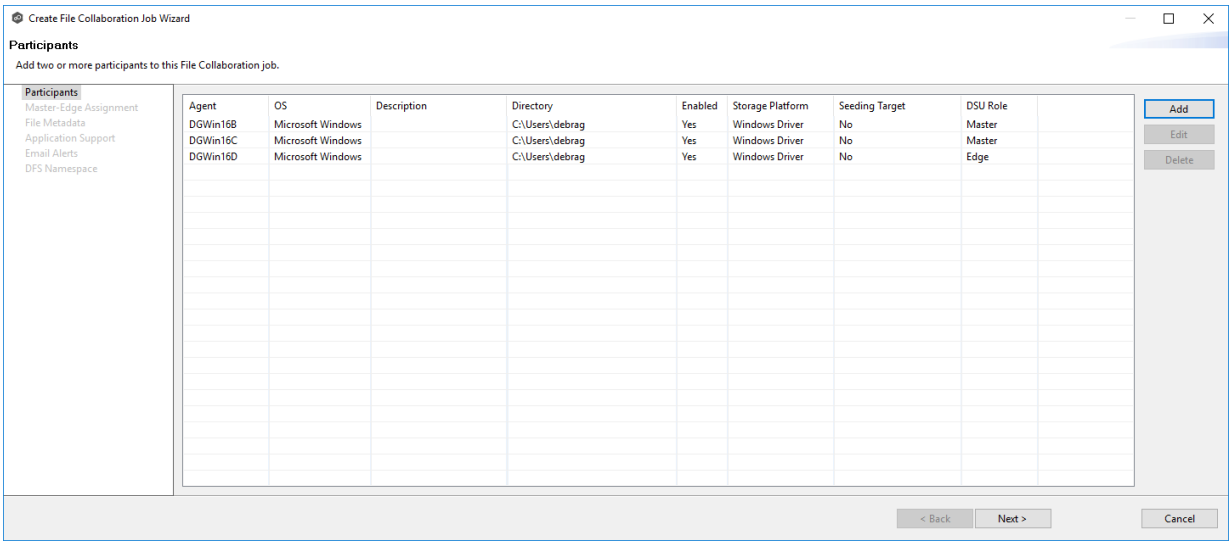

9. Continue adding more participants if applicable or continue with Step 4: Master-Edge Assignment

Step 4: Master-Edge Assignment

This step is optional.

The **Master-Edge Assignment** page appears only if you enabled Dynamic Storage Utilization for one or more participants in Step 2. The purpose of this page is to allow you to assign one or more master participants to each edge participant.

Every edge participant must have at least one master participant assigned to it, so that DSU will know which master participants to use when reading or rehydrating a stub file on an edge participant. For each edge participant, you want to assign the master participant that is the fastest and closest to it. This will increase the speed of rehydrating a stub file.

It is highly recommended that you assign, if possible, more than one master participant to an edge participant, so that if one master participant is not available, another master participant is able to provide the file content to the edge participant. You can set the order in which the master participants are consulted.

If you have only one edge participant and one master participant, the master participant is automatically assigned to the edge participant and listed in the Master Participants column. If you have multiple master participants, you need to explicitly assign master participants to each edge participant and then set the failover order.

1. Select an edge participant in the **Assignment** table.

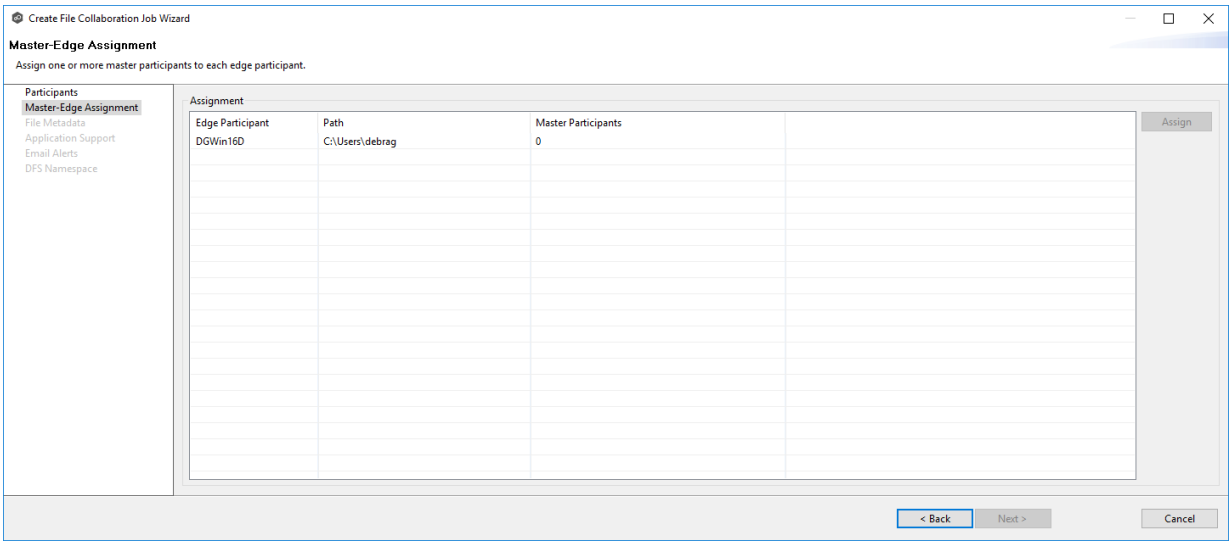

2. Click the **Assign** button.

   The **Assign Master Participants** dialog appears.

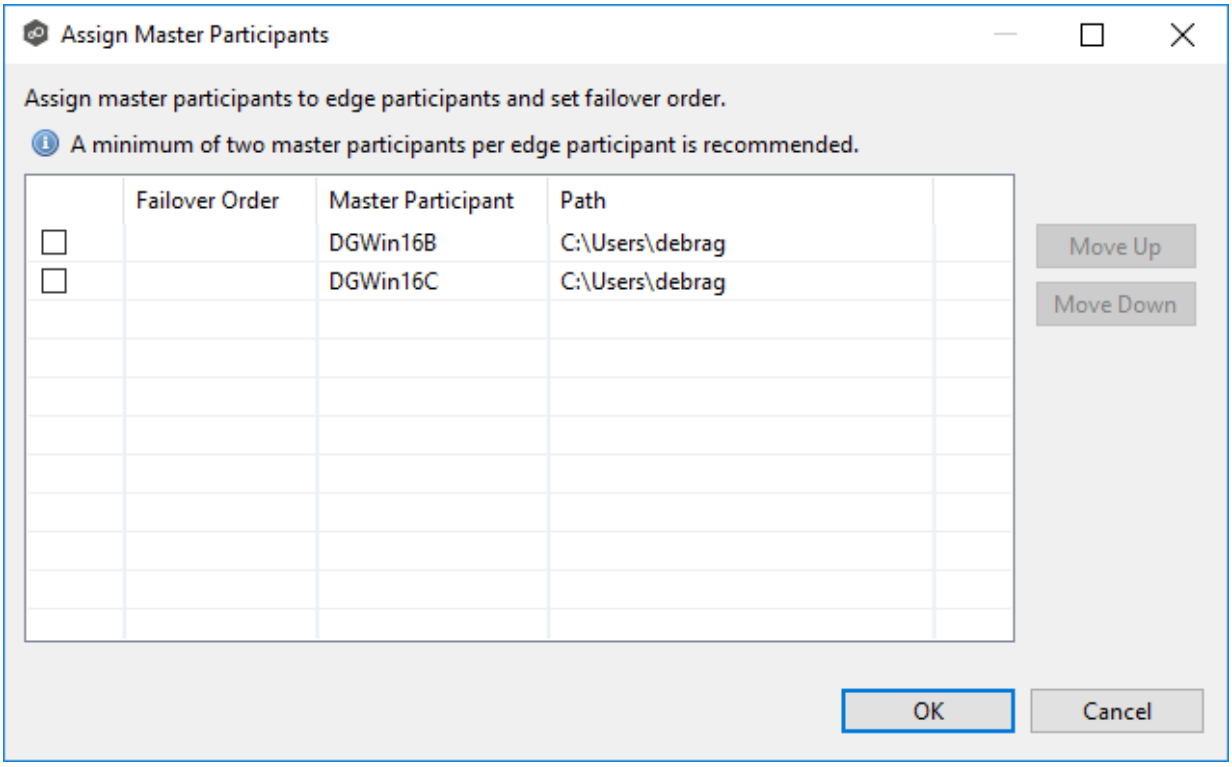

3. Select the master participants you want to assign to the edge participant.

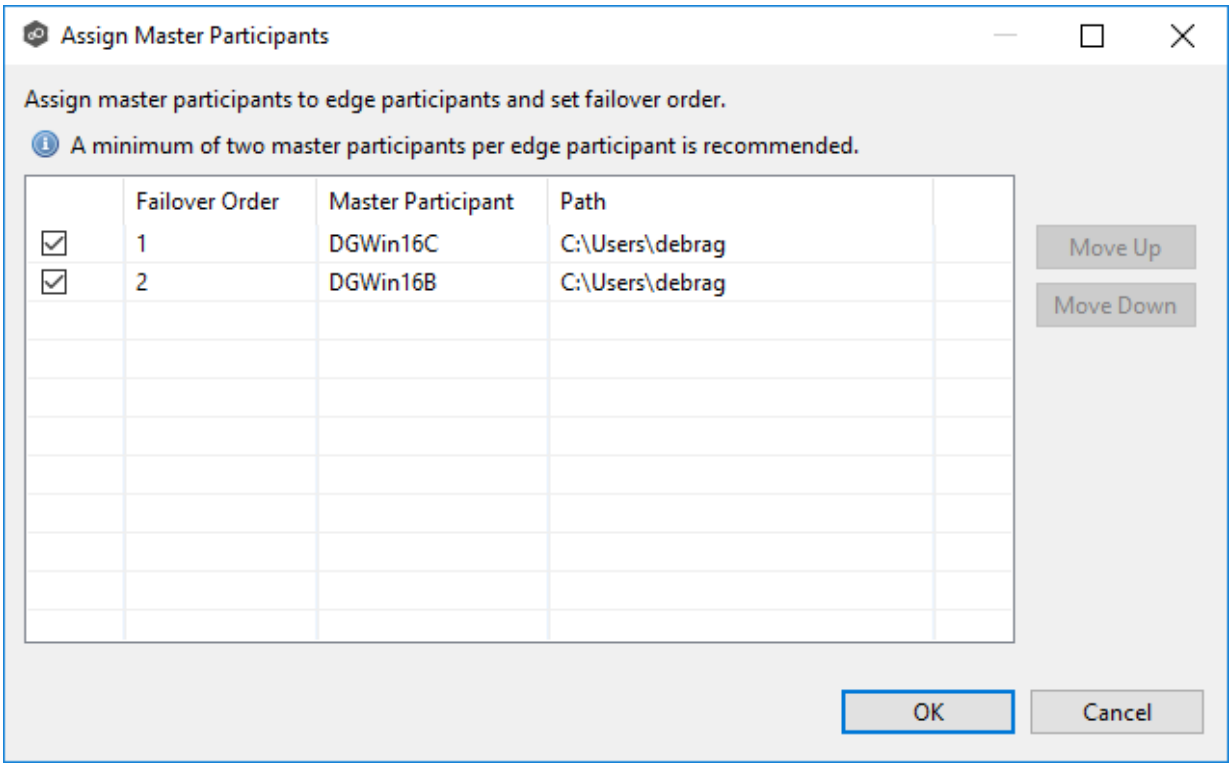

4. (Optional) Change the failover order by selecting a master participant and using the **Move Up** and **Move Down** buttons.

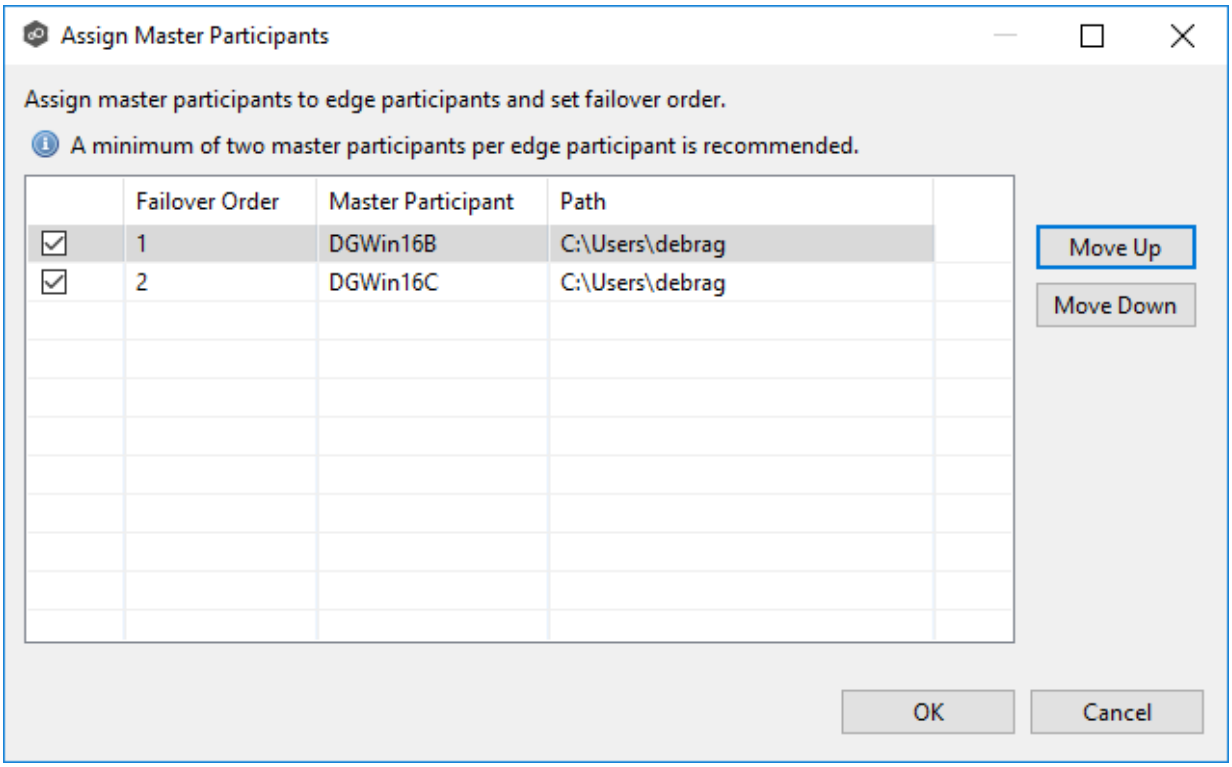

5. Click **OK**.

   The **Master Participants** column has been updated for that particpant.

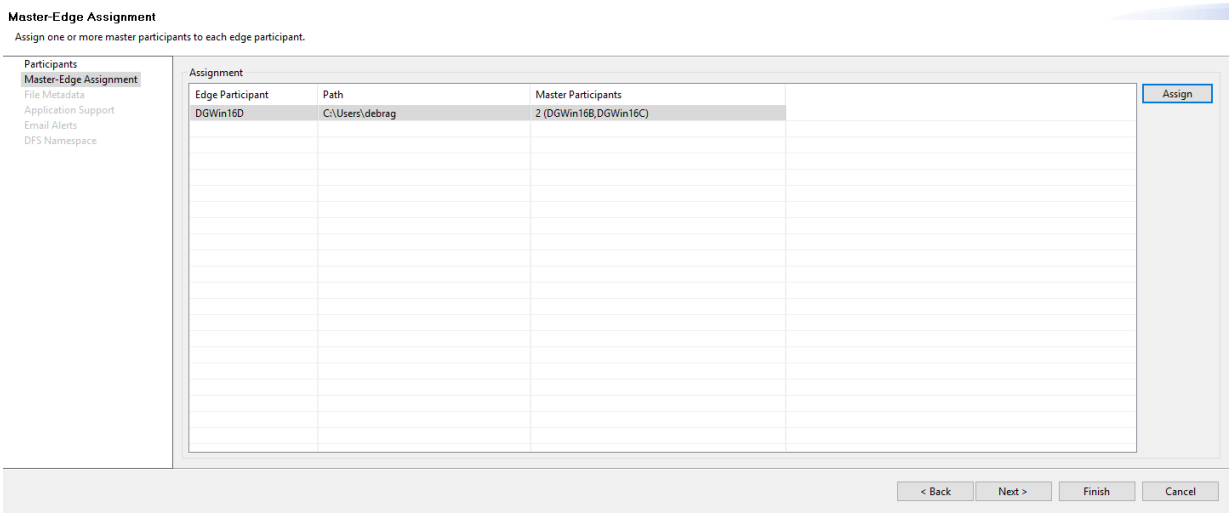

6. Repeat Steps 1-5 for each edge participant.

7. Click **Next**.

   The File Metadata page appears.

Step 5: File Metadata

This step is optional.

The **File Metadata** page allows you to specify whether you want to synchronize NTFS security permissions metadata and the types of metadata. It also allows you to specify which volume/share/folder's metadata should be used if there is a conflict during the initial synchronization. The volume/share/folder used if there is a conflict is referred to as the master host.

For more information on synchronizing NTFS metadata, see File Metadata Synchronization in the Advanced Topics section.

To enable file metadata synchronization:

1. Select when you want the metadata synchronized (you can select one or both options):

   - **Enable synchronizing NTFS security descriptors (ACLs) in real-time** - Select this option if you want the metadata synchronized in real-time. If enabled, changes

to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized to all participants as they occur.

- **Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan** - Select this option if you want the metadata synchronized during the initial scan. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized during the initial scan.

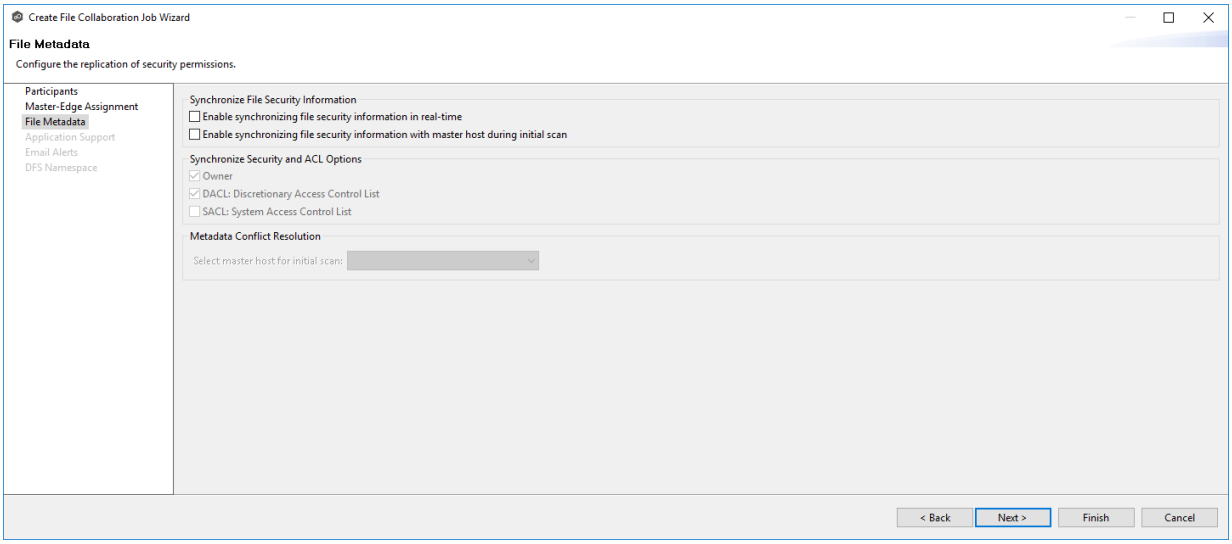

2. Click **OK** in the message that appears after selecting a metadata option.

3. If you selected either of the first two options in the **Synchronize Security Descriptor Options** section, select the security descriptor components (Owner, DACL, and SACL) to be synchronized.

4. If you selected the option for metadata synchronization during the initial scan, select the host to be used as the master host in case of metadata conflict.

   If a master host is not selected, then no metadata synchronization will be performed during the initial scan. If one or more security descriptors do not match across participants during the initial scan, conflict resolution will use permissions from the designated master host as the winner. If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

5. Click **Next**.

   The Application Support page is displayed.

Step 6: Application Support

This step is optional.

A File Collaboration job can be automatically optimized to work with specific applications. Optimization is performed automatically for all watch sets of all participants in the job.

The **Application Support** page lists the file types for which Peer Software has known best practices, which include filtering recommendations, the prioritization of certain file types, and the enabling or disabling of file locking. However, if an application is not listed, this does not mean that the application is not supported.

For details about how an application is optimized, contact support@peersoftware.com.

1. Select the applications that have files in the job's watch set.

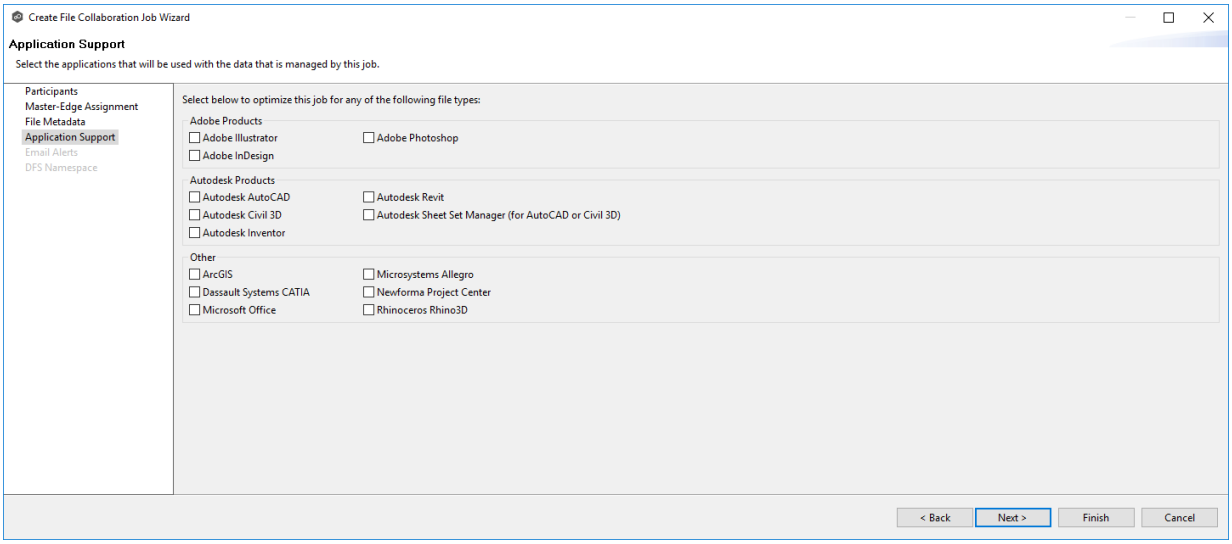

2. Click **Next**.

   The Email Alerts page is displayed.

Step 7: Email Alerts

This step is optional.

An email alert notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have

been applied to the job. When you first create a job, this list is empty. Email alerts are defined in Preferences and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply.

To create a new alert, see Email Alerts in the Preferences section.

To apply an existing email alert to the job.

1. Click the **Select** button.

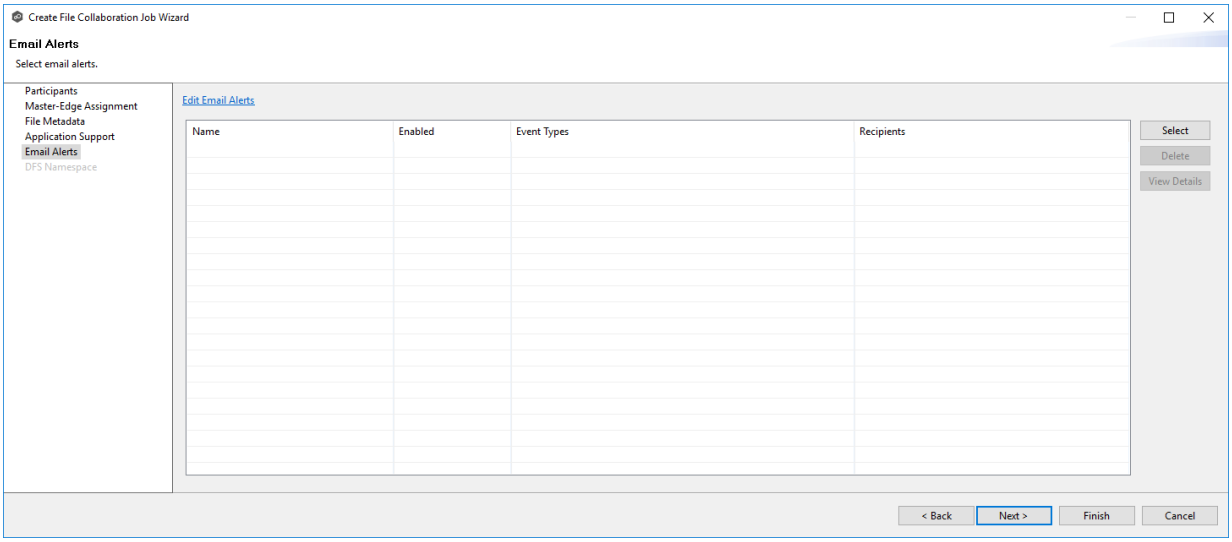

The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.

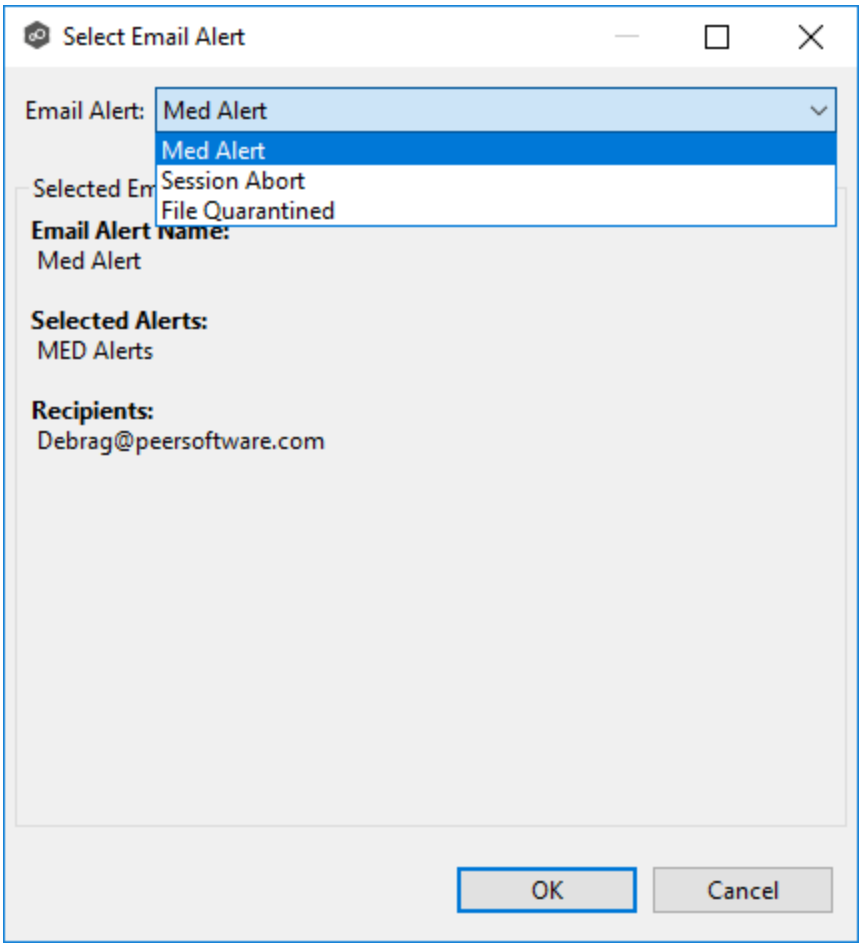

3. Click **OK**.

   The alert is listed in the **Email Alerts** page.

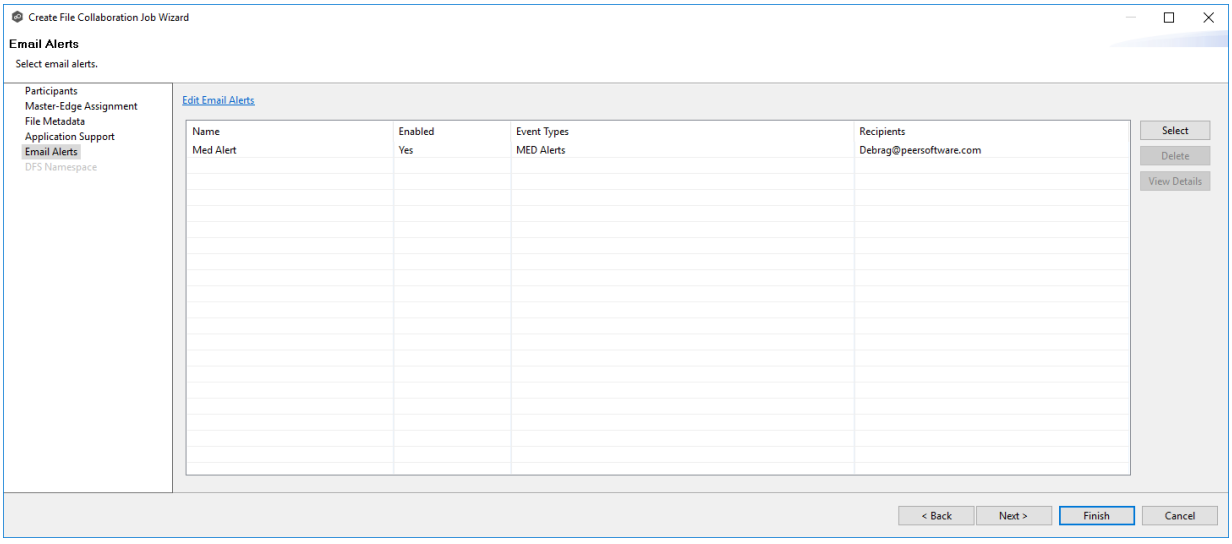

4. (Optional) Repeat steps 1-3 to apply additional alerts.

5. Continue to Step 8: DFS Namespace.

Step 8: DFS Namespace

This step is optional.

The **DFS Namespace** page presents three options for linking a DFS namespace folder to this File Collaboration job.

To link a namespace to this job:

1. Click the **Enable DFS Linking** checkbox.

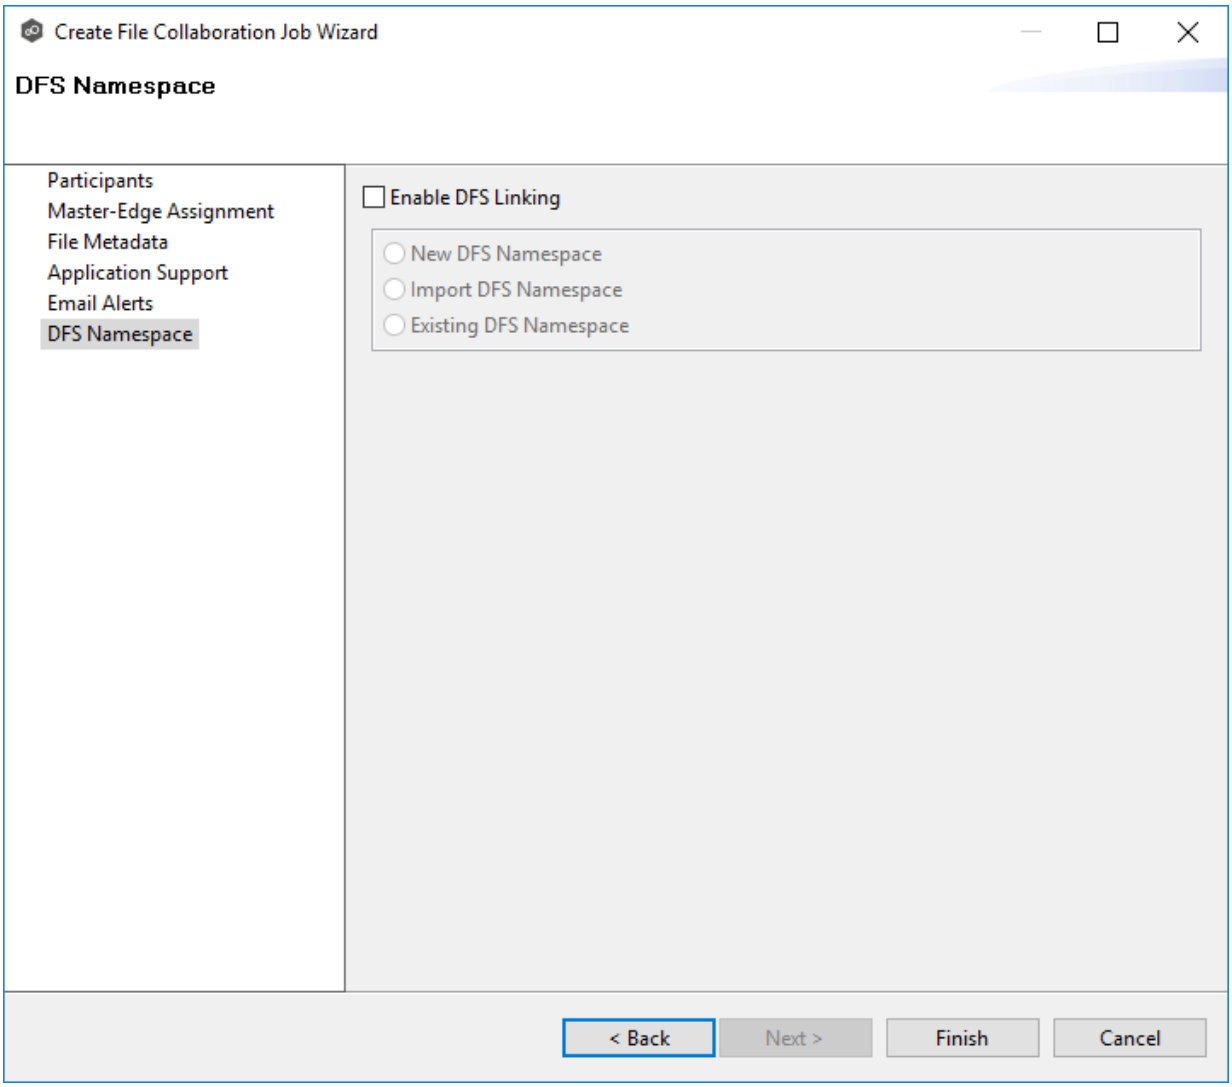

The three options are enabled.

2. Select one of the three options:

   - **New DFS Namespace** - Select this option if you want to create a new namespace. If you select this option, the **Create DFS-N Management Job Wizard** opens. Follow these steps to create a new namespace.

   - **Import DFS Namespace** - Select this option if you have a namespace that was created using the Microsoft DFS Management tool and is not currently being managed by a DFS-N Management job. If you select this option, the **Import Existing Namepaces** Wizard opens. For detailed instructions, follow these steps to import an existing namespace.

- **Existing DFS Namespace** - Select this option if you want to use an existing namespace that is being managed by a DFS-N Management job. If you select this option, it will display the namespace folder and folders associated with namespace.

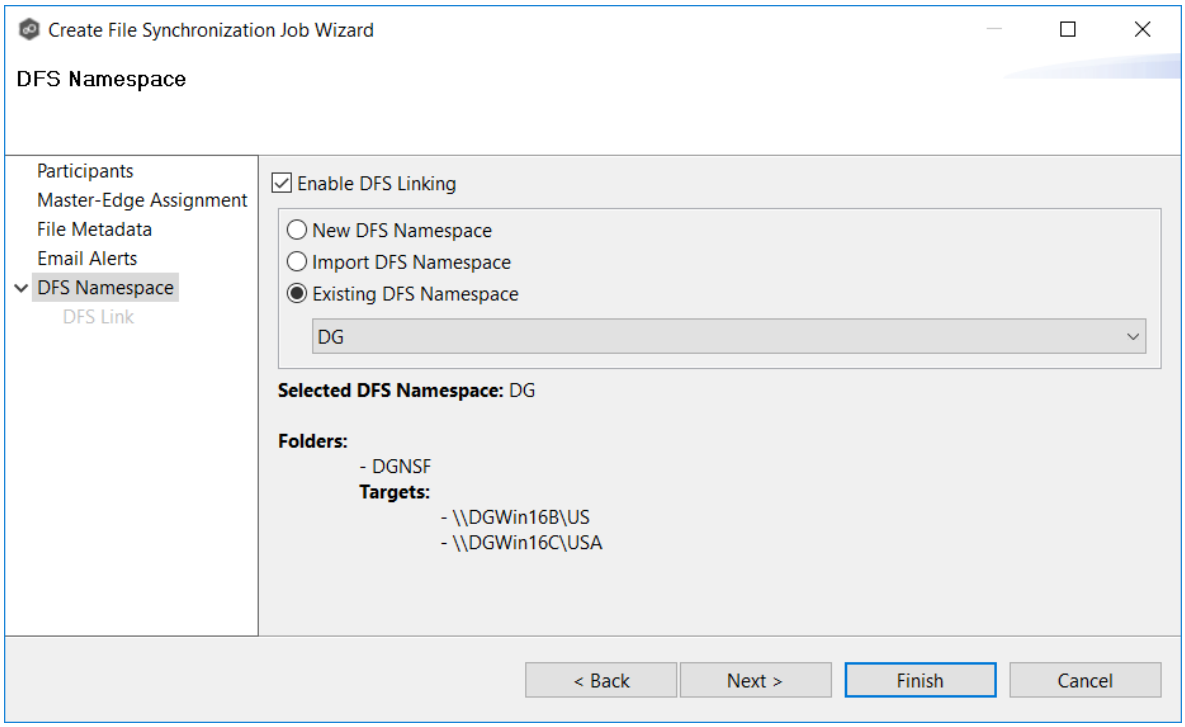

Click **Next** if you want to link participants with folder targets on the **DFS Link** page; otherwise continue with Step 3.

For more information about linking participants to folder targets, see Linking an Existing Namespace Folder to an Existing File Collaboration or File Synchronization Job.

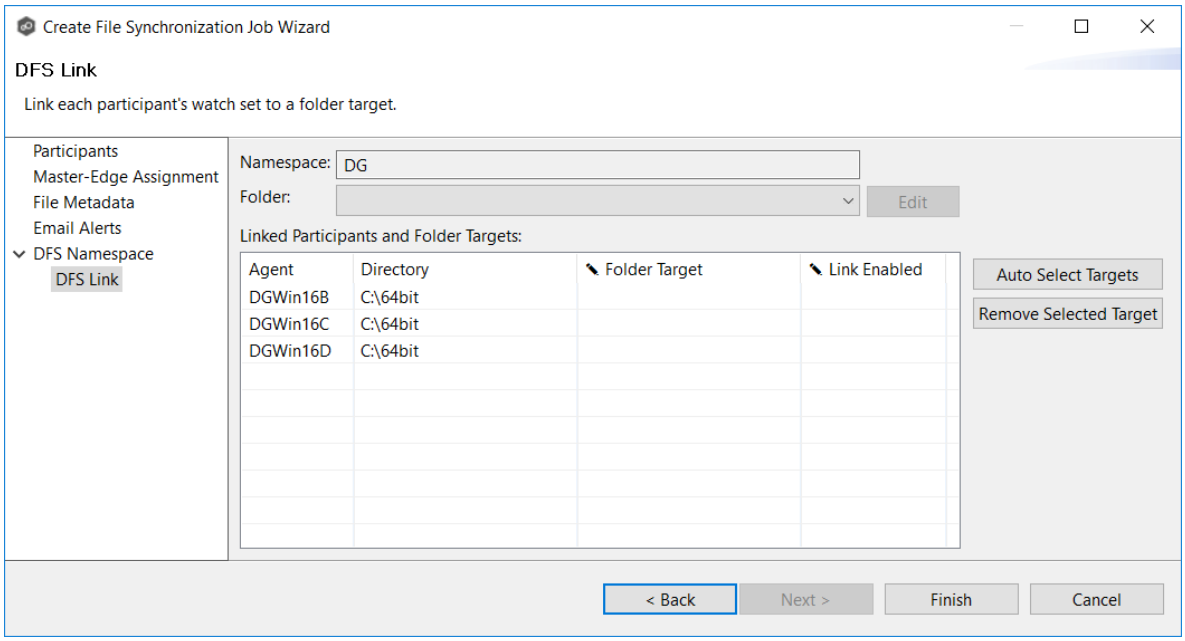

3. Continue to Step 9: Save Job.

Step 9: Save Job

You are now ready to save the job configuration.

1. If you are satisfied with your job configuration, click **Finish** to save your job. Otherwise, click the **Back** button and make any necessary changes.

   Congratulations! You have created a File Synchronization job. It is now listed in the **Jobs** view under **File Synchronization**.

2. Click the **Summary** tab to view a summary of the job.

   See Running and Managing a File Collaboration job for more information.

Editing a File Collaboration Job

You can edit a File Collaboration job while it is running; however, any changes will not take effect until the job is restarted.

Overview

When you create a File Collaboration job, the **Create Job** wizard guides you through the process, presenting the most common options for configuration. When editing a job, you have access to all options, allowing you to fine-tune the job configuration. Options not included in the initial job creation include:

- Delta Replication
- File and Folder Filters
- File Metadata - Some options are available only when editing the job.
- File Locking
- General
- Logging and Alerts
- Scheduled Replication
- SNMP Notifications
- Target Protection
- Tags
- DFS-N

You can edit multiple File Collaboration jobs simultaneously. For information about simultaneously editing multiple jobs, see Editing Multiple Jobs.

Editing a Job

To edit a File Collaboration job:

1. Select the job in the **Jobs** view.
2. Right-click and select **Edit Job**.

   The **Edit File Collaboration** dialog appears.

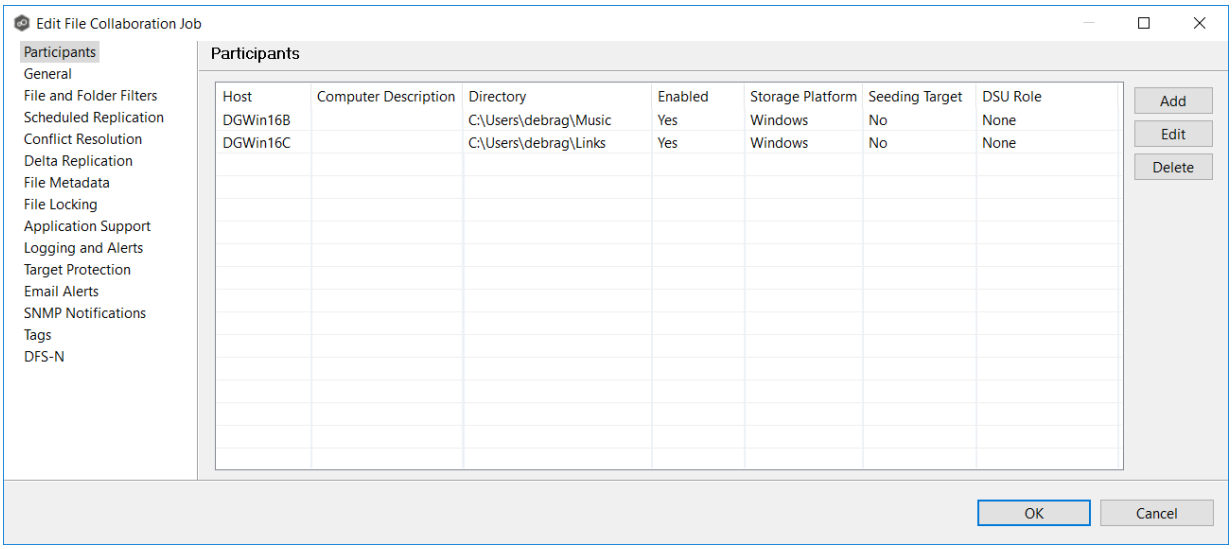

3. Select a configuration item in the navigation tree and make the desired changes:

- Participants
- General
- File and Folder Filters
- Scheduled Replication
- Conflict Resolution
- Delta Replication
- File Metadata
- File Locking
- Application Support
- Logging and Alerts
- Target Protection
- Email Alerts
- SNMP Notifications
- Tags

- DFS-N

4. Click **OK** when finished.

Participants

The **Participants** page in the **Edit File Collaboration Configuration** dialog allows you to:

- Add and delete participants from a job.
- Edit a participant.

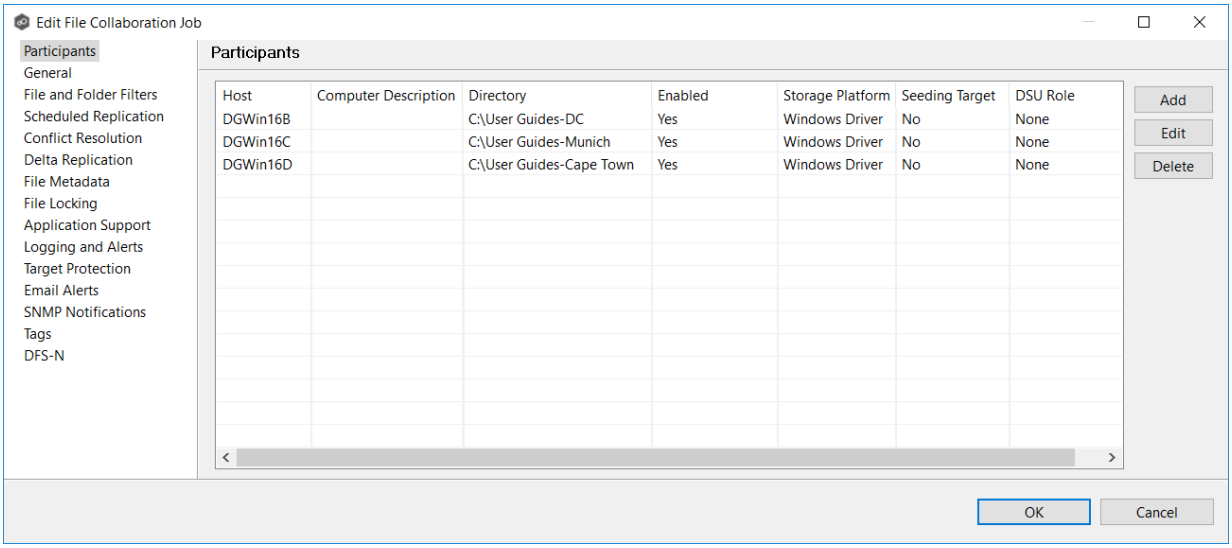

This topic describes adding and deleting participants in a File Collaboration job.

Adding a Participant to a File Collaboration Job

To add a participant to a file collaboration job:

1. Select the job in the **Jobs** view; right-click and select **Edit Job**.

The **Edit File Collaboration** dialog open; the **Participants** page displays the current job participants.

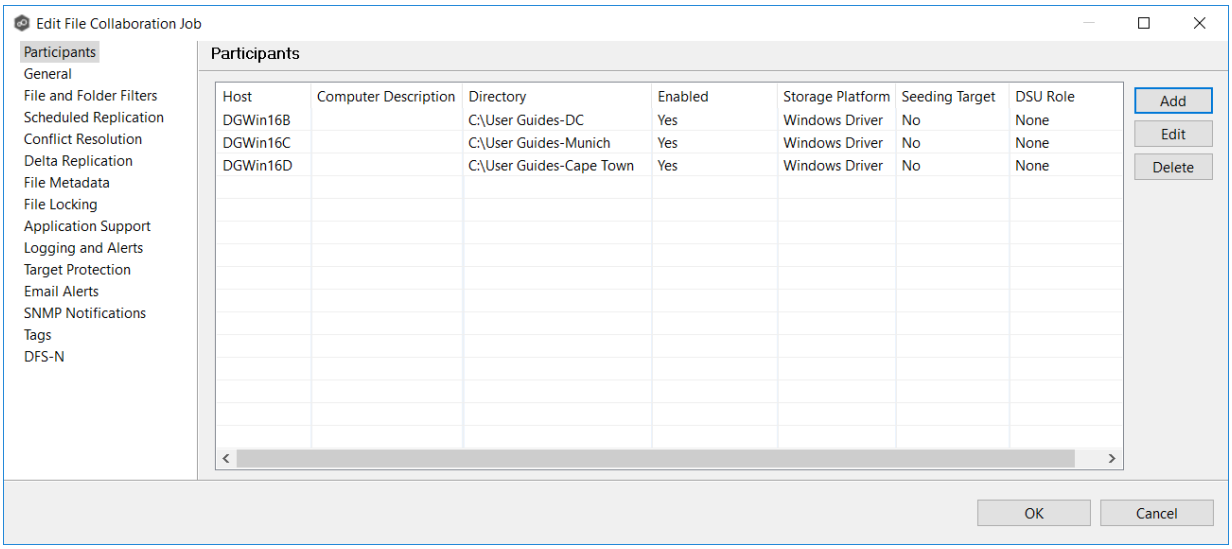

2. Click the **Add** button.

   The **Add New Participant** wizard opens; the **Management Agent** page lists the Agents available to be added.

   Tip: If the Agent you want is not listed, try restarting the Peer Agent Windows Service on that host. If it successfully connects to the Peer Management Broker, then the list is updated with that Agent.

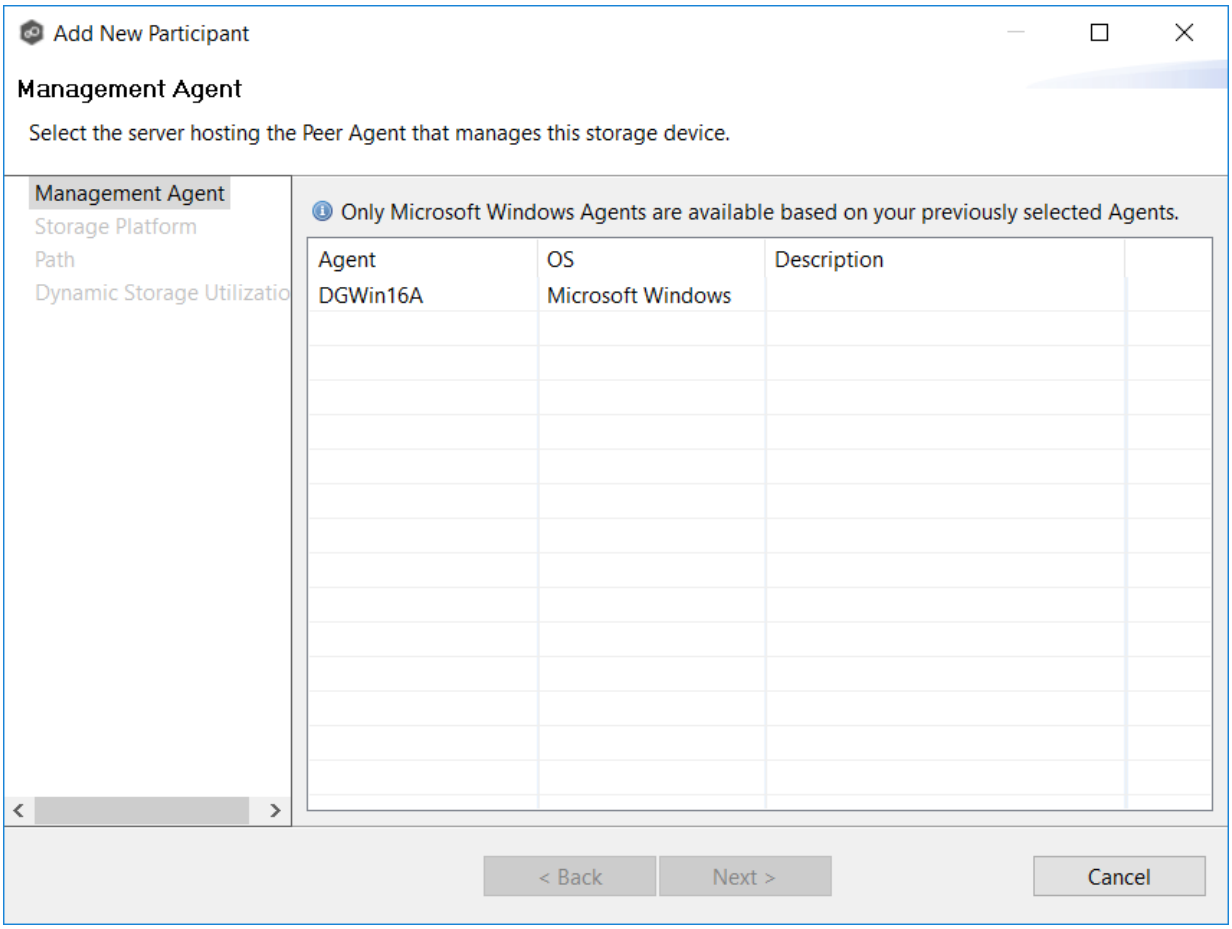

3. Select a Management Agent, and then click **Next**.

   The **Storage Platform** page appears.

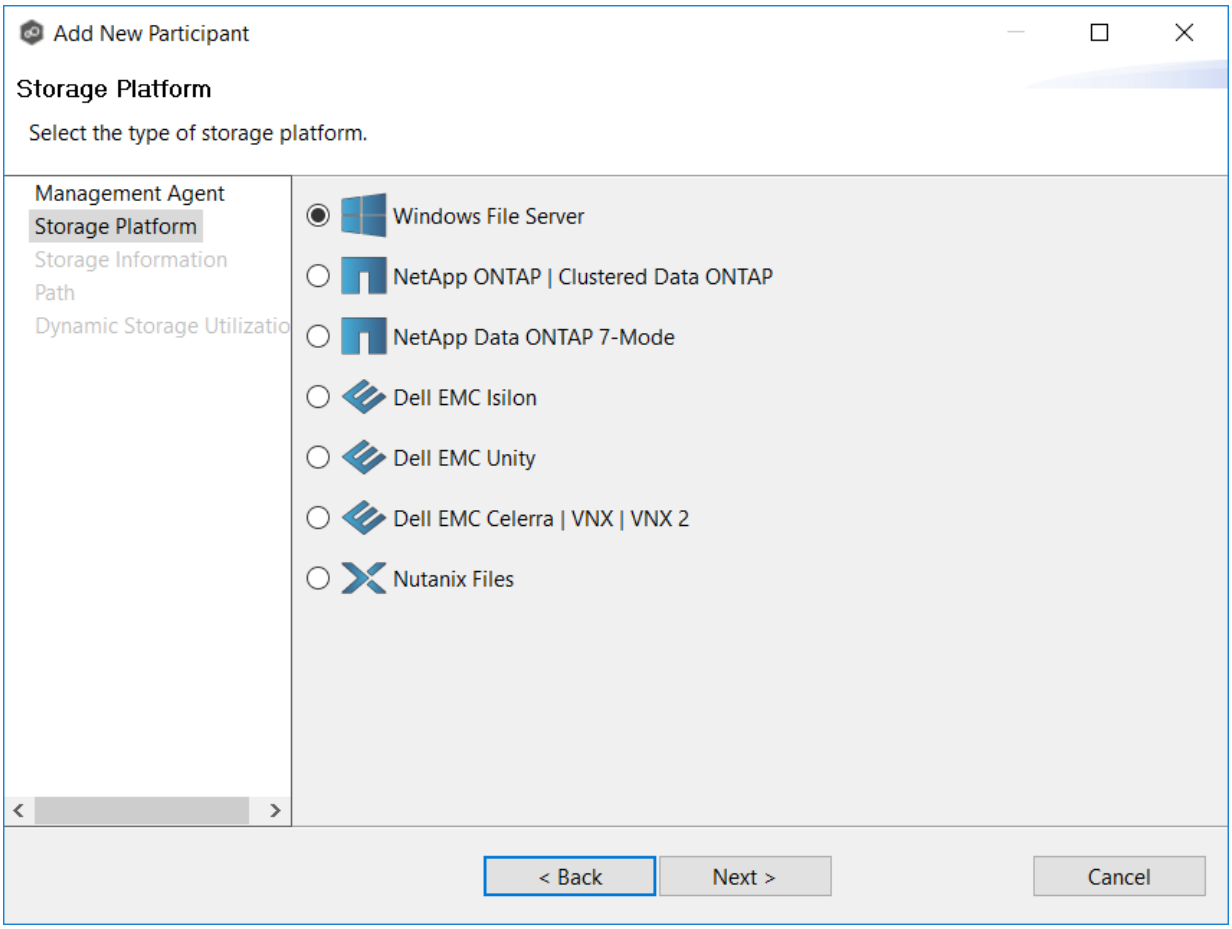

4. Select the type of storage platform that hosts the data you want to collaborate on, and then click **Next**.

   The **Storage Information** page appears; the choices available depend on your selection in the **Storage Platform** page.

5. Enter the requested information.

   Windows File Server

   NetApp ONTAP | Clustered Data ONTAP

   NetApp Data ONTAP 7-Mode

   Dell EMC Isilon

   Dell EMC Unity

   Dell EMC Celerra | VNX | VNX 2

Nutanix Files

6. Click **Next**.

   The **Path** page appears.

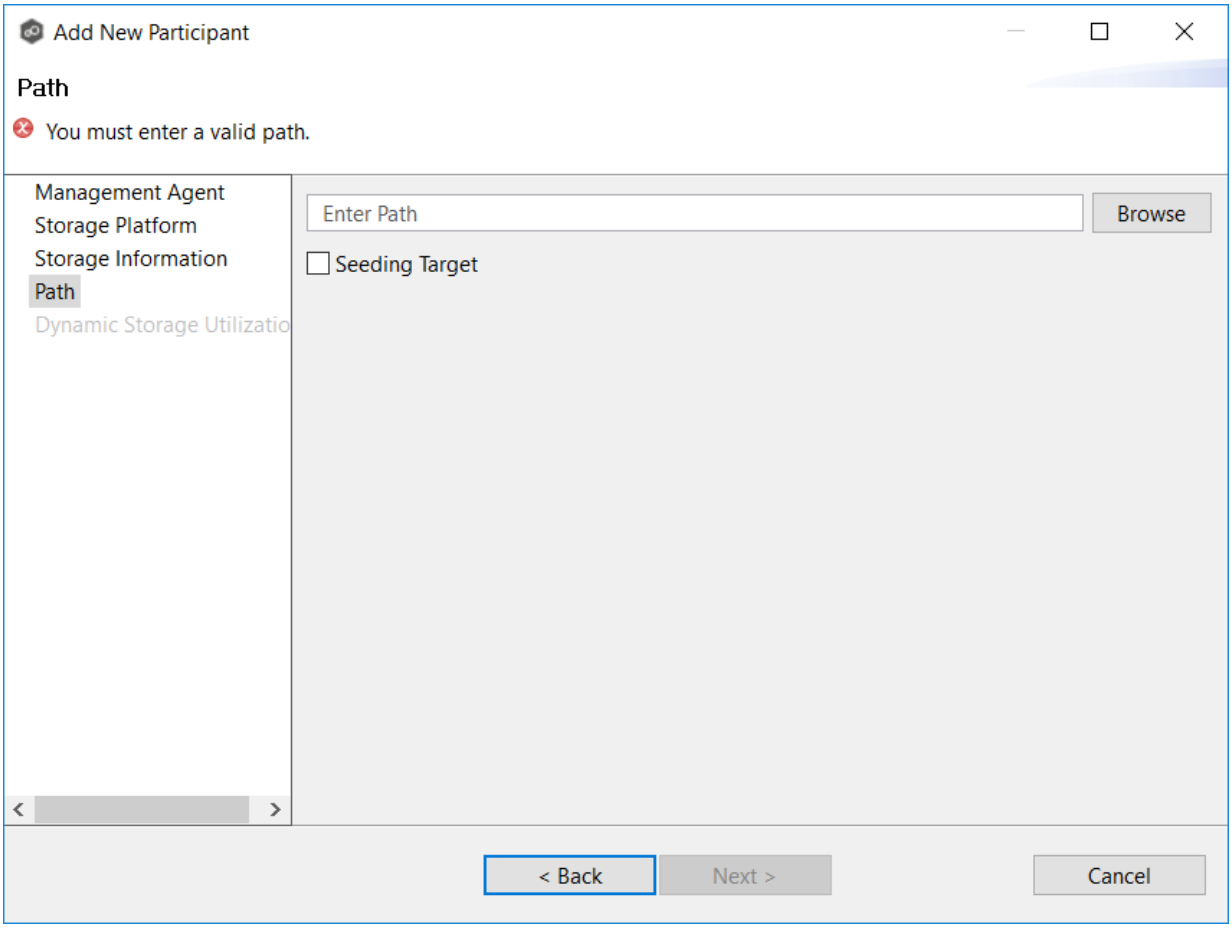

7. Browse to or enter the path to the watch set.

8. (Optional) Select the **Seeding Target** checkbox, and then click **OK** in the dialog that appears.

9. Click **Next**.

   The **Dynamic Storage Utilization** page appears.

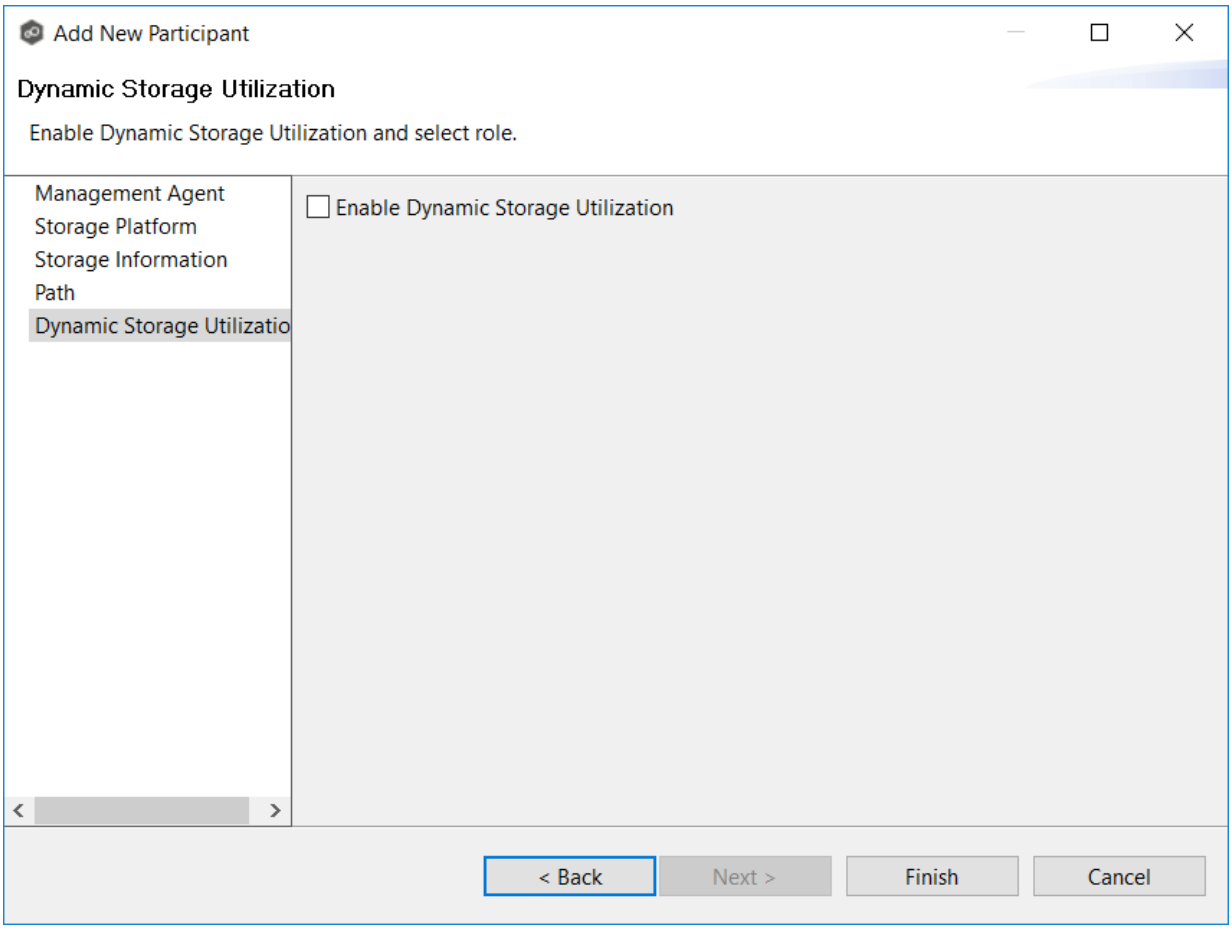

10. (Optional) Select the **Enable Dynamic Storage Utilization** checkbox if you want this participant to be able to use Dynamic Storage Utilization; otherwise, click **Finish**.

11. If you enabled Dynamic Storage Utilization, follow the steps outlined in Step3: Dynamic Storage Utilization in Creating a File Collaboration Job.

    For more information about DSU, see Dynamic Storage Utilization in Advanced Topics.

12. Click **Finish** to complete the wizard.

    The new participant appears in the **Participants** table.

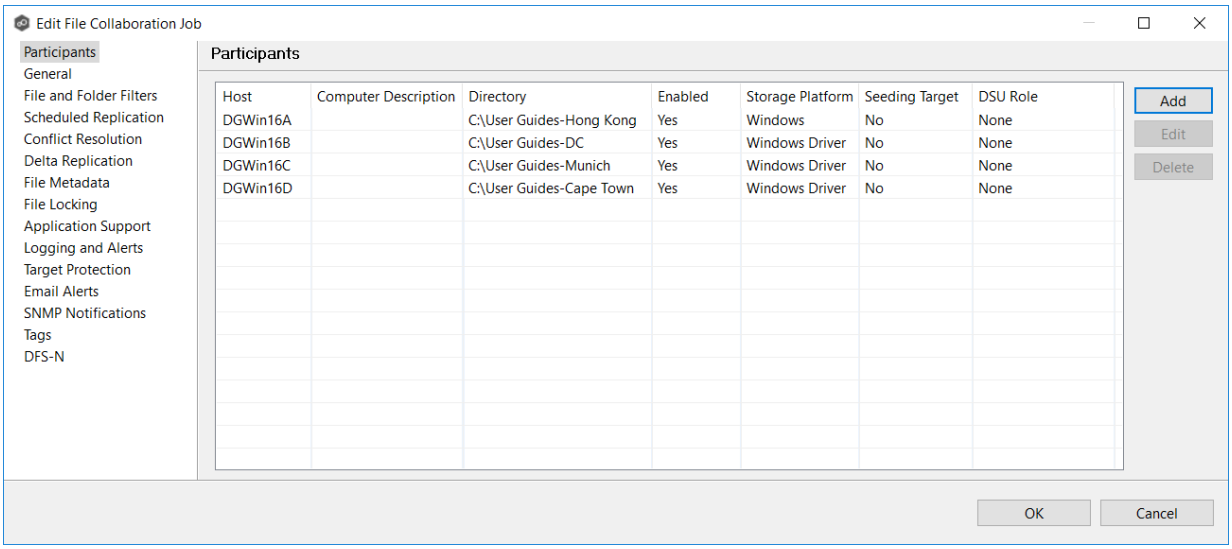

13. Click **OK** to close the Edit wizard or select another configuration item to modify.

Deleting a Participant from a File Collaboration Job

To delete a participant from a File Collaboration job:

1. In the **Edit File Collaboration** dialog, select the participant in the **Participants** table you want to remove from the job.

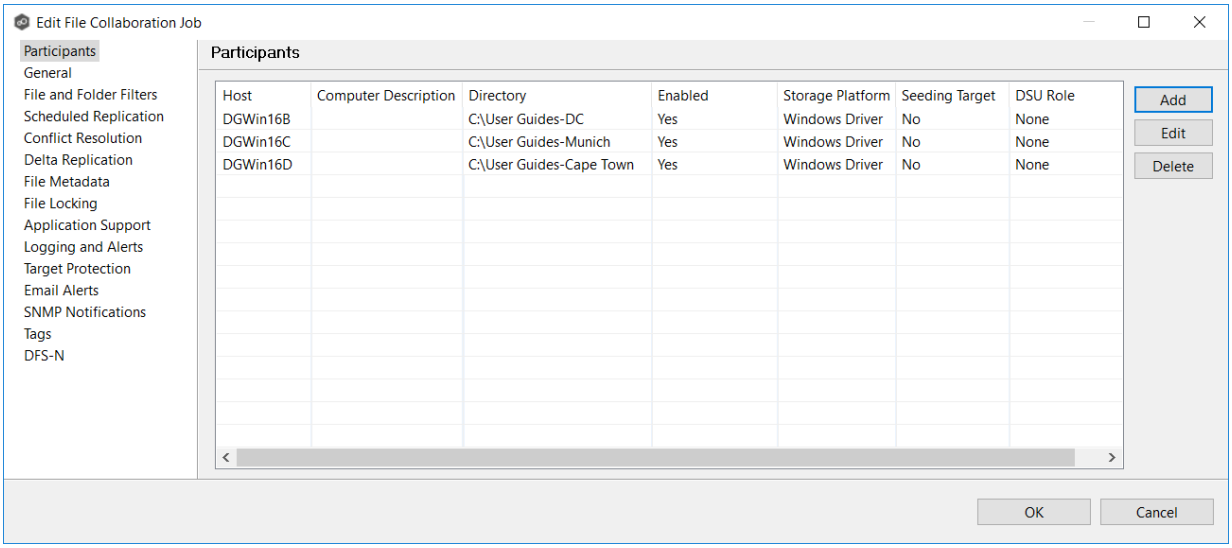

2. Click the **Delete** button

3. Click **OK** in the **Delete Confirmation** dialog.

   The participant is removed from the **Participants** table.

Note: A File Collaboration job must have at least two participants, so if after deleting a participant, there is only a single participant, you must add another participant to the job.

To edit a participant:

1. In the **Edit File Collaboration** dialog, select the participant in the **Participants** table you want to edit.

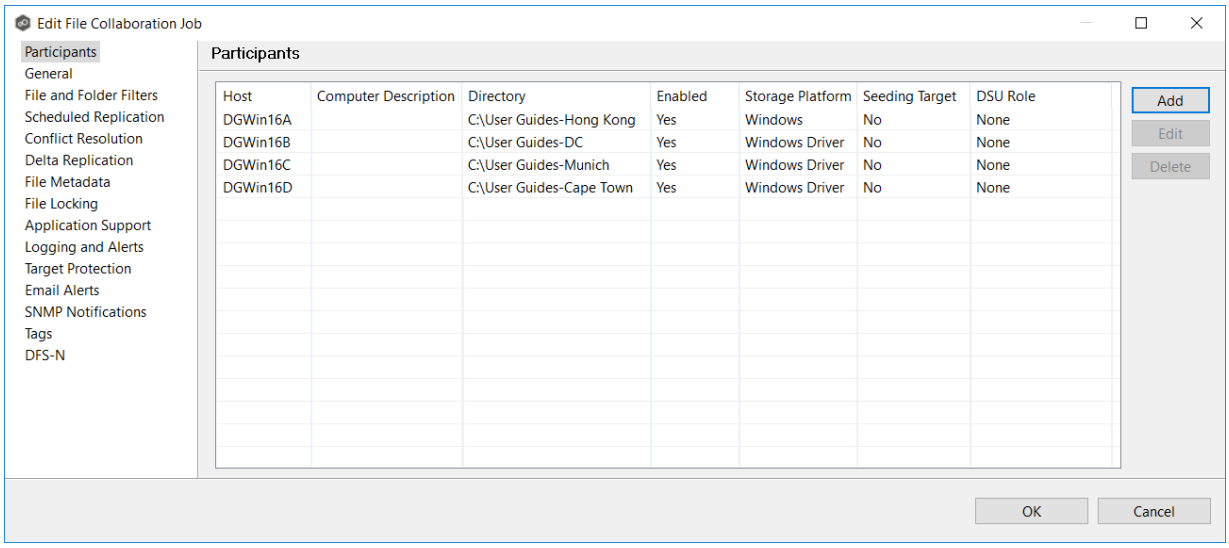

2. Click **Edit**.

   The **Edit Participant** dialog appears.

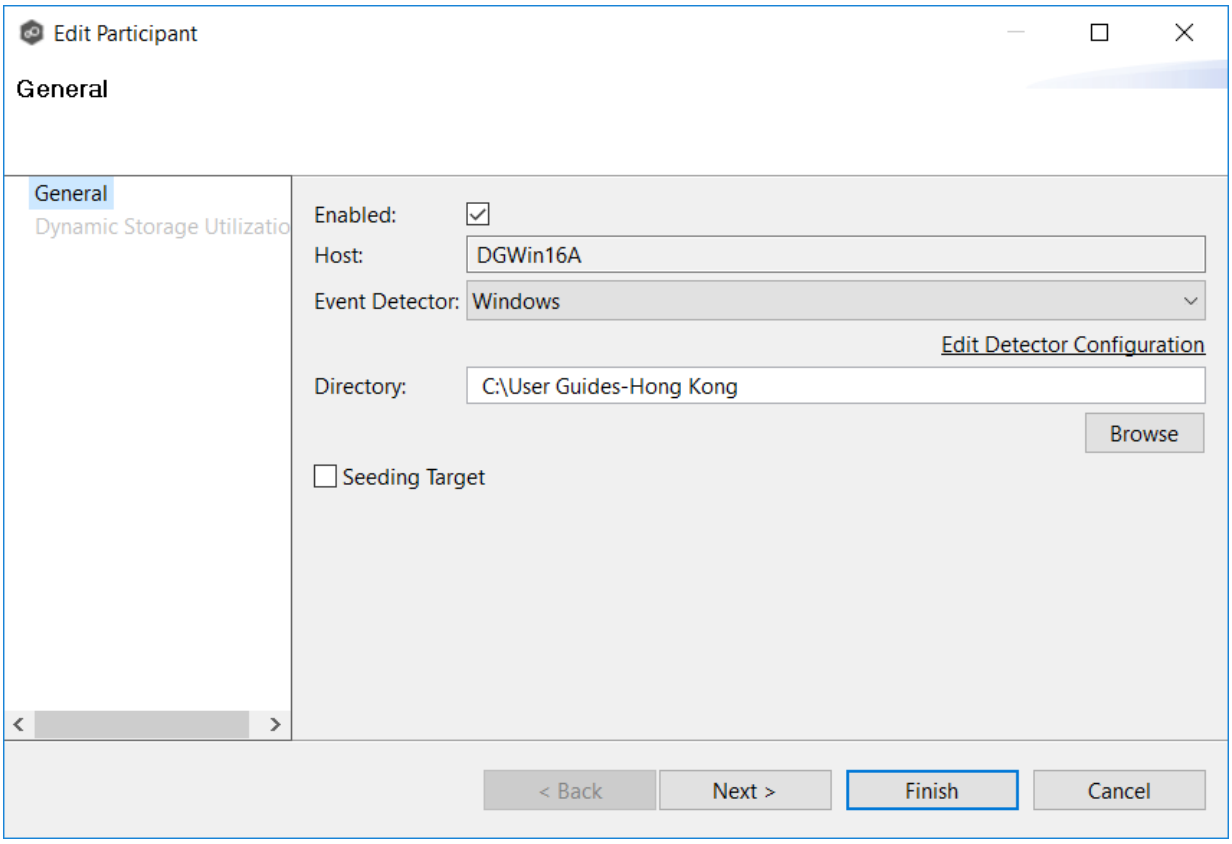

3. To enable or disable the Agent, select or deselect the **Enabled** checkbox.

4. To change the directory/folder/share that is replicated, enter the path or browse to the new watch set in the **Directory** field.

5. If the settings required to connect to the storage device have changed, click **Edit Detector Configuration**, and then make the necessary modifications.

6. To change whether the participant is a seeding target, select or deselect the **Seeding Target** checkbox.

7. Click **Next** to edit Dynamic Storage Utilization options; otherwise, click **Finish**. and continue with Step 10.

   If you click **Next**, the Dynamic Storage Utilization page appears.

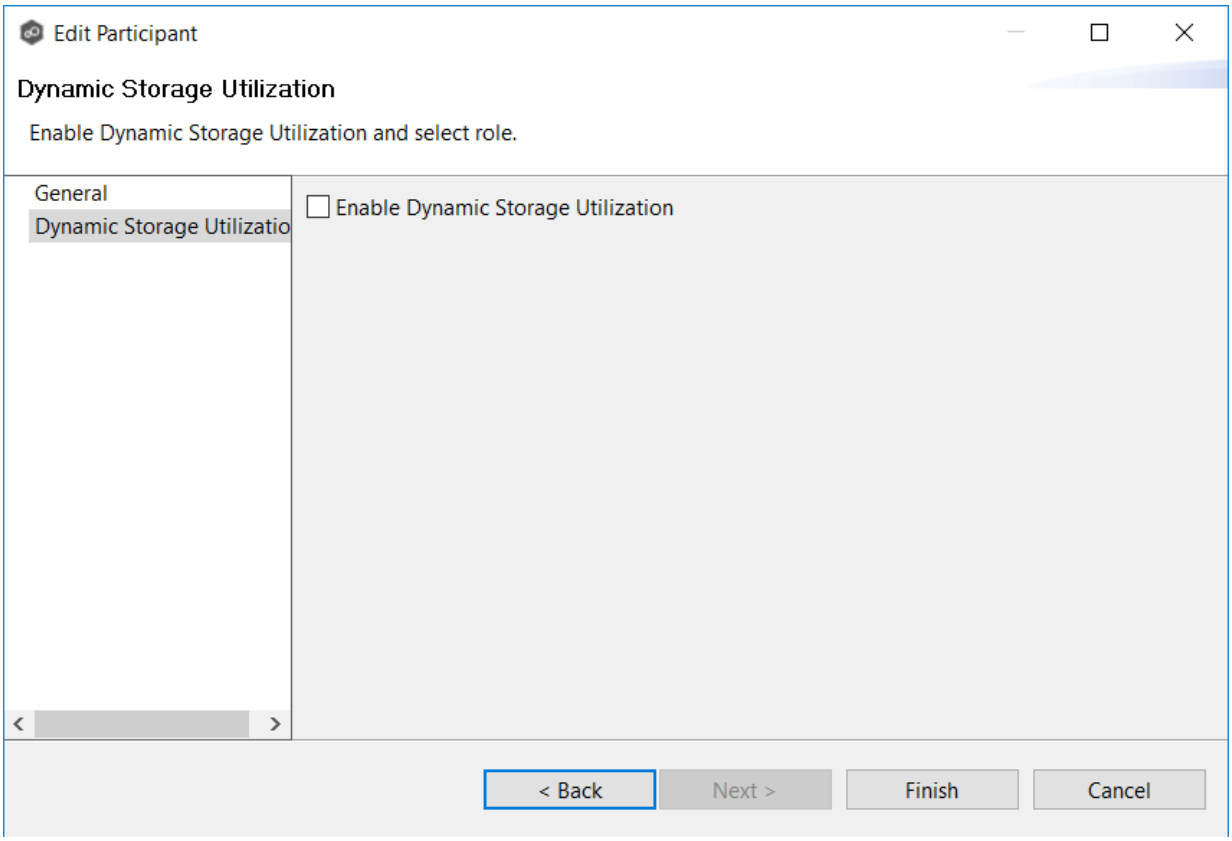

8. (Optional) Select the **Enable Dynamic Storage Utilization** checkbox if you want this participant to be able to use Dynamic Storage Utilization.

9. If you enabled Dynamic Storage Utilization, follow the steps outlined in Step 2: Dynamic Storage Utilization in Creating a File Collaboration Job.

10. Click **OK** to close the Edit wizard or select another configuration item to modify.

General

The **General** page in the **Edit File Collaboration Job** dialog presents miscellaneous settings pertaining to a File Collaboration job. You may want to consult with Peer Software's Technical Support team before modifying these values.

To modify these settings:

1. Enter the values recommended by Peer Software Support.

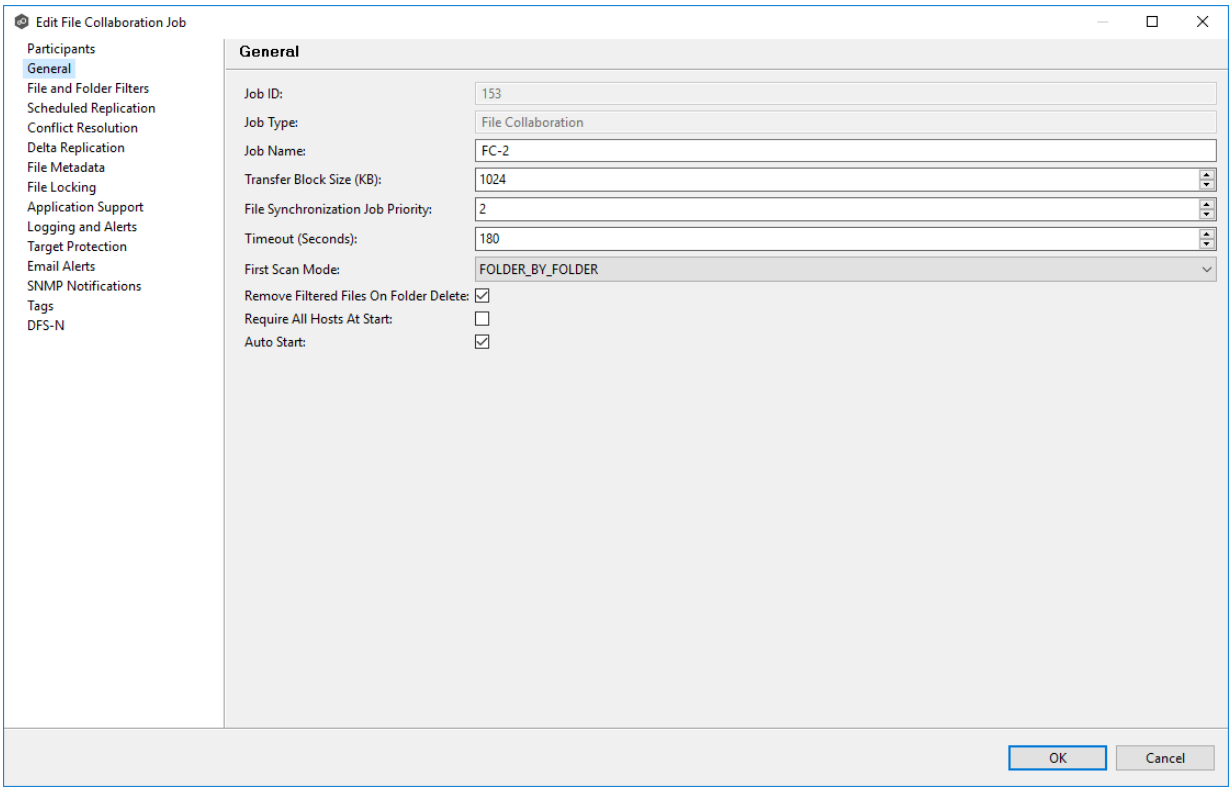

Option	Description
Job ID	Unique, system-generated job identifier that cannot be edited.
Job Type	Identifies the job type. This cannot be modified.
Job Name	Name of this File Collaboration job. This name must be unique.
Transfer Block Size (KB)	The block size in Kilobytes used to transfer files to hosts. Larger sizes will yield faster transfers on fast networks but will consume more memory in the Peer Management Broker and Peer Agents.
File Synchronization Job Priority	Use this to increase or decrease a job's file synchronization priority relative to other configured job priorities. Jobs are serviced in a round-robin fashion, and this number determines the maximum number of synchronization tasks that will be executed sequentially before yielding to another job.

