Peer Global File Service
User Guide

Updated Monday, March 18, 2019

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

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 Sync 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	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 job type that enables the creation and management of DFS namespaces.
Event	A single operation performed by a user on a file server.
Failback	The process of redirecting previously displaced users from a secondary file server back to the primary after a failure state has been resolved.

Term	Definition
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	A 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 job type 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 locking job	A job type that prevents users on different file servers (usually at different locations) from opening different copies of the same file at the same time.
File Replication job	A job type that involves real-time and/or scheduled copying of files and folders from one file server to another.
File Synchronization job	A job type 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.

Term	Definition
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 the 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 the 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 (also host participant or participating Host)	A server that a Peer Agent is installed upon. When active in a job, it is called the host participant.
Initialization process	The steps executed whenever a job is started in the Peer Management Center. Each job type has a different set of steps for its initialization process.
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 the Peer 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

Term	Definition
	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. The server that the Agent is installed upon is called the host participant or simply host. Applies to File Collaboration, File Locking, 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 the 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. The 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.
Peer Management Broker	The central messaging system of Peer Global File Service. The Peer Management Broker serves to connect the Peer Management Center and the Peer Agents, forming a Peer Management Center "network" that can be cast over local or wide-area networks via TCP/IP. A Peer Management Center environment will deploy one or more Peer Management Brokers.
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,

Term	Definition
	quorum will not be met. Applies to File Collaboration, File Locking, File Replication, and File Synchronization jobs.
Real-time Event Detection	A key technology that backs all job types in the Peer Management Center. The PMC 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 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, 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.

Term	Definition
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 the Peer Management Center's user interface, each providing unique information and control. Examples: Main View, Jobs View, Peer 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 configured root folder and all subfolders on a file server that are being scanned and/or monitored by a FC, FS, FL, FR or Cloud Sync job.

Installation and Configuration

Peer Management Center can be installed in numerous ways based on your needs and environment. The Peer Management Center installation consist of two separate installers, both of which are available for download from our website:

- Peer Management Center installer, containing the Peer Management Center and the Peer Management Broker

- Peer Agent installer

Requirements

Before you get started, review the environmental requirements and platform prerequisites:

- Environmental requirements

- EMC Isilon prerequisites
- EMC Unity prerequisites
- EMC VNX/Celerra prerequisites
- NetApp 7-Mode prerequisites
- NetApp cDOT/ONTAP9+ prerequisites
- Nutanix Files prerequisites

Installing the Peer Management Center and the Peer Management Broker

Both the Peer Management Center and Peer Management Broker are packaged with the main Peer Management Center installer and by default, will be installed on the same server.

Requirements

See Requirements for more detailed information.

Software Installation and Launching

1. Run the **PMC_Installer_64win.exe** installer and follow all instructions.

2. After the installation finishes, both the Peer Management Center and Peer Management Broker will be installed. The Peer Management Broker will automatically be installed as a running Windows service and set to auto-start. The Peer Management Center is installed in three parts:

   - A Windows service that is set to auto-start.
   - A web service for granting access to the Windows service via web browsers.
   - A rich client for interacting with the Windows service. The rich client is started as a normal Windows application.

3. Start the Peer Management Center Client by launching the **PL-Hub.exe** executable located in the base installation directory. If both the Peer Management Broker and Peer Management Center Service are up and running as background services, then the Peer Management Center should successfully start. If not, please make sure that both the Peer Management Broker Service and Peer Management Center Service are running as Windows services via the Windows Service Panel (services.msc)

Secure Encrypted TLS Connections

See Enabling a Secure Encrypted TLS Connection for information.

Uninstalling

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.

Enabling a Secure Encrypted TLS Connection

By default, the Peer Management Center and Peer Management Broker are installed on the same host machine that does not require secure TLS communication between each other. However, if they are not installed on the same host machine, then a secure TLS connection between the Peer Management Center and Peer Management Broker needs to be enabled.

To enable a secure TLS connection between the Peer Management Center and Peer Management Broker:

1. Stop the Peer Management Center Service via the Windows Service Panel (services.msc).

2. Once stopped, navigate to the directory **Hub\workspace\prefs**, relative to the installation directory.

3. Within this directory, open the **com.ci.pl.hub.runtime.prefs** file in a text editor.

4. If the file does not contain a line starting with **hub.jms.providerURL**, then add the following line in its entirety:

   hub.jms.providerURL=failover\:(ssl\://localhost\:61617)?jms.alwaysSyncSend\=true

5. Otherwise, make the following changes to the line starting with **hub.jms.providerURL** (changes are bold):

   From: hub.jms.providerURL=failover\:(**tcp**\://localhost\:**61616**)?jms.alwaysSyncSend\=true

   To: hub.jms.providerURL=failover\:(**ssl**\://localhost\:**61617**)?jms.alwaysSyncSend\=true

6. Once these changes are complete, save the file, and then restart the Peer Management Center Service.

Installing a Peer Agent

You will need to install a Peer Agent on each server you plan to include in any of your jobs.

Basic Requirements

See Requirements for more detailed requirements.

Software Installation and Launching

1. Run the **P-Agent_win.exe** or **P-Agent_64win.exe** installer on the target server and follow all instructions.

2. During installation you will need to specify the Peer Management Broker Host Name (computer name, fully qualified domain name, or IP Address) of the server where the Peer Management Broker is running, as well as the configured TCP/IP port number (the default port for TLS communication is 61617).

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

4. Make sure that the Peer Agent was able to successfully connect to the Peer Management Broker by going to the Peer Agent installation folder, opening the **output.log** text file, and making sure that **Ready** is displayed on the first line.

Secure Encrypted TLS Connections

By default, the Peer Agent is installed with TLS encryption enabled, where the Peer Agent connects to the 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.

If AES-256 support is required, please contact support@peersoftware.com to obtain the necessary installers.

Uninstalling

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.

Configuring the Web Interface

This section describes:

- Setting Up the Web Interface
- Securing Access to the Web Interface
- Modifying the Web Interface Configuration
- Accessing the Web Interface
- Managing Web Interface User Accounts

Setting Up the Web Interface

During the Peer Management Center installation process, you configure access to the web interface. You can modify these options at a later date. See Modifying the Web Interface Configuration.

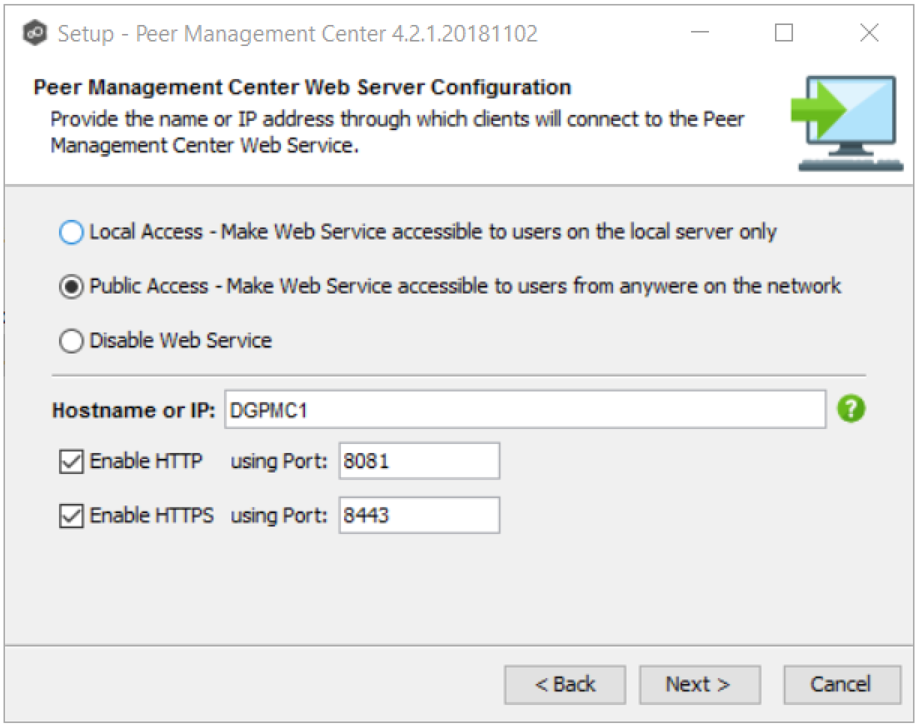

Peer Management Center Web Server Configuration

Provide the name or IP address through which clients will connect to the Peer Management Center Web Service.

○ Local Access - Make Web Service accessible to users on the local server only

● Public Access - Make Web Service accessible to users from anywere on the network

○ Disable Web Service

Hostname or IP: DGPMC1

☑ Enable HTTP using Port: 8081

☑ Enable HTTPS using Port: 8443

The web interface configuration options are shown in the following table.

Local Access	Allows access to the web interface only when remotely connected into and using a web browser on the local Peer Management Center server.
Public Access	Allows access to the web interface via the configured hostname or IP address. Note that **Public Access** does not necessarily mean that anyone on the Internet will be able to access the web interface. This access should be further limited via NAT and network firewall policies. As an option, "0.0.0.0" can be used in the **Hostname** or **IP** field in conjunction with the **Public Access** option to fully open up web access on your network.
Disable Web Service	Completely disables the web interface and sets the Peer Management Center Web Service to manual.
Hostname or IP	Identifies the hostname or IP address via which clients can access the web interface. If **Local Access** is set, this will be forced to use "localhost".

Enable HTTP (using Port)	Enables HTTP access to the web interface using the specified port.
Enable HTTPS (using Port)	Enables HTTPS access to the web interface using the specified port and a built-in TLS certificate. See Securing Access to the Web Interface for details on changing TLS certificates.

Modifying the Web Interface Configuration

You can modify the configuration of the PMC web interface by modifying a config.ini file.

To modify the web interface configuration:

1. Stop the Peer Management Center Web Service via the Windows Service Panel (services.msc).

2. Use Notepad to modify the **config.ini** file.

   The **config.ini** file is located in Peer Management Center_INSTALL_FOLDER\Hub\web-configuration\ (where PMC_INSTALL_FOLDER represents the root installation directory of Peer Management Center).

3. Modify the settings. See Modifying the Web Interface Configuration for details.

4. Once modifications are complete, save the file

5. Restart the **Peer Management Center Web Service**.

Modifying the config.ini File

The important items to configure in the config.ini file are:

org.eclipse.equinox.http.jetty.http.enabled	Set to "true" to enable HTTP access to the web interface. If set to "true", the **org.eclipse.equinox.http.jetty.http.host** and **org.osgi.service.http.port** items must also be configured in order to enable HTTP access to the web interface. If set to "false", HTTP access is disabled and the other HTTP-related settings are ignored.

org.eclipse.equinox.http.jetty.http.host	Set to the hostname or IP address via which the web interface can be accessed using HTTP. Set to "localhost" to enable local access only for HTTP.
org.osgi.service.http.port	Set to the port to be used for HTTP access.
org.eclipse.equinox.http.jetty.https.enabled	Set to "true" to enable HTTPS access to the web interface. If set to "true", the **org.eclipse.equinox.http.jetty.https.host** and **org.osgi.service.http.port.secure** items must also be configured in order to enable HTTPS access to the web interface. If set to "false", HTTPS access is disabled and the other HTTPS-related settings are ignored.
org.eclipse.equinox.http.jetty.https.host	Set to the hostname or IP address via which the web interface can be accessed using HTTPS. Set to "localhost" to enable local access only for HTTPS.
org.osgi.service.http.port.secure	Set to the port to be used for HTTPS access.

- All settings listed above must be followed by an "=" and a value. For example, to enable HTTP access, the line in the **config.ini** file with **org.eclipse.equinox.http.jetty.http.enabled** should look like:

```
org.eclipse.equinox.http.jetty.http.enabled=true
```

- HTTP and HTTPS are configured independently of one another in the **config.ini** file and as such, can be set to different modes. For example, HTTPS could be configured in a public mode, while HTTP is set to private ("localhost").

- DO NOT modify any other settings in the **config.ini**. Doing so may result in the inability of the web interface to start**.**

- Duplicate entries in the **config.ini** file may also result in the inability of the web interface to start.

Accessing the Web Interface

Once Peer Management Center has been installed and all services have been started, you can access the web interface.

To access the PMC web interface:

1. Open a web browser and enter the following URL:

   http://localhost:8081

Note that the exact URL will vary depending on the settings you selected during the installation process (for example: http vs https, appropriate hostname or IP, and appropriate port). These options are described in Setting Up the Web Interface.

2. In the page that appears, select the **Peer Management Center Portal** link.

   The login page is displayed.

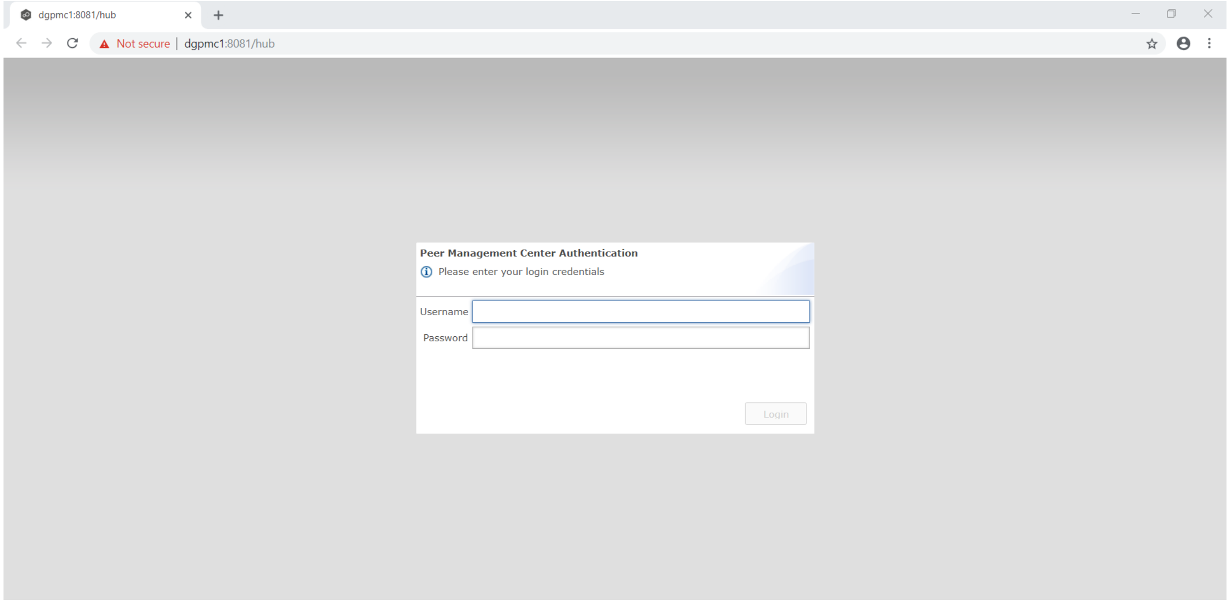

Peer Management Center Authentication

Please enter your login credentials

3. Enter a user name and password.

   The default user name is **admin**; the default password is **password**. We highly recommend that you change the password. See User Management for more information on changing account passwords.

4. Click **Login**.

   If logged in with an **admin** account, the following is displayed:

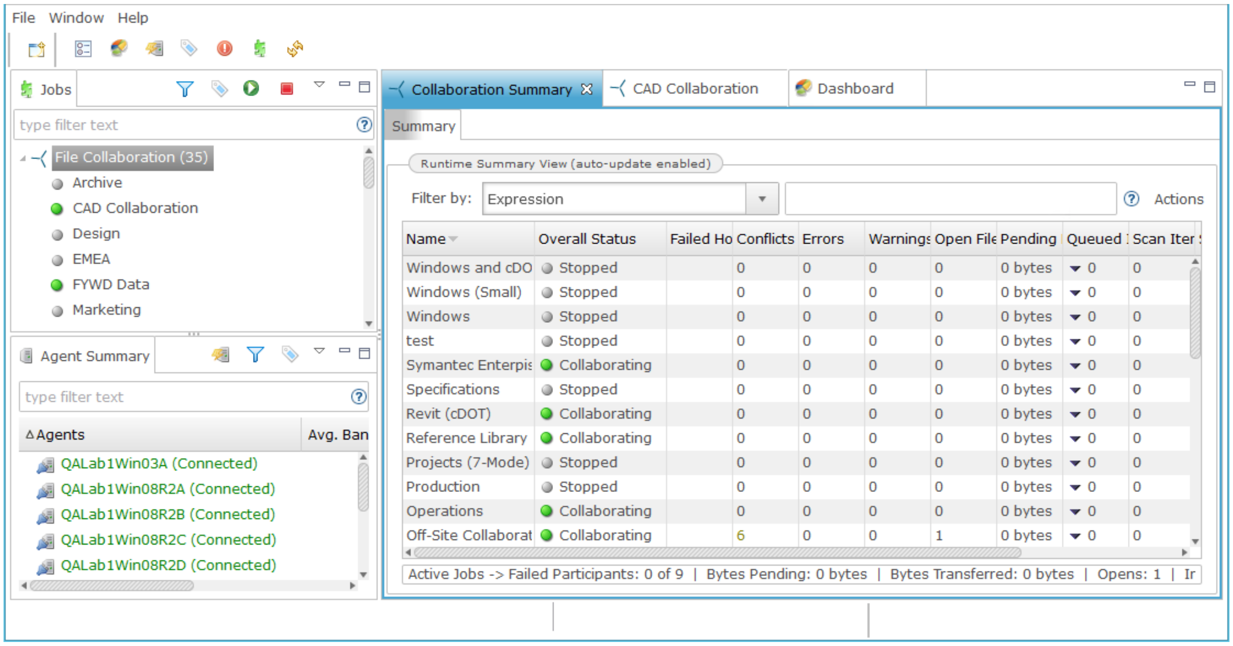

File Collaboration (35)
- ● Archive
- ● CAD Collaboration
- ● Design
- ● EMEA
- ● FYWD Data
- ● Marketing

Agents	Avg. Ban
QALab1Win03A (Connected)	
QALab1Win08R2A (Connected)	
QALab1Win08R2B (Connected)	
QALab1Win08R2C (Connected)	
QALab1Win08R2D (Connected)	

Name	Overall Status	Failed Ho	Conflicts	Errors	Warnings	Open File	Pending	Queued	Scan Iter
Windows and cDO	● Stopped		0	0	0	0	0 bytes	▾ 0	0
Windows (Small)	● Stopped		0	0	0	0	0 bytes	▾ 0	0
Windows	● Stopped		0	0	0	0	0 bytes	▾ 0	0
test	● Stopped		0	0	0	0	0 bytes	▾ 0	0
Symantec Enterpis	● Collaborating		0	0	0	0	0 bytes	▾ 0	0
Specifications	● Stopped		0	0	0	0	0 bytes	▾ 0	0
Revit (cDOT)	● Collaborating		0	0	0	0	0 bytes	▾ 0	0
Reference Library	● Collaborating		0	0	0	0	0 bytes	▾ 0	0
Projects (7-Mode)	● Stopped		0	0	0	0	0 bytes	▾ 0	0
Production	● Stopped		0	0	0	0	0 bytes	▾ 0	0
Operations	● Collaborating		0	0	0	0	0 bytes	▾ 0	0
Off-Site Collaborat	● Collaborating		6	0	0	1	0 bytes	▾ 0	0

Active Jobs -> Failed Participants: 0 of 9 | Bytes Pending: 0 bytes | Bytes Transferred: 0 bytes | Opens: 1 | Ir

Securing Access to the Web Interface

There are several important things to keep in mind when it comes to securing access to Peer Management Center's web interface:

- The default **admin** account password should be changed immediately. For details, see User Management.

- Access to the web interface can be in the form of both HTTP and HTTPS. The latter will ensure that all communication between the client browser and the service hosting the web interface is encrypted. Regardless of which is enabled, the hostname or IP address through which clients can reach the web interface can be configured to limit access. See Setting Up the Web Interface for more details.

- While HTTPS access to the web interface is secured out of the box with a built-in certificate, this certificate can be swapped for a custom one. For more details on this process, please contact Peer Software's support team via email: support@peersoftware.com

Managing Web Interface User Accounts

Management of users with access to Peer Management Center's web interface can be performed through either the Peer Management Center's rich client or through an admin account logged into the web interface. See User Management for more information about managing web interface accounts.

Peer Management Center User Interface

The Peer Management Center is a container for configuring and deploying jobs. The Peer Management Center graphical user interface enables you to create, view, edit, and delete your File Collaboration and Cloud Sync jobs, as well as view runtime information for running jobs.

The Peer Management Center has two graphical user interface options:

- A rich client installed and run on the server running the Peer Management Center.
- A web service that, when configured, can be accessed from remote systems via a web browser.

For more information, see:

- Web Interface
- Menus
- Views
- Job Alerts View
- Peer Agent Summary View
- Alerts View

Rich Client Interface

After starting up the Peer Management Center Client, the interface below is displayed. The interface can be divided into four quadrants: Jobs, Agents, Summaries, and Alerts. Each quadrant displays information in sub-panels called views. The interface also contains menus and a toolbar.

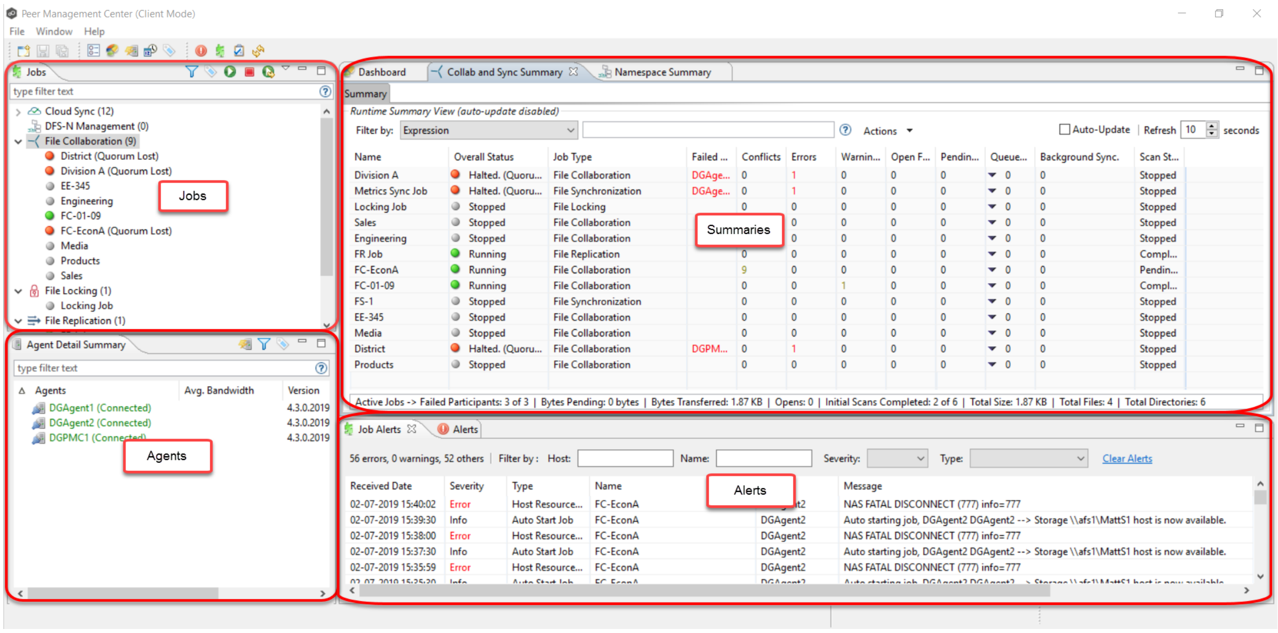

type filter text

- Cloud Sync (12)
- DFS-N Management (0)
- File Collaboration (9)
  - District (Quorum Lost)
  - Division A (Quorum Lost)
  - EE-345
  - Engineering
  - FC-01-09
  - FC-EconA (Quorum Lost)
  - Media
  - Products
  - Sales
- File Locking (1)
  - Locking Job
- File Replication (1)

Agents	Avg. Bandwidth	Version
DGAgent1 (Connected)		4.3.0.2019
DGAgent2 (Connected)		4.3.0.2019
DGPMC1 (Connected)		4.3.0.2019

Runtime Summary View (auto-update disabled)

Filter by: Expression   Actions   ☐ Auto-Update | Refresh 10 seconds

Name	Overall Status	Job Type	Failed ...	Conflicts	Errors	Warnin...	Open F...	Pendin...	Queue...	Background Sync.	Scan St...
Division A	Halted. (Quoru...	File Collaboration	DGAge...	0	1	0	0	0	▼ 0	0	Stopped
Metrics Sync Job	Halted. (Quoru...	File Synchronization	DGAge...	0	1	0	0	0	▼ 0	0	Stopped
Locking Job	Stopped	File Locking		0	0	0	0	0	▼ 0	0	Stopped
Sales	Stopped	File Collaboration			0	0	0	0	▼ 0	0	Stopped
Engineering	Stopped	File Collaboration			0	0	0	0	▼ 0	0	Stopped
FR Job	Running	File Replication		0	0	0	0	0	▼ 0	0	Compl...
FC-EconA	Running	File Collaboration		9	0	0	0	0	▼ 0	0	Pendin...
FC-01-09	Running	File Collaboration		0	0	1	0	0	▼ 0	0	Compl...
FS-1	Stopped	File Synchronization		0	0	0	0	0	▼ 0	0	Stopped
EE-345	Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stopped
Media	Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stopped
District	Halted. (Quoru...	File Collaboration	DGPM...	0	1	0	0	0	▼ 0	0	Stopped
Products	Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stopped

Active Jobs -> Failed Participants: 3 of 3 | Bytes Pending: 0 bytes | Bytes Transferred: 1.87 KB | Opens: 0 | Initial Scans Completed: 2 of 6 | Total Size: 1.87 KB | Total Files: 4 | Total Directories: 6

56 errors, 0 warnings, 52 others | Filter by : Host: Name: Severity: Type: Clear Alerts

Received Date	Severity	Type	Name		Message
02-07-2019 15:40:02	Error	Host Resource...	FC-EconA	t2	NAS FATAL DISCONNECT (777) info=777
02-07-2019 15:39:30	Info	Auto Start Job	FC-EconA	DGAgent2	Auto starting job, DGAgent2 DGAgent2 --> Storage \\afs1\MattS1 host is now available.
02-07-2019 15:38:00	Error	Host Resource...	FC-EconA	DGAgent2	NAS FATAL DISCONNECT (777) info=777
02-07-2019 15:37:30	Info	Auto Start Job	FC-EconA	DGAgent2	Auto starting job, DGAgent2 DGAgent2 --> Storage \\afs1\MattS1 host is now available.
02-07-2019 15:35:59	Error	Host Resource...	FC-EconA	DGAgent2	NAS FATAL DISCONNECT (777) info=777
02-07-2019 15:35:30	Info	Auto Start Job	FC-EconA	DGAgent2	Auto starting job, DGAgent2 DGAgent2 --> Storage \\afs1\MattS1 host is now available.