Option	Description
Timeout (Seconds)	Number of seconds to wait for a response from any host before performing retry logic.
First Scan Mode	Determines which scan type will be used when the job is first started. For environments where most data is NOT seeded, the FOLDER_BY_FOLDER method will be best. For environments where most data IS seeded, the BULK_CHECKSUM method will result in a faster first scan.
Remove Filtered Files On Folder Delete	If selected, then all child files on target hosts will be deleted when its parent folder is deleted on another source host. Otherwise, filtered files will be left intact on targets when a parent folder is deleted on another source host.
Require All Hosts At Start	If selected, requires all participating hosts to be online and available at the start of the File Collaboration job in order for the job to successfully start.
Auto Start	If selected, then this file collaboration session will automatically be started when the Peer Management Center Service is started.

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

File and Folder Filters

The **File and Folder Filters** page in the **Edit File Collaboration Job** dialog displays a list of file and folder filters. A file or folder filter enables you to exclude and/or include files and folders from the job based on file type, extension, name, or directory path. Any file or folder that matches the filter is excluded or included from replication, depending on the filter's definition. By default, all files and folders selected in the **Source Paths** page will be replicated.

1. Select the file and folder filters you want to apply to the job.

   If you want to create a new file or folder filter or modify an existing one, click **Edit File Filters**. See File and Folder Filters in the Preferences section for information about creating or modifying a file filter.

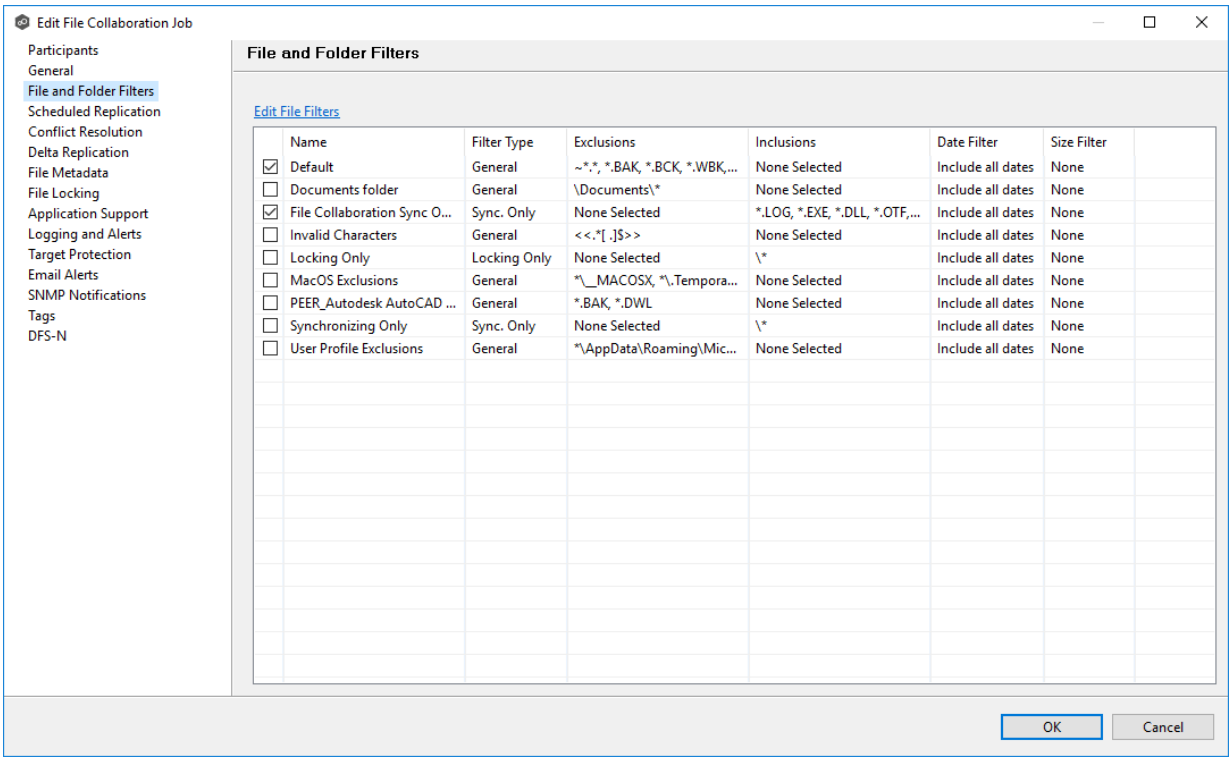

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

Scheduled Replication

The **Scheduled Replication** page in the **Edit File Collaboration Job** dialog displays a list of scheduled replication filters. A scheduled replication filter enables you to identify files and folders that you don't want replicated in real-time.

1. Select the scheduled replication filters you want to apply to the job.

   If you want to create a new filter or modify an existing one, click **Edit Scheduled Replication Filters**. See Scheduled Replication in the Preferences section for information about creating or modifying a scheduled replication filter.

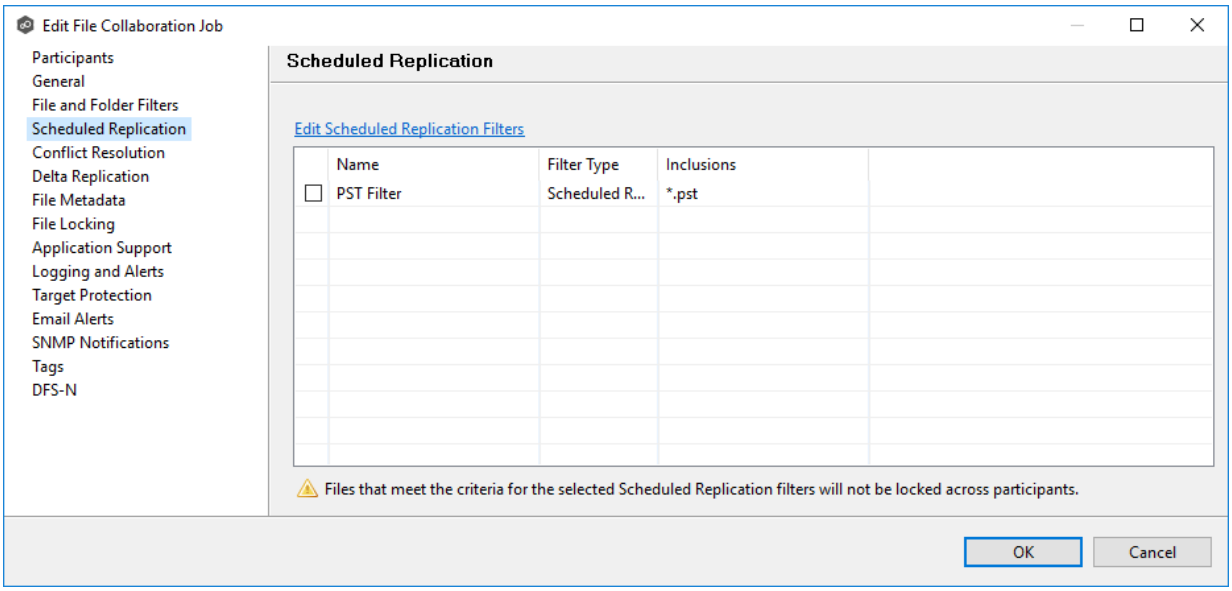

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

Conflict Resolution

By default, any file conflicts that are encountered during the initial synchronization process are automatically resolved by Peer Management Center. Peer Management Center resolves the conflict by selecting the file with the most recent modification time. Conflicts that cannot automatically be resolved result in the files being quarantined. The **Conflict Resolution** page in the **Edit File Collaboration Job** allows you to select options for resolving file conflicts and quarantines.

However, if you want to resolve the conflicts yourself, you can contact Peer Software to enable manual resolution. With manual resolution, you can select the host with the correct version of the file.

For more information about the cause of file conflicts, see Conflicts, Retries, and Quarantines.

To modify conflict resolution settings for the File Collaboration job:

1. In the **File Conflict Resolution** section, select the **Truncate Milliseconds** option if you want the millisecond value truncated from each time stamp when comparing the time stamps of a file on two or more hosts.

   As some storage platforms and applications track milliseconds slightly differently, selecting this option will prevent subtle millisecond differences from causing an otherwise in-sync file to be replicated or quarantined.

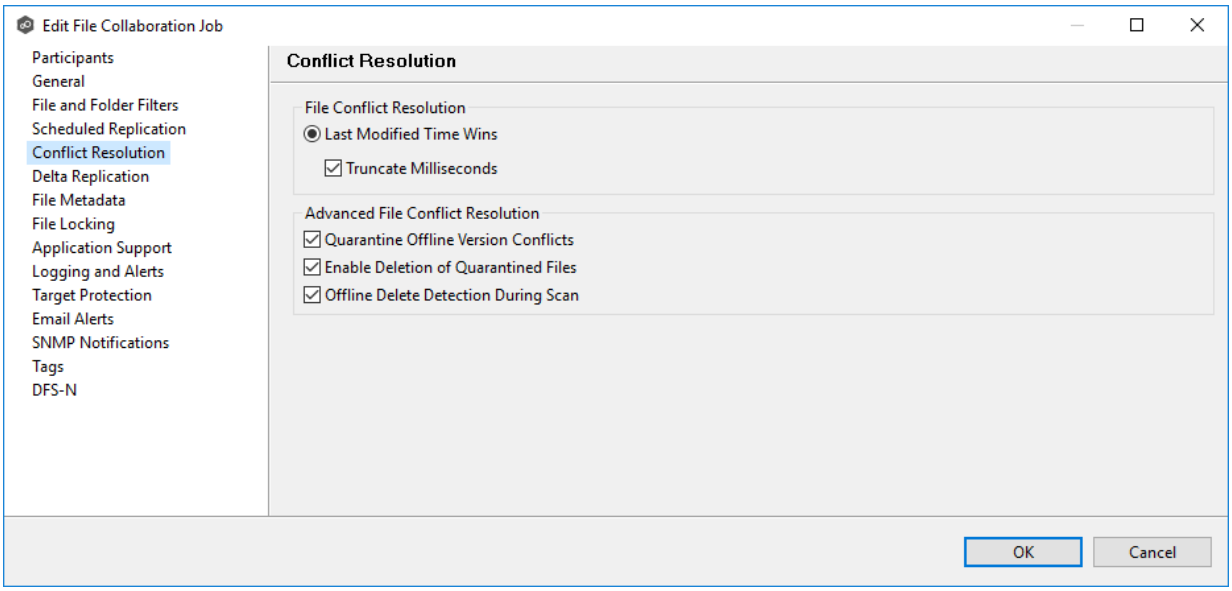

2. Select the **Advanced File Conflict Resolution** options you want applied:

Option	Description
Quarantine Offline Version Conflicts	Select this option if you want Peer Management Center to quarantine a file that was updated in two or more locations while the collaboration session was not running. If it is not selected, the file with the most recent last modified time will be replicated to all other participants.
Enable Deletion of Quarantined Files	Select this option if you want Peer Management Center to process a delete event for a quarantined file. If it is not selected, the quarantined file is not deleted and remains quarantined.
Offline Delete Detection During Scan	Select this option (and enable target protection), if you want any files or folders that were deleted while the job was stopped to be deleted from all participants when the job is restarted. If it is not selected, then the deleted file or folder is restored from a participant with the file or folder to any participant where it no longer exists.

3. Click **OK** to close the Edit wizard or select another configuration item to modify.

Delta Replication

The **Delta Replication** page in the **Edit File Collaboration Job** dialog allows you to specify the delta-replication options to use for the selected File Collaboration job. Delta-level replication is a byte replication technology that enables block/byte level synchronization for a File Collaboration job. Through this feature, Peer Management Center is able to transmit only the bytes/blocks of a file that have changed instead of transferring the entire file. This results in much lower network bandwidth utilization, which can be an enormous benefit if you are transferring files across a slow WAN or VPN, as well as across a high-volume LAN.

Delta-level replication is enabled on a per File Collaboration job basis and generally affects all files in the watch set. You will only benefit from delta-level replication for files that do not change much between file modifications, which includes most document editing programs.

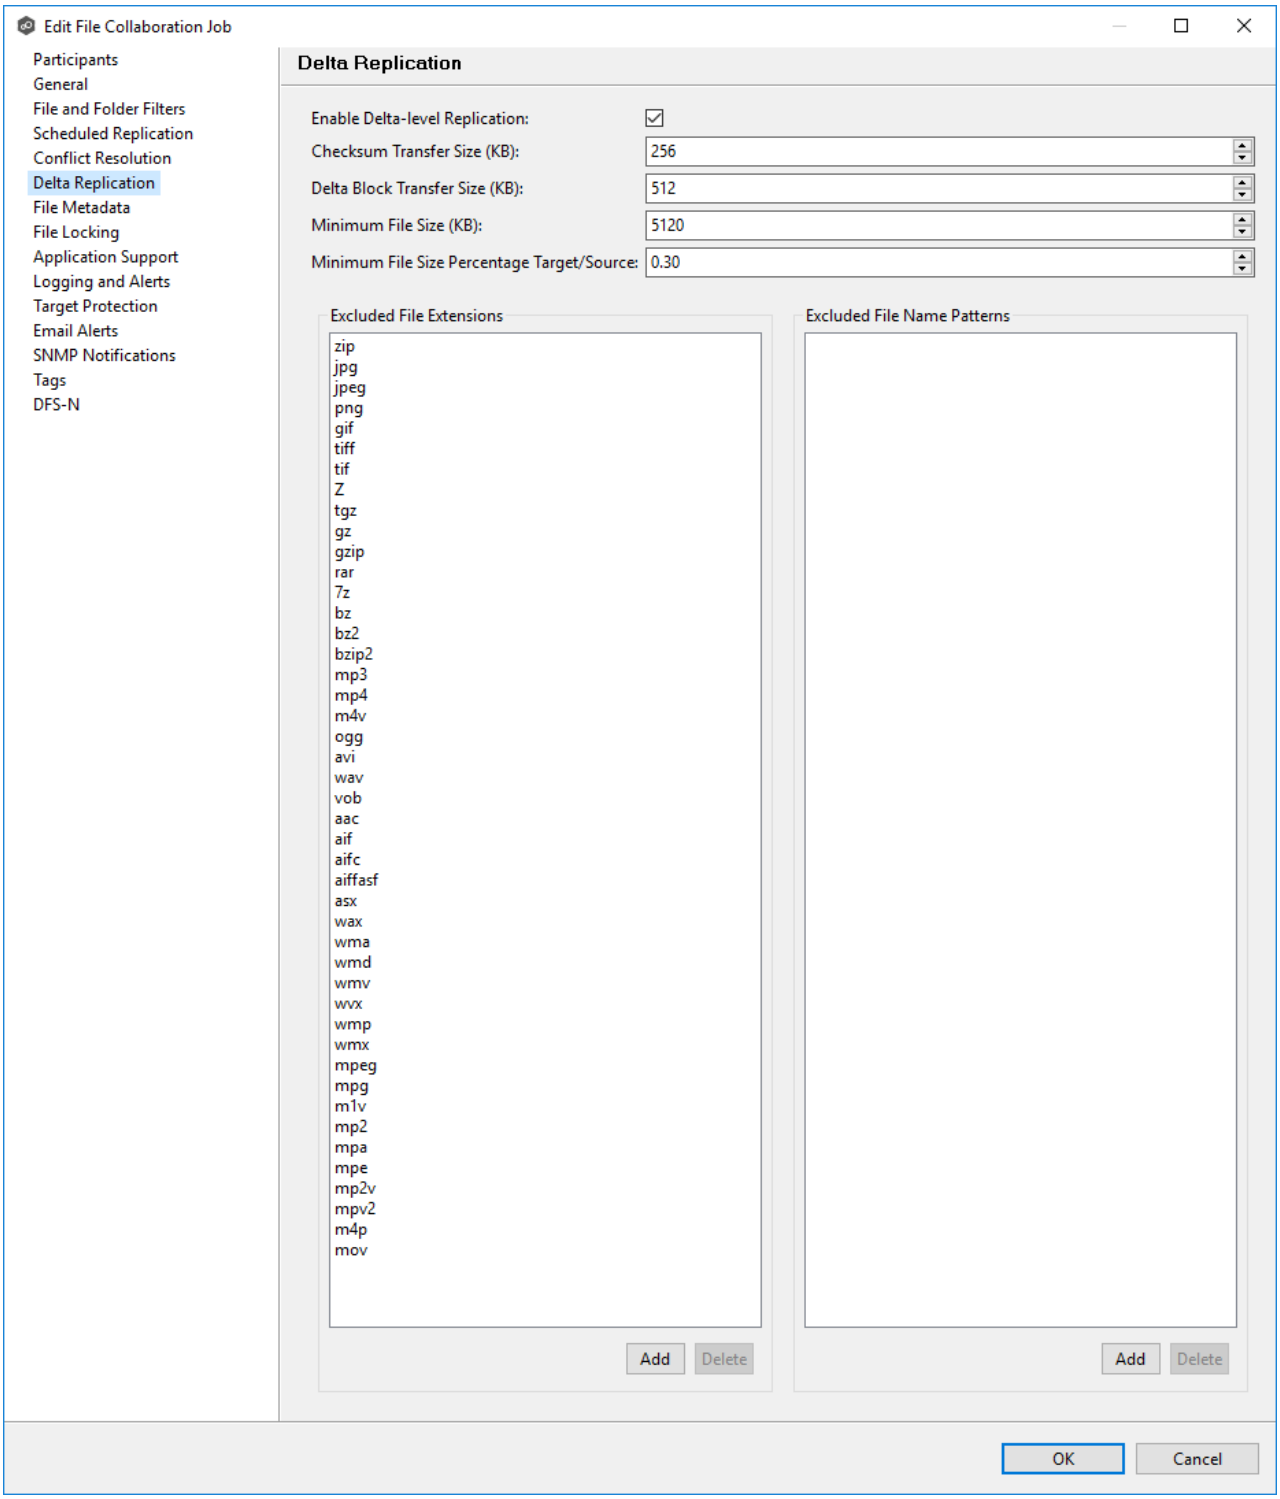

To modify delta-level replication options:

1. Modify the following the fields as necessary.

Field	Description
Enable Delta-Level Replication	Select to enable delta encoded file transfers which only sends the file blocks that are different between source and target(s). If this is disabled, the standard file copy method will be used to synchronize files.
Checksum Transfer Size (KB)	Enter the block in kilobytes used to transfer checksums from target to source at one time. Larger sizes will result in faster checksum transfer, but will consume more memory on the Peer Agents
Delta Block Transfer Size (KB)	Enter the block size in kilobytes used to transfer delta encoded data from source to target at one time. Larger sizes will result in faster overall file transfers but will consume more memory on the Peer Agents.
Minimum File Size (KB)	Enter the minimum size of files in kilobytes to perform delta encoding for. If a file is less than this size, then delta encoding will not be performed.
Minimum File Size Percentage Target/Source	Enter the minimum allowed file size difference between source and target, as a percentage, to perform delta encoding. If the target file size is less than this percentage of the source file size, then delta encoding will not be performed.
Excluded File Extensions	Enter a comma-separated list of file extension patterns to be excluded from delta encoding, e.g., zip, jpg, png. In general, compressed files should be excluded from delta encoding and the most popular compressed file formats are excluded by default.
Excluded File Name Patterns	Enter a list of file name patterns to be excluded from delta encoding. If a file name matches any pattern in this list, then it will be excluded from delta encoding transfers and a regular file transfer will be performed. See File and Folder Filters for more information on specifying wildcard expressions.

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

File Metadata

The **File Metadata** page in the **Edit File Collaboration Job** dialog allows you to modify your file metadata synchronization settings and provides some additional options not available when creating the job. See File Metadata Synchronization in Advanced Topics for more information about file metadata replication.

To enable file metadata synchronization:

1. Select when you want the metadata synchronized (you can select one or both options):

   - **Enable synchronizing NTFS security descriptors (ACLs) in real-time** - Select this option if you want the metadata replicated in real-time. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be transferred to the target host file(s) as they occur.

   - **Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan** - Select this option if you want the metadata replicated during the initial scan. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized during the initial scan.

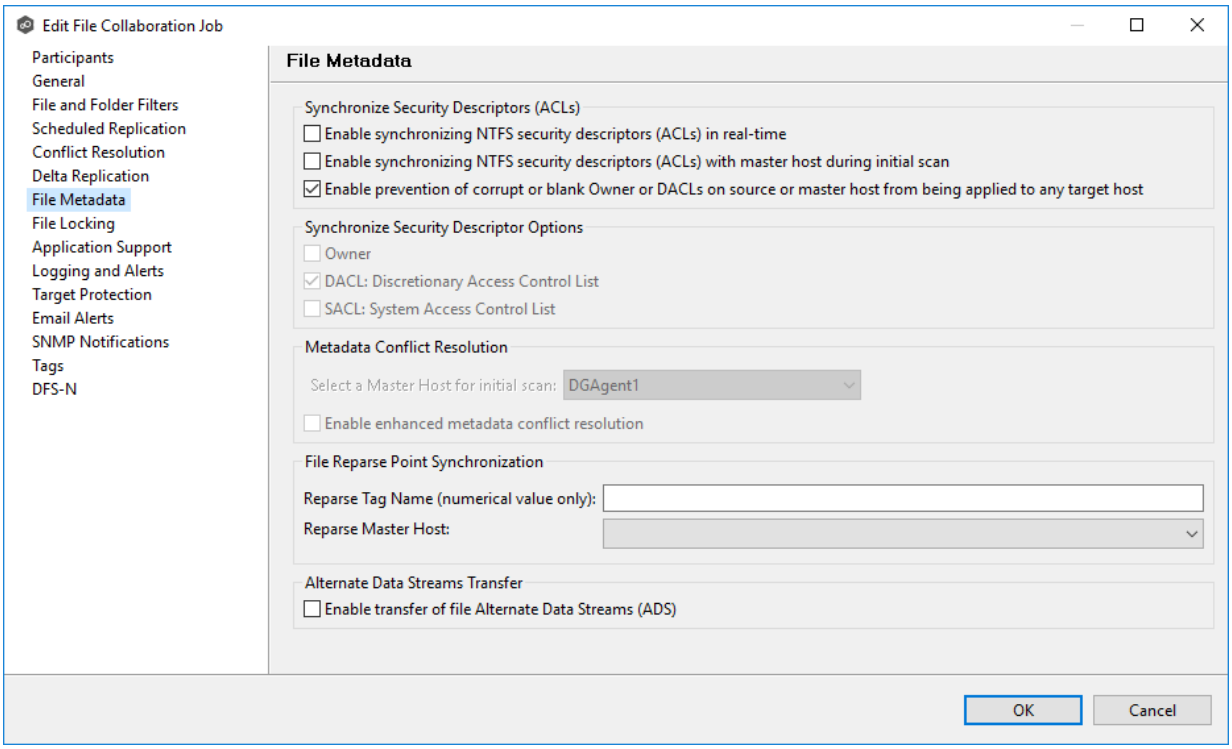

2. Click **OK** in the message that appears after selecting a metadata option.

3. If you selected either of the first two options in the **Synchronize Security Descriptor (ACLs)** section, select the security descriptor components (**Owner**, **DACL**, and **SACL**) to be synchronized.

4. If you selected the option for metadata synchronization during the initial scan (second option in the **Synchronize Security Descriptor (ACLs)** section, select the host to be used as the master host in case of file metadata conflict.

   If a master host is not selected, then no metadata synchronization will be performed during the initial scan. If one or more security descriptors do not match across participants during the initial scan, conflict resolution will use permissions from the designated master host as the winner. If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

5. (Optional) Select the **Enable enhanced metadata conflict resolution** checkbox.

   This option is only available when both of the first two options in the **Synchronize Security Descriptor (ACLs)** section are enabled, and **Owner** is selected under **Synchronize Security Descriptor Options.**

   If you select **Enable enhanced metadata conflict resolution**, this will prevent the Peer Agent service account from being assigned as the owner of that file or folder when a metadata conflict occurs, and a file or folder is written to a target. If the Peer Agent service account were to be assigned as the owner, the user may not have permission to access the file or folder.

   Note: The Peer Agent service account cannot be a local or system administrator. As described in Peer Global File Service - Environmental Requirements, the Peer agent service account should be an actual user.

6. (Optional) Enter values for one or both file reparse point data synchronization options:

   - **Reparse Tag Name** - Enter a single numerical value. Must be either blank (if blank, reparse synchronization will be disabled) or greater than or equal to 0. The default for Symantec Enterprise Vault is 16. A value of 0 enables reparse point synchronization for all reparse file types. If you are unsure as to what value to use, then contact Peer Software technical support, or you can use a value of 0 if you are sure that you are only utilizing one vendor's reparse point functionality.

   - **Reparse Master Host** - Select a master host. If a master host is selected, then when the last modified times and file sizes match on all hosts, but the file reparse attribute differs (e.g., archived/offline versus unarchived on file server), then the file reparse data will be synchronized to match the file located on the master host. For Enterprise Vault, this should be the server where you run the archiving task on. If the value is left blank, then no reparse data synchronization will be performed, and the files will be left in their current state.

   Note: Use this option only if you are utilizing archiving or hierarchical storage solutions that make use of NTFS file reparse points to access data in a remote location, such as Symantec's Enterprise Vault. Enabling this option allows synchronization of a file's

reparse data, and not the actual offline content, to target hosts, and prevents the offline file from being recalled from the remote storage device.

7. (Optional) Select the **Enable transfer of file Alternate Data Stream (ADS)** checkbox.

   If enabled, Alternate Data Streams (ADS) of updated files will be transferred to the corresponding files on target participants as a post process of the normal file synchronization.

   Known limitation: ADS information is transferred only when a modification on the actual file itself is detected. ADS will not be compared between participants. The updated file's ADS will be applied to the corresponding files on target participants.

8. Click **OK** to close the Edit wizard or select another configuration item to modify.

File Locking

The **File Locking** page in the **Edit File Collaboration Job** dialog presents options pertaining to how source and target files are locked by Peer Management Center.

To modify file locking options:

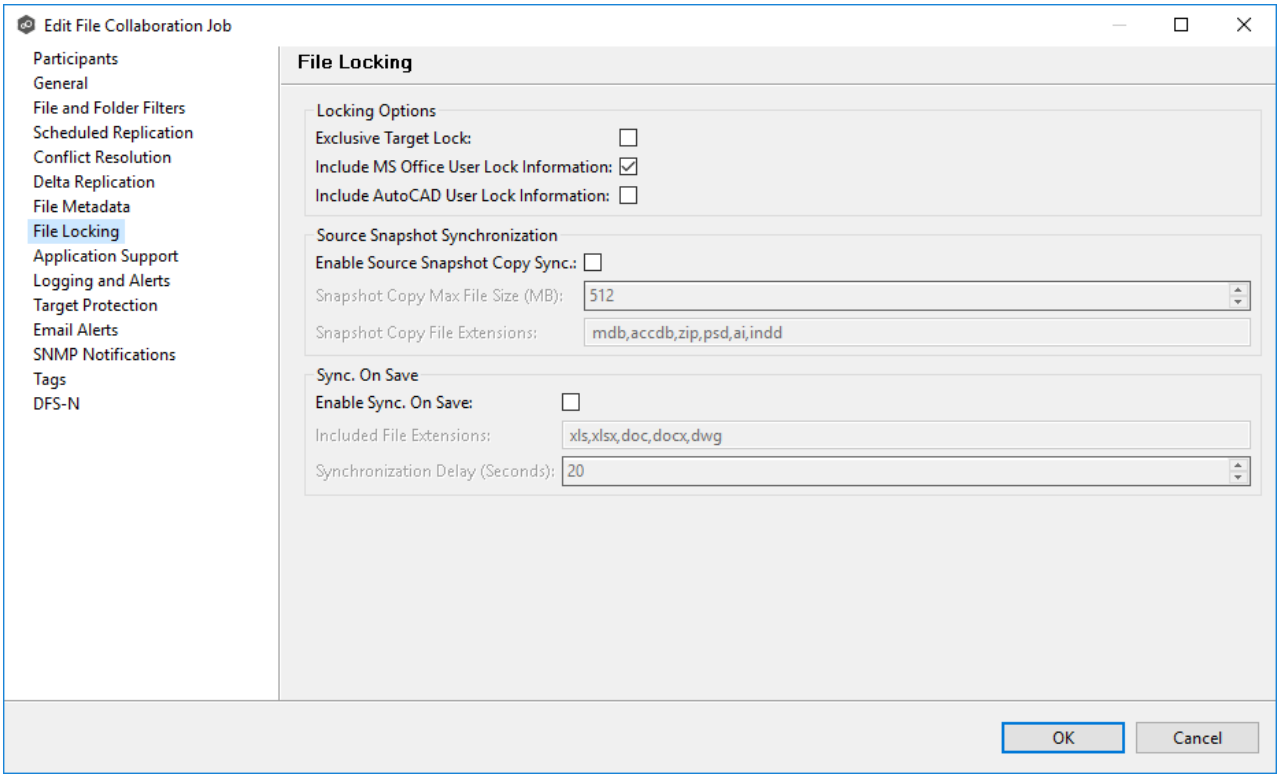

Modify these fields as needed:

Option	Description
Exclusive Target Lock	If enabled, then whenever possible, an exclusive lock will be obtained on target file handles, which will prevent users from opening the file (even in read-only mode) while a user has the file opened on the source host. When this option is disabled, then users will be allowed to open files for read-only if the application allows for this.
Include MS Office User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported Microsoft Office files (e.g., Word, Excel, and PowerPoint).
Include AutoCad User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported AutoCAD files.
Enable Source Snapshot Copy Sync.	If enabled, a snapshot copy of the source file will be created for files that meet the snapshot configuration criteria below, and this copy will be used for synchronization purposes. In addition, no file handle will be held on the source file except while making a copy of the file.
Snapshot Copy Max File Size (MB)	The maximum file size for which source snapshot synchronization will be utilized.
Snapshot Copy File Extensions	A comma-separated list of file extensions for which source snapshot synchronization will be utilized.
Enable Sync. On Save	If enabled, this feature will allow supported file types to be synchronized after a user saves a file, rather than waiting for the file to close.
Included File Extensions	A comma-separated list of file extensions for which to enable the Sync. On Save feature.
Synchronization Delay (Seconds)	The number of seconds to wait after a file has been saved before initiating a synchronization of the file.

Application Support

When you create a File Collaboration job, you have the option of selecting applications to be automatically optimized. When editing the job, you can modify your selections in the **Application Support** page in the **Edit File Collaboration Job** dialog.

To modify which applications are optimized:

1. Select the applications to be optimized.

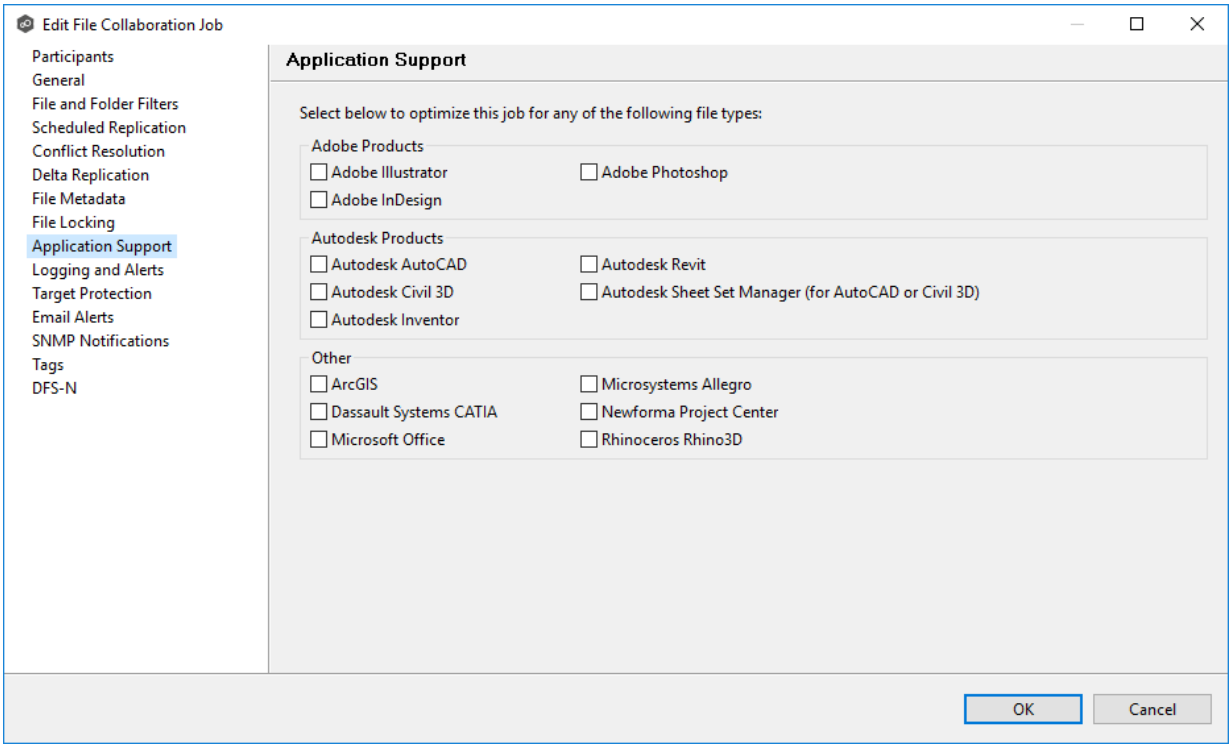

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

Logging and Alerts

Overview of File Event Logging

Various types of file collaboration events can be written to a log file and to the Event Log tab located within the File Collaboration runtime view for the selected File Collaboration job. Each

job will log to the **fc_event.log** file located in the **Hub\logs** subdirectory within the Peer Management Center installation directory. All log files are stored in a tab-delimited format that can easily be read by Microsoft Excel or other database applications.

Log Entry Severity Levels

Informational	Informational log entry, e.g., a file was opened.
Warning	Some sort of warning occurred that did not produce an error but was unexpected or may need further investigation.
Error	An error occurred performing some type of file activity.
Fatal	A fatal error occurred that caused a host to be taken out of the session, a file to be quarantined, or a session to become invalid.

Configuration

By default, all file collaboration activity is logged for all severity levels. You can enable or disable file event logging as well as select the level of granularity.

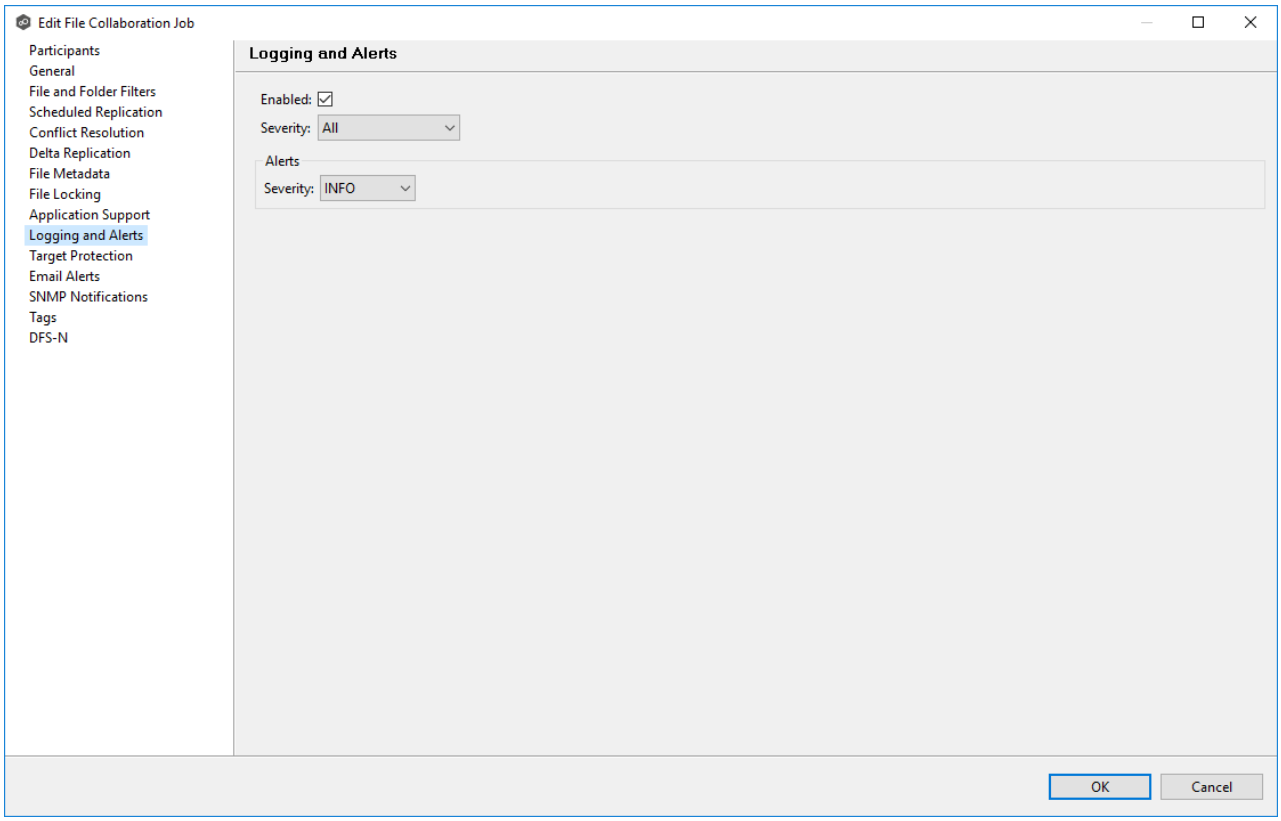

Logging Fields

Below is a list of logging fields and their descriptions:

Field	Description
Enabled	Selecting this option will enable file event logging based on the other settings. Deselecting this option will completely disable all logging.
Severity	Determines what severity levels will be logged. There are two options: • All (Informational, Warnings, Error, Fatal) • Errors & Warnings (Warnings, Error, Fatal)
Event Types	If checked, the corresponding event type will be logged.
File Open	A file was opened by a remote application on a source host.

Field	Description
File Lock	A file lock was acquired on a target host by the File Collaboration job.
File Close	A file was closed.
File Add	A file was added to the watch set.
File Modify	A file was modified in the watch set.
File Delete	A file was deleted.
File Rename	A file was renamed.
Attribute Change	A file attribute was changed.
Security (ACL) Change	The security descriptor of a file or folder was changed.
Directory Scan	Indicates when a directory was scanned as a result of the initial synchronization process.
File ADS Transfer	The Alternate Data Stream of a modified file was synced to target host(s).

Alerts

Various types of alerts can be logged to a log file and to the Alerts table located within the File Collaboration runtime view for the selected job. Each File Collaboration job will log to the **fc_alert.log** file located in the **Hub\logs** subdirectory within the Peer Management Center installation directory. All log files are stored in a tab-delimited format that can easily be read by Microsoft Excel or other database applications.

The default log level is WARNING, which will show any warning or error alerts that occur during a running session. Depending on the severity of the alert, the session may need to be restarted.

Target Protection

Target protection is used to protect files on target hosts by saving a backup copy before a file is either deleted or overwritten on the target host. If enabled, then whenever a file is deleted or modified on the source host but before the changes are propagated to the targets, a copy of the existing file on the target is moved to the Peer Management Center trash bin.

The trash bin is located in a hidden folder named **.pc-trash_bin** found in the root directory of the watch set of the target host. A backup file is placed in the same directory hierarchy location as the source folder in the watch set within the recycle bin folder. If you need to restore a previous version of a file, you can copy the file from the trash bin into the corresponding location in the watch set and the changes will be propagated to all other collaboration hosts.

You can configure target protection in the **Target Protection** page in the **Edit File Collaboration Job** dialog.

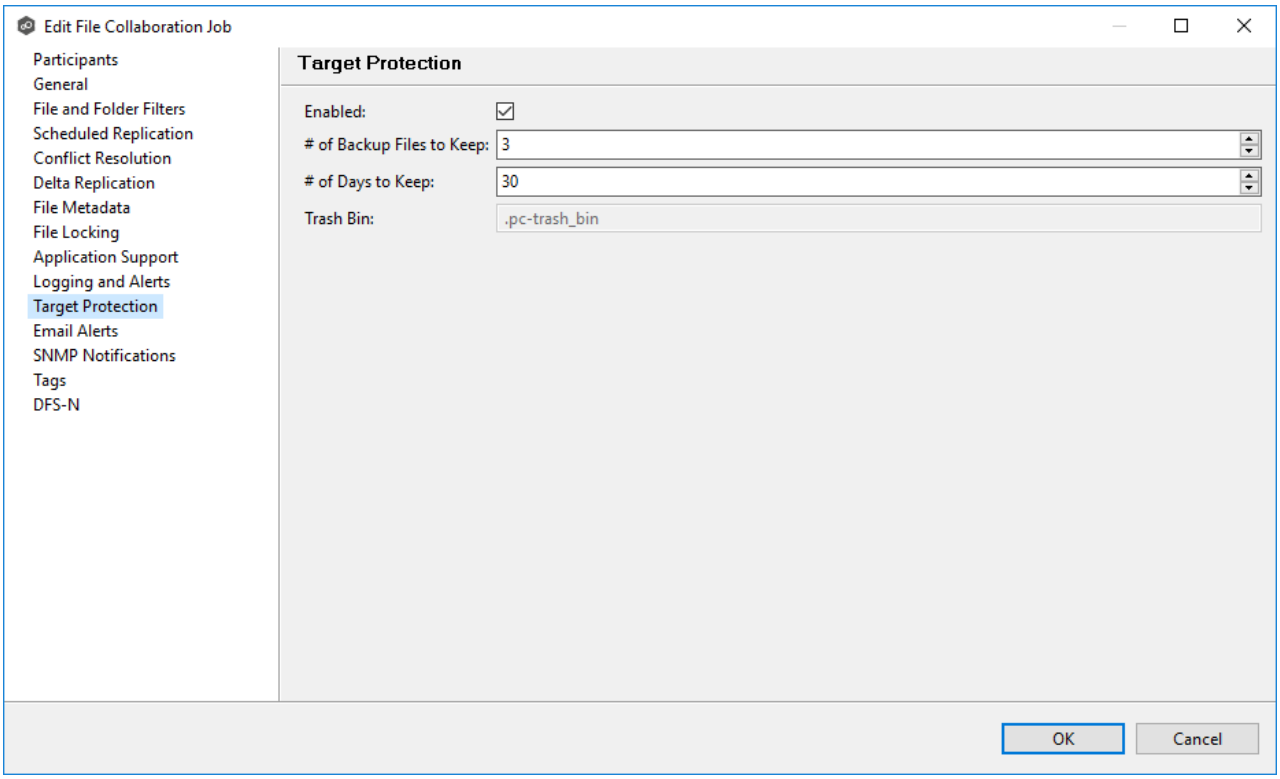

Modify the fields as needed:

Field	Description
Enabled	Enables target protection.
# of Backup Files to Keep	The maximum number of backup copies of an individual file to keep in the trash bin before purging the oldest copy.
# of Days to Keep	The number of days to keep a backup archive copy around before deleting from disk. A value of 0 will disable purging any files from archive.
Trash Bin	The trash bin folder name located in the root directory of the watch set. This is a hidden folder and the name cannot be changed by the end-user.

Email Alerts

The **Email Alerts** page in the **Edit File Collaboration Job** dialog allows you to select which email alerts to apply to a File Collaboration job. Email alerts are defined in the Preferences dialog and can then be applied to individual jobs. See Email Alerts in the **Preferences** section for information about creating an email alert for a File Collaboration job.

To apply email alerts to a File Collaboration job while editing the job:

1. Click the **Select** button.

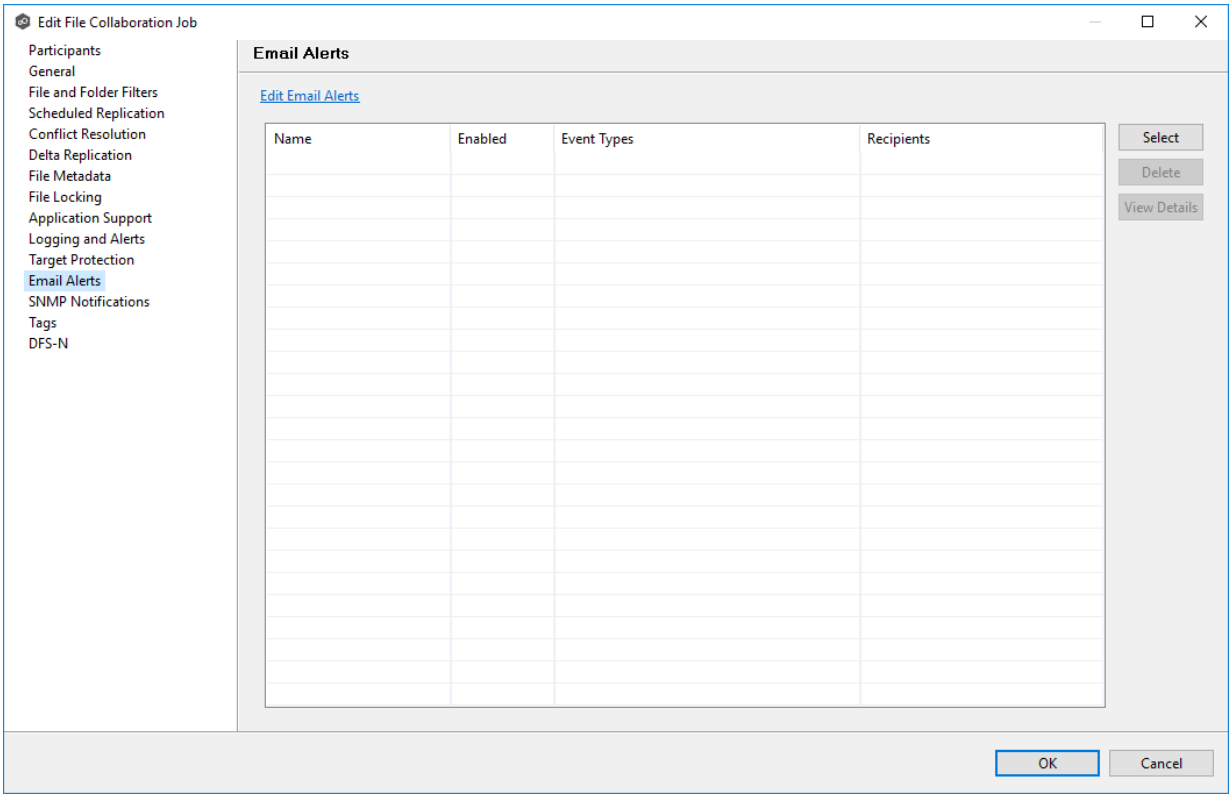

The **Select Email Alert** dialog opens.

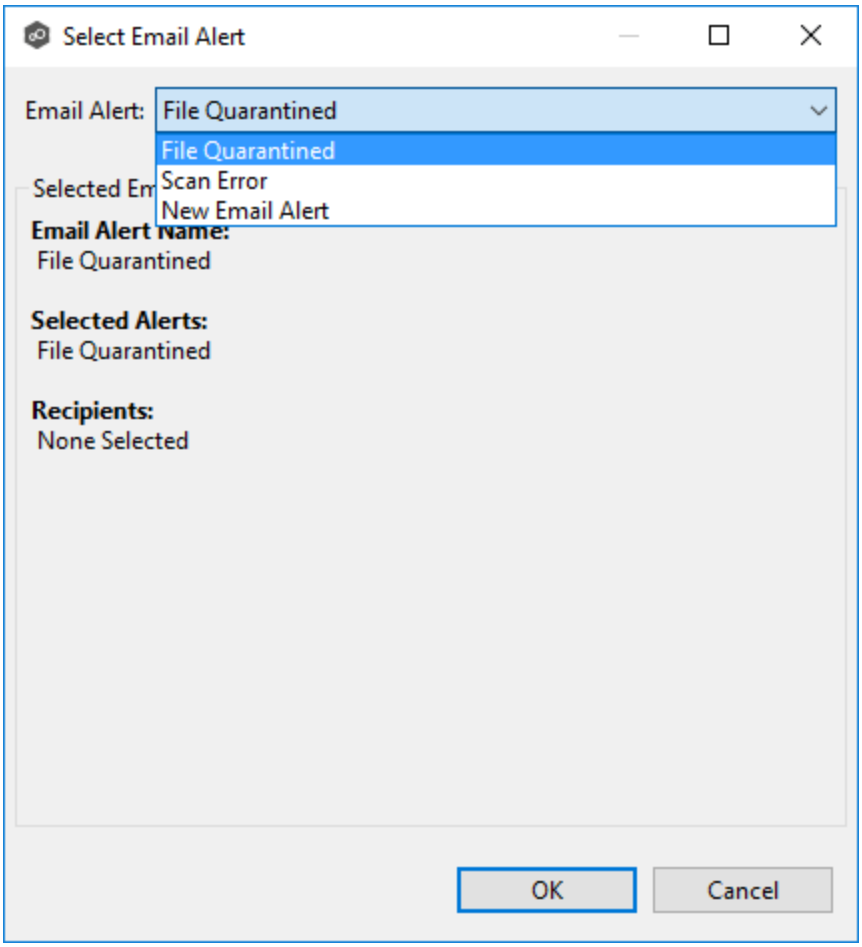

2. Select the email alert from the drop-down list, and then click **OK**.

   The newly added email alert appears in the **Email Alerts** table.

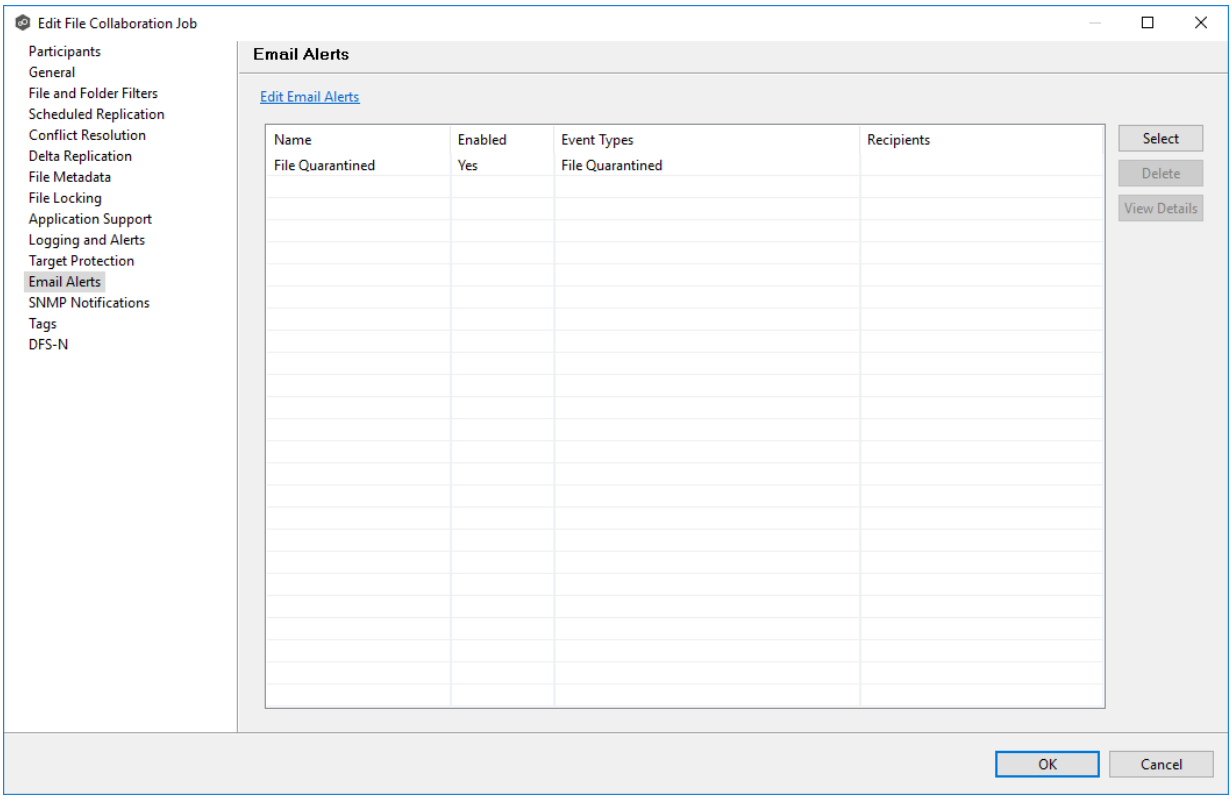

3. Repeat to add additional alerts to the job.

4. Click **OK** to close the Edit wizard or select another configuration item to modify.

SNMP Notifications

The **SNMP Notifications** page in the **Edit File Collaboration Job** dialog allows you to apply SNMP notifications to a File Collaboration job.

SNMP notifications, like email alerts and file filters, are configured at a global level in the Preferences dialog, then applied to individual jobs. For more information about SMNP Notifications, see SNMP Notifications in the **Preferences** section.

To enable or disable SNMP notifications for a File Collaboration job:

1. To enable, select an SNMP notification from the drop-down list.

   To disable, select **None - Disabled**.

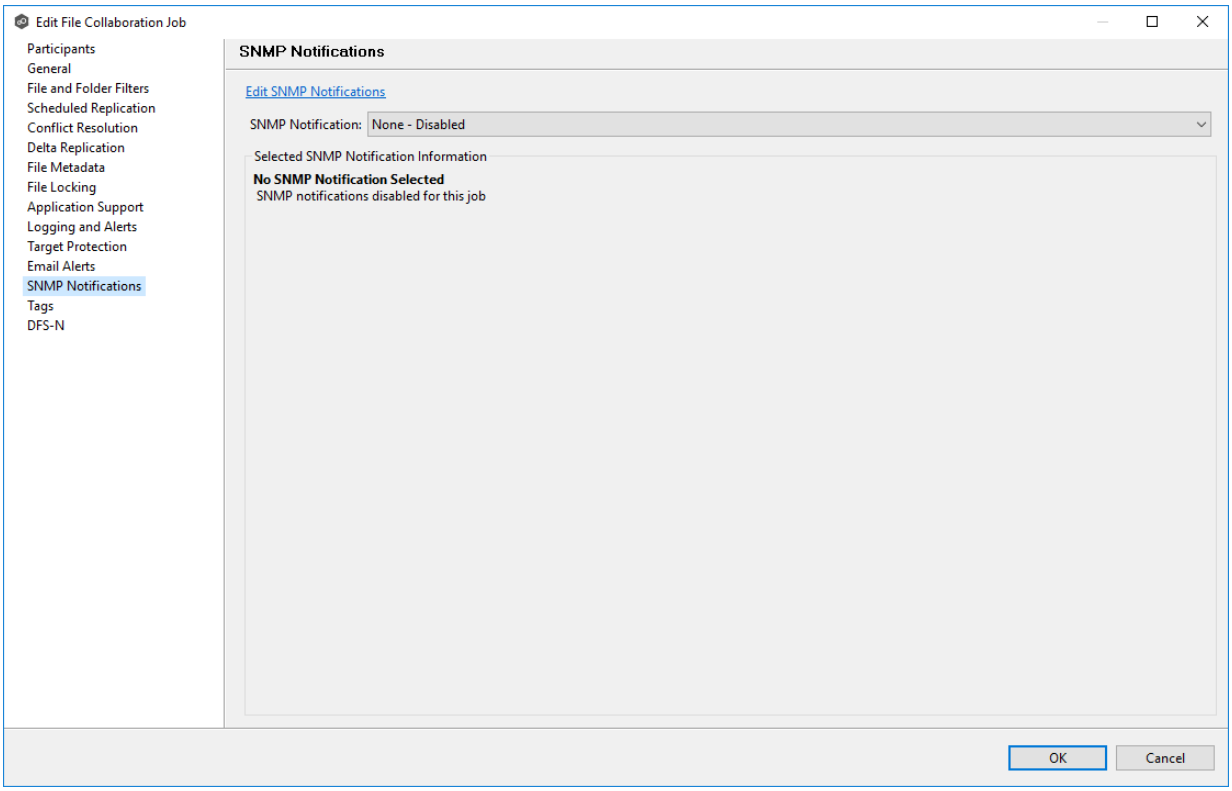

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

Tags

The **Tags** page in the **Edit File Collaboration Job** dialog allows you to assign existing tags and categories to the selected job. This page is not available in Multi-Job Editing mode. For more information about tags, see Tags in the Basic Concepts section.

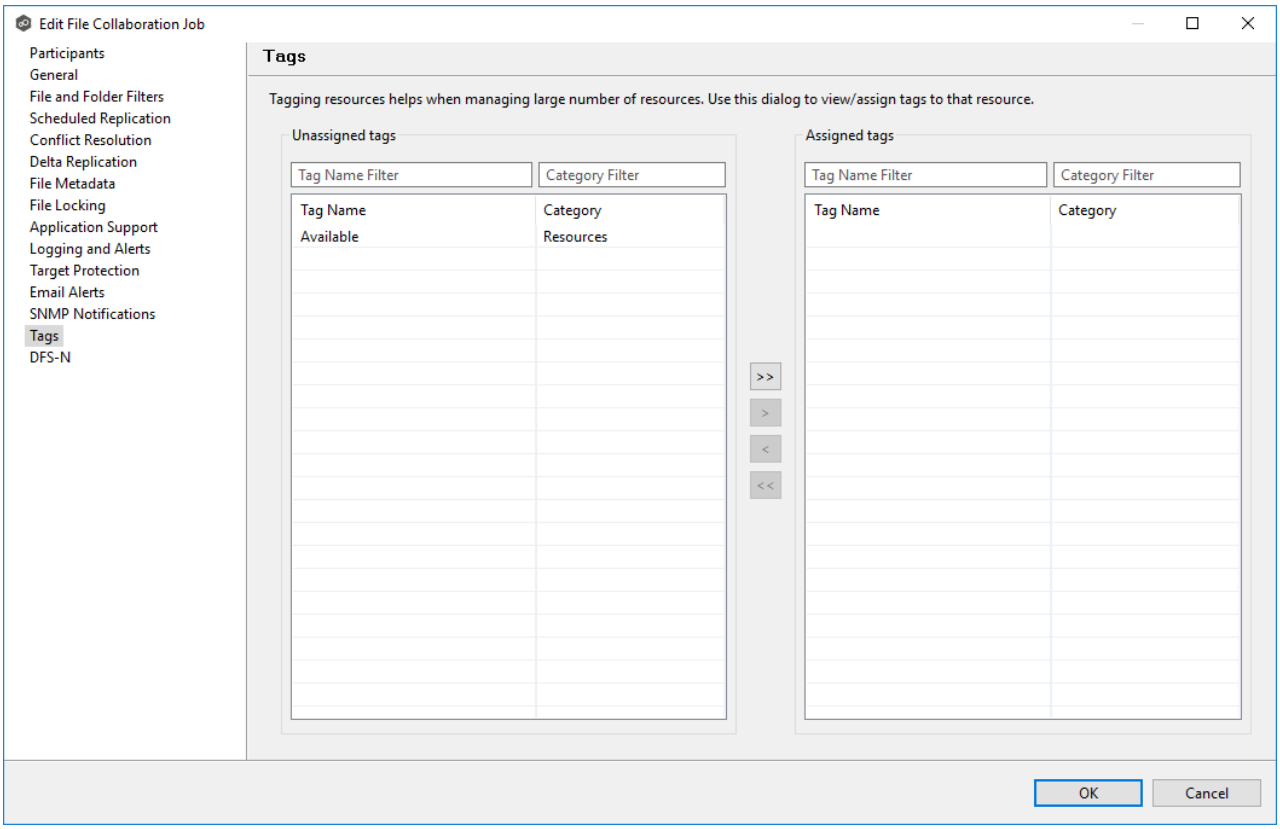

DFS-N

The **DFS-N** page in the **Edit File Collaboration Job** dialog presents options for linking a DFS namespace folder to this job. See Linking a Namespace Folder to an Existing File Collaboration or Synchronization Job for more information.

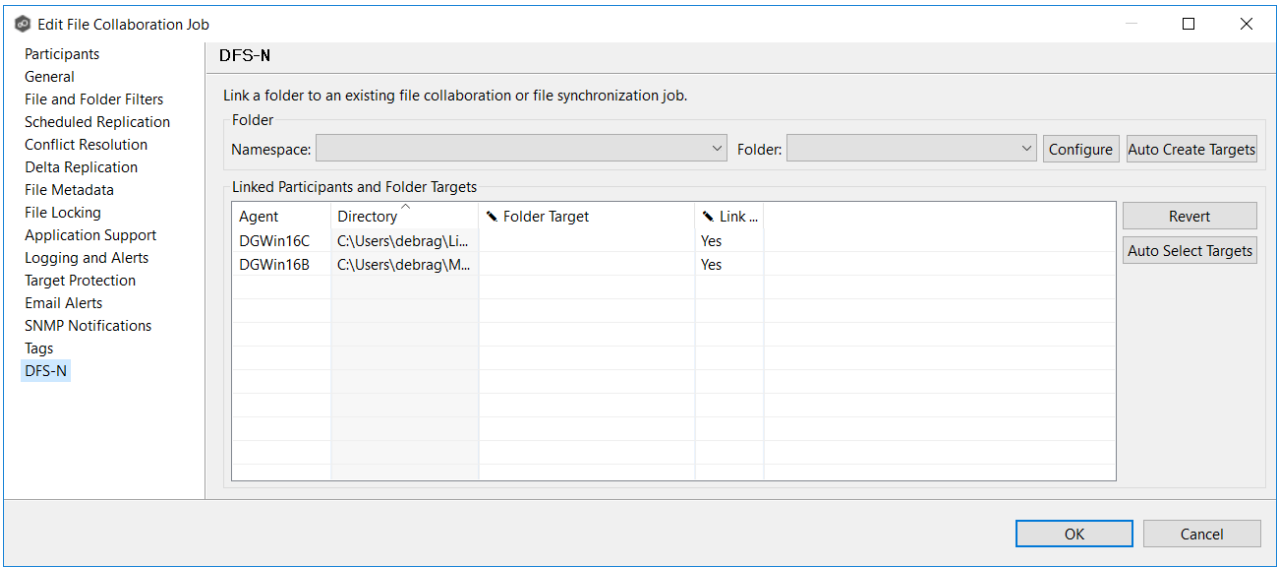

Editing Multiple Jobs

Peer Management Center supports multi-job editing, allowing you to quickly and effectively manipulate multiple File Collaboration jobs simultaneously. For example, you can use this feature to change a single configuration item such as **Auto Start** for any number of already configured jobs in one operation instead of having to change the item individually for each.

While this feature does cover most of the options available on a per-job basis, certain options are unavailable in multi-job edit mode, specifically ones tied to participants. Configuration of participants must be performed on a per job basis.

To edit multiple jobs simultaneously:

1. Open Peer Management Center.

2. Select the jobs you want to edit in the **Jobs** view.

3. Right-click and select **Edit Jobs**.

   For the most part, the original configuration dialog remains the same with a few minor differences depending on similarities between the selected File Collaboration jobs. A sample dialog is as follows:

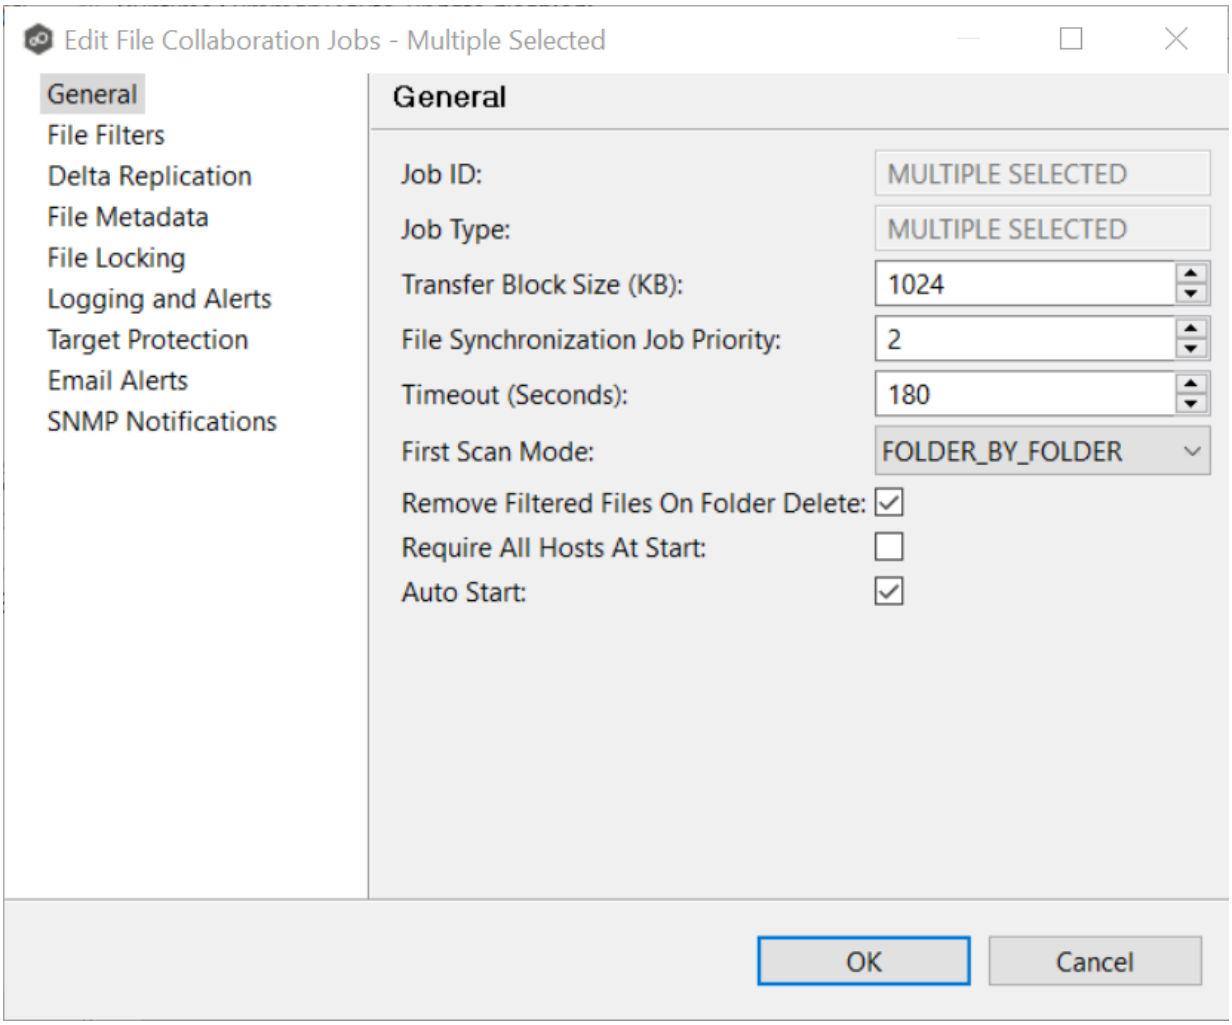

In this dialog, any discrepancies between multiple selected jobs will generally be illustrated by a read-only text field with the caption **Multiple Values - Click to Edit**. Clicking this field will bring up a dialog similar to the following:

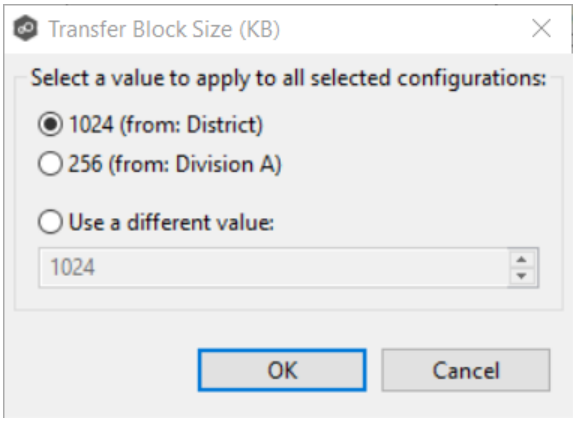

This dialog gives you the option of choosing a value that is already used by one or more selected File Collaboration jobs, in addition to the ability to use your own value. Notice that variances in the look and feel of this pop-up dialog above depend on the type of information it is trying to represent (for example, text vs. a checkbox vs. a list of items).

Upon clicking **OK**, the read-only text field you originally clicked will be updated to reflect the new value. Any fields that have changed will be marked by a small caution sign. On saving in this multi-job edit dialog, the changed values will be applied to all selected jobs.

Note: Read all information on each configuration page carefully when using the multi-job edit dialog. A few pages operate in a slightly different manner then mentioned above. All the necessary information is provided at the top of these pages in bold text.

Running and Managing a File Collaboration Job

The topics in this section provide some basic information about starting, stopping, and managing File Collaboration jobs:

- Overview
- Starting a File Collaboration Job
- Stopping a File Collaboration Job
- Auto-Restarting a File Collaboration Job
- Host Connectivity Issues
- Removing a File from Quarantine
- Manual Retries

Overview

This topic describes:

- The initialization process for a File Collaboration job: What occurs the first time you run a File Collaboration job.

- The initial synchronization process: How files are synchronized the first time you run a File Collaboration job.

The initialization process for a File Collaboration job consists of the following steps:

1. All participating hosts are contacted to make sure they are online and properly configured.

2. Real-time event detection is initialized on all participating hosts where file locks and changes will be propagated in real-time to all participating hosts. You can view real-time activity and history via the various Runtime Job views for the open job.

3. The initial synchronization process is started; all the configured root folders on the participating hosts are scanned in the background, and a listing of all folders and files are sent back to the running job.

4. The background directory scan results are analyzed, and directory structures compared to see which files are missing from which hosts. In addition, file conflict resolution is performed to decide which copy to use as the master for any detected file conflicts based on the File Conflict Resolution settings.

5. After the analysis is performed, all files that need to be synchronized are copied to the pertinent host(s).

Before you start a File Collaboration job for the first time, you need to decide how you would like the initial synchronization to be performed.

During the initial synchronization process:

- The watch set is recursively scanned on all participating hosts.

- File conflict resolution is performed.

- Any files that require updating are synchronized with the most current copy of the file.

The two primary options are:

- Have the File Collaboration job perform the initial synchronization based on the Conflict Resolution settings.

- Pre-seed all participating hosts with the correct folder and file hierarchy for the configured root folders before starting the session.

If you have a large data set, we strongly recommend that you perform the initial synchronization manually by copying the data from a host with the most current copy to all other participating hosts. This needs to be done only once--before the first time that you run the File Collaboration job.

If you choose the first option, click the **Start** button to begin collaboration session initialization. Otherwise, pre-seed each participating host with the necessary data, then click the **Start** button.

Starting a File Collaboration Job

Before starting a File Collaboration job for the first time, make sure that you have decided how you want the initial synchronization to be performed.

When running a File Collaboration job for the first time, you must manually start it. After the initial run, a job will automatically start, even when the Peer Management Center server is rebooted.

Note: You cannot run two jobs concurrently on the same volume if the watch sets contain an overlapping set of files and folders.

To manually start a job:

1. Choose one of three options:

   - Right-click the job name in the **Jobs** view.

   - Right-click the job name in the **Summary** tab of the **Collab, Sync, and Repl Summary** summary view, and then choose **Start** from the context menu.

   - Open a job and then click the **Start/Stop** button to the left of the **Status** field in the bottom left corner of the job's runtime view (shown below).

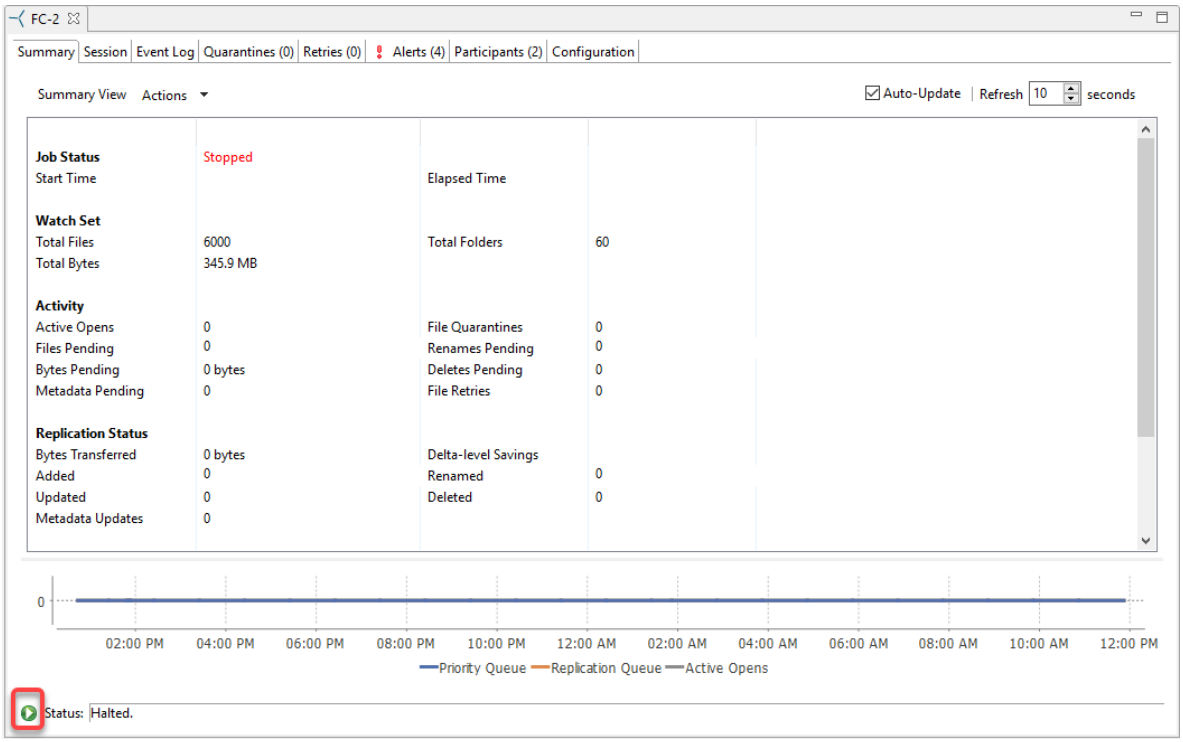

2. Click **Yes** in the confirmation dialog.

   After the job initialization has completed, the job will run. Once the job starts, the icon next to the job name in the **Jobs** view changes from gray to green.

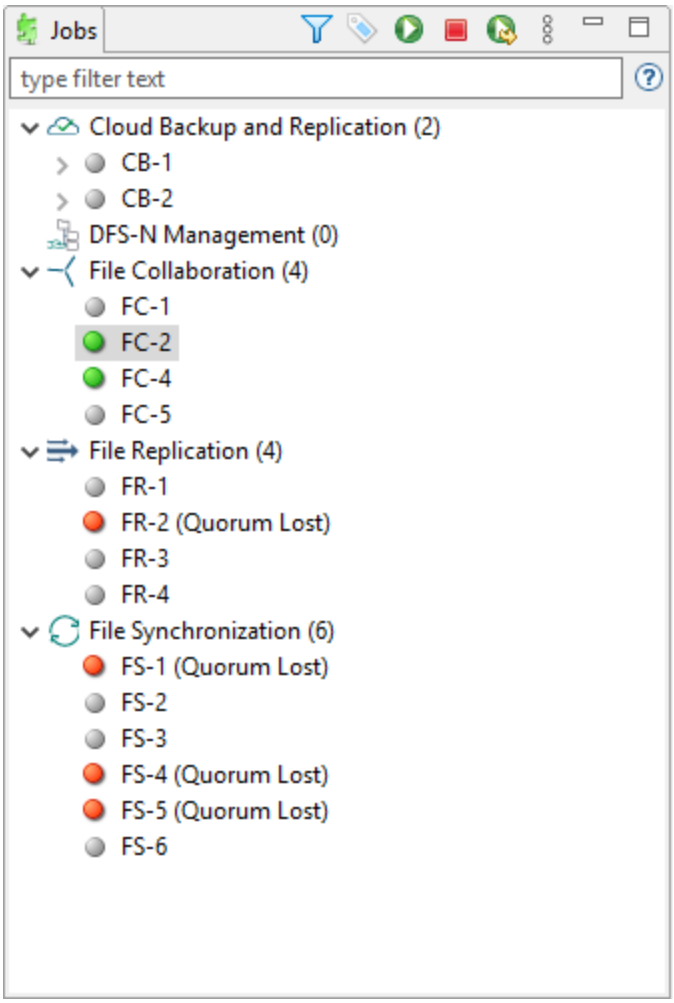

Stopping a File Collaboration Job

You can stop a File Collaboration job at any time by selecting the job in the **Jobs** view and clicking the **Stop** button. Doing this shuts down the real-time file event detection and close all running operations (e.g., file transfers).

Auto-Restarting a File Collaboration Job

Peer Management Center includes support for automatically restarting File Collaboration jobs that include participating hosts that have been disconnected, have reconnected, and are once again available.

After a host becomes unavailable and the quorum is lost on a running File Collaboration job, the job automatically stops running and enters a pending state, waiting for one or more hosts to become available again so that the quorum can be met. Once the quorum is met, the pending job will automatically be restarted, beginning with a scan of all root folders.

In a job where a host becomes unavailable but quorum is not lost, the remaining hosts will continue collaborating. If the unavailable host becomes available once again, it is brought back into the running job and a background scan begins on all participating hosts, similar in fashion to the initial background scan at the start of a job.

You can enable all File Collaboration jobs to auto-restart. You can also disable auto-restart File Collaboration jobs on a per-job and per-host instance. For more information on disabling auto-restart at the job level, see Participants Tab.

To enable all File Collaboration jobs to auto-restart:

1. Select **Preferences** from the **Window** menu.

2. Select **Collab, Sync, and Repl Summary** in the navigation tree.

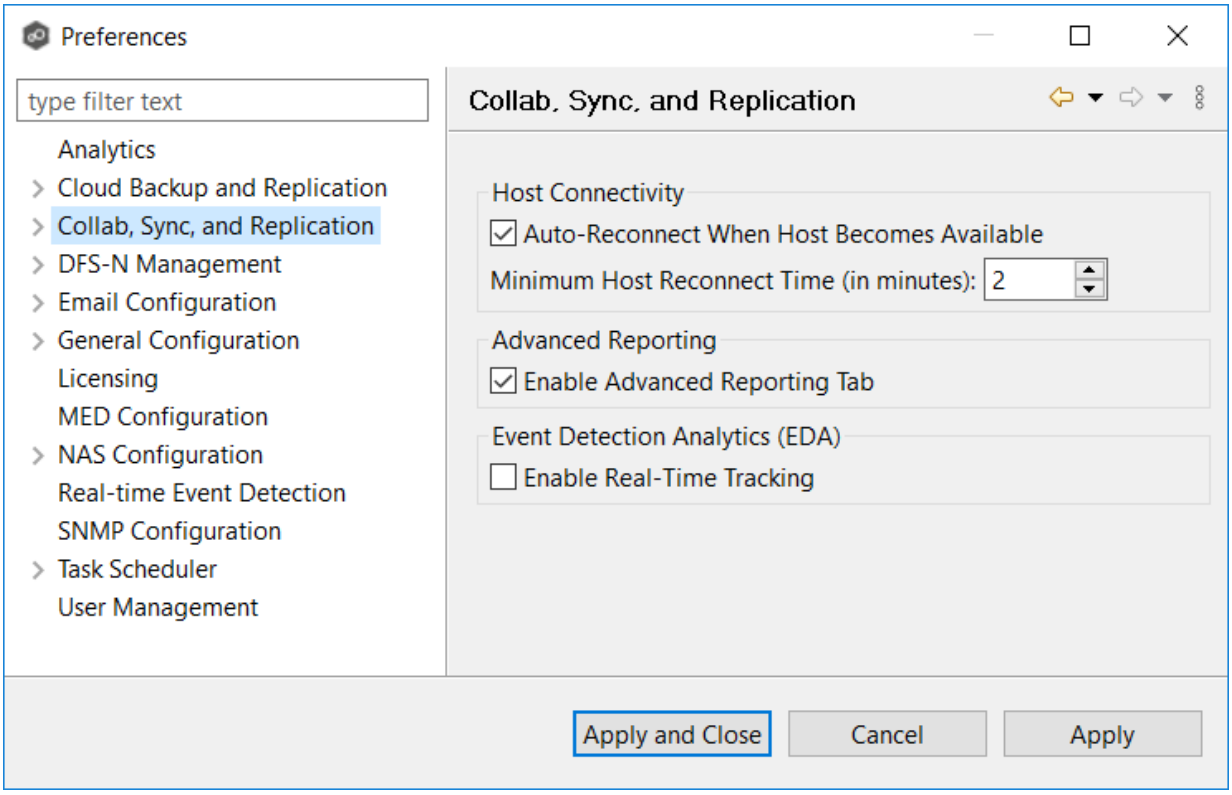

3. Select the **Auto Reconnect when Host Becomes Available** checkbox.

4. Enter the minimum number of minutes to wait after an Agent reconnects before re-enabling it in any associated jobs in the **Minimum Host Reconnect Time** field.

5. Click **OK**.

Host Connectivity Issues

Peer Management Center is designed to be run in an environment where all participating hosts are highly available and on highly available networks. The two primary connectivity issues result from:

- Unavailable Hosts
- Quorum Not Met

Unavailable Hosts

If a host becomes unavailable while a File Collaboration job is running and is unreachable within the configured timeout period (specified in the job's General settings), it may be removed from collaboration. If no response is received while performing a file collaboration operation within the timeout period, Peer Management Center pings the host; if still no response, the host is taken out of the running session, a FATAL event is logged , and the Participants tab for the job is updated to indicate that the host has failed. In addition, if email alerts and/or SNMP notifications are configured and enabled for **Host Timeouts**, then the appropriate message(s) are sent.

If auto-restart not enabled, you must stop and start the File Collaboration job to bring any failed hosts back into the session. As a result, all root folders on all hosts will need to be scanned again to detect any inconsistencies. Therefore, if you are operating over a WAN with low bandwidth, you will want to set the timeout to a higher value on each related job.

Quorum Not Met

For a File Collaboration job to run correctly, a quorum of available hosts must be met. When a quorum is lost, a message appears after the job name in the **Jobs** view.

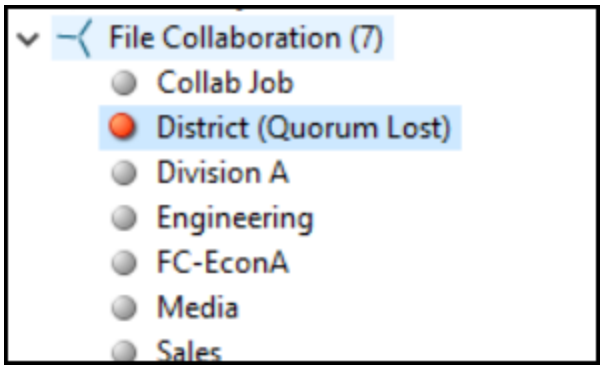

Quorum is currently set to at least two hosts, and if quorum is not met, then the collaboration session is automatically be terminated. If email alerts and/or SNMP notifications are configured and enabled for **Session Aborts**, then the appropriate message(s) are sent.

Removing a File from Quarantine

Quarantines are a key feature of Peer Global File Service, used for resolving version conflicts. For more detailed information on how quarantines work, see Conflicts, Retries, and Quarantines in the Advanced Topics section.

You must explicitly remove a file from quarantine in order to have it participate in the collaboration session once again.

You may also choose to perform no action, in which case, the file is removed from the **Quarantines** table but none of the file versions are modified. Therefore, if the files are not currently in-sync, then the next time the file is modified, changes will be propagated to the other hosts. If an error occurs while removing the quarantine, then the **Status** field in the **Quarantines** table is updated to reflect the error.

To remove a file (or multiple files) from quarantine:

1. Open the job.
2. The runtime view for the job appears.
3. Click the Quarantines tab.
4. Select the file(s) in the **Quarantines** table.
5. Select the host with the correct version.
6. Click the **Release Conflict** button.

   After doing this, all hosts are checked to make sure the file is not currently locked by anyone. If no locks are found, then locks are obtained on all versions of the file and the targets that are out-of-date are synchronized with the selected source host.

Manual Retries

Retries are a key feature of Peer Global File Service, used for automatically handling errors in the collaboration environment that would have otherwise led to a quarantine. For more detailed

information on how retries work, see Conflicts, Retries, and Quarantines in the Advanced Topics section.

When a file is put in the retry list, it will be automatically retried based on the settings defined in File Retries in Preferences. If need be, you can also manually force the retry of a file. This can be done from the Retries list of a specific File Collaboration job.

You may also choose to perform no action, in which case, the file is removed from the Retries list but none of the file versions are modified. Therefore, if the files are not currently in-sync, then the next time the file is modified, changes will be propagated to the other hosts. If an error occurs while forcing the retry, then the **Status** field in the **Retries** table is updated to reflect the error.

To manually force the retry of a file (or multiple files):

1. Select the file(s) in the **Retries** list.
2. Select the host with the correct version.
3. Click the **Release Conflict** button.

   After doing this, all hosts are checked to make sure the file is not currently locked by anyone. If no locks are found, then locks are obtained on all versions of the file and the targets that are out-of-date are synchronized with the selected source host.

File Replication Jobs

This section provides information about creating a File Replication job.

- Overview
- Before You Create Your First File Replication Job
- Creating a File Replication Job

Overview

A File Replication job is designed to push files one way from a single file server (known as the source) to another single file server (known as the destination or target). This job type requires two Agents, although only the Agent at the source location will register with its local storage

platform for real-time activity. The destination Agent will simply act as a relay to the destination file server.

Before You Create Your First File Replication Job

We strongly recommend that you configure global settings such as SMTP configuration and email alerts before configuring your first File Replication job. See Preferences for details on what and how to configure these settings.

Creating a File Replication Job

The **Create Job Wizard** walks you through the process of creating a File Replication job:

Step 1: Job Type and Name

Step 2: Source Platform

Step 3: Source Agent

Step 4: Storage Information

Step 5: Source Path

Step 6: Destination Agent

Step 7: Destination Path

Step 8: File Metadata

Step 9: Email Alerts

Step 10: Save Job

Step 1: Job Type and Name

1. Open Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **Create New Job** wizard displays a list of job types you can create.

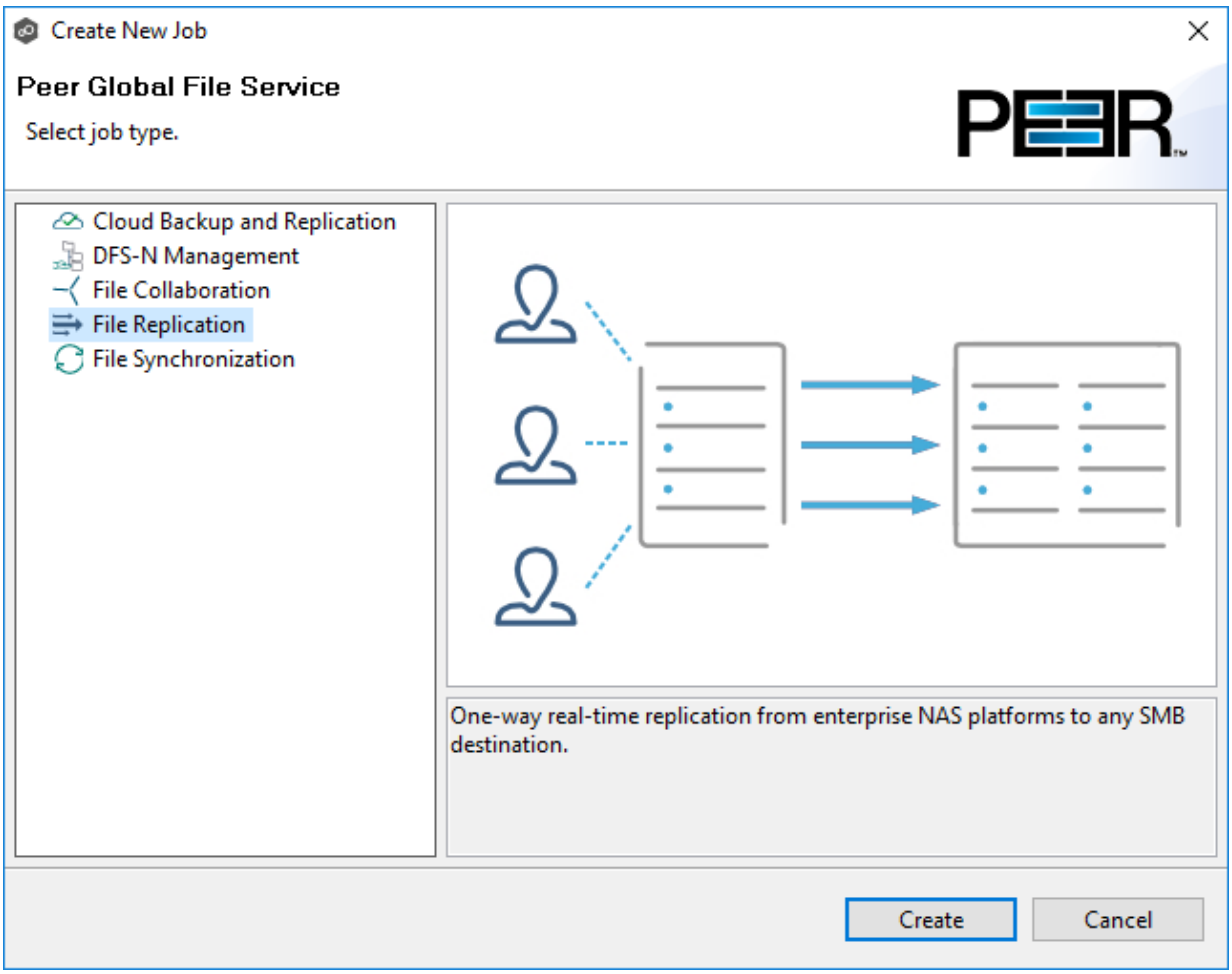

3. Click **File Replication**, and then click **Create**.
4. Enter a name for the job in the dialog that appears.

   The job name must be unique.

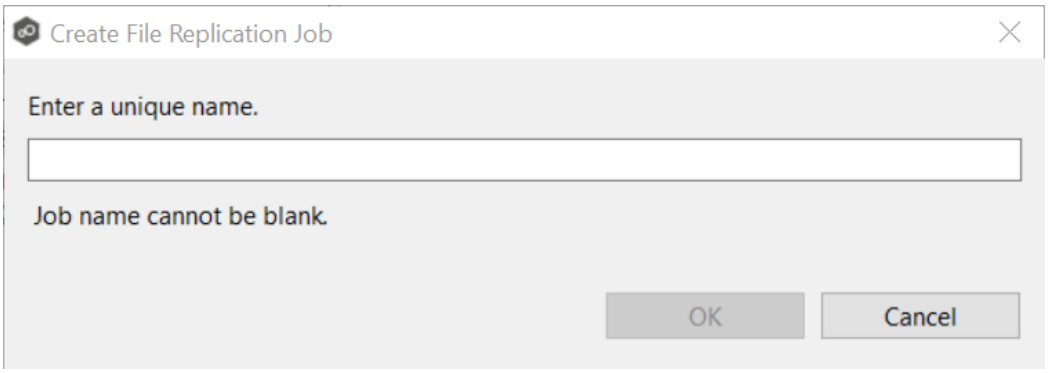

5. Click **OK**.

   The Storage Platform page is displayed.

Step 2: Storage Platform

The **Storage Platform** page lists the types of source storage platforms that File Replication supports. The source storage device hosts the data you want to replicate.

1. Select the type of storage platform you want to replicate.

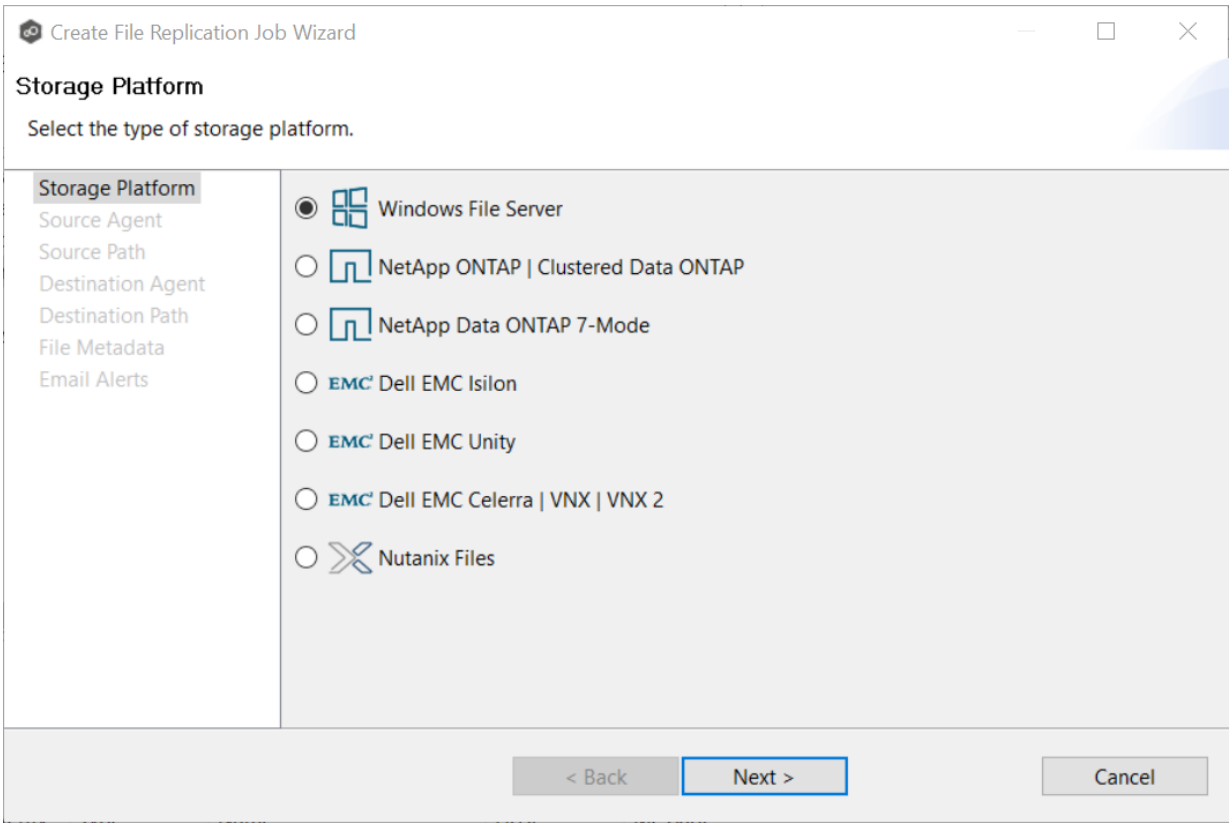

2. Click **Next**.

   The Source Agent page is displayed.

Step 3: Source Agent

The **Source Agent** page lists available Agents. You can have more than one Agent managing a storage device—however, the Agents must be managing different volumes/shares/folders. You should select the Agent that manages the volumes/shares/folders you want to replicate in this job.

1. Select the Source Agent for the volume/share/folder you want replicated.

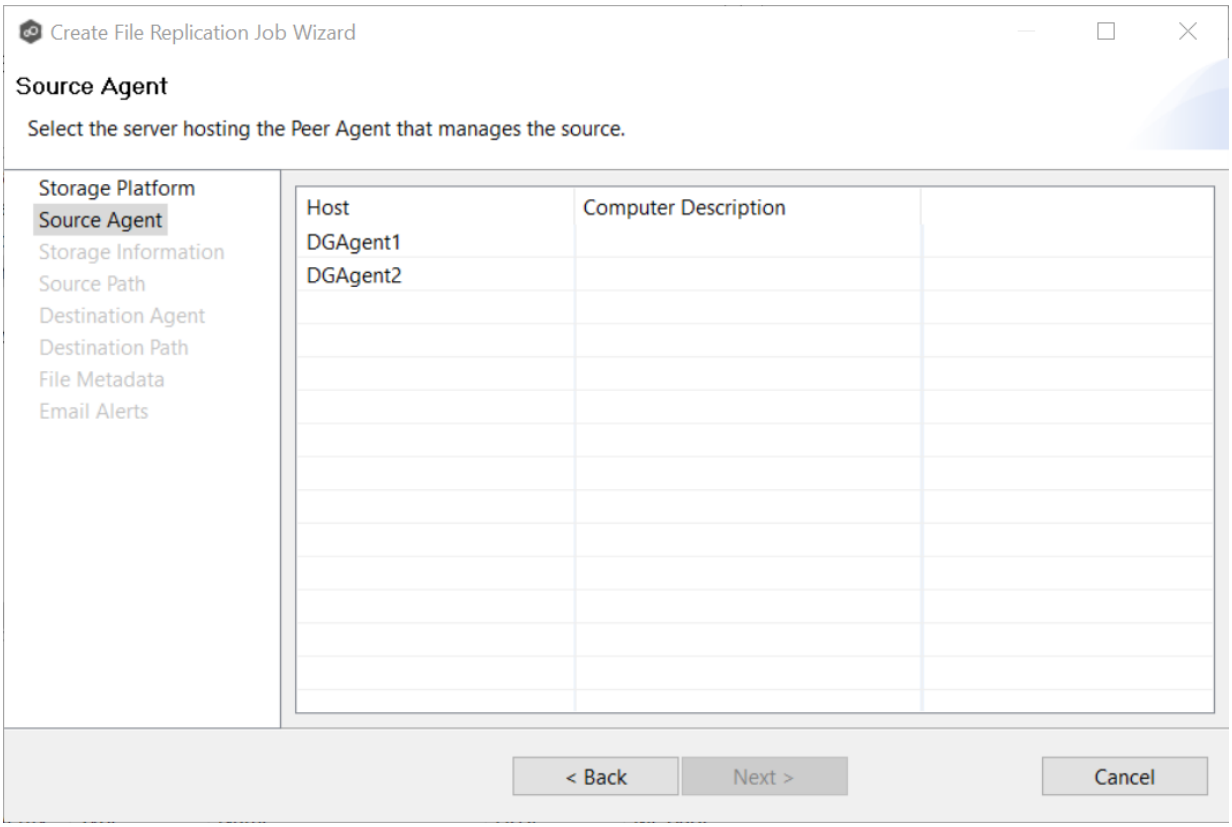

2. Click **Next**.

   The Storage Information page is displayed.

Step 4: Storage Information

The **Storage Information** page varies, depending on your selection in the Storage Platform page:

- If you selected **Windows File Server**, you are prompted for configuration relating to Windows File Server. See Windows File Server.
- If you selected any other type of storage device other than Windows File Server, the **Storage Information** page requests the credentials necessary to connect to that storage device. Continue with the following steps.

For storage platforms other than Windows File Server:

1. Select **New Credentials** or **Existing Credentials**.

2. If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next**. Continue with Step 5. Source Path.

   If you selected **New Credentials**, enter the credentials for connecting to the storage platform. The information you are prompted to enter varies, depending on the type of storage platform:

   NetApp ONTAP | Clustered Data ONTAP

   NetApp Data ONTAP 7-Mode

   Dell EMC Isilon

   Dell EMC Unity

   Dell EMC Celerra | VNX | VNX 2

   Nutanix Files

3. Click **Validate** to test the credentials, and then click **OK** in the confirmation message that appears.

4. Click **Next**.

   The Source Path page is displayed.

1. Select the **Detector Type**.

   Note: Currently the only option is **Windows**.

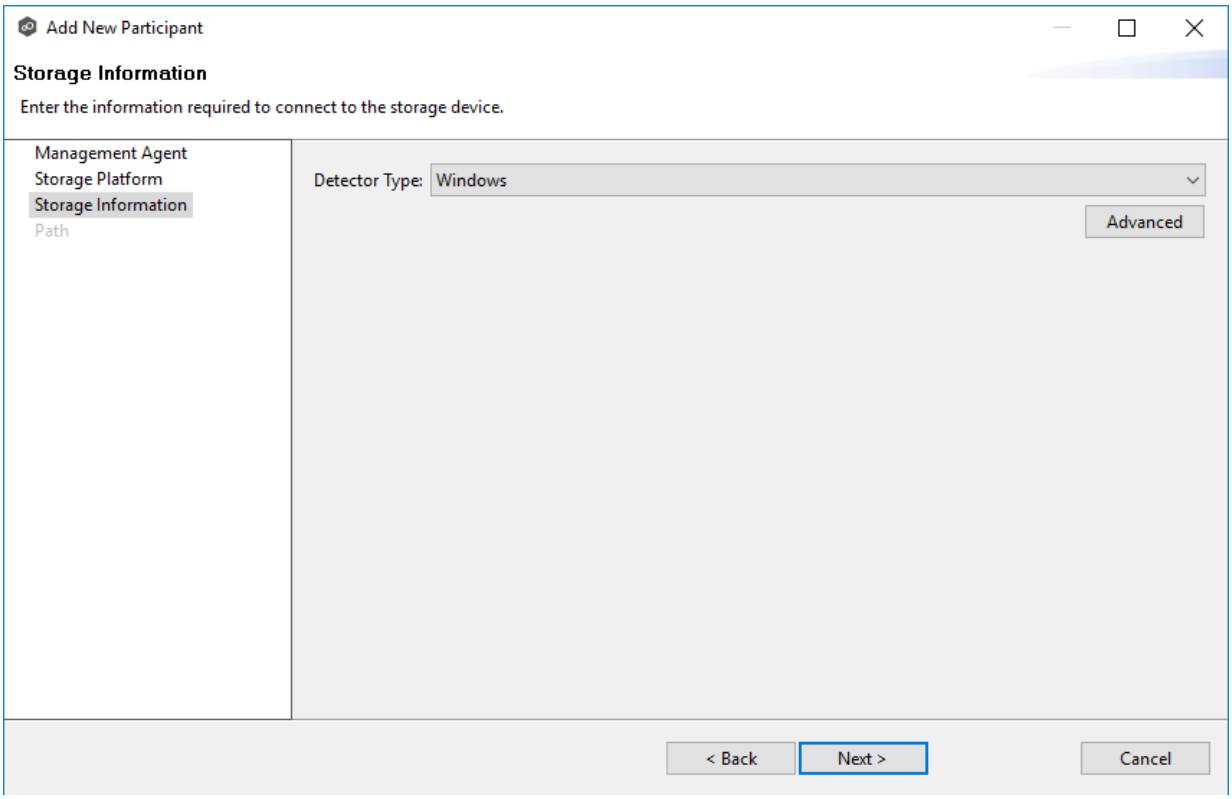

2. Click **Advanced** if you want to set advanced options.
3. Click **Next**.

Windows File Server Advanced Options

1. Modify the options as desired.

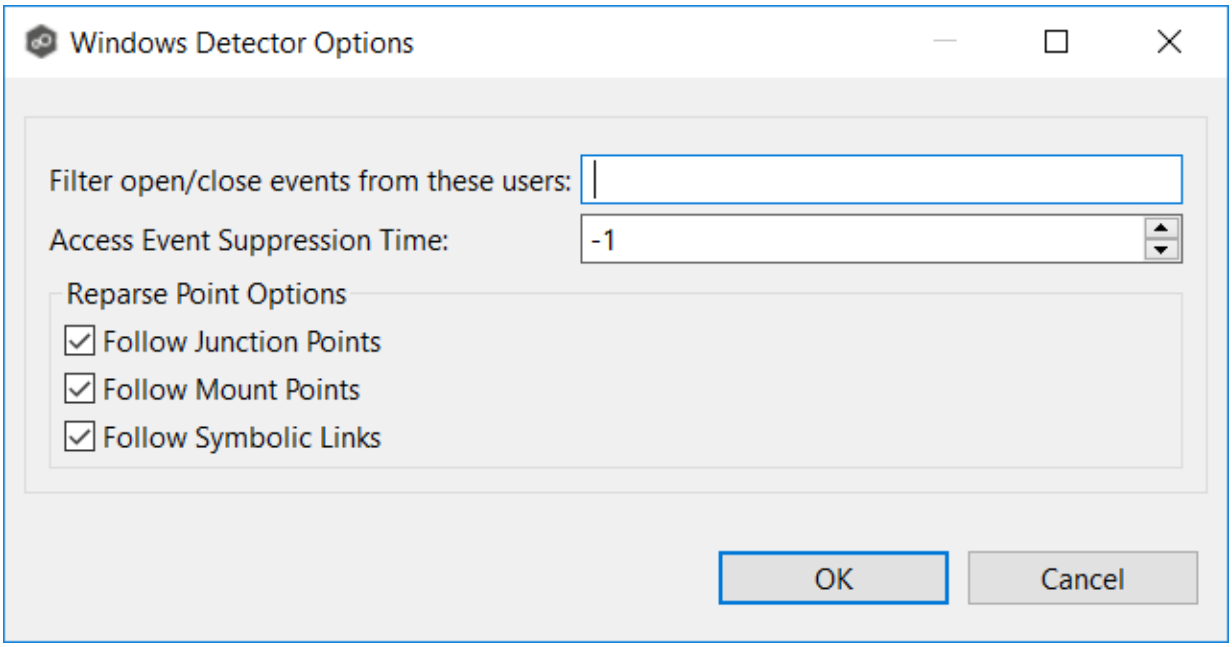

Option	Description
Filter open/close events from these users	A comma-separated list of user account names from which all opens and closes will be ignored. Ideal for filtering out events from backup and/or archival services by filtering on the username under which a backup and/or archival service is running.
Access Event Suppression Time	Represents number of seconds an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default vault is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time an open event will be delayed based on the load of the system.
Follow Junction Points	Enables junction point support for the selected Windows File Server.
Follow Mount Points	Enables mount point support for the selected Windows File Server.
Follow Symbolic Links	Enables symbolic link support for the selected Windows File Server.

For more information about junction points or symbolic links, contact <% SUPPORT_EMAIL%

2. Click **OK**.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

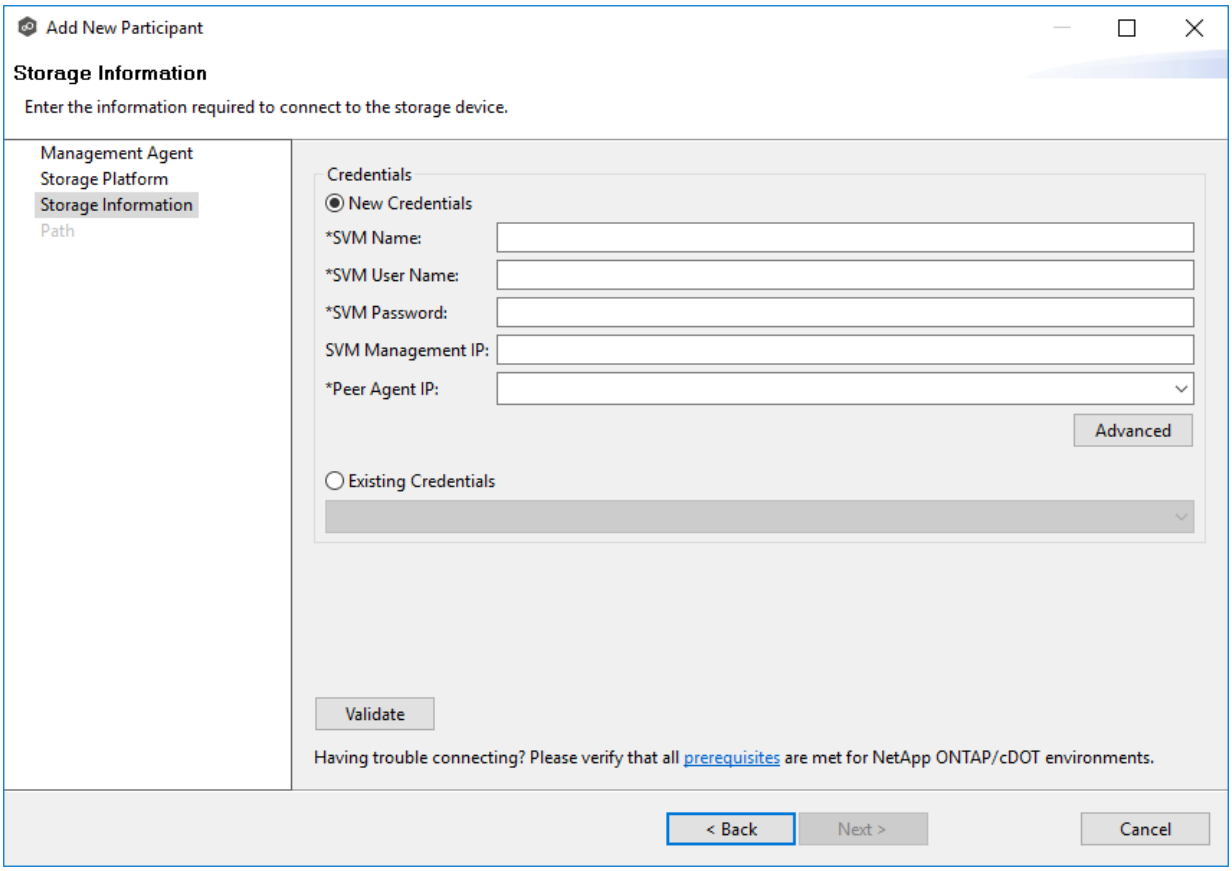

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information

Field	Description
SVM Name	Enter the name of the Storage Virtual Machine hosting the data to be replicated.
SVM User Name	Enter the user name for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	(Optional) Enter the IP address used to access the management API of the NetApp Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already allow management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Source Path page is displayed.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

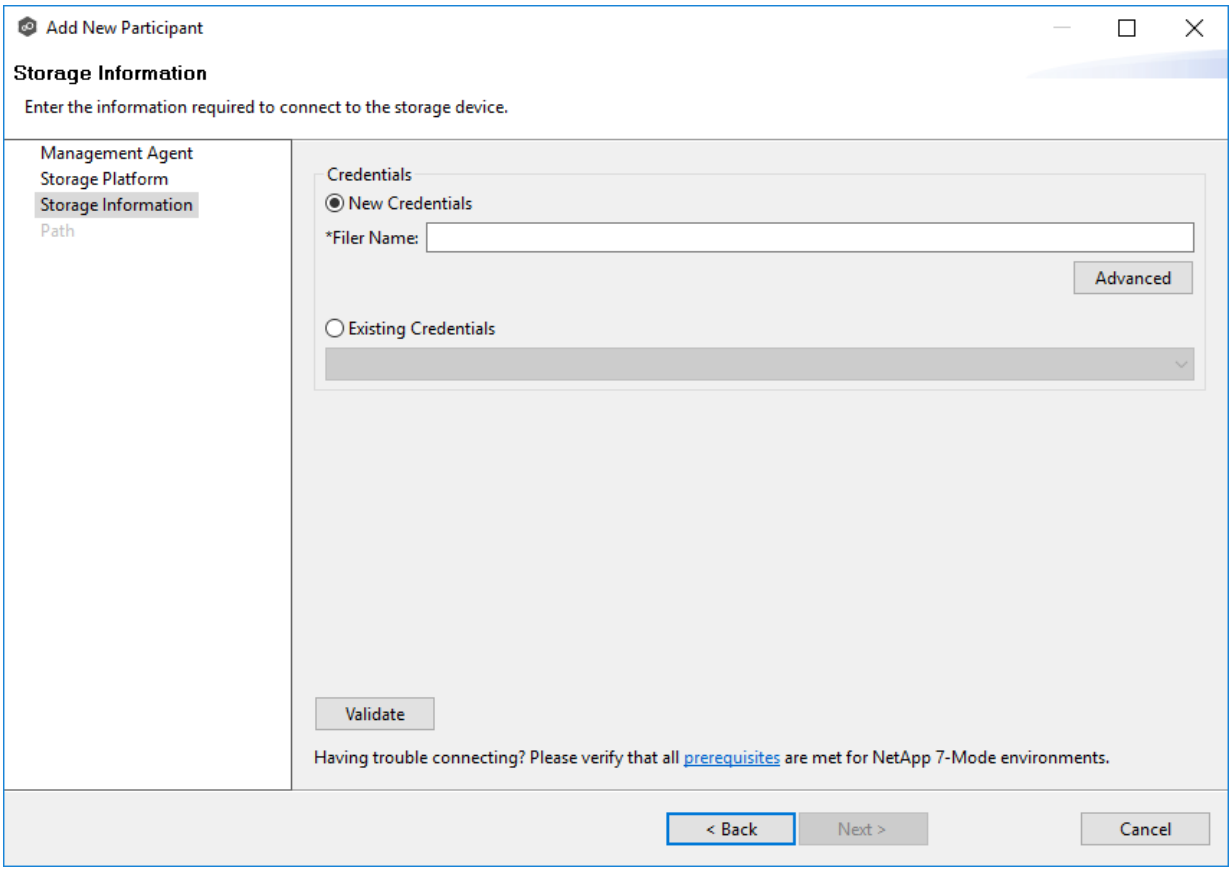

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
Filer Name	Enter the name of the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

3. Click **Advanced** if you want to set advanced options
4. Click **Validate**.
5. Click **Next**.

   The Source Path page is displayed.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the EMC Isilon cluster hosting the data to be replicated.

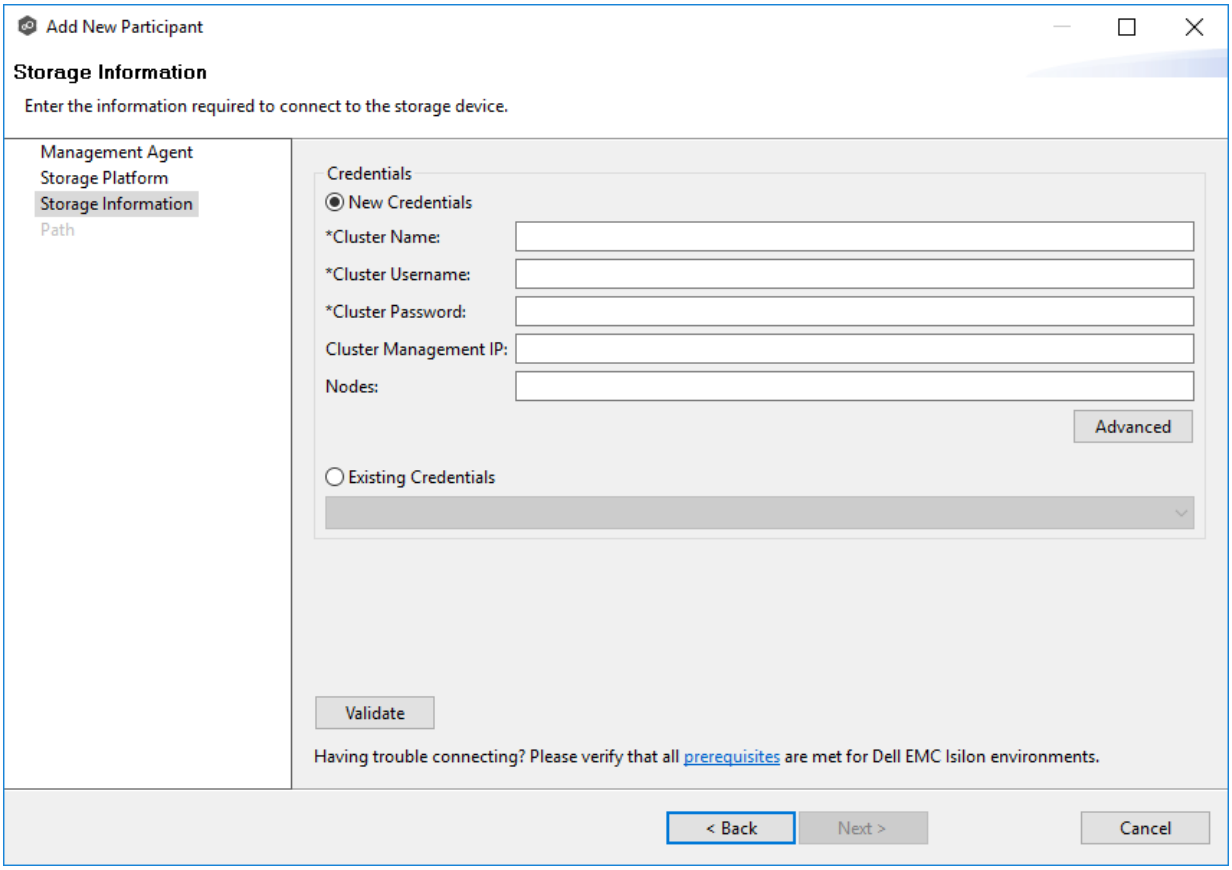

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
Cluster Name	Enter the name of the Dell EMC Isilon cluster hosting the data to be replicated.
Cluster Username	Enter the user name for the account managing the Dell EMC Isilon cluster.
Cluster Password	Enter the password for account managing the Dell EMC Isilon cluster.
Cluster Management IP	(Optional) Enter the IP address of the system used to manage the Dell EMC Isilon cluster. Required only if multiple Access Zones are in use on the cluster.
Nodes	(Optional) Enter one IP from each node in the Isilon cluster that the Agent can access to perform open file lookups. Use commas to separate nodes.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Source Path page is displayed.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the CIFS Server hosting the data to be replicated.

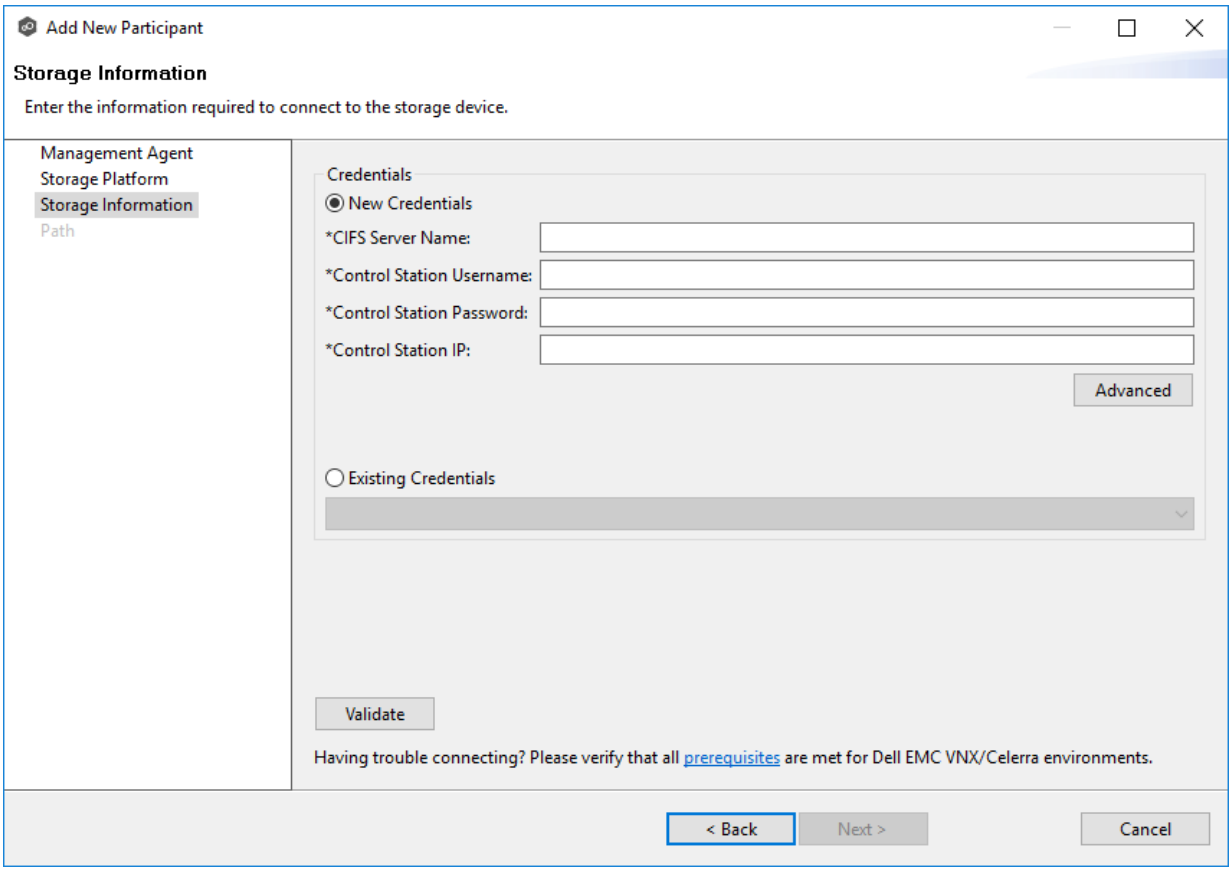

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
CIFS Server Name	Enter the name of the CIFS Server hosting the data to be replicated.
Control Station Username	Enter the user name for the Control Station account managing the Celerra/VNX storage device.
Control Station Password	Enter the password for the Control Station account managing the Celerra/VNX storage device.
Control Station IP	Enter the IP address of the Control Station system used to manage the Celerra/VNX storage device. This should not point to the CIFS Server.

3. Click **Advanced** if you want to set advanced options
4. Click **Validate**.
5. Click **Next**.

   The Source Path page is displayed.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the CIFS Server hosting the data to be replicated.

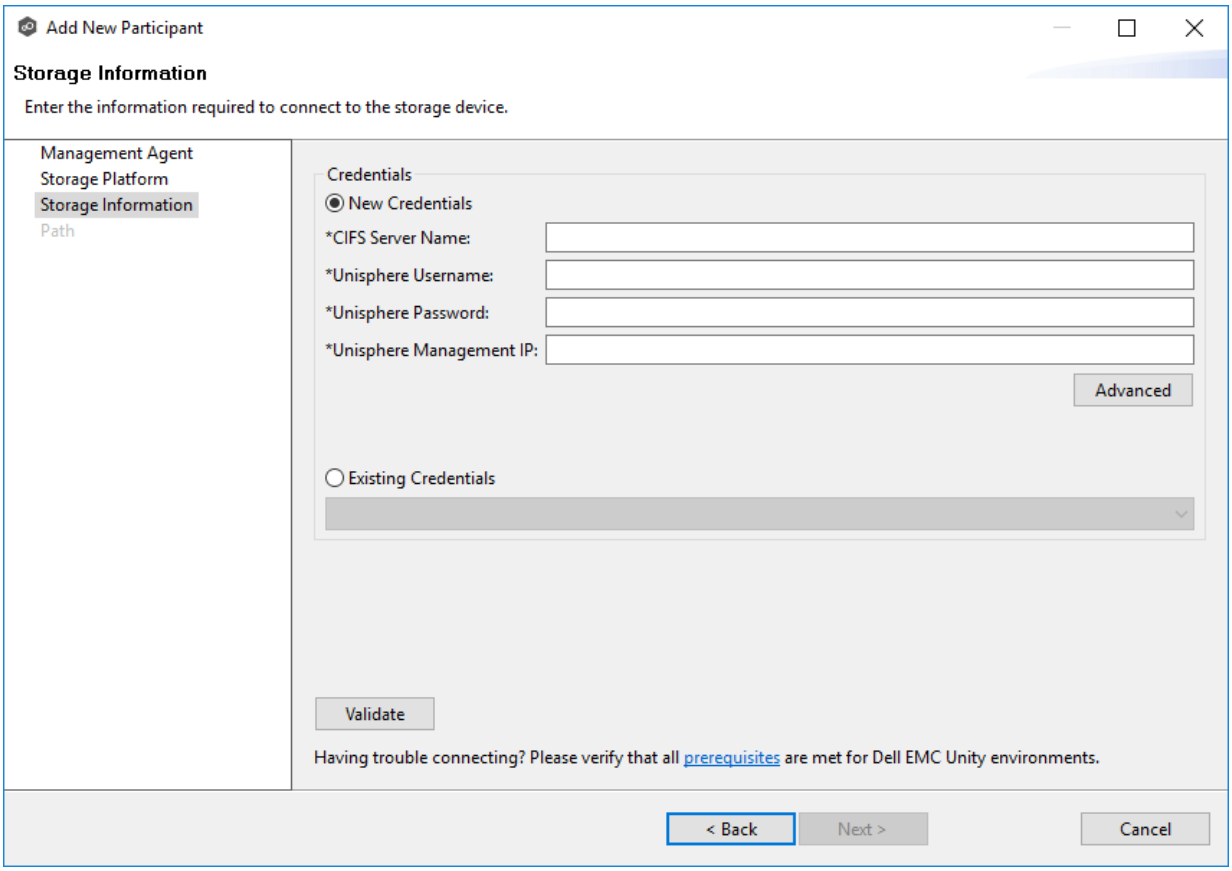

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
CIFS Server Name	Enter the name of the NAS server hosting the data to be replicated.
Unisphere Username	Enter the user name for the Unisphere account managing the Unity storage device.
Unisphere Password	Enter the password for the Unisphere account managing the Unity storage device.
Unisphere Management IP	Enter the IP address of the Unisphere system used to manage the Unity storage device. This should not point to the NAS server.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Source Path page is displayed.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated.

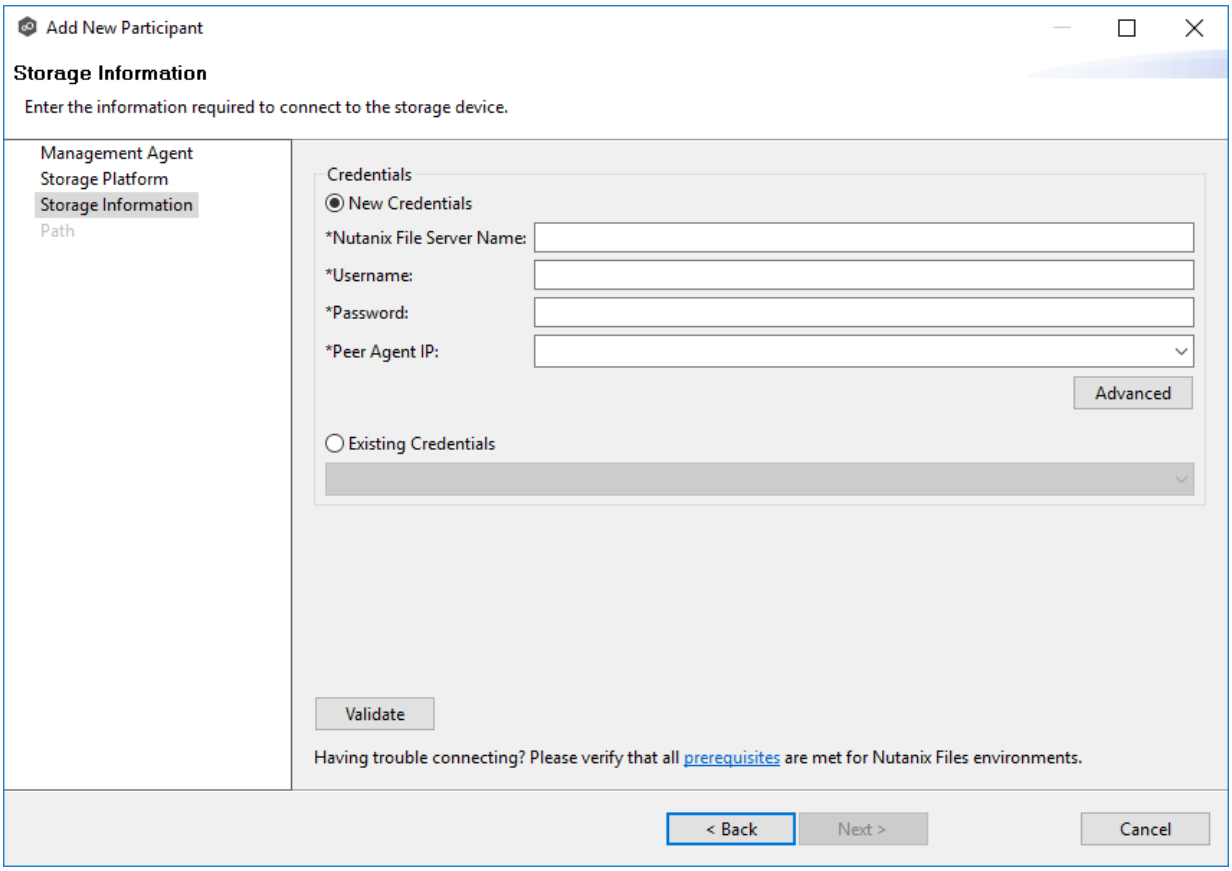

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address. This should not point to the Files cluster itself.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Source Path page is displayed.

Step 5: Source Path

The **Source Path** page is where you specify the path to the volume/share/folder you want to replicate. This volume/share/folder is referred to as the watch set. The watch set can contain a single volume/share/folder. If you want to replicate multiple volumes/shares/folders, you need to create a separate job for each one.

1. Browse to or enter the path to the watch set.

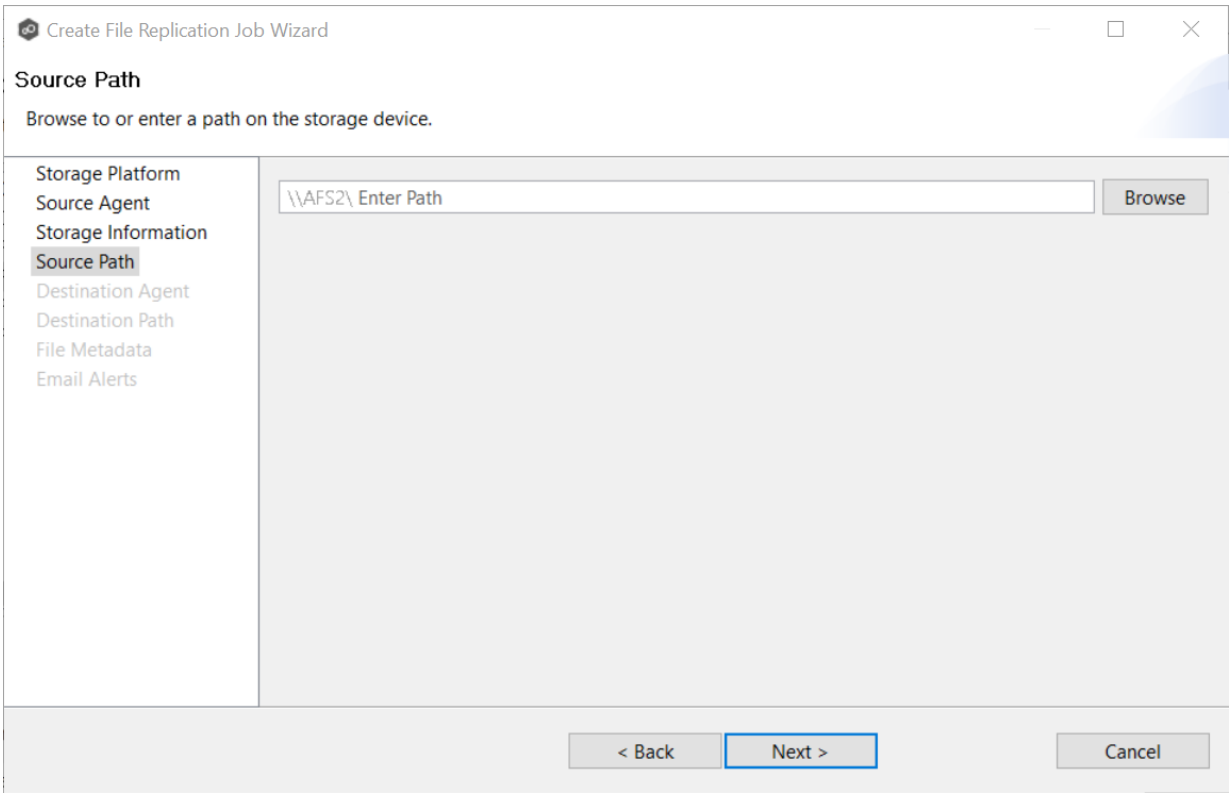

If you selected **Browse**, the **Folder Browser** dialog appears:

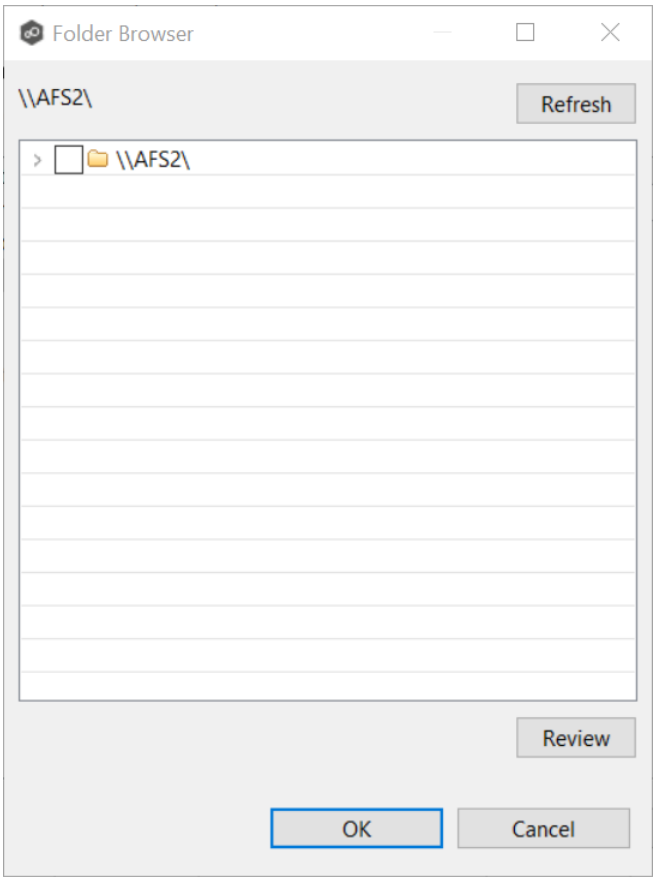

a. Expand the folder tree.

b. Select the appropriate volume/share/folder.

c. (Optional) Click the **Review** button to see your selection.

d. Click **OK**.

2. Click **Next**.

   The Destination Agent page is displayed.

Step 6: Destination Agent

The **Destination Agent** page lists available Agents, not including the Agent used as the Source Agent. This Destination Agent will be responsible for writing files and metadata to the destination storage device. No credentials are required for this Agent as it will not be monitoring anything in real-time.

1. Select the Agent that manages the destination storage device. If the destination is a Windows file server, the Agent should be installed on it.

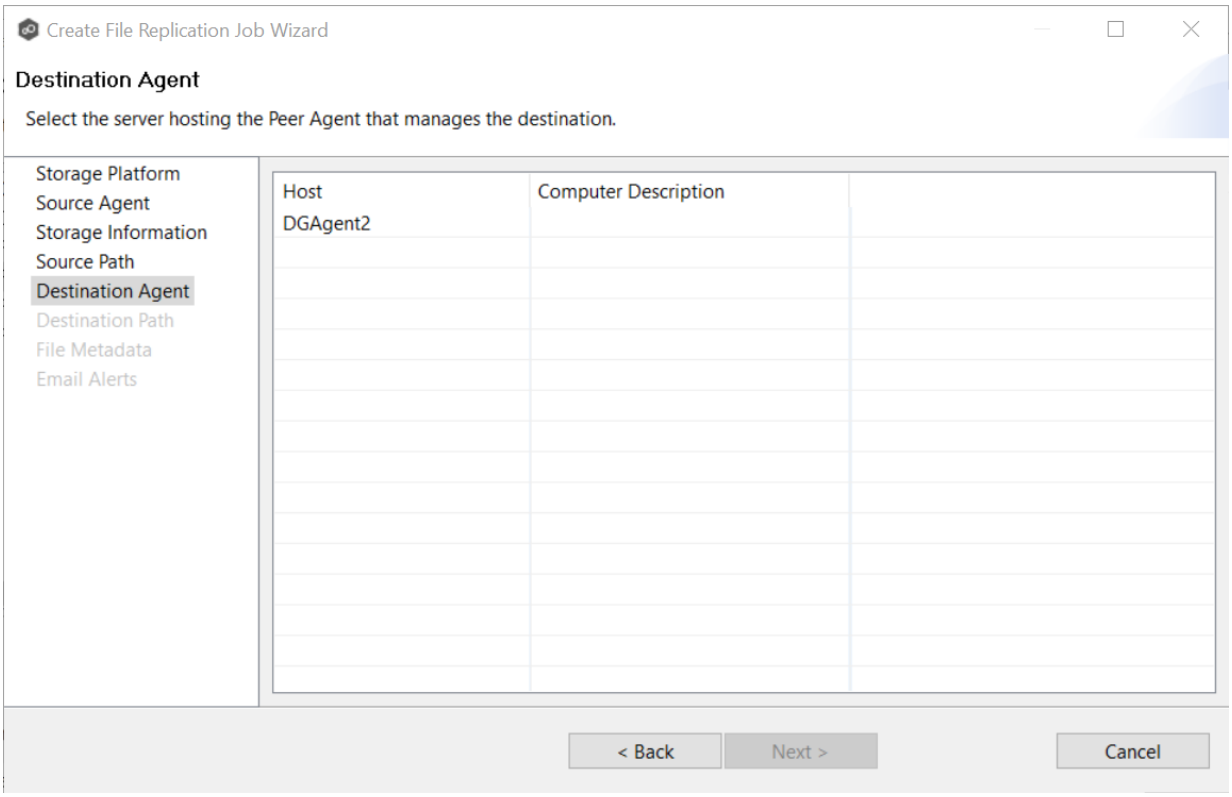

Tip: If the Agent you want is not listed, try restarting the Peer Agent Windows Service on that host. If it successfully connects to the Peer Management Broker, then the list is updated with that Agent.

2. Click **Next**.

   The Destination Path page is displayed.

Step 7: Destination Path

The **Destination Path** page is where you specify the volume/share/folder that you want to replicate to. If the destination storage device is a Windows file server, this path should be a local path such as D:\Data. This path can also be the UNC path to any SMB-capable file server.

1. Browse to or enter the destination path:

   - If the path field is empty when you click **Browse**, the **Folder Browser** dialog will present a list of local drives and folders on the Agent server itself.

- If you enter the start of a UNC path and click **Browse**, the **Folder Browser** dialog will attempt to present a list of the available shares on the file server specified in the path.

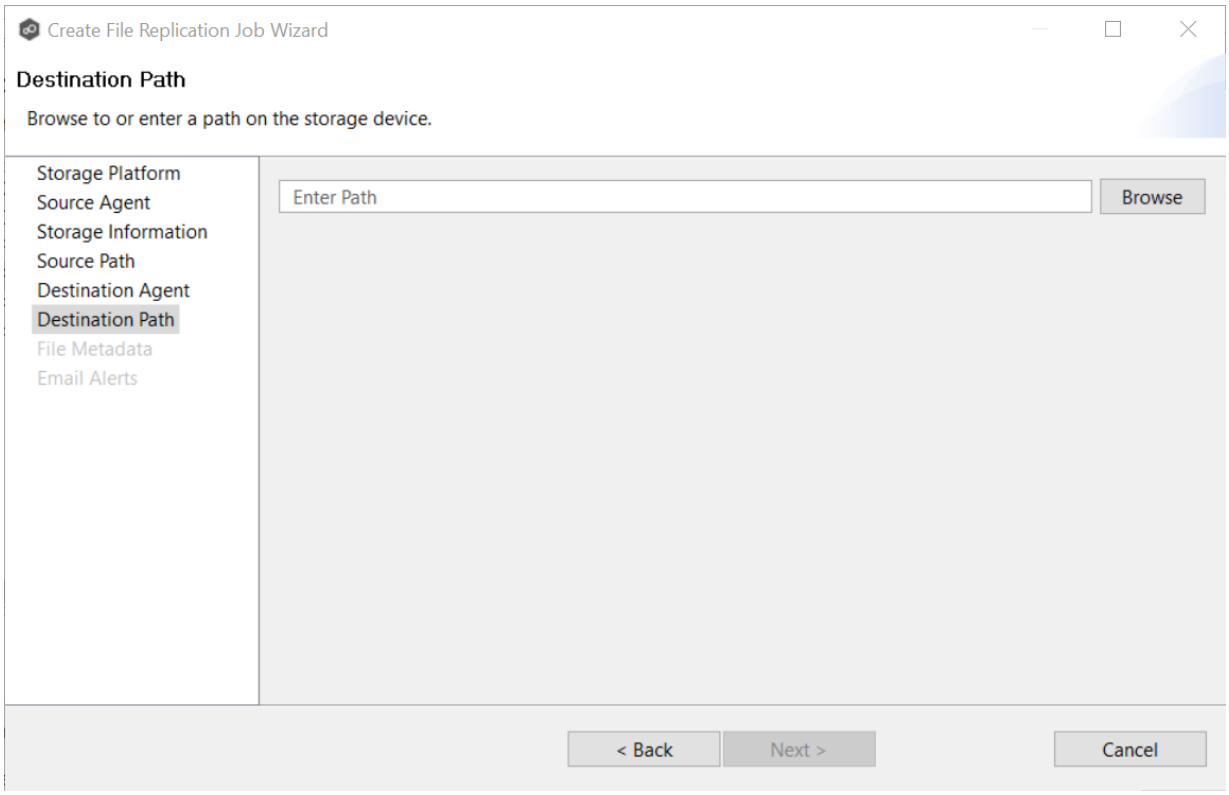

2. Click **Next**.

   The File Metadata page is displayed.

Step 8: File Metadata

This step is optional.

The **File Metadata** page allows you to specify whether you want to replicate NTFS security permissions metadata and the types of metadata to synchronize. It also allows you to specify which volume/share/folder's metadata should be used if there is a conflict during the initial synchronization. The volume/share/folder used if there is a conflict is referred to as the master host.

For more information on synchronizing NTFS metadata, see File Metadata Synchronization in the Advanced Topics section.

To enable file metadata synchronization:

1. Select when you want the metadata replicated (you can select one or both options):

   - **Enable synchronizing NTFS security descriptors (ACLs) in real-time** - Select this option if you want the metadata synchronized in real-time. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be transferred to the target host file(s) as they occur.

   - **Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan** - Select this option if you want the metadata replicated during the initial scan. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized during the initial scan.

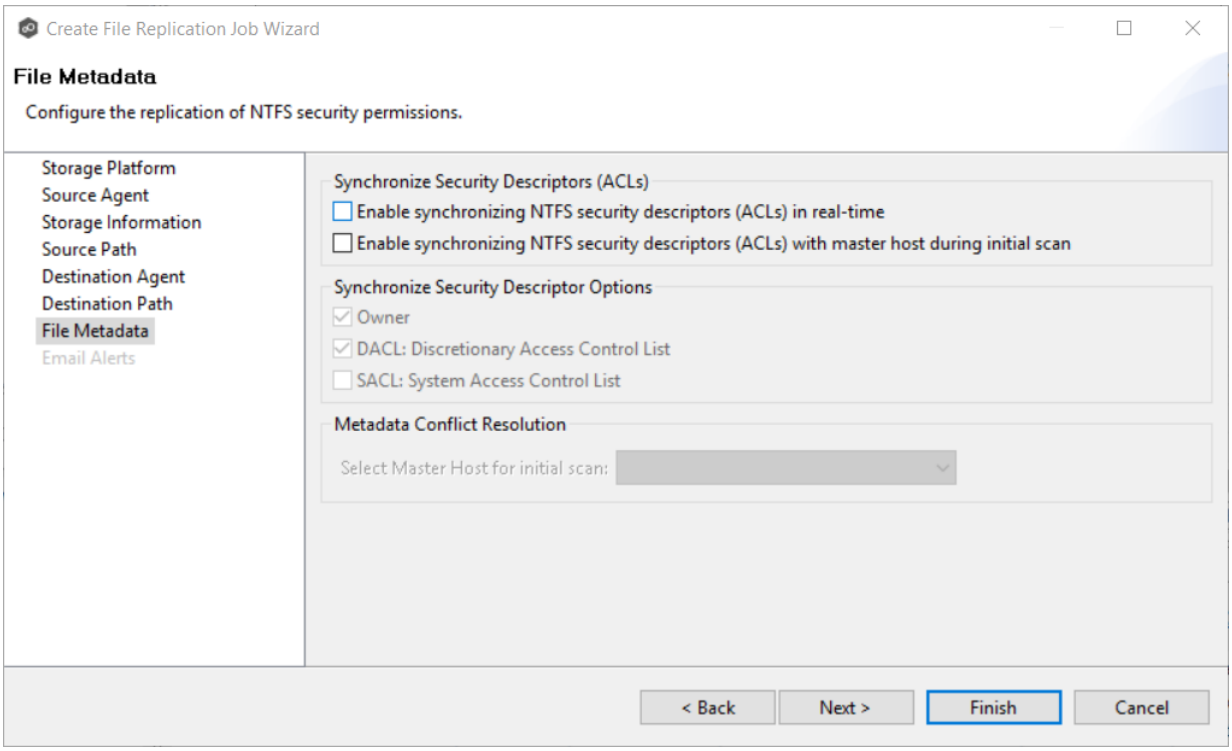

2. Click **OK** in the message that appears after selecting a metadata option.

3. If you selected either of the first two options in the **Synchronize Security Descriptor Options** section, select the security descriptor components (Owner, DACL, and SACL) to be synchronized.

4. If you selected the option for metadata synchronization during the initial scan, select the host to be used as the master host in case of file metadata conflict.

   If a master host is not selected, then no metadata synchronization will be performed during the initial scan. If one or more security descriptors do not match across participants during the initial scan, conflict resolution will use permissions from the

designated master host as the winner. If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

5. Click **Next**.

   The Email Alerts page is displayed.

Step 9: Email Alerts

This step is optional.

An email alert notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in Preferences and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply.

To create a new alert, see Email Alerts in the Preferences section.

To apply an existing email alert to the job.

1. Click the **Select** button.

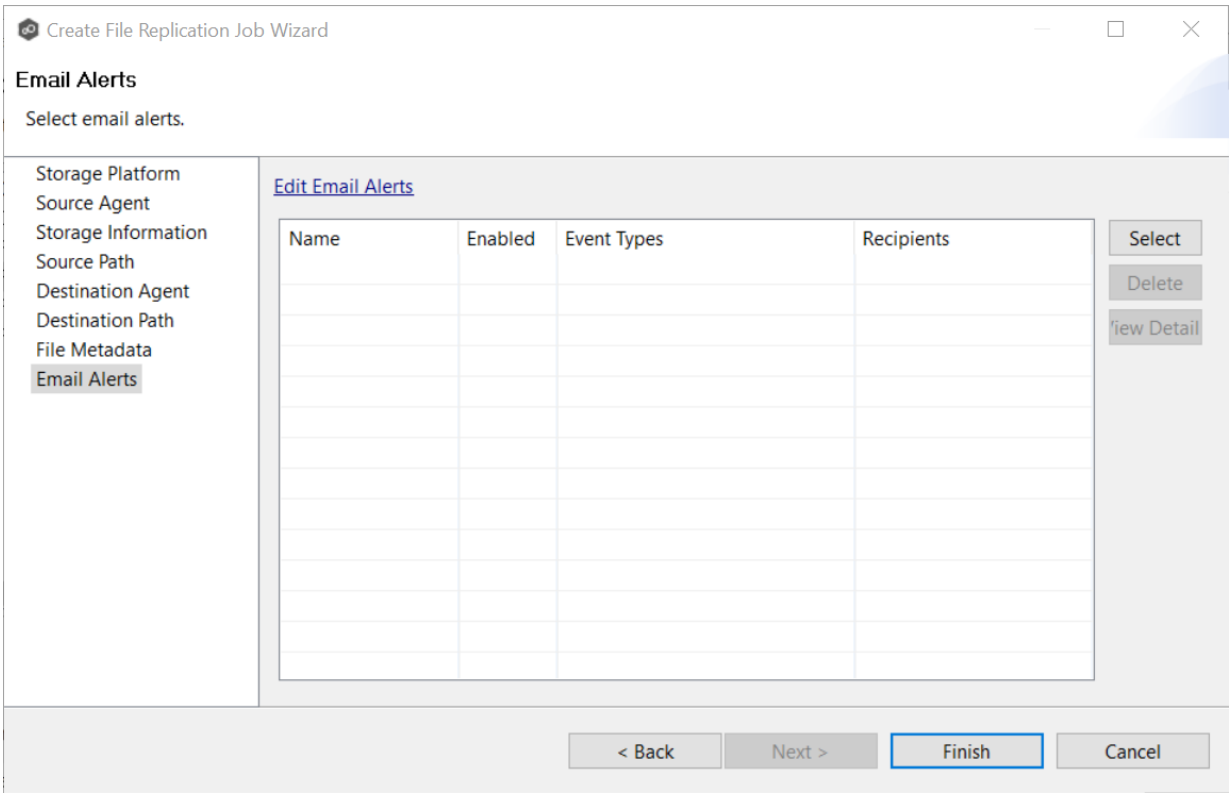

The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.

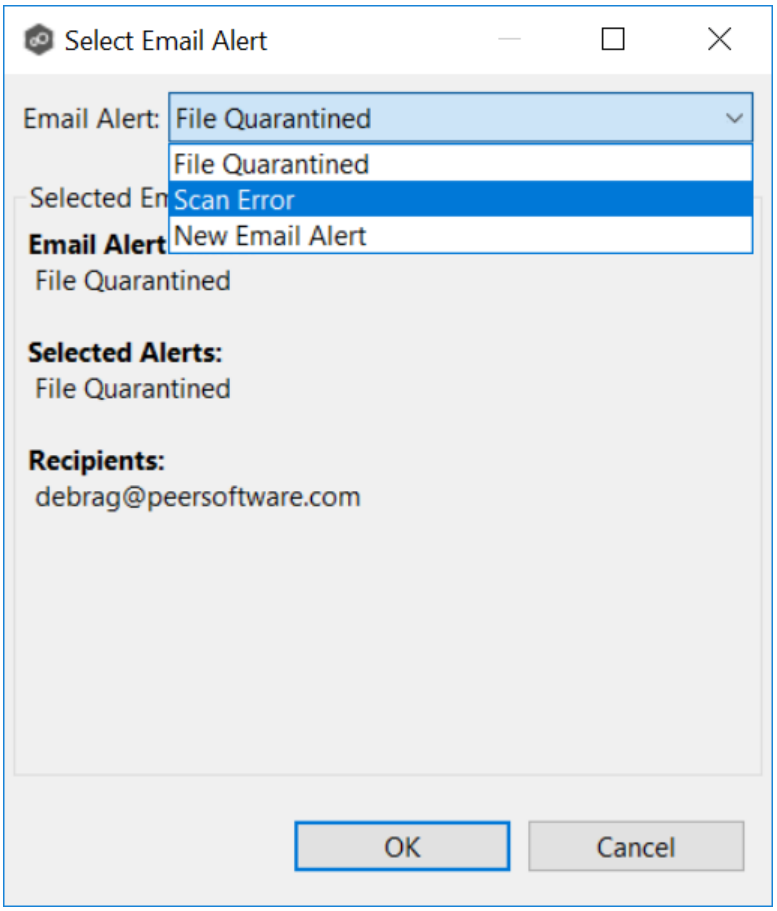

3. Click **OK**.

   The alert is listed on the **Email Alerts** page.

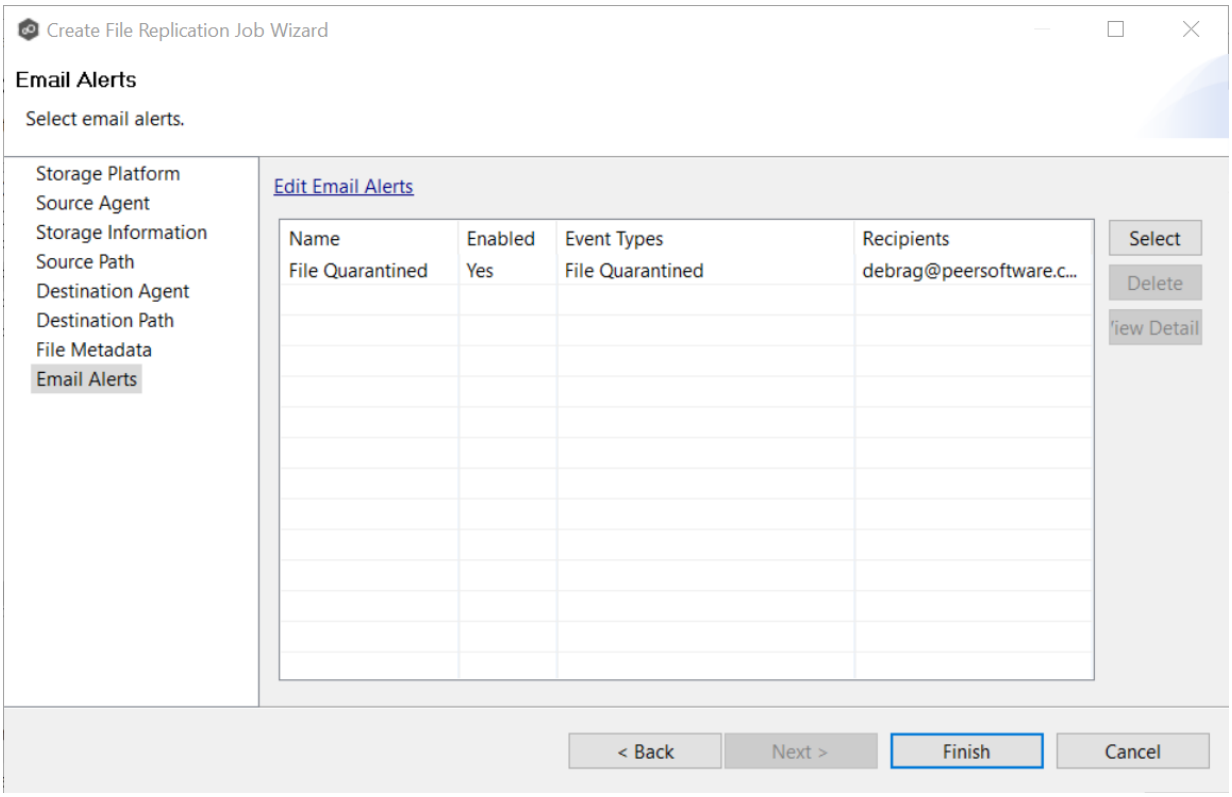

4. (Optional) Repeat steps 1-3 to apply additional alerts.

5. Continue to Step 10: Save Job.

Step 10: Save Job

Now that you have completed the first nine steps of the wizard, you are ready to save the job configuration.

1. If you are satisfied with your job configuration, click **Finish** to save your job. Otherwise, click the **Back** button and make any necessary changes.

   Congratulations! You have created a File Replication job. It is now listed in the **Jobs** view under **File Replication** and a job run-time view appears in the **Runtime Summaries** area. You can start the job from either place.

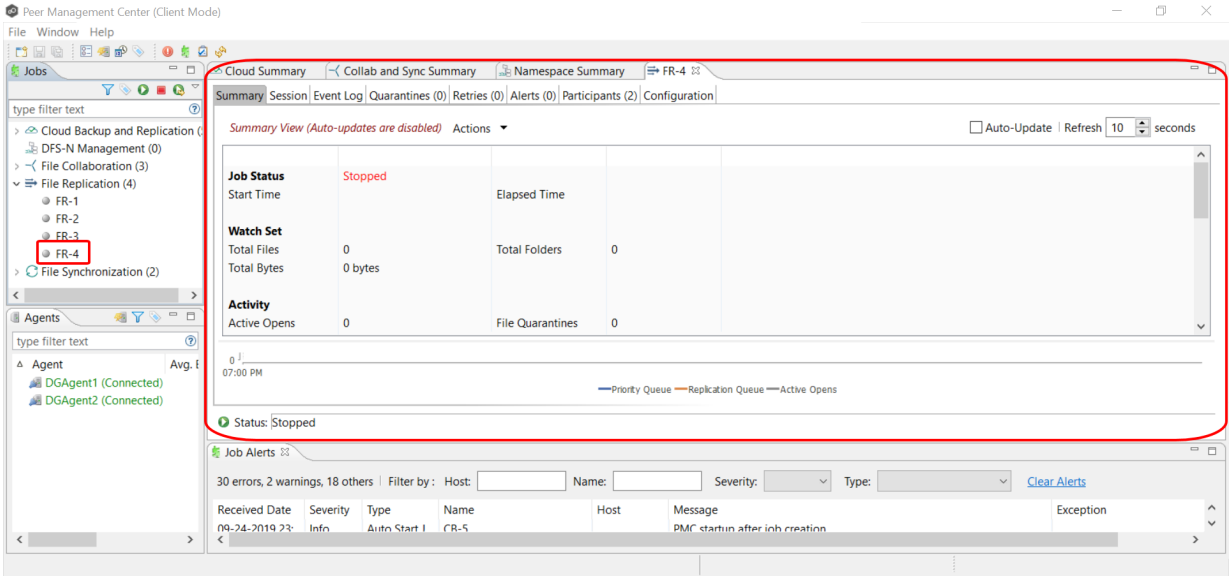

File Synchronization Jobs

This section provides information about creating a File Synchronization job:

- Overview
- Before You Create Your First File Synchronization Job
- Creating a File Synchronization Job
- Editing a File Synchronization Job
- Running and Managing a File Synchronization Job

Overview

A File Synchronization job provides real-time, multi-directional synchronization between various storage platforms and across locations. It is designed to handle non-collaborative workloads where files still need to be kept in-sync at multiple locations in real-time without locking. This job type is specifically optimized for use with user home directories and profiles.

Before You Create Your First File Synchronization Job

We strongly recommend that you configure global settings such as SMTP configuration and email alerts before configuring your first File Synchronization job. See Preferences for details on what and how to configure these settings.

Creating a File Synchronization Job

The **Create Job Wizard** walks you through the process of creating a File Synchronization job:

Step 1: Job Type and Name

Step 2: Participants

Step 3: File Metadata

Step 4: Email Alerts

Step 5: Save Job

Step 1: Job Type and Name

1. Open Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **Create New Job** wizard displays a list of job types you can create.
3. Click **File Synchronization**, and then click **Create**.

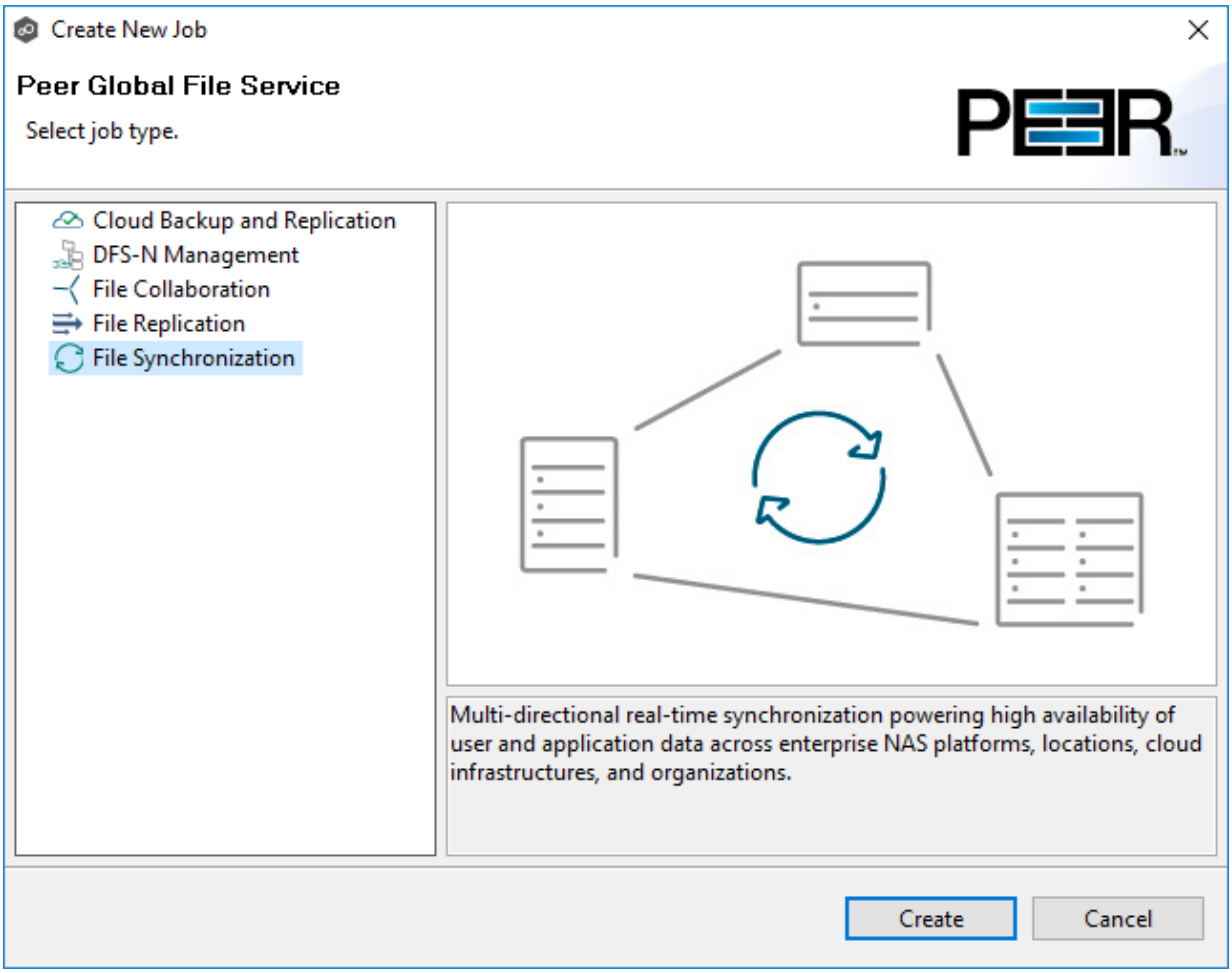

4. Enter a name for the job in the dialog that appears.

   The job name must be unique.