Web Interface

You can manage and monitor jobs via a robust web interface. Unlike many other web management consoles, Peer Management Center's web interface 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 properly configured, the web interface allows system administrators to manage Peer Management Center's jobs from any location without the need to remotely log in to the Peer Management Center server.

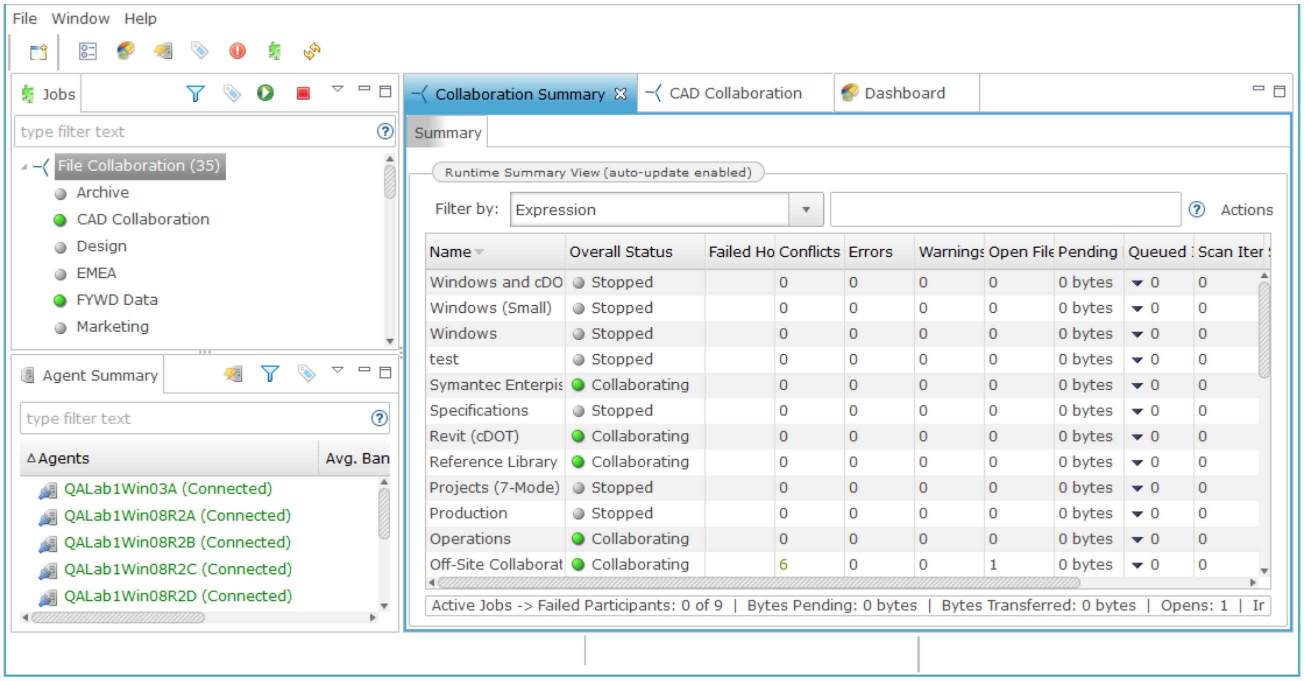

type filter text

- File Collaboration (35)
  - Archive
  - CAD Collaboration
  - Design
  - EMEA
  - FYWD Data
  - Marketing

type filter text

Agents	Avg. Ban
QALab1Win03A (Connected)	
QALab1Win08R2A (Connected)	
QALab1Win08R2B (Connected)	
QALab1Win08R2C (Connected)	
QALab1Win08R2D (Connected)	

Runtime Summary View (auto-update enabled)

Filter by: Expression   Actions

Name	Overall Status	Failed Ho	Conflicts	Errors	Warnings	Open File	Pending	Queued	Scan Iter
Windows and cDO	Stopped		0	0	0	0	0 bytes	▼ 0	0
Windows (Small)	Stopped		0	0	0	0	0 bytes	▼ 0	0
Windows	Stopped		0	0	0	0	0 bytes	▼ 0	0
test	Stopped		0	0	0	0	0 bytes	▼ 0	0
Symantec Enterpis	Collaborating		0	0	0	0	0 bytes	▼ 0	0
Specifications	Stopped		0	0	0	0	0 bytes	▼ 0	0
Revit (cDOT)	Collaborating		0	0	0	0	0 bytes	▼ 0	0
Reference Library	Collaborating		0	0	0	0	0 bytes	▼ 0	0
Projects (7-Mode)	Stopped		0	0	0	0	0 bytes	▼ 0	0
Production	Stopped		0	0	0	0	0 bytes	▼ 0	0
Operations	Collaborating		0	0	0	0	0 bytes	▼ 0	0
Off-Site Collaborat	Collaborating		6	0	0	1	0 bytes	▼ 0	0

Active Jobs -> Failed Participants: 0 of 9 | Bytes Pending: 0 bytes | Bytes Transferred: 0 bytes | Opens: 1 | Ir

Roles

The web interface includes a role-based log-in system with two predefined roles:

- **admin** - This role has complete access to all functionality found in the Peer Management Center's rich client.

- **helpdesk** - This role has read-only view of jobs and the ability to release conflicts for any running jobs.

  Specifically, users with **helpdesk** accounts are limited to read-only access of the following:

  - The Jobs view
  - The Collaboration Summary view
  - The Summary and Session tabs of each job.

  In addition, **helpdesk** accounts have read-write access to the File Conflicts tab of each job, with the ability to release conflicts.

For more information about the web interface, see:

- Setting Up the Web Interface
- Using the Web Interface
- Securing Access to the Web Interface
- Managing User Accounts

Menus

The following menu options are available in the main window of the Peer Management Center Client:

File Menu

New	Presents a list of job types, from which you can create a new job.

Save/Save All	This button will be enabled if any of the open jobs have been modified. Selecting **Save** will result in the currently open and selected job to be saved to disk. **Save All** saves all open and modified jobs to disk.
Exit	Selecting this option will close the Peer Management Center Client application. Note that as long as the Peer Management Center Service remains running, all running jobs will continue to operate.

Window Menu

Open Perspective	Open a predefined layout of views geared towards a specific purpose. For example, one perspective is for job creation and management, while another is for managing Peer Agents.
Reset Perspective	Selecting this option will reset all current windows, views, and editors to their default size and layout.
View Dashboard	Shows the Dashboard Summary view, which displays metrics and key performance indicators from all running File Collaboration Jobs, File Synchronization Jobs, and Agents.
View Agent Detail Summary	The Peer Agent Detail Summary is a panel that 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.
View Alerts	Opens the Alerts view, which displays Peer Management Center alerts such as Peer Agent connection status changes.
View Job Alerts	Opens the Job Alerts view, which displays alerts such as job restarts.
View Progress	Opens the Progress view, which displays information pertaining to any running background tasks within the Peer Management Center.
Preferences	Opens the Preferences dialog, allowing the user to configure settings for the Peer Management Center, as well as global settings for File Collaboration and Cloud Sync jobs.

Refresh	Refreshes all current views and tabs.

Help Menu

User Guide	Selecting this option will open this help system.
Download Peer Agent Installer	This operation takes you to our website where you can download the Peer Agent installer compatible with this version of the Peer Management Center.
Retrieve PMC/Agent Logs	This operation will collect and retrieve all useful log files for specified Peer Agents, the Peer Management Center, and all configured jobs. All of this information will be assembled into a single encrypted zip file that can optionally be uploaded to our technical support team. The collection and retrieval of the log and support files will be performed in the background, which might take awhile depending on content size and network speed. Upon completion, you will be notified and will be able to view the zip file yourself.
Retrieve Broker Statistics	This will display detailed statistical information about all messaging that has transpired for all connections (Peer Agents and the Peer Management Center) to the Peer Management Broker.
Thread Dump	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 of these can be used by our technical support to debug certain issues.
Generate Memory Dump File	This will generate a memory dump of the running Peer Management Center Client and Service, which can be used by our technical support to debug certain issues.
Compress DB on Restart	Check this option in cases where the database consumes a large amount of disk space. This option will compress the database upon restart of the Peer Management Center Service.
About Peer Management Center	Displays version information about the Peer Management Center along with which components are installed.

Tables

Table Detail Viewer

Most tables shown throughout the Peer Management Center support double-clicking on any row. This action brings up a pop-up dialog containing all of the details pertaining to the information in that row, for example:

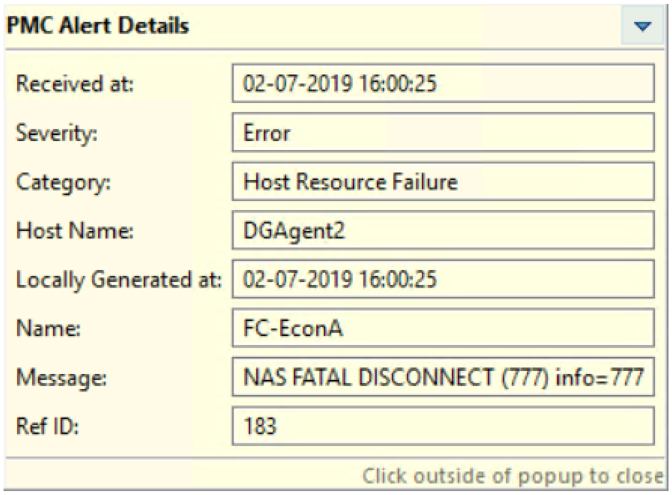

PMC Alert Details

Received at:	02-07-2019 16:00:25
Severity:	Error
Category:	Host Resource Failure
Host Name:	DGAgent2
Locally Generated at:	02-07-2019 16:00:25
Name:	FC-EconA
Message:	NAS FATAL DISCONNECT (777) info=777
Ref ID:	183

Click outside of popup to close

In addition, most right-click context menus contain the ability to copy this detailed information on one or more rows all at the same time. This information can then be pasted into any document editor.

Toolbar

The Peer Management Center toolbar buttons include:

New Job	Initiates the process of creating a new job.
User Preferences	Opens the Preferences window, allowing the user to configure global settings for the Peer Management Center, as well as settings for individual job types.

View Dashboard	Displays the Dashboard Summary view, which displays metrics and key performance indicators from all running File Collaboration jobs, Cloud Sync jobs, File Synchronization jobs, and Agents.
View Agent Detail Summary	Shows the Peer Agent Detail Summary view, which displays a list of all known Peer Agents deployed and their detailed status information, and can be used to assess the health of the environment.
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.
View Alerts	Opens the Alert view, which displays Peer Management Center alerts such as Peer Agent connection status changes.
View Job Alerts	Opens the Job Alerts view, which displays alerts such as job restarts.
Refresh	Refreshes all current views and tabs.

Views

The Peer Management Center Client interface can be divided into four quadrants: Jobs, Agents, Summaries, and Alerts. Each quadrant information in sub-panels called views.

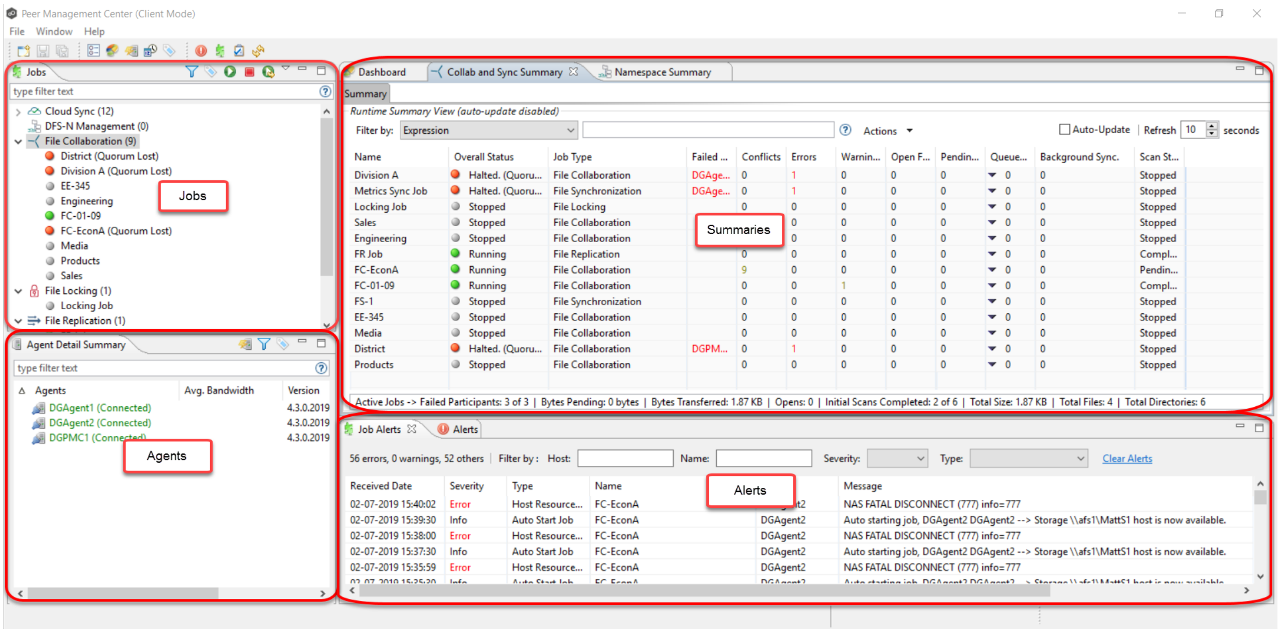

Agents	Avg. Bandwidth	Version
DGAgent1 (Connected)		4.3.0.2019
DGAgent2 (Connected)		4.3.0.2019
DGPMC1 (Connected)		4.3.0.2019

Runtime Summary View (auto-update disabled)

Filter by: Expression | Actions | ☐ Auto-Update | Refresh 10 seconds

Name	Overall Status	Job Type	Failed ...	Conflicts	Errors	Warnin...	Open F...	Pendin...	Queue...	Background Sync.	Scan St...
Division A	Halted. (Quoru...	File Collaboration	DGAge...	0	1	0	0	0	▼ 0	0	Stopped
Metrics Sync Job	Halted. (Quoru...	File Synchronization	DGAge...	0	1	0	0	0	▼ 0	0	Stopped
Locking Job	Stopped	File Locking		0	0	0	0	0	▼ 0	0	Stopped
Sales	Stopped	File Collaboration			0	0	0	0	▼ 0	0	Stopped
Engineering	Stopped	File Collaboration			0	0	0	0	▼ 0	0	Stopped
FR Job	Running	File Replication		0	0	0	0	0	▼ 0	0	Compl...
FC-EconA	Running	File Collaboration		9	0	0	0	0	▼ 0	0	Pendin...
FC-01-09	Running	File Collaboration		0	0	1	0	0	▼ 0	0	Compl...
FS-1	Stopped	File Synchronization		0	0	0	0	0	▼ 0	0	Stopped
EE-345	Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stopped
Media	Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stopped
District	Halted. (Quoru...	File Collaboration	DGPM...	0	1	0	0	0	▼ 0	0	Stopped
Products	Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stopped

Active Jobs -> Failed Participants: 3 of 3 | Bytes Pending: 0 bytes | Bytes Transferred: 1.87 KB | Opens: 0 | Initial Scans Completed: 2 of 6 | Total Size: 1.87 KB | Total Files: 4 | Total Directories: 6

56 errors, 0 warnings, 52 others | Filter by : Host: Name: Severity: Type: Clear Alerts

Received Date	Severity	Type	Name		Message
02-07-2019 15:40:02	Error	Host Resource...	FC-EconA	...t2	NAS FATAL DISCONNECT (777) info=777
02-07-2019 15:39:30	Info	Auto Start Job	FC-EconA	DGAgent2	Auto starting job, DGAgent2 DGAgent2 --> Storage \\afs1\MattS1 host is now available.
02-07-2019 15:38:00	Error	Host Resource...	FC-EconA	DGAgent2	NAS FATAL DISCONNECT (777) info=777
02-07-2019 15:37:30	Info	Auto Start Job	FC-EconA	DGAgent2	Auto starting job, DGAgent2 DGAgent2 --> Storage \\afs1\MattS1 host is now available.
02-07-2019 15:35:59	Error	Host Resource...	FC-EconA	DGAgent2	NAS FATAL DISCONNECT (777) info=777

The primary views are are described in the following table.

Jobs	The Jobs view displays a list of all jobs, grouped by job type. The following buttons are available within this panel: • **Start** and **Stop** buttons allow you to start and stop any selected jobs. • View Runtime Summary button displays a table of summary information for all jobs of a selected job type.
Agent Summary	The Agent Summary 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.
Alerts	The Alerts view 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 will be reported here.
Job Alerts	The Jobs Alerts view displays a list of all job-specific alerts that have occurred (including those for file collaboration sessions with detailed information about each alert. Alerts relating to the automatic stopping and restarting of jobs will be reported here.

Runtime views (tabbed view in center of screen)	Displays runtime and configuration information for all open jobs.
Dashboard	Shows the **Dashboard Summary View** panel, which displays metrics and key performance indicators from all running File Collaboration jobs, File Synchronization jobs, and Agents.
Peer Agent Detail Summary	The Peer Agent Detail Summary is a panel that 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.

Jobs View

The **Jobs** view is located in the top left quadrant of the Peer Management Center interface and lists all the jobs, grouped by type.

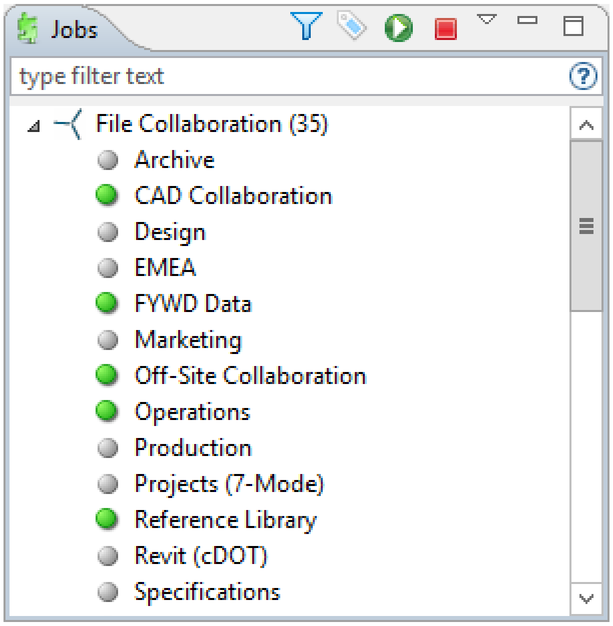

File Collaboration (35)
- ○ Archive
- ● CAD Collaboration
- ○ Design
- ○ EMEA
- ● FYWD Data
- ○ Marketing
- ● Off-Site Collaboration
- ● Operations
- ○ Production
- ○ Projects (7-Mode)
- ● Reference Library
- ○ Revit (cDOT)
- ○ Specifications

Double-clicking any job will open the selected job and display a view in the Summaries quadrant. For example, double-clicking a File Collaboration job will display the Collaboration Summary view.

Context Menu

Right-clicking any job will open a context pop-up menu with the following options:

Open	Open the selected job in an already open tab within the run-time view. Otherwise, a new tab will be opened for the selected job.
Open in New Tab	Open the selected job in a new tab within the File Collaboration Runtime View.
Start	Start the selected job if it is not already running.
Stop	Stop the selected job if it is already running.
Delete	Delete the selected job from the Peer Management Center and from disk.
Edit Jobs(s)	Edit the configuration of the selected job.
Copy	Copy the selected job while assigning it a unique name.
Rename	Rename the selected job.

Selecting multiple jobs and right-clicking will show a subset of the above context pop-up menus. Doing so will allow you to open, start, stop, and edit multiple jobs at once. For more information, see Editing Multiple Jobs.

Toolbar

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

Manage, Save and Load Filters	Allows for the selection of built-in or user-defined filters and to save/manage filter expressions. Default job filters include Failed Jobs, Jobs with Backlog, and Running Scans. For example, filter: "running Scans."
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.
Start Job	Start one or more selected and currently stopped jobs.
Stop Job	Stop one or more selected and currently running jobs.

View Collaboration Summary	View a table of summary information for all jobs of a selected job type. The view is defined and opened by simply clicking on a selected job (such as **Document Collaboration** in the image above) or its parent job type (or **File Collaboration** in the image above), then pressing the **View Runtime Summary** button.

Filtering Using List Filters

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 resources, see List Filters.

Agent Summary View

The **Agent Summary** view is located below the Jobs view. This view contains a list of all known Peer Agents installed in your environment and displays the current connection status for each.

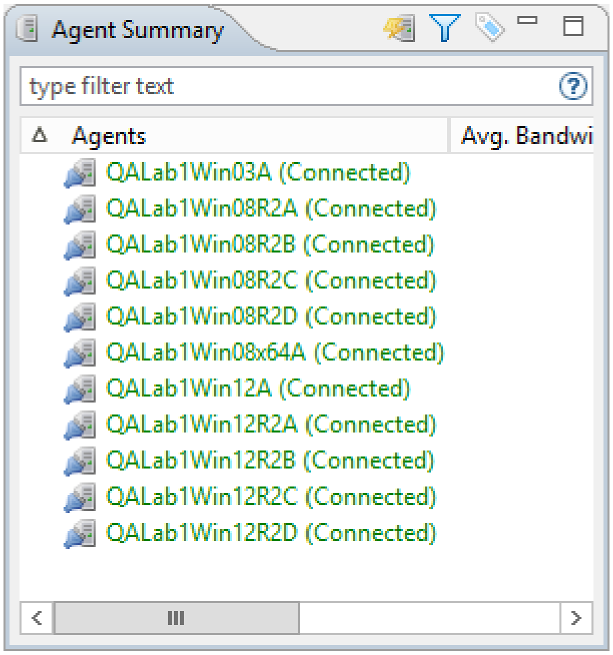

QALab1Win03A (Connected)
QALab1Win08R2A (Connected)
QALab1Win08R2B (Connected)
QALab1Win08R2C (Connected)
QALab1Win08R2D (Connected)
QALab1Win08x64A (Connected)
QALab1Win12A (Connected)
QALab1Win12R2A (Connected)
QALab1Win12R2B (Connected)
QALab1Win12R2C (Connected)
QALab1Win12R2D (Connected)

Toolbar

The following buttons are available on the toolbar within the Peer Agent Summary view:

View Agent	Opens the Agent Detail Summary View, which provides details for all known Agents and their status.

Detail Summary	
Manage, Save and Load Filters	Allows for the selection of built-in or user-defined filters and to save and manage list 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.

Filtering

To filter a large list of Agents, use the **Filter** field located below the toolbar buttons in the Agent Summary view panel. See Peer Agent Connection Statuses for information about Agent connection statuses.

For more details on how to filter resources, see List Filters.

Peer Agent Menu Options

Right-clicking one or more host names in the Peer Agent list will open a context pop-up menu with the following options:

Remove	This will remove the selected Peer Agent(s) from the 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	Displays properties for the selected Peer Agent, e.g., heartbeat information, host machine configuration, messaging statistics, performance statistics. See View Peer Agent Properties Dialog for more details.
Edit Configuration	Clicking this menu item will display a dialog where you can edit user configurable properties for the selected Peer Agent.
Restart Agent	If the selected Peer Agent is connected, this menu item will restart the Peer Agent Windows service running on the corresponding host. In the event

Service	that the Peer Agent is not connected to the Peer Management Center Broker, an attempt will be made to restart the Peer Agent Windows service using the Windows sc command. Please note that this will only work 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	Launch a Windows Remote Desktop connection to the selected Peer Agent.
Edit Agent Configuration	This action 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 Advanced Peer Agent Configuration.
Retrieve Log Files	This action retrieves log files for the selected Peer Agent containing information used by our technical support staff to assist in debugging issues. The log files are encrypted and will be located in the support folder of the Peer Management Center installation directory. They can optionally be uploaded to our technical support team.
Test Agent Bandwidth Speed	If the selected Peer Agent is connected, this menu item will start a bandwidth speed test to be performed in the background. You will be notified at completion with the results of the test.
Generate Thread Dump	This will generate a thread dump for the selected Peer Agent, which can be used by our technical support to debug certain issues. The debug file will be located in the Peer Agent installation directory.
Generate Memory Dump	This will generate a memory dump for the selected Peer Agent, which can be used by our technical support to debug certain issues. The debug file will be located in the Peer Agent installation directory.
Memory Garbage Collection	Force a garbage collection operation to attempt to reclaim memory that is no longer used within the Peer Agent's JVM.
Copy File	This action copies a specified file from the Peer Management Center to the designated target folder on each selected Peer Agent. The target folder is relative to the Peer Agent installation directory.
Transfer Rate Report	This action displays a time series performance chart of average transfer rate for the selected Peer Agent over the last 24 hours.

(not available on Web Client)	

Peer Agent Updates

Additionally, if the Peer Agent software running on a host is out of date, the host will be shown as having a pending update in the Peer Agent Summary view. When right-clicking the host, the option to automatically update the Peer Agent software will also be available. This process can be done right from the Peer Management Center and usually does not require any additional actions on the host server itself.

Agent Detail Summary View

The **Agent Detail 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.

To view the Agent Detail Summary view, use of the following methods:

- Select **View Agent Detail Summary** from the **Window** menu.
- Click the **View Agent Detail Summary** icon in the Peer Management Center or Peer Agent Summary View toolbars.

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

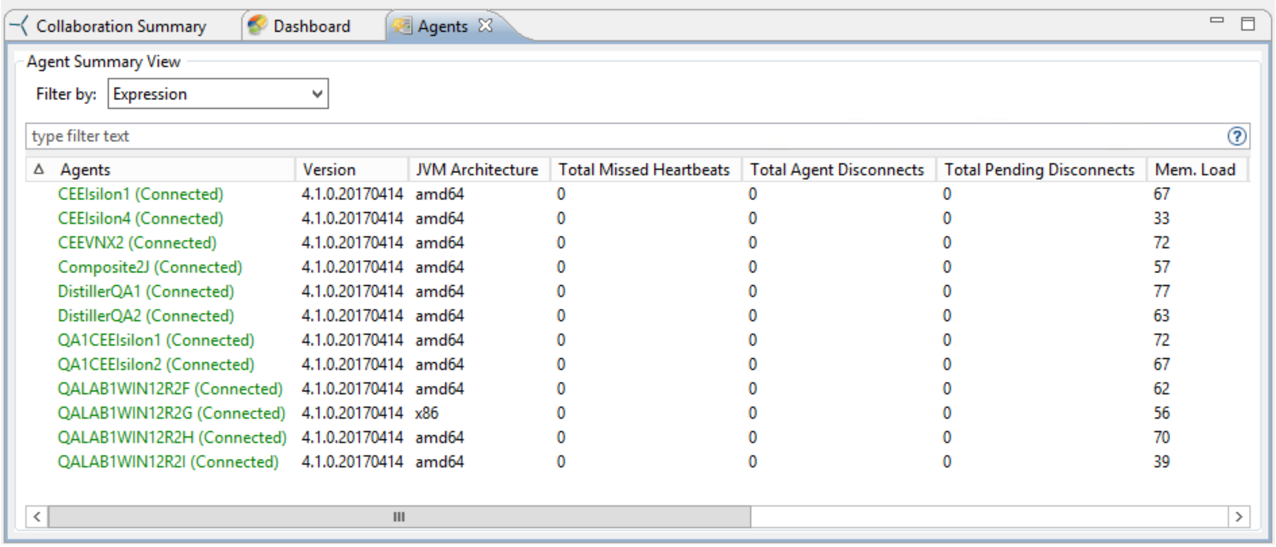

Agent Summary View

Filter by: Expression

type filter text

Agents	Version	JVM Architecture	Total Missed Heartbeats	Total Agent Disconnects	Total Pending Disconnects	Mem. Load
CEEIsilon1 (Connected)	4.1.0.20170414	amd64	0	0	0	67
CEEIsilon4 (Connected)	4.1.0.20170414	amd64	0	0	0	33
CEEVNX2 (Connected)	4.1.0.20170414	amd64	0	0	0	72
Composite2J (Connected)	4.1.0.20170414	amd64	0	0	0	57
DistillerQA1 (Connected)	4.1.0.20170414	amd64	0	0	0	77
DistillerQA2 (Connected)	4.1.0.20170414	amd64	0	0	0	63
QA1CEEIsilon1 (Connected)	4.1.0.20170414	amd64	0	0	0	72
QA1CEEIsilon2 (Connected)	4.1.0.20170414	amd64	0	0	0	67
QALAB1WIN12R2F (Connected)	4.1.0.20170414	amd64	0	0	0	62
QALAB1WIN12R2G (Connected)	4.1.0.20170414	x86	0	0	0	56
QALAB1WIN12R2H (Connected)	4.1.0.20170414	amd64	0	0	0	70
QALAB1WIN12R2I (Connected)	4.1.0.20170414	amd64	0	0	0	39

Alerts View

The **Alerts** view is automatically displayed when a critical system (Error or Fatal) alert is received. By default, the **Alerts** view is displayed under the runtime views.

Alert severity is broken down into four categories: Informational (containing Info, Debug, and Trace), Warning, Error and Fatal. An example of an Informational alert is when an Peer Agent connects to the Peer Management Broker. If an 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** sub directory 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.

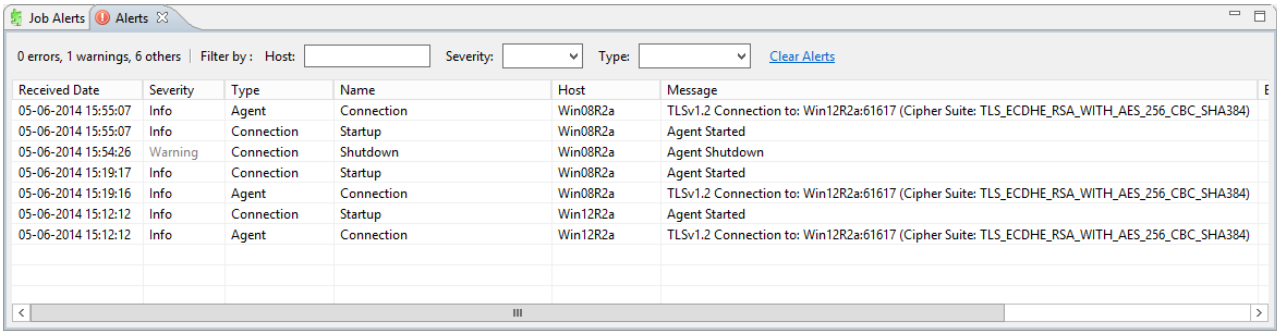

Received Date	Severity	Type	Name	Host	Message	E
05-06-2014 15:55:07	Info	Agent	Connection	Win08R2a	TLSv1.2 Connection to: Win12R2a:61617 (Cipher Suite: TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384)	
05-06-2014 15:55:07	Info	Connection	Startup	Win08R2a	Agent Started	
05-06-2014 15:54:26	Warning	Connection	Shutdown	Win08R2a	Agent Shutdown	
05-06-2014 15:19:17	Info	Connection	Startup	Win08R2a	Agent Started	
05-06-2014 15:19:16	Info	Agent	Connection	Win08R2a	TLSv1.2 Connection to: Win12R2a:61617 (Cipher Suite: TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384)	
05-06-2014 15:12:12	Info	Connection	Startup	Win12R2a	Agent Started	
05-06-2014 15:12:12	Info	Agent	Connection	Win12R2a	TLSv1.2 Connection to: Win12R2a:61617 (Cipher Suite: TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384)	

You can close the view at any time by clicking on the **X** (Close) button on the **Alerts** tab. 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 the **Window** menu, then the **Show View** submenu, followed by the **View Alerts** menu item.

You can also 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 on the Alerts tab again.

Job Alerts View

The **Job Alerts** view is automatically displayed when a critical selected job-related (Error or Fatal) alert is received. By default, the Job Alerts view is displayed under the runtime views, next to the Alerts view.

Job alert severity is broken down into four primary categories: Informational (containing Info, Debug, and Trace), Warning, Error, and 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.

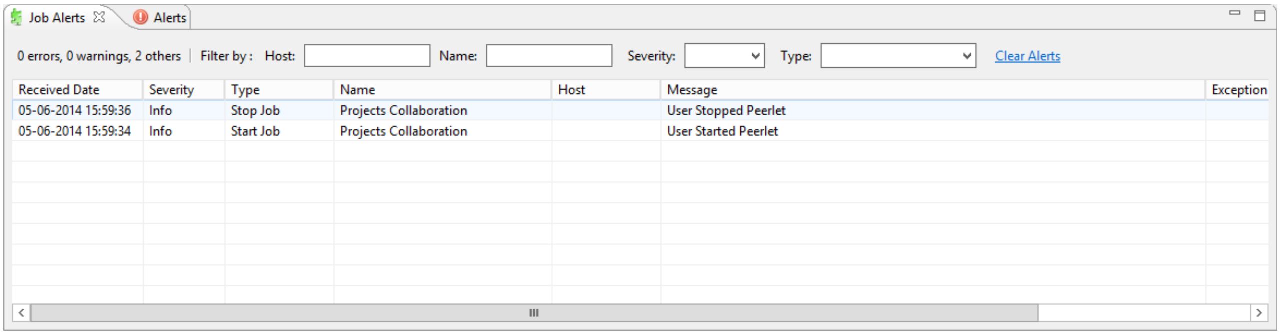

0 errors, 0 warnings, 2 others | Filter by :

Received Date	Severity	Type	Name	Host	Message	Exception
05-06-2014 15:59:36	Info	Stop Job	Projects Collaboration		User Stopped Peerlet	
05-06-2014 15:59:34	Info	Start Job	Projects Collaboration		User Started Peerlet	

You can close the view at anytime by clicking on the **X** (Close) button on the Job Alerts tab. 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 also 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 maximum the view. You can restore the view to its original, non-maximized size by double-clicking the **Job Alerts** tab again.

Basic Concepts

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

- Email Alerts
- File and Folder Filters
- List Filters
- Tags

Email Alerts

Overview

An email alert notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. An email is sent to all listed recipients whenever a selected event type is triggered by that job.

An email alert consists of a unique name, a selection of event types, and a list of email addresses. The event types depend on the job type. You create email alerts for a job type in Preferences. After you create an alert, you apply it to a job when creating or editing a job of the same type. For example, an email alert created in the Preferences for File Collaboration job can be applied only to File Collaboration jobs.

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

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

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.

Defining and Applying File 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 Filter Patterns. For more information on defining a filter pattern that can be used to filter folders, see Folder Filters.

You define a file filter in the Preferences for a job type; the filter can then be applied to individual jobs of the same type. For example, a file filter defined in File Collaboration Preferences can be applied to any File Collaboration 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.

Types of File Filters

There are three types of file filters:

- **Any** - 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.

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 on 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:

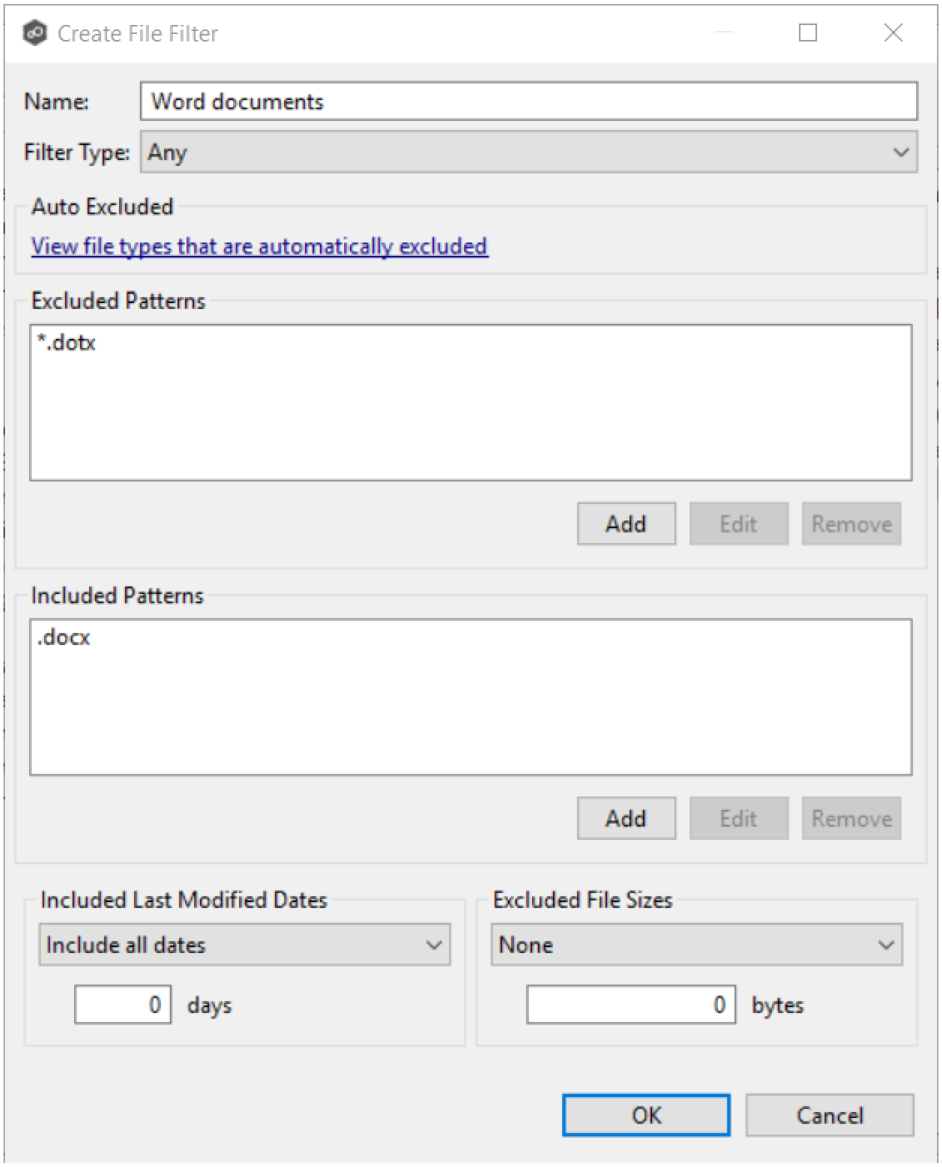

Create File Filter

Name: Word documents

Filter Type: Any

Auto Excluded

View file types that are automatically excluded

Excluded Patterns

*.dotx

Included Patterns

.docx

Included Last Modified Dates: Include all dates

0 days

Excluded File Sizes: None

0 bytes

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
```

A good reference on regular expressions can be found here: http://www.regular-expressions.info/reference.html

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**

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 files, you would add these extensions (with wildcards) to the **Excluded Patterns** field:

1. Click the **Add** button under the **Excluded Patterns** field.
2. Enter ***.AC$**, and then click **OK**.
3. Repeat Step 1 to add ***.SV$**, ***.DWL*** and ***.BAK**.

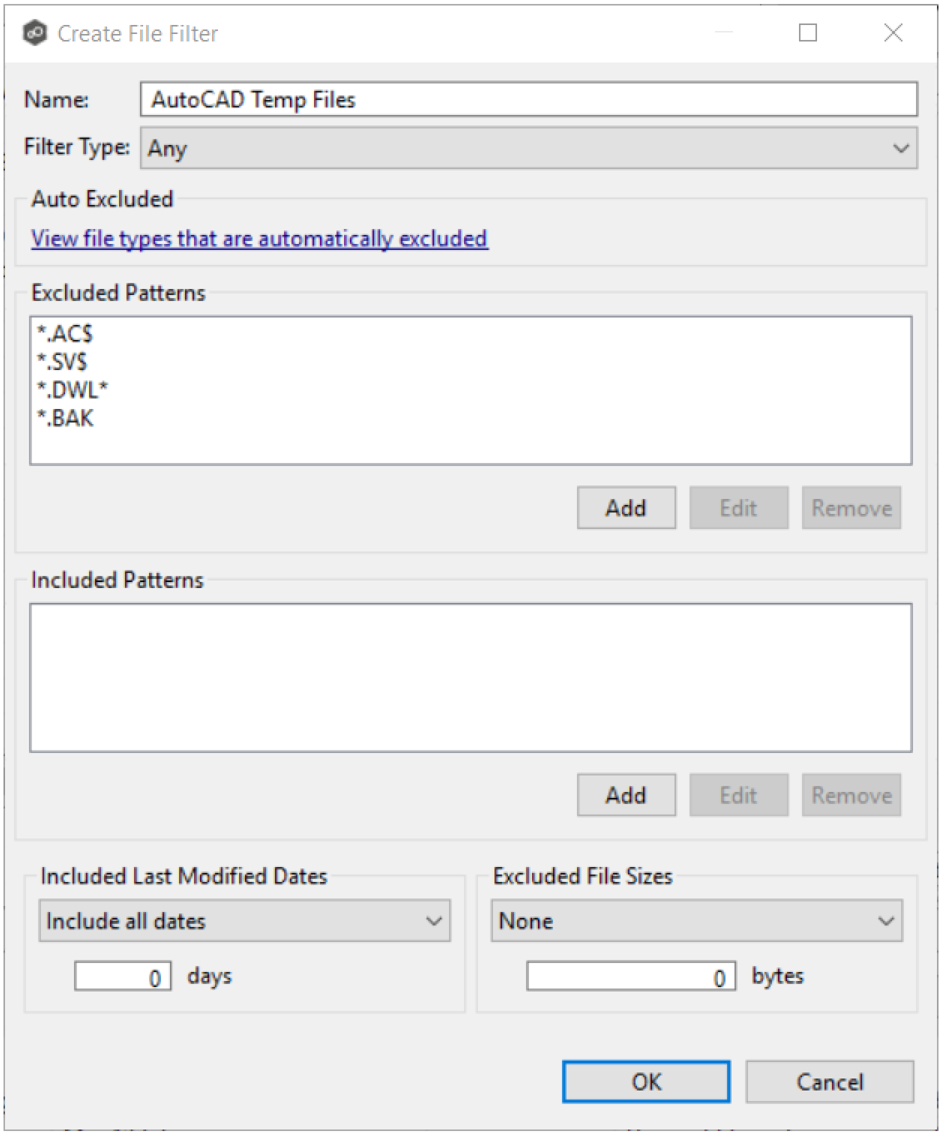

Create File Filter

Name: AutoCAD Temp Files

Filter Type: Any

Auto Excluded

View file types that are automatically excluded

Excluded Patterns

*.AC$
*.SV$
.DWL
*.BAK

Included Patterns

Included Last Modified Dates: Include all dates

0 days

Excluded File Sizes: None

0 bytes

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

Filtering on Last Modified Date

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 only supports filtering on a file's last modified date and does not support filtering on a folders last modified date. In addition, 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.

Note that a file filter cannot combine filtering on last modifed 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 then 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 then the specified number days.

Filtering on File Size

You can also filter based on an individual file's size, excluding files that are greater or less than a specified size. Peer Management Center does not support filtering on a folder's total size. In addition, 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.

However, you cannot define a file filter that combines filtering on file size with inclusion or exclusion patterns or last modifed 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 1GB (1073741824 bytes).

Exclude files less than	Select this option to exclude files whose size is less than the specified number of bytes.

Predefined File Filters

In addition to defining your own file filters, there are six predefined file filters that can be applied to jobs.

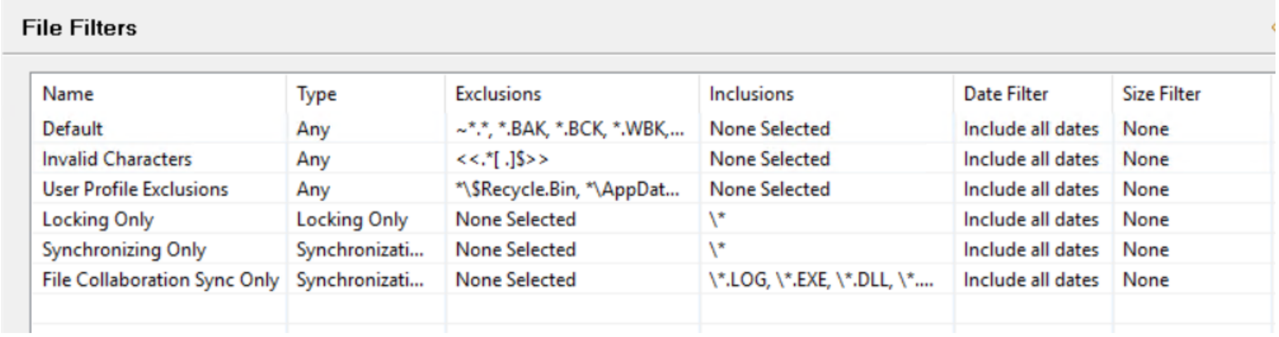

File Filters

Name	Type	Exclusions	Inclusions	Date Filter	Size Filter
Default	Any	~*.*, *.BAK, *.BCK, *.WBK,...	None Selected	Include all dates	None
Invalid Characters	Any	<<.*[.]$>>	None Selected	Include all dates	None
User Profile Exclusions	Any	*\$Recycle.Bin, *\AppDat...	None Selected	Include all dates	None
Locking Only	Locking Only	None Selected	*	Include all dates	None
Synchronizing Only	Synchronizati...	None Selected	*	Include all dates	None
File Collaboration Sync Only	Synchronizati...	None Selected	*.LOG, *.EXE, *.DLL, *....	Include all dates	None

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.

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

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

Filtering Folders

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	Job 1	D:\General		D:\General	
Jobs	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	Job 1	D:\	D:\	\General*	All other files
Jobs				\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

The 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) (i.e 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.

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.

Removing List Filters

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

Tags

Tags can be used to categorize resources and customize a user's workspace or perspective. Examples of resources include Agents, jobs, and web roles. Tagging helps when managing large number of resources. Tagging resources helps when managing large number of resources.

You can assign tags to:

- Jobs
- Resources
- User roles

- Agents

You can also assign resources to tags.

Creating Tags and Categories

Tags and categories are created globally in the Tags Configuration dialog. 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 user roles
- Assigning resources to one or more tags

Assigning Tags to Jobs

- During job creation - You can assign tags during the creation of a brand new 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 Configuration(s)**, and navigating to the **Tags** page of the job edit wizard.

Assigning Tags to Resources

To assign tags to one or more resources:

1. Click the **Assign Tags** button from a Summary 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 to the right.

5. From the **Unassigned Resources** table on the left side, select the resources that need to be assigned the selected tag and 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 resources to select more than one).

6. Click the **Save** button to commit your changes and close the dialog.

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

Assigning Tags to User Roles

User roles can be assigned tags that customize a user's Jobs perspective when they log in via the Web Interface. 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 achieve this:

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 dialog.

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 r and 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 close the **Preferences** dialog.

   The user will only see the jobs which 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 toolbars.

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 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 close the dialog.
8. Repeat the steps above 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.

Advanced Topics

This section discusses the following topics:

- DFS Namespace Failover and Failback
- File Conflict Resolution
- File Metadata Synchronization
- Managing Peer Agents
- Smart Data Seeding
- TLS Certificates

DFS Namespace Failover and Failback

You can manage DFS namespaces through a dedicated job type in the Peer Management Center, the DFS-N Management job. The PMC controls failover and failback by automatically disabling and enabling DFS namespace folder targets.

Failover

The 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 the PMC detects a failure, it 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. If configured to do so, DFS Namespaces will automatically re-direct clients to another available folder target.

Failback

While DFS Namespaces itself can automate the disabling of folder targets, it does not automate the re-enabling of disabled targets. When configured to talk to DFS namespaces, the Peer Management Center can automate this process. When the PMC 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 re-scan 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.

File Conflict Resolution

File conflict resolution allows you to specify the type of file conflict resolution to use during a scan when a file conflict exists for a file between two or more hosts. Configuration is available when editing a File Collaboration job.

Overview

Conflict resolution is a key feature of file collaboration that is in effect at the start of a job. When a File Collaboration job begins, the host participants' configured watch sets are synchronized by a scan and merge phase, during which conflicts can be detected.

This topic describes:

- How file conflicts are defined.

- The scheme, and the

- Options for resolving conflicts.

Defining a Conflict

When a job begins, the participants' folders are first scanned and then merged to form a collective view of the content of all participants. All files found under the designated folders are subject to collaboration or synchronization, except for those excluded by file filters (see File Filters for more details).

A potential conflict occurs when a file path is found to exist on more than one host in a File Collaboration job. For example, the following files are found to be in conflict:

\\Host-A\FC-Session-UserGuide\release-1.0\readme.txt

\\Host-B\FileCollab-UG\release-1.0\readme.txt

\\Host-C\FCS-UserGuide\release-1.0\readme.txt

In this example, the file **\release-1.0\readme.txt** is found to be in conflict across three hosts. Note that each host can designate varying root folders. Conflicts may occur across a partial or total set of participants.

A file conflict 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.

Resolving a File Conflict

The goal of file conflict resolution is to designate one instance of a conflicted file as the "winning" copy or the one designated as the source for synchronization. The criteria for resolving conflicts are based on the file's metadata, such as size, modification time or host name.

It is important to note that conflict resolution must select a single instance of a file, although it is quite possible that several copies of a file are potential candidates. Drawing from the examples listed in the previous section, if our session was configured to resolve conflicts based on a files last modified time and all instances of **\release-1.0\readme.txt** had the same size and last modified time, then all three would be resolution candidates. In this case, the winner would be arbitrarily selected from the candidate set. This concept applies to all resolution types that are prone to multiple candidate selection.

Once the merge and conflict resolution phases have completed for the session, synchronization transfers begin to distribute the source content. This includes all source copies of conflict winners as well as files that are missing from participants.

See File Conflicts View for a more detailed explanation on how the file conflict process works and how to remove file conflicts and quarantines.

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:

- **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.

- **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.

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.

- In order 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 the Peer Management Center.

Right-clicking one or more host names in the Peer Agent list will open a context pop-up menu with the following options:

Remove	Removes the selected Peer Agent(s) from the 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	Displays properties for the selected Peer Agent, for example, heartbeat information, host machine configuration, messaging statistics, performance statistics. See View Peer Agent Properties for more details.
Edit Configuration	Clicking this menu item will display a dialog where you can edit user configurable properties for the selected Peer Agent.
Restart Agent Service	If the selected Peer Agent is connected, restarts the Peer Agent Windows service running on the corresponding host. 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 option 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	Launch 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 Advanced Peer Agent Configuration.
Retrieve Log Files	Retrieves log files for the selected Peer Agent containing information used by our technical support staff to assist in debugging issues. The log files are encrypted and will be located in the support folder of the Peer Management Center installation

	directory. Log files can optionally be uploaded to our technical support team.
Test Agent Bandwidth Speed	If the selected Peer Agent is connected, starts a bandwidth speed test to be performed in the background. You will be notified at completion with the results of the test.
Generate Thread Dump	Generates a thread dump for the selected Peer Agent that can be used by our technical support to debug certain issues. The debug file will be located in the Peer Agent installation directory.
Generate Memory Dump	Generates a memory dump for the selected Peer Agent that can be used by our technical support to debug certain issues. The debug file will be located in the Peer Agent installation directory.
Memory Garbage Collection	Forces a garbage collection operation to attempt to reclaim memory that is no longer used within the Peer Agent's JVM.
Copy File	Copies a specified file from the Peer Management Center to the designated target folder on each selected Peer Agent. The target folder is relative to the Peer Agent installation directory.
Transfer Rate Report	Displays a time series performance chart of average transfer rate for the selected Peer Agent over the last 24 hours. Note: Not available in web client.

Peer Agent Connection Statuses

Valid connection statuses are:

Connected	Indicates Peer Agent is currently connected to the Peer Management 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 Disconnect	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 the Peer Management Center was started, or the first heartbeat message has not been received from that host.

Editing an Agent Configuration

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

To access, right-click any connected Peer Agent, and then select **Edit Agent Configuration**. The **Peer Agent Configuration** dialog will be displayed, it contains five tabs of configuration items. In order 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 configuration dialog.

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

For more information, see:

- Broker Configuration
- General
- Logging
- VM Options

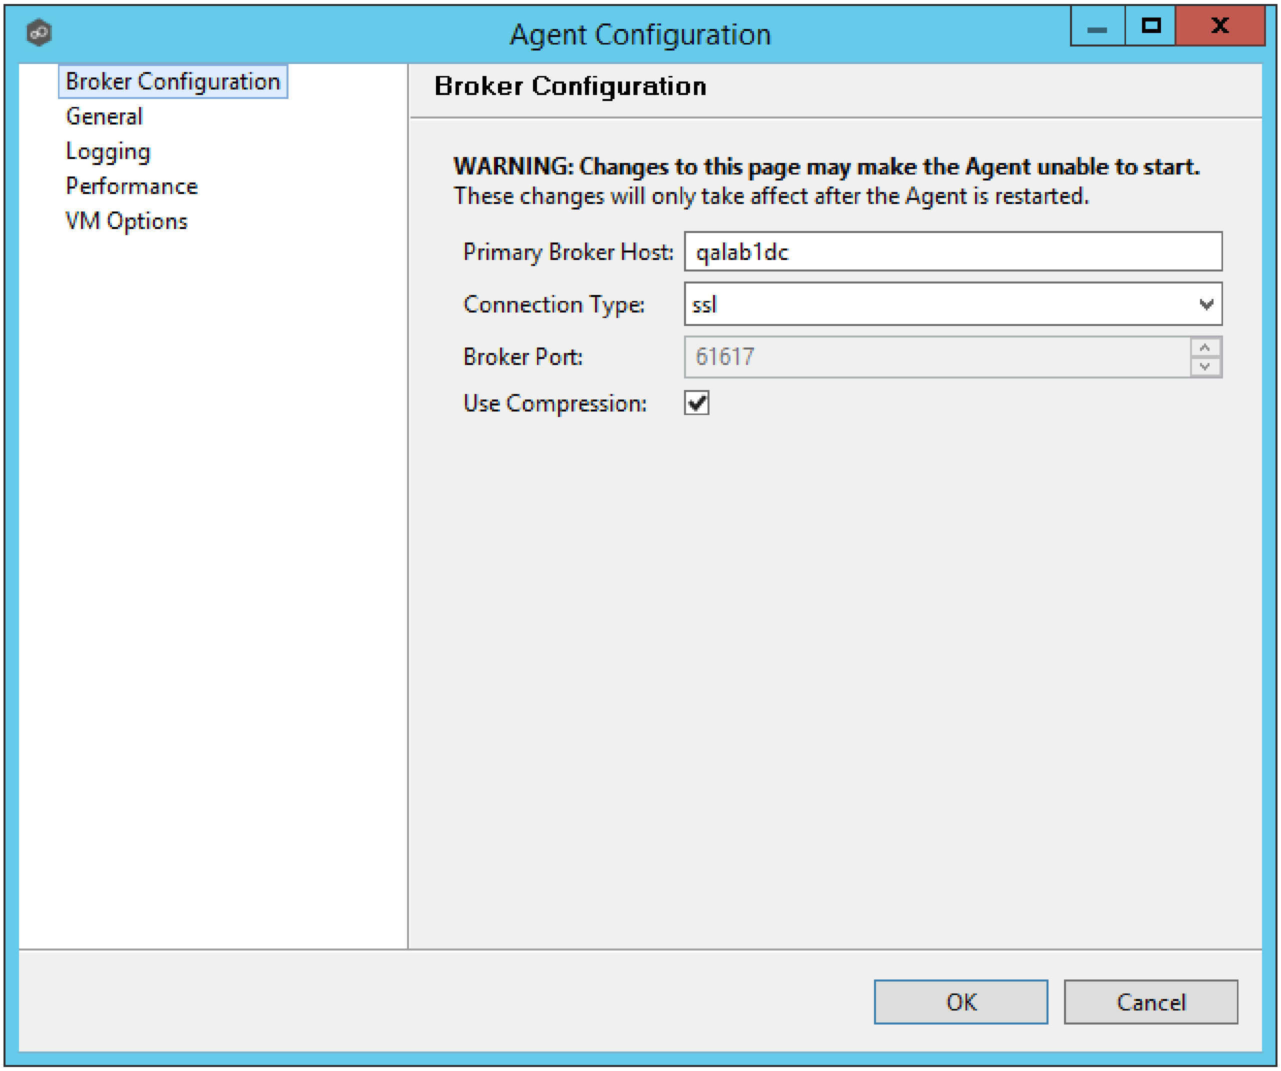

Broker Configuration

WARNING: Changes to this page may make the Agent unable to start.
These changes will only take affect after the Agent is restarted.

Primary Broker Host: qalab1dc

Connection Type: ssl

Broker Port: 61617

Use Compression: ☑

Please note that these settings only apply to communication between the selected Peer Agent and Peer Management Broker and not to communication between the Peer Management Center and Peer Management Broker.

Primary Broker Host	The IP address or fully qualified host name of the server running the Peer Management Center Broker.
Connection Type	The type of connection to use when communicating with the Peer Management Center Broker. Types include ssl (encrypted) and tcp (not encrypted).
Broker Port	The port on which to communicate with the Peer Management Center Broker.

Use Compression	When enabled, all communication between the selected Peer Agent and the Peer Management Center Broker will be compressed.

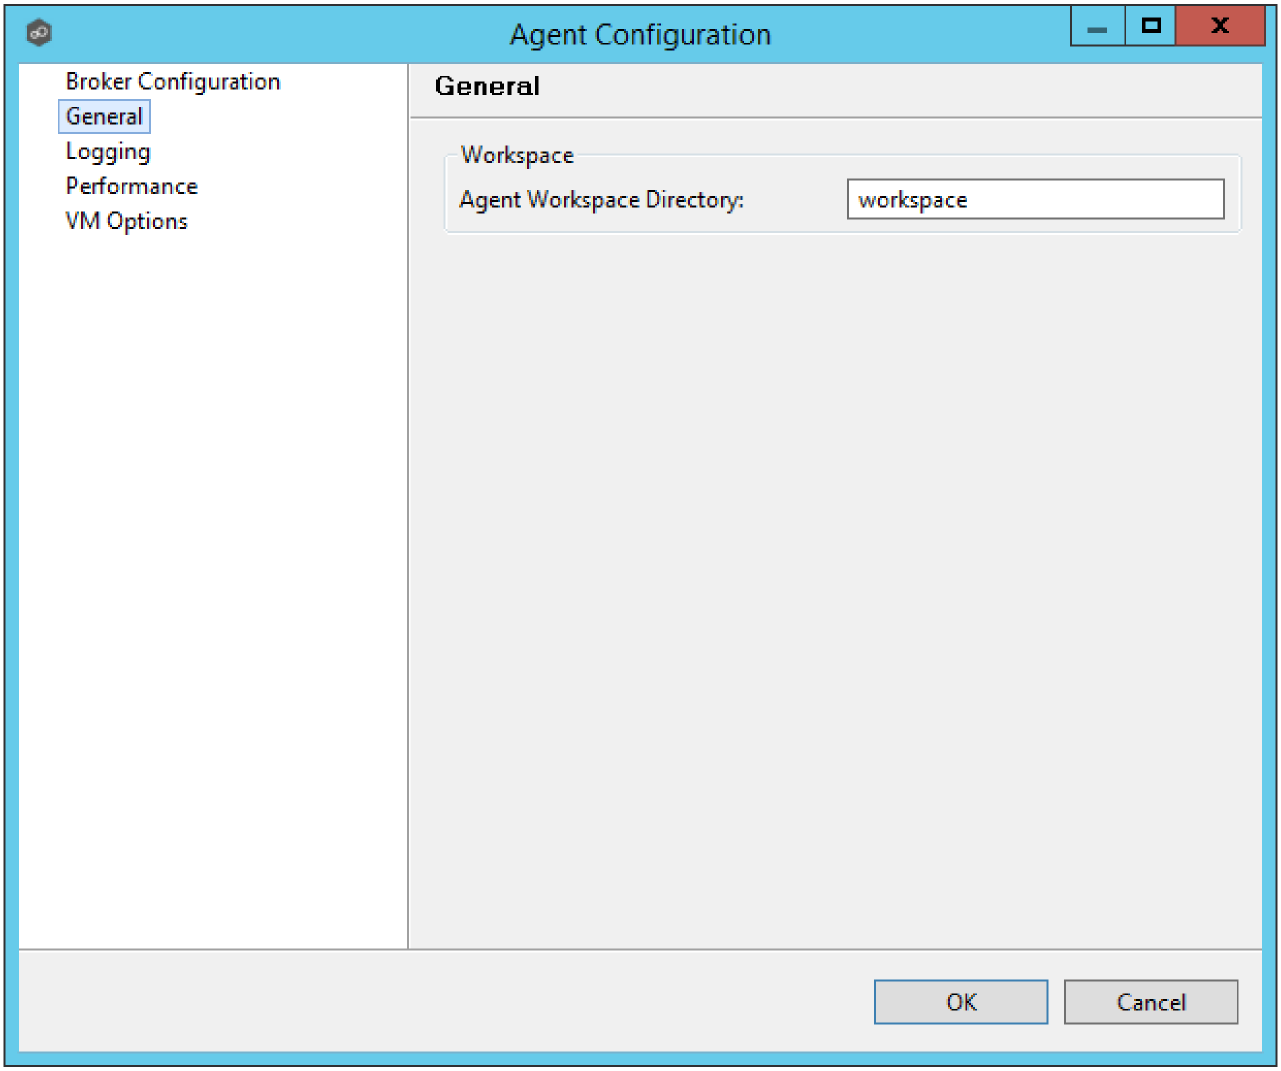

Agent Workspace Directory: workspace

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

Alerts

Notification and response settings for when the selected Peer Agent runs low on memory.

Low Memory Alert Percentages	Memory percentages at which the Peer Agent with post an alert to the Peer Management Center's Alert list. Multiple percentages can be set, separated by commas. For example: .85,.90,.99
Enable Low Memory Auto-Restart	When enabled, the Peer Agent will attempt to restart itself when its memory usage hits a certain threshold.
Restart Memory Percentage	If **Enable Low Memory Auto-Restart** is enabled, the Peer Agent will attempt to restart itself at this memory threshold. For example: .98

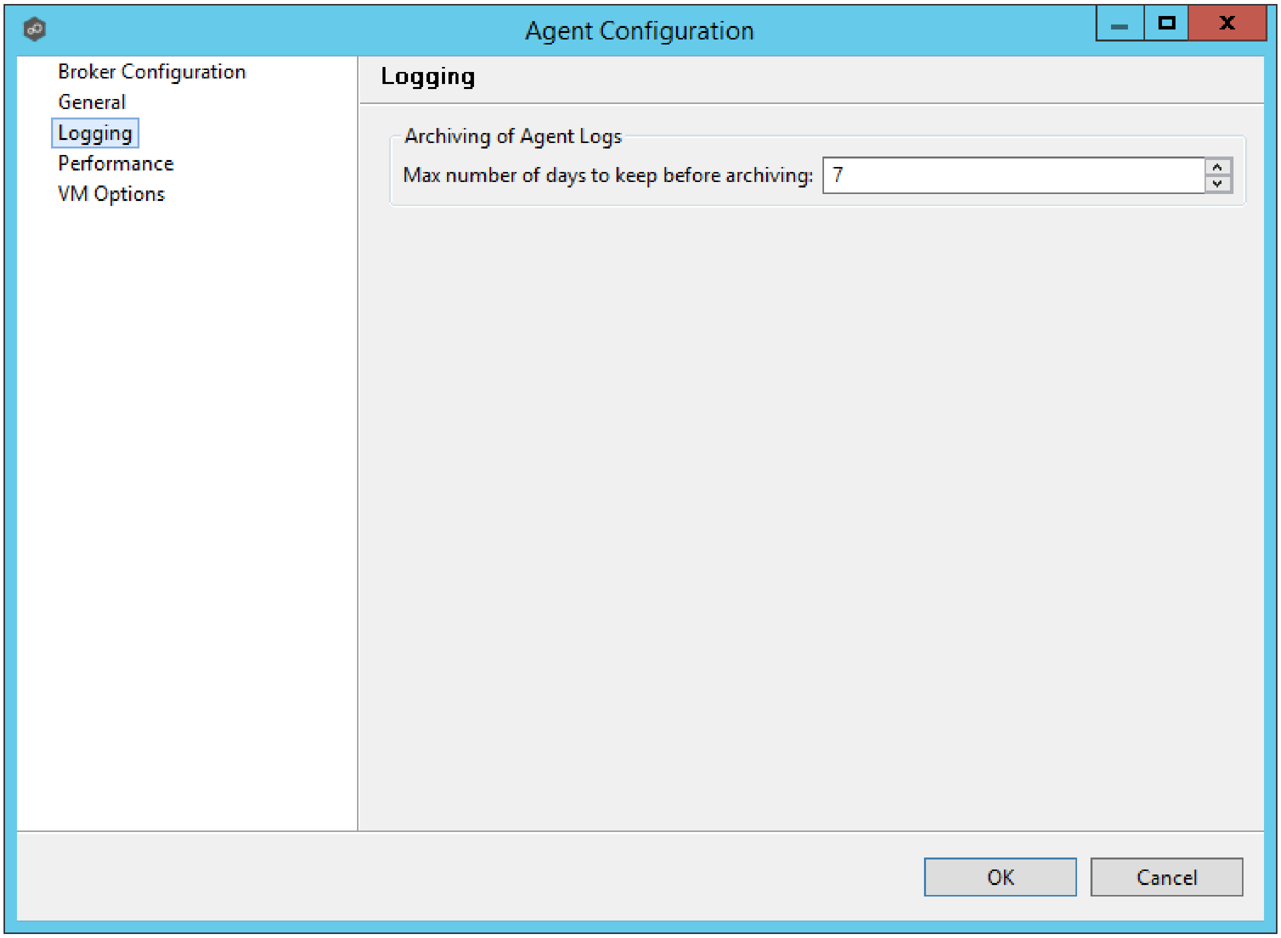

Logging

Archiving of Agent Logs

Max number of days to keep before archiving: 7

Max number of days to keep before archiving	Log files that are older than this date will automatically be zipped up and archived to reduce required space on disk.

Agent Logging

Settings for tuning Peer Agent logging. Depending on these settings, large log files may be produced.

Agent Logging	Peer Agent logging directory relative to the Peer Agent's installation directory. This can also be set to an explicit full

Directory	path. Selected folder must already exist before the Peer Agent is restarted.
Agent Log File Size (in MB)	The maximum size to which each Peer Agent .log file will grow before rolling over to a new file.
Max number of Agent log files	The maximum number of rolling Peer Agent .log files to keep.
STDOUT Log File Size (in MB)	The maximum size to which each output log file will grow before rolling over to a new file.
Max number of STDOUT log files	The maximum number of rolling output log files to keep.
STDERR Log File Size (in MB)	The maximum size to which each error log file will grow before rolling over to a new file.
Max number of STDERR log files	The maximum number of rolling error log files to keep.
JMS Messages Log File Size (in MB)	The maximum size to which each JMS message log file will grow before rolling over to a new file.
Max number of JMS Message log files	The maximum number of rolling JMS message log files to keep.
Profiler Log File Size (in MB)	The maximum size to which each profiler log file will grow before rolling over to a new file.
Max number of Profiler log files	The maximum number of rolling profiler log files to keep.

JMS Message Logging

Settings for enabling and tuning JMS Message logging. These settings are useful for debugging purposes but will affect performance and produce large log files. Changes to these settings should only be made at the request of the Peer Support Team.

The first option on the page allows for the ability to tune the maximum amount of system memory that the Peer Agent service will use on the server where it is installed. The maximum amount is 1.5GB. We strongly recommend that this value not be set below 512MB.

The text field below this option should only be used under the direction of the Peer Support Team.

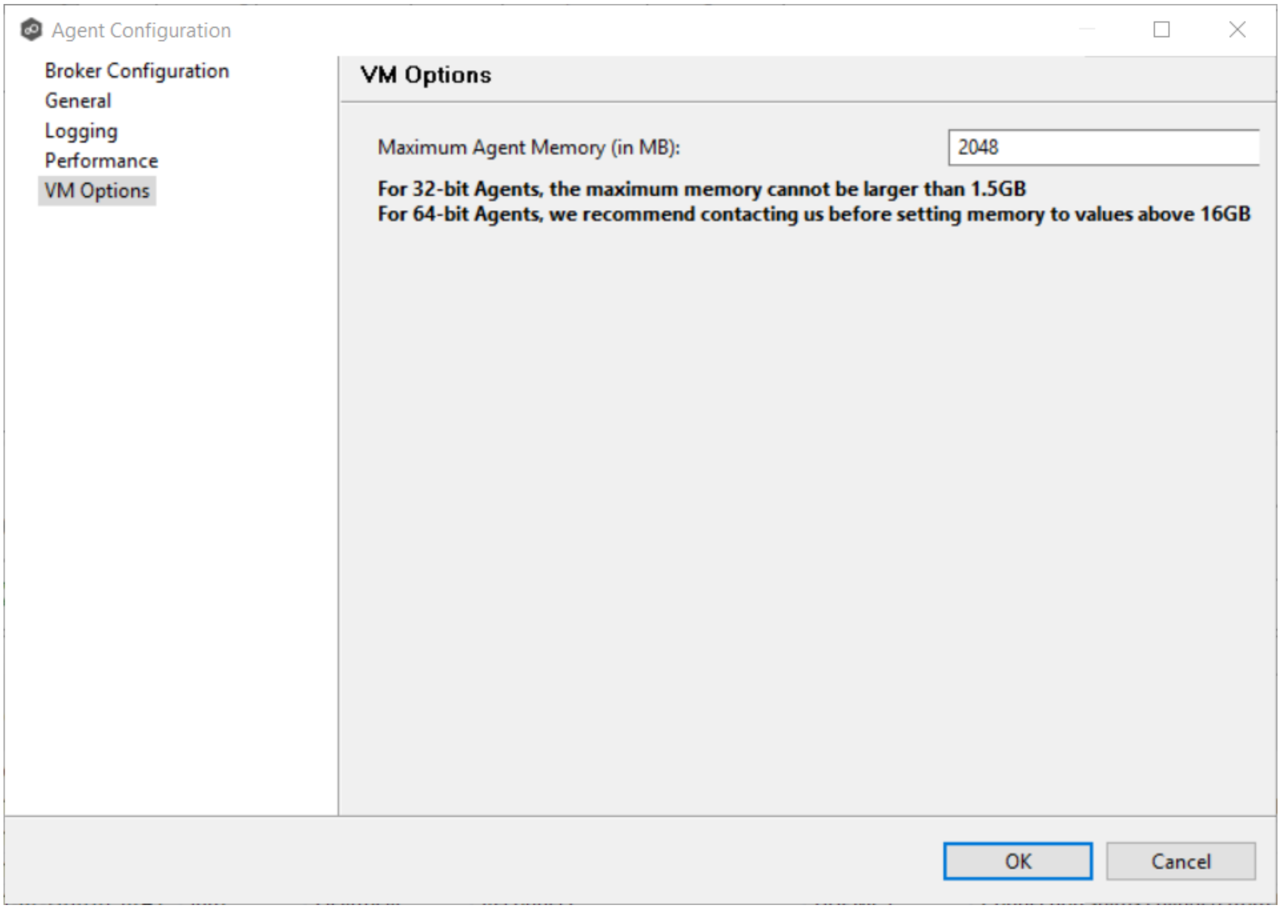

Maximum Agent Memory (in MB): 2048

For 32-bit Agents, the maximum memory cannot be larger than 1.5GB
For 64-bit Agents, we recommend contacting us before setting memory to values above 16GB

Viewing Agent Properties

To view properties of an Agent:

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

   The **View Agent Properties** dialog opens.

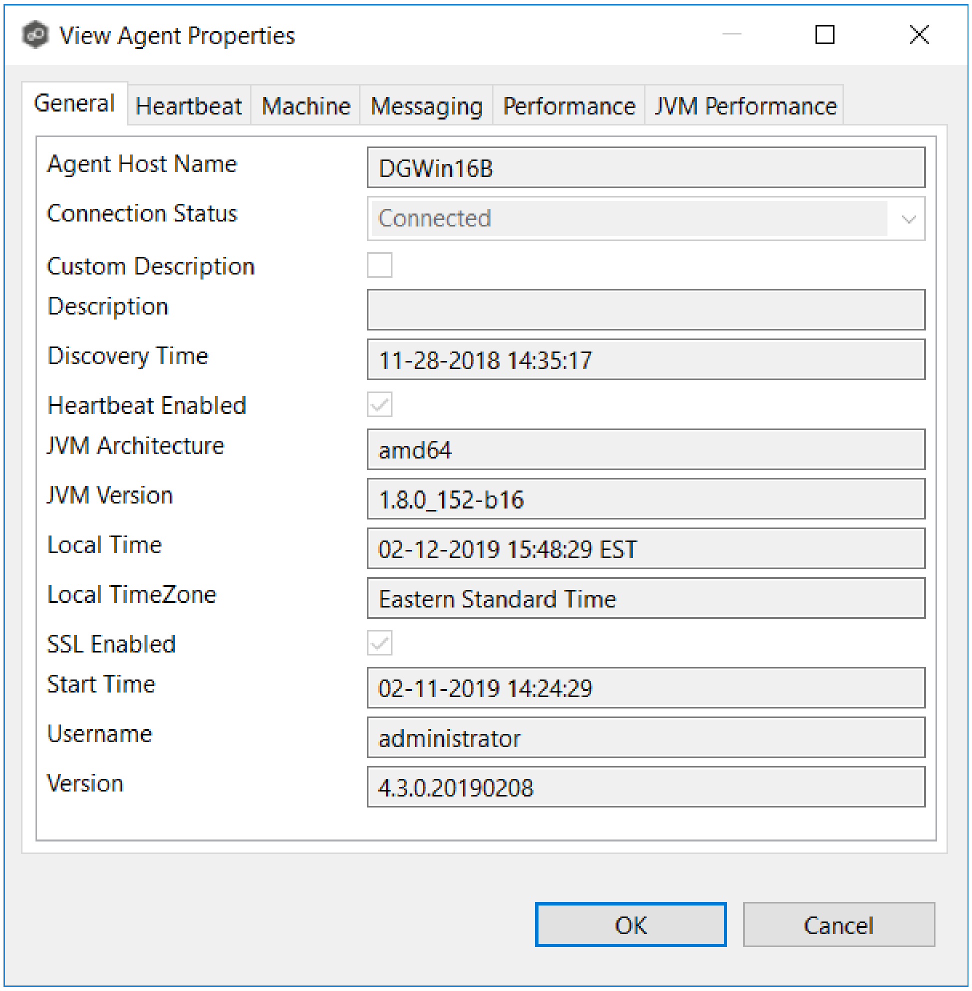

Agent Host Name	DGWin16B
Connection Status	Connected
Custom Description	□
Description	
Discovery Time	11-28-2018 14:35:17
Heartbeat Enabled	☑
JVM Architecture	amd64
JVM Version	1.8.0_152-b16
Local Time	02-12-2019 15:48:29 EST
Local TimeZone	Eastern Standard Time
SSL Enabled	☑
Start Time	02-11-2019 14:24:29
Username	administrator
Version	4.3.0.20190208

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

General	Displays general Peer Agent run-time information such as discovery time, local time, TLS use, Peer Agent start up time, Peer Agent version, 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.

General	Displays general Peer Agent run-time information such as discovery time, local time, TLS use, Peer Agent start up time, Peer Agent version, user name Peer Agent service is running as.
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 Agent Properties** menu item for a selected host will result in the opening of the following Peer Agent **Properties** dialog:

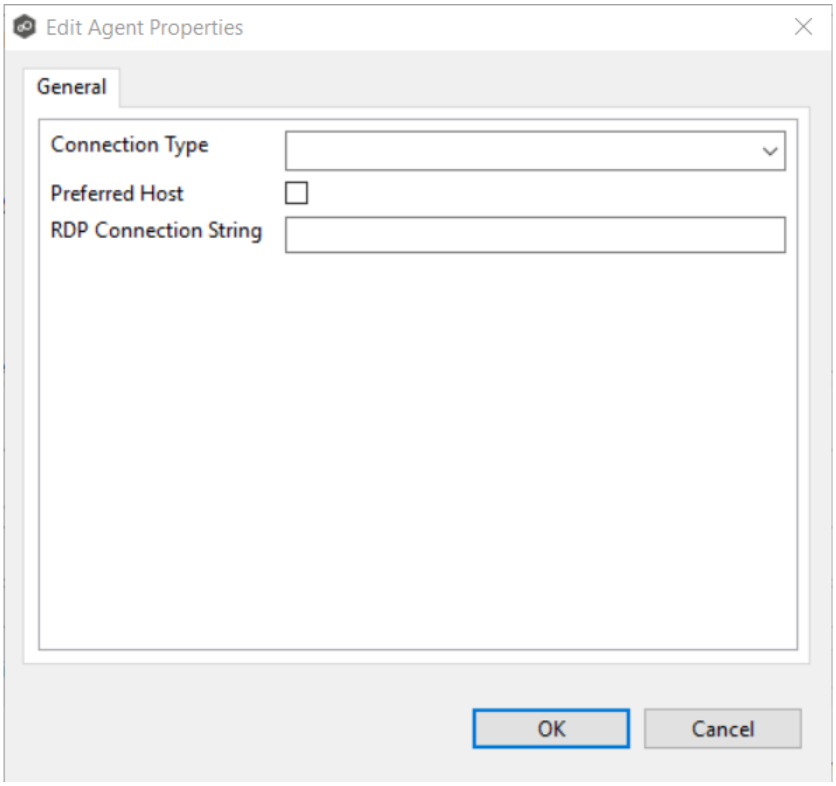

Edit Agent Properties

General

Connection Type

Preferred Host □

RDP Connection String

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

Connection Type	Allows for the selection of a connection type between selected Peer Agent and its associated Peer Management Center 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 Center Broker (or in appropriate cases, for selecting which Peer Agent are on the same subnet as the Peer Management Center Broker).
RDP Connection String	The connection string to use when activating an RDP session to this Peer Agent.

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 view of the job.

If an error is disabled because of a MED action, it will look like the following:

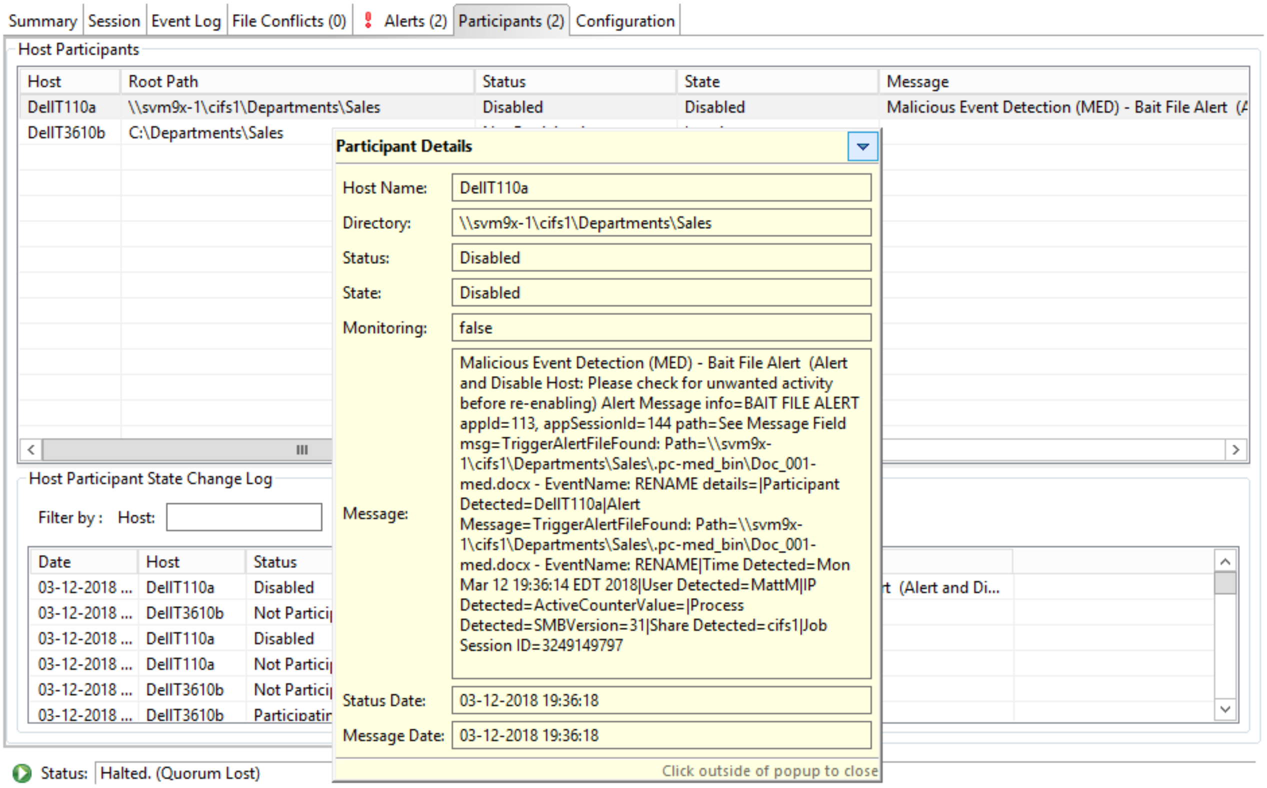

Host	Root Path	Status	State	Message
DellT110a	\\svm9x-1\cifs1\Departments\Sales	Disabled	Disabled	Malicious Event Detection (MED) - Bait File Alert (A
DellT3610b	C:\Departments\Sales			

Participant Details

Host Name: DellT110a

Directory: \\svm9x-1\cifs1\Departments\Sales

Status: Disabled

State: Disabled

Monitoring: false

Message: Malicious Event Detection (MED) - Bait File Alert (Alert and Disable Host: Please check for unwanted activity before re-enabling) Alert Message info=BAIT FILE ALERT appId=113, appSessionId=144 path=See Message Field msg=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\.pc-med_bin\Doc_001-med.docx - EventName: RENAME details=|Participant Detected=DellT110a|Alert Message=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\.pc-med_bin\Doc_001-med.docx - EventName: RENAME|Time Detected=Mon Mar 12 19:36:14 EDT 2018|User Detected=MattM|IP Detected=ActiveCounterValue=|Process Detected=SMBVersion=31|Share Detected=cifs1|Job Session ID=3249149797

Status Date: 03-12-2018 19:36:18

Message Date: 03-12-2018 19:36:18

Date	Host	Status	
03-12-2018 ...	DellT110a	Disabled	t (Alert and Di...
03-12-2018 ...	DellT3610b	Not Partici	
03-12-2018 ...	DellT110a	Disabled	
03-12-2018 ...	DellT110a	Not Partici	
03-12-2018 ...	DellT3610b	Not Partici	
03-12-2018 ...	DellT3610b	Participatir	

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

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 Peer Agent Summary view.

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

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 in-sync. 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 users on the target hosts (Active-Active). 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 the Peer Management Center Client:

```
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**:

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
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 if 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.

TLS Certificates

You can use custom or private TLS certificates to connect Peer Agent to the Peer Management Center 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 the Peer Management Center and Peer Agent.

For more information, see:

- Creating New Certificates
- Using Existing Certificates

Creating New Certificates

Perform the necessary commands using the keytool application bundled with your Peer Management Center or Peer Agent installation.

Keytool location on Peer Management Center system:	PMC_INSTALLATION_FOLDER\jre\bin
Keytool location on Peer Agent system:	PEER_AGENT_INSTALLATION_FOLDER\jre\bin

Broker Keystore Generation

Step 1. Using keytool, create a certificate for the 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:

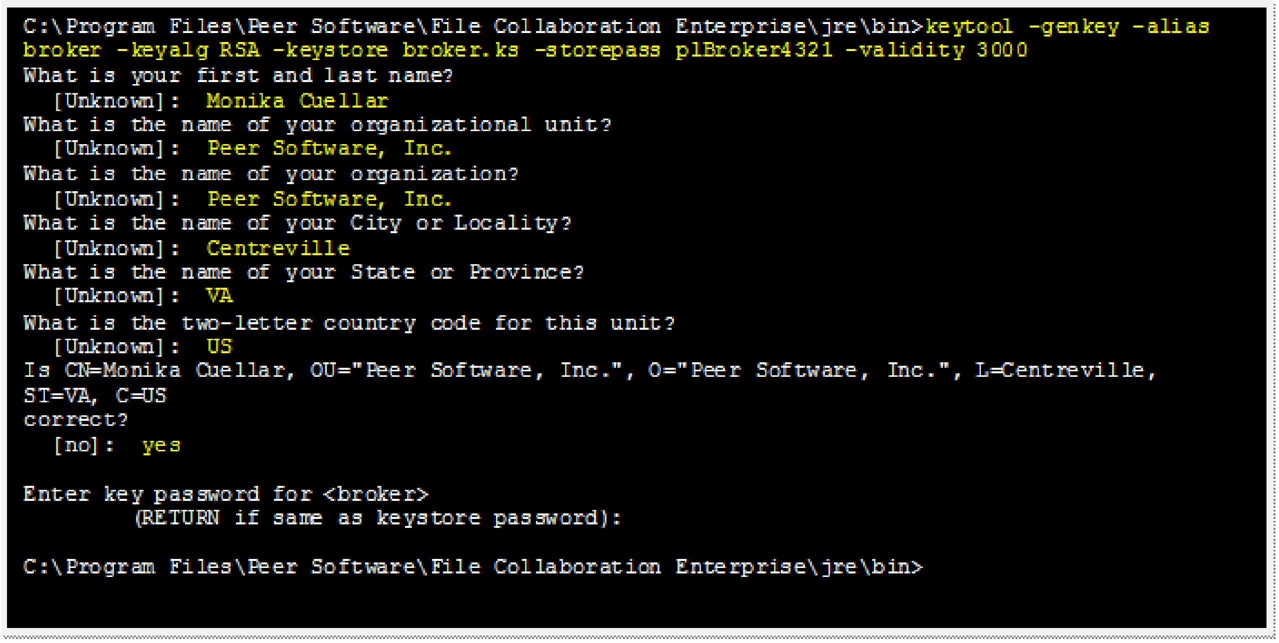

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -genkey -alias
broker -keyalg RSA -keystore broker.ks -storepass plBroker4321 -validity 3000
What is your first and last name?
  [Unknown]:  Monika Cuellar
What is the name of your organizational unit?
  [Unknown]:  Peer 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:

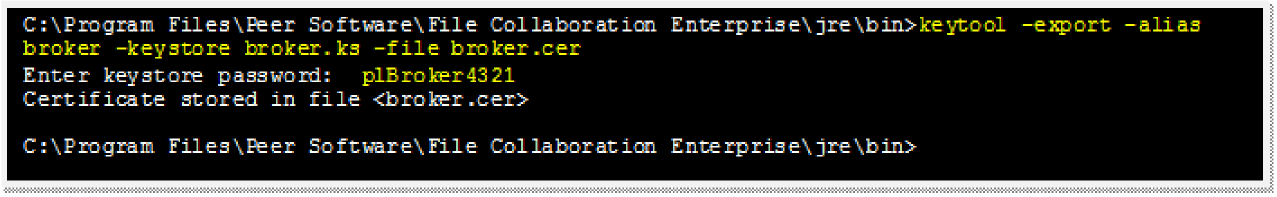

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -export -alias
broker -keystore broker.ks -file broker.cer
Enter keystore password:  plBroker4321
Certificate stored in file <broker.cer>

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:

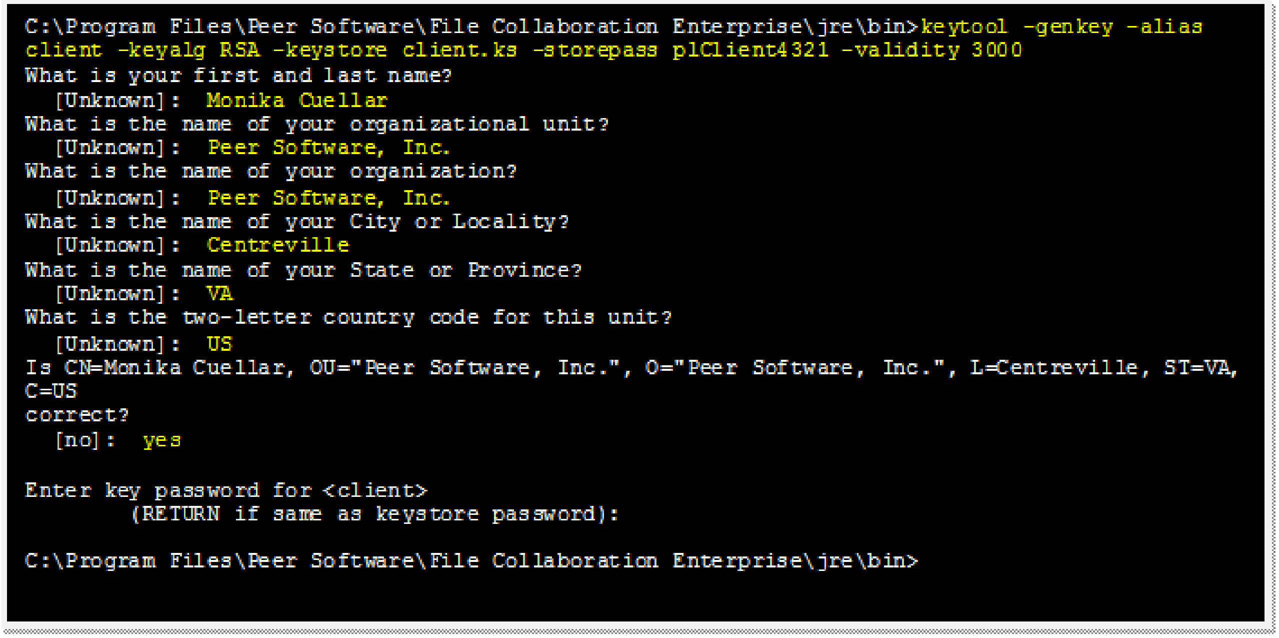

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -genkey -alias
client -keyalg RSA -keystore client.ks -storepass plClient4321 -validity 3000
What is your first and last name?
  [Unknown]:  Monika Cuellar
What is the name of your organizational unit?
  [Unknown]:  Peer 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 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:

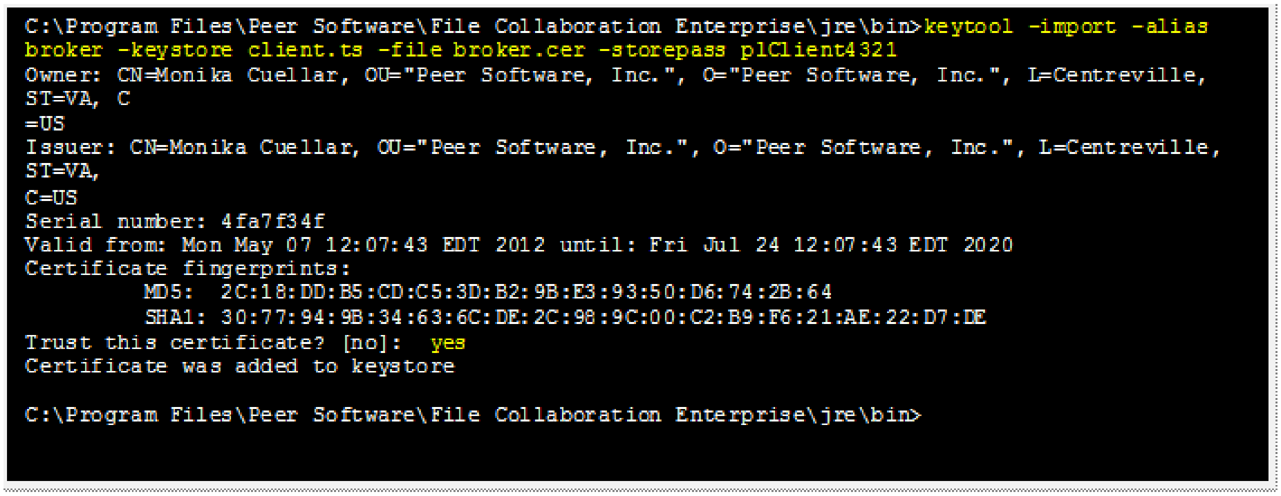

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -import -alias
broker -keystore client.ts -file broker.cer -storepass plClient4321
Owner: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=VA, C
=US
Issuer: CN=Monika Cuellar, OU="Peer 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:

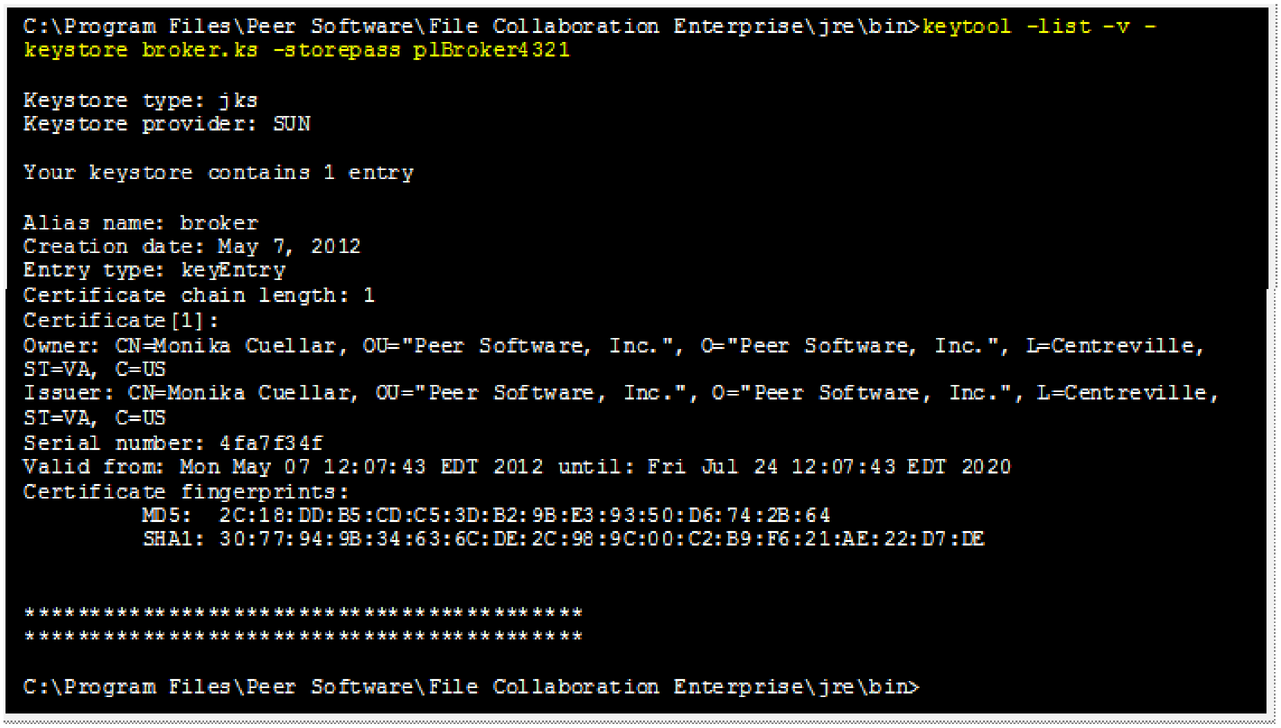

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -list -v -
keystore broker.ks -storepass plBroker4321

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:

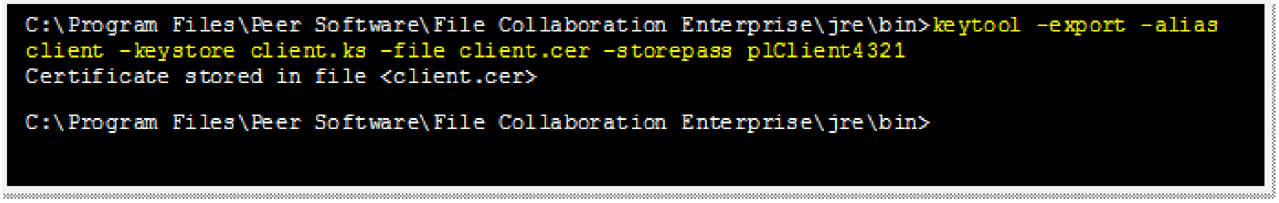

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -export -alias
client -keystore client.ks -file client.cer -storepass plClient4321
Certificate stored in file <client.cer>

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:

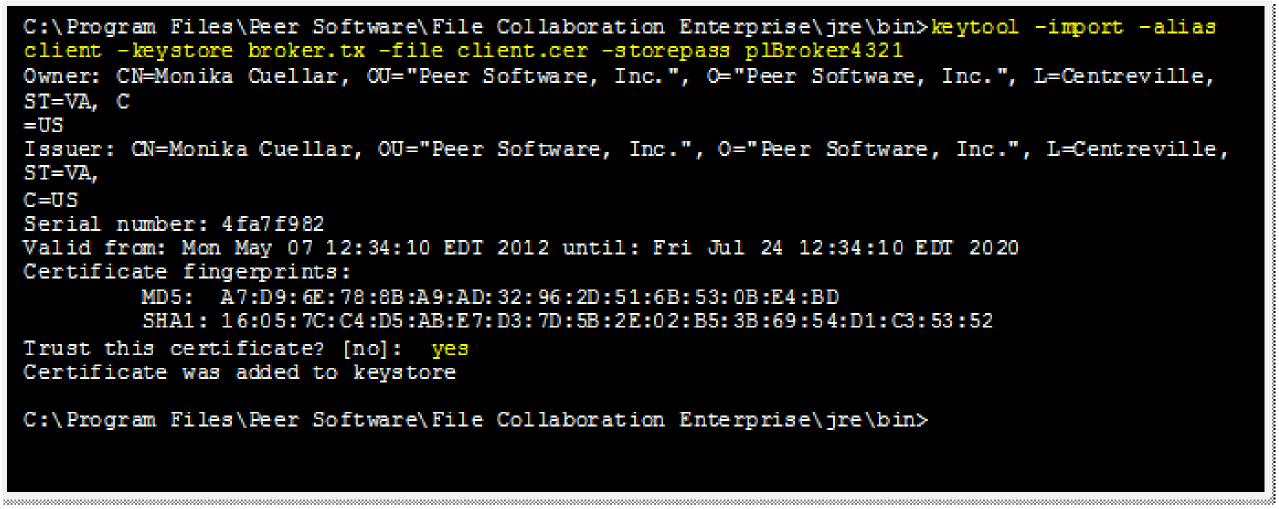

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -import -alias
client -keystore broker.tx -file client.cer -storepass plBroker4321
Owner: CN=Monika Cuellar, OU="Peer Software, Inc.", O="Peer Software, Inc.", L=Centreville,
ST=VA, C
=US
Issuer: CN=Monika Cuellar, OU="Peer 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:

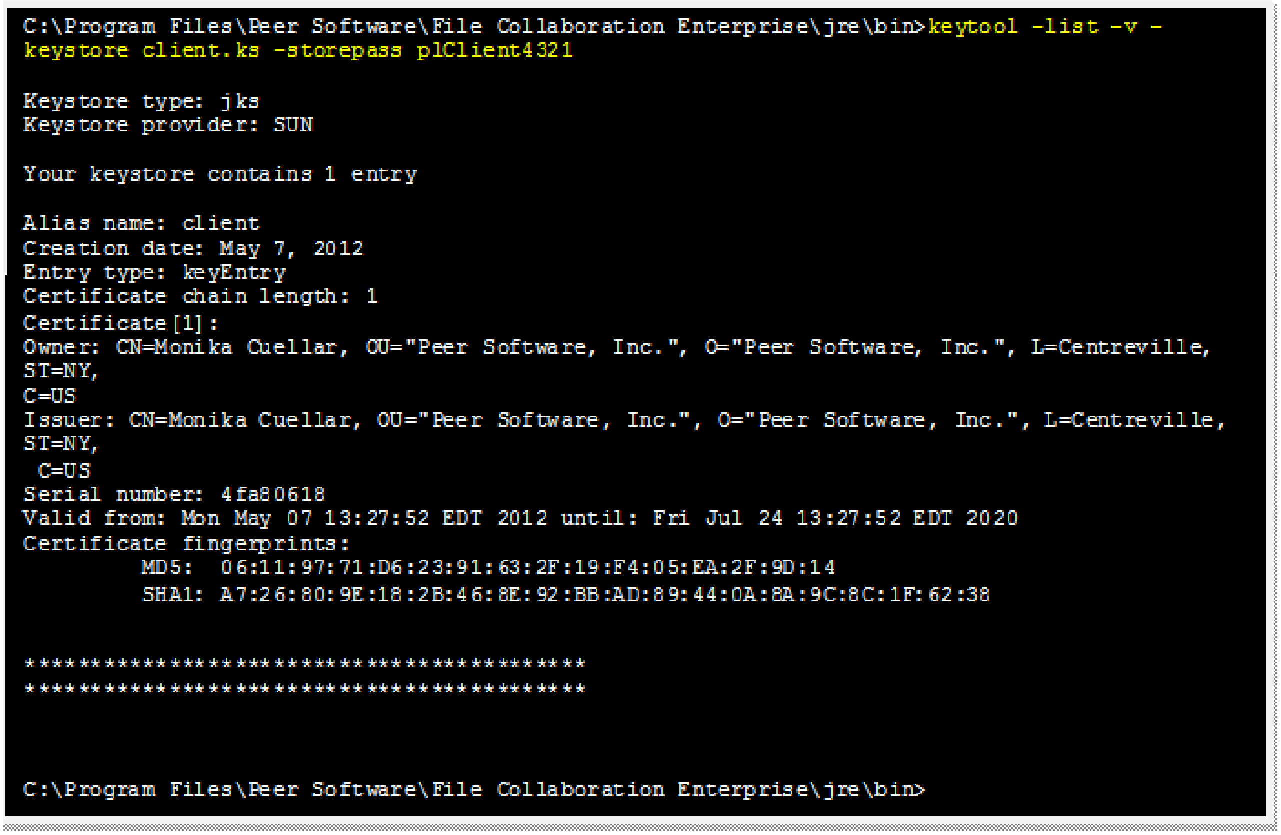

```
C:\Program Files\Peer Software\File Collaboration Enterprise\jre\bin>keytool -list -v -
keystore client.ks -storepass plClient4321

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

Perform the necessary commands using the keytool application bundled with your Peer Management Center or Peer Agent installation.

Keytool location on Peer Management Center system:	PMC_INSTALLATION_FOLDER\jre\bin
Keytool location on 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. View/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 the 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 application 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
```

plClient4321	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 application 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 (i.e. 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 (i.e., 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 trustore 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 trustore 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 "PMC_INSTALLATION_FOLDER\jre\bin" directory into the "PMC_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 the "PMC_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. The steps outlining this process will be posted shortly.

Preferences

Overview

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.

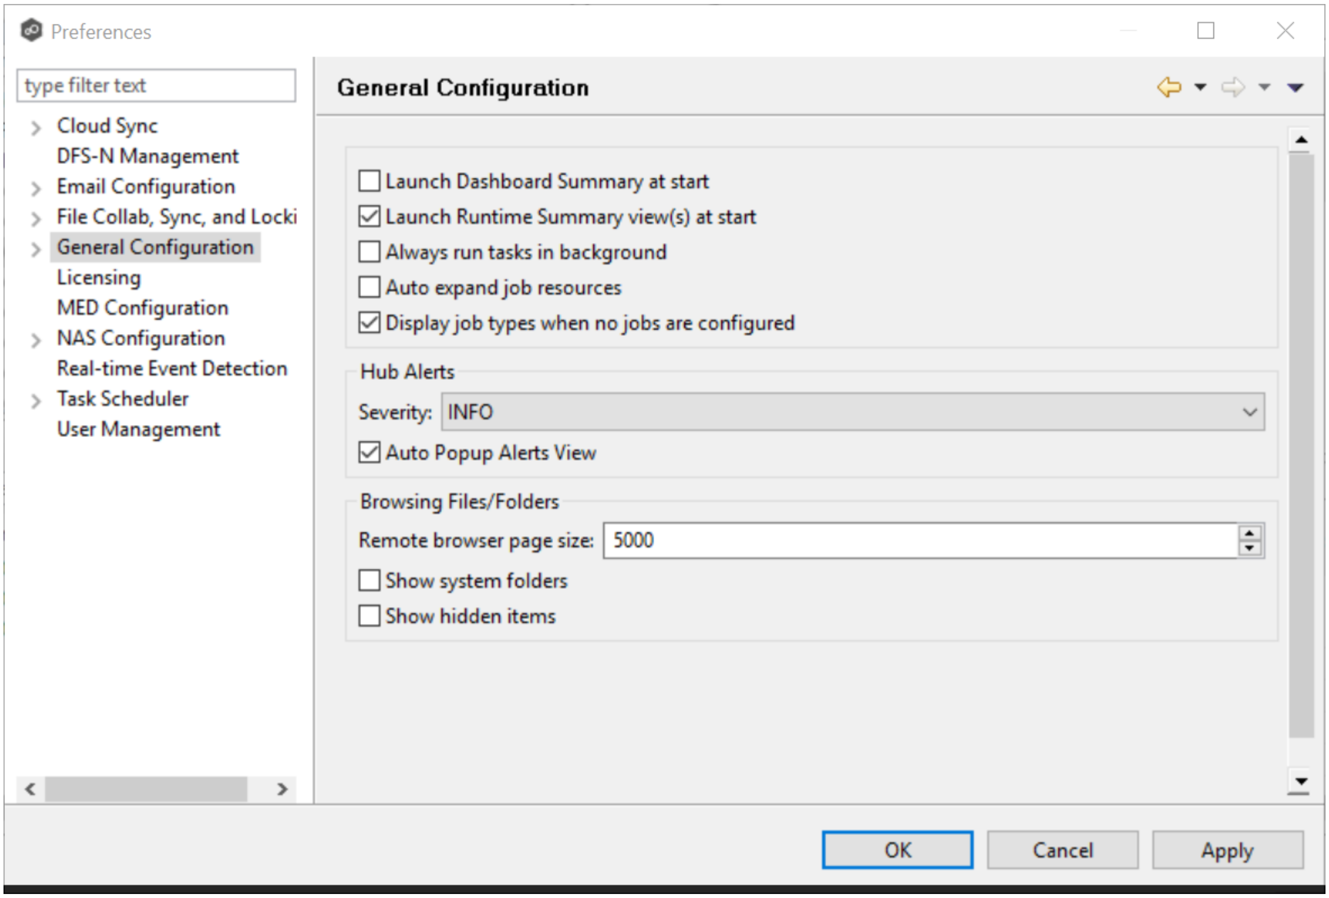

General Configuration

☐ Launch Dashboard Summary at start
☑ Launch Runtime Summary view(s) at start
☐ Always run tasks in background
☐ Auto expand job resources
☑ Display job types when no jobs are configured

Hub Alerts

Severity: INFO
☑ Auto Popup Alerts View

Browsing Files/Folders

Remote browser page size: 5000
☐ Show system folders
☐ Show hidden items

Configuring Global Settings

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

- SMTP 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

Cloud Sync

- Email

- File Filter

File Collaboration, File Synchronization, File Locking, and File Replication

- Email

- File Filter

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.

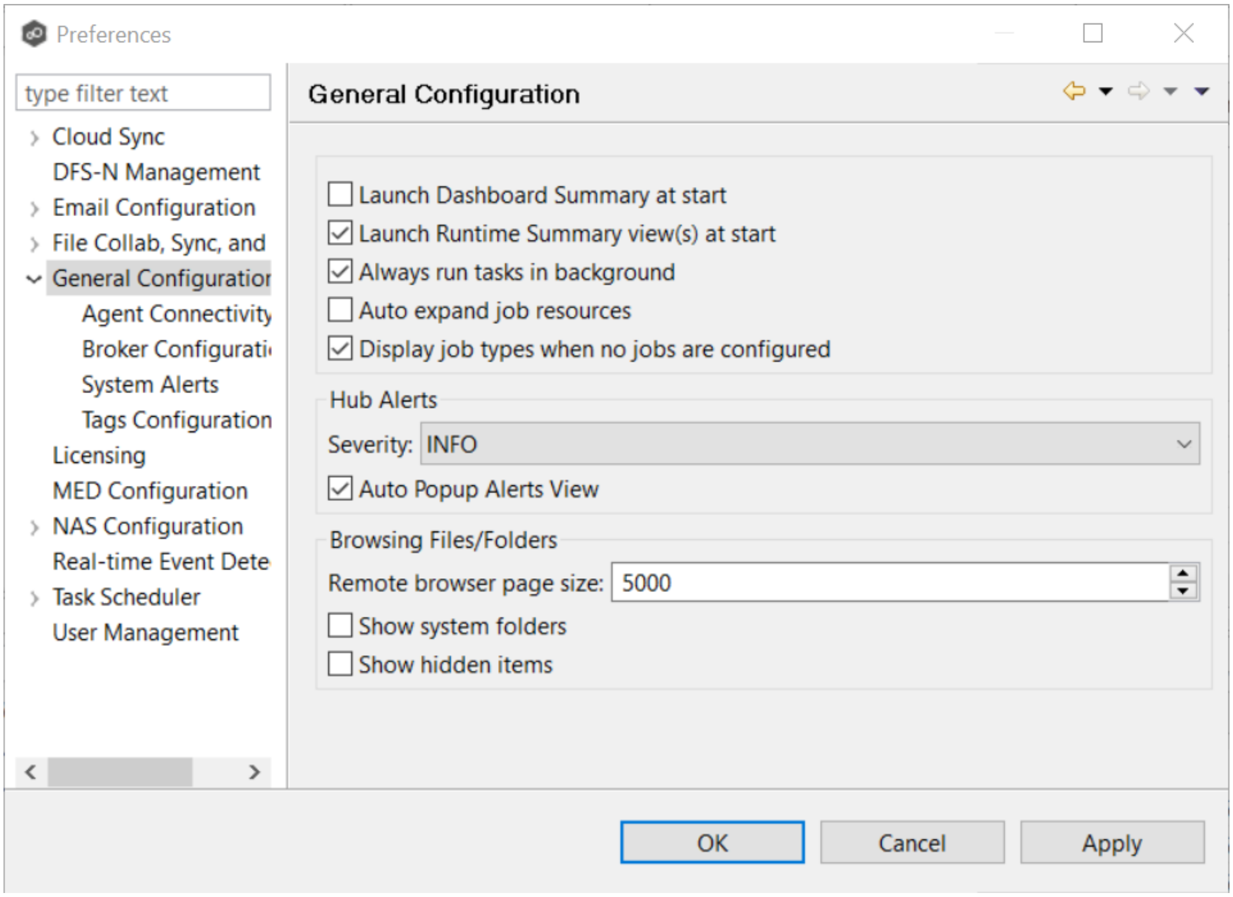

General Configuration

☐ Launch Dashboard Summary at start
☑ Launch Runtime Summary view(s) at start
☑ Always run tasks in background
☐ Auto expand job resources
☑ Display job types when no jobs are configured

Hub Alerts

Severity: INFO

☑ Auto Popup Alerts View

Browsing Files/Folders

Remote browser page size: 5000

☐ Show system folders
☐ Show hidden items

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

- **OK** 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.

Cloud Sync

You can modify the following Cloud Sync preferences:

- Cloud Sync
- Cloud Platform Credentials
- Database Connections

- Email Alerts
- File Filters
- Performance
- Scan Manager
- Sync and Retention Policies

Cloud Sync

Cloud Sync settings control the overall performance of all Cloud Sync jobs.

To modify these settings:

1. Select **Preferences** from the **Window** menu.
2. Select **Cloud Sync** in the navigation tree.

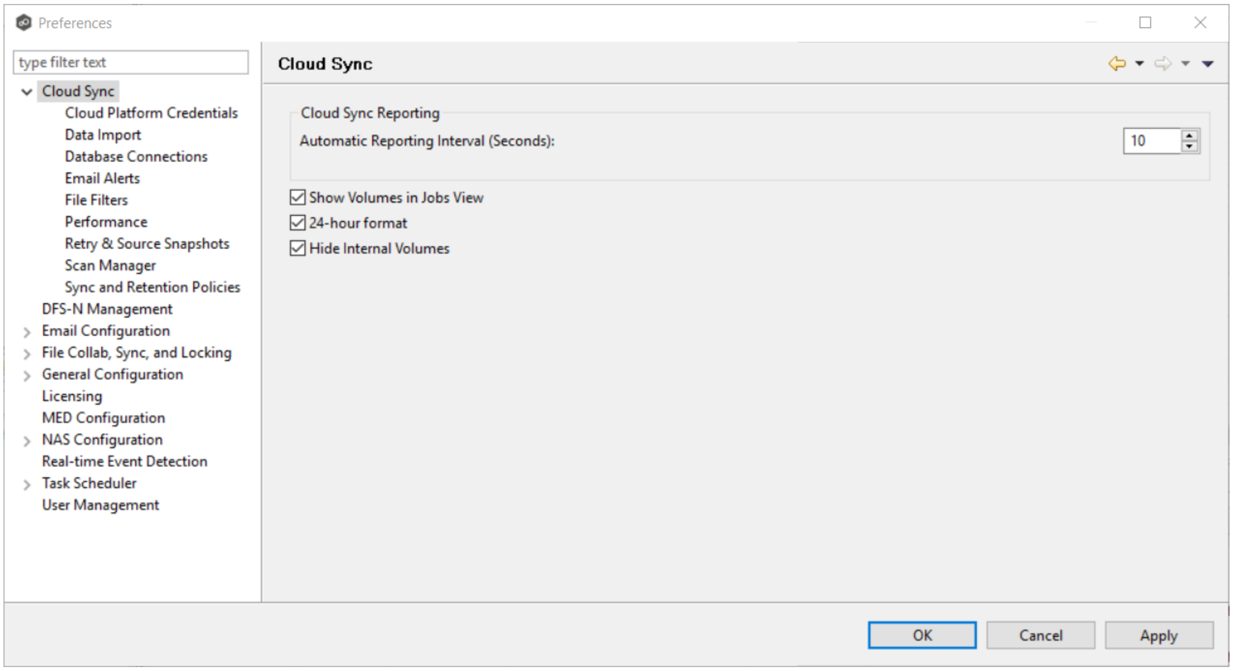

Cloud Sync

Cloud Sync Reporting

Automatic Reporting Interval (Seconds): 10

☑ Show Volumes in Jobs View

☑ 24-hour format

☑ Hide Internal Volumes

3. Modify the settings as needed.
4. Click **OK**.

Automatic Reporting Interval (Seconds)	Each Peer Agent automatically reports its statistics to the 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 check box if you want volumes to be displayed in the Jobs view.
Use 24-hour format	Select this check box if you want times to be displayed in a 24-hour format rather than a 12-hour format.
Hide Internal Volumes	Select this check box if you don't want internal volumes displayed when choosing which volumes to replicate.

Cloud Platform Credentials

Cloud Platform credentials are defined globally in **Preferences** in the **Cloud Platform Credentials** page. The page lists the existing user names and passwords used to grant access to your Azure Storage account. You can view, add, edit, and remove credentials.

When you create a job, you can select existing credentials to apply to the job or you can create a new filter and apply it to the job. You cannot modify or delete credentials when they are applied to a job.

To create new cloud platform credentials:

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

2. Expand **Cloud Sync** in the navigation tree, and then select **Cloud Platform Credentials**.

   The existing credentials are listed in the **Cloud Platform Credentials** table

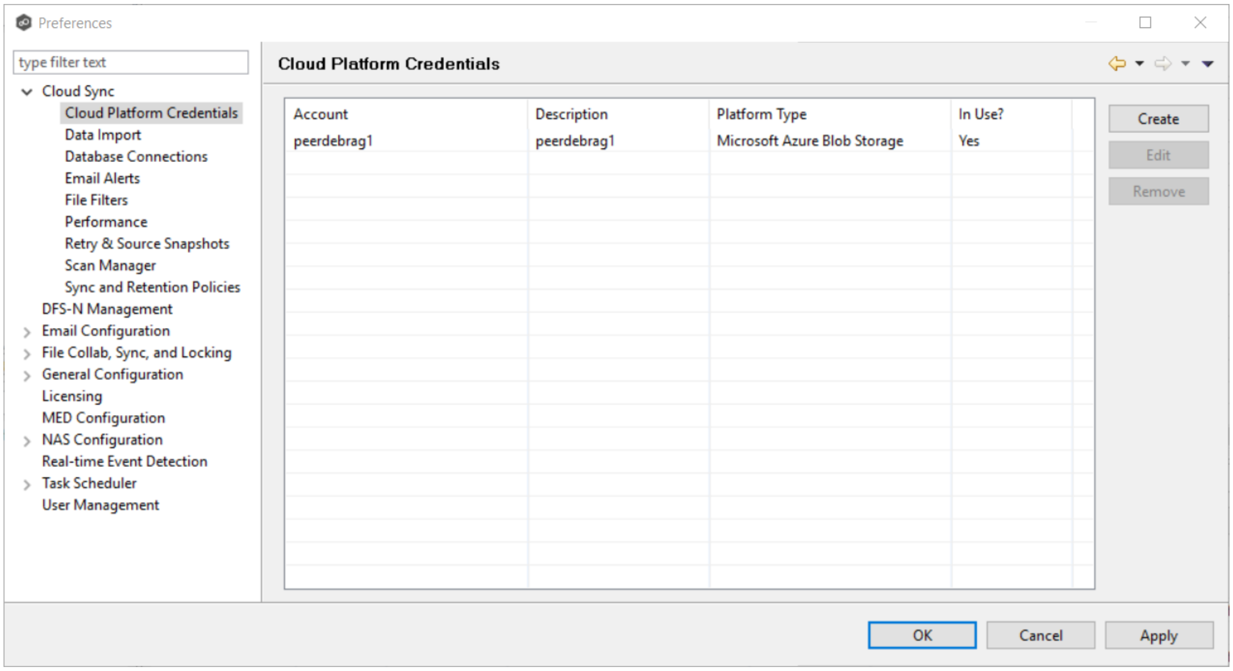

Cloud Platform Credentials

Account	Description	Platform Type	In Use?
peerdebrag1	peerdebrag1	Microsoft Azure Blob Storage	Yes

3. Click the **Add** button.

   The **Storage Account** dialog appears.

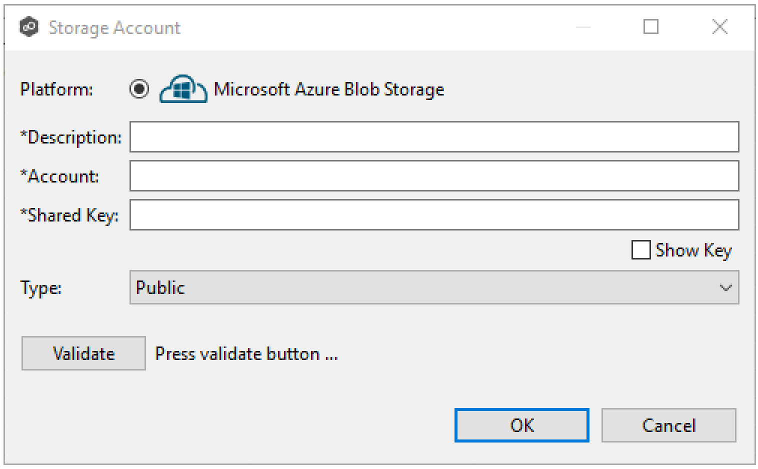

Platform: ● Microsoft Azure Blob Storage

*Description:

*Account:

*Shared Key:

□ Show Key

Type: Public

Press validate button ...

Enter the required values.

4. Click **Validate** to test the credentials.

5. Click **OK**.

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 of the Azure Storage account, which can be found in the Azure Portal.

Database Connections

Databases are used to track files and folders that have been replicated, individual file versions, and snapshots. Connections to your Microsoft SQL Server databases are defined globally in **Preferences** in the **Data Sources** page. You can view, add, edit, and remove database connections. You cannot modify or delete a database connection when it applied to a job.

To create a new database connection:

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

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

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

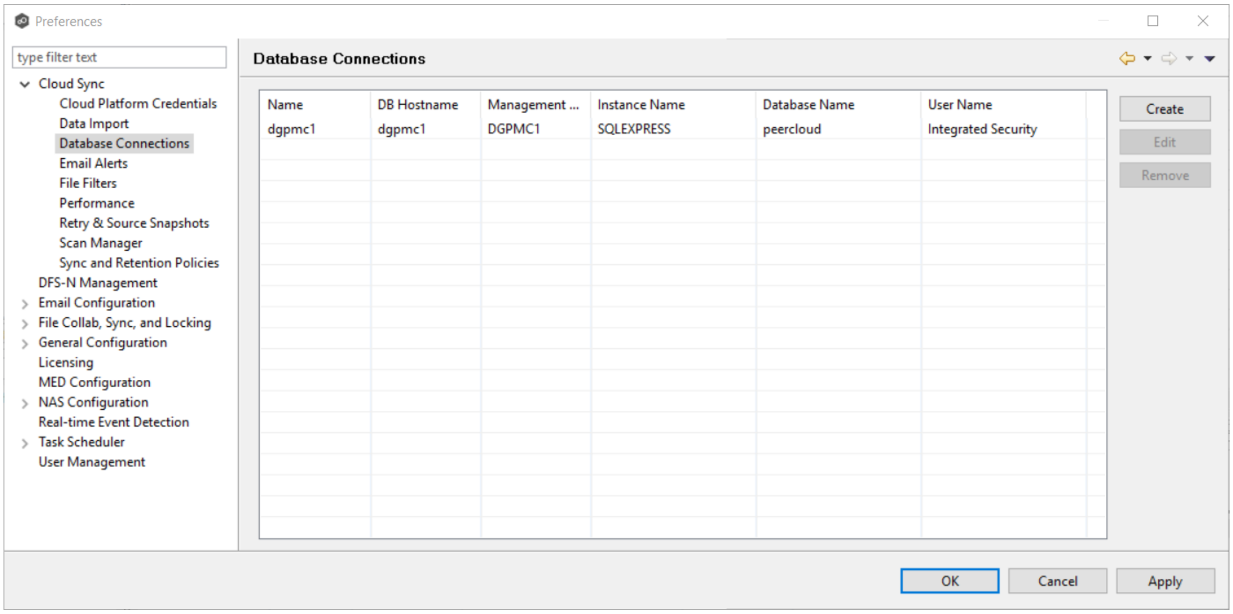

Name	DB Hostname	Management ...	Instance Name	Database Name	User Name
dgpmc1	dgpmc1	DGPMC1	SQLEXPRESS	peercloud	Integrated Security

3. Click the **Create** button.

The **Add Database Connection** dialog appears.