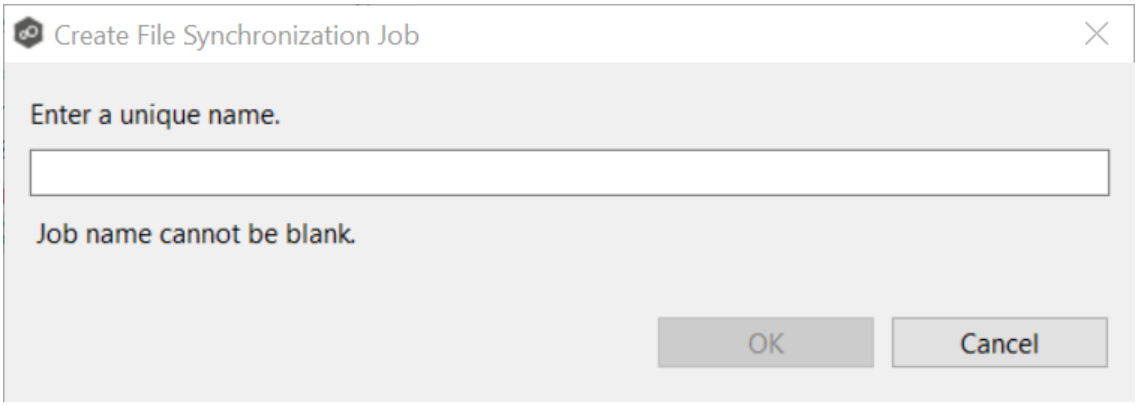

5. Click **OK**.

   The Participants page is displayed.

Step 2: Participants

After selecting the job type and naming the job, the **Participants** page is displayed. It contains a table that will display the job participants once you have added them. A File Synchronization job must have two or more participants. A **participant** consists of an Agent and the volume/share/folder to be replicated. A File Synchronization job synchronizes the files of participants in real-time.

1. Click the **Add** button to start the process of adding a participant.

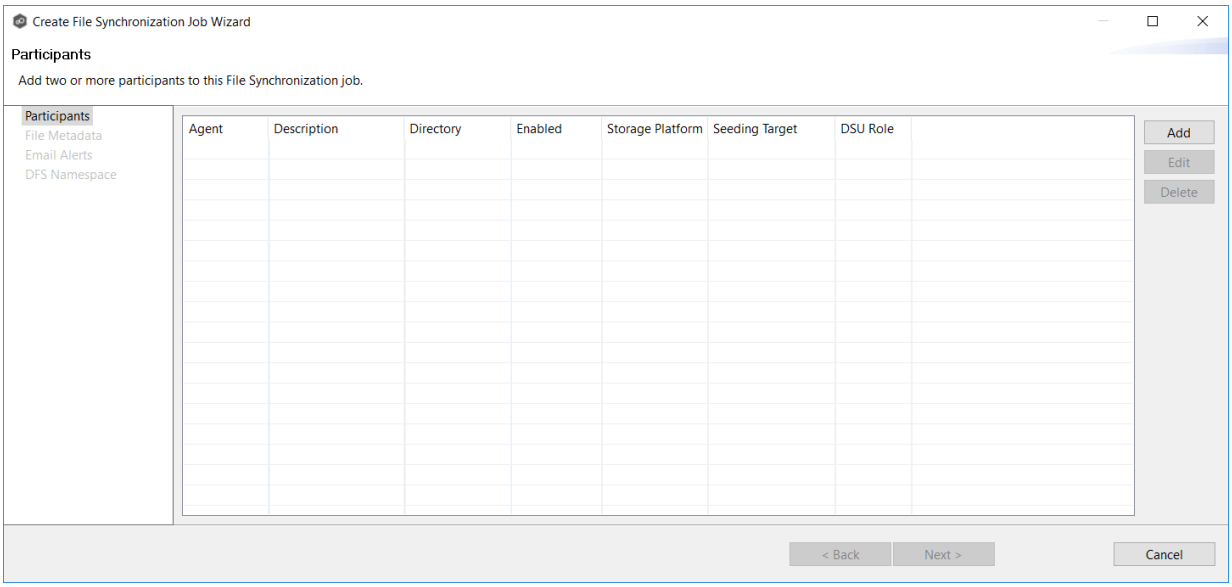

The **Add New Participant** wizard opens; it walks you through the steps for adding a participant:

a. Selecting a Management Agent, which is the Agent that will manage the storage device that hosts data you want to replicate.

b. Selecting the type of storage platform that hosts data you want to replicate.

c. Entering the credentials needed to access a specific storage device and providing other storage information.

d. Entering the path to the watch set (the data that you want to replicate) and selecting whether participant will be a seeding target.

e. (Optional) Enabling Dynamic Storage Utilization for the participant.

Once you have added a participant, it is listed in the **Participants** table.

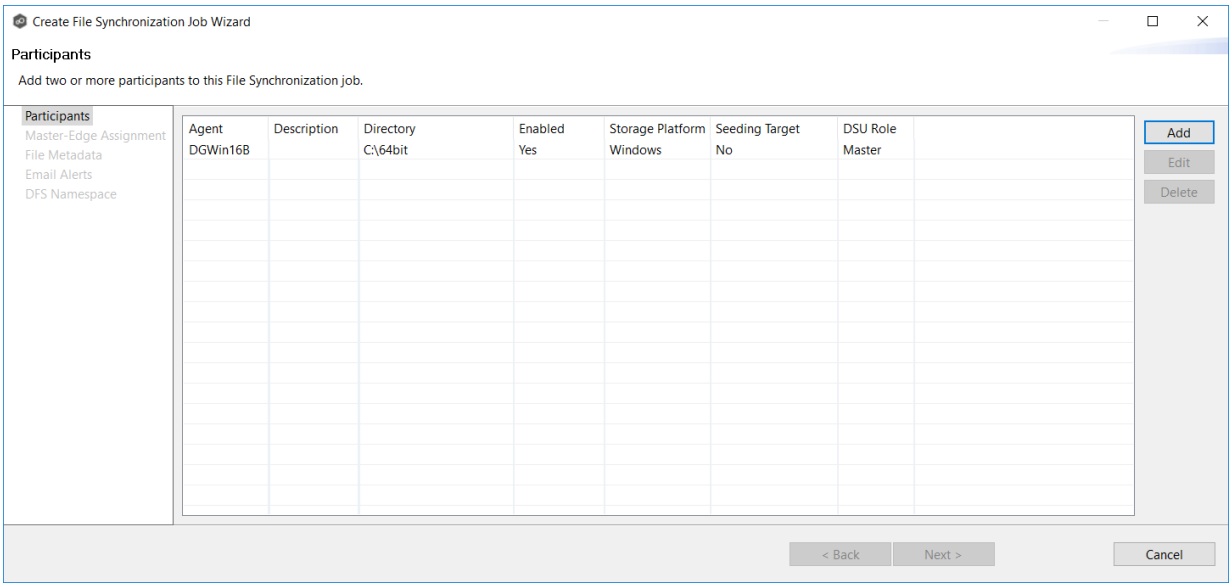

2. To add more participants, click **Add** and repeat the steps for each participant you want to add to the job.

3. Once you have added all the participants, click **Next**.

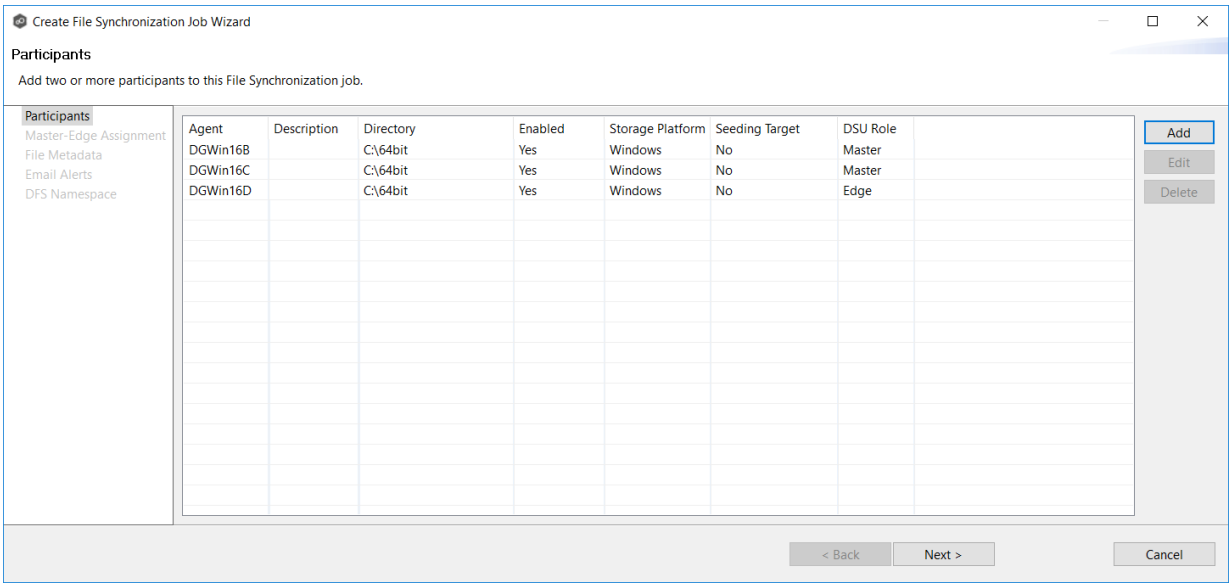

The page that appears next depends on options you selected when adding a participant:

- if you have enabled Dynamic Storage Utilization for a participant, the Dynamic Storage Utilization page appears.

- Otherwise, the File Metadata page appears.

The **Management Agent** page lists available Agents. You can have more than one Agent managing a storage device—however, the Agents must be managing different volumes/shares/folders on the storage device. For your File Synchronization job, you should select a Management Agent to manages the volumes/shares/folders you want to synchronize in this job.

1. Select an Agent to manage the storage device.

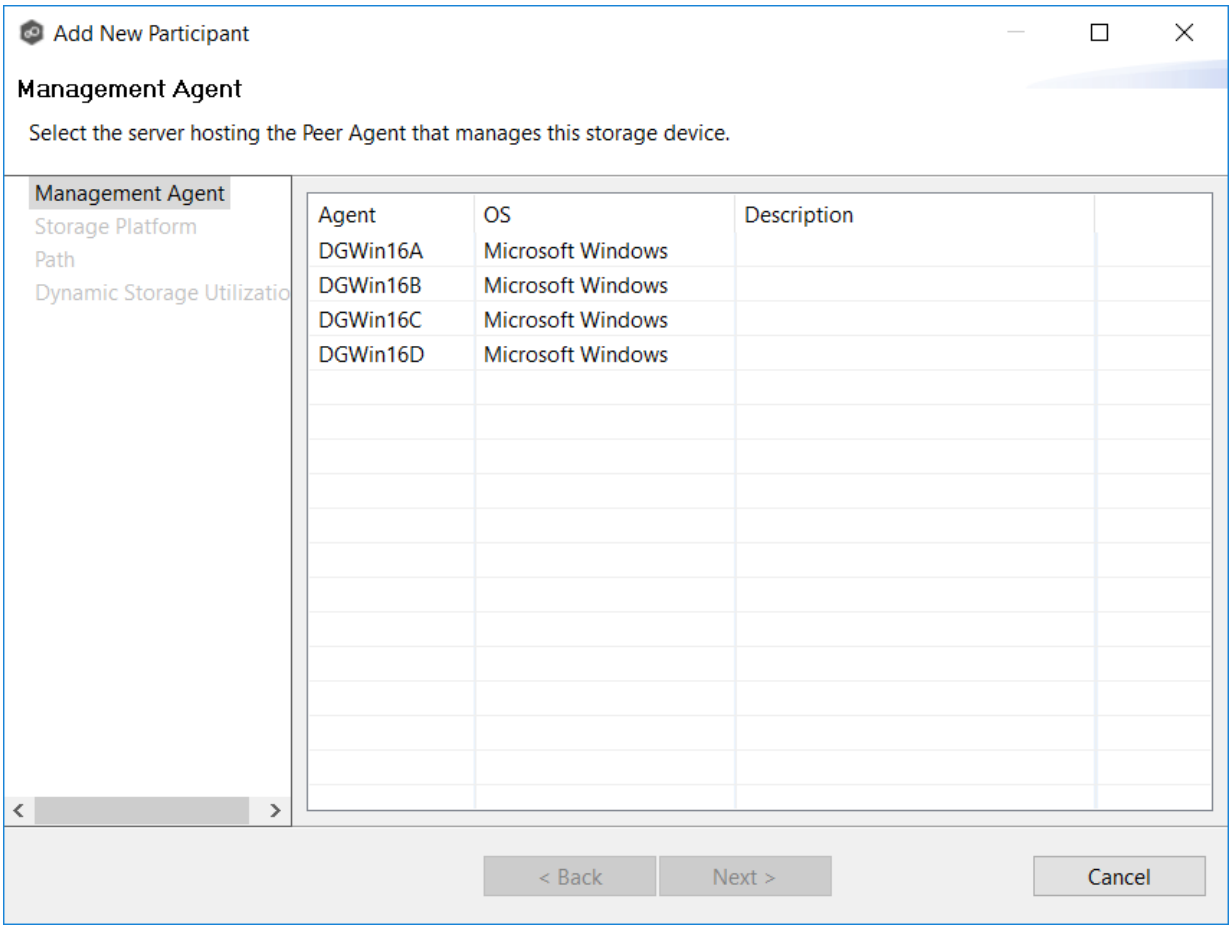

Tip: If the Agent you want is not listed, the Peer Agent Service may not be running on the server hosting the Agent. Try restarting the Peer Agent Service. If the service successfully connects to the Peer Management Broker, then the list is of available agents will be updated with that Agent.

2. Click **Next**.

The Storage Platform page is displayed.

The **Storage Platform** page lists the types of storage platforms that File Synchronization supports.

1. Select the type of storage platform that hosts the data you want to synchronize.

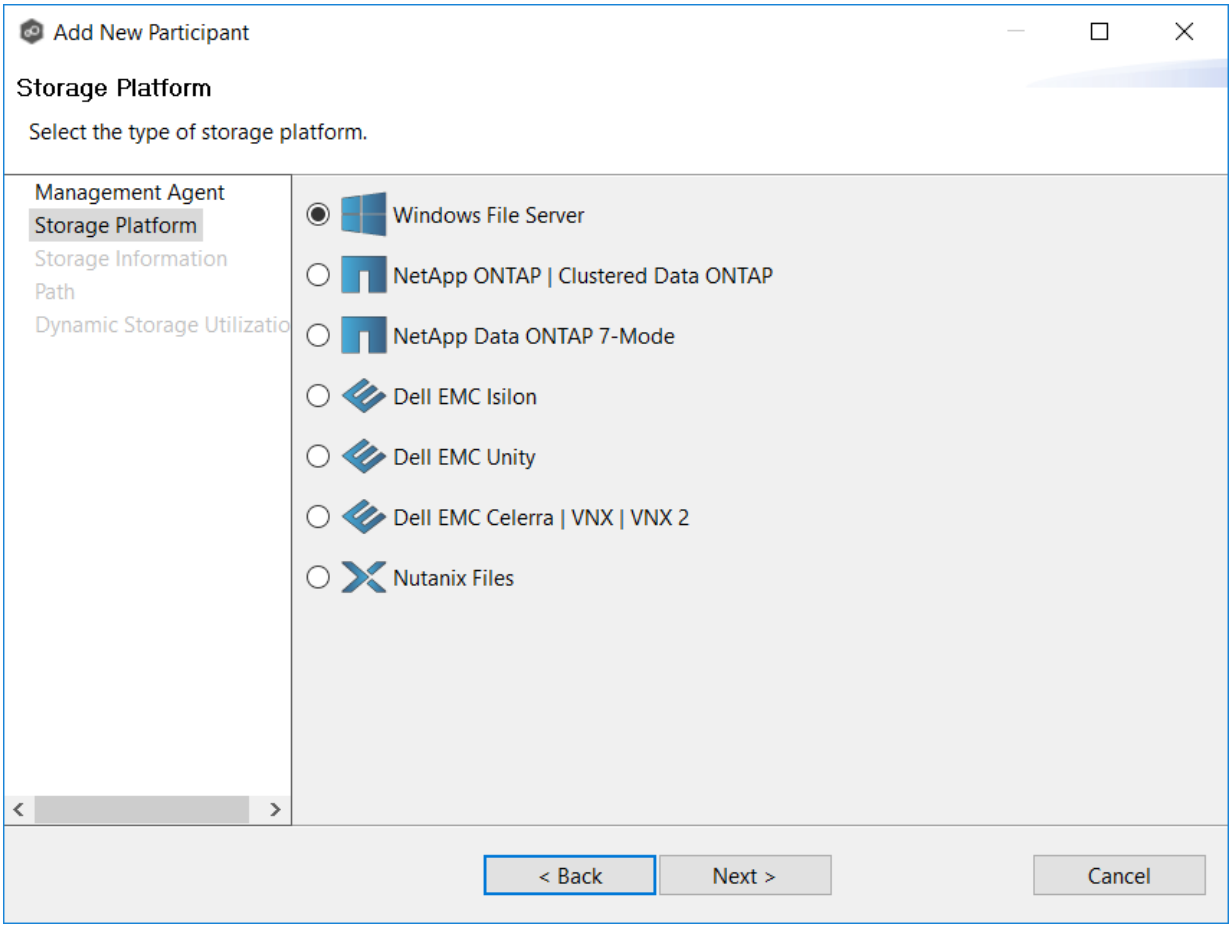

2. Click **Next**.

The Storage Information page is displayed.

On the **Storage Information** page, you will select a specific storage device containing data that you want to synchronize and enter other information about the storage device.

The contents of the **Storage Information** page varies, depending on your selection in the Storage Platform page:

- If you selected **Windows File Server**, you are prompted for configuration information relating to Windows File Server. Continue with the Windows File Server page.
- If you selected any other type of storage device other than Windows File Server, the **Storage Information** page requests the credentials necessary to connect to that storage device. Continue with the steps below.

Storage Platforms Other than Windows File Server

For storage platforms other than Windows File Server:

1. Select **New Credentials** or **Existing Credentials**.
2. If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the Path page.

   If you selected **New Credentials**, enter the credentials for connecting to the storage device. The information you are prompted to enter varies, depending on the type of storage platform:

   NetApp ONTAP | Clustered Data ONTAP

   NetApp Data ONTAP 7-Mode

   Dell EMC Isilon

   Dell EMC Unity

   Dell EMC Celerra | VNX | VNX 2

   Nutanix Files
3. Click **Validate** to test the credentials.

   After the credentials are validated, a success message appears.
4. Click **Next**.

   The Path page is displayed.

Windows File Server

1. Select the **Detector Type**.

   - Select **Windows Driver** for more robust logging and better performance (Recommended).
   - Select **Windows** if suggested by Peer Technical Support.

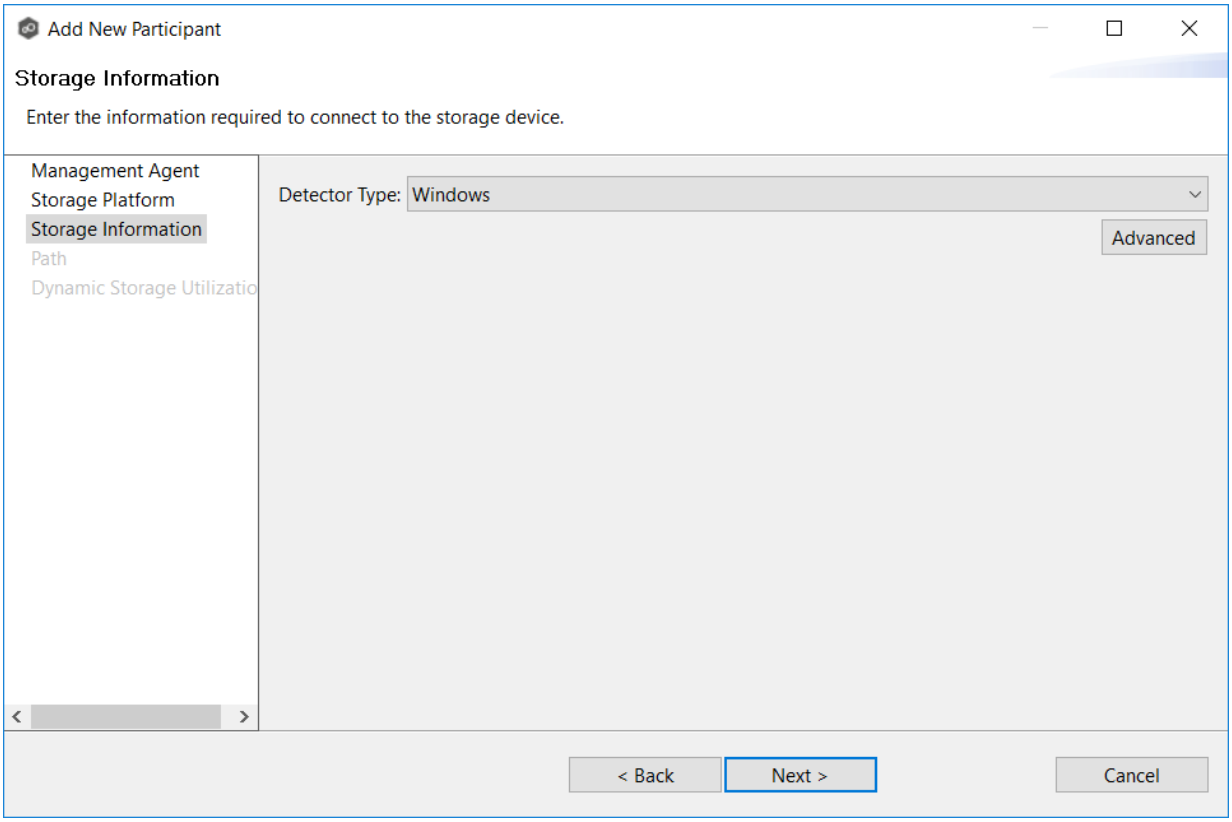

2. Click **Next**.

1. Modify the options as desired.

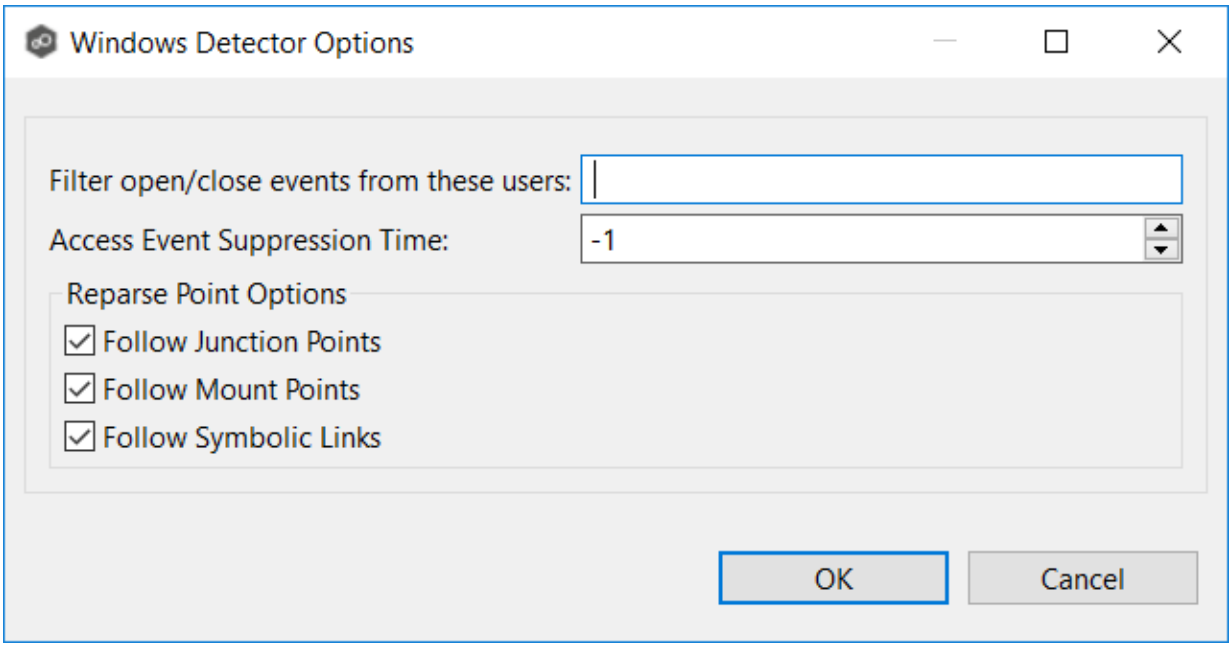

Option	Description
Filter open/close events from these users	A comma-separated list of user account names from which all opens and closes will be ignored. Ideal for filtering out events from backup and/or archival services by filtering on the username under which a backup and/or archival service is running.
Access Event Suppression Time	Represents number of seconds an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time an open event will be delayed based on the load of the system.
Follow Junction Points	Enables junction point support for the selected Windows File Server.
Follow Mount Points	Enables mount point support for the selected Windows File Server.
Follow Symbolic Links	Enables symbolic link support for the selected Windows File Server.

For more information about junction points or symbolic links, contact <% SUPPORT_EMAIL%

2. Click **OK**.

NetApp ONTAP | Clustered Data ONTAP

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine hosting the data to be synchronized.

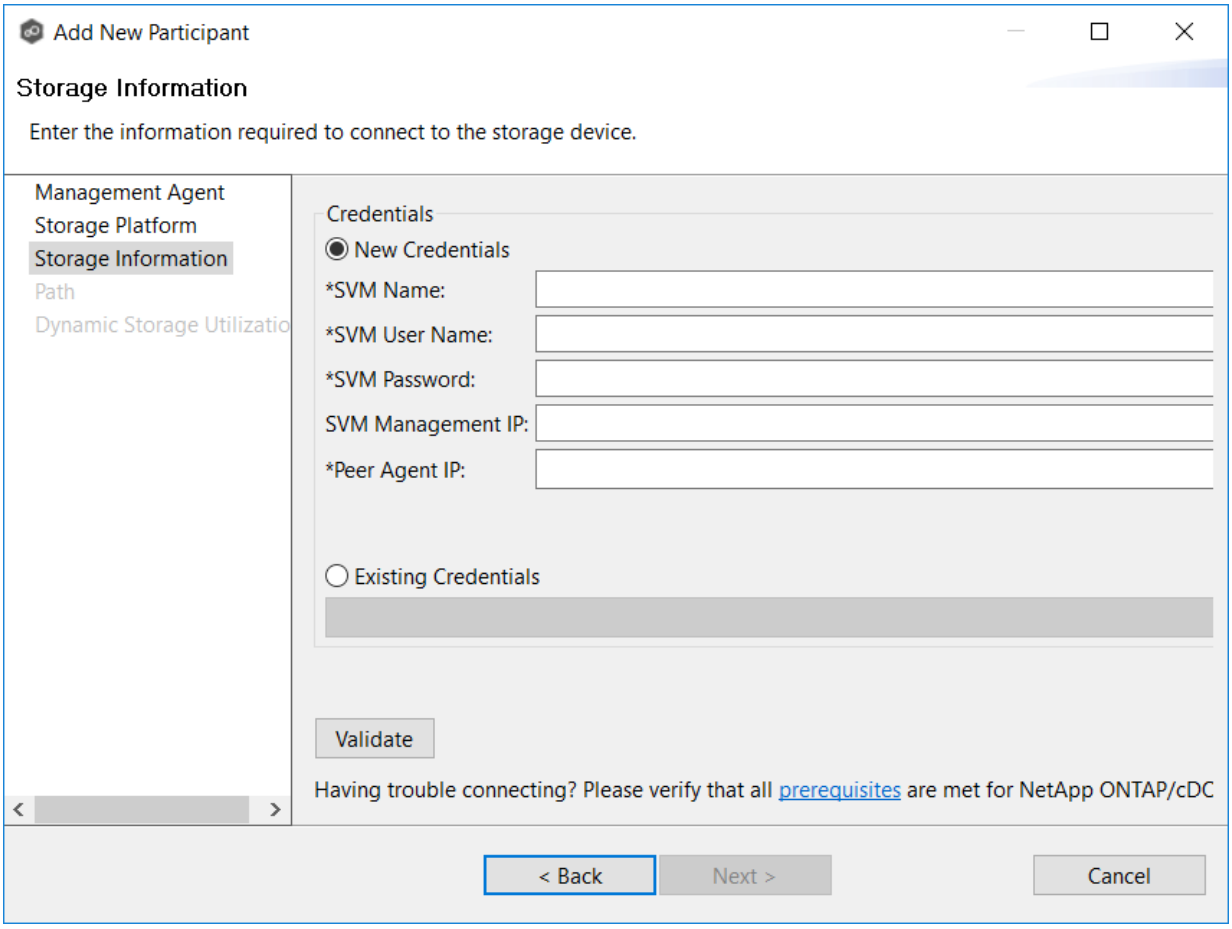

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
SVM Name	Enter the name of the Storage Virtual Machine hosting the data to be replicated.
SVM User Name	Enter the user name for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	(Optional) Enter the IP address used to access the management API of the NetApp Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already allow management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

NetApp Data ONTAP 7-Mode

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the NetApp 7-Mode filer or vFiler hosting the data to be synchronized.

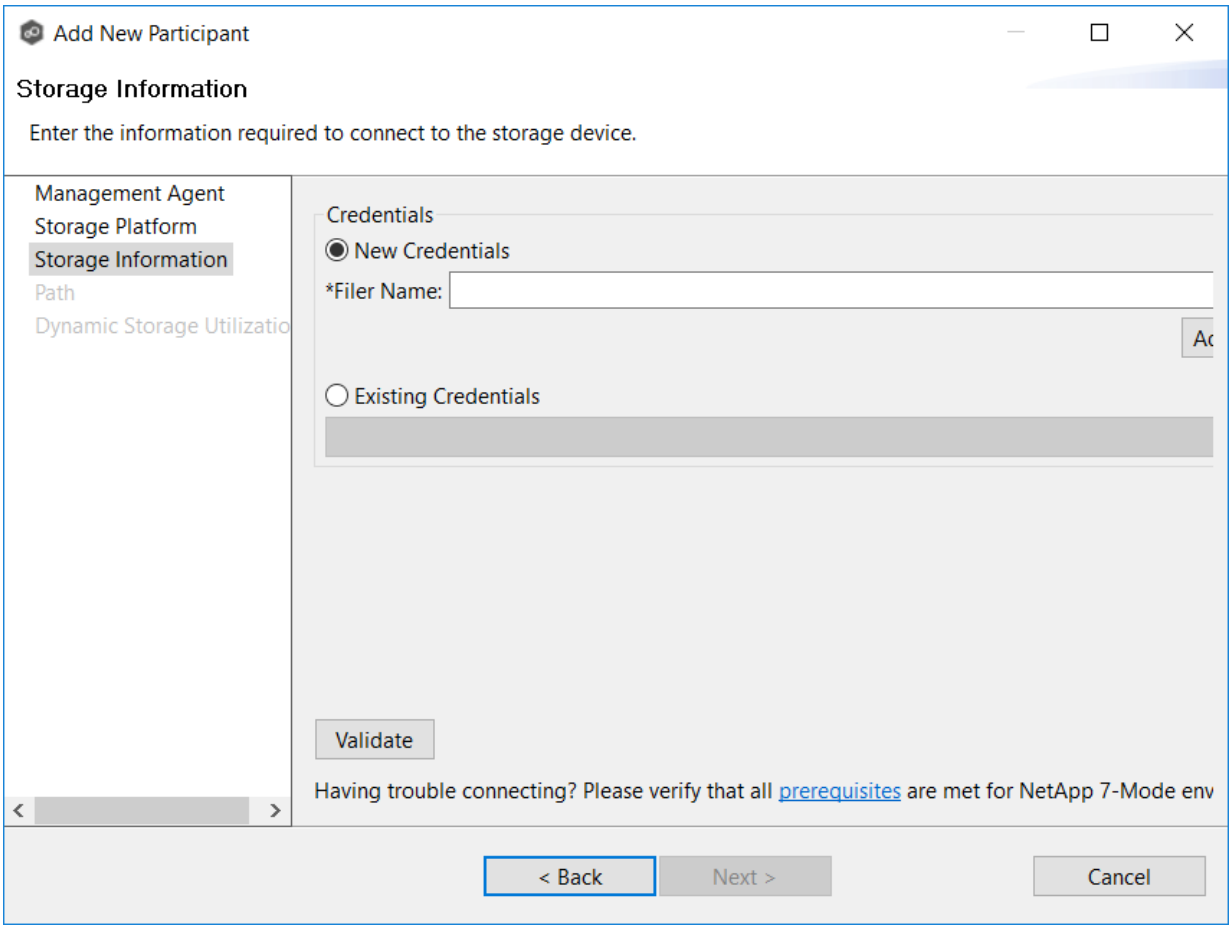

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
Filer Name	Enter the name of the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

The Path page is displayed.

Dell EMC Isilon

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect the EMC Isilon cluster hosting the data to be synchronized.

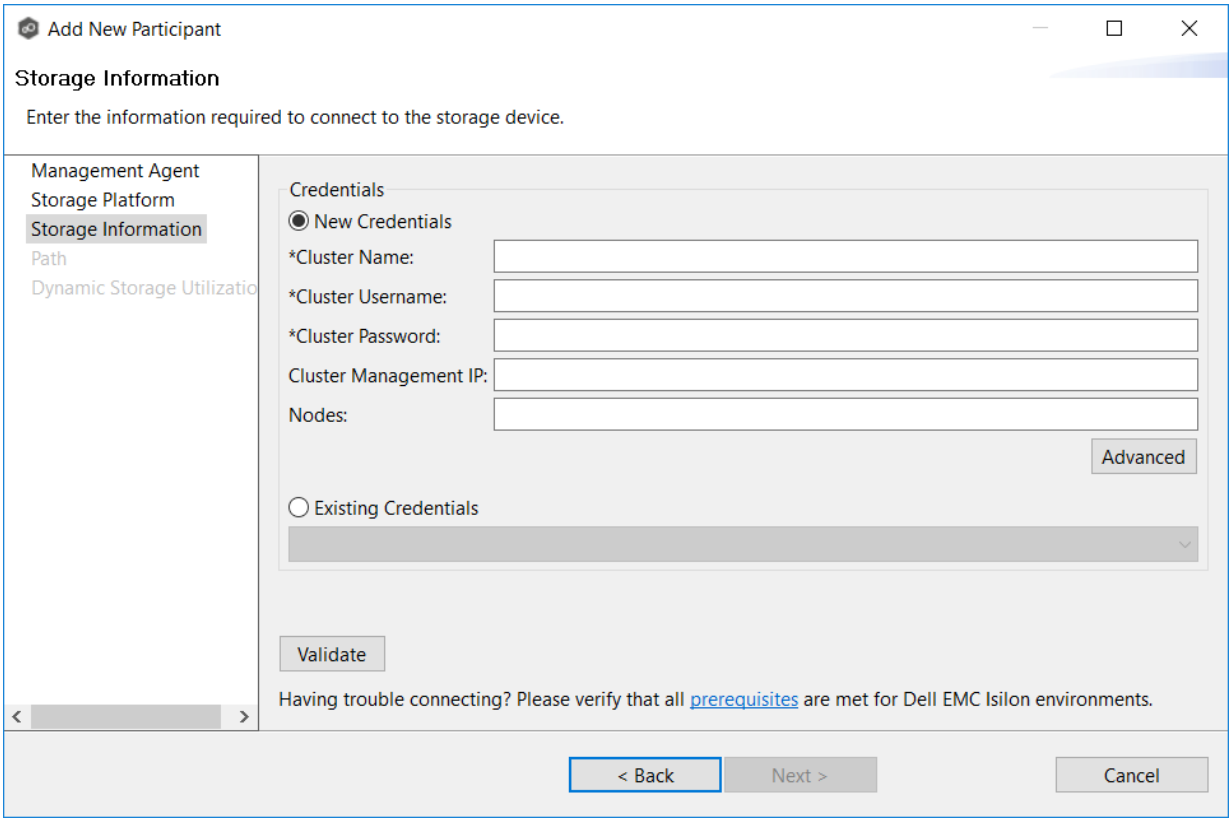

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
Cluster Name	Enter the name of the Dell EMC Isilon cluster hosting the data to be replicated.
Cluster Username	Enter the user name for the account managing the Dell EMC Isilon cluster.
Cluster Password	Enter the password for account managing the Dell EMC Isilon cluster.
Cluster Management IP	(Optional) Enter the IP address of the system used to manage the Dell EMC Isilon cluster. Required only if multiple Access Zones are in use on the cluster.
Nodes	(Optional) Enter one IP from each node in the Isilon cluster that the Agent can access to perform open file lookups. Use commas to separate nodes.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

Dell EMC Unity

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the CIFS Server hosting the data to be synchronized.

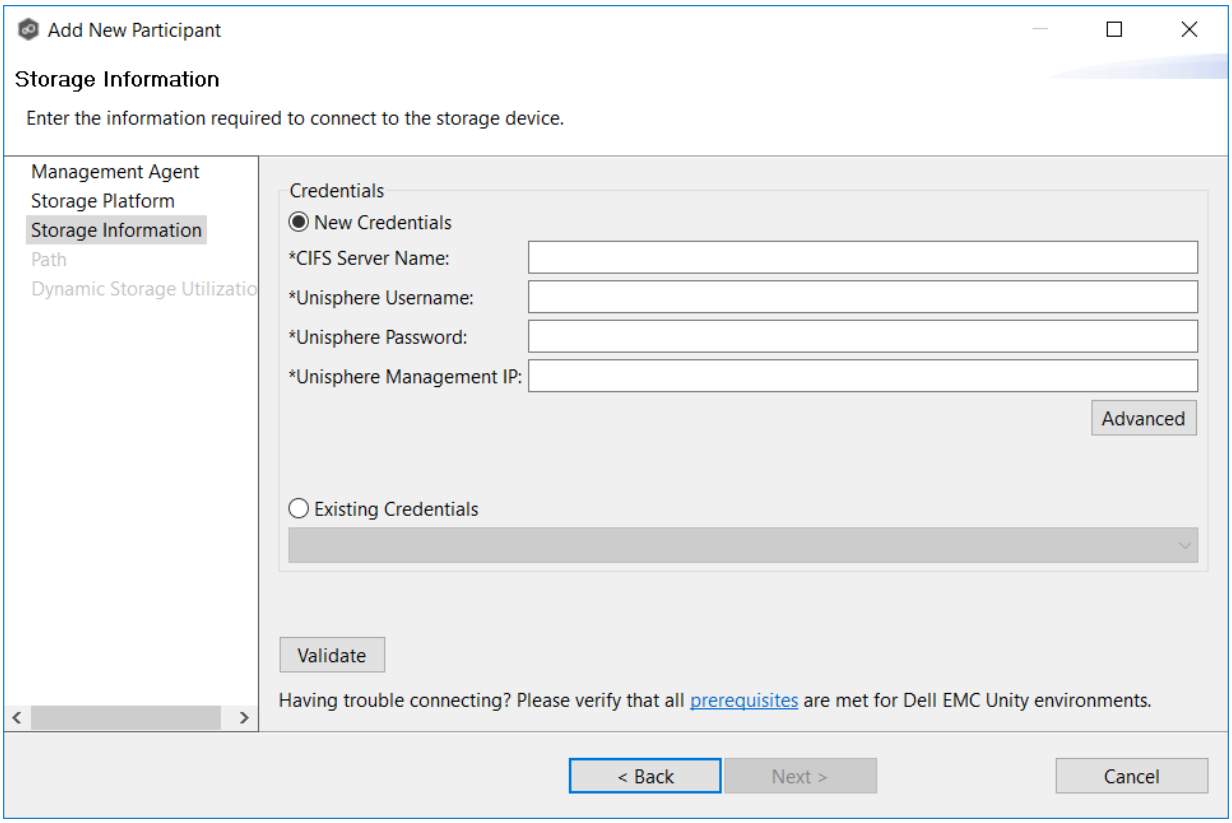

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
CIFS Server Name	Enter the name of the NAS server hosting the data to be replicated.
Unisphere Username	Enter the user name for the Unisphere account managing the Unity storage device.
Unisphere Password	Enter the password for the Unisphere account managing the Unity storage device.
Unisphere Management IP	Enter the IP address of the Unisphere system used to manage the Unity storage device. This should not point to the NAS server.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

Dell EMC Celerra | VNX | VNX 2

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the CIFS Server hosting the data to be synchronized.

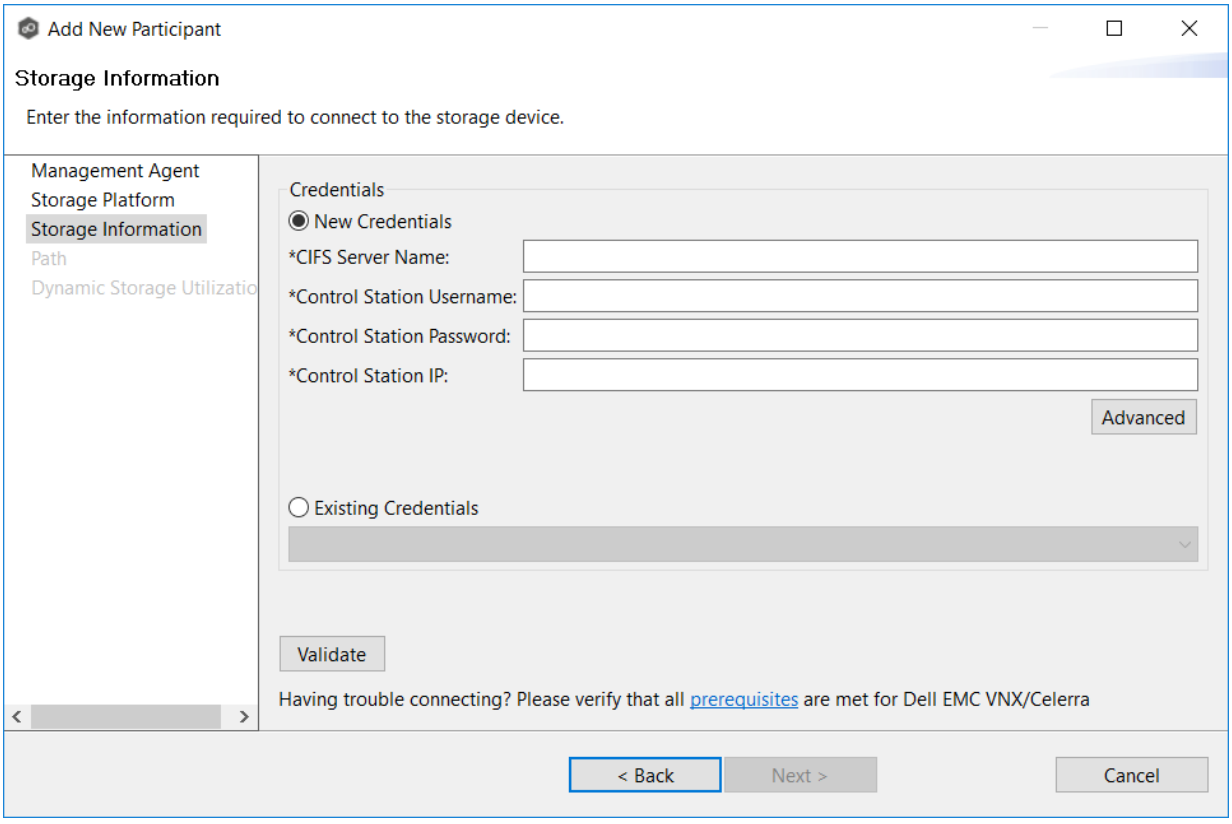

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
CIFS Server Name	Enter the name of the CIFS Server hosting the data to be replicated.
Control Station Username	Enter the user name for the Control Station account managing the Celerra/VNX storage device.
Control Station Password	Enter the password for the Control Station account managing the Celerra/VNX storage device.
Control Station IP	Enter the IP address of the Control Station system used to manage the Celerra/VNX storage device. This should not point to the CIFS Server.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

Nutanix Files

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Nutanix Files cluster hosting the data to be synchronized.

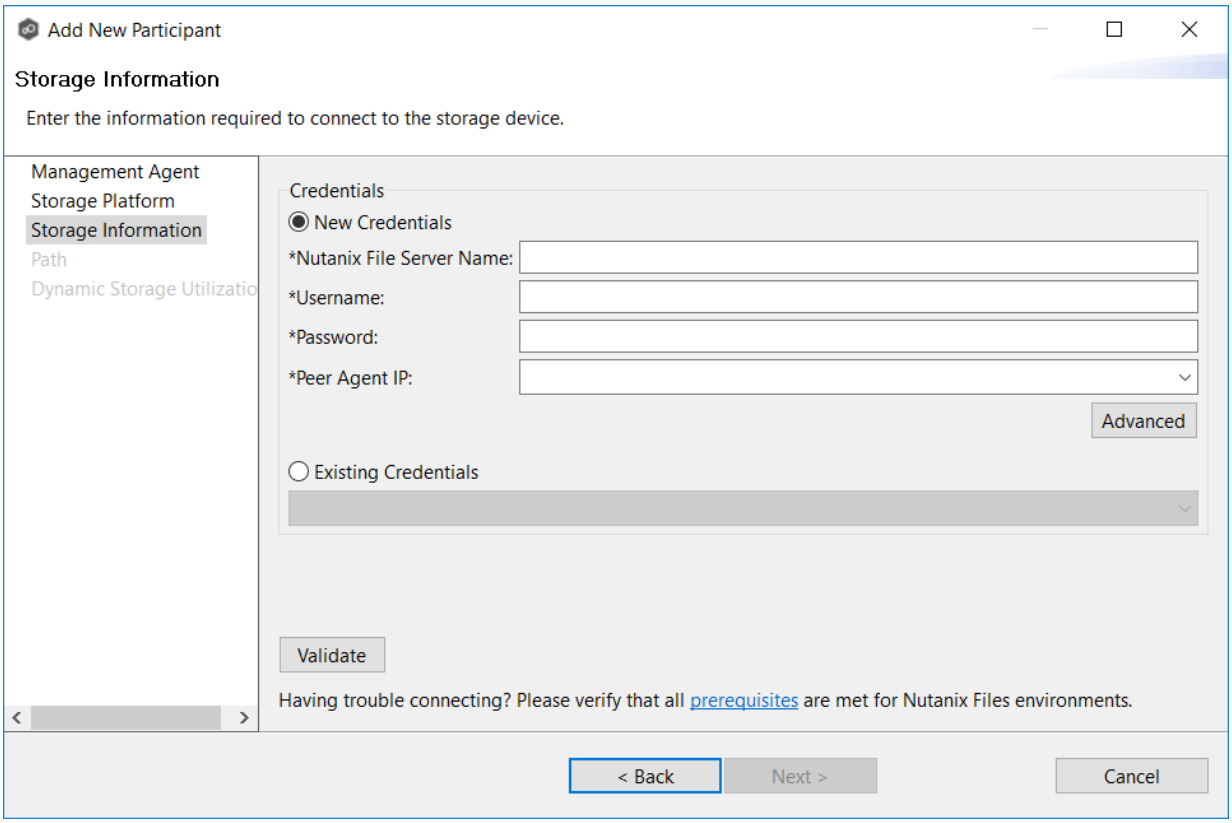

2. If you selected **Existing Credentials**, select the credentials, and then click **Next**.

   If you selected **New Credentials**, supply the required information.

Field	Description
Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to this IP address. If the IP address you want does not appear, manually enter the address. This should not point to the Files cluster itself.

3. Click **Advanced** if you want to set advanced options.
4. Click **Validate**.
5. Click **Next**.

   The Path page is displayed.

The **Path** page is where you specify the path to the volume/share/folder you want to replicate. This volume/share/folder is referred to as the watch set. The watch set can contain a single volume/share/folder. If you want to replicate multiple volumes/shares/folders, you need to create a separate job for each one.

1. Browse to or enter the path to the watch set.

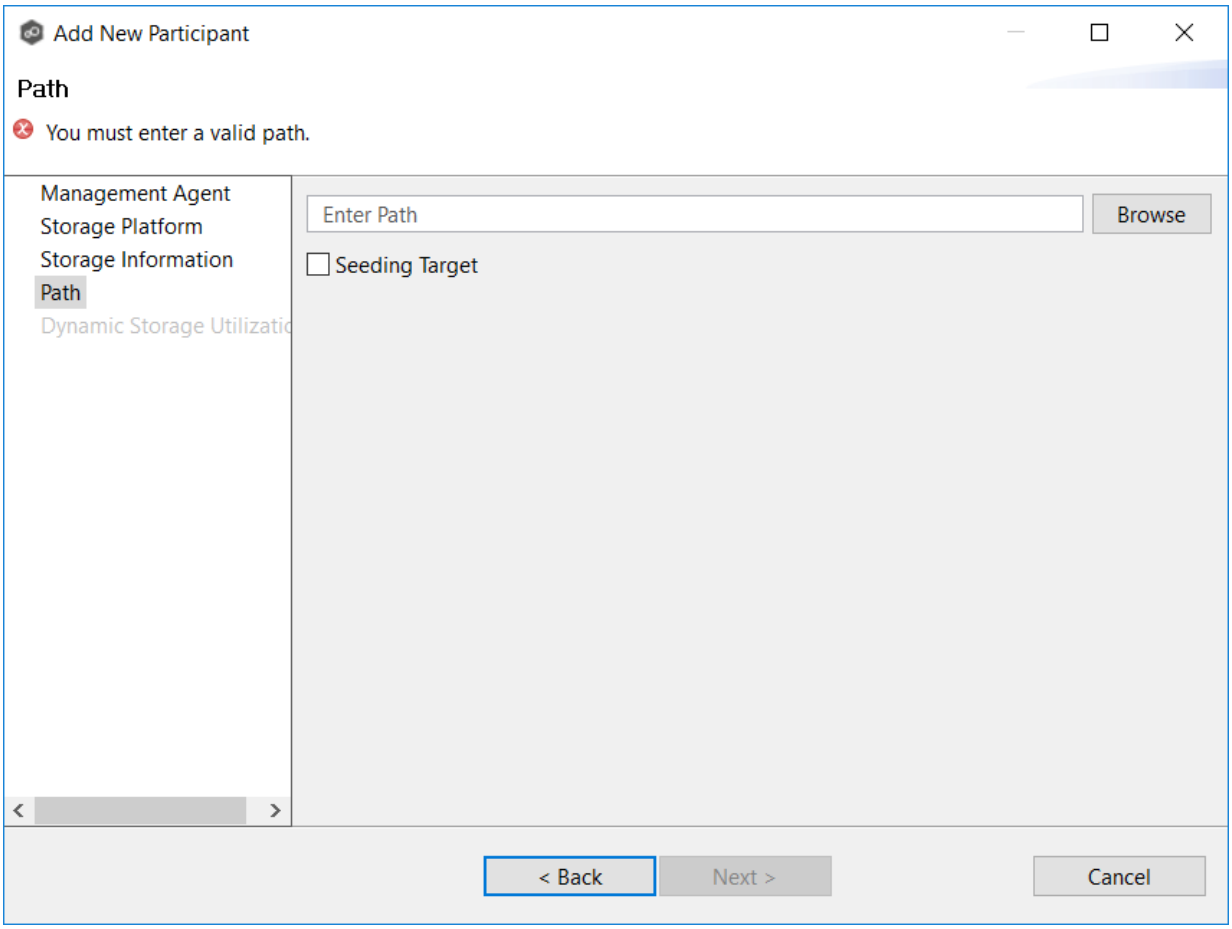

2. (Optional) Select the **Seeding Target** checkbox, and then click **OK** in the dialog that appears.

   If you select this option, a message describing seeding behavior is displayed. Multiple participants in a File Synchronization job can be set as smart data seeding targets; however, at least one participant should not be set as a smart data seeding target. This participant will be acting as the "master" source for the smart data seeding targets. For more information about smart data seeding, see Smart Data Seeding or contact support@peersoftware.com.

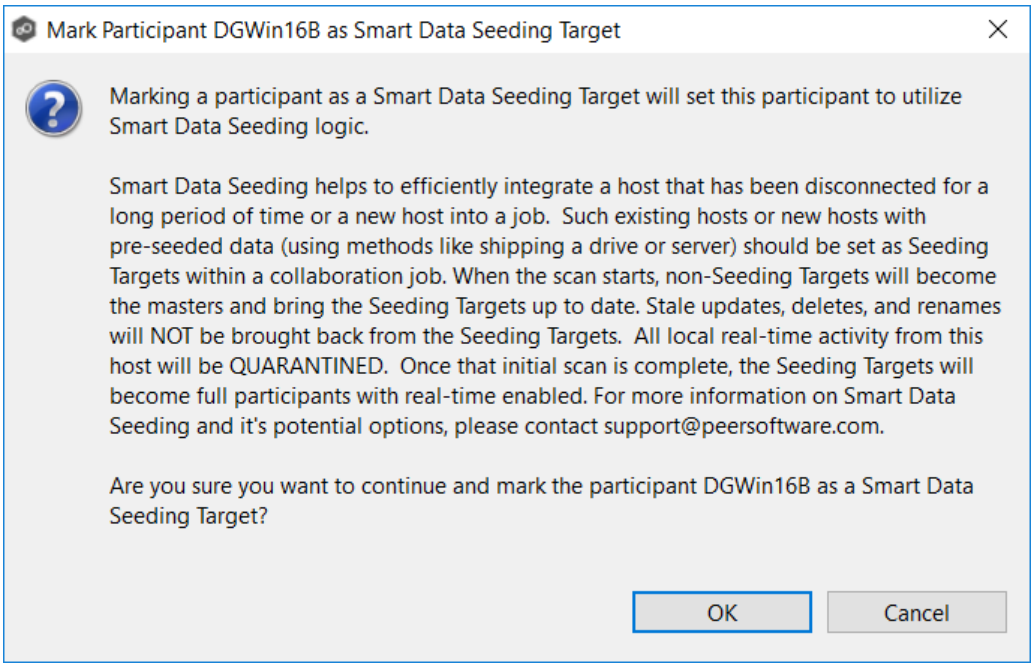

3. Click **Finish** to complete the wizard for this participant.

4. Return to Step 2: Participants to add more participants, if applicable. A File Synchronization job must have at least two participants. If you have added all the participants, continue with Step 3: File Metadata.

Dynamic Storage Utilization (DSU) is a method for conserving space on storage devices by caching files until needed. DSU saves space by stubbing files and rehydrating them as needed. DSU is optional; if you don't need to conserve space on the storage device managed by the Agent, then you do not need to select this option.

If you enable Dynamic Storage Utilization (DSU) for a participant, you must designate the participant as either a **master** or **edge** participant.

- **Master participant** - A master participant always has complete set of files for that job. None of the files are stubbed; they are stored physically on that device.

- **Edge participant** - A subset of the files stored on a master participant are physically stored on an edge participant; the rest of the files on an edge participant are stub files that take up minimal space. DSU allows users to seamlessly retrieve stubbed files directly from a master participant as needed; when retrieved, the local stub file is rehydrated so that the full file is stored locally on the edge participant.

A job can have master and edge participants, as well as participants that don't have either role. If you do not choose to enable DSU for a participant, it will always have a full set of files like a master participant but will not be used to serve file content to any edge participants.

Notes:

- A participant can be a master participant for some jobs and an edge participant for other jobs.
- A job needs at least one master participant that isn't a seeding target. If there is only one master participant for the job, it should not be a seeding target.

For more information about DSU, see Dynamic Storage Utilization in Advanced Topics.

1. Select the **Enable Dynamic Storage Utilization** checkbox if you want this participant to be able to use Dynamic Storage Utilization; otherwise, click **Finish**.

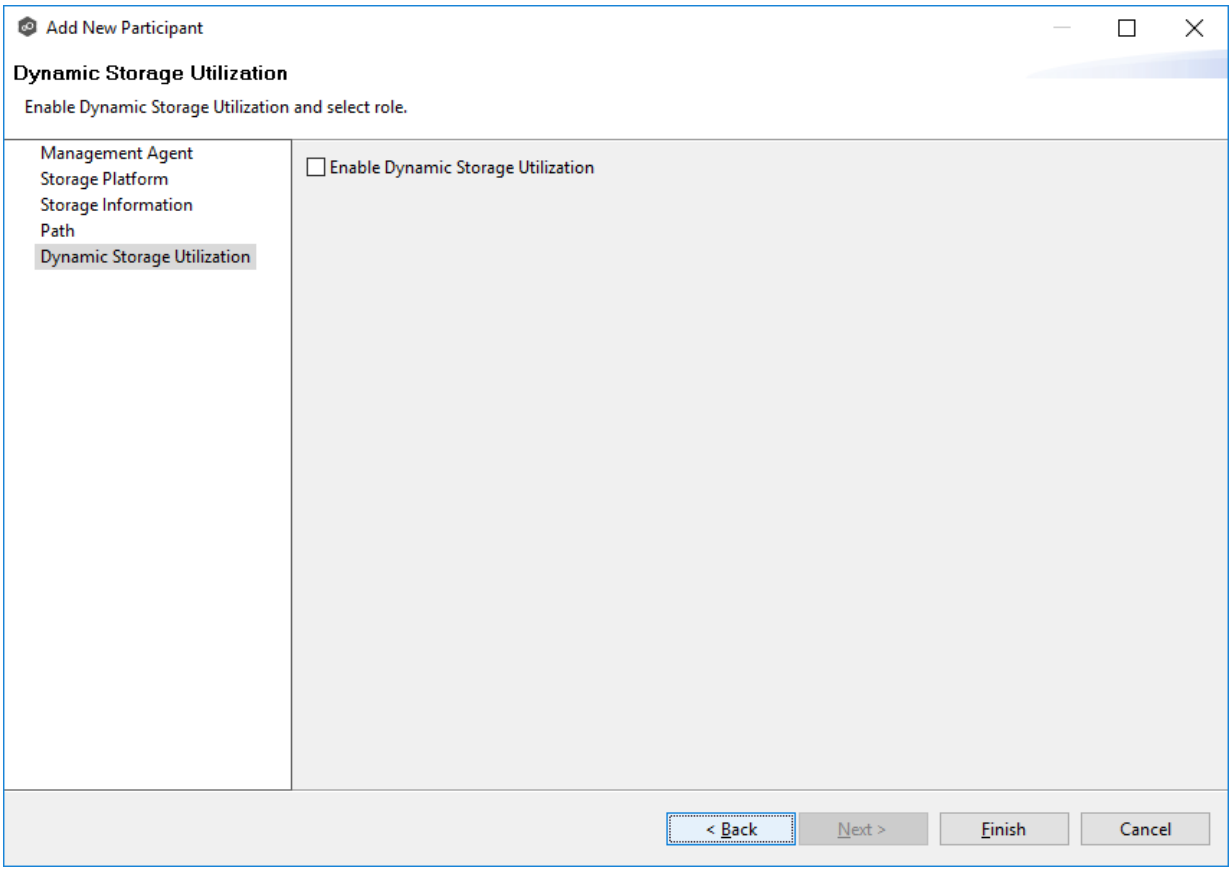

   If you enable DSU, the DSU role options are displayed; the **Master** role is selected by default.

2. Choose a DSU role for the participant:

- Choose **Master** if the storage device managed by the Agent will contain complete copies of all files for this job. Any type of storage platform can be a master participant.

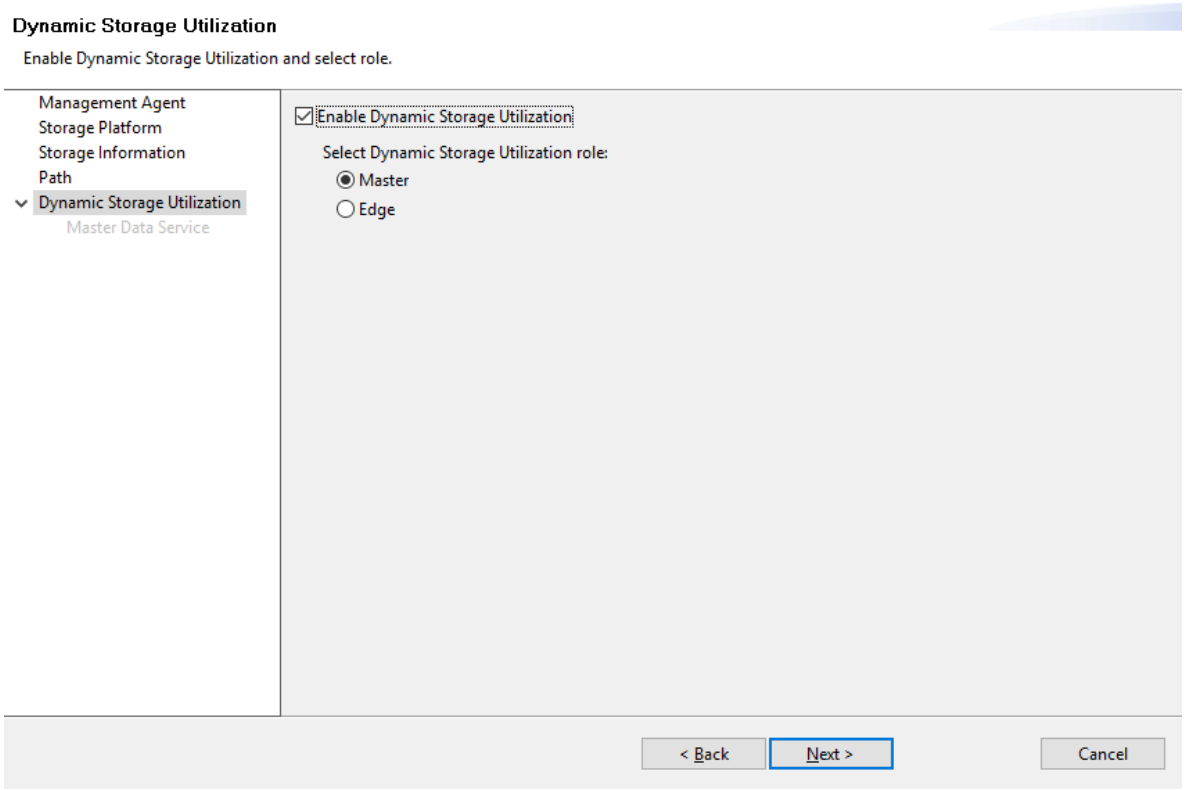

- Choose **Edge** if you want to conserve space on the storage device managed by the Agent. Only Windows File Servers can be an edge participant.

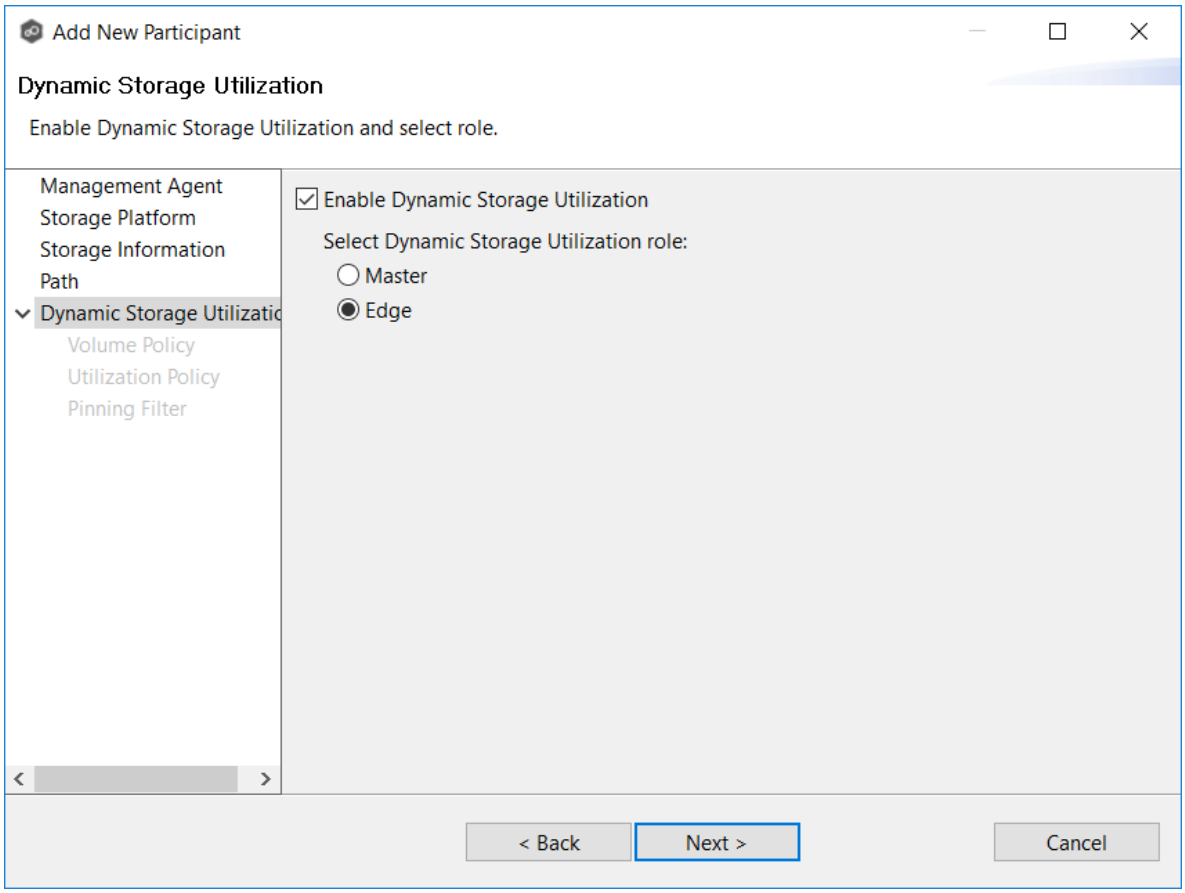

3. Click **Next**.

   - If you selected **Master**, continue with the Master Data Service page.
   - If you selected **Edge**, continue with the Volume Policy page.

Master Data Service

The **Master Data Service** page appears if you chose the master role for the participant. The Master Data Service handles requests from edge participants for files on a master participant. The Master Data Service is installed on the Agent server as part of the Peer Agent installation process.

The first two fields on this page are automatically populated:

- **Protocol**: This field lists the protocol that will be used to transfer file content between master participants and edge participants. HTTPS is currently the only available option as

it requires only a single open firewall port on each master participant and does a better job of handling latency over WAN-based connections.

- **Agent Name**: This field lists the name of the Management Agent that you selected at the beginning of Step 2.

1. (Optional) Enter a value for **Agent Alias**. The value can be a hostname, FDQN, or IP address.

   A value for this field is required only if the name of the Agent cannot be converted to an IP address via DNS. If an alias is entered, it will be used by the Edge Service on edge participants to connect with the Master Data Service. If no alias is entered, the name in the Agent Name will be used.

2. (Optional) Modify the port number that the Master Data Service will listen on for this master participant.

   A default value for the port number, 8446, is set when the Agent is installed. If you modify the port number, the Master Data Service is started with the new port number.

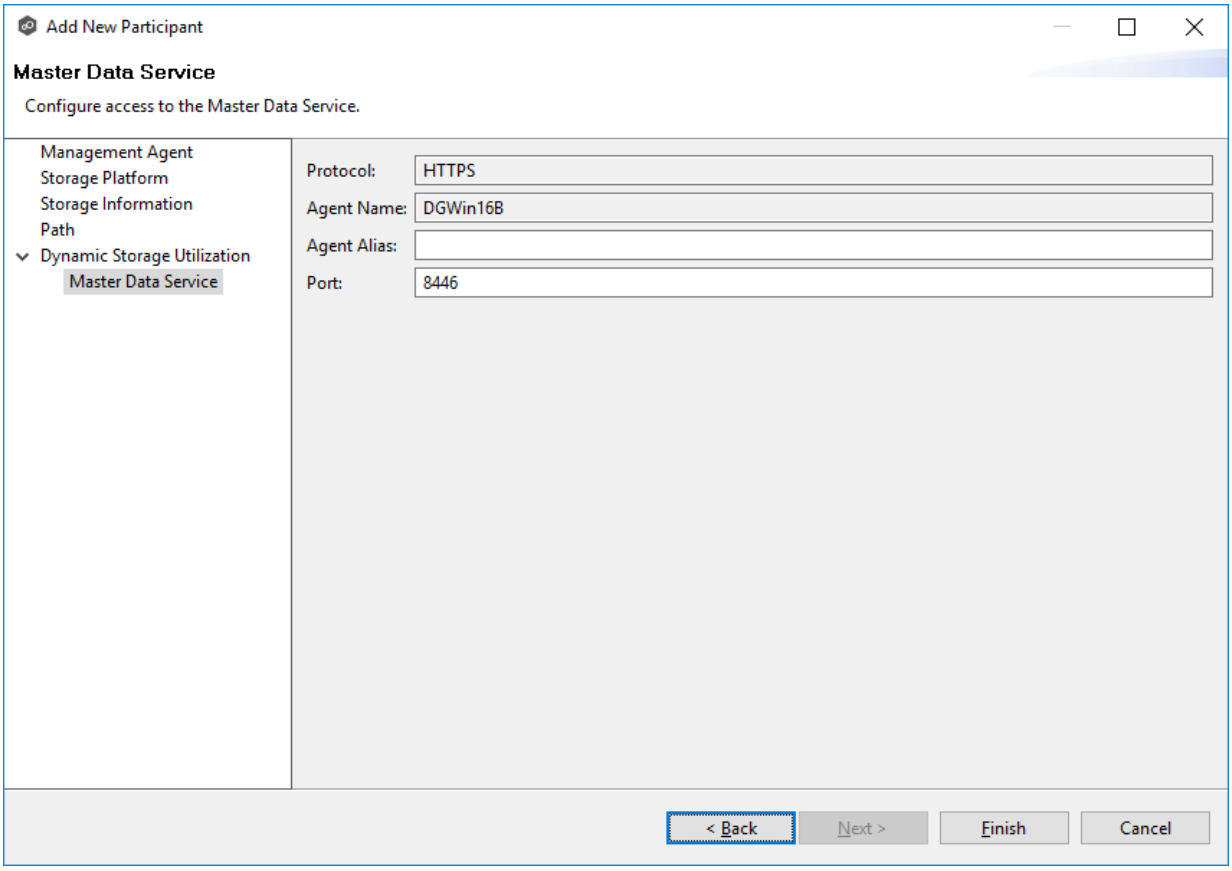

Note: If the Agent you selected is already being used as a master participant in another job utilizing DSU, then the existing Master Data Service parameters will be displayed. You can edit the values by clicking the **Edit Master Data Service** link. If you modify the

port number, the Master Data Service will be restarted and the new port number will take effect immediately. Any modifications you apply will be applied to every other job that use this Agent as a master participant.

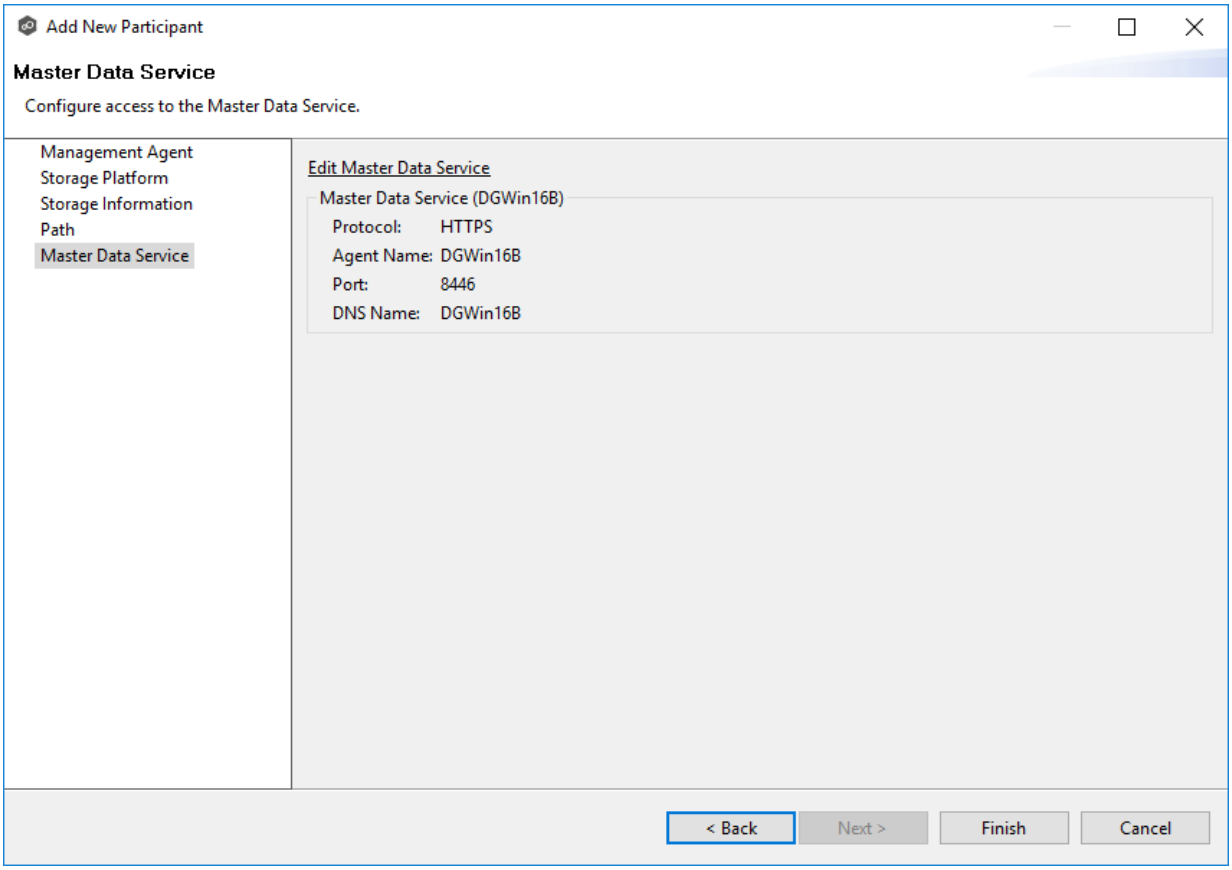

3. Click **Finish**.

The **Participants** page reappears. The participant is listed in the **Participants** table with the **Master** role.

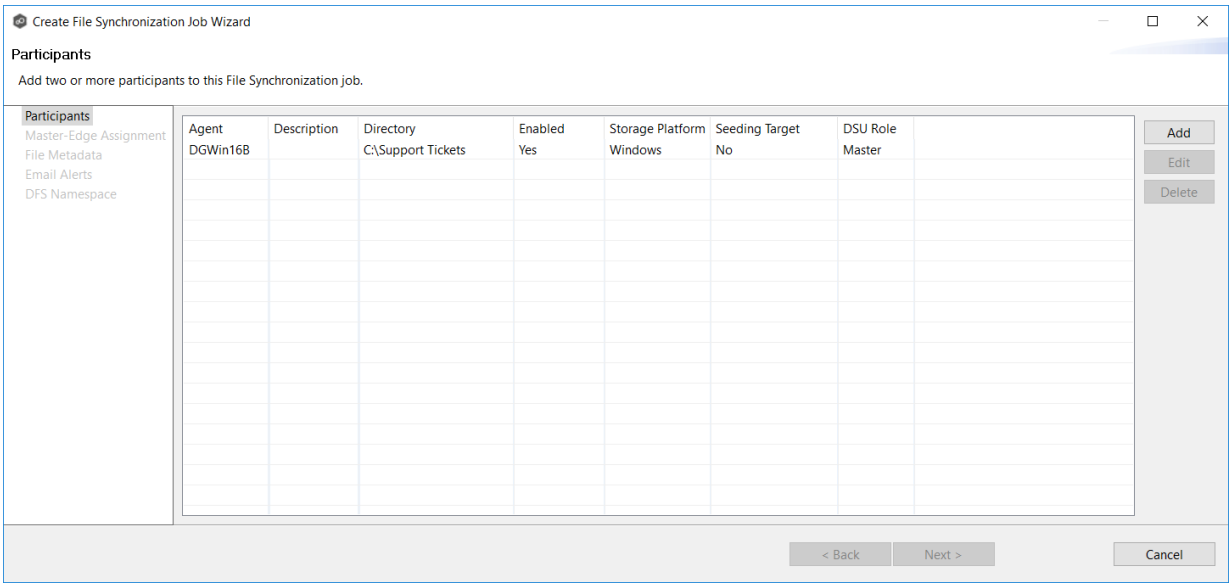

4. Continue adding more participants if applicable or continue with Step 4: Master-Edge Assignment.

Volume Policy

The **Volume Policy** page appears if you chose the edge role for the participant.

A volume policy is applied when a caching scan is run. The primary purpose of a **volume policy** is to specify how much space is available to DSU on a specific volume (or drive letter), i.e, to define the **cache size**. The cache size specifies the maximum amount of disk space you want to allocate to DSU for fully hydrated files on the volume specified by the path on the **Path** page. For example, if the participant is configured to monitor D:\Data, the volume policy for this participant would apply to the D volume.

The cache size can be specified as a percentage of the volume disk space or as a fixed size. For example, if an edge participant is configured to monitor a volume that has 1 TB of disk space, and you tell DSU to use 75% of that volume, then up to 750 GB of files could be locally available on the volume monitored by that edge participant. For optimal performance, we recommend that this cache be dedicated to DSU's use on this volume.

A volume policy applies to each job where the following three elements are true:

- DSU is enabled for the job.
- The participant is an edge participant.
- The paths specified for each job share the same volume.

To create a volume policy:

1. In the **Cache Size** section, choose an option for setting the cache size:

   - Use up to X % of this volume
   - Use up to X size of this volume

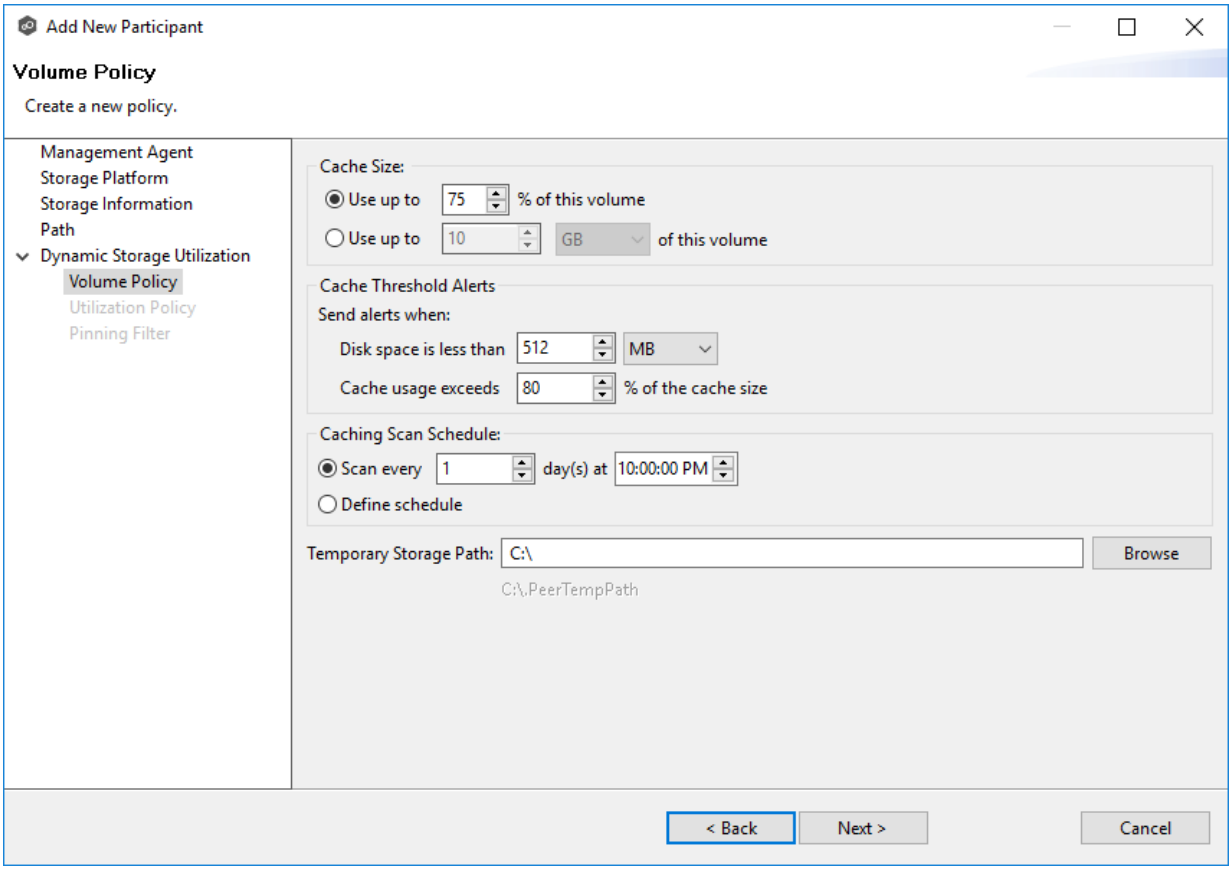

2. In the **Cache Threshold Alerts** section, set threshold values for automatic alerts about free disk space and cache usage.

   Peer Management Center will automatically display alerts in the **Alerts** tab when:

   - The amount of free disk space on the volume falls below the specified value. For example, if a 1 TB volume has 500 MB of free space and the threshold is set to 512 MB, an alert will be sent.
   - Cache usage on the volume exceeds the specified percentage of the cache size. For example, if the cache size is set to 80%, equating to 750 GB, DSU will start sending alerts when it has used 600 GB.

You can also send cache threshhold alerts via email alerts and SNMP notifications. You configure these in Dynamic Storage Utilization preferences for File Collaboration and File Synchronization jobs.

3. In the **Caching Scan Schedule**, set the frequency that the volume should be scanned to optimize the balance of fully hydrated vs stubbed files.

   This scan can be run daily at a specified time or you can define a more customized schedule.

4. In the **Temporary Storage Path** field, enter a path or browse to the location you want to be used as temporary storage space.

   The temporary storage space will be used to store the content of stub files as they are are being rehydrated. The content of files undergoing rehydration are referred to as **file blocks**. File blocks are fixed-length chunks of data that are read into memory when requested by an application. DSU will create a subfolder named **.PeerTempPath** under the location that you specify and temporarily store the file blocks in that subfolder.

   For optimal performance, we recommend that the temporary storage space be on the same volume as the watch set. If that is not possible, it should be on a high performance disk.

5. Click **Next**.

   The Utilization Policy page appears.

Note: If the Agent you selected is already being used as an edge participant in another job utilizing DSU, the existing volume policy will be displayed on this page. You can edit the existing volume policy; however, be aware that any modifications to the volume policy will be applied to every other job that uses this Agent as an edge participant and "touches" the same volume.

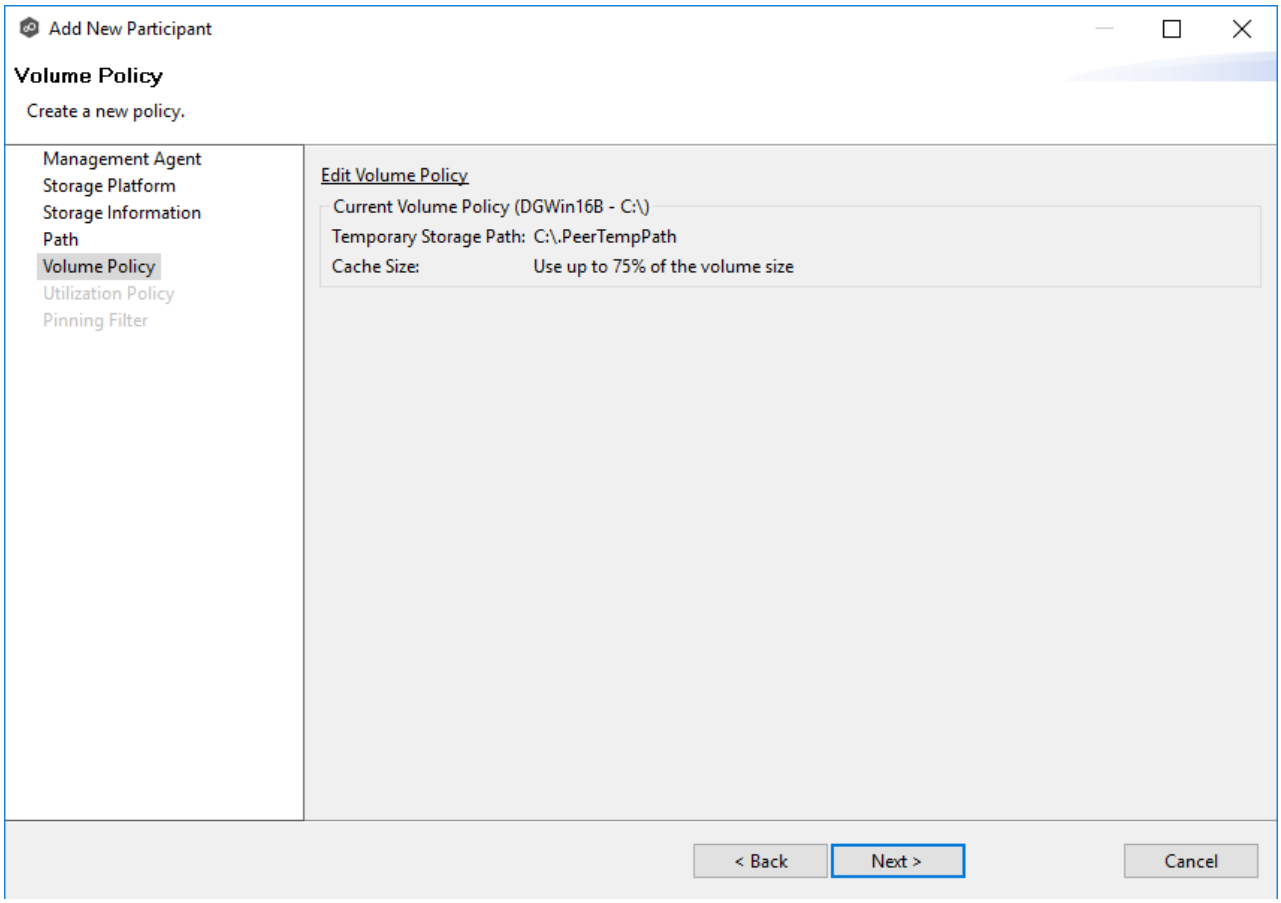

Utilization Policy

The **Utilization Policy** page appears if you chose the edge role for the participant. The primary purpose of a **utilization policy** is to specify the parameters that govern when files on this edge participant should be stubbed or fully hydrated. Whereas the volume policy controls how much space is available to DSU on a specific volume (or drive letter), the utilization policy controls whether to stub or hydrate a file.

Utilization policy parameters are based on the size of the files to be potentially stubbed and when they were last accessed and modified. A utilization policy enables you to balance getting the best performance while keeping the cache as full as possible.

You can select an existing utilization policy to apply to the job or create a new utilization policy. Whereas a volume policy is specific to a volume, a utilization policy can be reused for multiple jobs.

1. Select **New Policy** or **Existing Policy**.

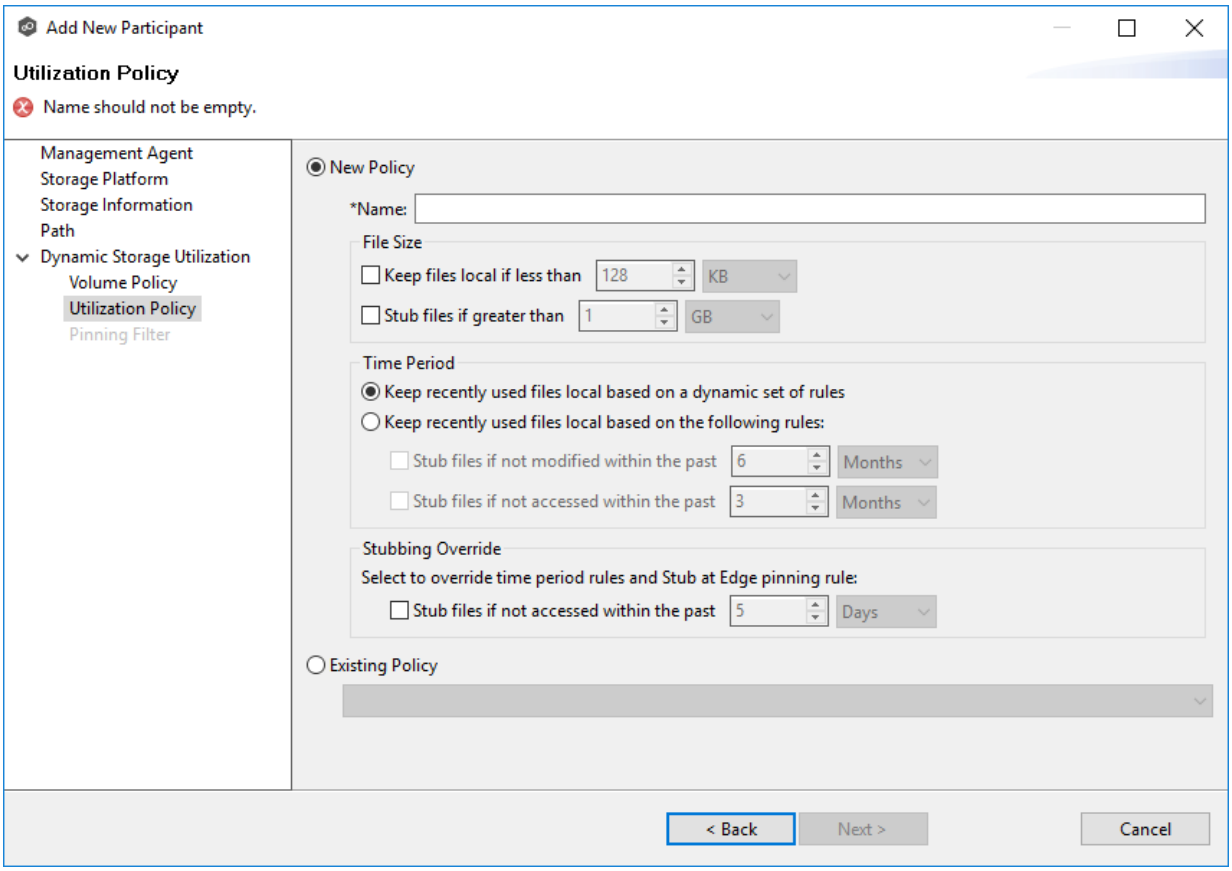

2. If you selected **Existing Policy**, select the policy, and then click **Next**.

   If you selected **New Policy**, enter a name for the policy.

3. (Optional) In the **File Size** section, select one or both options:

Field	Description
Keep files local if less than X size	Select this option if you want files under a specified size to remain local.
Stub files if greater than X size	Select this option if you want files over a specified size to be stubbed.

4. (Optional) In the **Time Period** section, select one of the options:

Field	Description
Keep recently used files local based on a dynamic set of rules	Select this option if you want DSU to control when to stub files based on last accessed and last modified times. DSU dynamically adjusts the accessed and modified time periods it uses based on the total amount of space that DSU is actively using on a volume.
Keep recently used files local based on the following rules	Select this option if you want to specify the dates used when stubbing files based on the last accessed and last modified.

5. If you selected the **Keep recently used files local based on the following rules** option in **Time Period**, two additional options are enabled; select the options to be used and specify the time periods to be used:

Field	Description
Stub files if not modified within the past X time period	Select this option and specify a time range to use the last modified date of a file to determine if it should be kept local or stubbed.
Stub files is not accessed within the past X time period	Select this option and specify a time range to use the last accessed date of a file to determine if it should be kept local or stubbed.

6. (Optional) In the **Stubbing Override** section, select the checkbox and a time period if you want to use the last accessed date of all recently hydrated files to determine if they should be kept local or re-stubbed.

7. Click **Next** or **Finish**.

   If you click **Next**, the Pinning Filter page appears.

Pinning Filter

The **Pinning Filter** page allows you to create a new pinning filter or select an existing pinning filter to apply to the job. A **pinning filter** specifies whether specific files or files in a particular

directory are always stubbed or always local on an edge participant. A pinning filter similar is to a utilization policy—it can be applied to multiple jobs. If there is a conflict between a pinning filter and utilization policy (where, for example, you might have something set to be always stubbed), the pinning filter will take precedence. Pinning filters are optional.

1. Select one of the options: **No Filter**, **New Filter**, or **Existing Filter**.

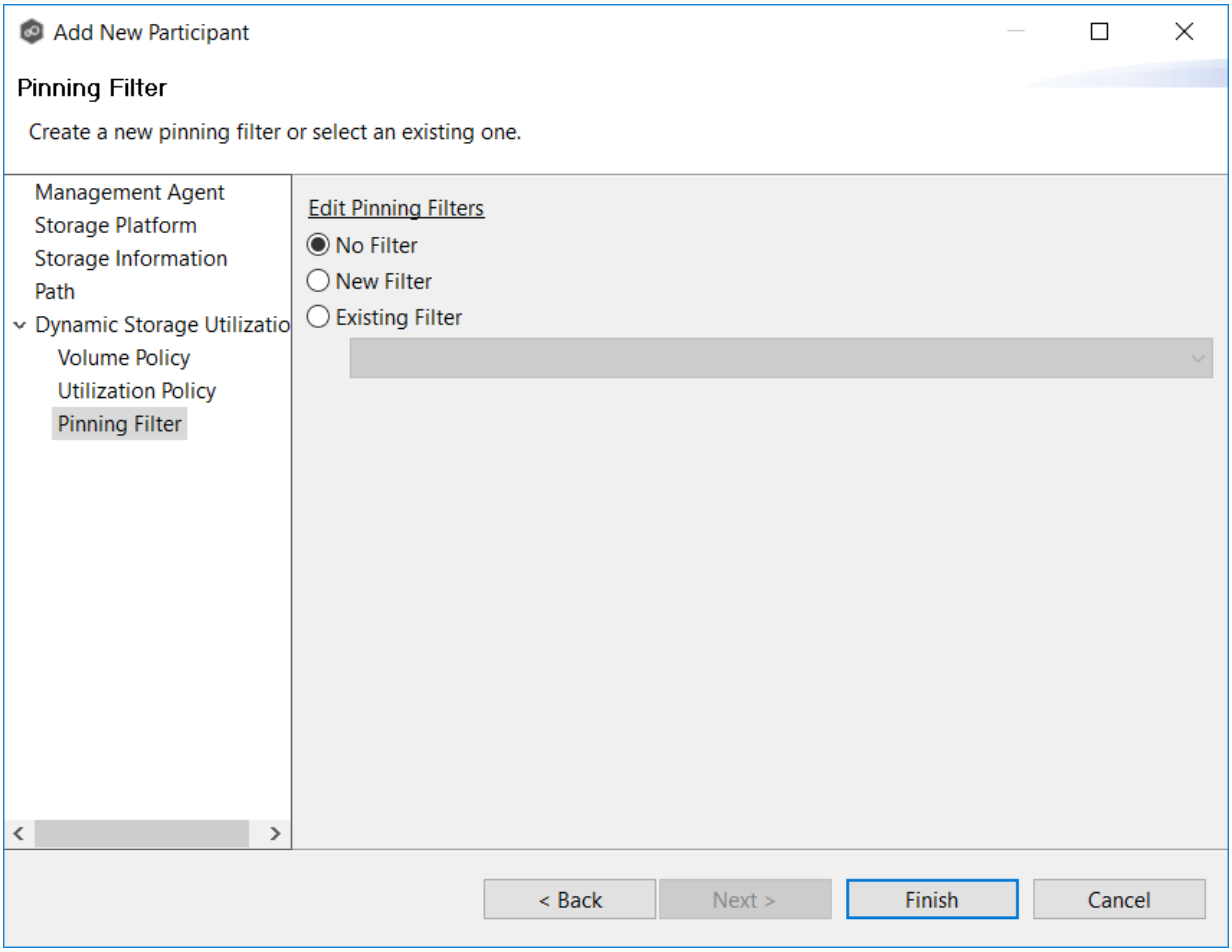

2. If you selected **No Filter**, click **Finish**; if you selected **Existing Filter**, select the filter, and then click **Finish**.

   If you selected **New Filter**, enter a name for the filter.

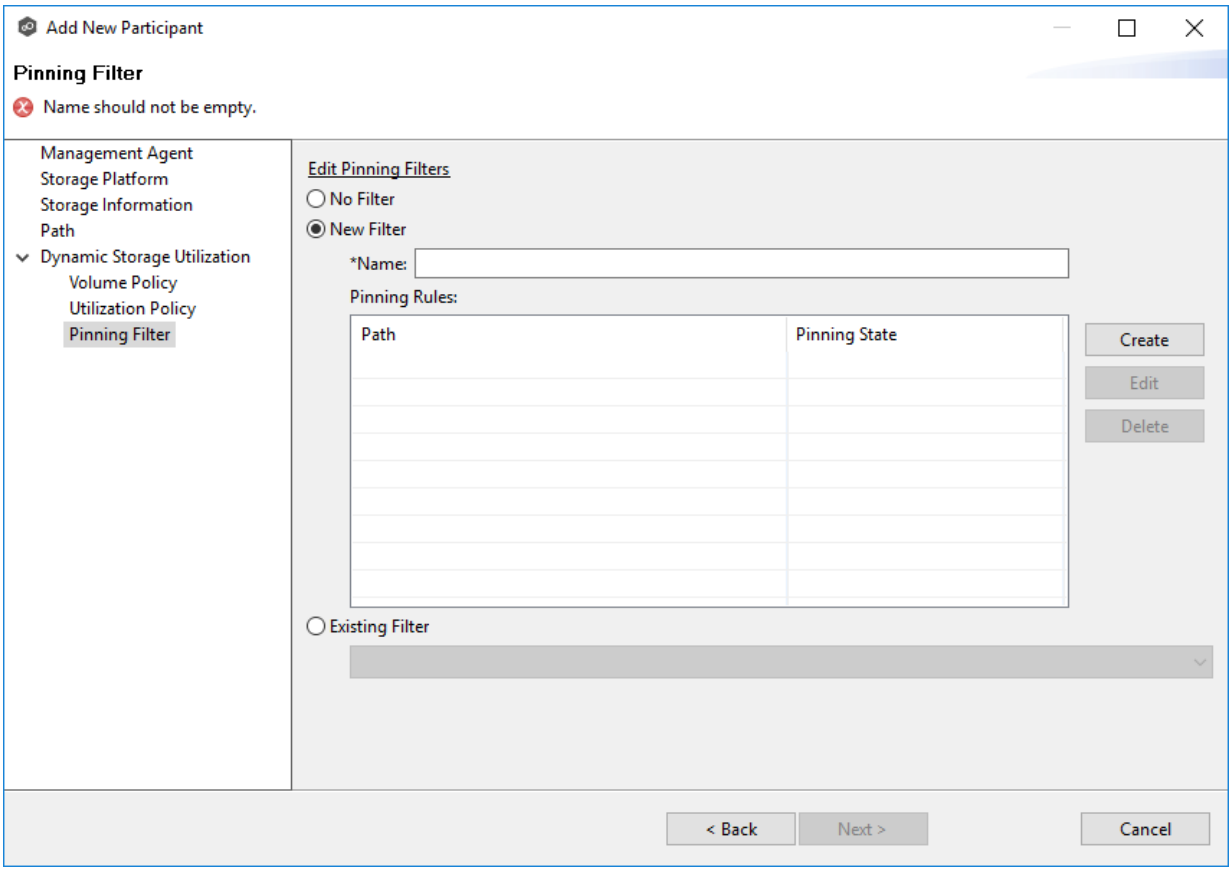

3. Enter a name for the filter.
4. Click **Create**.

   The **Create Pinning Rule** dialog appears.

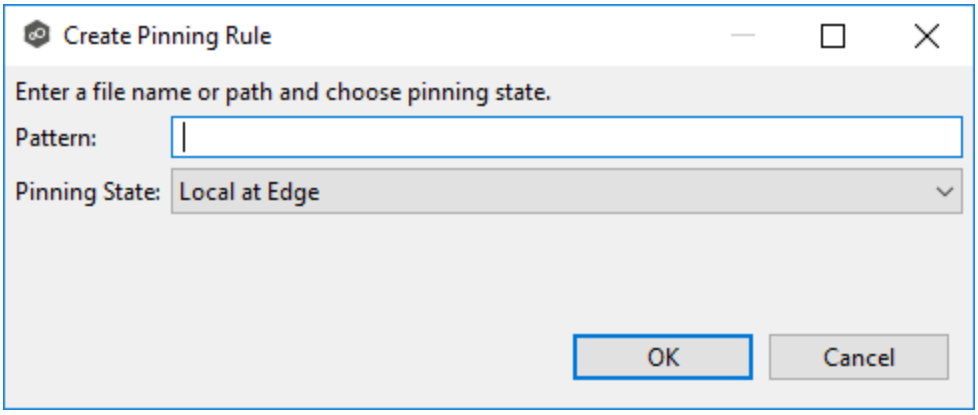

5. Enter a file name or path in the **Pattern** field and then choose a pinning state: **Local at Edge** or **Stubbed at Edge**.

6. Click **OK**.

   The rule appears in the filter table.

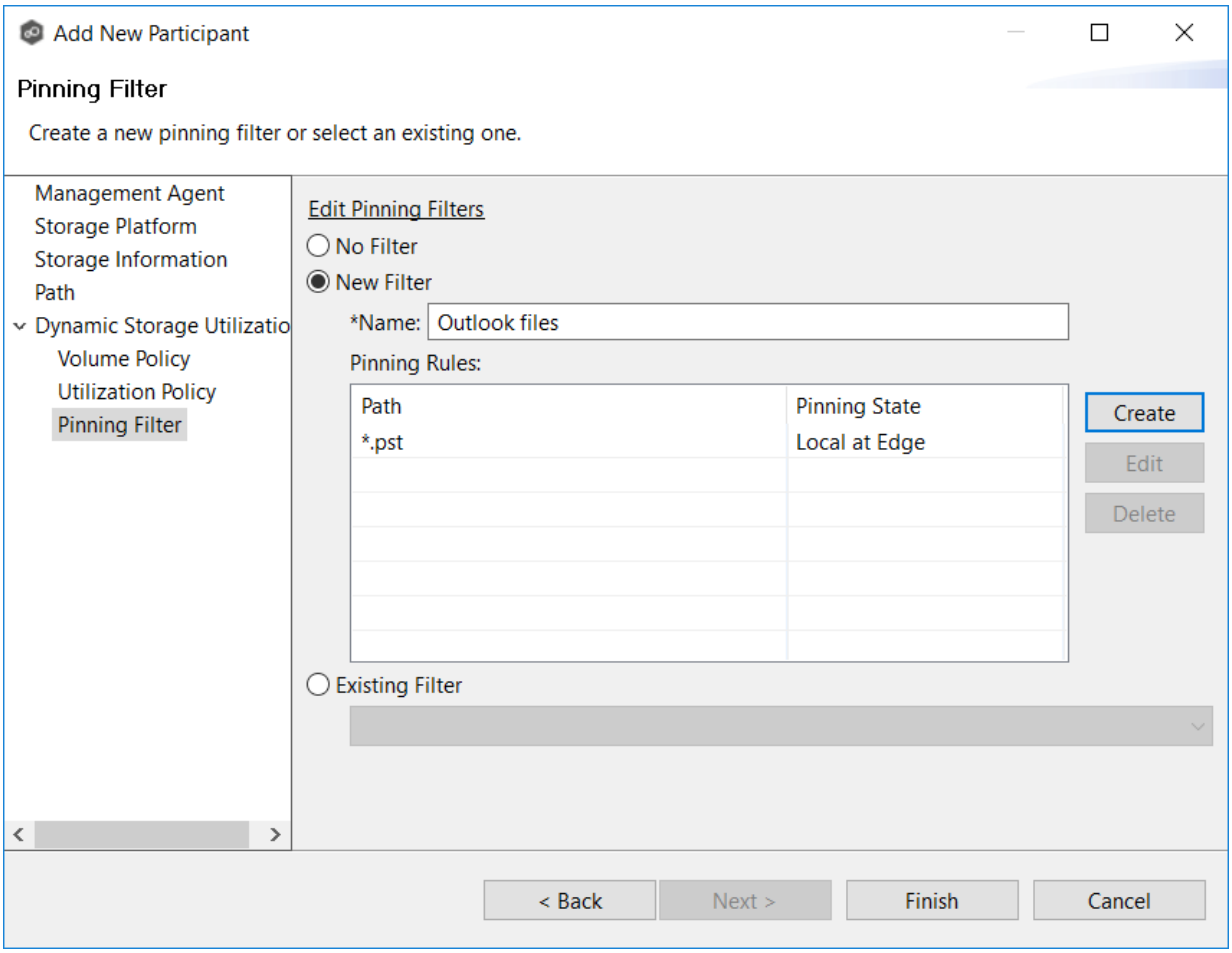

7. (Optional) Create additional pinning rules.

8. Click **Finish**.

   The **Participants** page reappears. The participant is listed in the **Participants** table with the **Edge** role.

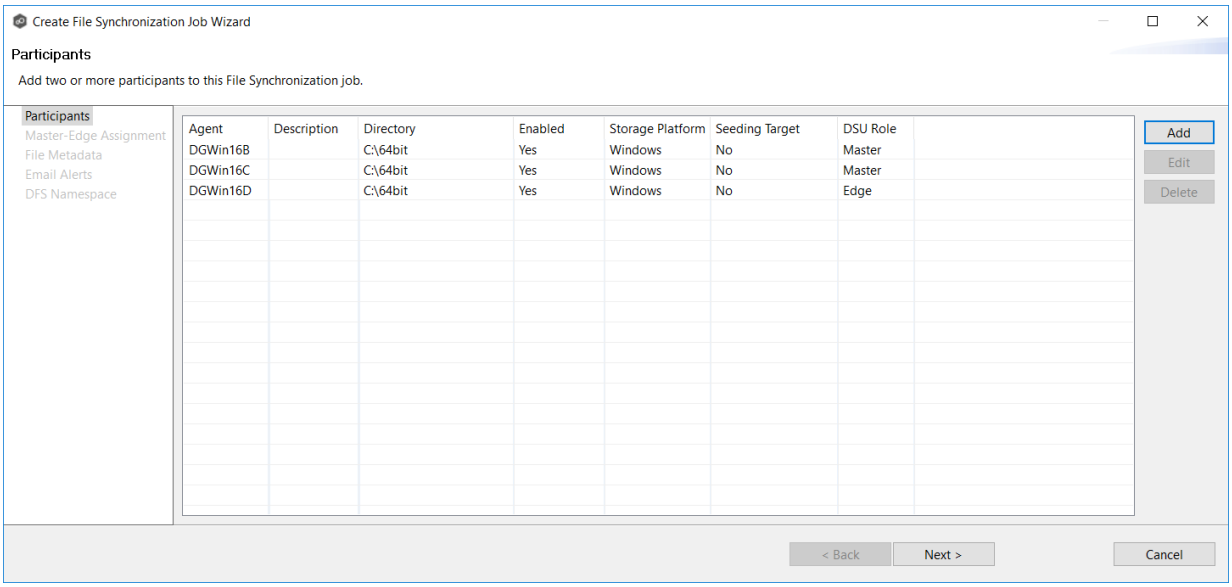

9. Continue adding more participants if applicable or continue with Step 4: Master-Edge Assignment

Step 3: Master-Edge Assignment

This step is optional.

The **Master-Edge Assignment** page appears only if you enabled Dynamic Storage Utilization for one or more participants in Step 2. The purpose of this page is to allow you to assign one or more master participants to each edge participant.

Every edge participant must have at least one master participant assigned to it, so that DSU will know which master participants to use when reading or rehydrating a stub file on an edge participant. For each edge participant, you want to assign the master participant that is the fastest and closest to it. This will increase the speed of rehydrating a stub file.

It is highly recommended that you assign, if possible, more than one master participant to an edge participant, so that if one master participant is not available, another master participant is able to provide the file content to the edge participant. You can set the order in which the master participants are consulted.

If you have only one edge participant and one master participant, the master participant is automatically assigned to the edge participant and listed in the Master Participants column. If you have multiple master participants, you need to explicitly assign master participants to each edge participant and then set the failover order.

1. Select an edge participant in the **Assignment** table.

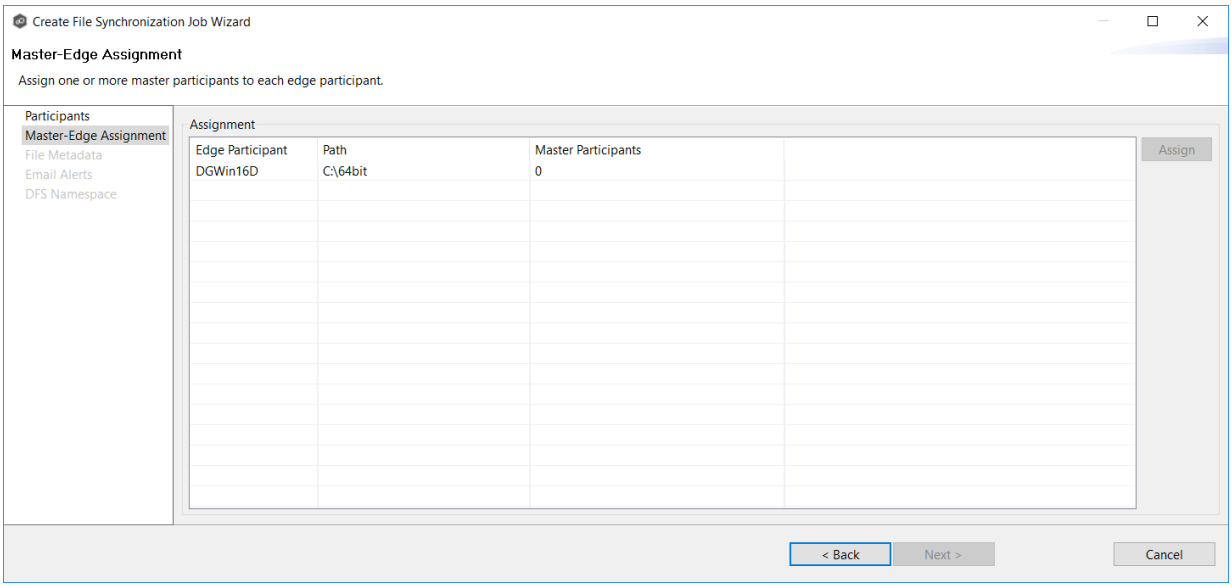

2. Click the **Assign** button.

   The **Assign Master Participants** dialog appears.

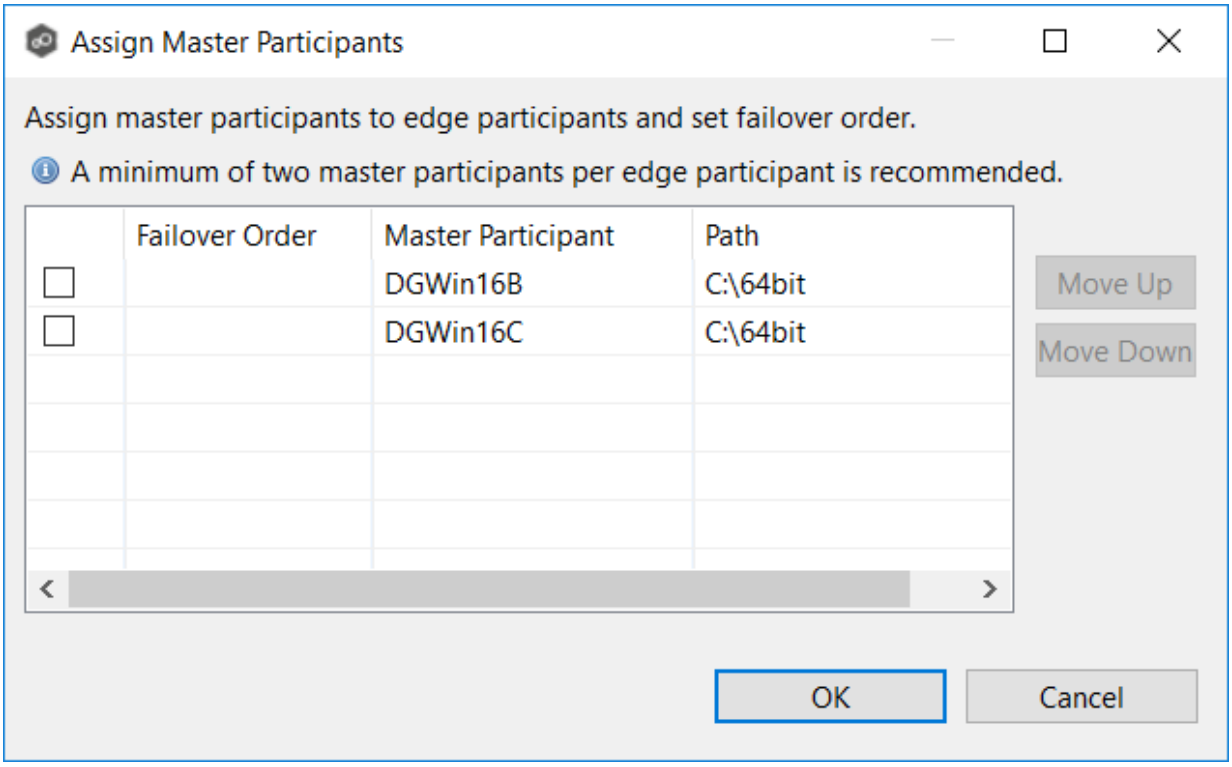

3. Select the master participants you want to assign to the edge participant.

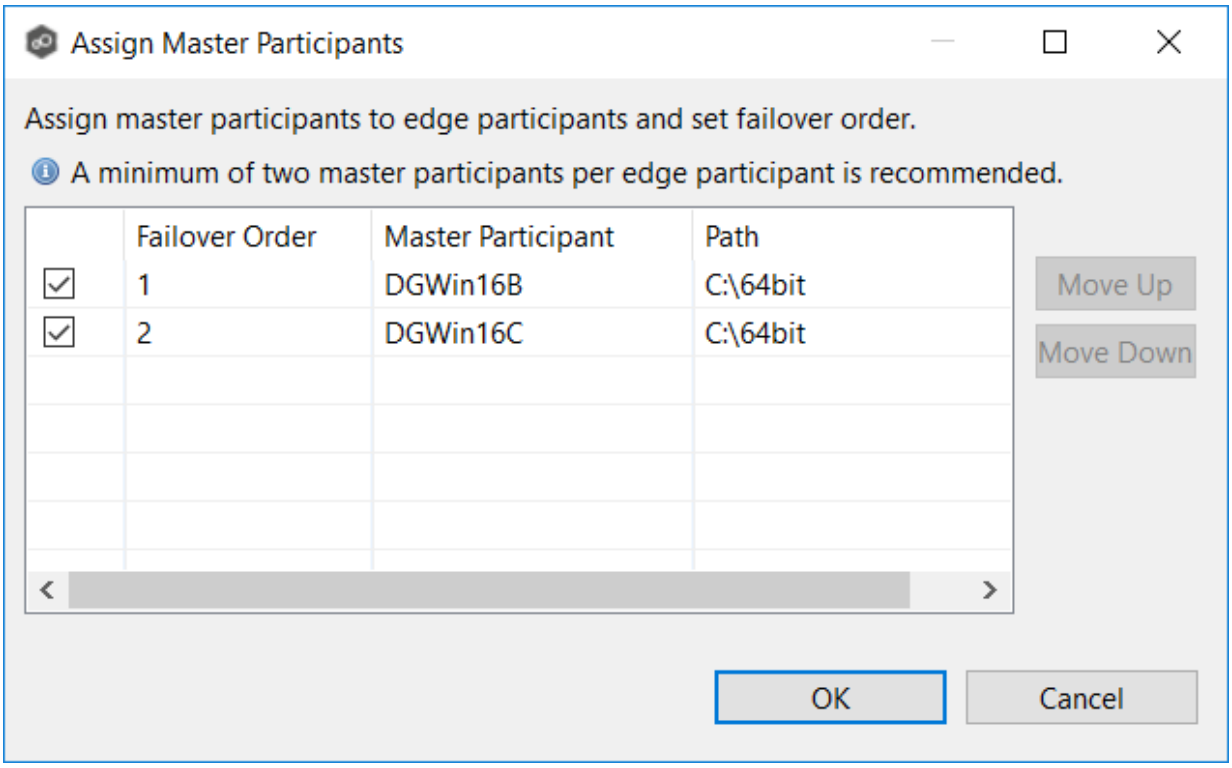

4. (Optional) Change the failover order by selecting a master participant and using the **Move Up** and **Move Down** buttons.

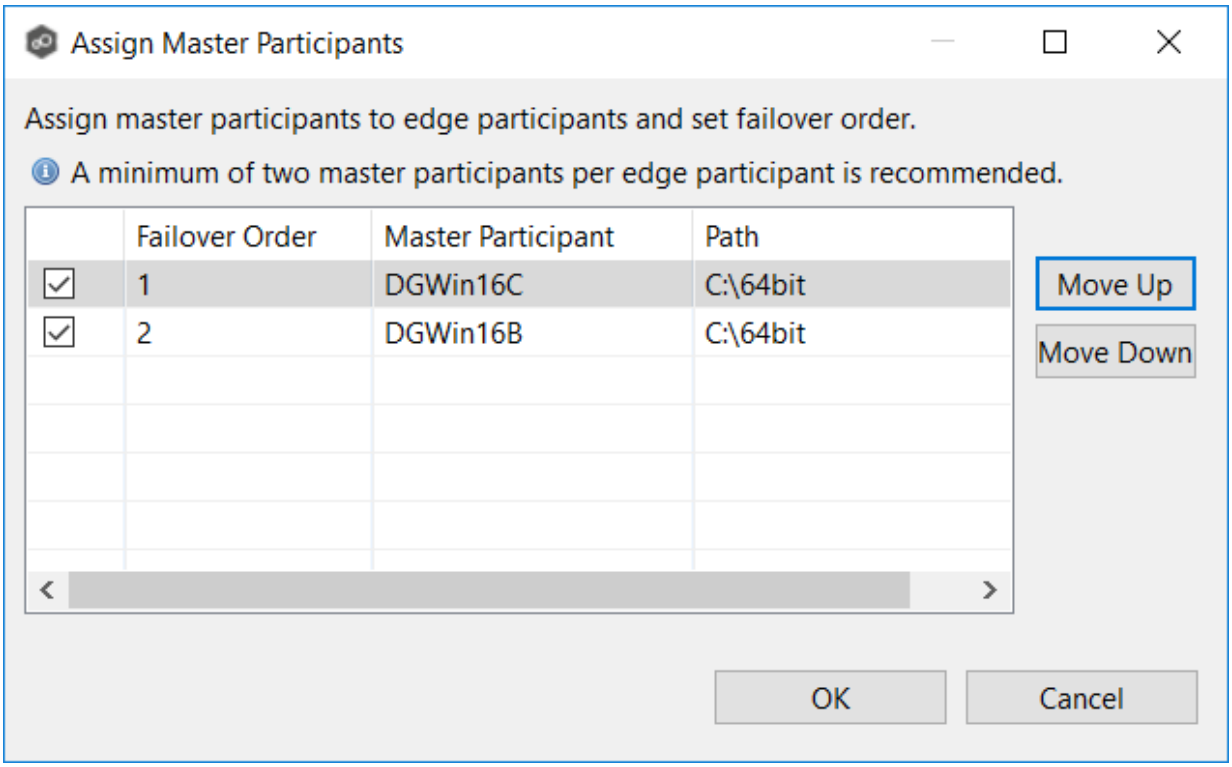

5. Click **OK**.

   The **Master Participants** column has been updated for that particpant.

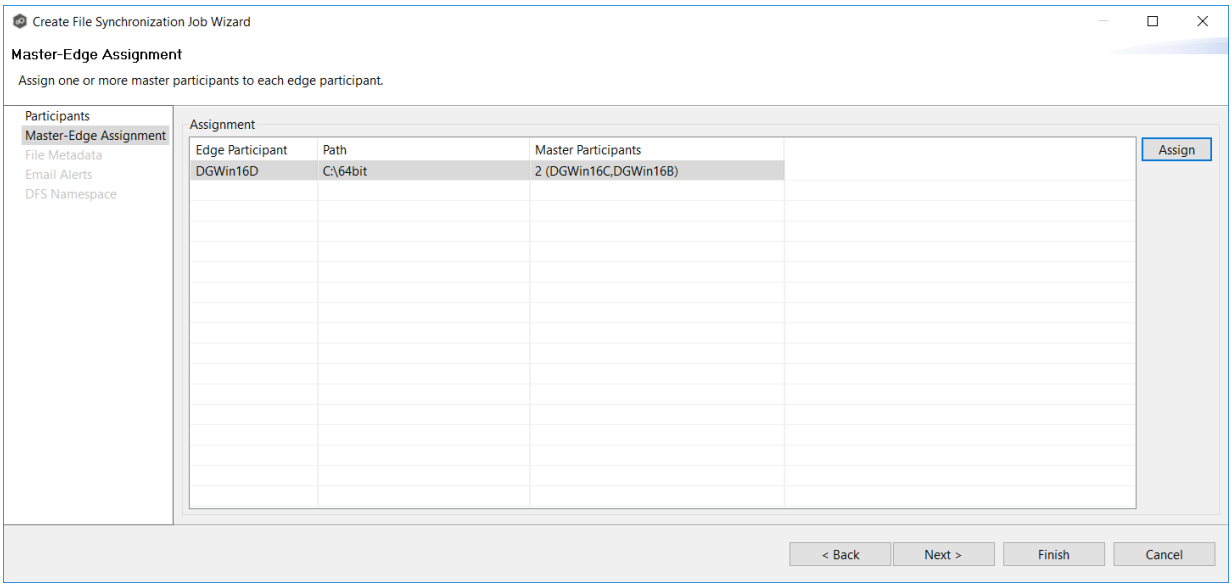

6. Repeat Steps 1-5 for each edge participant.

7. Click **Next**.

   The File Metadata page appears.

Step 4: File Metadata

This step is optional.

The **File Metadata** page allows you to specify whether you want to synchronize NTFS security permissions metadata and the types of metadata. It also allows you to specify which volume/share/folder's metadata should be used if there is a conflict during the initial synchronization. The volume/share/folder used if there is a conflict is referred to as the master host.

For more information on synchronizing NTFS metadata, see File Metadata Synchronization in the Advanced Topics section.

To enable file metadata synchronization:

1. Select when you want the metadata synchronized (you can select one or both options):

- **Enable synchronizing NTFS security descriptors (ACLs) in real-time** - Select this option if you want the metadata synchronized in real-time. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized to all participants as they occur.

- **Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan** - Select this option if you want the metadata replicated during the initial scan. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized during the initial scan.

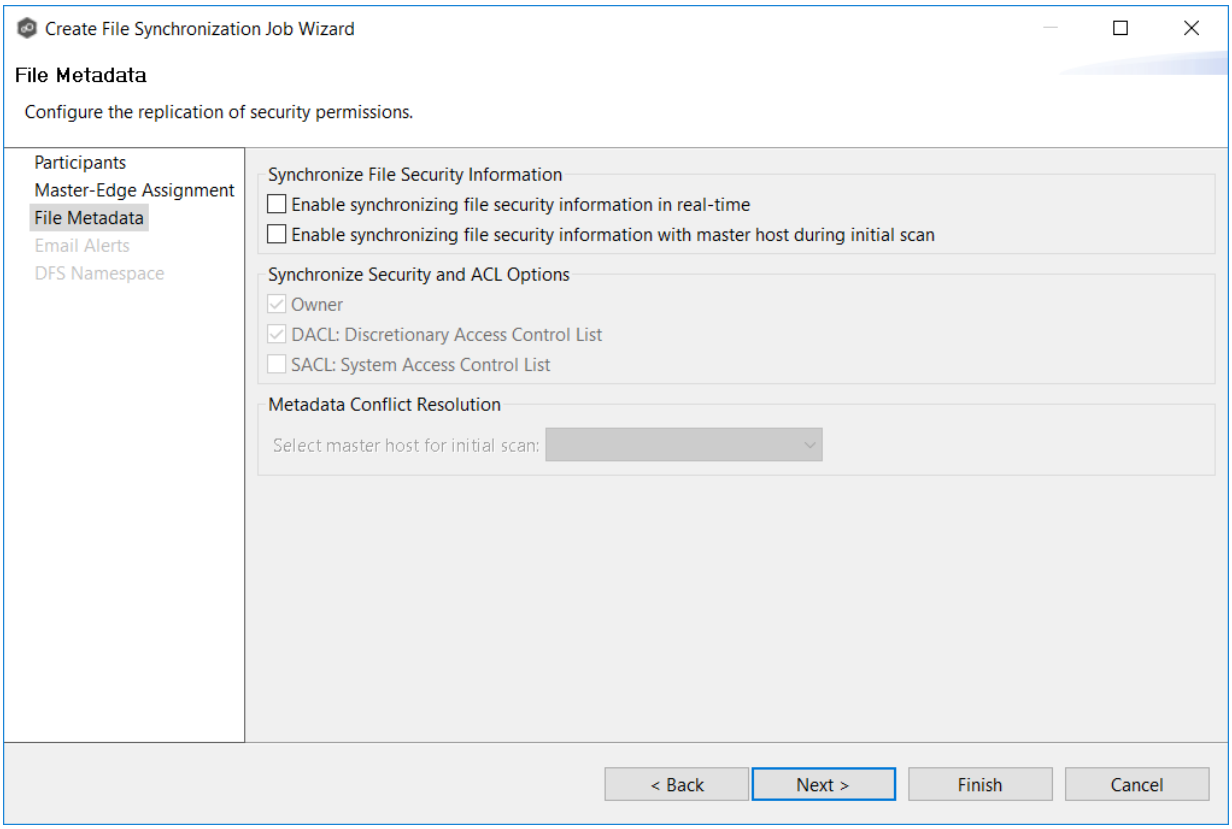

2. Click **OK** in the message that appears after selecting a metadata option.

3. If you selected either of the first two options in the **Synchronize Security Descriptor Options** section, select the security descriptor components (Owner, DACL, and SACL) to be synchronized.

4. If you selected the option for metadata synchronization during the initial scan, select the host to be used as the master host in case of file metadata conflict.

   If a master host is not selected, then no metadata synchronization will be performed during the initial scan. If one or more security descriptors do not match across participants during the initial scan, conflict resolution will use permissions from the designated master host as the winner. If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

5. Click **Next**.

   The Email Alerts page is displayed.

Step 5: Email Alerts

This step is optional.

An email alert notifies recipients when a certain type of event occurs, for example, session abort, host failure, or system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in Preferences and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply.

To create a new alert, see Email Alerts in the Preferences section.

To apply an existing email alert to the job.

1. Click the **Select** button.

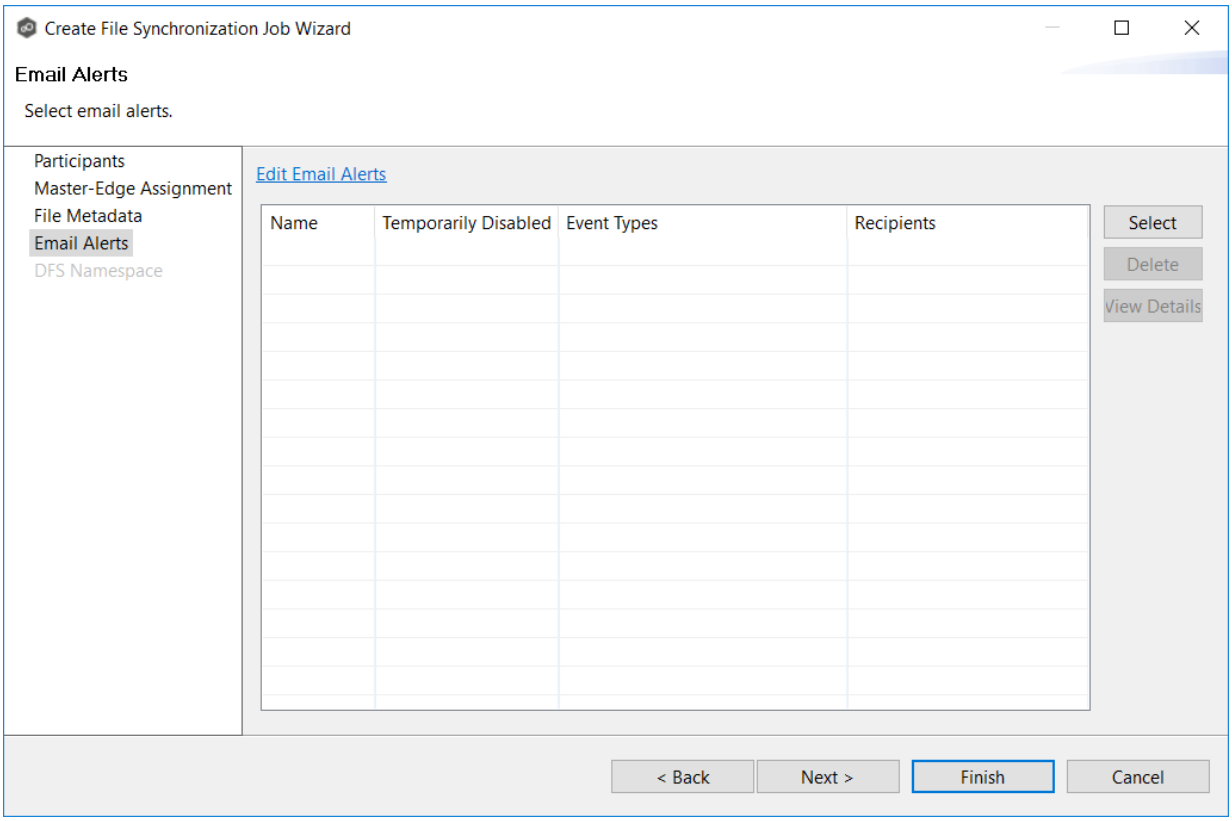

The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.

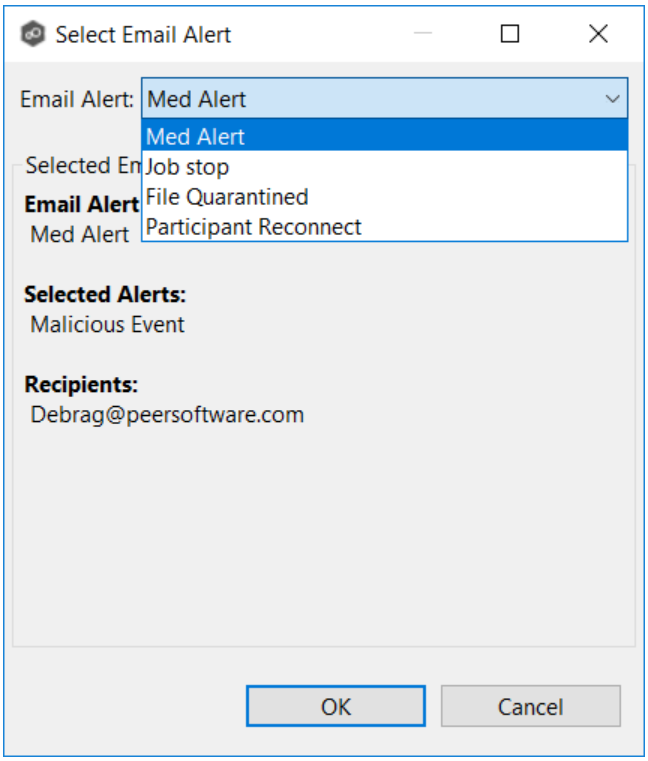

3. Click **OK**.

   The alert is listed on the **Email Alerts** page.

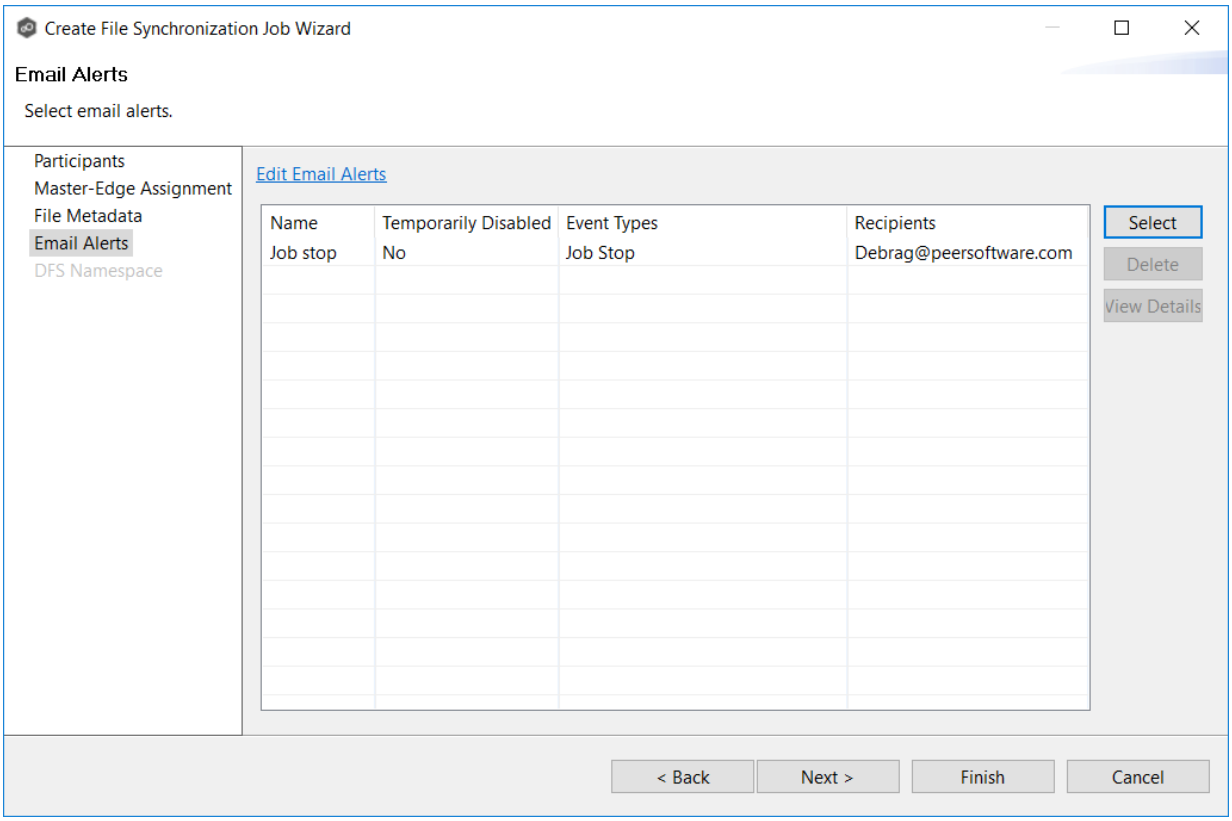

4. (Optional) Repeat steps 1-3 to apply additional alerts.

5. Continue to Step 6: DFS Namespace.

Step 6: DFS Namespace

This step is optional.

The **DFS Namespace** page presents three options for linking a DFS namespace folder to a File Synchronization job.

To link a namespace to this job:

1. Click the **Enable DFS Linking** checkbox.

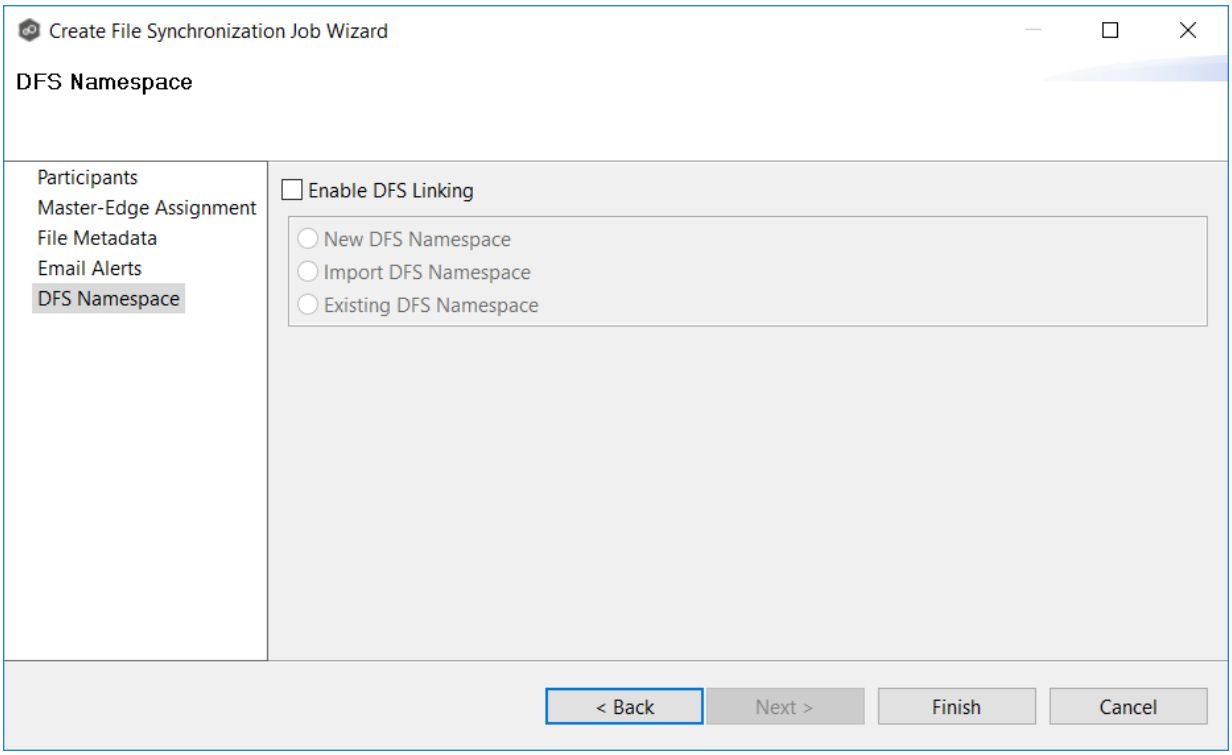

The three options are enabled.

2. Select one of the three options:

   - **New DFS Namespace** - Select this option if you want to create a new namespace that will automatically be linked to this job. If you select this option, the **Create DFS-N Management Job Wizard** opens. Follow these steps to create a new namespace.

   - **import DFS Namespace** - Select this option if you have a namespace that was created using the Microsoft DFS Management tool and is not currently being managed by a DFS-N Management job. If you select this option, the **Import Existing Namespaces** wizard opens. For detailed instructions, follow these steps to import an existing namespace.

   - **Existing DFS Namespace** - Select this option if you want to use an existing namespace that is being managed by a DFS-N Management job. If you select this option, it will display the namespace folder and folders associated with namespace. If you want to make changes to the namespace, you can edit the DFS-N Management job managing that namespace.

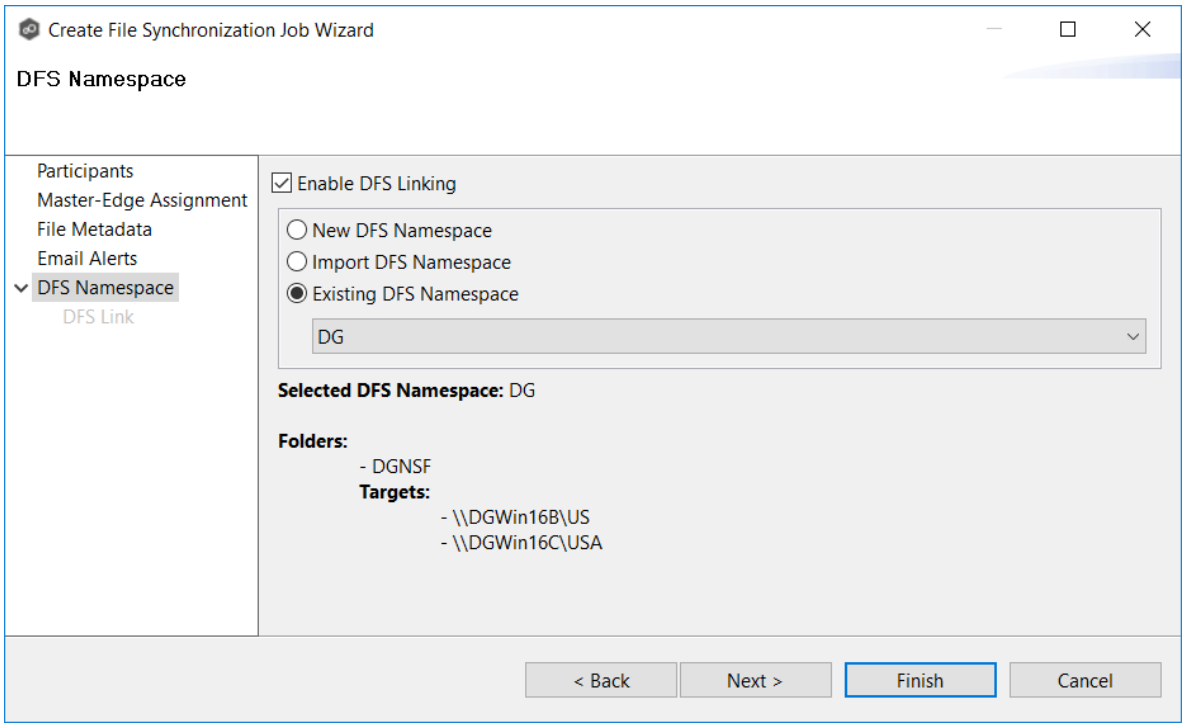

Click **Next** if you want to link participants to folder targets on the **DFS Link** page; otherwise continue with Step 3.

For more information about linking participants to folder targets, see Linking an Existing Namespace Folder to an Existing File Collaboration or File Synchronization Job.

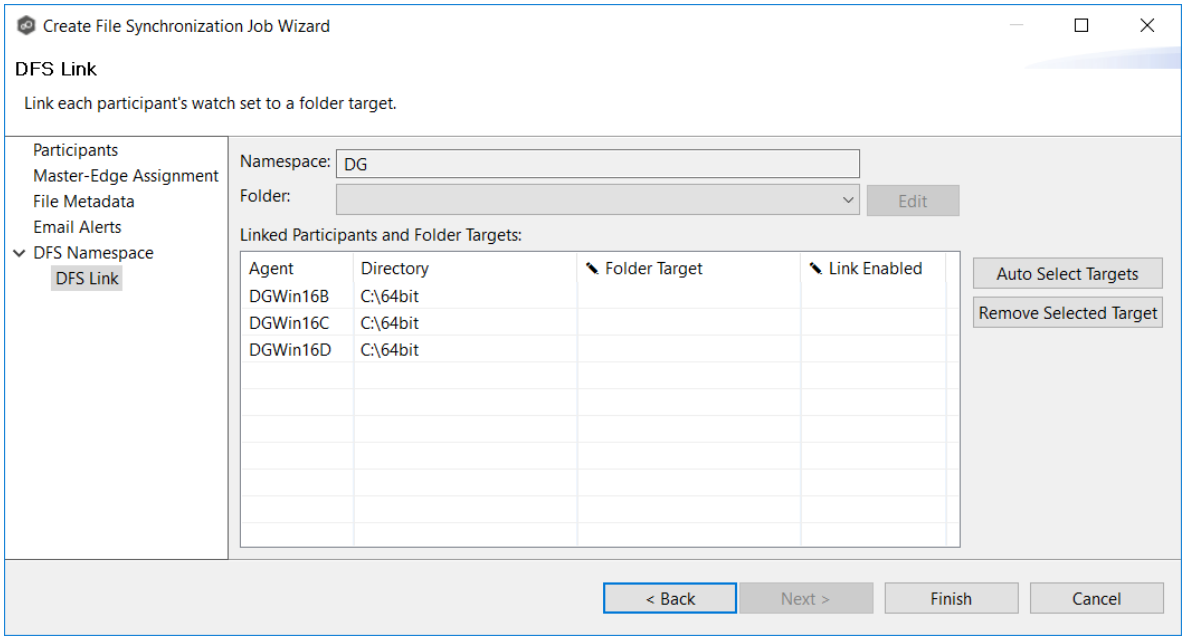

3. Continue to Step 6: Save Job.

Step 7: Save Job

You are now ready to save the job configuration.

1. If you are satisfied with your job configuration, click **Finish** to save your job. Otherwise, click the **Back** button and make any necessary changes.

   Congratulations! You have created a File Synchronization job. It is now listed in the **Jobs** view under **File Synchronization**.

2. Click the **Summary** tab to view the summary information about the job.

   See Running and Managing a File Synchronization Job Running for more information.

Editing a File Synchronization Job

You can edit a File Synchronization job while it is running; however, any changes will not take effect until the job is restarted.

Overview

When you create a File Synchronization job, the **Create Job** wizard guides you through the process, presenting the most common options for configuration. When editing a job, you have access to all options, allowing you to fine-tune the job configuration. Options not included in the initial job creation include:

- Application Support
- Conflict Resolution
- Delta Replication
- DFS-N
- File and Folder Filters
- File Locking

- General
- Logging and Alerts
- Scheduled Replication
- SNMP Notifications
- Target Protection
- Tags

You can edit multiple File Synchronization jobs simultaneously. For information about simultaneously editing multiple jobs, see Editing Multiple Jobs.

Editing a Job

To edit a File Synchronization job:

1. Select the job in the **Jobs** view.
2. Right-click and select **Edit Job**.

   The **Edit File Synchronization Configuration** dialog appears.

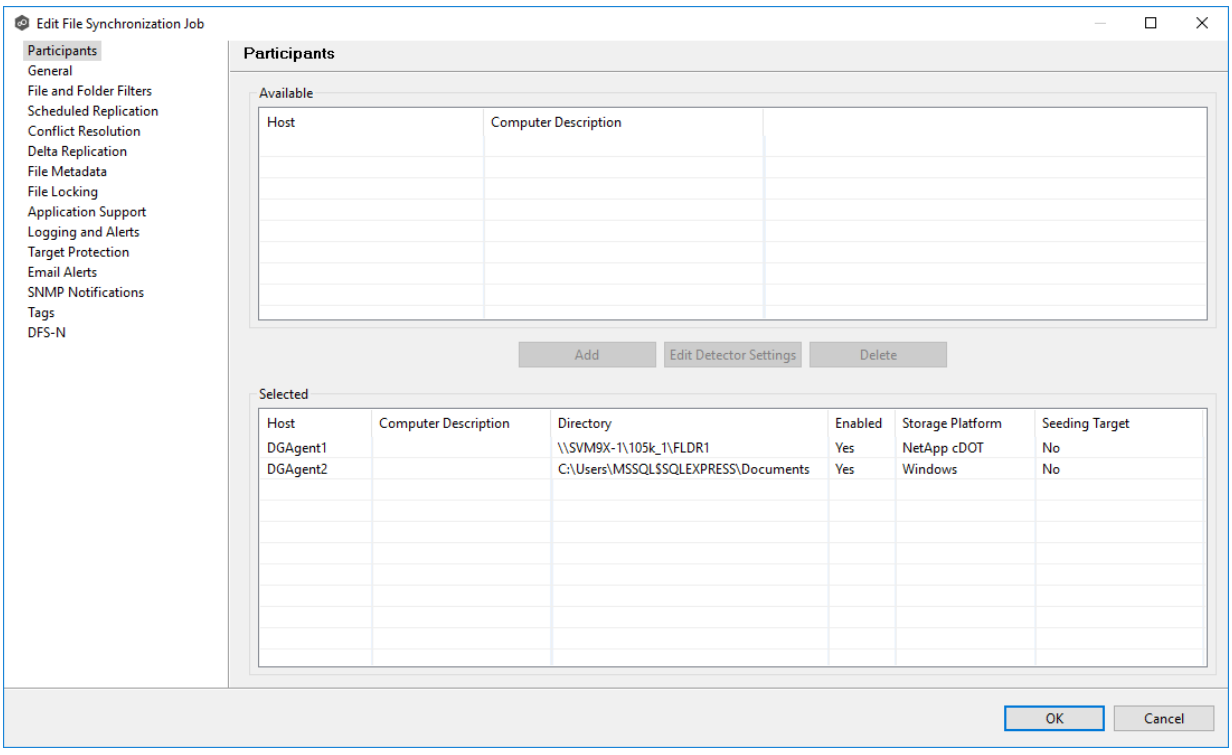

3. Select a configuration item in the navigation tree and make the desired changes:
   - Participants
   - General
   - File and Folder Filters
   - Scheduled Replication
   - File Conflict Resolution
   - Delta Replication
   - File Metadata
   - File Locking
   - Logging and Alerts
   - Application Support
   - Target Protection
   - Email Alerts
   - SNMP Notifications
   - Tags
   - DFS-N
4. Click **OK** when finished.

Participants

The **Participants** page in the **Edit File Synchronization Job** dialog allows you to:

- Add and remove participants from a job.
- Edit a participant.

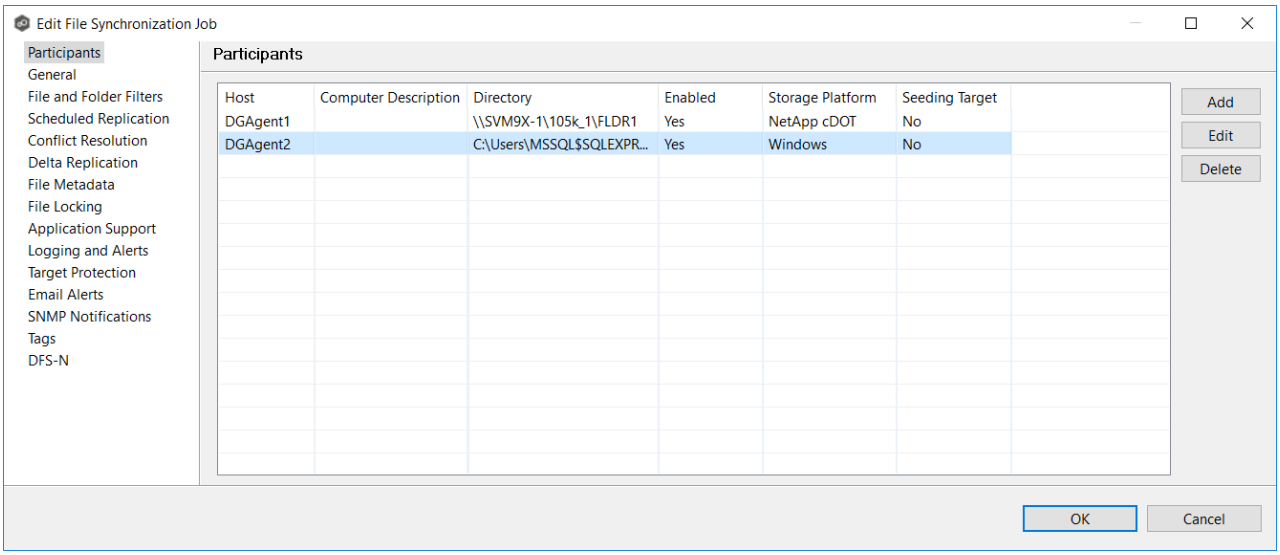

This topic describes adding and deleting participants in a File Synchronization job.

Adding a Participant to a File Synchronization Job

To add a participant to a File Synchronization job:

1. Select the job in the Jobs view; right click and select **Edit Job**.

   The **Edit File Synchronization** dialog opens; the **Participants** page displays the current job participants.

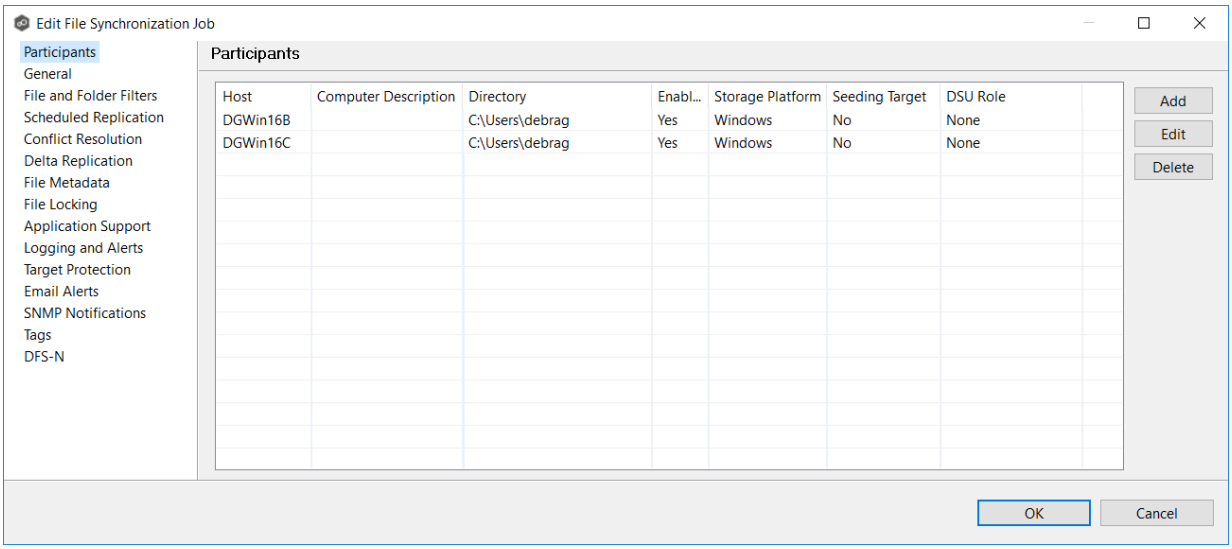

2. Click the **Add** button.

   The **Add New Participant** wizard opens; the **Management Agent** page lists the Agents available to be added.

   Tip: If the Agent you want is not listed, try restarting the Peer Agent Windows Service on that host. If it successfully connects to the Peer Management Broker, then the list is updated with that Agent.

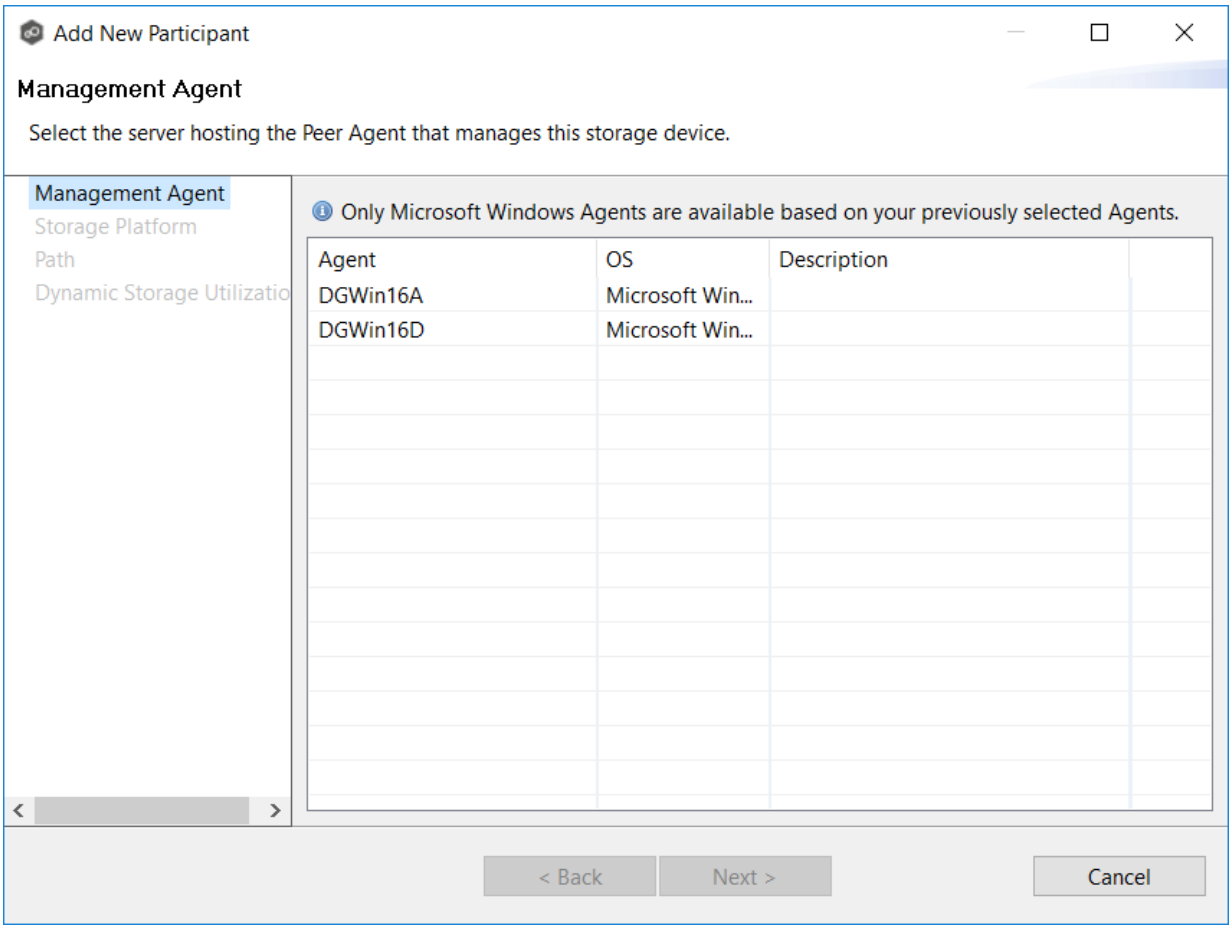

3. Select a Management Agent, and then click **Next**.

   The **Storage Platform** page appears.

4. Select the type of storage platform that hosts the data you want to synchronize, and then click **Next**.

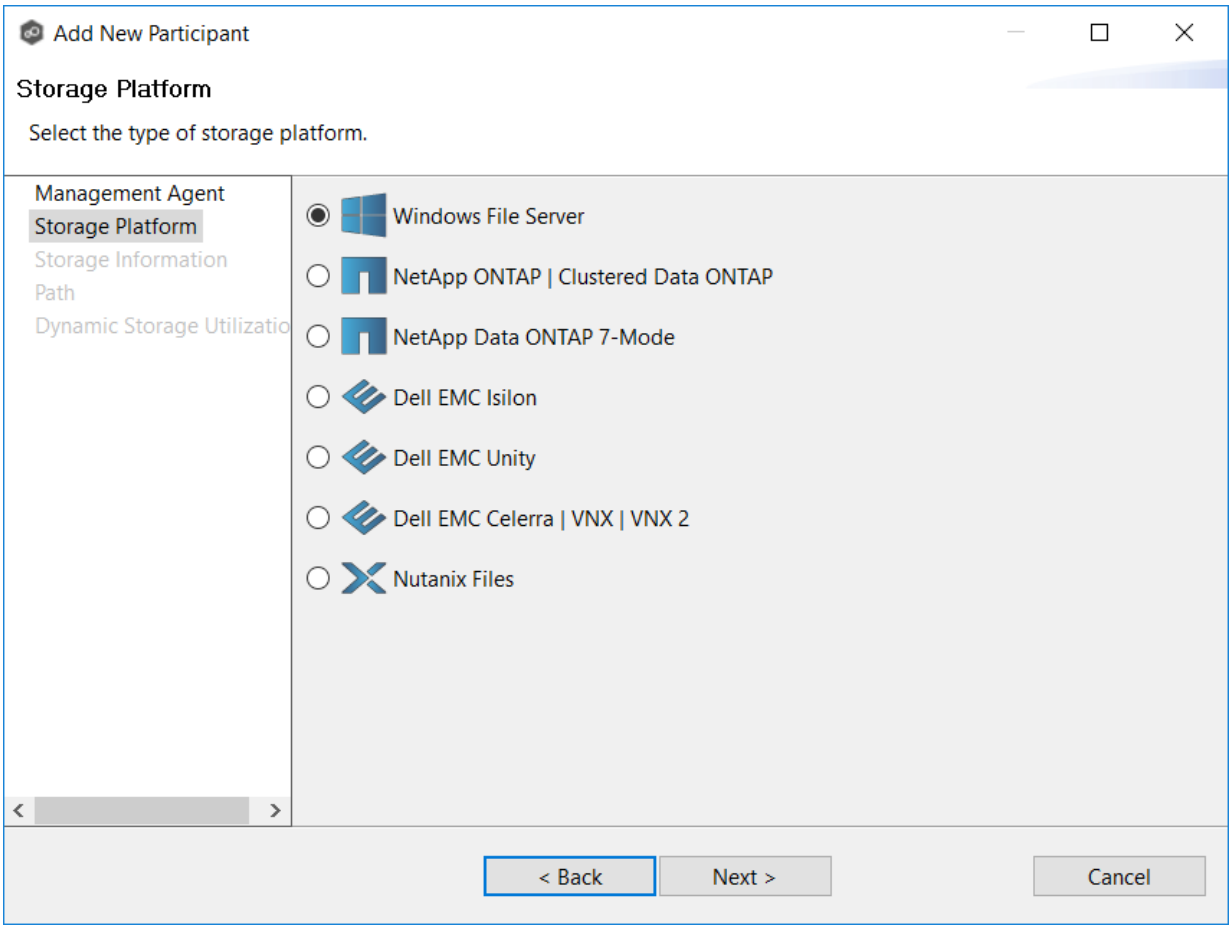

The **Storage Information** page appears; the choices available depend on your selection in the **Storage Platform** page.

5. Enter the requested information.

   Windows File Server

   NetApp ONTAP | Clustered Data ONTAP

   NetApp Data ONTAP 7-Mode

   Dell EMC Isilon

   Dell EMC Unity

   Dell EMC Celerra | VNX | VNX 2

   Nutanix Files

6. Click **Next**.

The **Path** page appears.