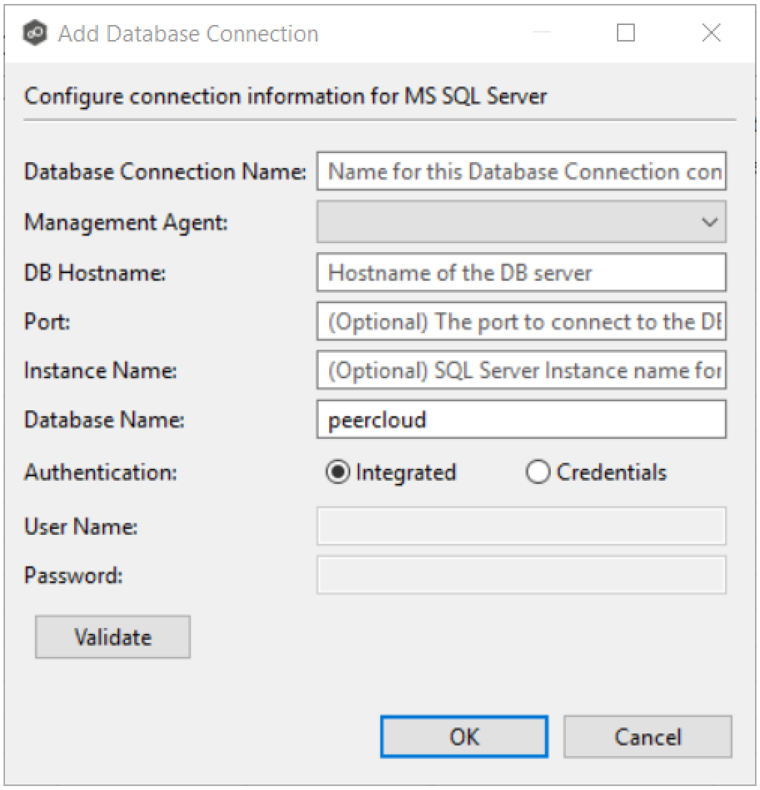

Add Database Connection

Configure connection information for MS SQL Server

Database Connection Name: Name for this Database Connection con

Management Agent:

DB Hostname: Hostname of the DB server

Port: (Optional) The port to connect to the DI

Instance Name: (Optional) SQL Server Instance name foi

Database Name: peercloud

Authentication: ● Integrated ○ Credentials

User Name:

Password:

4. Enter the required values.

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.

Database Connection Name	Enter a name for this database connection.
Management Agent	Select the Agent that will use this connection.
DB Hostname	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.
Port	Enter the port to be used to communicate with the specified SQL Server. If not defined, the connection defaults to port 1433.
Instance Name	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 empty.
Database Name	Enter the name of the database that Cloud Sync 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.
User Name	Enter the user name of the database administrator account to be used by Cloud Sync. This can be a locally-defined account such as "sa" or a domain account.
Password	Enter the password of the database administrator account to be used by Cloud Sync.

Email Alerts

Email alerts can be sent to notify users when a certain type of event occurs, for example, session abort, host failure, system alert. You can view, add, edit, and remove alerts. See Email Alerts in PMC Concepts for more information about email alerts.

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. You cannot modify or delete an email alert when it applied to a job.

To create an email alert:

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

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

   The existing Cloud Sync email alerts are listed in the **Email Alerts** table.

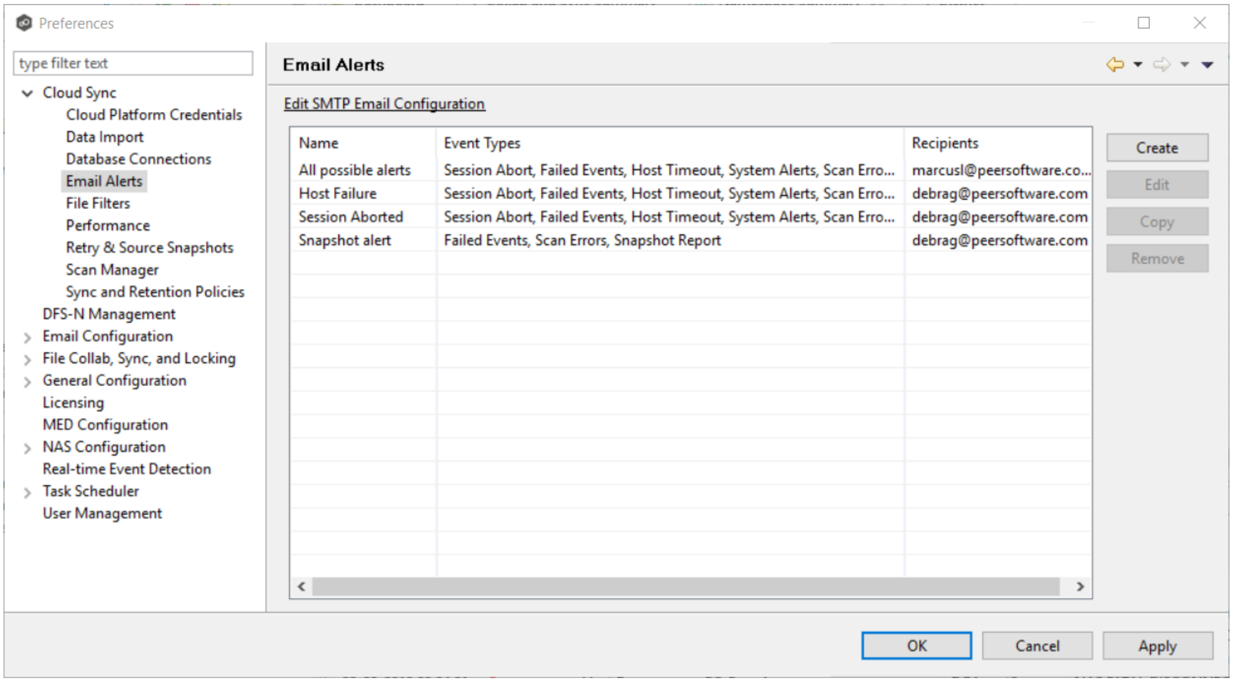

Name	Event Types	Recipients
All possible alerts	Session Abort, Failed Events, Host Timeout, System Alerts, Scan Erro...	marcusl@peersoftware.co...
Host Failure	Session Abort, Failed Events, Host Timeout, System Alerts, Scan Erro...	debrag@peersoftware.com
Session Aborted	Session Abort, Failed Events, Host Timeout, System Alerts, Scan Erro...	debrag@peersoftware.com
Snapshot alert	Failed Events, Scan Errors, Snapshot Report	debrag@peersoftware.com

3. Click the **Create** button.

   The **Add Email Alert** dialog appears.

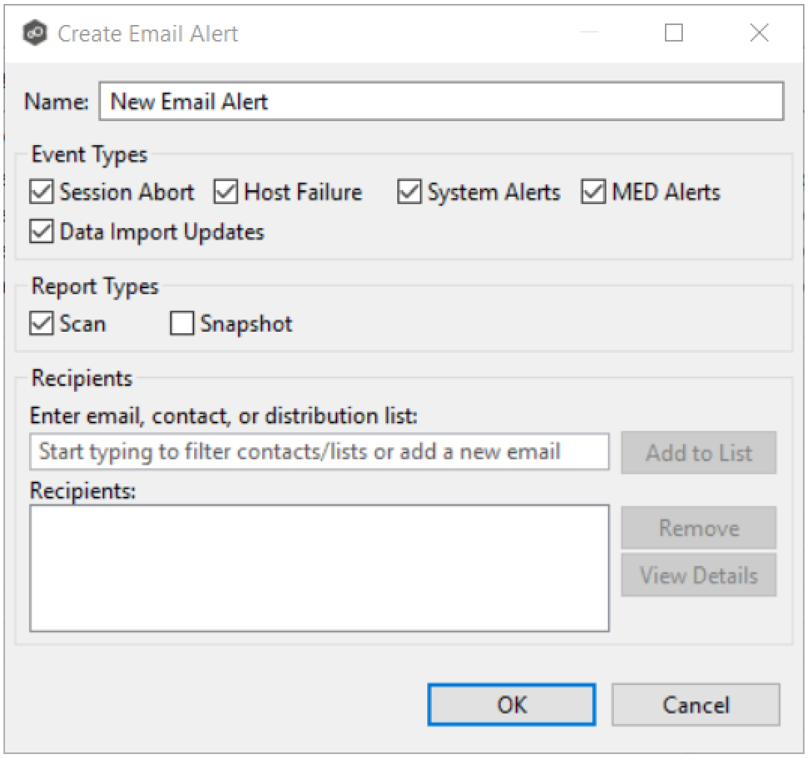

Name: New Email Alert

Event Types

☑ Session Abort ☑ Host Failure ☑ System Alerts ☑ MED Alerts

☑ Data Import Updates

Report Types

☑ Scan □ Snapshot

Recipients

Enter email, contact, or distribution list:

Start typing to filter contacts/lists or add a new email

Recipients:

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.

6. Select the report types to be sent.

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

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

8. Click **OK**.

Event Types

Session Abort	Sends an alert when the Cloud Sync job stops unexpectedly.
Host Failure	Sends an alert when the Management Agent of a Cloud Sync job disconnects or stops responding.

System Alerts	Sends an alert when a system event such as low memory or low hub disk space occurs.

Report Types

Scan	Sends scan statics after a scan has completed.
Snapsho t	Sends the status of the snapshot after the snapshot is triggered.

File Filters

File filters allow you to specify files and folders to exclude from or include in a job. You can view, add, edit, and remove filters. See File Filters in PMC Concepts for more information about file filters.

When you create a job, you can select an existing file filter to apply to the job or you can create a new filter and apply it to the job. You cannot modify or delete a file filter when it is being used by a job.

To create a file filter:

1. Select **Preferences** from the **Window** menu.
2. Expand **Cloud Sync** in the navigation tree, and then select **File Filters**.

   The existing Cloud Sync file filters are listed in the **File Filters** table

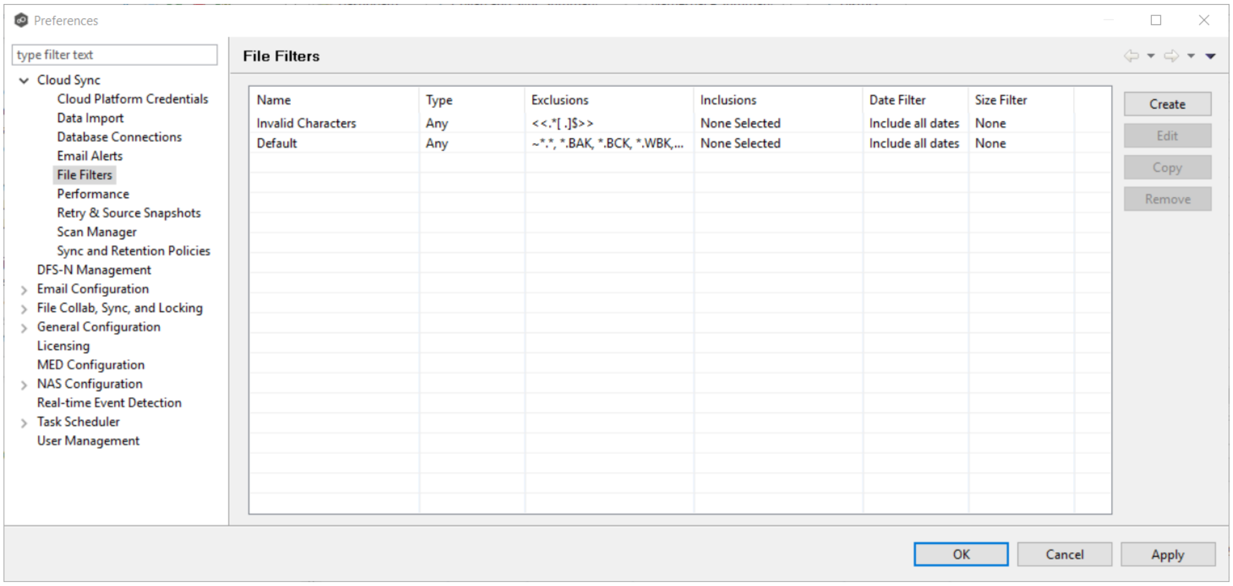

Name	Type	Exclusions	Inclusions	Date Filter	Size Filter
Invalid Characters	Any	<<.*[.]$>>	None Selected	Include all dates	None
Default	Any	~*.*, *.BAK, *.BCK, *.WBK,...	None Selected	Include all dates	None

3. Click the **Create** button.

   The **Create File Filter** dialog appears.

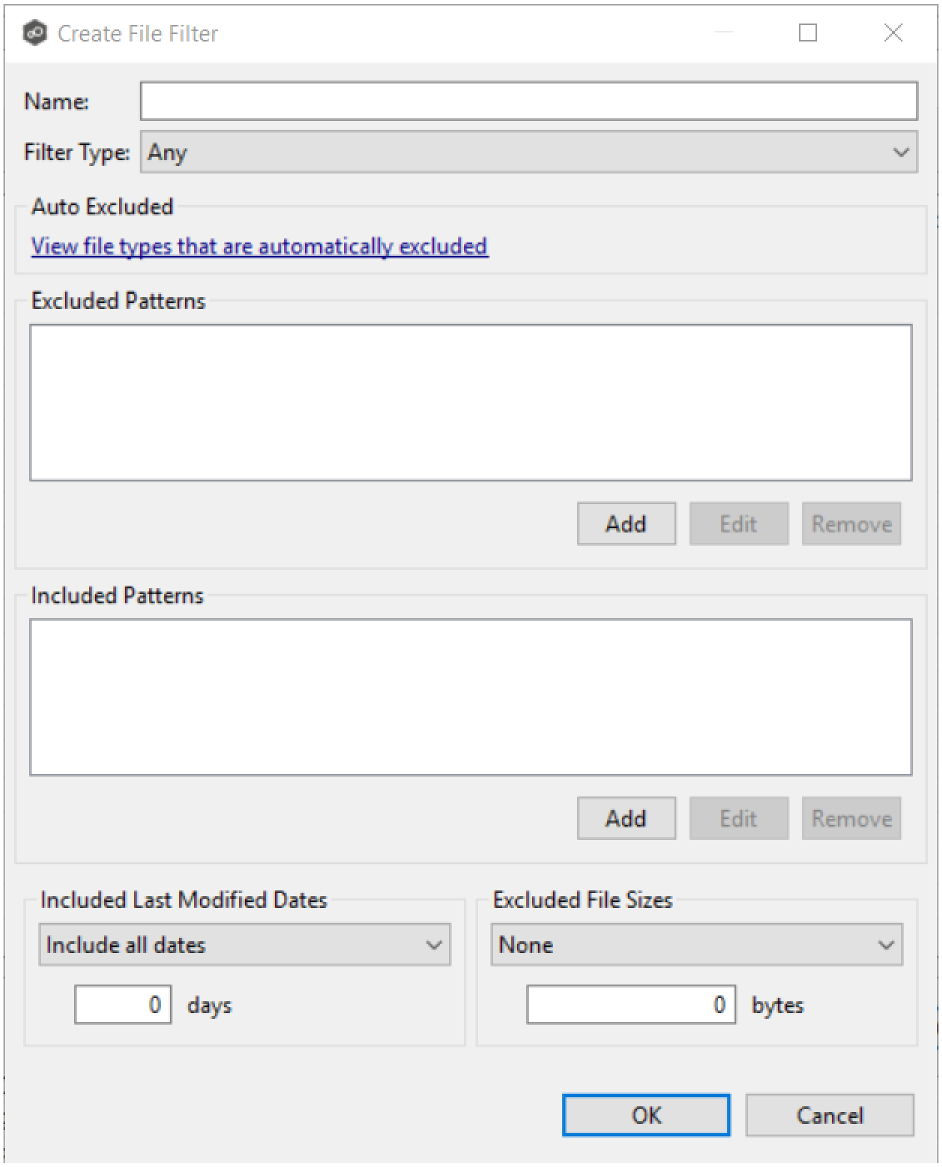

Name:

Filter Type: Any

Included Last Modified Dates: Include all dates

0 days

Excluded File Sizes: None

0 bytes

4. Enter a unique name for the filter.

5. Select the filter type.

6. (Optional) Click **Add** to enter filter patterns for files that you want excluded from the job.

7. (Optional) Click **Add** to enter filter patterns for files that you want included from the job.

8. (Optional) Select a value for Included Last Modified Dates.

9. (Optional) Select a value for Excluded File Sizes.

10. Click **OK**.

    See File Filters in the **Concepts** section for information on entering values for the fields.

Performance

Performance settings allow you to adjust Cloud Sync performance.

To modify the Cloud Sync performance settings:

1. Select **Preferences** from the **Window** menu.
2. Expand **Cloud Sync** in the navigation tree, and then select **Performance**.

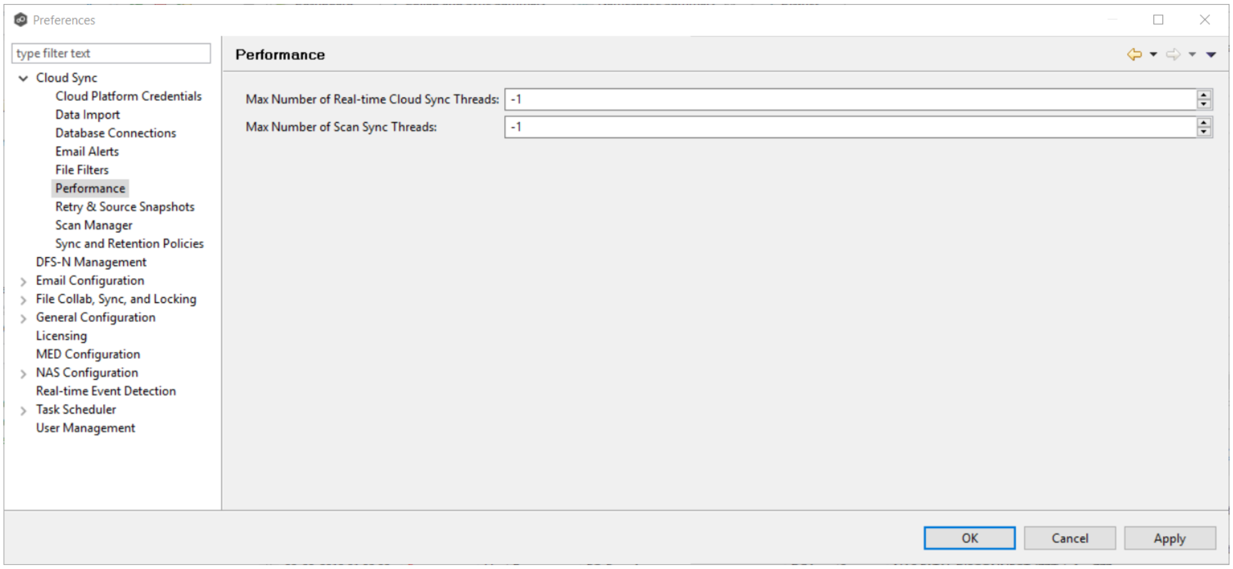

Max Number of Real-time Cloud Sync Threads: -1

Max Number of Scan Sync Threads: -1

3. Modify the retry settings as needed:

Max Number of Real-time Cloud Sync Threads	Enter the maximum number of threads available for replicating files as they are updated in real-time on the source storage device.
Max Number of Scans Synch. 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 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.

4. Modify the VSS settings as needed:

VSS Retry Interval in minutes	Enter the number of minutes to wait in between retries of the failed replication of a file or folder.
Max Number of VSS Scans Synch. 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 VSS 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.

5. Click **OK**.

Scan Manager

The 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 Sync jobs:

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

2. Expand **Cloud Sync** in the navigation tree, and then select **Scan Manager**.

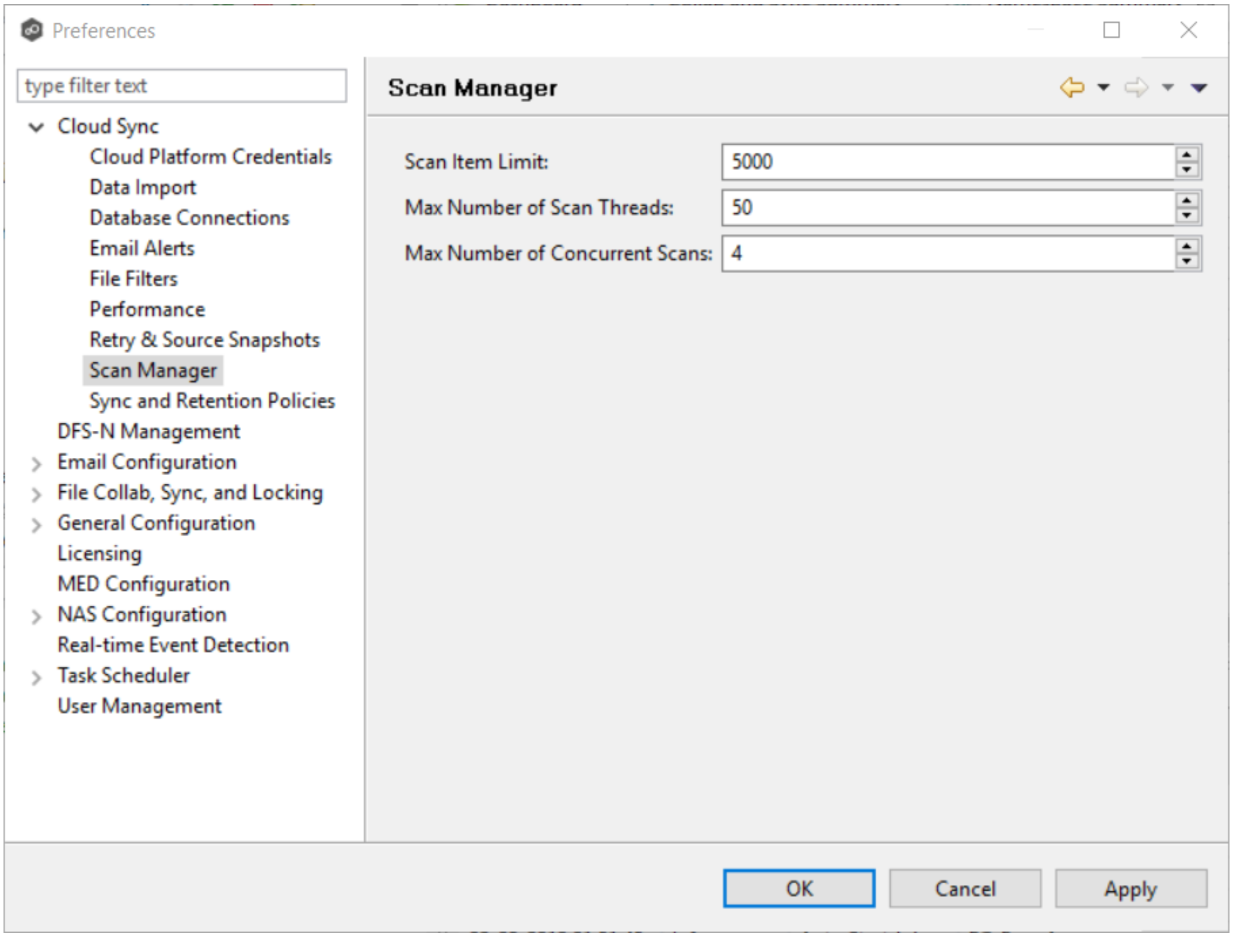

Scan Manager

Scan Item Limit: 5000

Max Number of Scan Threads: 50

Max Number of Concurrent Scans: 4

3. Modify the settings as needed.

Scan Item Limit	Enter the maximum number of files and folders to obtain from a folder structure at a time during a scan.
Max Number of Scan Threads	Enter the maximum number of threads available for scanning files and folders. This number should be set 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 **OK** or **Apply**.

Sync and Retention Policies

Each Cloud Sync job must have a Sync and Retention Policy applied to it. Policies are defined globally in **Preferences** in the **Sync and Retention Policies** page. You can view, add, edit, and remove policies.

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 cannot modify or delete a policy when it is being used by a job.

To create a new policy:

1. Select **Preferences** from the **Window** menu.
2. Expand **Cloud Sync** in the navigation tree, and then select **Scan and Retention Policies**.