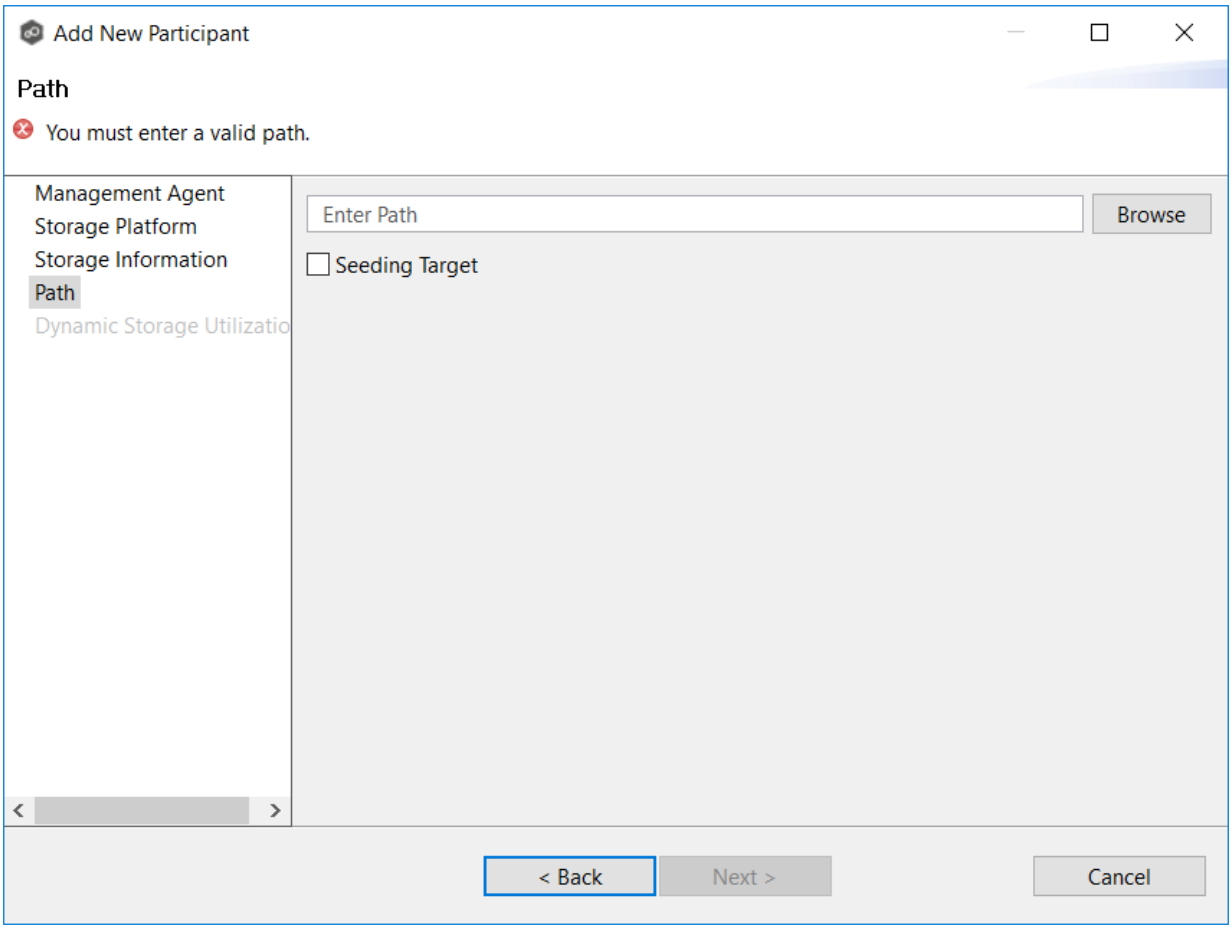

7. Browse to or enter the path to the watch set.

8. (Optional) Select the **Seeding Target** checkbox, and then click **OK** in the dialog that appears.

9. Click **Next**.

   The **Dynamic Storage Utilization** page appears.

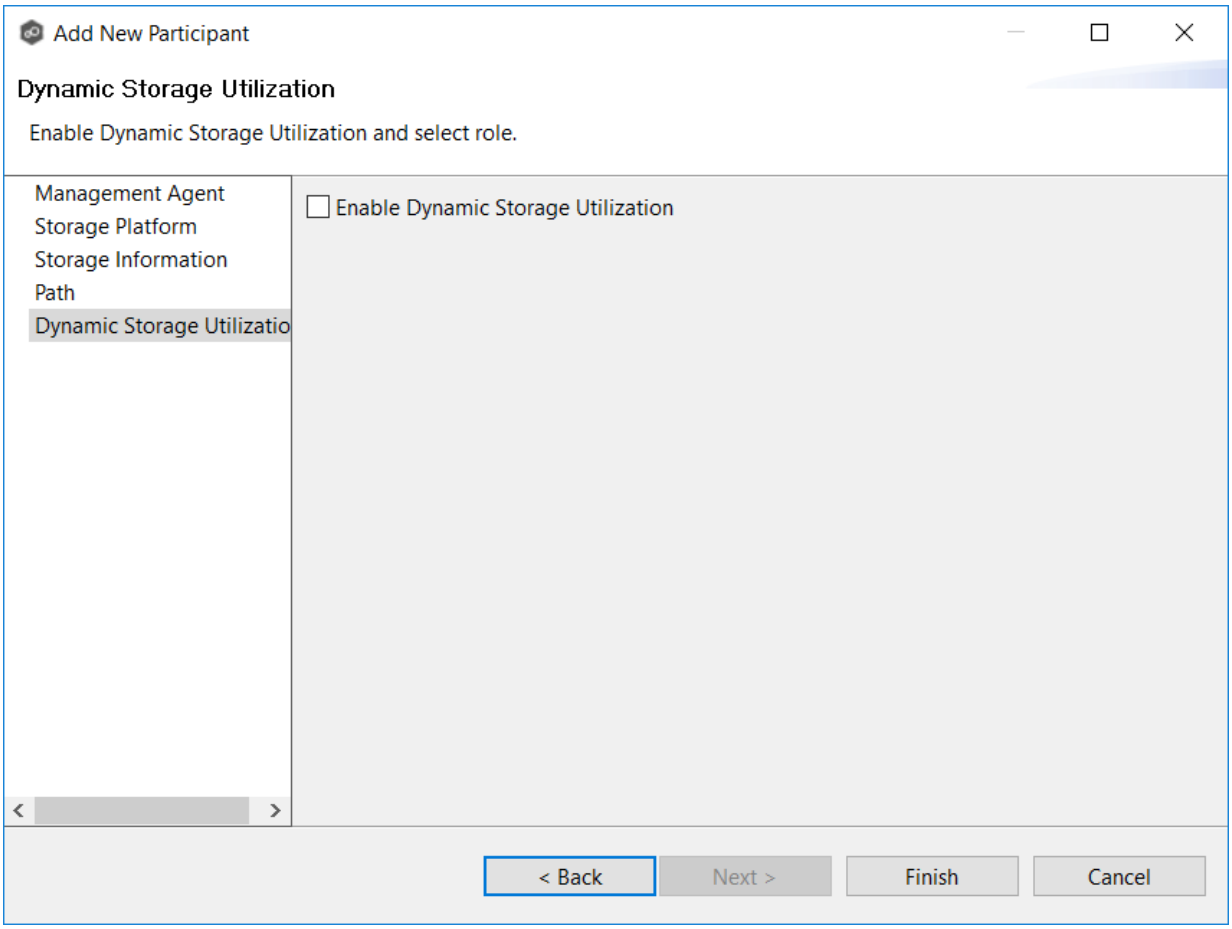

10. (Optional) Select the **Enable Dynamic Storage Utilization** checkbox if you want this participant to be able to use Dynamic Storage Utilization; otherwise, click **Finish**.

11. If you enabled Dynamic Storage Utilization, follow the steps outlined in Step 2: Dynamic Storage Utilization in Creating a File Synchronization Job.

    For more information about DSU, see Dynamic Storage Utilization in Advanced Topics.

12. Click **Finish** to complete the wizard.

    The new participant appears in the **Participants** table.

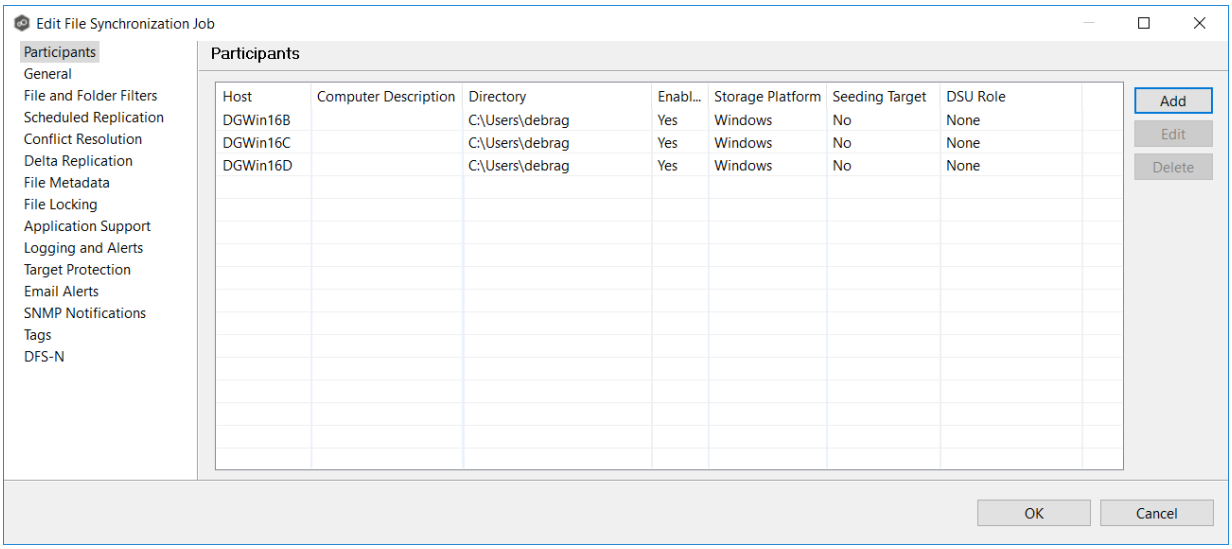

Deleting a Participant from a File Synchronization Job

To delete a participant from a File Synchronization job:

1. In the **Edit File Synchronization** dialog, select the participant in the **Participants** table you want to remove from the job.

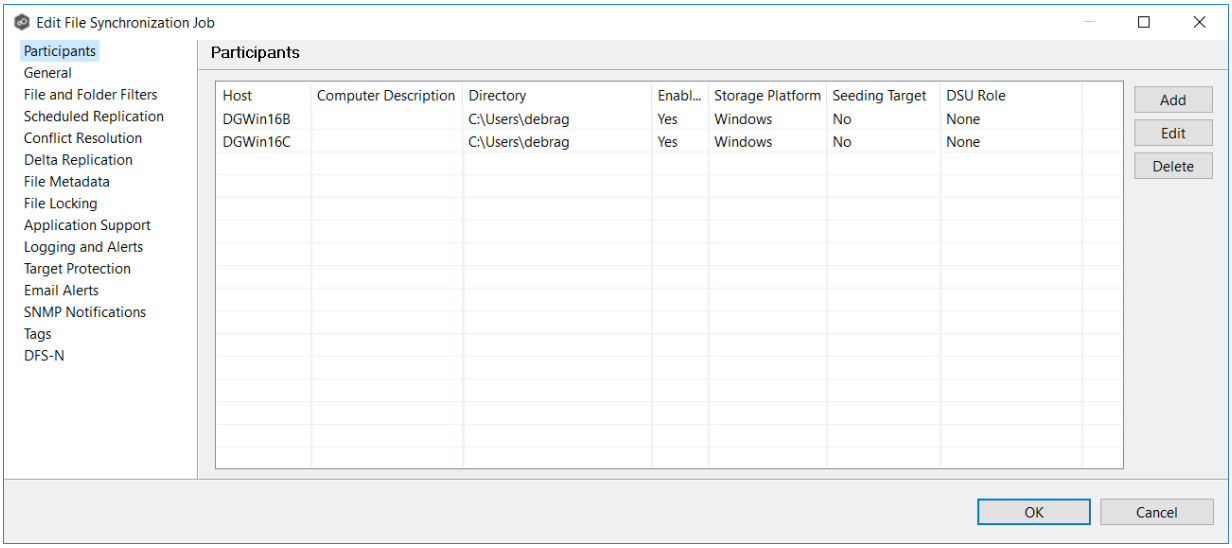

2. Click the **Delete** button

3. Click **OK** in the **Delete Confirmation** dialog.

   The participant is removed from the **Participants** table.

Note: A File Synchronization job must have at least two participants, so if after deleting a participant, there is only a single participant, you must add another participant to the job.

To edit a participant:

1. In the **Edit File Synchronization** dialog, select the participant in the **Participants** table you want to edit.

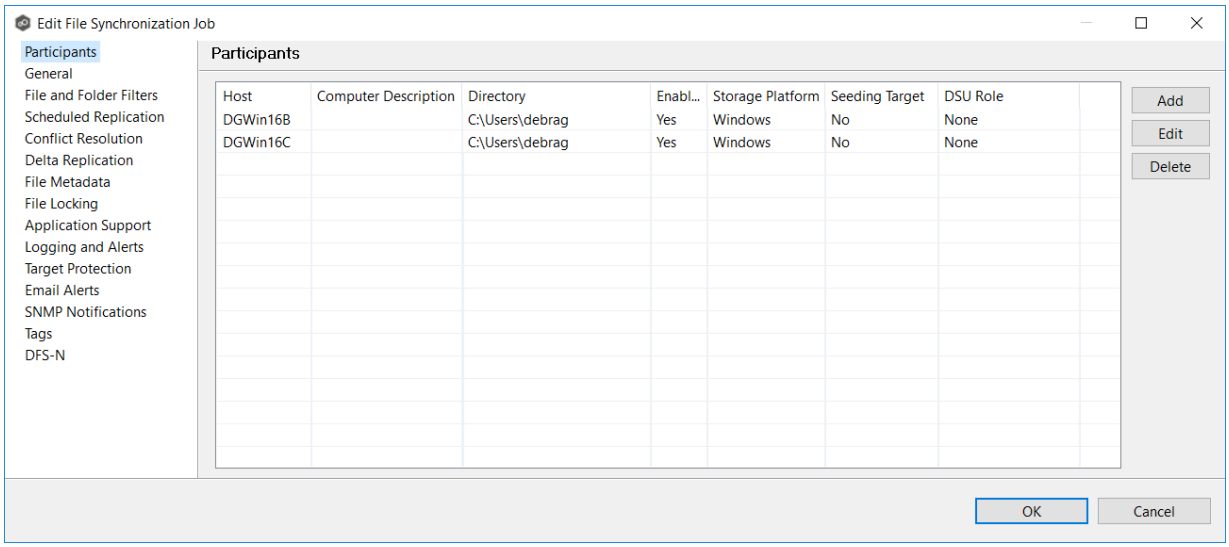

2. Click **Edit**.

   The **Edit Participant** dialog appears.

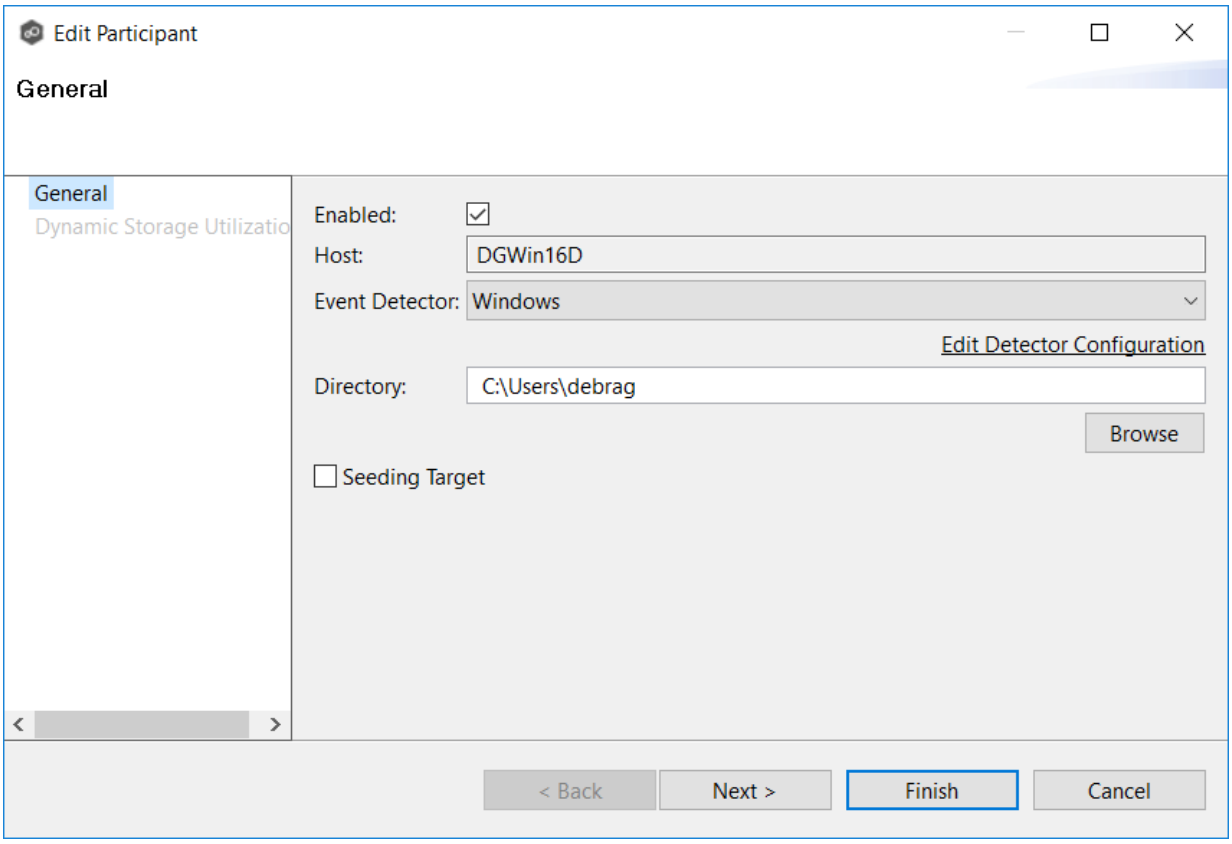

3. To enable or disable the agent, select or deselect the **Enabled** checkbox.

4. To change the directory/folder/share that is replicated, enter the path or browse to the new watch set in the **Directory** field.

5. If the storage device that the agent is managing has changed to a different storage platform, click **Edit Detector Configuration**, and then make the necessary modifications.

6. To change whether the participant is a seeding target, select or deselect the **Seeding Target** checkbox.

   If you select the **Seeding Target** checkbox, review the information in the message dialog that appears and then click **OK**.

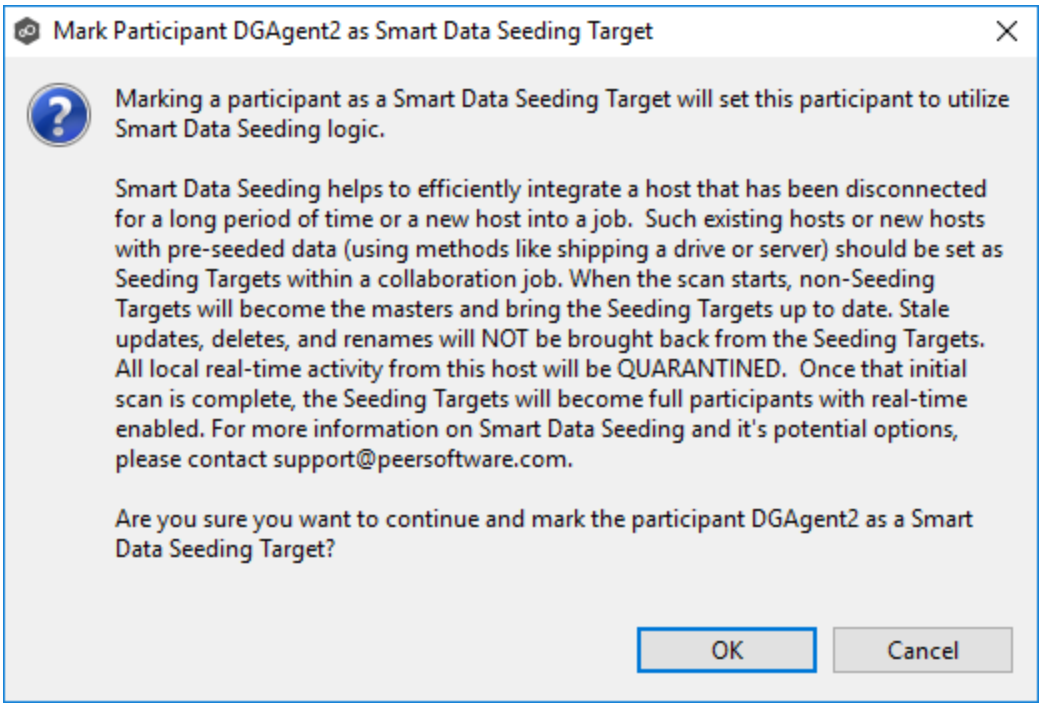

7. Click **Next** to edit Dynamic Storage Utilization options; otherwise, click **Finish**. and continue with Step 10.

8. If you click **Next**, the Dynamic Storage Utilization page appears.

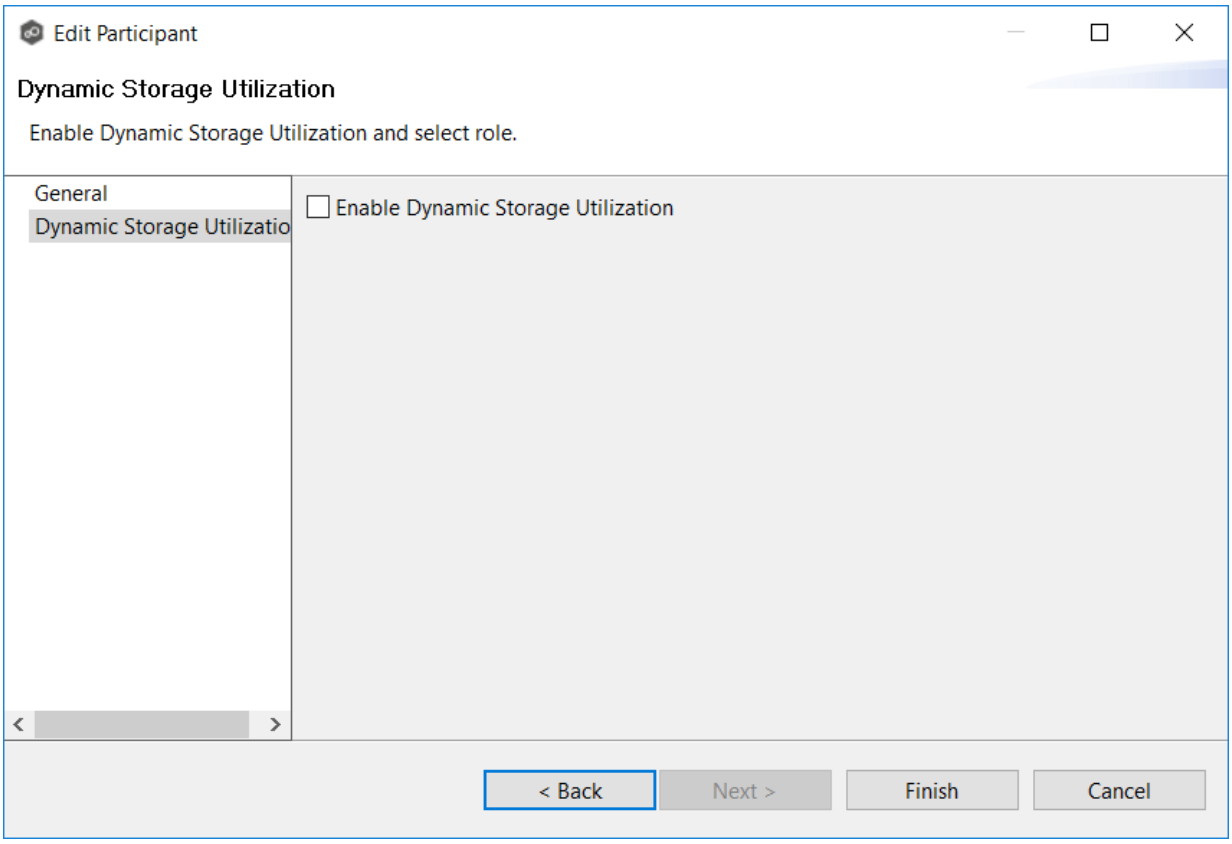

9. If you enabled Dynamic Storage Utilization, follow the steps outlined in Step 2: Dynamic Storage Utilization in Creating a File Synchronization Job.

10. Click **OK** to close the Edit wizard or select another configuration item to modify.

General

The **General** page in the **Edit File Synchronization Job** dialog presents miscellaneous settings pertaining to a File Synchronization job. You may want to consult with Peer Software's Technical Support team before modifying these values.

To modify these settings:

1. Enter the values recommended by Peer Software Support.

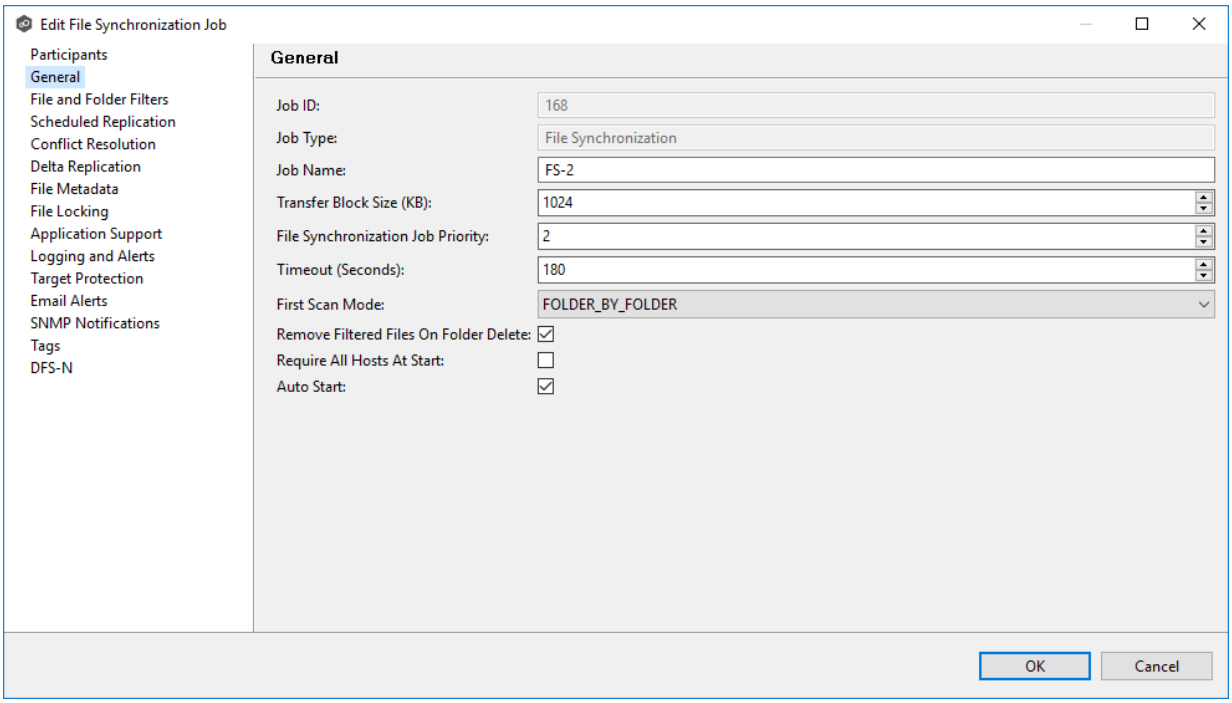

Option	Description
Job ID	Unique, system-generated job identifier that cannot be edited.
Job Type	Identifies the job type. This cannot be modified.
Job Name	Name of this File Synchronization job. This name must be unique.
Transfer Block Size (KB)	The block size in Kilobytes used to transfer files to hosts. Larger sizes will yield faster transfers on fast networks but will consume more memory in the Peer Management Broker and Peer Agents.
File Synchronization Job Priority	Use this to increase or decrease a job's file synchronization priority relative to other configured job priorities. Jobs are serviced in a round-robin fashion, and this number determines the maximum number of synchronization tasks that will be executed sequentially before yielding to another job.
Timeout (Seconds)	Number of seconds to wait for a response from any host before performing retry logic.

Option	Description
First Scan Mode	Determines which scan type will be used when the job is first started. For environments where most data is NOT seeded, the FOLDER_BY_FOLDER method will be best. For environments where most data IS seeded, the BULK_CHECKSUM method will result in a faster first scan.
Remove Filtered Files On Folder Delete	If selected, then all child files on target hosts will be deleted when its parent folder is deleted on another source host. Otherwise, filtered files will be left intact on targets when a parent folder is deleted on another source host.
Require All Hosts At Start	If selected, requires all participating hosts to be online and available at the start of the File Synchronization job in order for the job to successfully start.
Auto Start	If selected, then this file synchronization session will automatically be started when the Peer Management Center Service is started.

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

File and Folder Filters

The **File and Folder Filters** page in the **Edit File Synchronization Job** dialog displays a list of file and folder filters. A file or folder filter enables you to exclude and/or include files and folders from the job based on file type, extension, name, or directory path. Any file or folder that matches the filter is excluded or included from replication, depending on the filter's definition. By default, all files and folders selected in the **Source Paths** page will be replicated.

1. Select the file and folder filters you want to apply to the job.

   If you want to create a new file or folder filter or modify an existing one, click **Edit File Filters**. See File Filters in the Preferences section for information about creating or modifying a file filter.

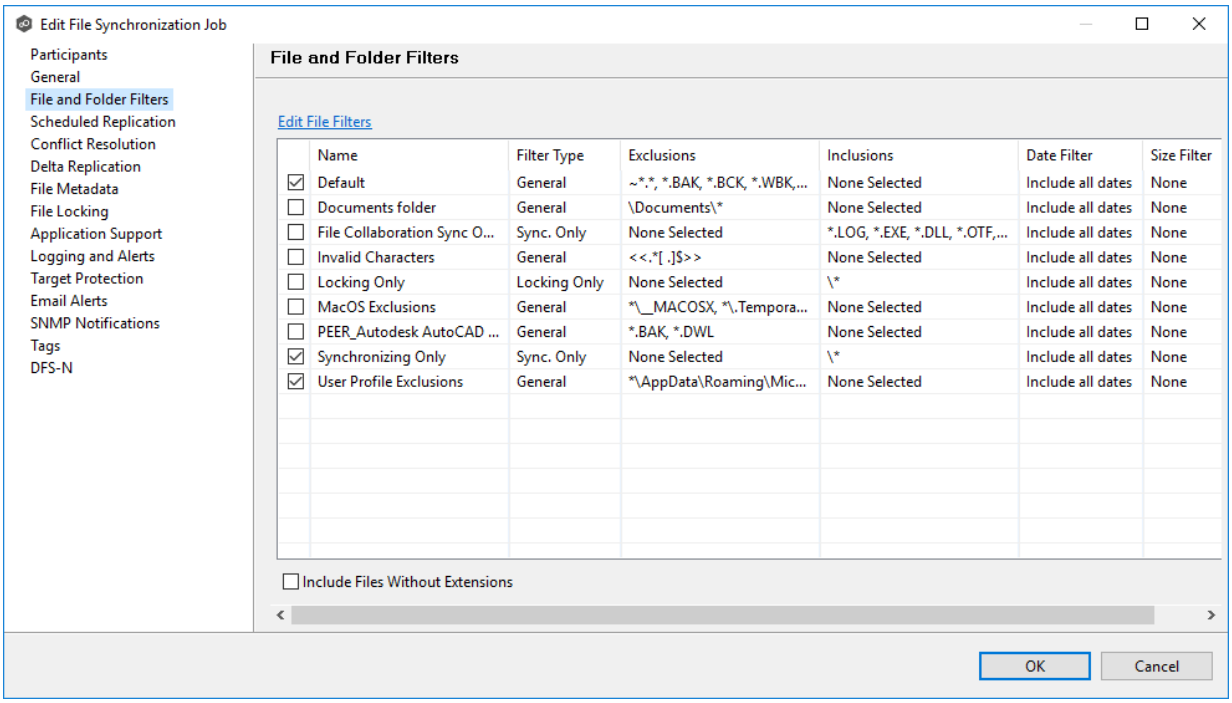

2. Select the **Include Files Without Extensions** checkbox if you want to replicate file that do not have extensions.

   Note: Files without extensions are ignored during replication unless you select this checkbox.

3. Click **OK** to close the Edit wizard or select another configuration item to modify.

Scheduled Replication

The **Scheduled Replication** page in the **Edit File Synchronization Job** dialog displays a list of scheduled replication filters. A scheduled replication filter enables you to identify files and folders that you don't want replicated in real-time.

1. Select the scheduled replication filters you want to apply to the job.

   If you want to create a new filter or modify an existing one, click **Edit Scheduled Replication Filters**. See Scheduled Replication in the Preferences section for information about creating or modifying a scheduled replication filter.

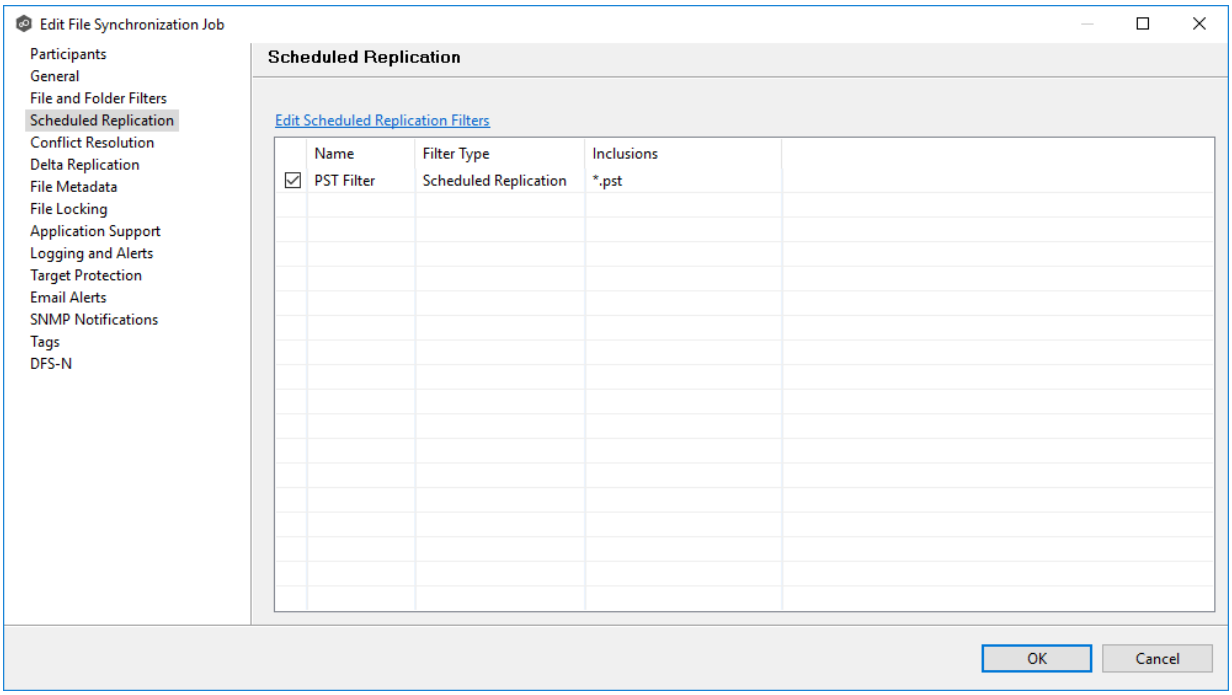

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

Conflict Resolution

By default, any file conflicts that are encountered during the initial synchronization process are automatically resolved by Peer Management Center. Peer Management Center resolves the conflict by selecting the file with the most recent modification time. Conflicts that cannot be automatically resolved result in the files being quarantined. The **Conflict Resolution** page in the **Edit File Synchronization Job** allows you to select options for resolving file conflicts and quarantines.

However, if you want to resolve the conflicts yourself, you can contact Peer Software to enable manual resolution. With manual resolution, you can select the host with the correct version of the file.

For more information about the cause of file conflicts, see Conflicts, Retries, and Quarantines.

To modify conflict resolution settings for the File Synchronization job:

1. In the **File Conflict Resolution** section, select the **Truncate Milliseconds** option if you want the millisecond value truncated from each time stamp when comparing the time stamps of a file on two or more hosts.

As some storage platforms and applications track milliseconds slightly differently, selecting this option will prevent subtle millisecond differences from causing an otherwise in-sync file to be replicated or quarantined.

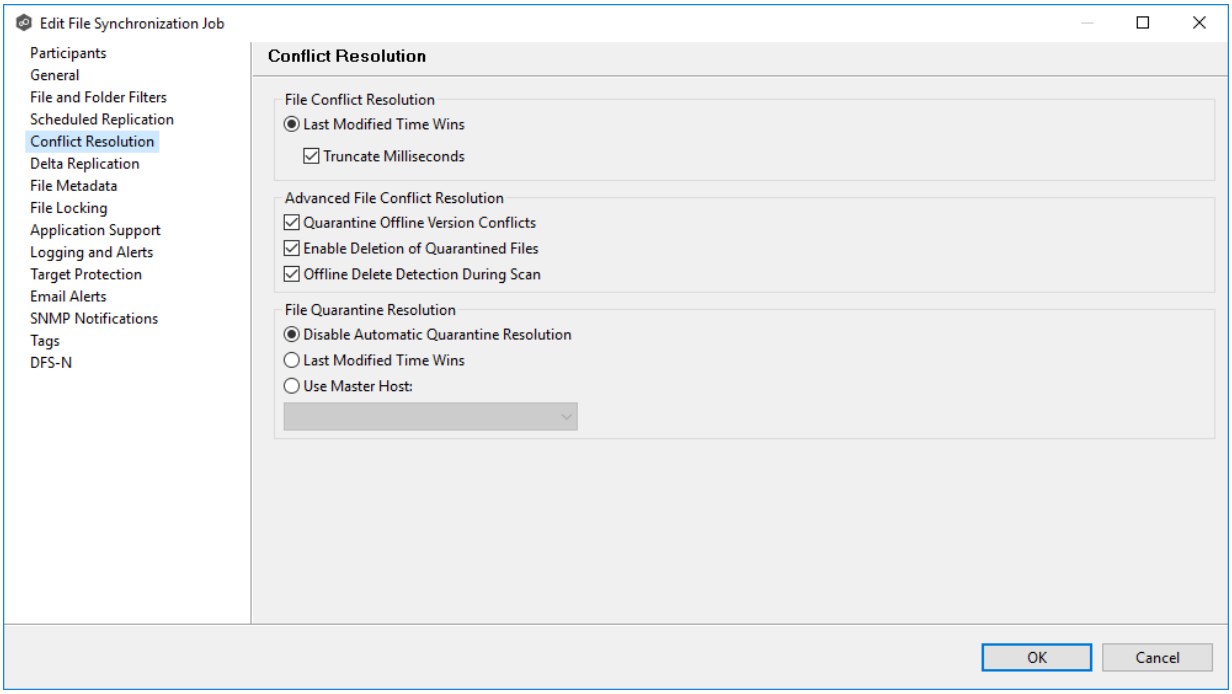

2. Select the **Advanced File Conflict Resolution** options you want applied:

Quarantine Offline Version Conflicts	Select this option if you want Peer Management Center to quarantine a file that was updated in two or more locations while the collaboration session was not running. If it is not selected, the file with the most recent last modified time will be replicated to all other participants.
Enable Deletion of Quarantined Files	Select this option if you want Peer Management Center to process a delete event for a quarantined file. If it is not selected, the quarantined file is not deleted and remains quarantined.
Offline Delete Detection During Scan	Select this option (and enabled target protection), if you want any files or folders that were deleted while the job was stopped to be deleted from all participants when the job is restarted. If it is not selected, then the deleted file or folder is restored from a participant with the file or folder to any participant where it no longer exists.

3. Select an option for automatically resolving quarantines (this option is intended to be used in environments where a single file server is active for a job):

Disable Automatic Resolution of Quarantines	Select this option if you want to manually resolve quarantines. For more information, see Removing a File from Quarantine.
Last Modified Time	Select this option if you want quarantines automatically resolved by selecting the file with the latest modification time.
Use Master Host	Select this option if you want quarantines automatically resolved by selecting the file on the Master Host.

4. Click **OK** to close the Edit wizard or select another configuration item to modify.

Delta Replication

The **Delta Replication** page in the **Edit File Synchronization Job** dialog allows you to specify the delta-replication options to use for the selected File Synchronization job. Delta-level replication is a byte replication technology that enables block/byte level synchronization for a File Synchronization job. Through this feature, Peer Management Center is able to transmit only the bytes/blocks of a file that have changed instead of transferring the entire file. This results in much lower network bandwidth utilization, which can be an enormous benefit if you are transferring files across a slow WAN or VPN, as well as across a high-volume LAN.

Delta-level replication is enabled on a per File Synchronization job basis and generally affects all files in the watch set. You will only benefit from delta-level replication for files that do not change much between file modifications, which includes most document editing programs.

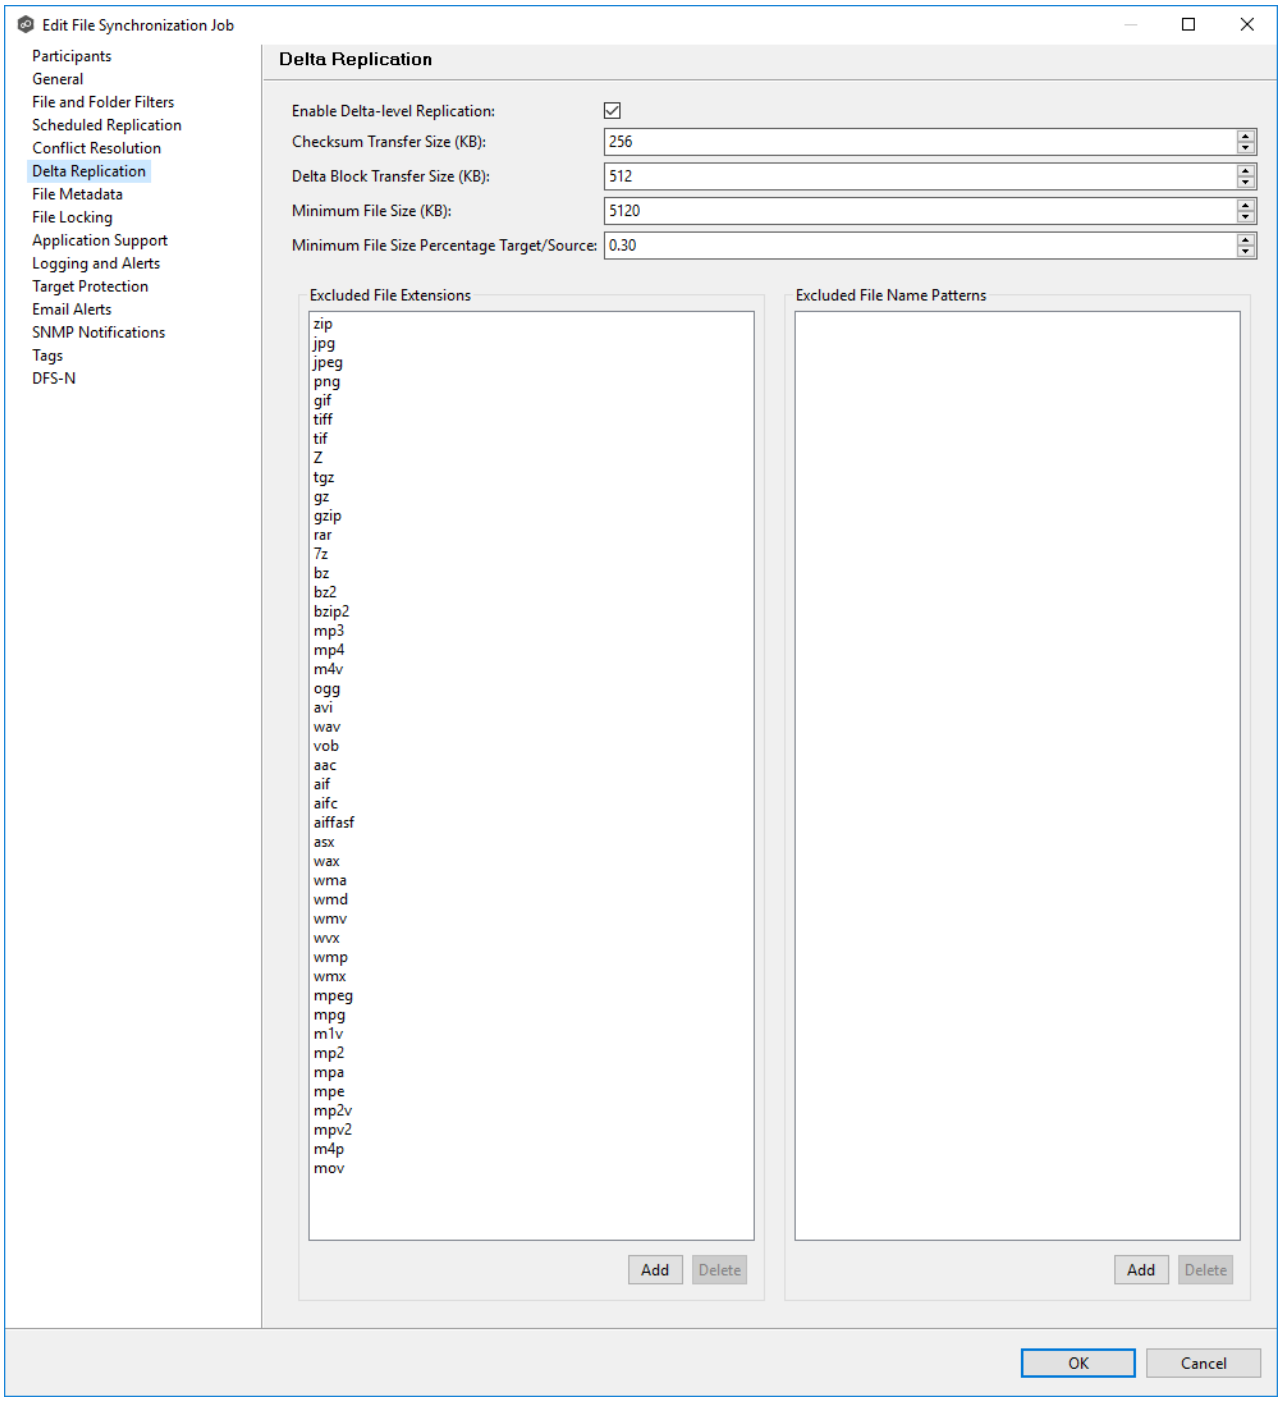

To modify delta-level replication options:

1. Modify the following the fields as necessary.

Enable Delta-Level	Select to enable delta encoded file transfers which only sends the file blocks that are different between source and

Replication	target(s). If this is disabled, the standard file copy method will be used to synchronize files.
Checksum Transfer Size (KB)	Enter the block in kilobytes used to transfer checksums from target to source at one time. Larger sizes will result in faster checksum transfer, but will consume more memory on the Peer Agents
Delta Block Transfer Size (KB)	Enter the block size in kilobytes used to transfer delta encoded data from source to target at one time. Larger sizes will result in faster overall file transfers but will consume more memory on the Peer Agents.
Minimum File Size (KB)	Enter the minimum size of files in kilobytes to perform delta encoding for. If a file is less than this size, then delta encoding will not be performed.
Minimum File Size Percentage Target/Source	Enter the minimum allowed file size difference between source and target, as a percentage, to perform delta encoding. If the target file size is less than this percentage of the source file size, then delta encoding will not be performed.
Excluded File Extensions	Enter a comma-separated list of file extension patterns to be excluded from delta encoding, e.g., zip, jpg, png. In general, compressed files should be excluded from delta encoding and the most popular compressed file formats are excluded by default.
Excluded File Name Patterns	Enter a list of file name patterns to be excluded from delta encoding. If a file name matches any pattern in this list, then it will be excluded from delta encoding transfers and a regular file transfer will be performed. See File and Folder Filters for more information on specifying wildcard expressions.

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

File Metadata

The **File Metadata** page in the **Edit File Synchronization Job** dialog allows you to modify your file metadata synchronization settings and provides some additional options not available when creating the job. See File Metadata Synchronization in Advanced Topics for more information about file metadata replication.

To enable file metadata synchronization:

1. Select when you want the metadata synchronized (you can select one or both options):

   - **Enable synchronizing NTFS security descriptors (ACLs) in real-time** - Select this option if you want the metadata replicated in real-time. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be transferred to the target host file(s) as they occur.

   - **Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan** - Select this option if you want the metadata replicated during the initial scan. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized during the initial scan.

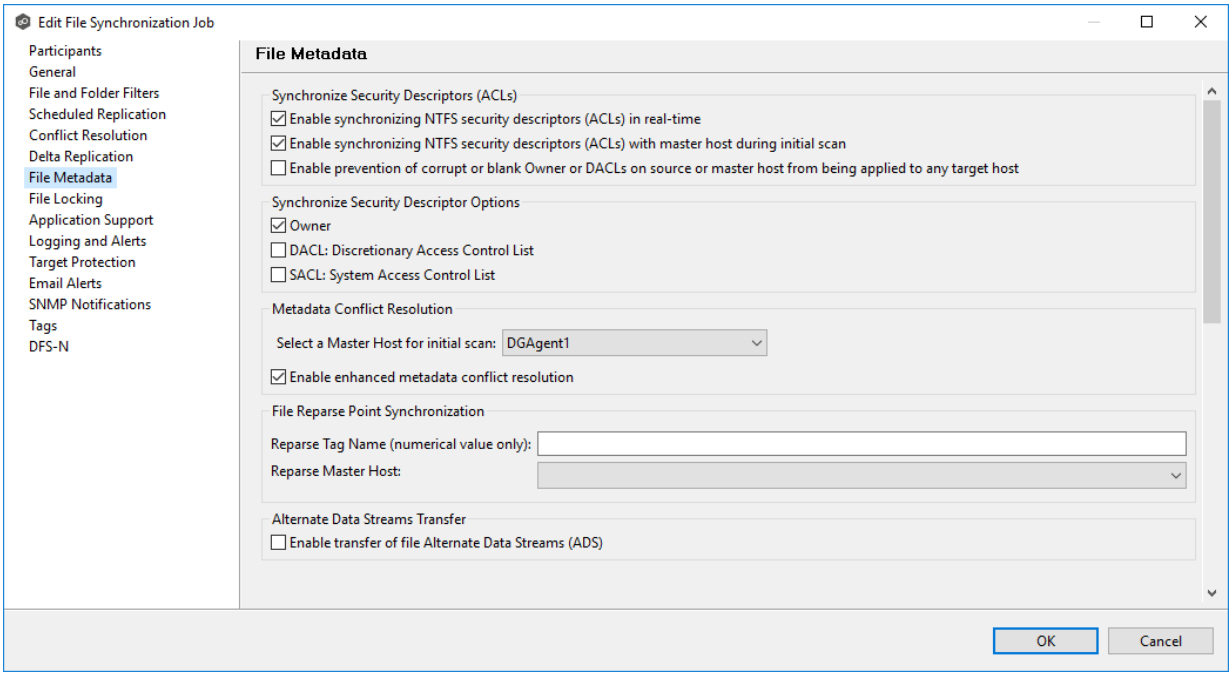

2. Click **OK** in the message that appears after selecting a metadata option.

3. If you selected either of the first two options in the **Synchronize Security Descriptor Options** section, select the security descriptor components (Owner, DACL, and SACL) to be synchronized.

   Note: To synchronize SACLs or Owner, the user that a Peer Agent service is run under on each participating host must have permission to read and write Owner and SACLs.

4. If you selected the option for metadata synchronization during the initial scan, select the host to be used as the master host in case of file metadata conflict.

   If a master host is not selected, then no metadata synchronization will be performed during the initial scan. If one or more security descriptors do not match across participants during the initial scan, conflict resolution will use permissions from the designated master host as the winner. If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

5. This option is only available when both of the first two options in the **Synchronize Security Descriptor (ACLs)** section are enabled, and **Owner** is selected under **Synchronize Security Descriptor Options.**

   If you select **Enable enhanced metadata conflict resolution**, this will prevent the Peer Agent service account from being assigned as the owner of that file or folder when a metadata conflict occurs and a file or folder is written to a target. If the Peer Agent service account were to be assigned as the owner, the user may not have permission to access the file or folder.

   Note: The Peer Agent service account cannot be a local or system administrator. As described in Peer Global File Service - Environmental Requirements, the Peer agent service account should be an actual user.

6. (Optional) Enter values for one or both file reparse point data synchronization options:

   - **Reparse Tag Name** - Enter a single numerical value. Must be either blank (if blank, reparse synchronization will be disabled) or greater than or equal to 0. The default for Symantec Enterprise Vault is 16. A value of 0 enables reparse point synchronization for all reparse file types. If you are unsure as to what value to use, then contact Peer Software technical support, or you can use a value of 0 if you are sure that you are only utilizing one vendor's reparse point functionality.

   - **Reparse Master Host** - Select a master host. If a master host is selected, then when the last modified times and file sizes match on all hosts, but the file reparse attribute differs (e.g., archived/offline versus unarchived on file server), then the file reparse data will be synchronized to match the file located on the master host. For Enterprise Vault, this should be the server where you run the archiving task on. If the value is left blank, then no reparse data synchronization will be performed, and the files will be left in their current state.

   Note: Use this option only if you are utilizing archiving or hierarchical storage solutions that make use of NTFS file reparse points to access data in a remote location, such as Symantec's Enterprise Vault. Enabling this option allows synchronization of a file's reparse data, and not the actual offline content, to target hosts, and prevents the offline file from being recalled from the remote storage device.

7. (Optional) Select the **Enable transfer of file Alternate Data Stream (ADS)** checkbox.

If enabled, Alternate Data Streams (ADS) of updated files will be transferred to the corresponding files on target participants as a post process of the normal file synchronization.

Known limitation: ADS information is transferred only when a modification on the actual file itself is detected. ADS will not be compared between participants. The updated file's ADS will be applied to the corresponding files on target participants.

8. Click **OK** to close the Edit wizard or select another configuration item to modify.

File Locking

The **File Locking** page in the **Edit File Synchronization Job** dialog presents options for managing how source and target files are locked by Peer Management Center.

To modify file locking options:

1. Modify the fields in the **Source Snapshot Synchronization** section as needed:

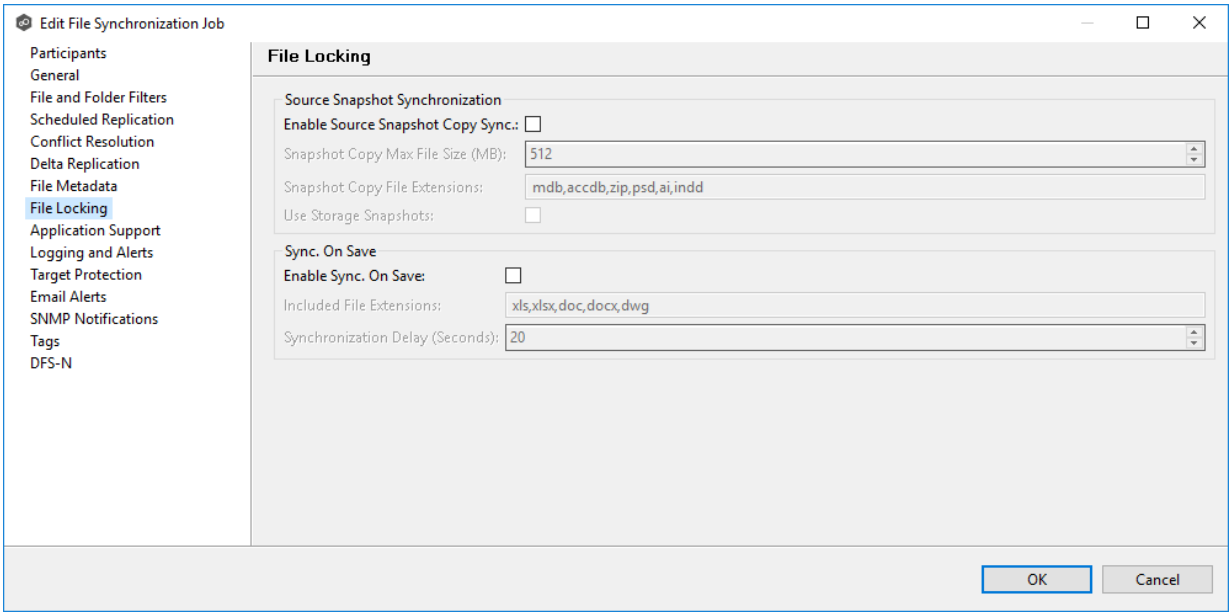

Field	Description
Enable Source	If enabled, a snapshot copy of the source file is created for files that meet the snapshot configuration criteria below, and this copy

Field	Description
Snapshot Sync.	is used for synchronization purposes. In addition, no file handle is held on the source file except while making a copy of the file.
Snapshot Copy Max File Size (MB)	The maximum file size for which source snapshot synchronization is utilized.
Snapshot Copy File Extensions	A comma-separated list of file extensions for which source snapshot synchronization is utilized.
Use Storage Snapshots	If enabled, a storage volume snapshot is created and used for synchronization purposes. As a result, no file handle is held on the source file. The snapshot is created using either VSS or storage-platform specific snapshot technologies. This option is in addition to the **Enable Source Snapshot Sync.** option above and will only apply to files with pst, mdf, ldf, and ndf extensions.

2. Modify the fields in the **Sync. on Save** section as needed.

Field	Description
Enable Sync. On Save	If enabled, this feature allows supported file types to be synchronized after a user saves a file, rather than waiting for the file to close.
Included File Extensions	A comma-separated list of file extensions for which to enable the Sync. On Save feature.
Synchronization Delay (Seconds)	The number of seconds to wait after a file has been saved before initiating a synchronization of the file.

3. Click **OK**.

Application Support

When you create a File Synchronization job, you have the option of selecting applications to be automatically optimized. When editing the job, you can modify your selections in the **Application Support** page in the **Edit File Synchronization Job** dialog.

To modify which applications are optimized:

1. Select the applications to be optimized.

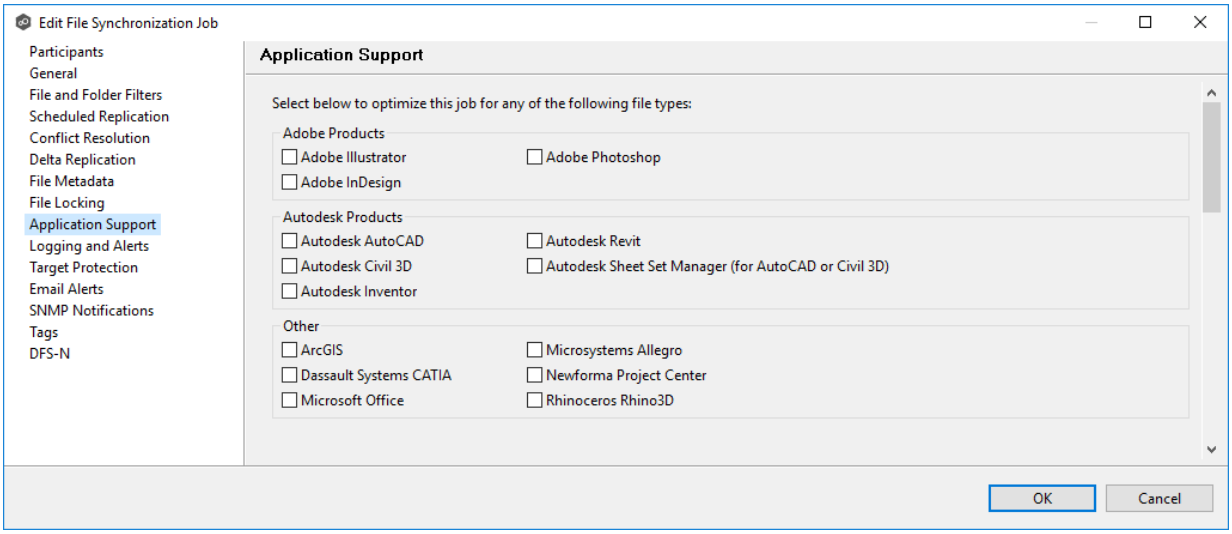

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

Logging and Alerts

Overview of File Event Logging

Various types of file synchronization events can be written to a log file and to the Event Log tab located within the File Synchronization runtime view for the selected File Synchronization job. Each job will log to the **fc_event.log** file located in the **Hub\logs** subdirectory within the Peer Management Center installation directory. All log files are stored in a tab-delimited format that can easily be read by Microsoft Excel or other database applications.

Log Entry Severity Levels

Info rmat	Informational log entry, e.g., a file was opened.

ional	
Warning	Some sort of warning occurred that did not produce an error but was unexpected or may need further investigation.
Error	An error occurred performing some type of file activity.
Fatal	A fatal error occurred that caused a host to be taken out of the session, a file to be quarantined, or a session to become invalid.

Configuration

By default, all file synchronization activity is logged for all severity levels. You can enable or disable file event logging as well as select the level of granularity.

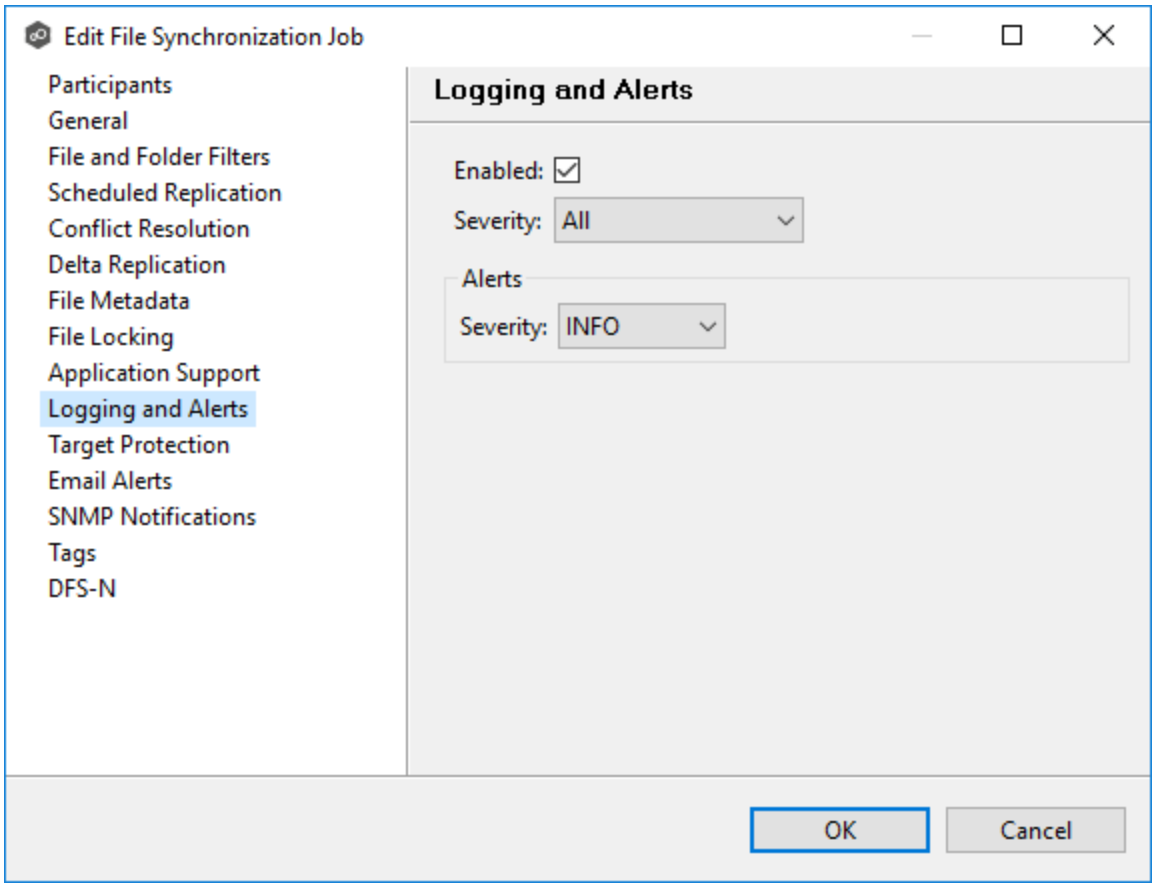

Logging Fields

Below is a list of logging fields and their descriptions:

Field	Description
Enabled	Selecting this option will enable file event logging based on the other settings. Deselecting this option will completely disable all logging.
Severity	Determines what severity levels will be logged. There are two options: • All (Informational, Warnings, Error, Fatal) • Errors & Warnings (Warnings, Error, Fatal)
Event Types	If checked, the corresponding event type will be logged.
File Open	A file was opened by a remote application on a source host.
File Lock	A file lock was acquired on a target host during synchronization of a file.
File Close	A file was closed.
File Add	A file was added to the watch set.
File Modify	A file was modified in the watch set.
File Delete	A file was deleted.
File Rename	A file was renamed.

Field	Description
Attribute Change	A file attribute was changed.
Security (ACL) Change	The security descriptor of a file or folder was changed.
Directory Scan	Indicates when a directory was scanned as a result of the initial synchronization process.
File ADS Transfer	The Alternate Data Stream of a modified file was synced to target host(s).

Alerts

Various types of alerts can be logged to a log file and to the **Alerts** table located within the File Synchronization runtime view for the selected job. Each File Synchronization job will log to the **fc_alert.log** file located in the **Hub\logs** subdirectory within the Peer Management Center installation directory. All log files are stored in a tab-delimited format that can easily be read by Microsoft Excel or other database applications.

The default log level is WARNING, which will show any warning or error alerts that occur during a running session. Depending on the severity of the alert, the job may need to be restarted.

Target Protection

Target protection is used to protect files on target hosts by saving a backup copy before a file is either deleted or overwritten on the target host. If enabled, then whenever a file is deleted or modified on the source host but before the changes are propagated to the targets, a copy of the existing file on the target is moved to the Peer Management Center trash bin.

The trash bin is located in a hidden folder named **.pc-trash_bin** found in the root directory of the watch set of the target host. A backup file is placed in the same directory hierarchy location as the source folder in the watch set within the recycle bin folder. If you need to restore a previous version of a file, you can copy the file from the trash bin into the corresponding location in the watch set and the changes will be propagated to all other collaboration hosts.

Target protection configuration is available by selecting **Target Protection** from the tree node within the Edit File Synchronization Configuration dialog.

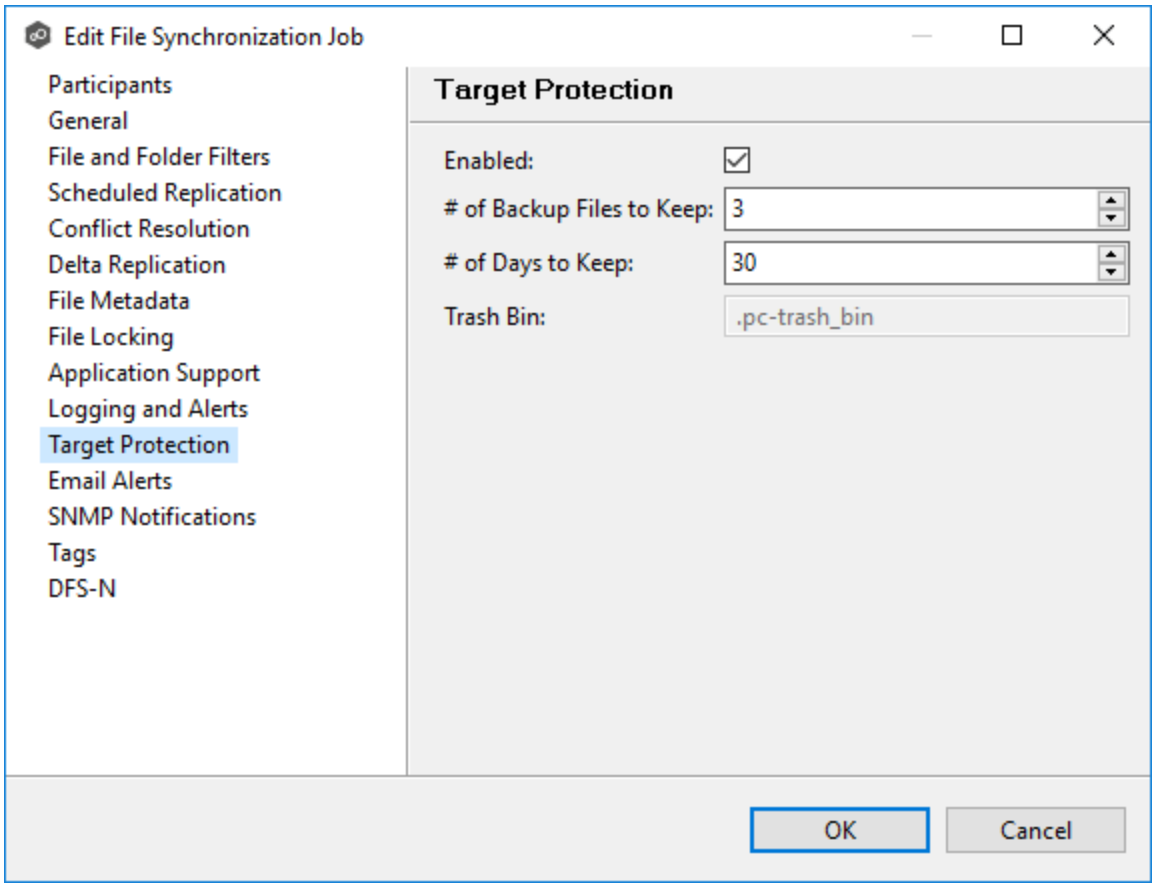

Modify the fields as needed:

Field	Description
Enabled	Enables target protection.
# of Backup Files to Keep	The maximum number of backup copies of an individual file to keep in the trash bin before purging the oldest copy.
# of Days to Keep	The number of days to keep a backup archive copy around before deleting from disk. A value of 0 will disable purging any files from archive.
Trash Bin	The trash bin folder name located in the root directory of the watch set. This is a hidden folder and the name cannot be changed by the end user.

Email Alerts

The **Email Alerts** page in the **Edit File Synchronization Job** dialog allows you to select which email alerts to apply to a File Synchronization job. Email alerts are defined in the Preferences dialog and can then be applied to individual jobs. See Email Alerts in the **Preferences** section for information about creating an email alert for a File Synchronization job.

To apply email alerts to a File Synchronization job while editing the job:

1. Click the **Select** button.

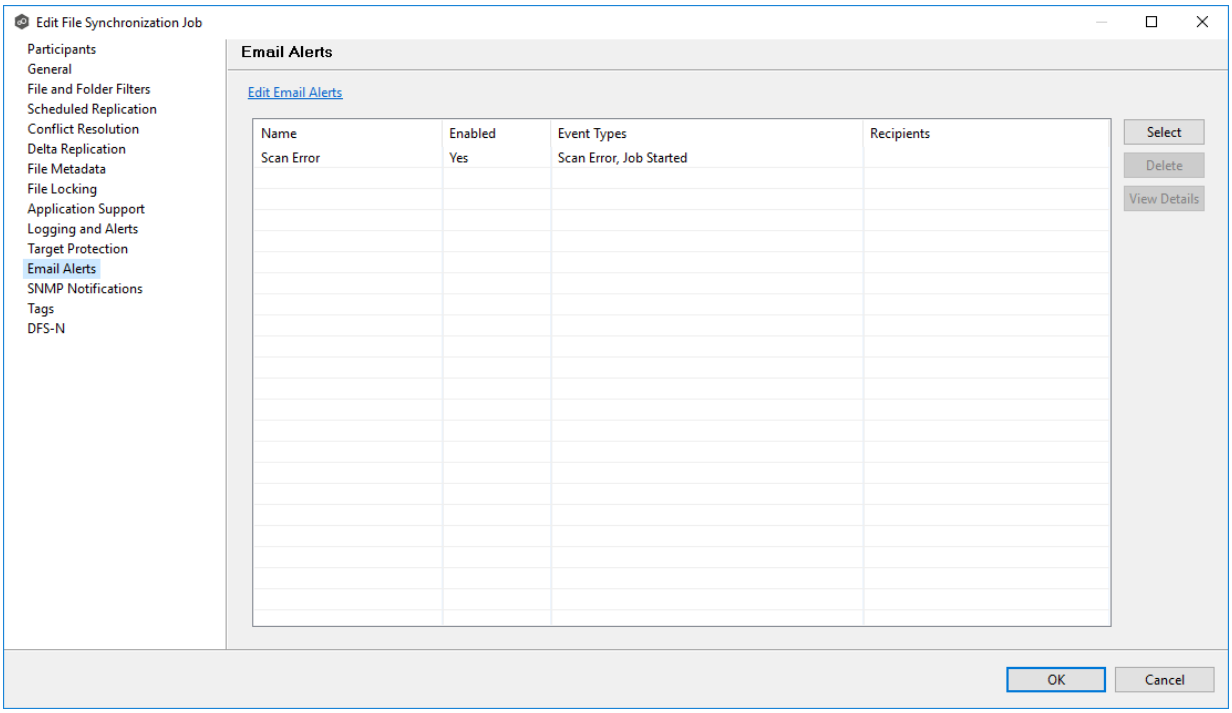

The **Select Email Alert** dialog opens.

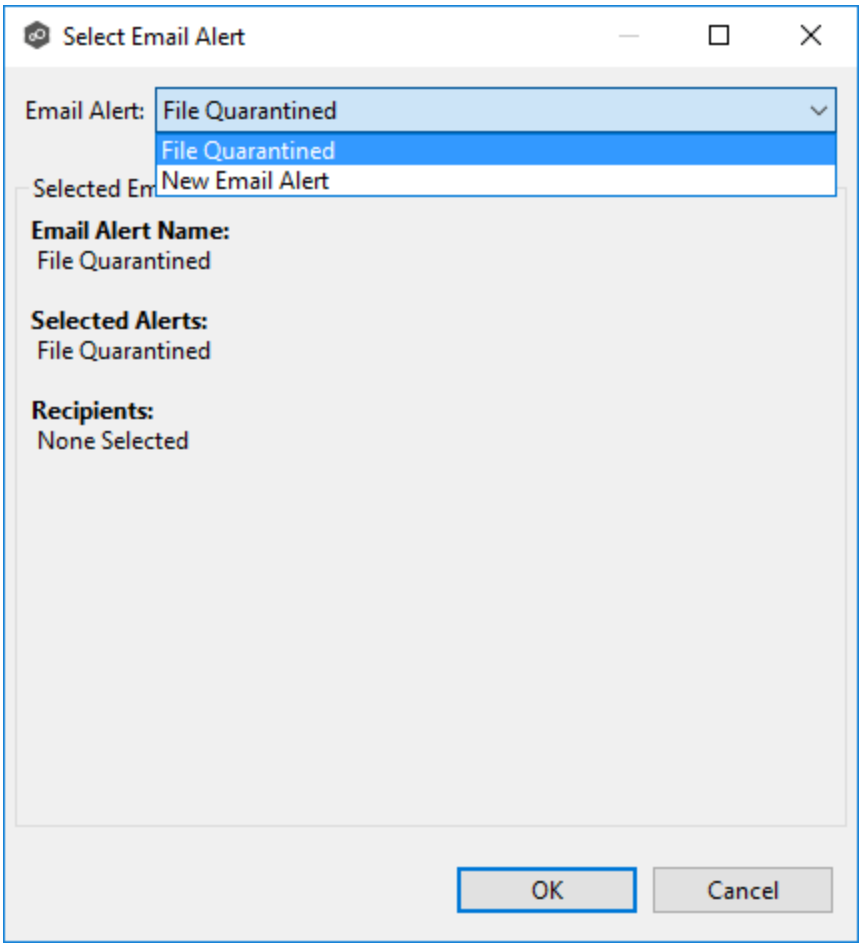

2. Select the email alert from the drop-down list, and then click **OK**.

   The newly added email alert appears in the **Email Alerts** table.

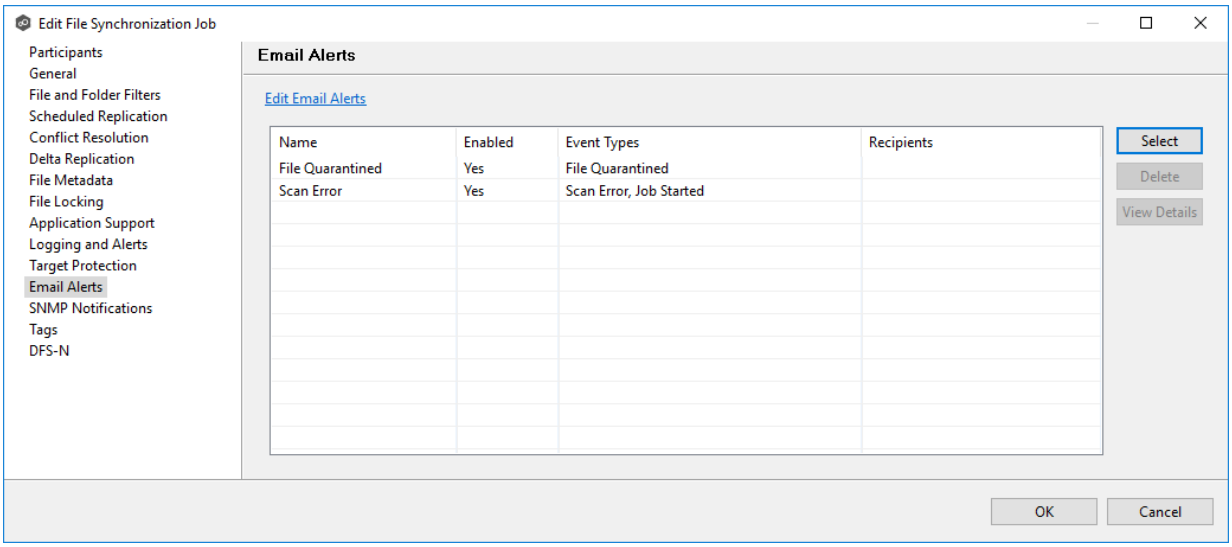

3. Repeat to add additional alerts to the job.

4. Click **OK** to close the Edit wizard or select another configuration item to modify.

SNMP Notifications

The **SNMP Notifications** page in the **Edit File Synchronization Job** dialog allows you to select which SNMP notification to apply to a File Synchronization job.

SNMP notifications, like email alerts and file filters, are configured at a global level in the Preferences dialog, then applied to individual jobs. For more information about SMNP Notifications, see SNMP Notifications in the **Preferences** section.

To enable or disable SNMP notifications for a File Synchronization job:

1. To enable, select an SNMP notification from the drop-down list.

   To disable, select **None - Disabled**.

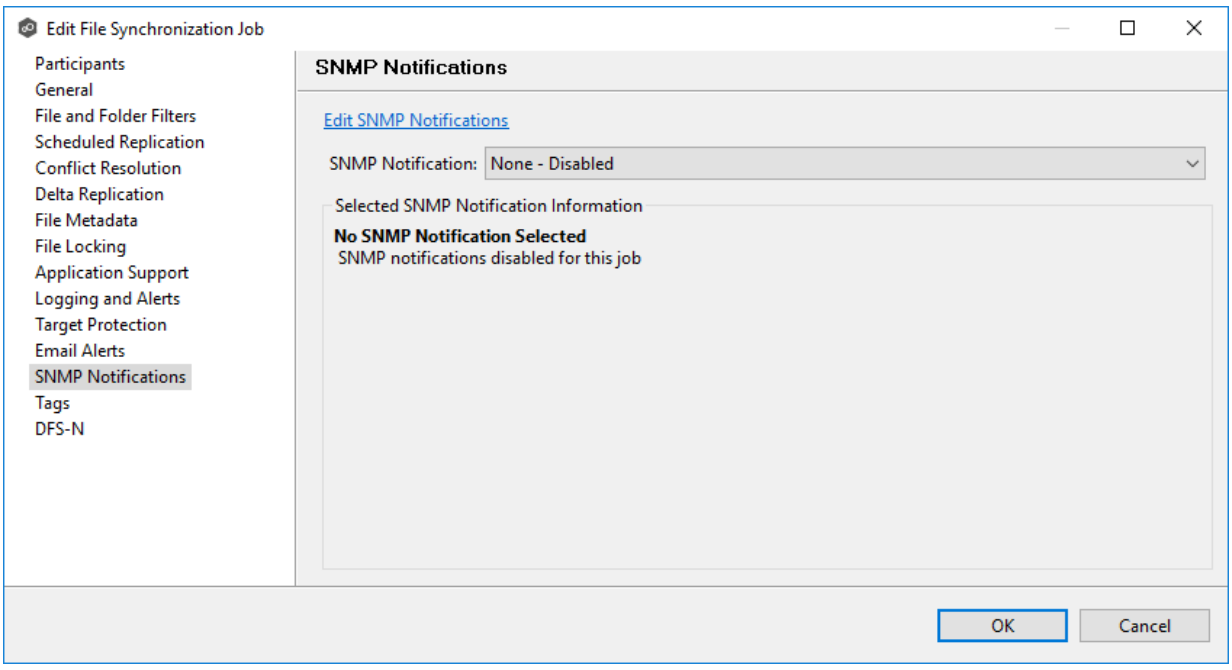

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

Tags

The **Tags** page in the **Edit File Synchronization Jobs** dialog allows you to assign existing tags and categories to the selected job. This page is not available in Multi-Job Editing mode. For more information about tags, see Tags in the Basic Concepts section.

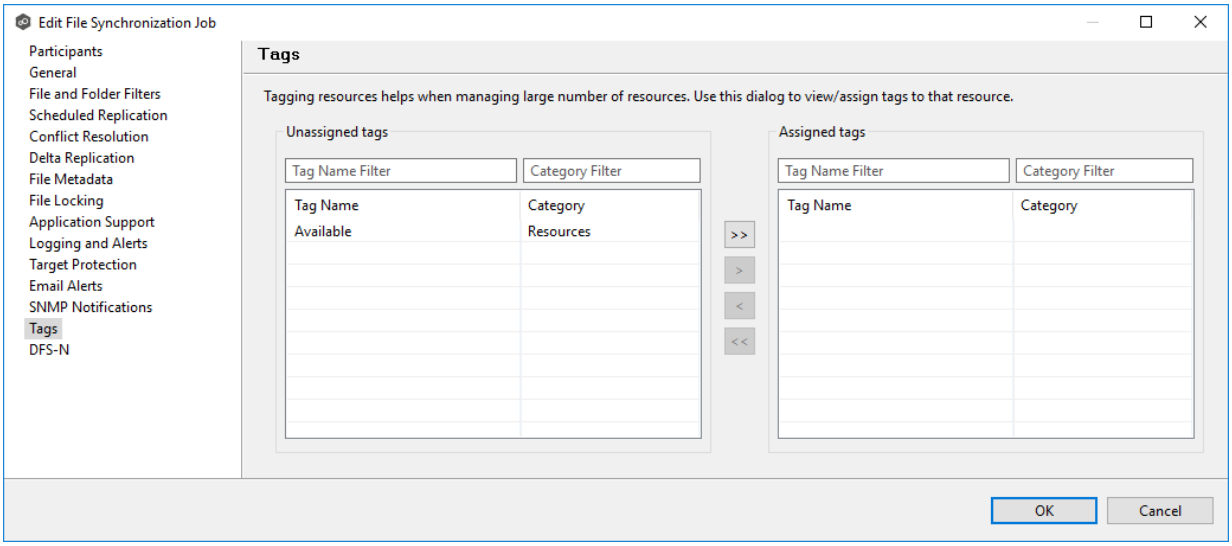

DFS-N

The **DFS-N** page in the **Edit File Synchronization Job** dialog presents options for linking a DFS namespace folder to this job. See Linking a Namespace Folder to an Existing File Collaboration or Synchronization Job for more information.

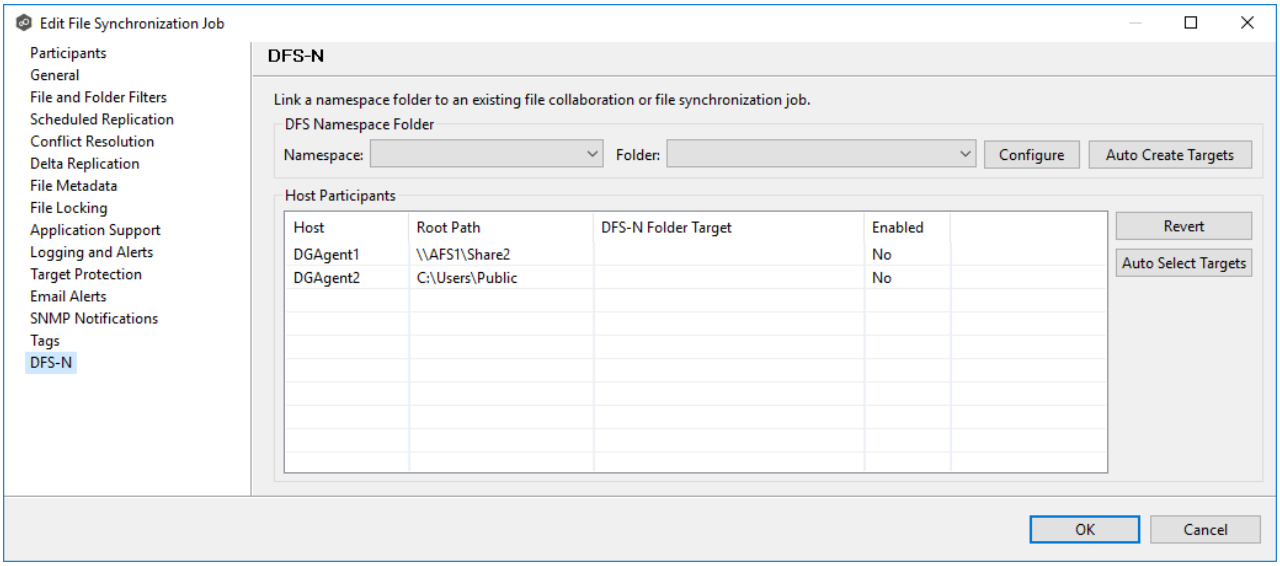

Editing Multiple Jobs

Peer Management Center supports multi-job editing, allowing you to quickly and effectively manipulate multiple File Synchronization jobs simultaneously. For example, you can use this feature to change a single configuration item such as **Auto Start** for any number of already configured jobs in one operation instead of having to change the item individually for each.

While this feature does cover most of the options available on a per-job basis, certain options are unavailable in multi-job edit mode, specifically ones tied to participants. Configuration of participants must be performed on a per job basis.

To edit multiple jobs simultaneously:

1. Open Peer Management Center.

2. Select the jobs you want to edit in the **Jobs** view.

3. Right-click and select **Edit Jobs**.

   For the most part, the original configuration dialog remains the same with a few minor differences depending on similarities between the selected File Synchronization jobs. A sample dialog is as follows:

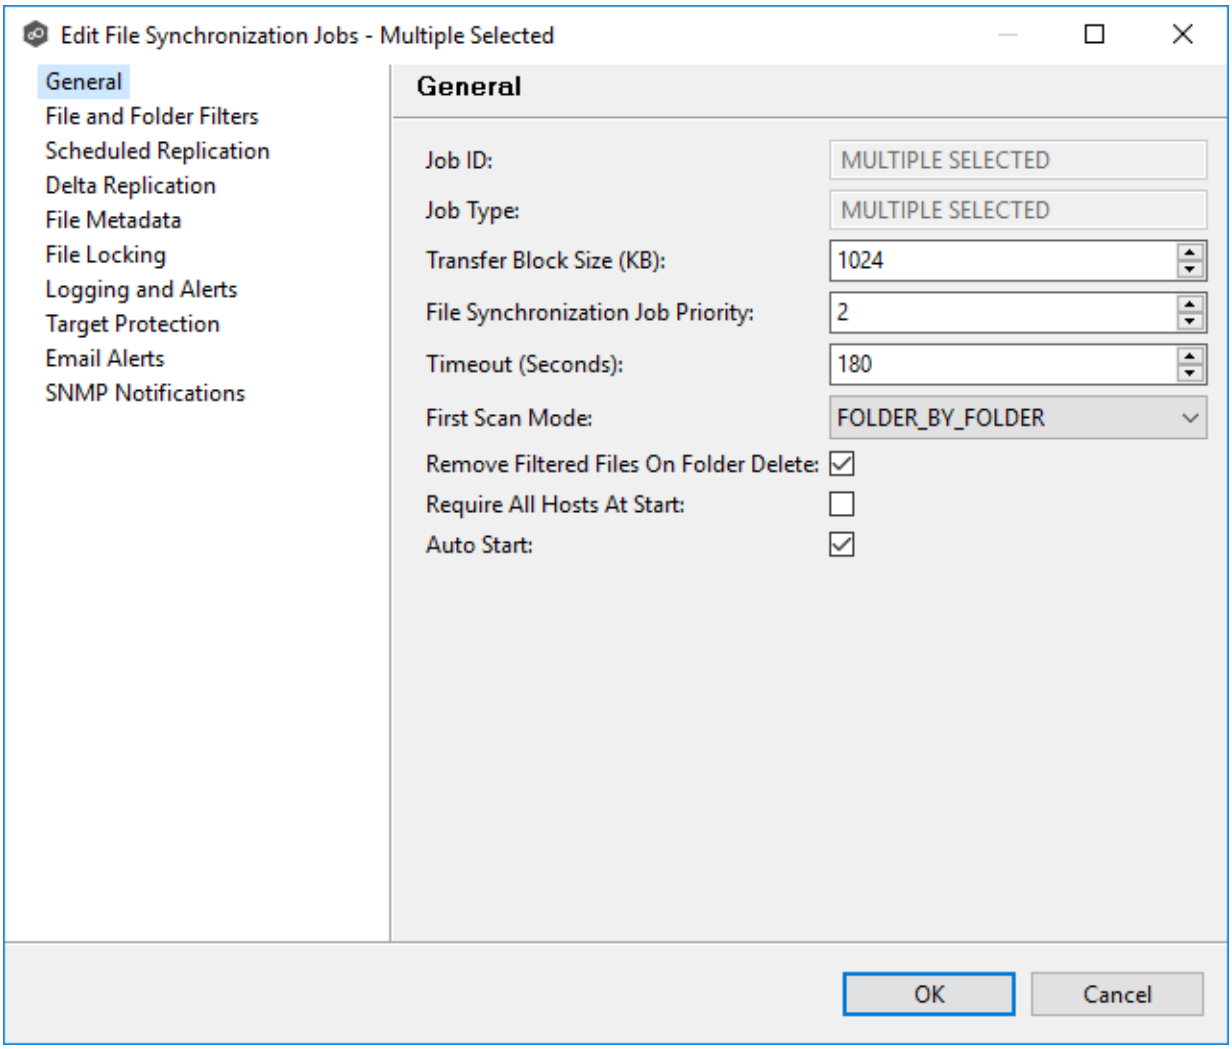

In this dialog, any discrepancies between multiple selected jobs will generally be illustrated by a read-only text field with the caption **Multiple Values - Click to Edit**. Clicking this field will bring up a dialog similar to the following:

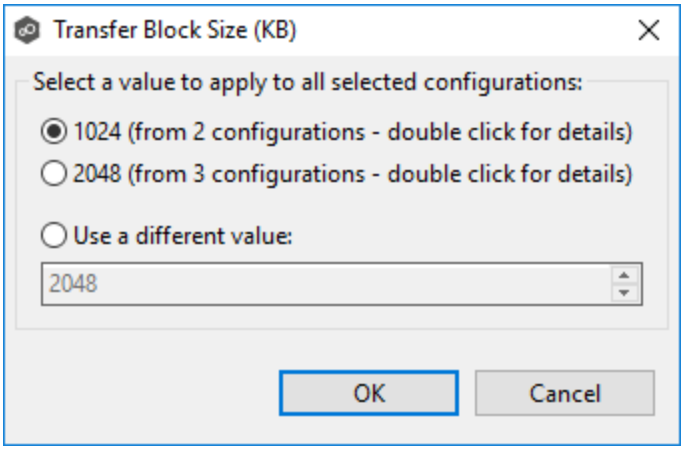

This dialog gives you the option of choosing a value that is already used by one or more selected File Synchronization jobs, in addition to the ability to use your own value. Notice that variances in the look and feel of this pop-up dialog above depend on the type of information it is trying to represent (for example, text vs. a checkbox vs. a list of items).

Upon clicking OK, the read-only text field you originally clicked will be updated to reflect the new value. Any fields that have changed will be marked by a small caution sign. On saving in this multi-job edit dialog, the changed values will be applied to all selected jobs.

Note: Read all information on each configuration page carefully when using the multi-job edit dialog. A few pages operate in a slightly different manner then mentioned above. All of the necessary information is provided at the top of these pages in bold text.

Running and Managing a File Synchronization Job

The topics in this section provide some basic information about starting, stopping, and managing File Synchronization jobs:

- Overview
- Starting a File Synchronization Job
- Starting a File Synchronization Job
- Auto-Restarting a File Synchronization Job
- Host Connectivity Issues
- Removing a File from Quarantine
- Manual Retries

Overview

This topic describes:

- The initialization process for a File Synchronization job: What occurs the first time you run a File Synchronization job.

- The initial synchronization process: How files are synchronized the first time you run a File Synchronization job.

The initialization process for a File Synchronization job consists of the following steps:

1. All participating hosts are contacted to make sure they are online and properly configured.

2. Real-time event detection is initialized on all participating hosts where file locks and changes will be propagated in real-time to all participating hosts. You can view real-time activity and history via the various Runtime views for the open job.

3. The initial synchronization process is started; all of the configured root folders on the participating hosts are scanned in the background, and a listing of all folders and files are sent back to the running job.

4. The background directory scan results are analyzed and directory structures compared to see which files are missing from which hosts. In addition, file conflict resolution is performed to decide which copy to use as the master for any detected file conflicts based on the File Conflict Resolution settings.

5. After the analysis is performed, all files that need to be synchronized are copied to the pertinent host(s).

Before you start a File Synchronization job for the first time, you need to decide how you would like the initial synchronization to be performed. During the initial synchronization process:

- The watch set is recursively scanned on all participating hosts.

- File conflict resolution is performed.

- Any files that require updating are synchronized with the most current copy of the file.

The two primary options are:

- Have the File Synchronization job perform the initial synchronization based on the Conflict Resolution settings.

- Pre-seed all participating hosts with the correct folder and file hierarchy for the configured root folders before starting the session.

If you have a large data set, we strongly recommend that you perform the initial synchronization manually by copying the data from a host with the most current copy to all other participating hosts. This needs to be done only once--before the first time that you run the File Synchronization job.

If you choose the first option, click the **Start** button to begin synchronization session initialization. Otherwise, pre-seed each participating host with the necessary data, then click the **Start** button.

Starting a File Synchronization Job

Before starting a File Synchronization job for the first time, make sure that you have decided how you want the initial synchronization to be performed.

When running a File Synchronization job for the first time, you must manually start it. After the initial run, a job will automatically start, even when the Peer Management Center server is rebooted.

Note: You cannot run two jobs concurrently on the same volume if the watch sets contain an overlapping set of files and folders.

To manually start a job:

1. Choose one of three options:

    - Right-click the job name in the **Jobs** view.

    - Right-click the job name in the **Summary** tab of the **Collab, Sync, and Repl Summary** summary view, and then choose **Start** from the pop-up menu.

    - Open a job and then click the **Start/Stop** button to the left of the **Status** field in the bottom left corner of the job's runtime view (shown below).

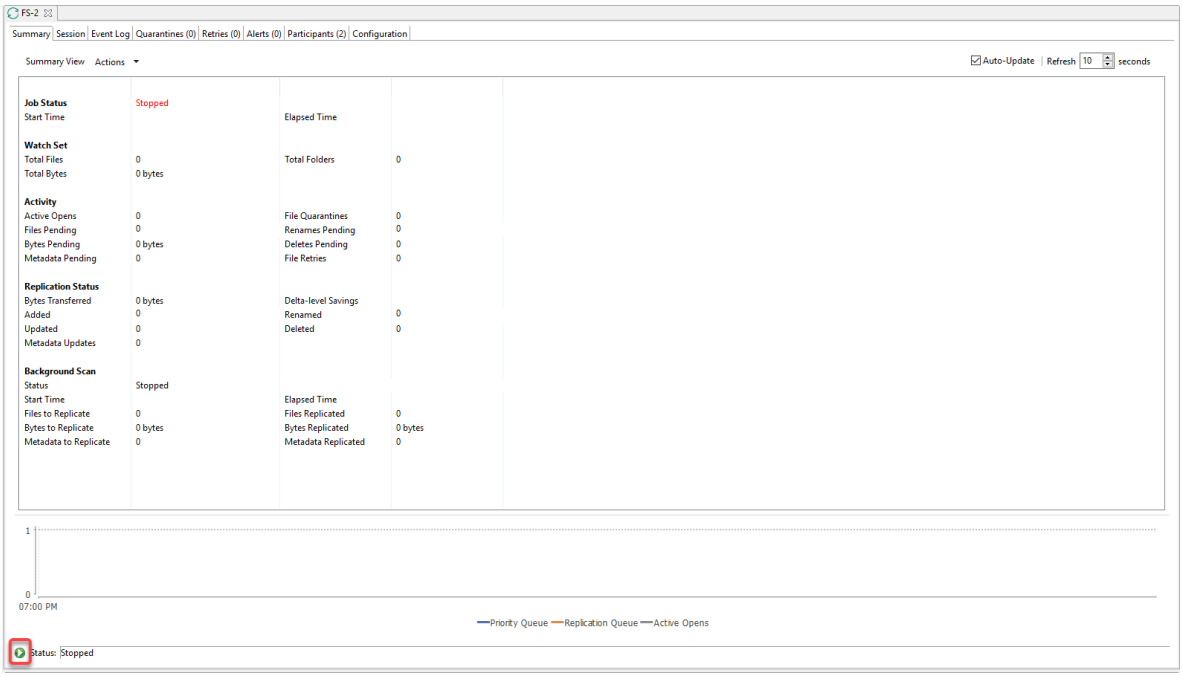

2. Click **Yes** in the confirmation dialog.

   After the job initialization has completed, the job will run. Once the job starts, the icon next to the job name in the **Jobs** view changes from gray to green.

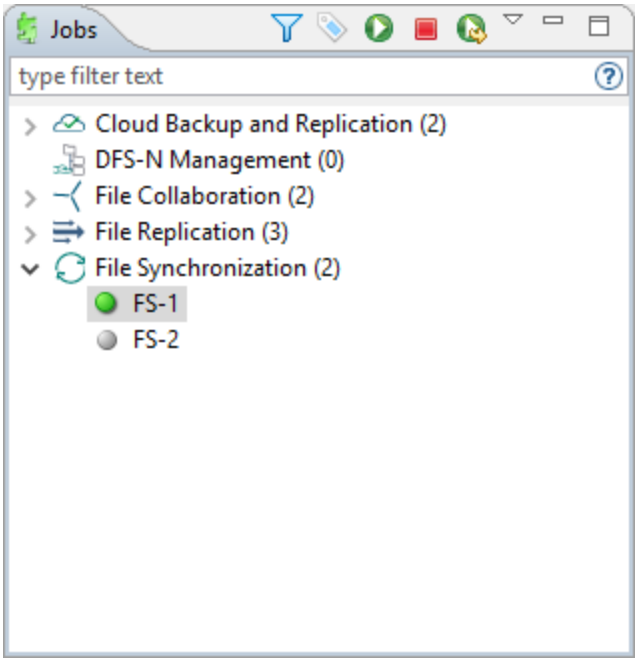

Stopping a File Synchronization Job

You can stop a File Synchronization job at any time by clicking the **Stop** button on the **Jobs** view toolbar. Doing this shuts down the real-time file event detection and close all running operations (e.g., file transfers).

Auto-Restarting a File Synchronization Job

Peer Management Center includes support for automatically restarting File Synchronization jobs that include participating hosts that have been disconnected, have reconnected, and are once again available.

After a host becomes unavailable and the quorum is lost on a running File Synchronization job, the job automatically stops running and enters a pending state, waiting for one or more hosts to become available again so that the quorum can be met. Once the quorum is met, the pending job will automatically be restarted, beginning with a scan of all root folders.

In a job where a host becomes unavailable but quorum is not lost, the remaining hosts will continue synchronizing. If the unavailable host becomes available once again, it is brought back into the running job and a background scan begins on all participating hosts, similar in fashion to the initial background scan at the start of a job.

You can enable all File Synchronization jobs to auto-restart. You can also disable auto-restart File Synchronization jobs on a per-job and host instance.

To enable all File Synchronization jobs to auto-restart:

1. Select **Preferences** from the **Window** menu.
2. Select **Collab, Sync, and Repl Summary** in the navigation tree.

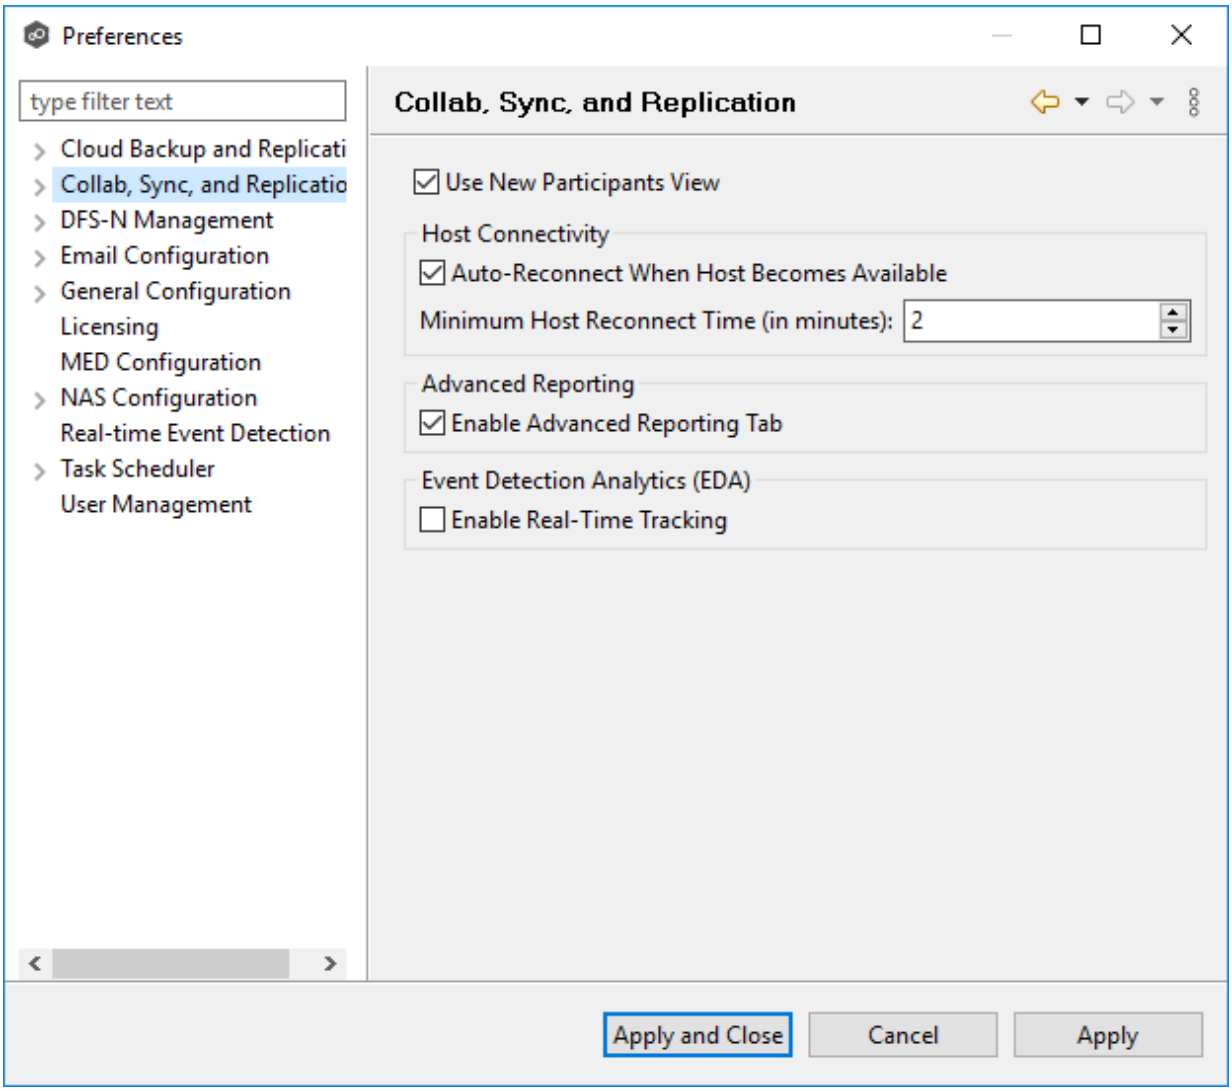

3. Select the **Auto Reconnect when Host Becomes Available** checkbox.

4. Enter the minimum number of minutes to wait after an Agent reconnects before re-enabling it in any associated jobs in the **Minimum Host Reconnect Time** field.

5. Click **OK**.

Host Connectivity Issues

Peer Management Center is designed to be run in an environment where all participating hosts are highly available and on highly available networks. The two primary connectivity issues result from:

- Unavailable Hosts
- Quorum Not Met

Unavailable Hosts

If a host becomes unavailable while a File Synchronization job is running and is unreachable within the configured timeout period (specified in the job's General settings), it may be removed from synchronization. If no response is received while performing a file synchronization operation within the timeout period, Peer Management Center pings the host; if still no response, the host is taken out of the running session, a FATAL event is logged, and the **Participants** tab for the job is updated to indicate that the host has failed. In addition, if email alerts and/or SNMP notifications are configured and enabled for **Host Timeouts**, then the appropriate message(s) are sent.

If auto-restart not enabled, you must stop and start the File Synchronization job to bring any failed hosts back into the session. As a result, all root folders on all hosts will need to be scanned again to detect any inconsistencies. Therefore, if you are operating over a WAN with low bandwidth you will want to set the timeout to a higher value on each related job.

Quorum Not Met

For a File Synchronization job to run correctly, a quorum of available hosts must be met. Quorum is currently set to at least 2 hosts, and if quorum is not met, then the synchronization session is automatically be terminated. If email alerts and/or SNMP notifications are configured and enabled for **Session Aborts**, then the appropriate message(s) are sent.

Removing a File from Quarantine

Quarantines are a key feature of Peer Global File Service, used for resolving version conflicts. For more detailed information on how quarantines work, see Conflicts, Retries, and Quarantines in the Advanced Topics section.

You must explicitly remove a file from quarantine in order to have it participate in the synchronization session once again.

You may also choose to perform no action, in which case, the file is removed from the **Quarantines** table but none of the file versions are modified. Therefore, if the files are not currently in-sync, then the next time the file is modified, changes will be propagated to the other hosts. If an error occurs while removing the quarantine, then the **Status** field in the **Quarantines** table is updated to reflect the error.

To remove a file (or multiple files) from quarantine:

1. Open the job.
2. Open the **Quarantines** tab.
3. Select the file(s) in the Quarantines table.
4. Select the host with the correct version.
5. Click the **Release Conflict** button.

   After doing this, all hosts are checked to make sure the file is not currently locked by anyone. If no locks are found, then locks are obtained on all versions of the file and the targets that are out-of-date are synchronized with the selected source host.

Manual Retries

Retries are a key feature of Peer Global File Service, used for automatically handling errors in the collaboration environment that would have otherwise led to a quarantine. For more detailed information on how retries work, see Conflicts, Retries, and Quarantines in the Advanced Topics section.

When a file is put in the retry list, it will be automatically retried based on the settings defined in File Retries in Preferences. If need be, you can also manually force the retry of a file. This can be done from the Retries list of a specific File Synchronization job.

You may also choose to perform no action, in which case, the file is removed from the Retries list but none of the file versions are modified. Therefore, if the files are not currently in-sync, then the next time the file is modified, changes will be propagated to the other hosts. If an error occurs while forcing the retry, then the **Status** field in the **Retries** table is updated to reflect the error.

To manually force the retry of a file (or multiple files):

1. Select the file(s) in the **Retries** list.
2. Select the host with the correct version.
3. Click the **Release Conflict** button.

   After doing this, all hosts are checked to make sure the file is not currently locked by anyone. If no locks are found, then locks are obtained on all versions of the file and the targets that are out-of-date are synchronized with the selected source host.

PeerSync Management Jobs

This section provides information about creating, editing, running, and managing a PeerSync Management job:

- Creating a PeerSync Management Job
- Running and Managing a PeerSync Management Job

Creating a PeerSync Management Job

The topics in this section provide some basic information about creating and editing PeerSync Management jobs.

Integrating Existing PeerSync Instances

To integrate existing PeerSync instances in Peer Management Center, follow the step-by-step instructions.

Creating and Deploying New PeerSync Instances

To create a new job and deploy the PeerSync installation to one or more hosts, click the **Create New** button in toolbar of Peer Management Center or select **New** from the **File** menu. A list of all installed job types will be displayed. Select the PeerSync Management option to open the PeerSync Management Configuration dialog. Go to the Step-by-Step instructions for more information.

When configuring Alerts, you will want to configure global settings like SMTP configuration, which is specific to Peer Management Center. Details on what and how to configure these global options can be found in the Before You Create Your First PeerSync Management Job section.

To edit the PeerSync Management configuration, right-click on the job in the Jobs view and select **Edit Job(s).** Within the **Edit PeerSync Management Job** dialog, select the **Associated Profile** node from the left. For step-by-step instructions, see Running and Managing PeerSync Management Jobs.

- Integrating Existing PeerSync Instances
- Deploying New PeerSync Instances

Before You Create Your First PeerSync Management Job

Before creating your first PeerSync Management job, we highly recommend preconfiguring a number of global options that can be applied to all PeerSync Management jobs.

The following configuration items are not always required, but highly recommended:

- Email Configuration
- Email Alerts

Overview

The Peer Management Center supports the concept of email alerts, where a single alert (consisting of a unique name, a selection of event types along with a list of email addresses) can be applied to multiple file synchronization jobs without requiring repeat entry for each job. When an email alert is applied to a job, an email is sent to all listed recipients anytime a selected event type is triggered by that job.

To manage email alerts, right-click any file synchronization job from the Jobs view and select the **Email Alerts** node from the **Monitoring** node. Click **Edit PeerSync Email Alerts**. The following screen represents the list of defined email alerts, along with buttons to add new ones and edit, copy, and remove existing ones.

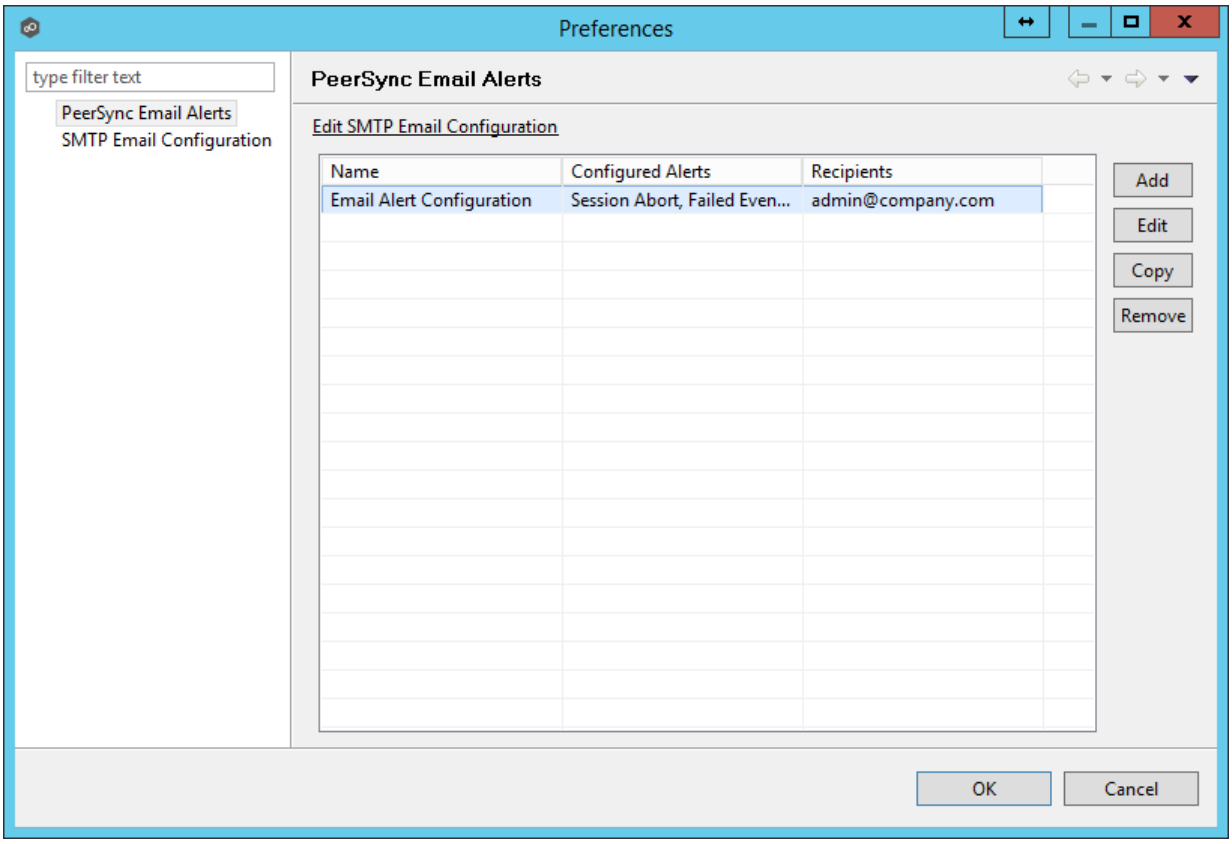

Upon adding or editing an email alert, the following dialog is displayed:

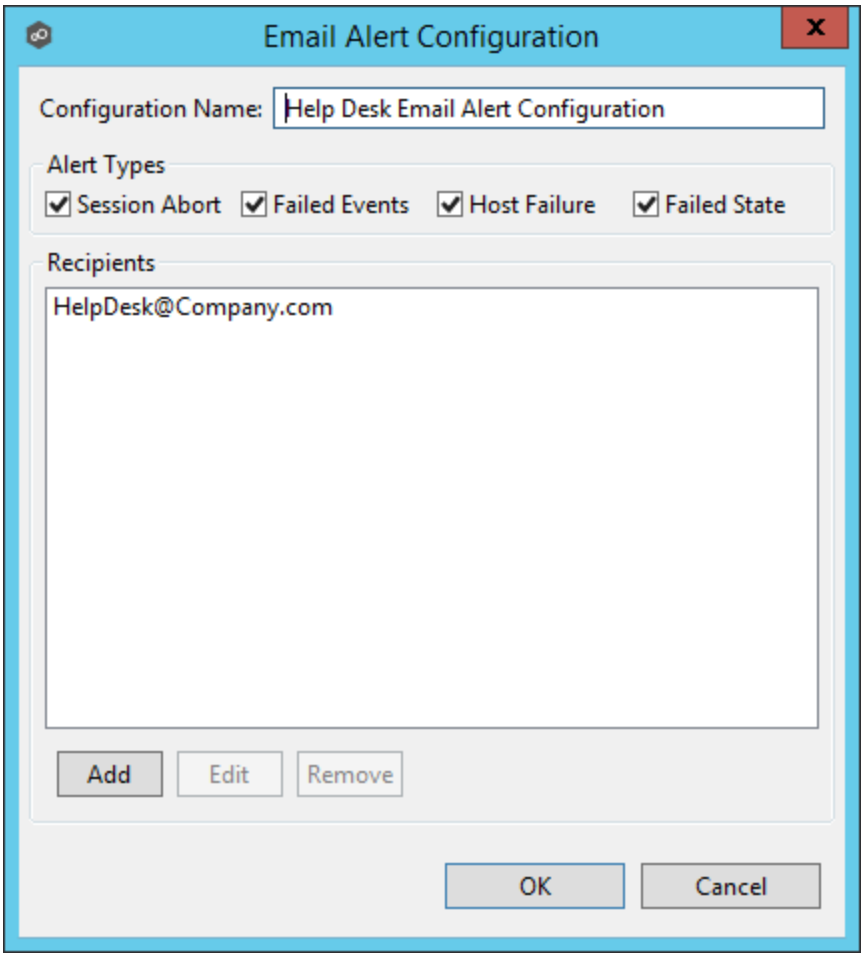

Within this dialog, you can select specific event triggers on which an email will be generated and configure the list of email recipients of the alert(s). Event types are defined below.

Event Types

Session Abort	Enables sending an alert when a session is aborted because of lack of quorum due to one or more failed host agents.
Failed Events	Enables sending an alert when a failed event is received from the PeerSync machine.
Host Failure	Enables sending an alert when a host agent timeout occurs or a PeerSync service timeout occurs.

Failed State	Enables sending an alert when the state of the File Synchronization machine changes from Active to "Failed State" indicating that either a failed scan or failed event was detected in the latest set of synchronization stats.

Integrating Existing PeerSync Instances

The topics in this section provide some basic information on how to integrate existing PeerSync instances within the Peer Management Center.

- Requirements
- How to Integrate Existing PeerSync Instances

- PeerSync must be installed as a Service and running version 9.3.0 or newer.
- Peer Agent must be installed on the PeerSync machine and connected to the Peer Management Center.

1. Open the profile on the PeerSync machine with the PeerSync Profiler.
2. Add the argument /LZTAI in Options/Command section.
3. Save the profile.
4. Restart the PeerSync Service.
5. Install the Peer Agent.
6. Start the Peer Agent.

Once the Peer Agent is started and connected to the Peer Management Center, PeerSync will be auto-detected, and a Peer Management Center file synchronization job will be generated with the name of the machine.

Optionally you can edit the job and add email alerts and save and restart the File Synchronization job for changes to take effect.

Deploying New PeerSync Instances

The topics in this section provide basic information on how to integrate existing PeerSync instances within the Peer Management Center:

- Requirements
- How To

- Peer Agent must be installed on the machine where PeerSync will be deployed to.
- It is recommended to run the Agent under a domain admin account or account with enough rights to modify registry and service configuration.

The topics in this section provide step-by-step instructions on how to create and deploy new PeerSync instances through Peer Management Center.

- Step 1: General Information
- Step 2: PeerSync Profile
- Step 3: Jobs Configuration List
- Step 4: Installation Settings

Step 1: General Information

Create a new PeerSync Management job by clicking the **Create New** button in the toolbar of the Peer Management Center, or by selecting the **New** from the **File** menu. A list of all available job types will be displayed. Selecting the **PeerSync Management** option will open the PeerSync Management **Configuration** dialog.

The first page of configuration will be for general information such as Host Participants and Job name tag.

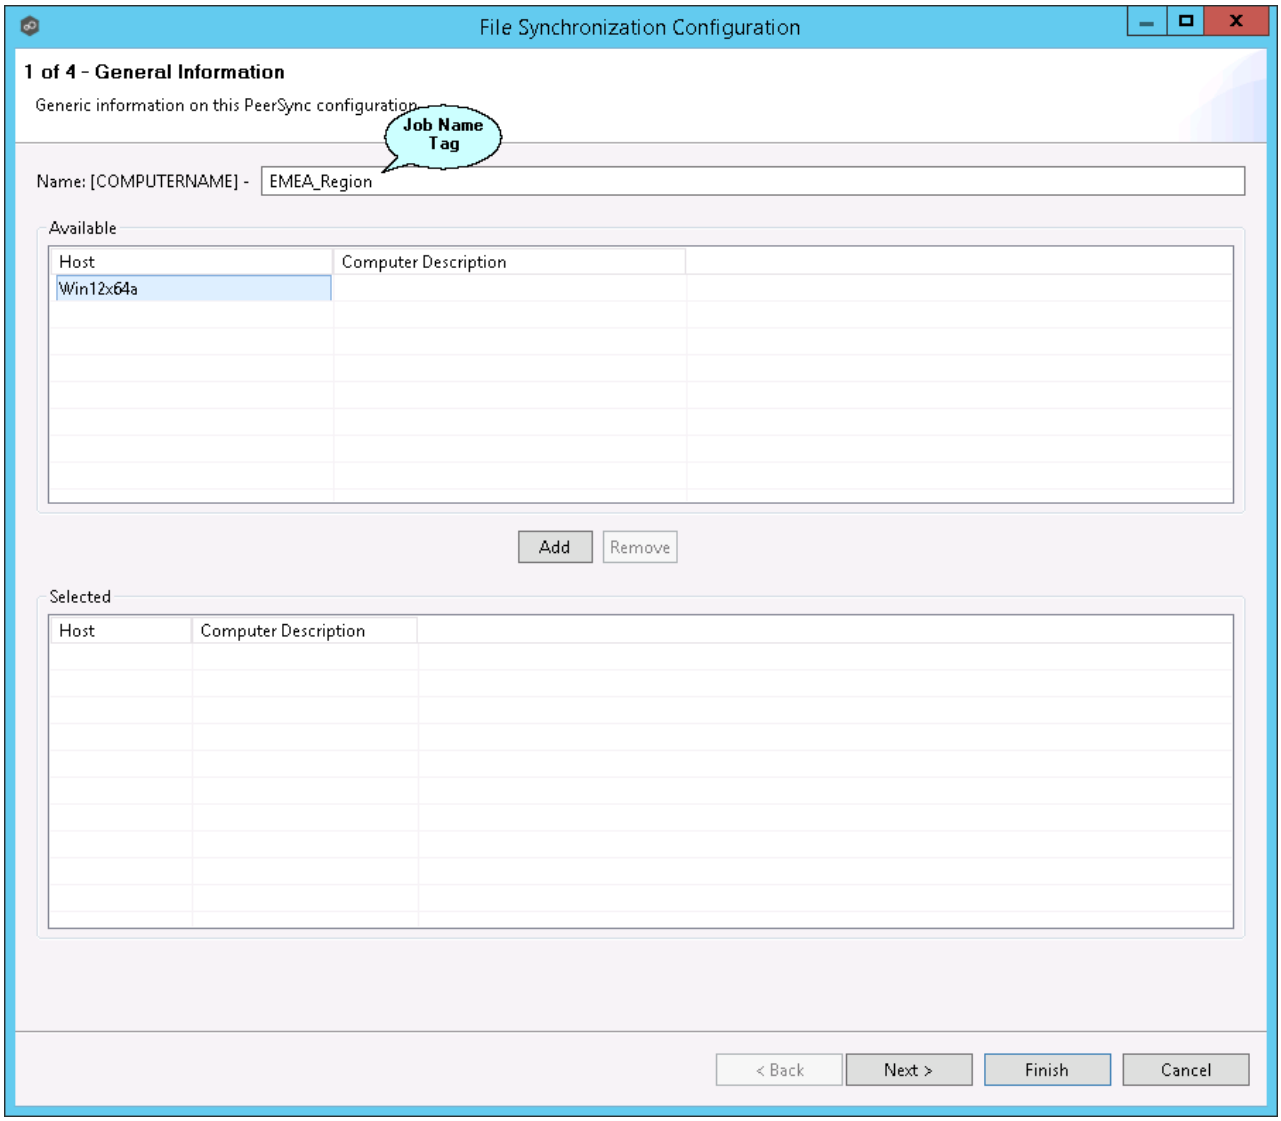

1. The job name will default to the computer name of the host participant. If you wish to group your computers, you can optionally add a name tag in the text box next to the job name (e.g., East Coast, EMEA, Region2). This will help in filtering machines by their given tag.

2. A list of all available hosts that have not yet been configured with a PeerSync installation, will appear in the **Available** table on the top of the page. Available hosts are any host with a Peer Agent installed that has successfully connected to the configured Peer Management Center Broker. The name that will be displayed is the computer name of the server that the Peer Agent is running on. If a particular host is not displayed in the list, then try restarting the Peer Agent Windows Service on that host, and if it successfully connects to Peer Management Center Broker, then the list will be updated with the computer name of that host.

   Note: Computer Description is defined through Windows on a per-computer basis.

3. Select one or more hosts from the **Available** table and click the **Add** button to add the hosts to the **Selected** table. These are the hosts you wish to deploy the PeerSync configuration and installation to.

Step 2: PeerSync Profile

In the second page, choose a preconfigured profile from the available templates, or browse to load a PeerSync profile you may have configured through the PeerSync Profiler and saved as a .snc file on this system.

You may also choose to start from scratch by choosing **Other** from the drop-down menu.

Enter or update the **Profile Description** and **Performance Options**.

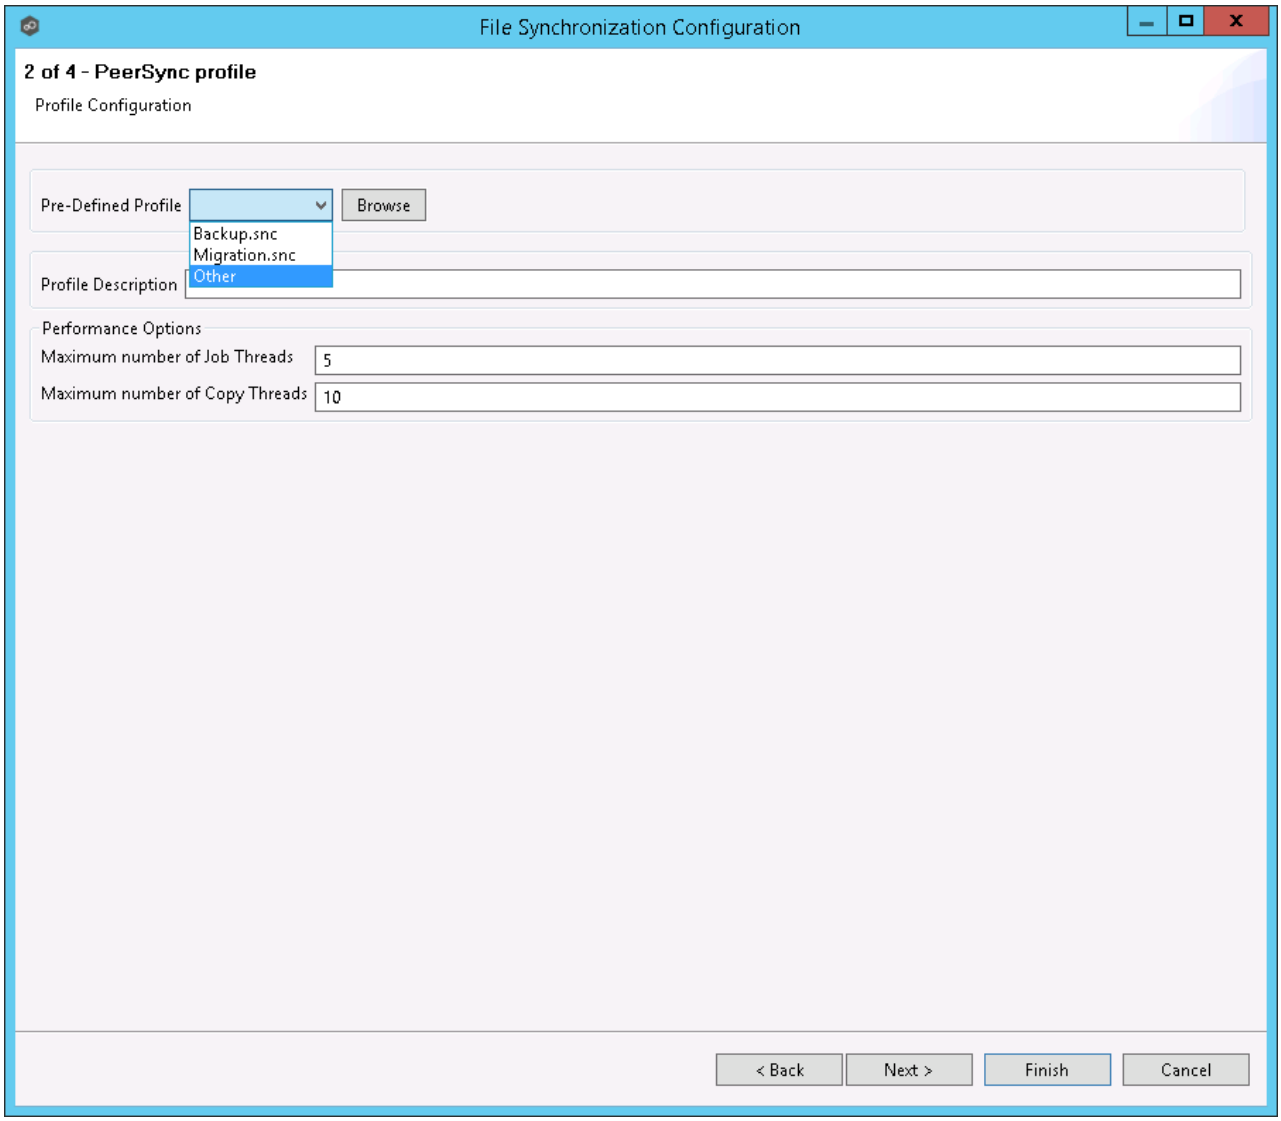

Profile Description	A textual description of the current profile.
Maximum number of Job Threads	Maximum number of job scans that can run parallel to one another.
Maximum number of Copy Threads	Maximum number of events that can be processed parallel to one another.

Step 3: Jobs Configuration List

In this page, modify the loaded PeerSync jobs and/or add new jobs by clicking on the **Add** or **Edit** button to the right of the view.

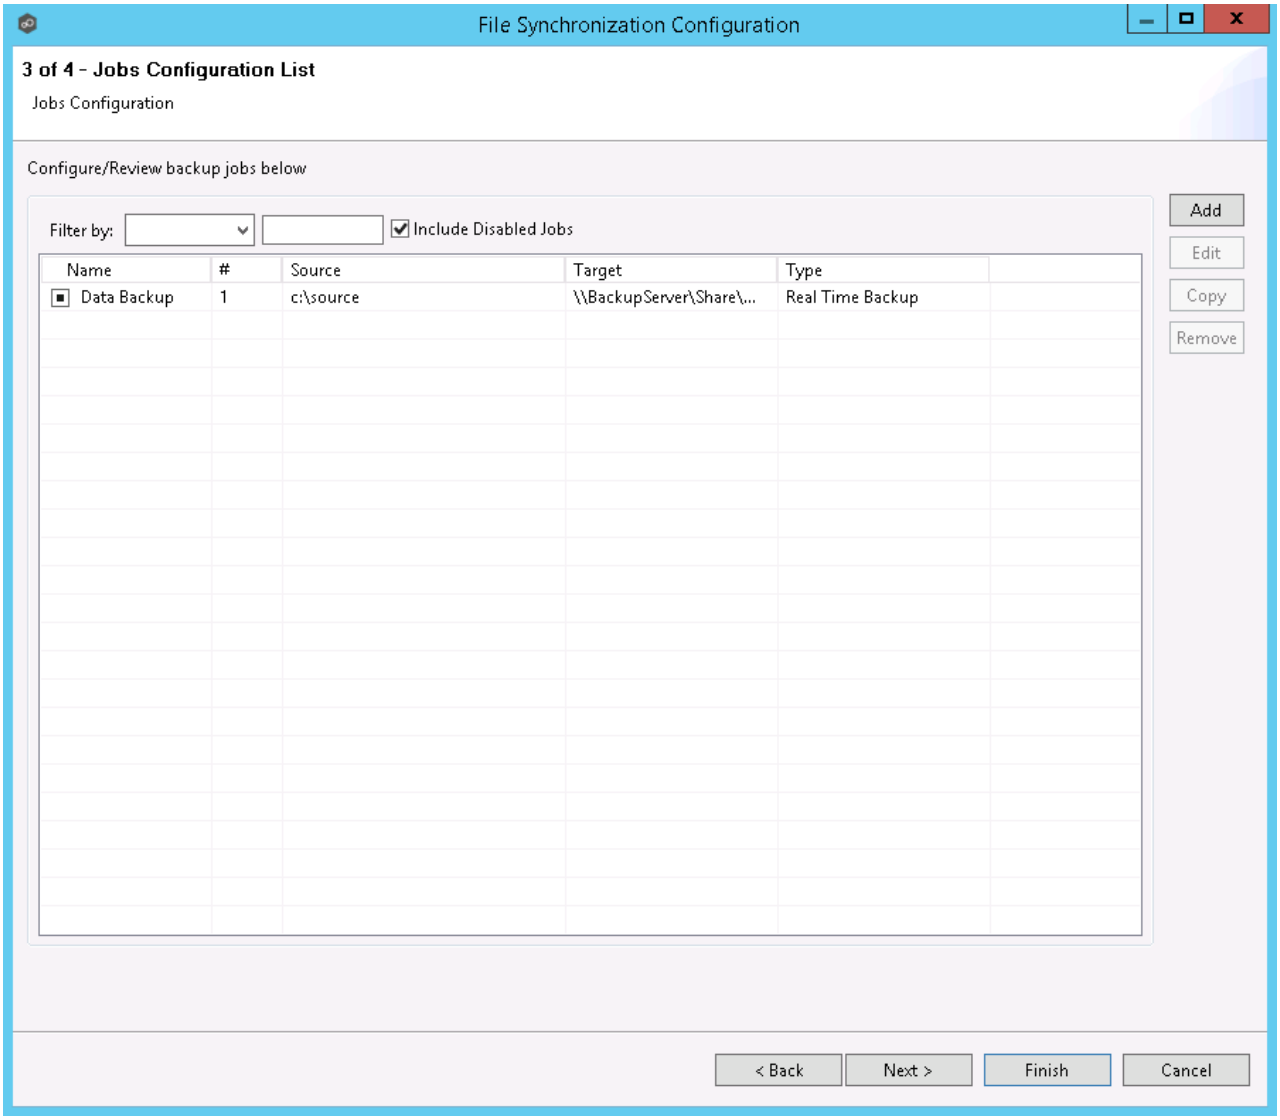

For more information, see Edit/Configure Jobs.

Step 4: Installation Settings

In the last page of the PeerSync Management Configuration wizard, enter the installation settings for this PeerSync instance.

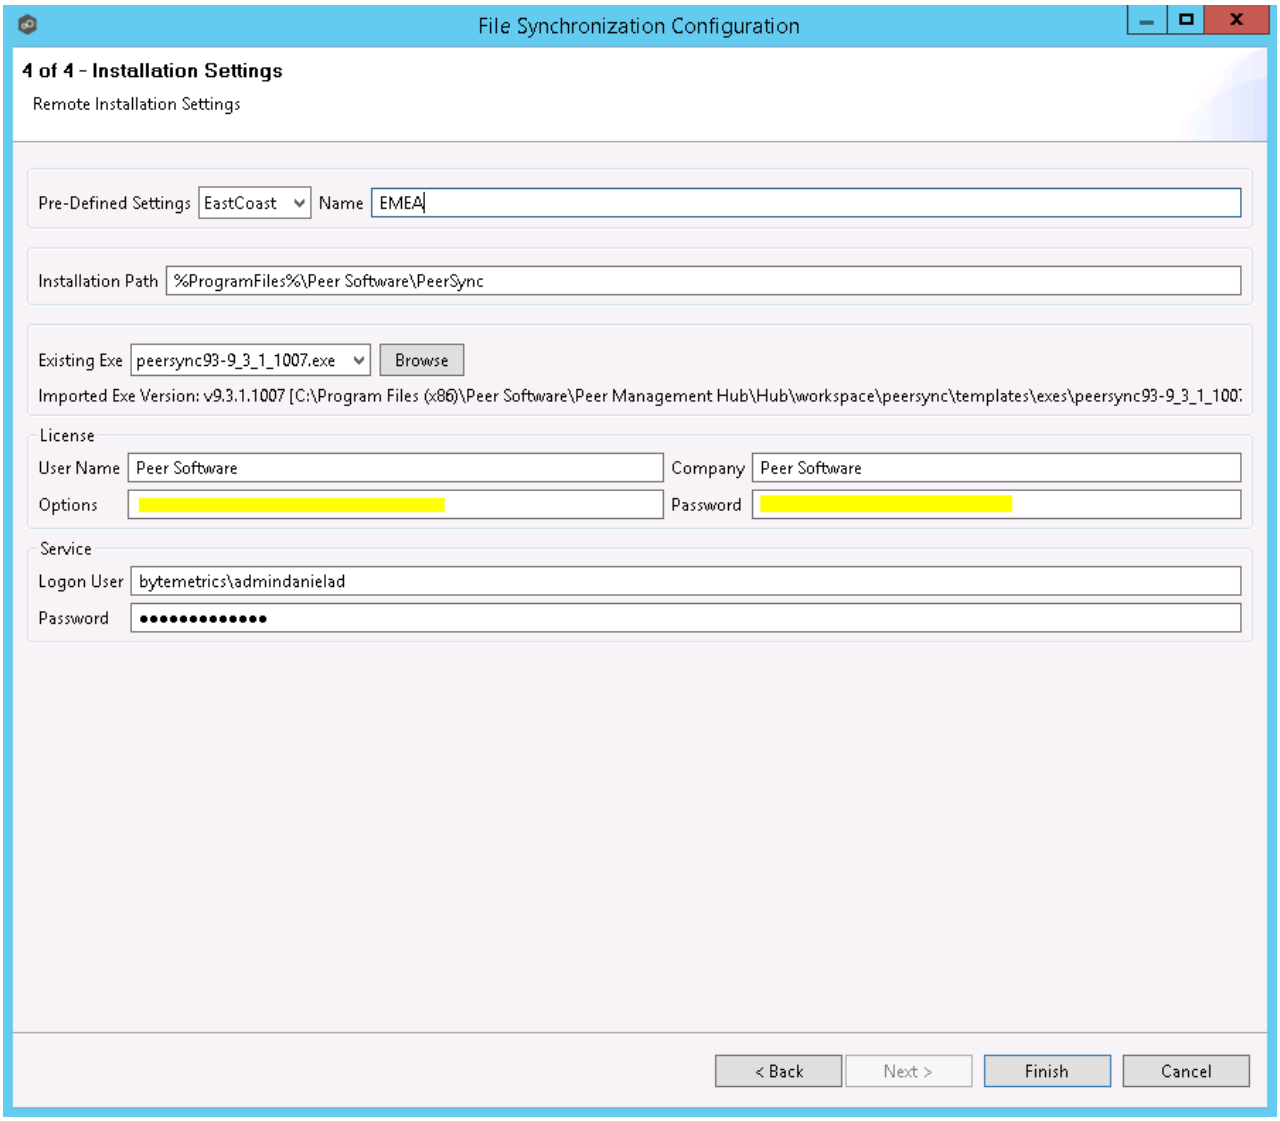

Predefined Settings	A list of previously used installation settings with the Name given at the time of use. Note: If the installation settings have the same path, service user name, license password, and installation exe, a new Installation record will not be created, regardless of whether a new name has been given to the Installation Settings.

Installation Path	Path where PeerSync will be installed. When using the %ProgramFiles% variable, PeerSync will install in the x86 Program Files directory for 64-bit systems, otherwise it will install in the Program Files base directory.
Existing Exe	A list of PeerSync executables available in the template folder or used in a past installation. This is the PeerSync executable that will be used to install PeerSync.
License User Name	License information provided by Peer Software. Cut and paste the User Name section in this field.
License Company	License information provided by Peer Software. Cut and paste the Company section in this field.
License Options	License information provided by Peer Software. Cut and paste the Options section in this field.
License Password	License information provided by Peer Software. Cut and paste the Password section in this field.
Service Logon User*	This is the Service account User ID used to run the PeerSync Service (DOMAIN\USER). Note: This account must be valid on all included participants for this File Synchronization Configuration.
Service Password	PeerSync Windows Service account password.

***Note:** When using a service account that has not been granted to run as a service on the machine, PeerSync will fail to start return the following Global Alert to the Peer Management Center. This will indicate that PeerSync could not start, and you will have to log on to that machine and confirm the credentials to grant access to that account to run as a service.

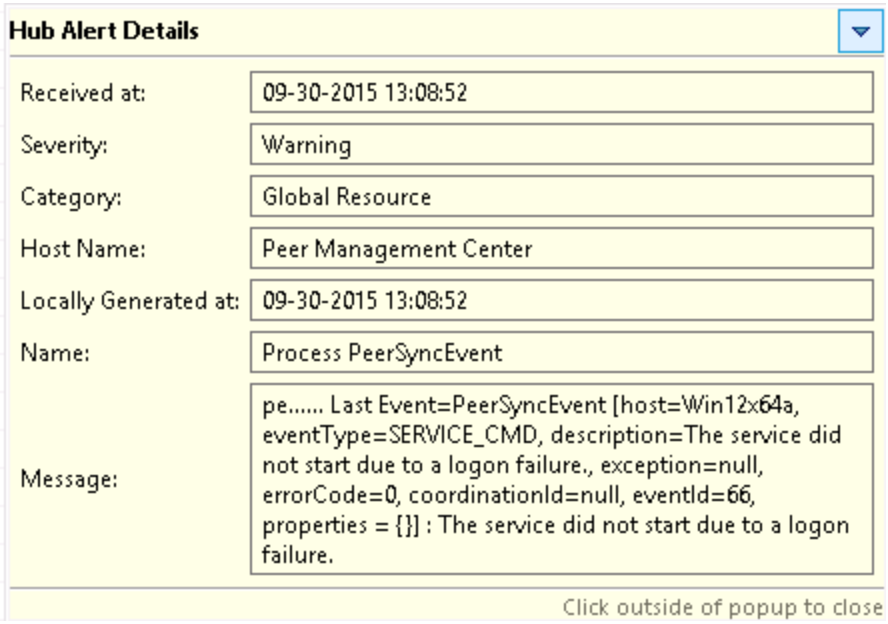

Once the Configuration Settings have completed, click **Finish** and the installation configuration will be sent to the selected Participants.

A File Synchronization job will be auto created for each Participant and set to be in a Pending Installation state. Once the installation completes and PeerSync reports to the Peer Management Center, the state will change to Running/Active.

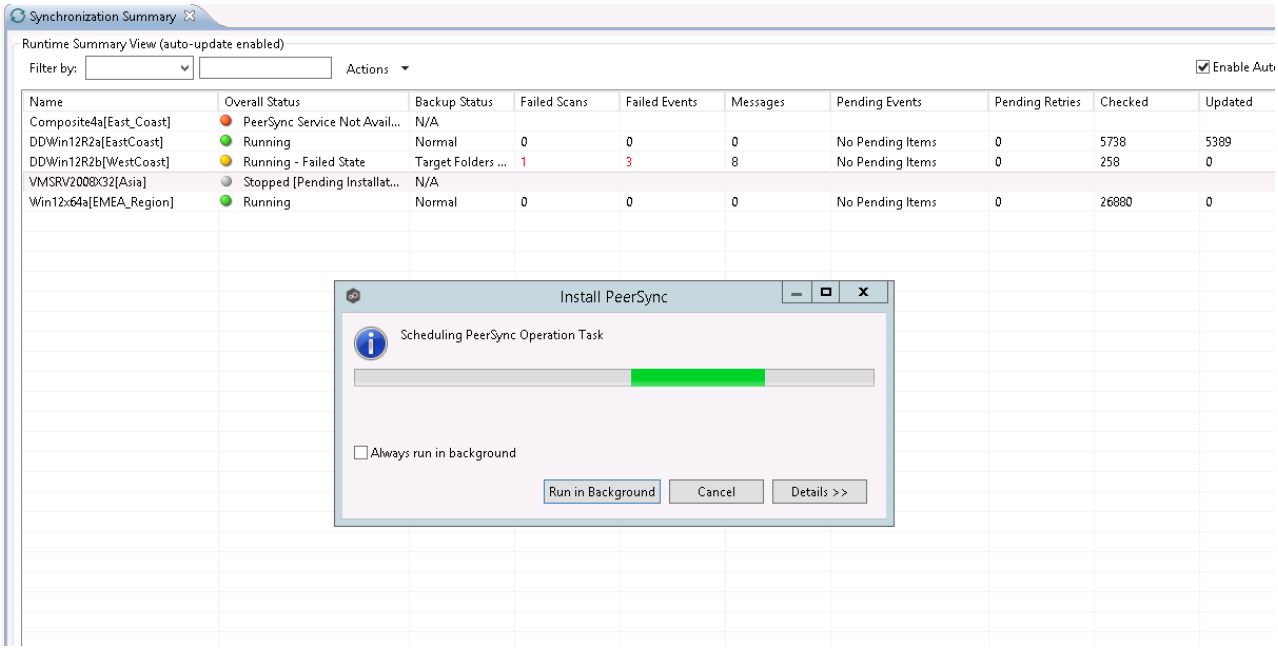

Logging and Alerts

Use the following dialog to enable or disable logging and alerts, including specifying event types to log.

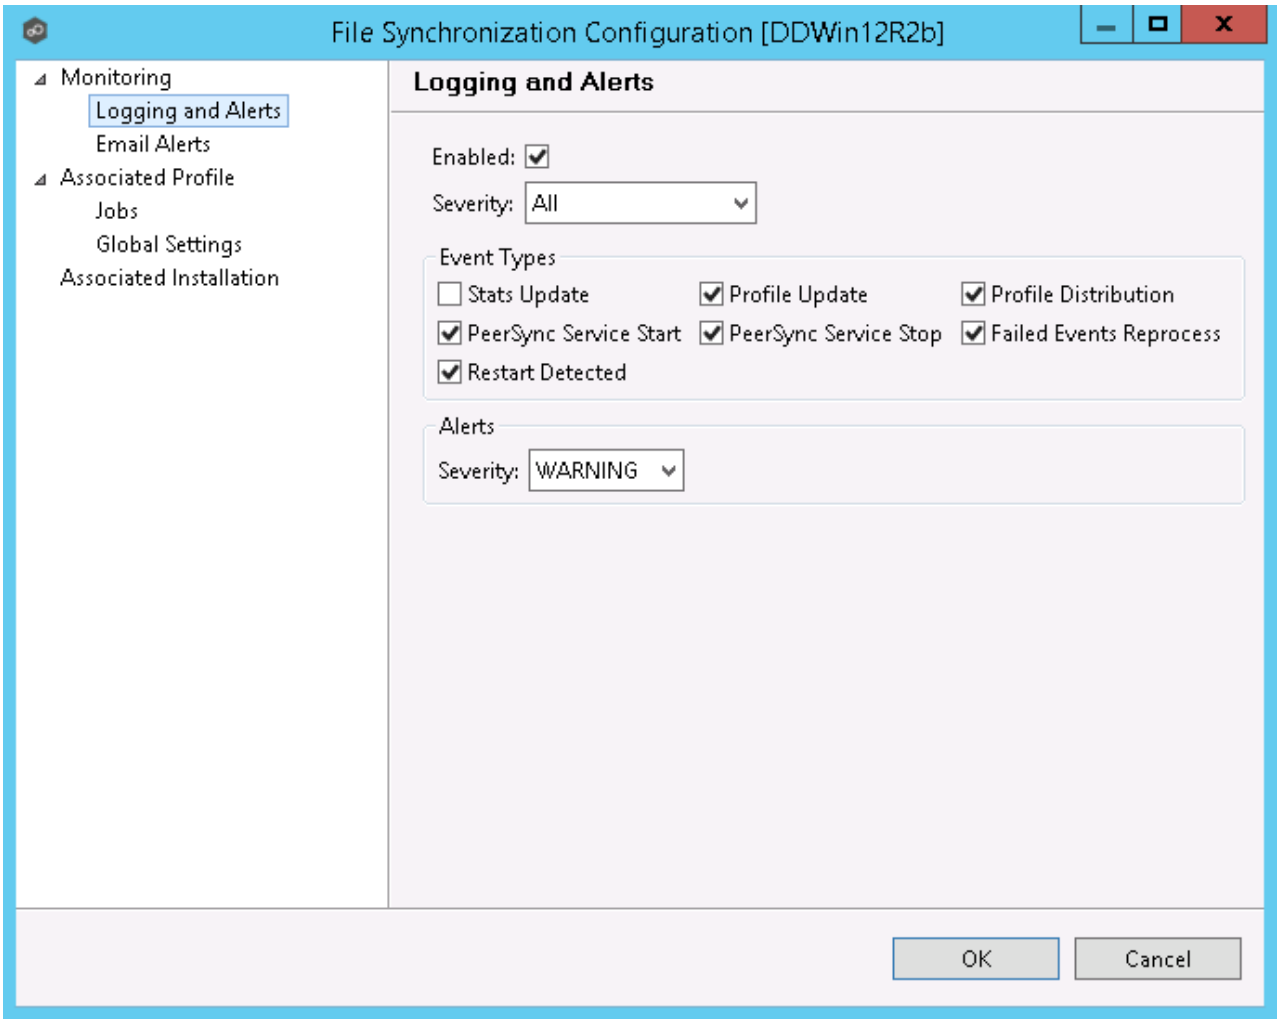

Stats Update	Log when PeerSync Stats are received (This could generate large amount of Log Entries).
Profile Update	Log whenever the PeerSync Profile configuration is updated.
Profile Distribution	Log when the PeerSync Profile is distributed to one or more hosts.

PeerSync Service Start	Log when a user initiates a PeerSync Service Start.
PeerSync Service Stop	Log when a user initiates a PeerSync Service Stop.
Failed Events Reprocess	Log when a user initiates a Failed Event Reprocess.
Restart Detected	Log when Peer Management Center detects that the PeerSync service has been restarted by comparing known Session Id with received one.

Email Alerts

Email alerts configuration is available by selecting **Email Alerts** from the tree node within the PeerSync Management Configuration dialog.

Email alerts are configured at a global level, then applied to individual jobs. The following screen shows how this is accomplished.

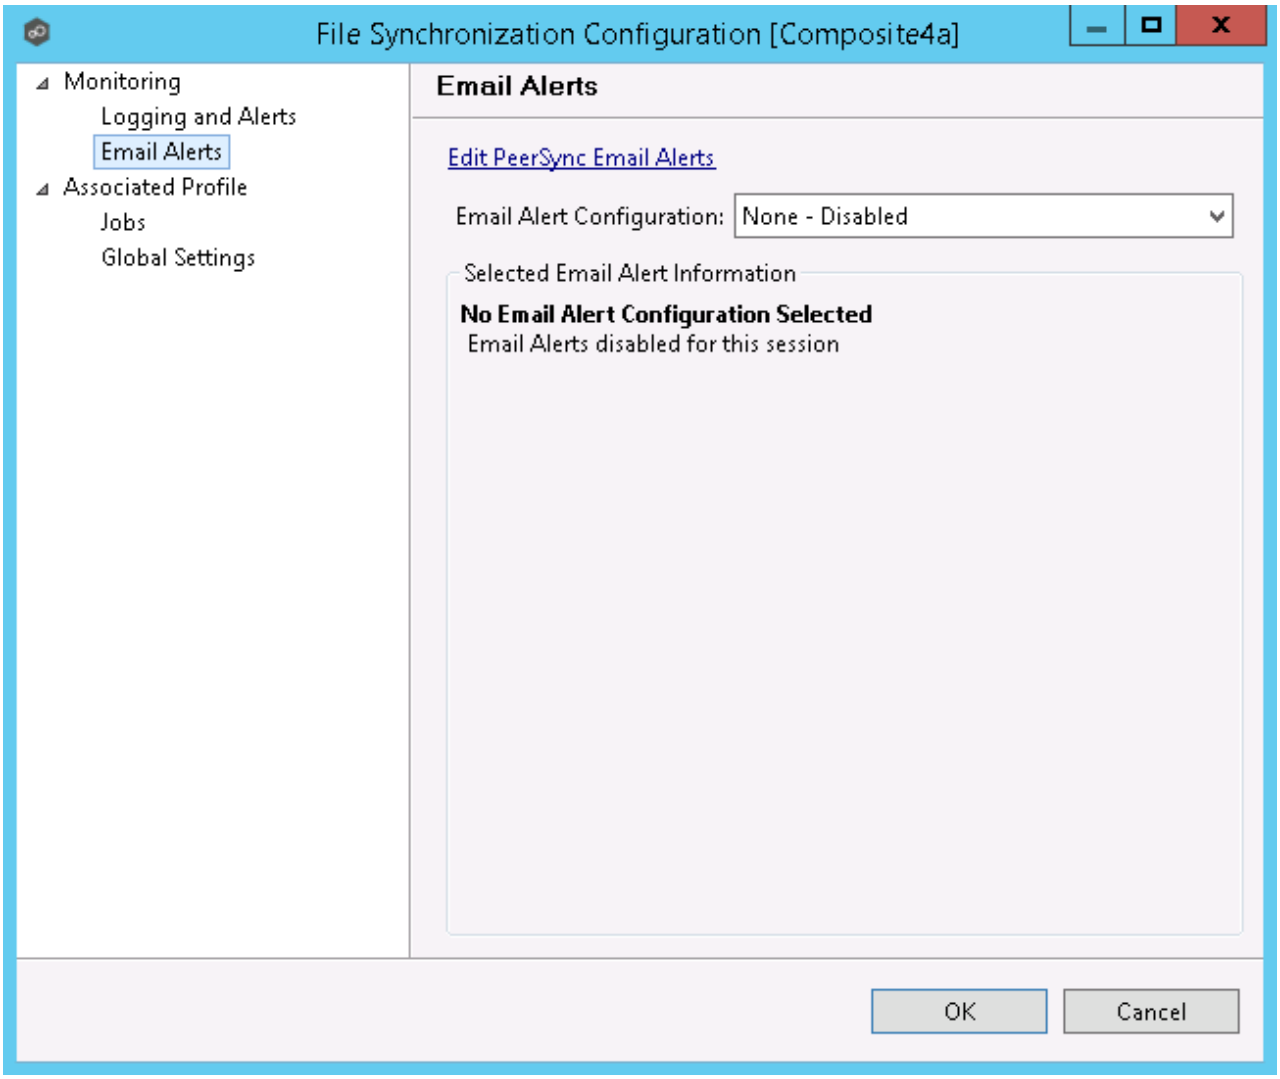

To enable email alerts for this job, select an email alert from the drop-down list. To disable, select **None - Disabled**.

Running and Managing a PeerSync Management Job

This section includes topics for managing your PeerSync Management Jobs.

- Starting and Stopping
- Synchronization Summary View
- Synchronization Dashboard Summary View

- PeerSync Profile Management
- PeerSync Service Management
- Runtime Job Views
- Upgrade/Reprocess Installation

Starting and Stopping

Starting a PeerSync Management Job

A PeerSync Management job is auto-started as soon as the Agent connects to the Peer Management Center; normally you will not need to manually start the job.

Click the **Start** button to begin the session.

Stopping a PeerSync Management Job

You can stop a PeerSync Management job at any time by pressing the **Stop** button. Doing this will shut down the monitoring of the specific PeerSync host(s).

Note: If the job is stopped and the participating host is still running an instance of the PeerSync software, the job will auto start the next time that host agent is restarted or a Reconnect is detected.

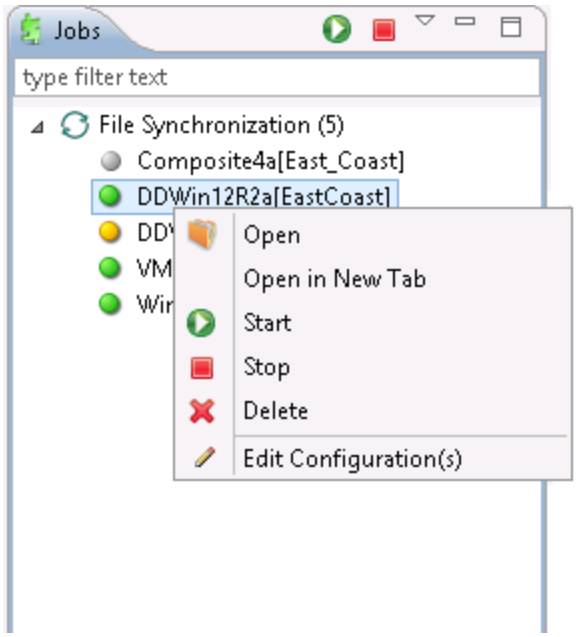

PeerSync Management Summary

The **Synchronization Summary** view aggregates critical status and statistical information from all configured PeerSync Management jobs in a single table view. It is automatically displayed when the Peer Management Center client is started and can be opened at any other time by double-clicking on the **PeerSync Management** parent tree node in the Job's View or by clicking the **View Runtime Summary** icon in the toolbar of the Jobs view.

Information in this view can be sorted and filtered. Operations such as starting, stopping, and editing multiple jobs at once are available, in addition to the ability to clear Monitoring Alerts, Start PeerSync, Stop PeerSync, Reprocess Failed Events, Request Support Info File and Reprocess/Upgrade Installation.

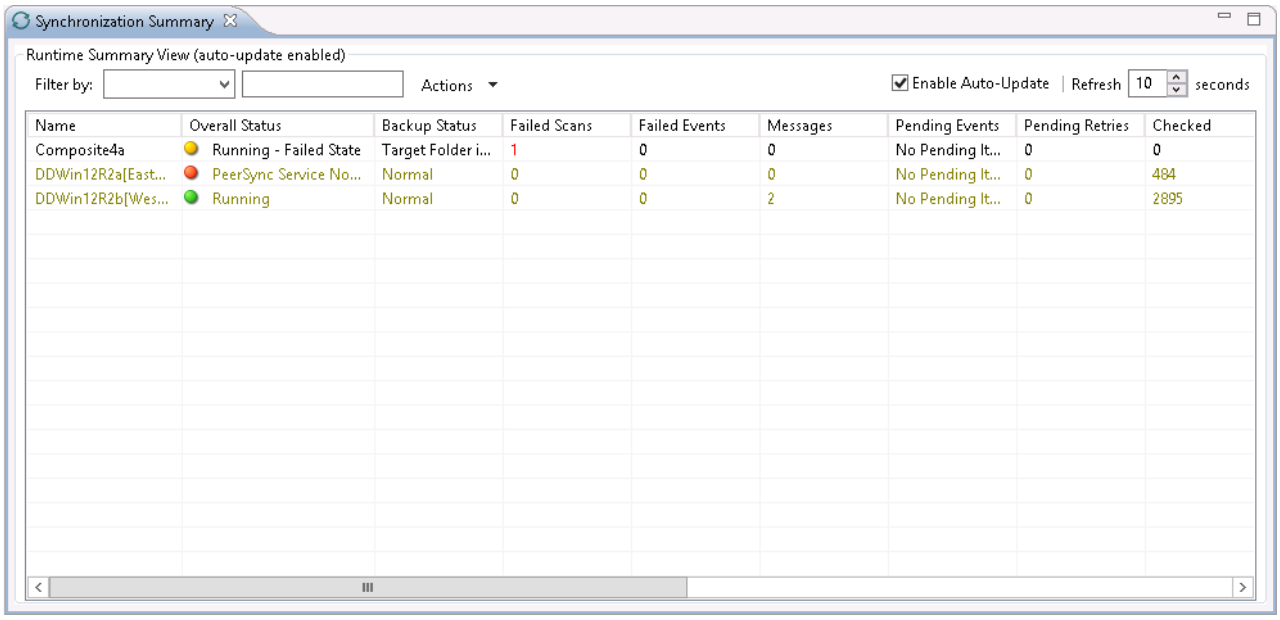

Unlike other views within the Peer Management Center, the **Synchronization Summary** view is not updated in real-time. This is done for performance reasons. Instead, the table can be set to automatically update itself every few seconds. Clicking **Enable Auto-Update** enables this functionality, while the refresh interval (in seconds) can be set right beside the checkbox. Additional columns can be added to and removed from the table from the right-click context menu.

Double-clicking any item in the table will automatically open the selected PeerSync Management job in a tab within the **Runtime Summary** view, allowing you to drill down and view specific information about that single job. Items in the summary table can be filtered by job name, overall status, activity state and host participant name.

Selecting one or more items in the table, then right-clicking will bring up a context menu of available actions that can be performed on the selected jobs. The actions that are unique to this table are as follows:

Clear Monitoring Alerts	Clears all monitoring alerts for the selected jobs. This can be performed while a job is running.

PeerSync multi-job global actions:

Clear Monitoring Alerts	Clears all monitoring alerts for the selected jobs. This can be performed while a job is running.

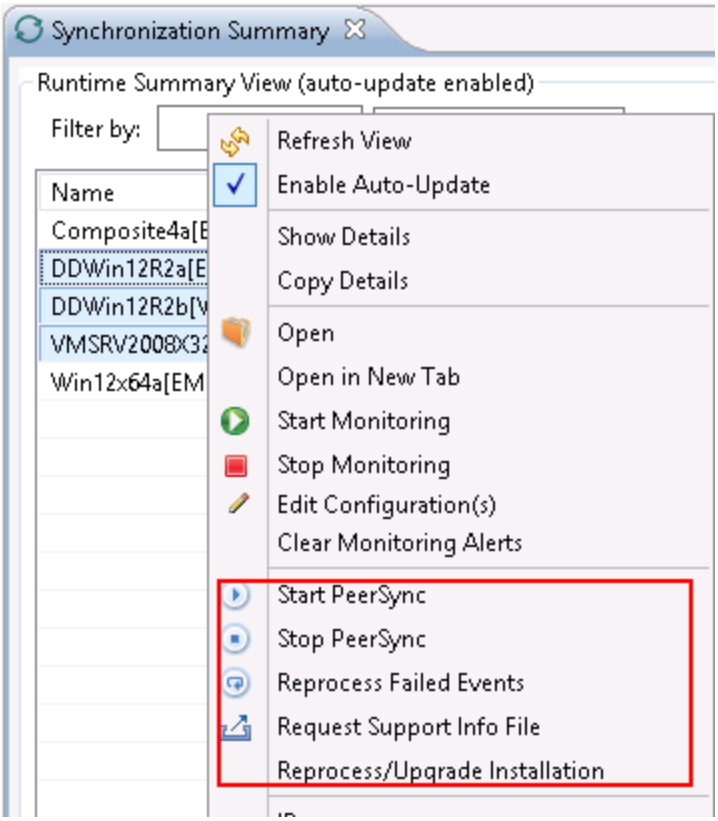

Start PeerSync	Send a Start command to the PeerSync service instance running on the selected jobs' participant. This can be performed while the associated File Synchronization Job is running.
Stop PeerSync	Send a Stop command to the PeerSync service instance running on the selected jobs' participant. This can be performed while the associated File Synchronization Job is running.
Reprocess Failed Events	Send a Reprocess Failed Events command to the PeerSync instance running on the selected jobs' participant. This can be performed while the associated File Synchronization Job is running.
Request Support Info File	Send a request to collect the Support info File from the PeerSync instance running on the selected jobs' participant. This can be performed while the associated File Synchronization Job is running.

Clear Monitoring Alerts	Clears all monitoring alerts for the selected jobs. This can be performed while a job is running.
Reprocess/Upgrade Installation	Deploy an upgrade or reprocess an existing installation for the selected File Synchronization Job(s). Upgrade/Reprocess Installation

Clicking the **Actions** table menu provides the following options:

Refresh View	Refresh all information provided in the table.
Copy All Filtered Statistics	Copy detailed information to the system clipboard for all items current displayed in the table, taking any filters into account. This information can then be pasted into a text editor.
Export Entire Table to File	Dump the entire contents of the table to a text file that can be viewed in any text editor.

PeerSync Management Dashboard Summary View

The **PeerSync Management Dashboard Summary** view is a view that displays metrics and key performance indicators from all running PeerSync Management jobs. It is automatically displayed when the Peer Management Center client is started and can be opened at any other time by selecting **View Dashb**oard from the **Window** menu or by clicking the **View Dashboard** icon in the Peer Management Center toolbar.

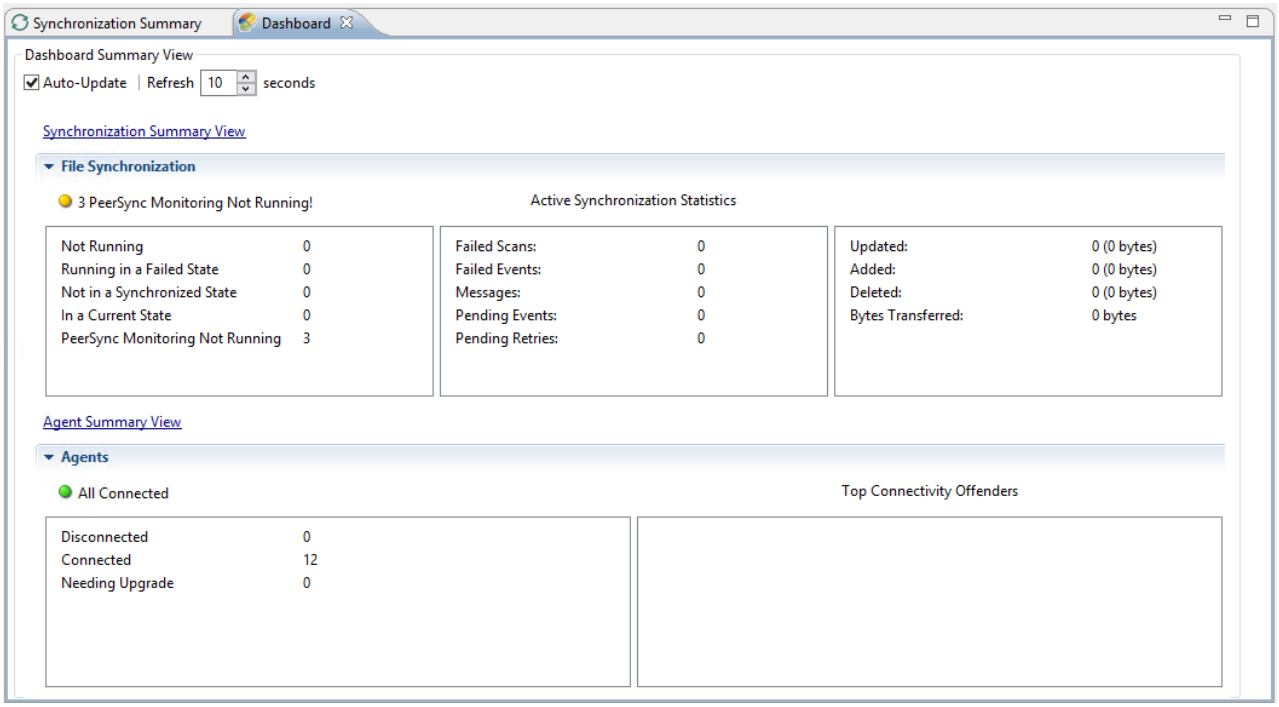

The Dashboard is not updated in real-time. This is done for performance reasons. Instead, the table can be set to automatically update itself every few seconds. Enabling the **Auto-Update** option will enable this functionality, while the **Refresh** interval (in seconds) can be set right beside the checkbox.

Entries in the first column of the **PeerSync Management Job** and **Agents** categories can be double-clicked, which will take the user to a filtered Runtime view of the selected item for additional details.

Managing the PeerSync Profile

The topics in this section provide some basic information about PeerSync Profile Management:

- Updating the Profile Configuration
- Importing an Existing Profile
- Distributing a Profile

This topic provides information on how to update a PeerSync profile from the Peer Management Center.

If using the Peer Management Center to manage the PeerSync instances, we recommend making changes through the Peer Management Center. If changes are made directly on the PeerSync machine, they should be imported in the Peer Management Center job manually to keep the Peer Management Center PeerSync Configuration in sync.

How to Update a PeerSync Profile through Peer Management Center

- From the **PeerSync Management Summary** runtime view (double-click the **PeerSync Management jobs** node from the left), right-click the machine you wish to modify the profile for and choose **Edit Configuration(s)**. Alternatively, you can right-click the machine job from the left menu under the **PeerSync Management** node and choose **Edit Configuration(s)**.

- You can update the Profile by importing an updated Profile through the **Import** button in the **Associated Profile** page, or manually update the configuration through **Jobs and/or Global Settings** section.

- If you wish to update the profile outside of the Peer Management Center, export the existing configuration using the **Export** button in **Associated Profile** page. Make your changes through the PeerSync Profiler and import the updated Profile back into Peer Management Center through the **Import** button.

- After having made your entire configuration changes either through Peer Management Center or by importing the updated Profile, choose **OK** and close the **Edit Configuration** dialog.

Your configuration changes will not reach the PeerSync machine until they are distributed. The updated profile will become active on the machine after the PeerSync service has been restarted.

- Import Existing Profile

- Edit/Configure Jobs

- Edit Global Settings

- Distribute Profile

Importing an Existing Profile

In the **Associated Profile** section of the **PeerSync Management Configuration** dialog, you can update the configured profile with one you have saved and configured outside of the Peer Management Center.

Note: If making changes outside of the Peer Management Center, we recommend exporting the profile from the Peer Management Center (by clicking the **Export** button), making necessary changes outside of the Peer Management Center, and finally importing the profile back into the Peer Management Center.

Click the **Import** button on the right of the dialog to import the profile. To propagate this new updated profile, close the **PeerSync Management** dialog, reopen it, and then distribute to the PeerSync host through the Distribute button.

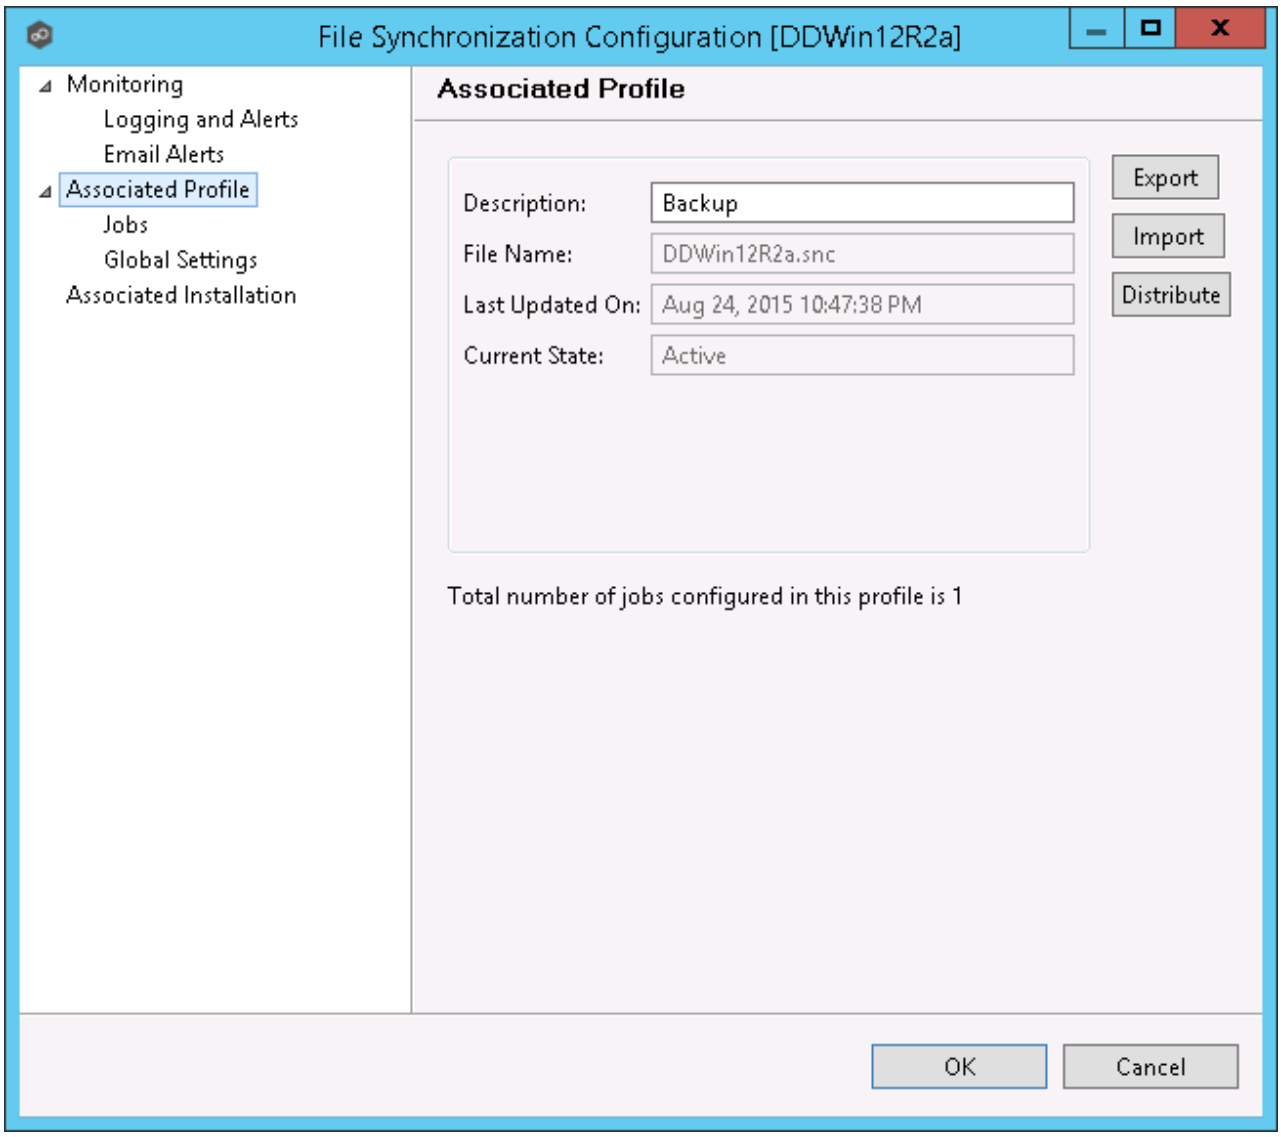

Editing and Configuring Jobs

From the **Jobs** view in the **PeerSync Management Configuration** dialog, you can make several configuration changes:

- Add New Job
- Edit Existing Job
- Enable/Disable Job
- Copy Job
- Remove Job

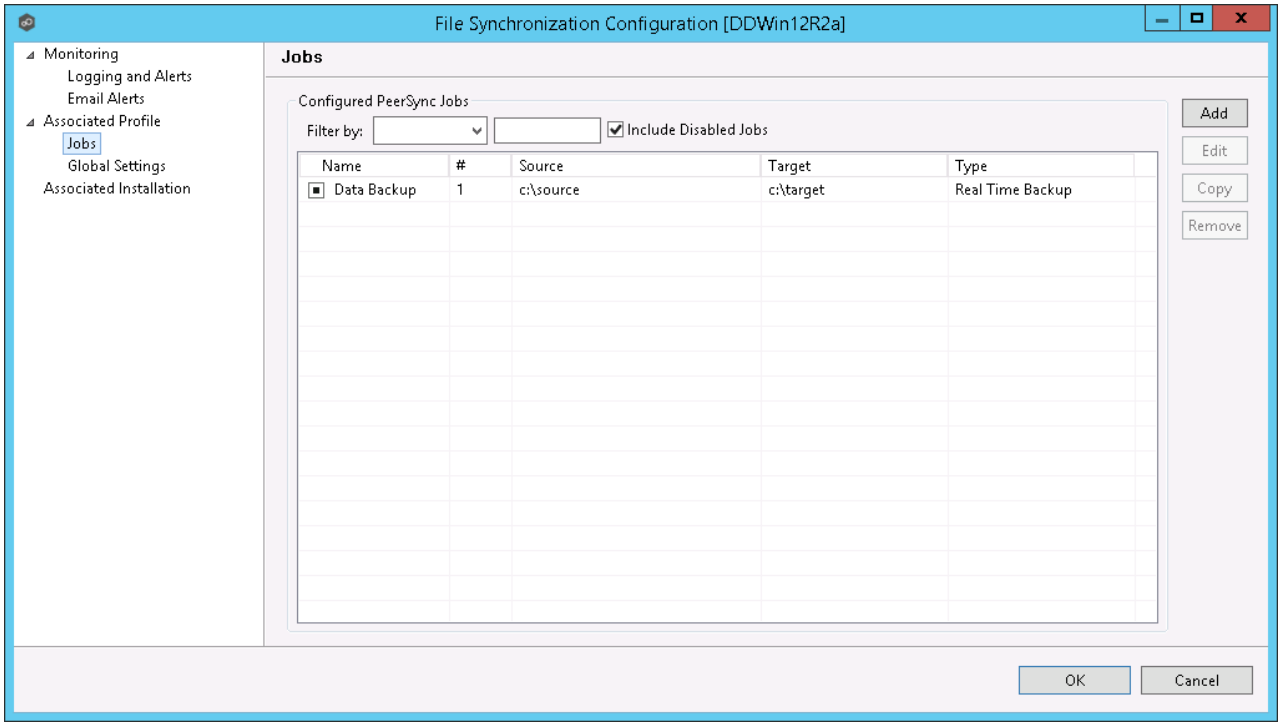

Add New Job

To add a new job, click the **Add** button on the right of the **Jobs** view and select one of the job types available from the drop-down list below.

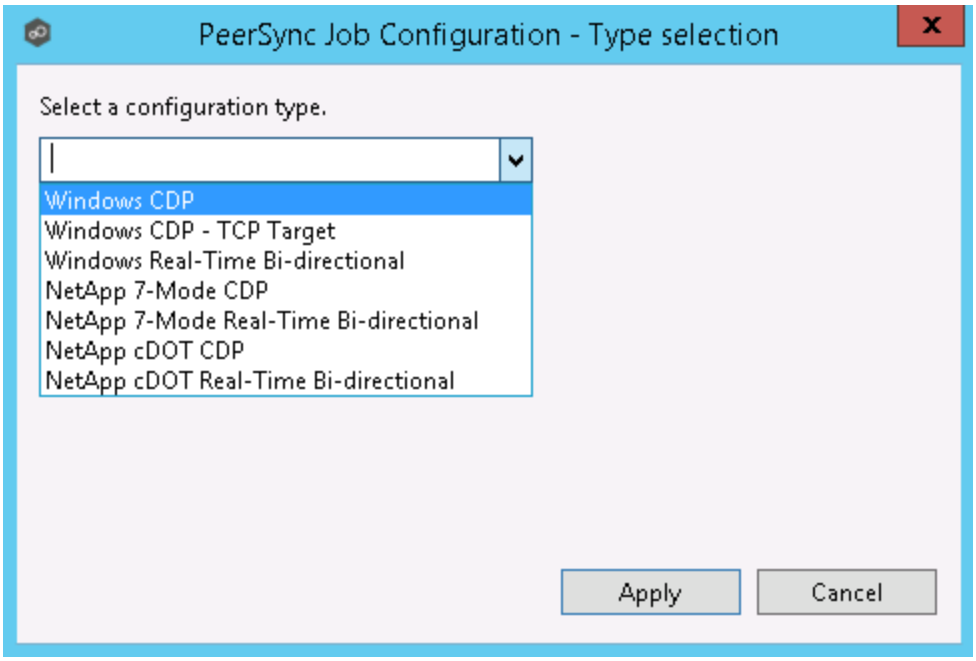

Once a job type has been selected, click **Apply** and complete the PeerSync Job Configuration wizard to complete the job configuration and add the job to the profile.

Edit Existing Job

To edit an existing job, select the job in the **Jobs** view, and then click the **Edit** button on the right. The **PeerSync Job Properties** dialog will open with all available settings grouped by category in the navigation tree.

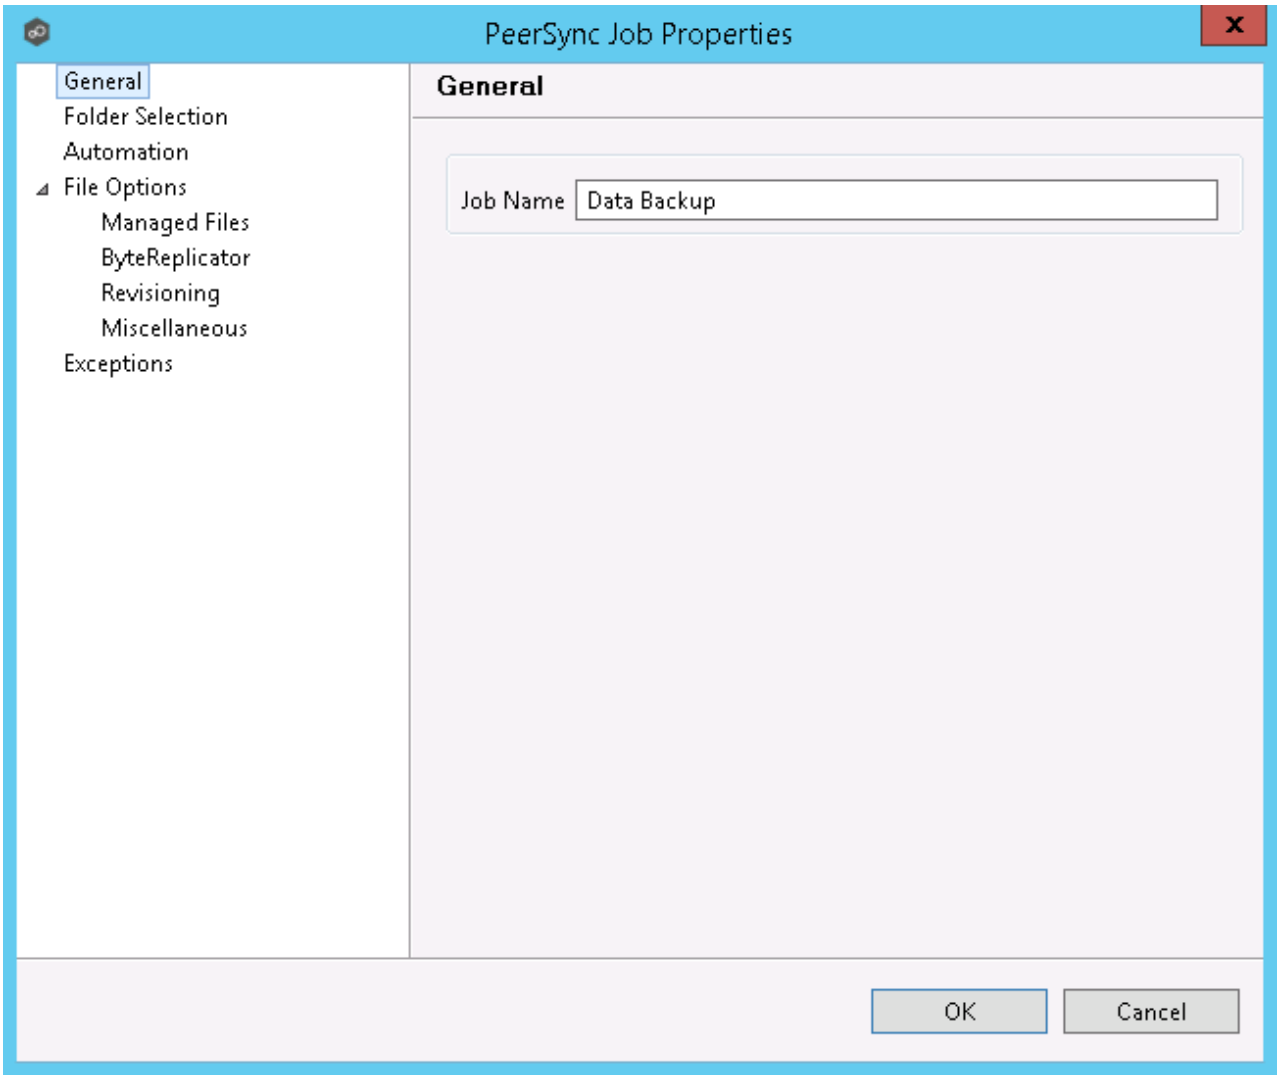

Enable/Disable Job

To enable or disable a job, click the checkbox to the left of the job name in the **Jobs** view.

To save these changes, click **OK** on the bottom right of the **PeerSync Management Configuration** dialog.

Copy Job

You can copy an existing job by selecting the job from the **Jobs** view and clicking the **Copy** button on the right. The **PeerSync Job Properties** dialog will open, allowing you to make changes to the copied job.

Note: You must make at least one change to the job settings. If the job settings remain identical, it will not be saved after the **OK** button is clicked.

Remove Job

To remove jobs from the PeerSync Configuration, select one job from the **Jobs** view, click the **Remove** button on the right. Repeat this for any additional jobs you wish to remove.

Editing Global Settings

In the **Global Settings** of the **PeerSync Management Configuration** dialog, you can make changes to settings that apply to all PeerSync Jobs within the profile.

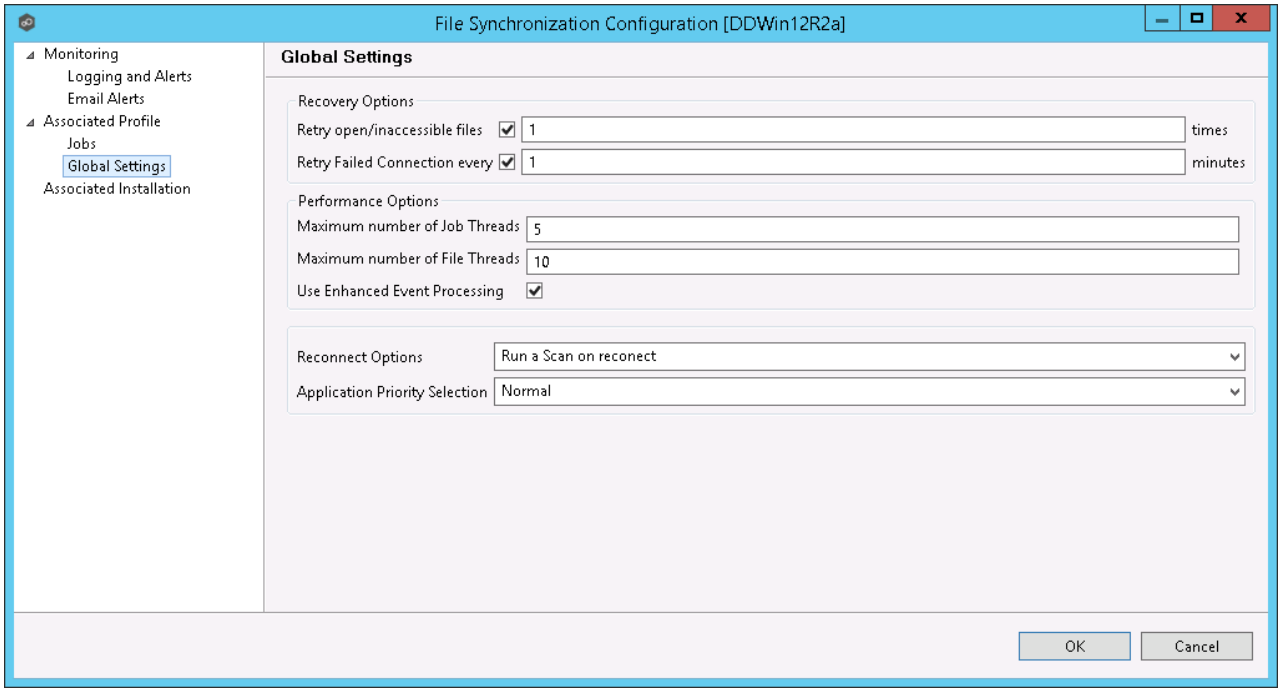

Recovery Options	These changes will update how we retry failed or inaccessible files as well as the interval in which we retry Failed Connections.
Performance Options	These settings allow you to change the maximum number of job scans that can run parallel to one another and the maximum number of events that can be processed parallel to one another.
Reconnect Options	This setting allows you to choose how PeerSync handles a re-established connection. Options are to **Run a Scan on reconnect**

	or **Store missed events and process on reconnect**.
Application Priority Selection	This setting enables to select the priority level you want PeerSync to have.

Distributing a Profile

This topic covers information on how to distribute changes to the PeerSync profile from the Peer Management Center

To distribute the PeerSync profile changes, right-click the PeerSync Management job from the **Jobs** view, and then click **Edit Configuration**.

In the **PeerSync Management Configuration** screen, click the **Associated Profile** node, and then click the **Distribute** button.

In the event that one or more of your jobs are configured to use a ByteReplicator Relay Server (usually used in NetApp source environments), the Distribute Profile process will also distribute your relay Server configurations by compiling all unique target Hosts and relay servers into a *%profilename%.pls* file. This file will be distributed to the PeerSync machine alongside the profile.

Note: This action will distribute the profile to the machine and attempt to stop and start PeerSync Service to commit those changes. If you do not wish to restart the PeerSync service, wait to distribute the profile until you are ready to have the service restart.

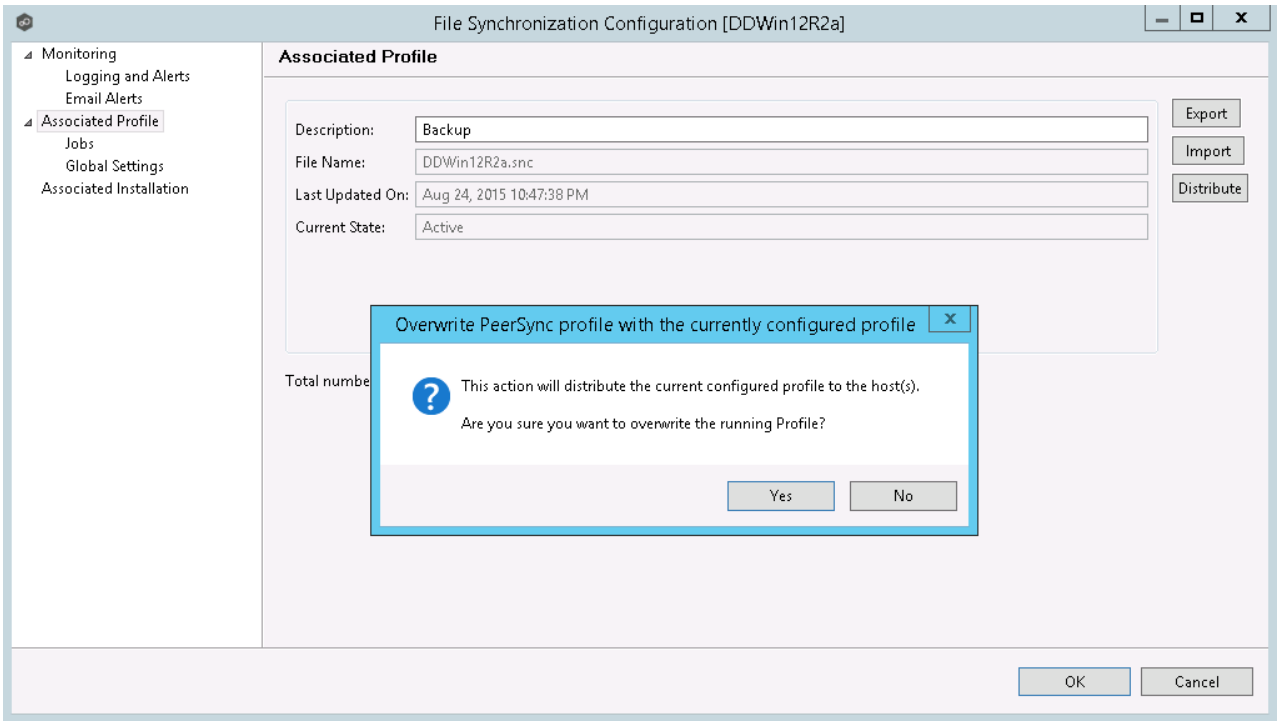

Managing the PeerSync Service

The following PeerSync service management actions are available from the Synchronization Summary view and the Summary view for a specific PeerSync Management job.

Starting the PeerSync Service

To start the PeerSync service associated with any PeerSync Management job, right-click the view and choose **Start PeerSync**.

Stopping the PeerSync Service

To stop the PeerSync service associated with any PeerSync Management job, right-click the view and choose **Stop PeerSync**.

Note: The associated PeerSync Management job must be running in order to successfully perform this action.

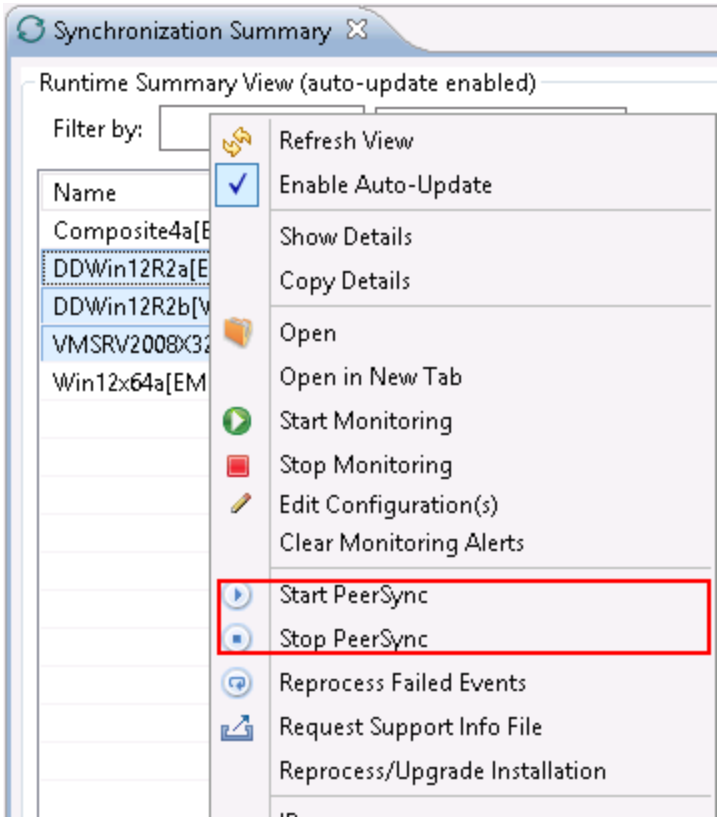

For information on the additional PeerSync multi-job global actions, see Synchronization Summary View.

Runtime Job Views

Double-clicking on the PeerSync Management job from the Synchronization Summary view will open the job-specific runtime views.

- Summary View
- Failed Events View
- Monitoring Log View
- Alerts View
- Participants View
- Configuration View

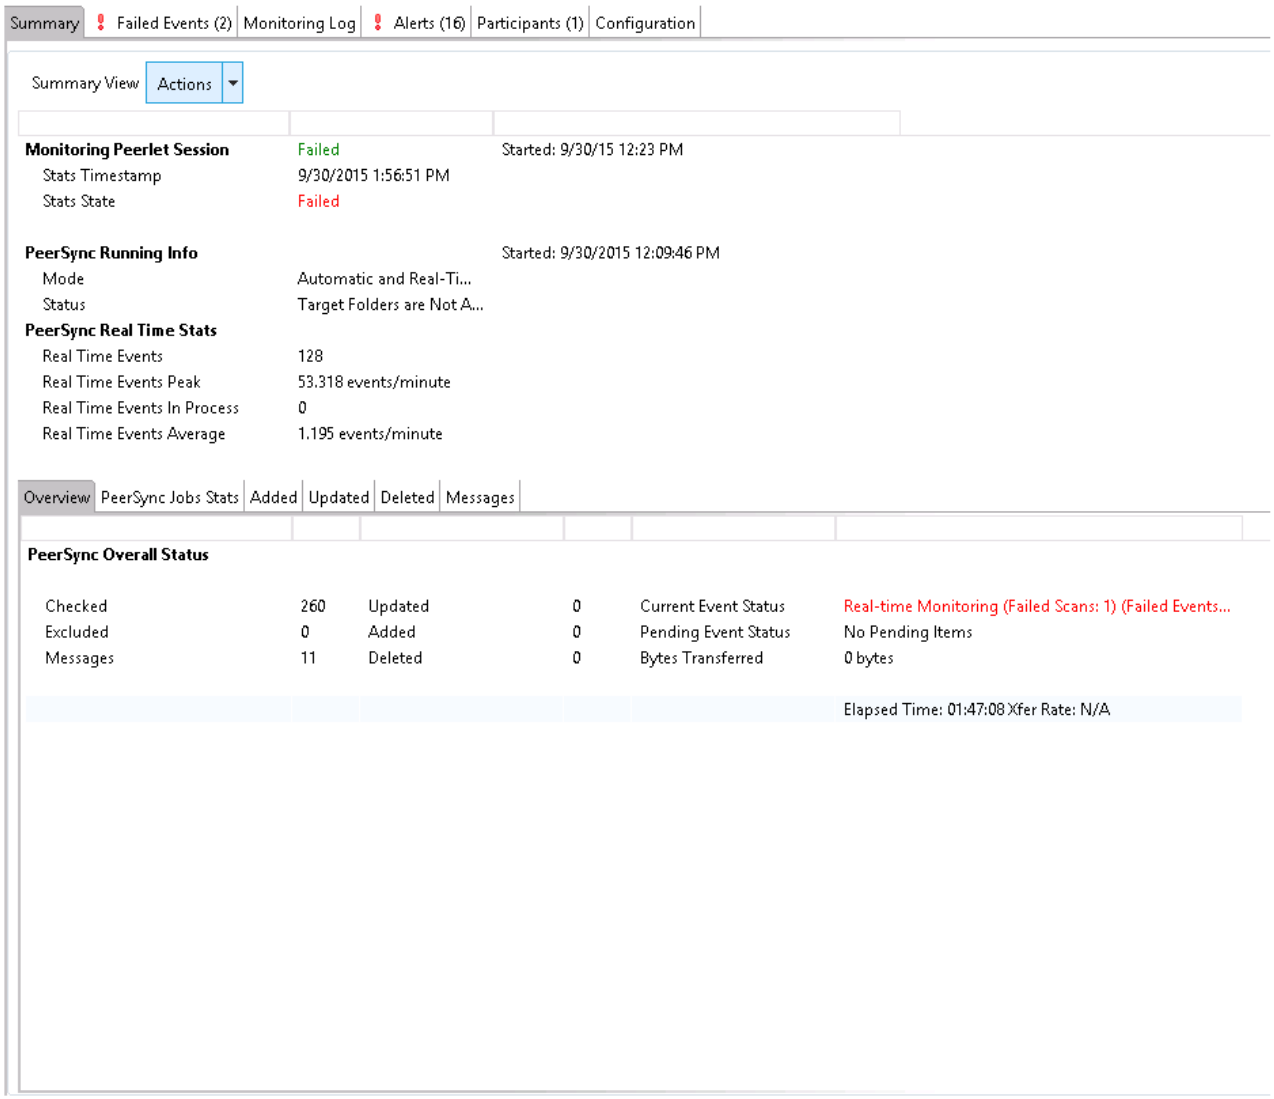

When double-clicking a PeerSync Management job, the default selected tab will be the **Summary** tab. This view will show information received by the PeerSync machine on the status of the PeerSync Management job.

Information found in this view is global to the PeerSync profile. To see PeerSync job-specific statistics, click the PeerSync Jobs Stats tab.

Information on this view is received whenever the information changes on the PeerSync machine, normally every 1 minute or so. To auto-refresh this view with the latest data, click **Enable Auto-Update** on the top right of the view, and choose a Refresh cycle. The cycle is the

not the cycle for receiving the information, just to refresh the view with the latest information received by PeerSync.

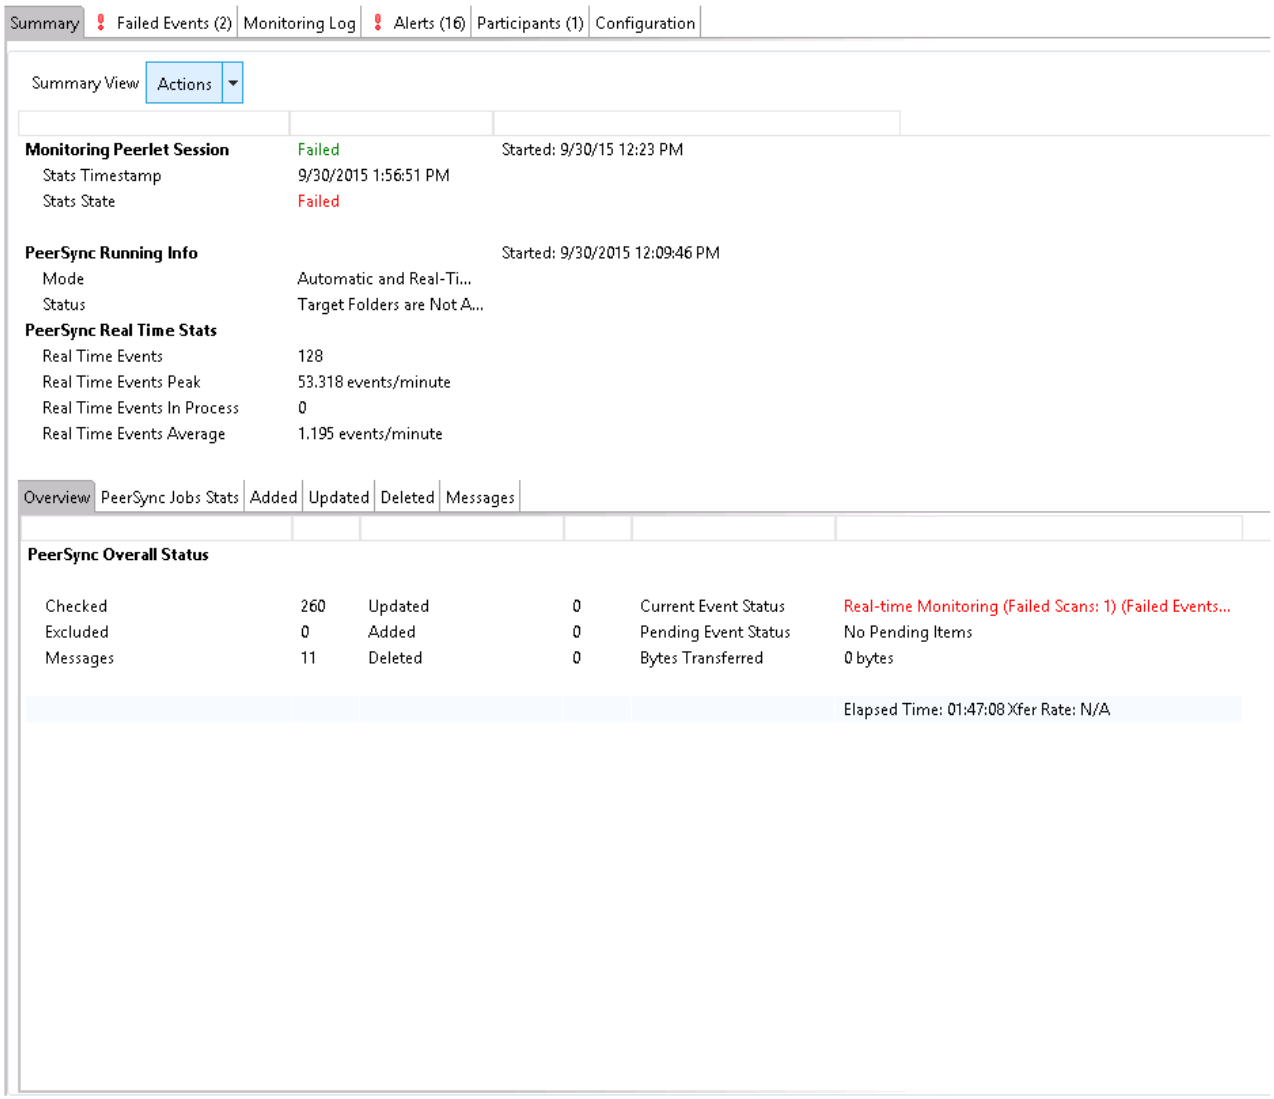

On this page, you can right-click to display the PeerSync **Actions** menu:

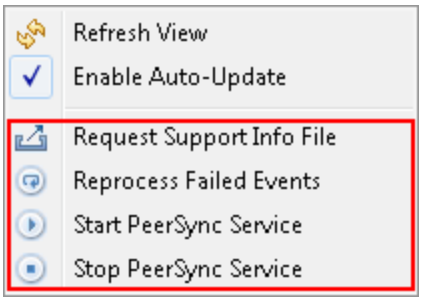

On the bottom half of the page, you will find a set of tabs showing more granular information regarding this PeerSync session.

- PeerSync Jobs Stats
- Added Files
- Updated Files
- Deleted Files
- Messages

PeerSync Jobs Stats

When clicking the **PeerSync Jobs Stats** view, a request goes out to the PeerSync machine to request job-specific statistics and return them to the Peer Management Center to be displayed. These statistics can be requested only if PeerSync is running on that machine and only if the PeerSync Management job is started on the Peer Management Center.

When the statistics are received, the view is updated with the job-specific statistics and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

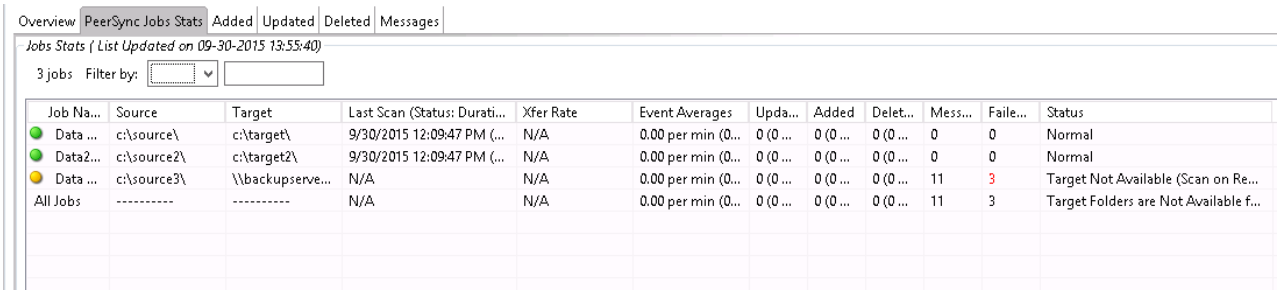

Added Files

When clicking the **Added** tab, a request goes out to the PeerSync machine to request a list of latest added files and return it to the Peer Management Center to be displayed. This list can be requested only if PeerSync is running on that machine and only if the PeerSync Management job is started on the Peer Management Center.

When the list is received, the view is updated with the latest added events processed by PeerSync, and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

This information in this table can be filtered by Path or by Job Name.

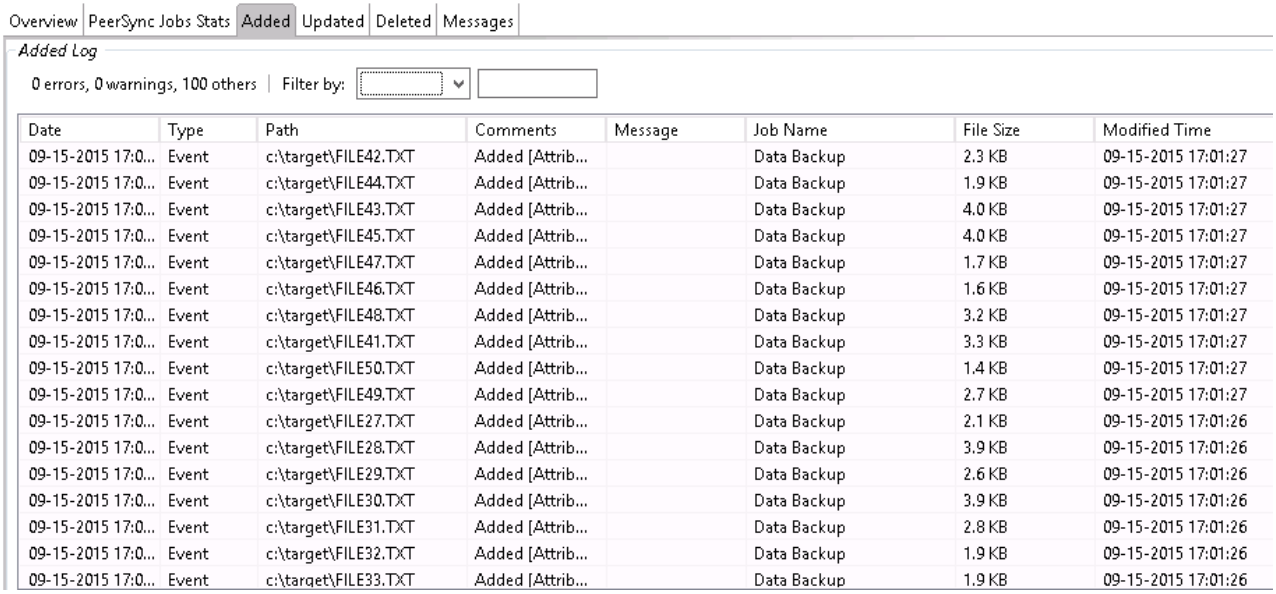

Date	Type	Path	Comments	Message	Job Name	File Size	Modified Time
09-15-2015 17:0...	Event	c:\target\FILE42.TXT	Added [Attrib...		Data Backup	2.3 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE44.TXT	Added [Attrib...		Data Backup	1.9 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE43.TXT	Added [Attrib...		Data Backup	4.0 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE45.TXT	Added [Attrib...		Data Backup	4.0 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE47.TXT	Added [Attrib...		Data Backup	1.7 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE46.TXT	Added [Attrib...		Data Backup	1.6 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE48.TXT	Added [Attrib...		Data Backup	3.2 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE41.TXT	Added [Attrib...		Data Backup	3.3 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE50.TXT	Added [Attrib...		Data Backup	1.4 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE49.TXT	Added [Attrib...		Data Backup	2.7 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE27.TXT	Added [Attrib...		Data Backup	2.1 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE28.TXT	Added [Attrib...		Data Backup	3.9 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE29.TXT	Added [Attrib...		Data Backup	2.6 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE30.TXT	Added [Attrib...		Data Backup	3.9 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE31.TXT	Added [Attrib...		Data Backup	2.8 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE32.TXT	Added [Attrib...		Data Backup	1.9 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE33.TXT	Added [Attrib...		Data Backup	1.9 KB	09-15-2015 17:01:26

Updated Files

When clicking the **Updated** tab, a request goes out to the PeerSync machine to request a list of latest updated files and return it to the Peer Management Center to be displayed. This list can be requested only if PeerSync is running on that machine and only if the PeerSync Management job is started on the Peer Management Center.

When the list is received, the view is updated with the latest updated events processed by PeerSync, and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

This information on this table can be filtered by Path or by Job Name.

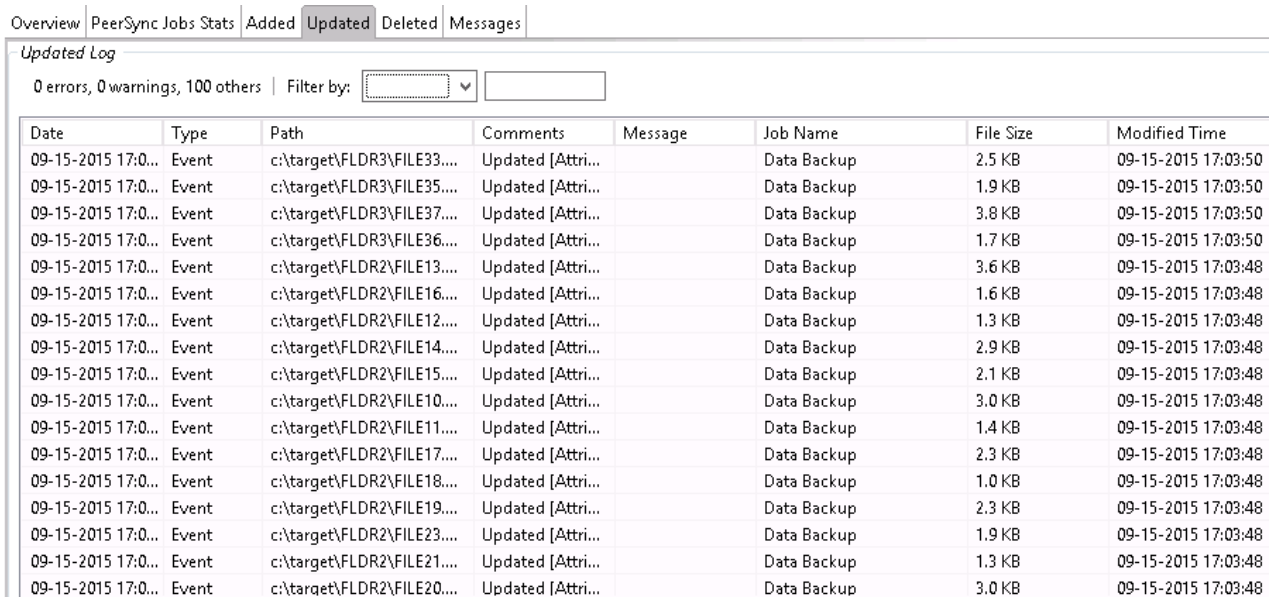

Date	Type	Path	Comments	Message	Job Name	File Size	Modified Time
09-15-2015 17:0...	Event	c:\target\FLDR3\FILE33....	Updated [Attri...		Data Backup	2.5 KB	09-15-2015 17:03:50
09-15-2015 17:0...	Event	c:\target\FLDR3\FILE35....	Updated [Attri...		Data Backup	1.9 KB	09-15-2015 17:03:50
09-15-2015 17:0...	Event	c:\target\FLDR3\FILE37....	Updated [Attri...		Data Backup	3.8 KB	09-15-2015 17:03:50
09-15-2015 17:0...	Event	c:\target\FLDR3\FILE36....	Updated [Attri...		Data Backup	1.7 KB	09-15-2015 17:03:50
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE13....	Updated [Attri...		Data Backup	3.6 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE16....	Updated [Attri...		Data Backup	1.6 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE12....	Updated [Attri...		Data Backup	1.3 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE14....	Updated [Attri...		Data Backup	2.9 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE15....	Updated [Attri...		Data Backup	2.1 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE10....	Updated [Attri...		Data Backup	3.0 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE11....	Updated [Attri...		Data Backup	1.4 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE17....	Updated [Attri...		Data Backup	2.3 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE18....	Updated [Attri...		Data Backup	1.0 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE19....	Updated [Attri...		Data Backup	2.3 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE23....	Updated [Attri...		Data Backup	1.9 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE21....	Updated [Attri...		Data Backup	1.3 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE20....	Updated [Attri...		Data Backup	3.0 KB	09-15-2015 17:03:48

Deleted Files

When clicking on the **Deleted** tab, a request goes out to the PeerSync machine to request a list of latest deleted files and return it to the Peer Management Center to be displayed. This list can be requested only if PeerSync is running on that machine and only if the PeerSync Management job is started on the Peer Management Center.

When the list is received, the view is updated with the latest deleted events processed by PeerSync, and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

This information on this table can be filtered by Path or by Job Name.

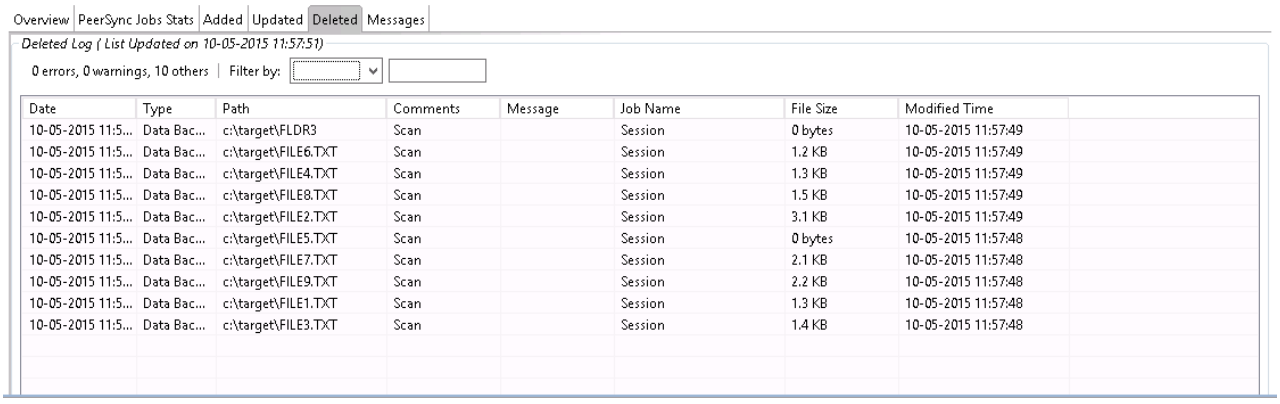

Date	Type	Path	Comments	Message	Job Name	File Size	Modified Time
10-05-2015 11:5...	Data Bac...	c:\target\FLDR3	Scan		Session	0 bytes	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE6.TXT	Scan		Session	1.2 KB	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE4.TXT	Scan		Session	1.3 KB	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE8.TXT	Scan		Session	1.5 KB	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE2.TXT	Scan		Session	3.1 KB	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE5.TXT	Scan		Session	0 bytes	10-05-2015 11:57:48
10-05-2015 11:5...	Data Bac...	c:\target\FILE7.TXT	Scan		Session	2.1 KB	10-05-2015 11:57:48
10-05-2015 11:5...	Data Bac...	c:\target\FILE9.TXT	Scan		Session	2.2 KB	10-05-2015 11:57:48
10-05-2015 11:5...	Data Bac...	c:\target\FILE1.TXT	Scan		Session	1.3 KB	10-05-2015 11:57:48
10-05-2015 11:5...	Data Bac...	c:\target\FILE3.TXT	Scan		Session	1.4 KB	10-05-2015 11:57:48

Messages

When clicking on the **Messages** tab, a request goes out to the PeerSync machine to request a list of messages/errors logged and return it to the Peer Management Center to be displayed. This list can be requested only if PeerSync is running on that machine and only if the File Synchronization job is started on the Peer Management Center.

When the info list is received the view is updated with the messages logged by PeerSync, and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

This information on this table can be filtered by Path, Job Name, or Message.

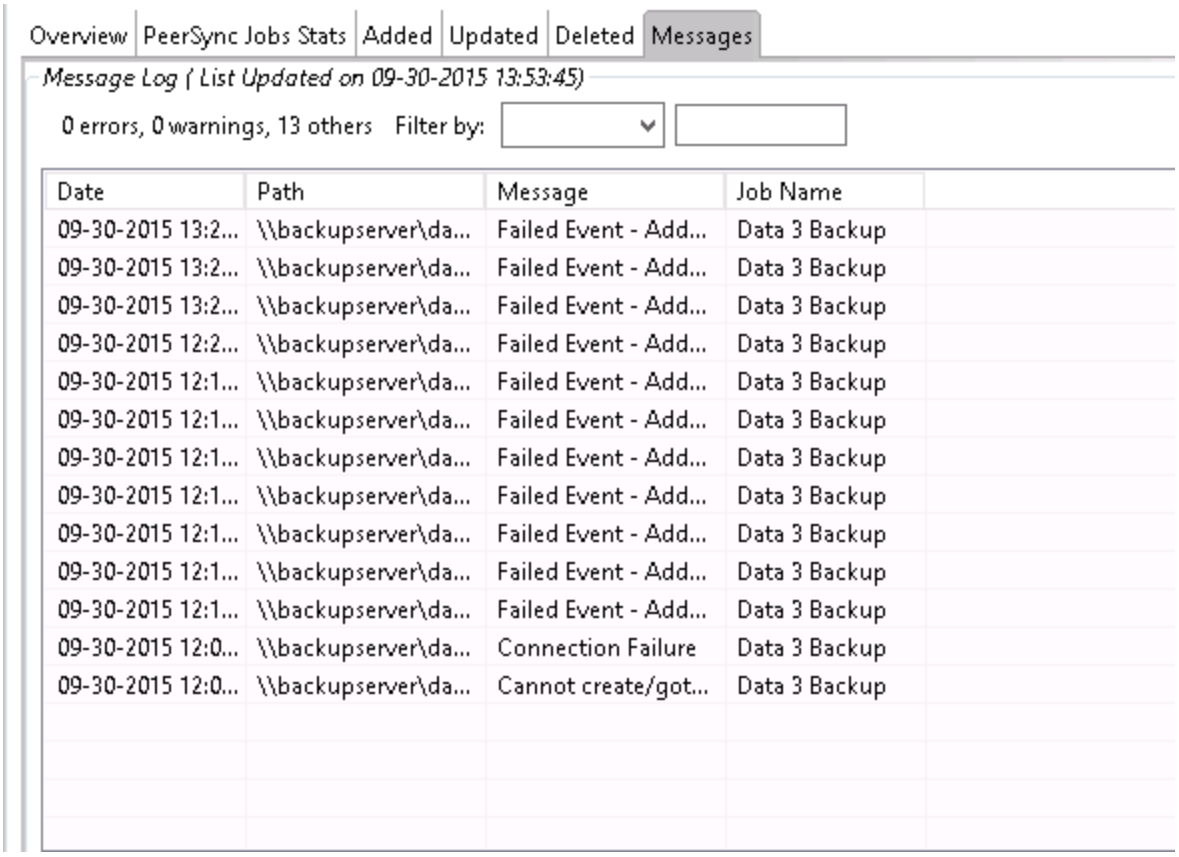

The **Failed Events** view allows you to see all those events that have failed to be processed by PeerSync. The list is populated when the PeerSync Management job starts, as well as in real-time as new failures occur. The information can be filtered by File Name.

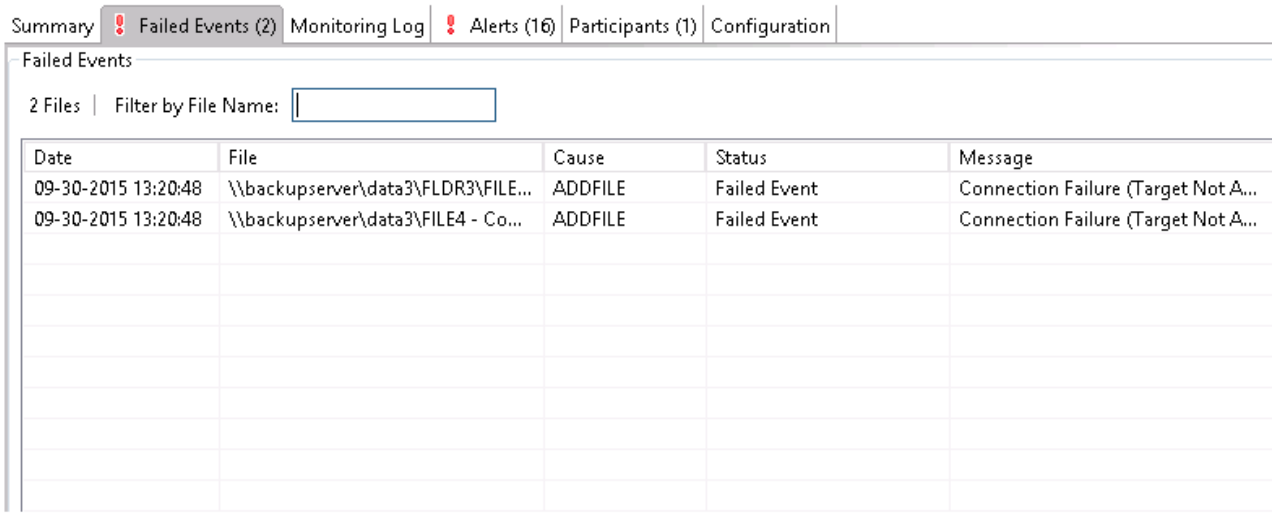

You can right-click the info table and choose to **Reprocess Failed Events.** This action will send a request to PeerSync to retry all the failed events in the list.

The **Monitoring Log** view allows you to view recent event history for the currently running PeerSync Management job based on your Logging and Alerts settings. You can specify the maximum number of events to store in the table by adjusting the **Display Events** spinner located in the top right corner of the panel. The maximum number of events that can be viewed is 3,000.

If you need to view more events or events from a prior session, then you can use the log files saved in the **Hub\logs** directory located in the installation directory. The event log files will start with **fs_event.log** and are written in a tab-delimited format. Microsoft Excel is a good tool to use to view and analyze a log file. See Logging and Alerts for more information about log files.

You can click on any column header to sort by the column. Warnings are displayed in light gray; errors are displayed in red; fatal errors are displayed in orange. Error records will also contain an error message in the **Message** column.

To change what is being logged, update the selected Event Types in the Logging and Alerts settings.

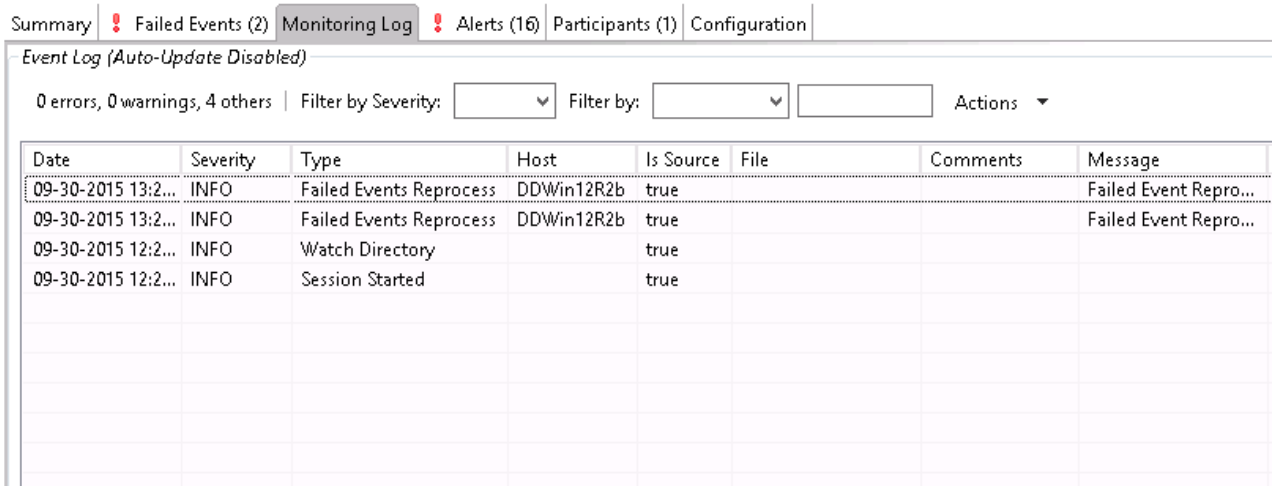

Clicking the **Actions** table menu provides the following options:

Refresh View	Refresh all information provided in the table.
Clear Events	Remove all items from the table.

The **Alerts** view allows you to view any alerts relevant to the running PeerSync Management job. Items shown here are based on the configured Alerts Severity setting on the Logging and Alerts configuration page. You can specify the maximum number of alerts to store in the table by adjusting the Display Alerts spinner located in the top right corner of the panel.

The alerts are also written to a tab-delimited file named **fs_alert.log** within the subdirectory **Hub/logs** within the installation directory of the Peer Management Center. See the Logging and Alerts settings for more information about log files.

You can click on any column header to sort by that column. For example, clicking on the **Severity** column will sort by alert severity. Warnings are displayed in light gray; errors and fatal alerts are displayed in red. A common error may be the PeerSync service is not running, which will trigger a PeerSync Quorum lost alert.

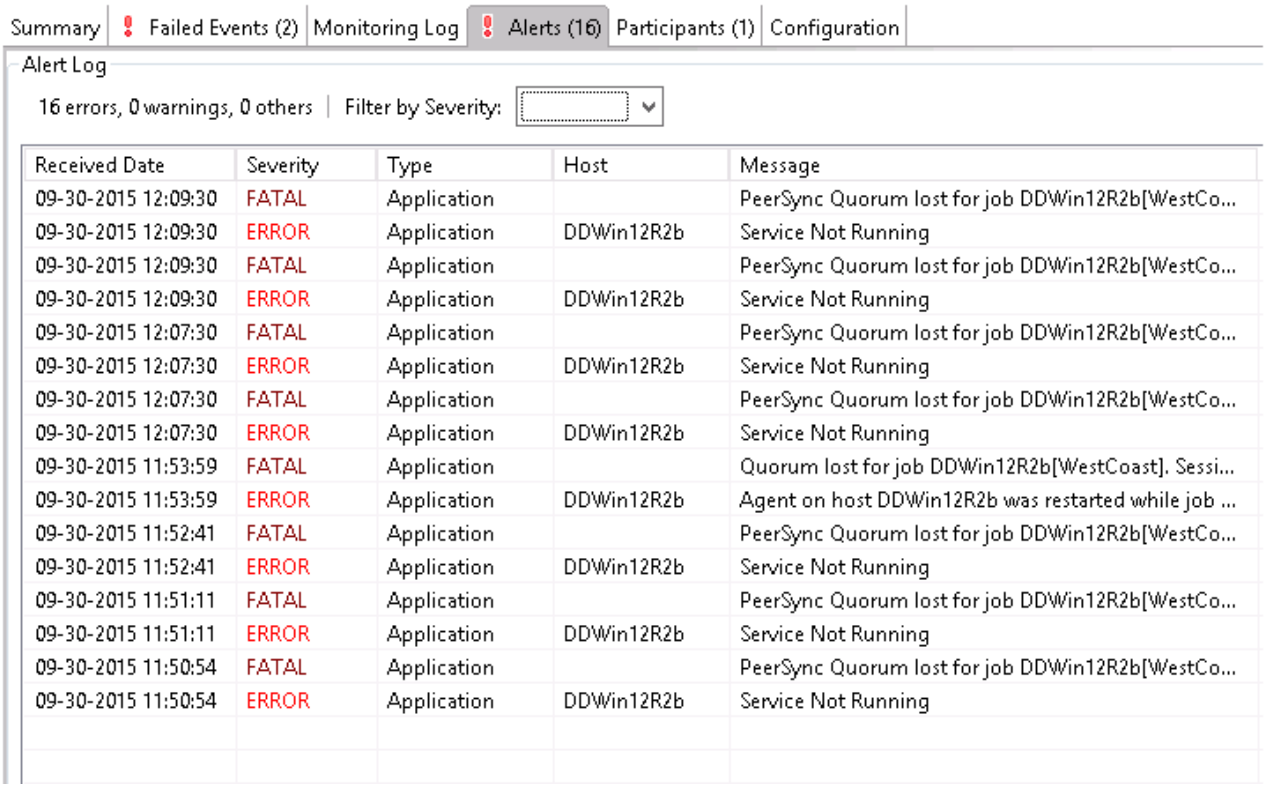

Received Date	Severity	Type	Host	Message
09-30-2015 12:09:30	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 12:09:30	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 12:09:30	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 12:09:30	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 12:07:30	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 12:07:30	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 12:07:30	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 12:07:30	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 11:53:59	FATAL	Application		Quorum lost for job DDWin12R2b[WestCoast]. Sessi...
09-30-2015 11:53:59	ERROR	Application	DDWin12R2b	Agent on host DDWin12R2b was restarted while job ...
09-30-2015 11:52:41	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 11:52:41	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 11:51:11	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 11:51:11	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 11:50:54	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 11:50:54	ERROR	Application	DDWin12R2b	Service Not Running

The following right-click menu items are unique to this table:

Refresh View	Refresh all information provided in the table. This can also be done from the right-click context menu of the table.
Clear Events	Remove all items from the table. This can also be done from the right-click context menu of the table.

The **Participants** view shows the currently configured host participant for the selected PeerSync Management job and contains a column used to display activity status occurring on the hosts. If a host has become unavailable or the PeerSync service stopped, an error message will be displayed in red next to the failed host.

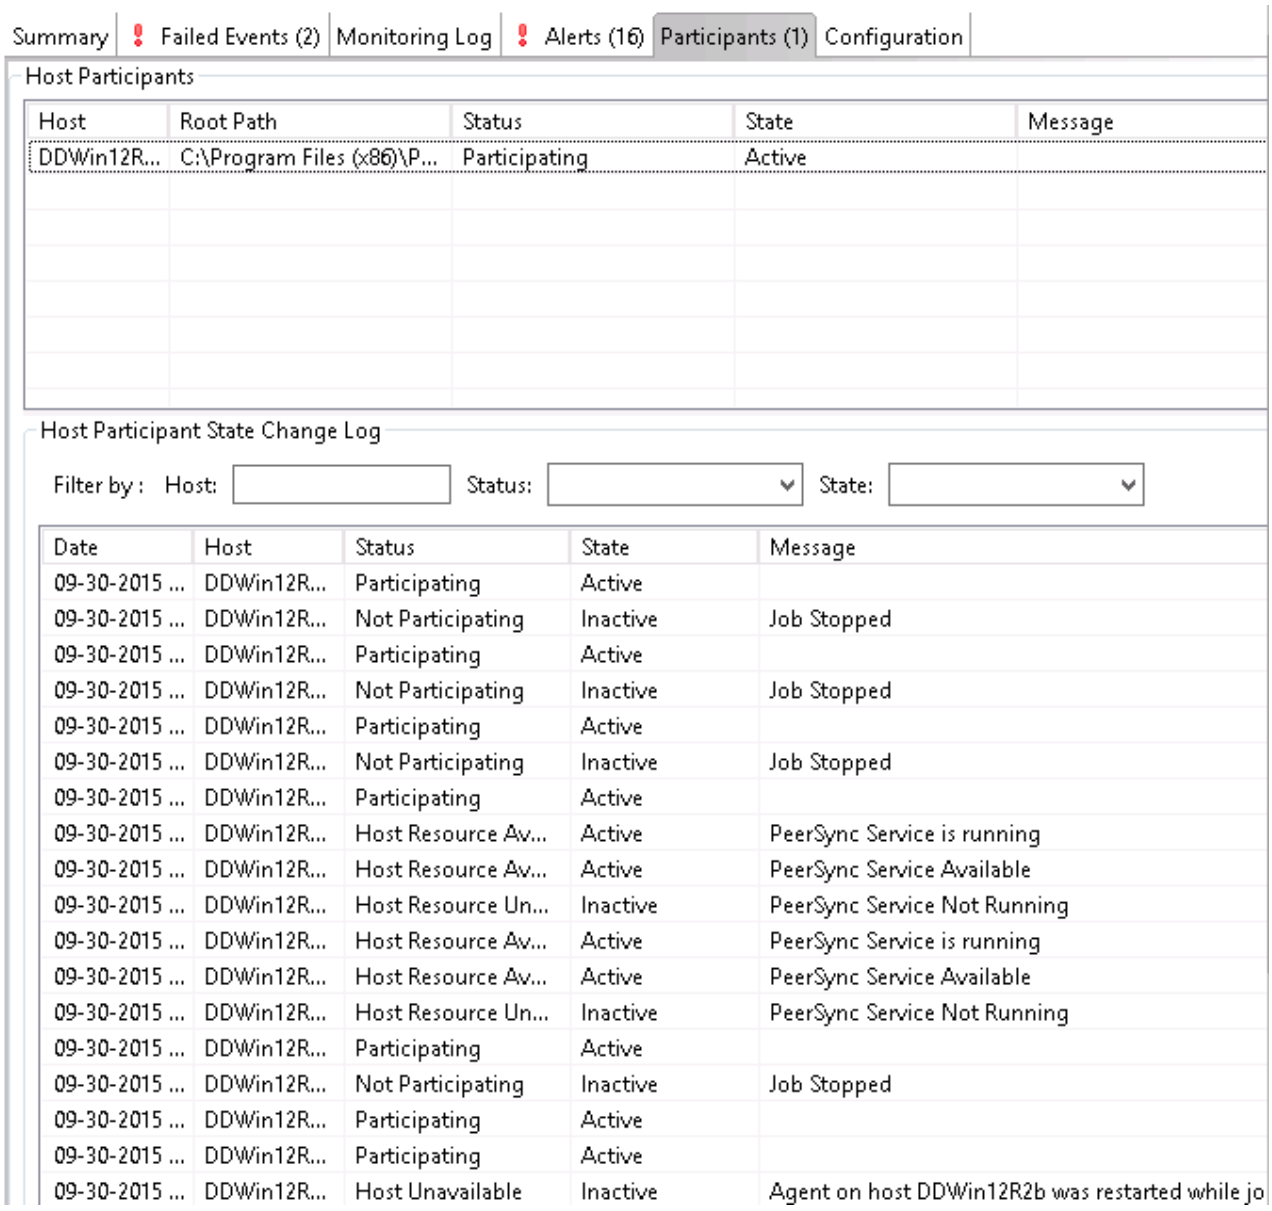

The **Participants** view also contains a table that displays the most recent host participant state changes, for example, when a host was removed from synchronization session, when a host came back online, or when the PeerSync service was stopped or started. This functionality is broken down into two parts: right-click context menu items and the **Host Participant State Change Log** view.

The **Host Participant State Change Log** is a log of all host participant status changes (e.g., Collaborating, Not Collaborating) and/or state changes (e.g., Active, Pending Restart) of a host participant. This table is currently limited to 250 rows and can be filtered by host, by status, and by state.

The following items are available in the right-click context menu for this table:

Refresh View	Refresh all information provided in the table.
Clear Events	Remove all items from the table.

The **Configuration** view displays a quick summary of all configurable items for the selected PeerSync Management job. Each page of the **File Synchronization Configuration** dialog is represented in its own part of the view and can be collapsed if desired.

Clicking **Edit this Configuration** will immediately bring you to the **PeerSync Management Configuration** dialog where you can edit the current monitoring configuration or the associated PeerSync profile.

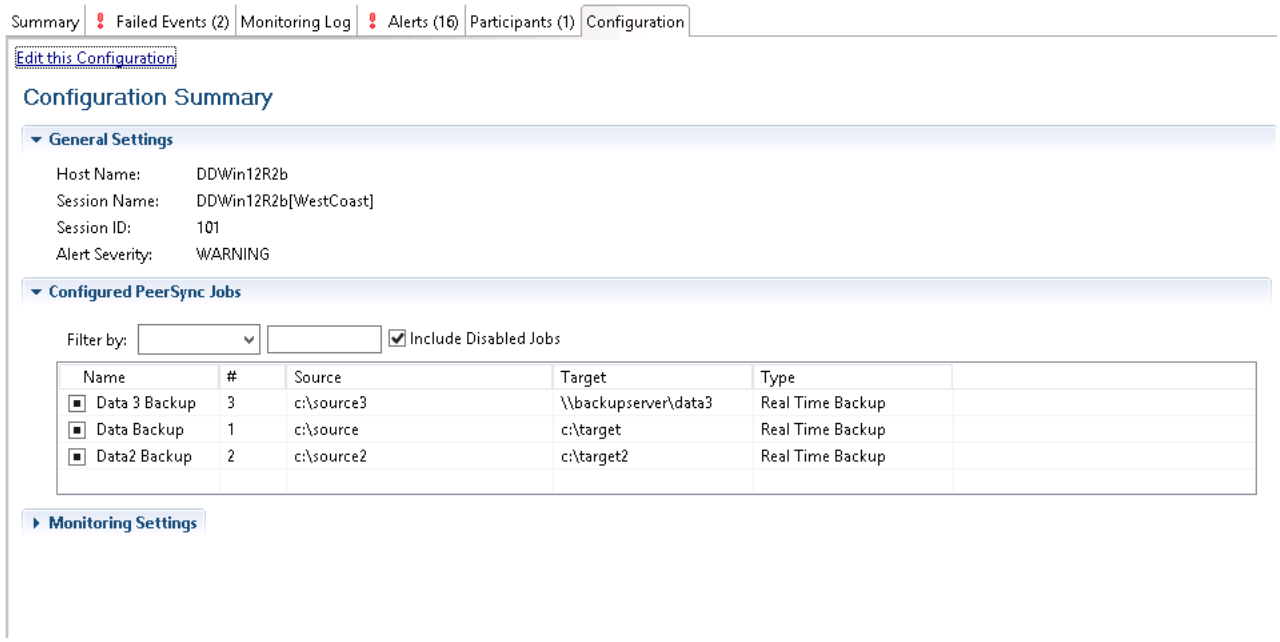

Upgrade/Reprocess Installation

From the Synchronization Summary view, you can click on one or more PeerSync Management jobs and choose **Reprocess/Upgrade Installation**. This option sends a request to the selected PeerSync instances to install/upgrade given the configured settings.

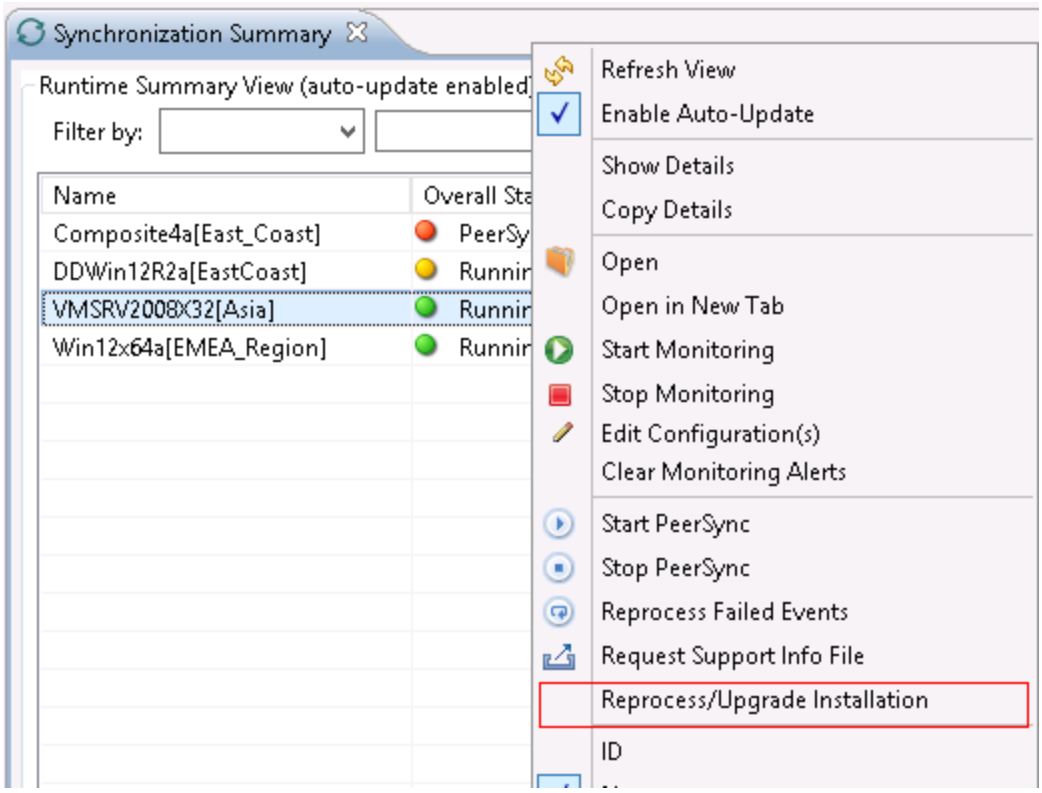

The installation settings should be common for ALL the PeerSync Management PeerSync instances in order to successfully install PeerSync.

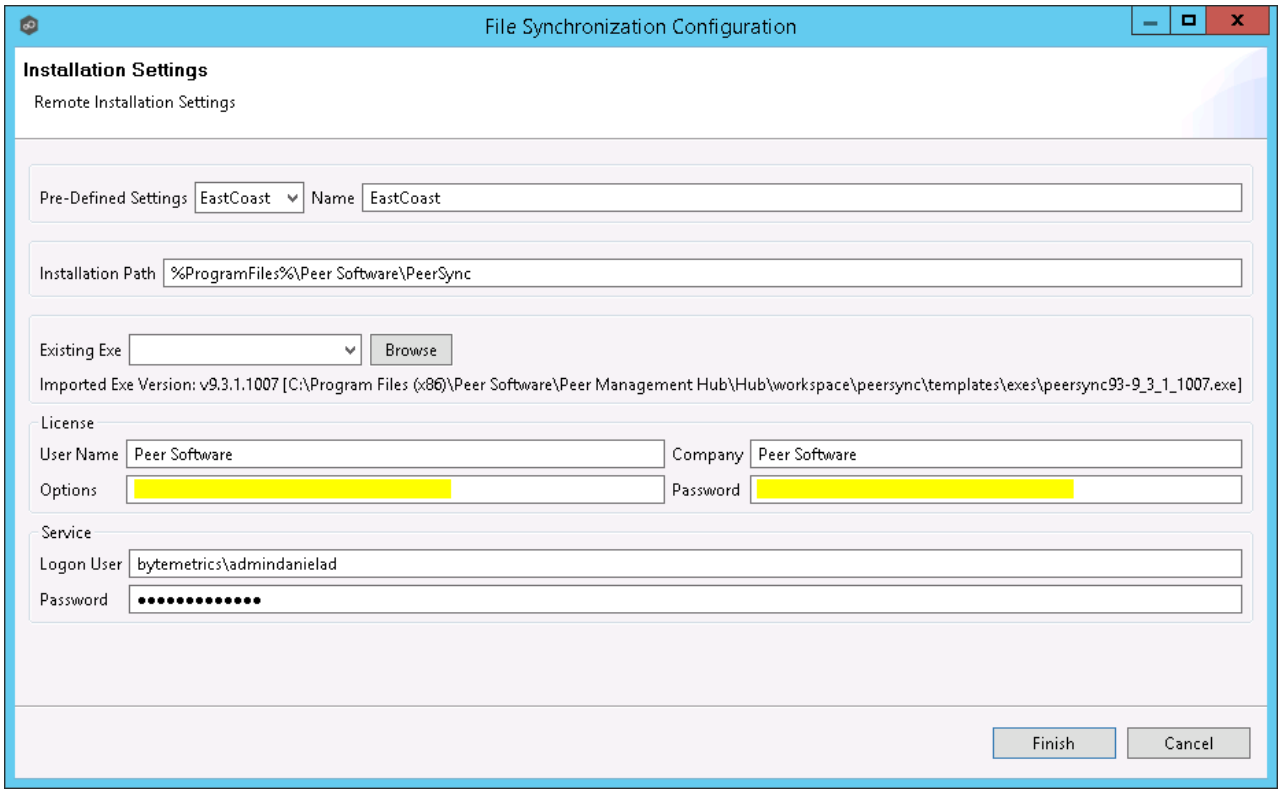

See Installation Settings for information on the settings on this page.

Index

- A -

- B -

- C -

- D -

- E -

- F -

- G -

- H -

- I -

- J -

- K -

- L -

- M -

- N -

- P -

- Q -

- R -

- S -

- T -

- U -

- V -

- W -

- Y -