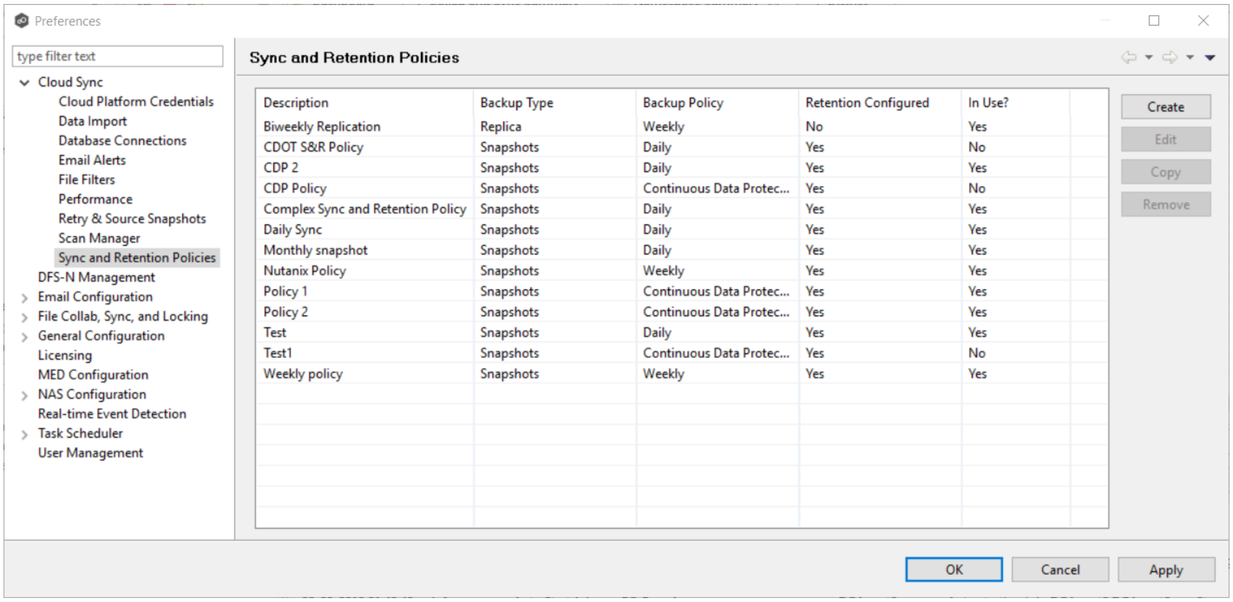

Sync and Retention Policies

Description	Backup Type	Backup Policy	Retention Configured	In Use?
Biweekly Replication	Replica	Weekly	No	Yes
CDOT S&R Policy	Snapshots	Daily	Yes	No
CDP 2	Snapshots	Daily	Yes	Yes
CDP Policy	Snapshots	Continuous Data Protec...	Yes	No
Complex Sync and Retention Policy	Snapshots	Daily	Yes	Yes
Daily Sync	Snapshots	Daily	Yes	Yes
Monthly snapshot	Snapshots	Daily	Yes	Yes
Nutanix Policy	Snapshots	Weekly	Yes	Yes
Policy 1	Snapshots	Continuous Data Protec...	Yes	Yes
Policy 2	Snapshots	Continuous Data Protec...	Yes	Yes
Test	Snapshots	Daily	Yes	Yes
Test1	Snapshots	Continuous Data Protec...	Yes	No
Weekly policy	Snapshots	Weekly	Yes	Yes

3. Click the **Create** button.

   The **Sync and Retention Policy Wizard** opens.

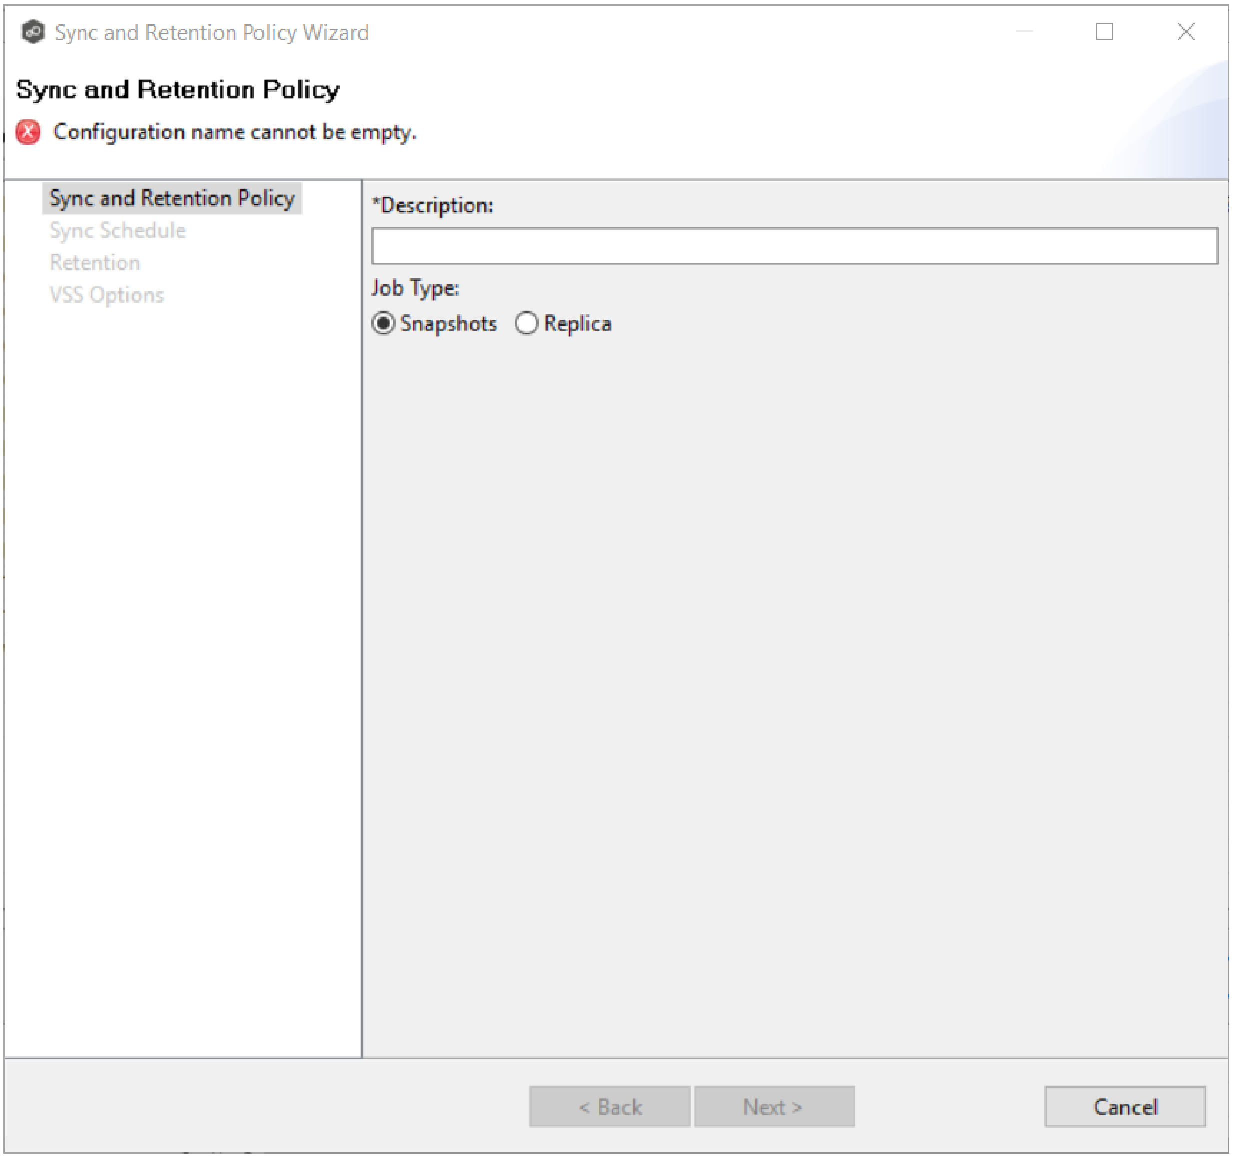
Sync and Retention Policy Wizard
Sync and Retention Policy
Sync and Retention Policy
Sync Schedule
Retention
VSS Options
< Back
Next >
Cancel

Configuration name cannot be empty.

*Description:

Job Type:

● Snapshots ○ Replica

4. Enter the required values, and then click **Finish**.

   See [Step 10: Sync and Retention Policy]() for assistance in completing the wizard.

DFS-N Management

To modify DFS-N Management settings:

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

The **Preferences** dialog appears.

2. Select **DFS-N Management** in the navigation tree.

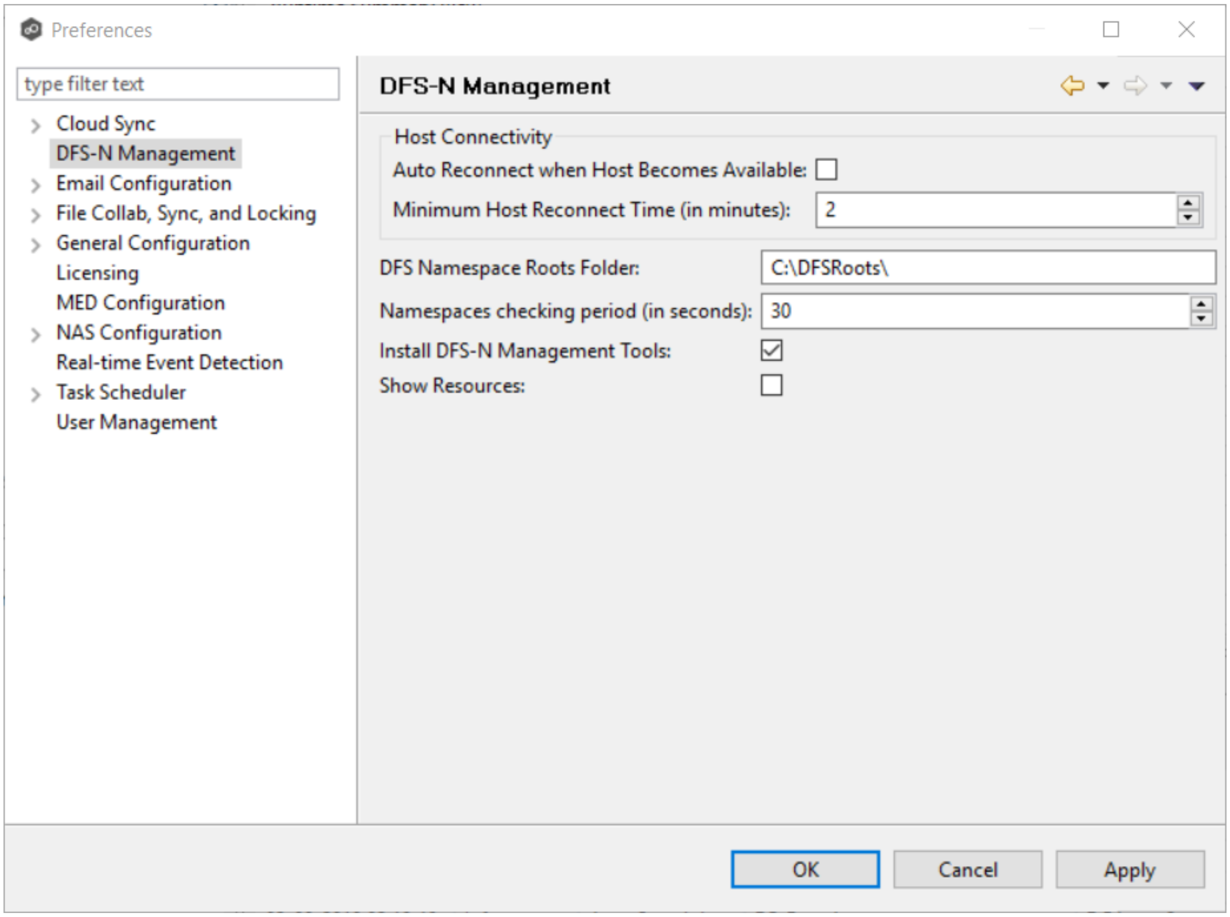

DFS-N Management

Host Connectivity

Auto Reconnect when Host Becomes Available: □

Minimum Host Reconnect Time (in minutes): 2

DFS Namespace Roots Folder: C:\DFSRoots\

Namespaces checking period (in seconds): 30

Install DFS-N Management Tools: ☑

Show Resources: □

3. Modify settings as needed.

Auto Reconnect when Host Becomes Available	If an Agent that is managing a DFS namespace goes offline, the DFS-N Management job will be stopped. If this auto-reconnect option is enabled, the DFS-N Management job will automatically restart when the Agent comes back online.
Minimum Host Reconnect Time (in minutes)	Tied to the auto-reconnect option above, the reconnect time controls how quickly the DFS-N Management job will restart after the Agent comes back online.
DFS Namespace Roots Folder	The default local parent folder for namespaces on each namespace server.
Namespaces checking period (in seconds)	Controls how often an Agent checks its assigned namespace using PowerShell. This check catches any changes made to a namespace from the Microsoft DFS Management tool.
Install DFS-N Management Tools	Enables the option to install the DFS Management PowerShell toolkit during while creating or importing a namespace.
Show Resources	Shows individual namespace folders under each namespace in the **Jobs** view.

4. Click **OK** or **Apply**.

Email Configuration

Before the Peer Management Center can send emails on behalf of any job, a few key SMTP settings must be configured. In addition, you can define contacts and distribution lists.

To configure SMTP settings:

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

   The **Preferences** dialog appears.

2. Select **Email Configuration** in the navigation tree.

   The following is displayed:

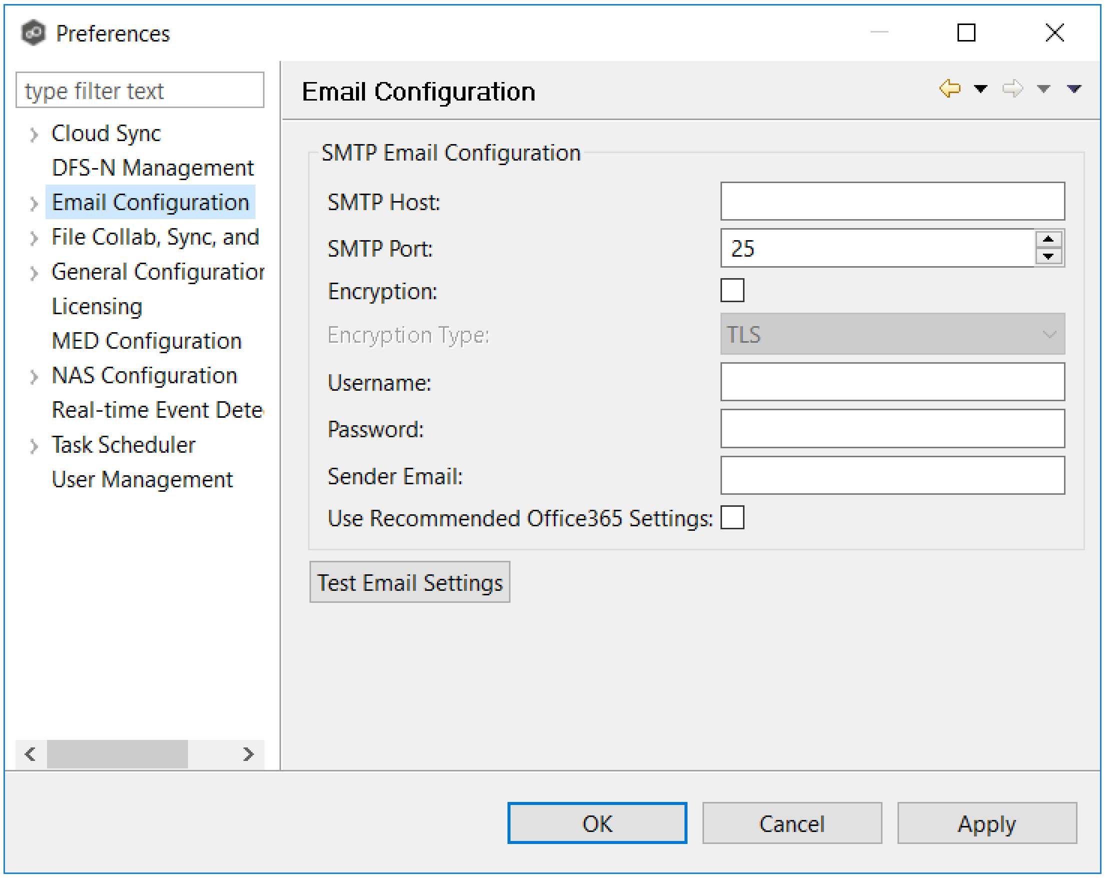

SMTP Host:

SMTP Port: 25

Encryption: □

Encryption Type: TLS

Username:

Password:

Sender Email:

Use Recommended Office365 Settings: □

3. Enter values for the following fields:

SMTP Host (required)	Enter the host name or IP address of the SMTP mail server through which the 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. T LS and SSL are the available options. If you do not know which one your mail server requires, try one, and then the other.
User	Enter the username to authenticate as on the SMTP mail server (optional).
Password	Enter the password of the username specified above (optional).
Sender Email (required)	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 Office365 Settings	Select this checkbox if you are connecting to an Office 365 SMTP server to use recommended settings for the connection. Follow Microsoft's **Direct Send** recommendations to set up Email configuration with an Office 365 SMTP server.

4. (Recommended) Click **Test Email Settings**.

   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, the Peer Management Center will attempt to send a test message using the specified settings.

5. Click **OK** or **Apply**.

File Collab, Sync, and Locking

You can modify the following preferences for File Collaboration, File Locking, and File Synchronization jobs:

- File Collab, Sync, and Locking
- DFS-N Management
- Email Alerts
- File Filters
- Locking
- Performance
- Real-time Event Detection
- Revit Enhancements
- SMNP Notifications
- Scan Manager

File Collab, Sync, and Locking

These settings control basic GUI and reconnect settings for all File Collaboration, File Synchronization, File Locking, and File Replication jobs.

To modify these settings:

1. Select **Preferences** from the **Window** menu.
2. Select **File Collab, Sync, and Locking** in the navigation tree.

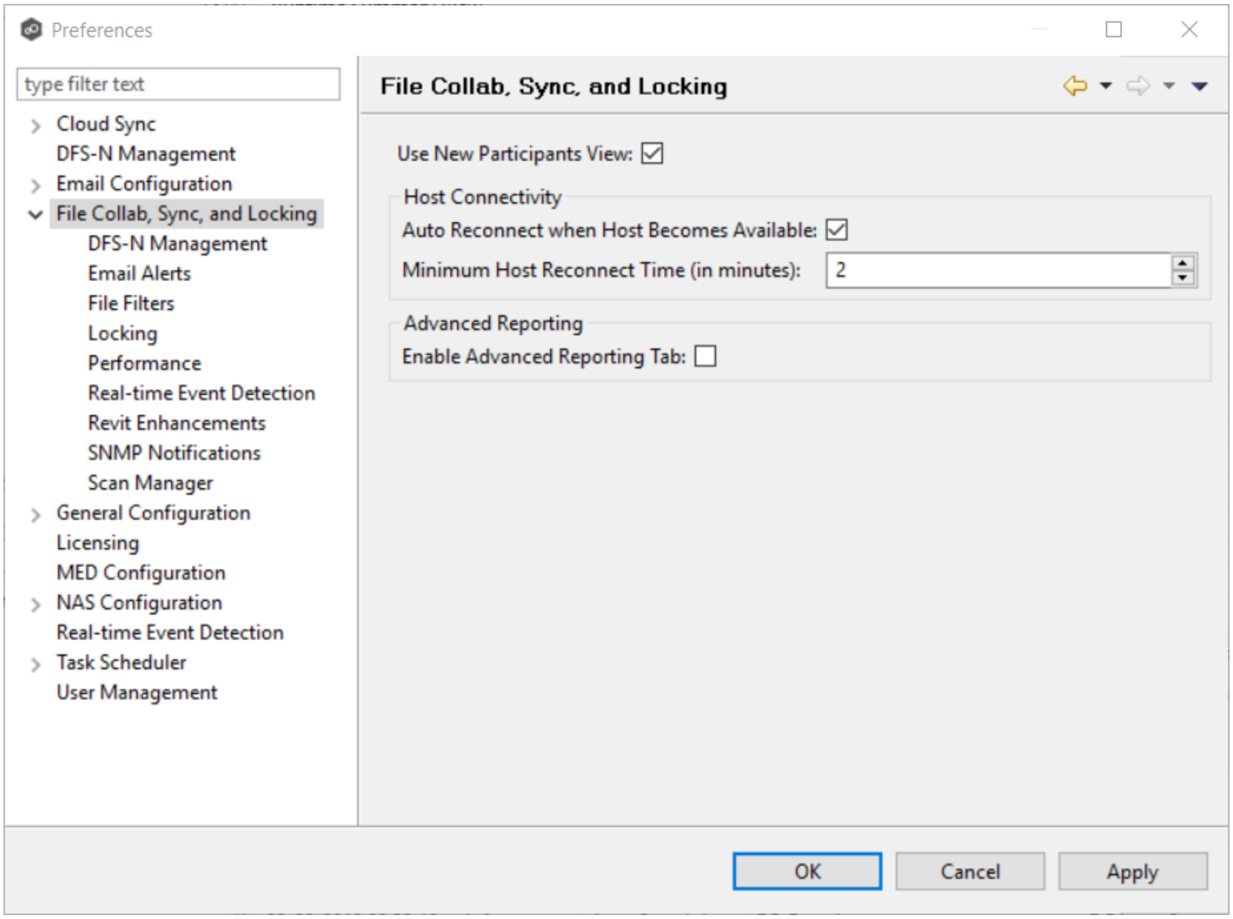

File Collab, Sync, and Locking

Use New Participants View: ☑

Host Connectivity

Auto Reconnect when Host Becomes Available: ☑

Minimum Host Reconnect Time (in minutes): 2

Advanced Reporting

Enable Advanced Reporting Tab: □

3. Modify the settings as needed.

Use New Participants View	When creating a new job, use the new **Add New Participant** wizard instead of the legacy participant view. ***Highly recommended.***
Auto Reconnect when Host Becomes Available	When an Agent reconnects to the PMC after a failure, automatically re-enable 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** sub-tab of the global **Collab and Sync Summary** view.

4. Click **OK** or **Apply**.

DFS-N Management

These settings control the basic interoperability of all DFS-N Management jobs with File Collaboration and File Synchronization jobs.

To modify these settings:

1. Select **Preferences** from the **Window** menu.
2. Select **File Collab, Sync, and Locking** in the navigation tree.
3. Select **DFS-N Management**.

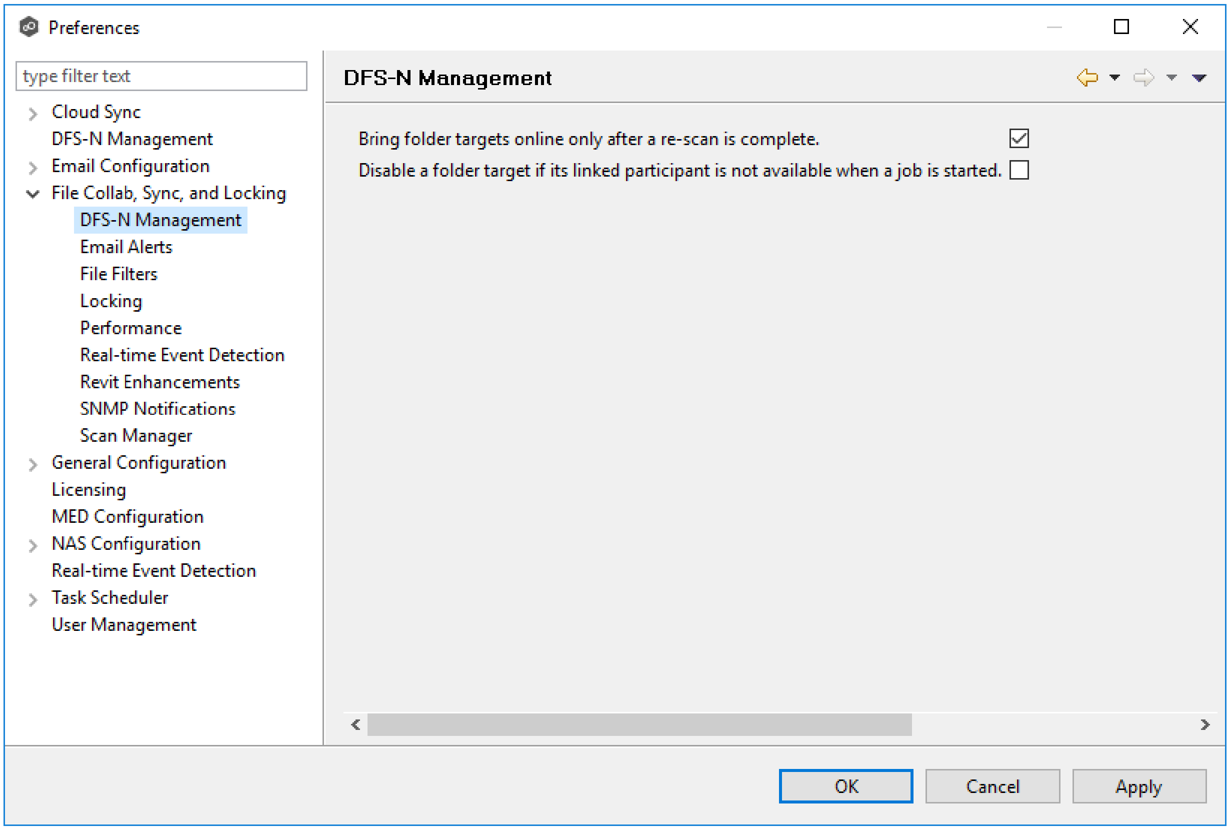

Bring folder targets online only after a re-scan is complete. ☑

Disable a folder target if its linked participant is not available when a job is started. □

4. Modify settings as needed.

Bring folder targets online only after a re-scan is complete	Re-enable a disabled folder target in a 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.

5. Click **OK** or **Apply**.

Email Alerts

For an overview of email alerts, see Email Alerts in the Basic Concepts section.

To create an email alert:

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

2. Expand **File Collab, Sync, and Locking** in the navigation tree, and then select **Email Alerts**.

   Any existing email alerts are listed in the **Email Alerts** table.

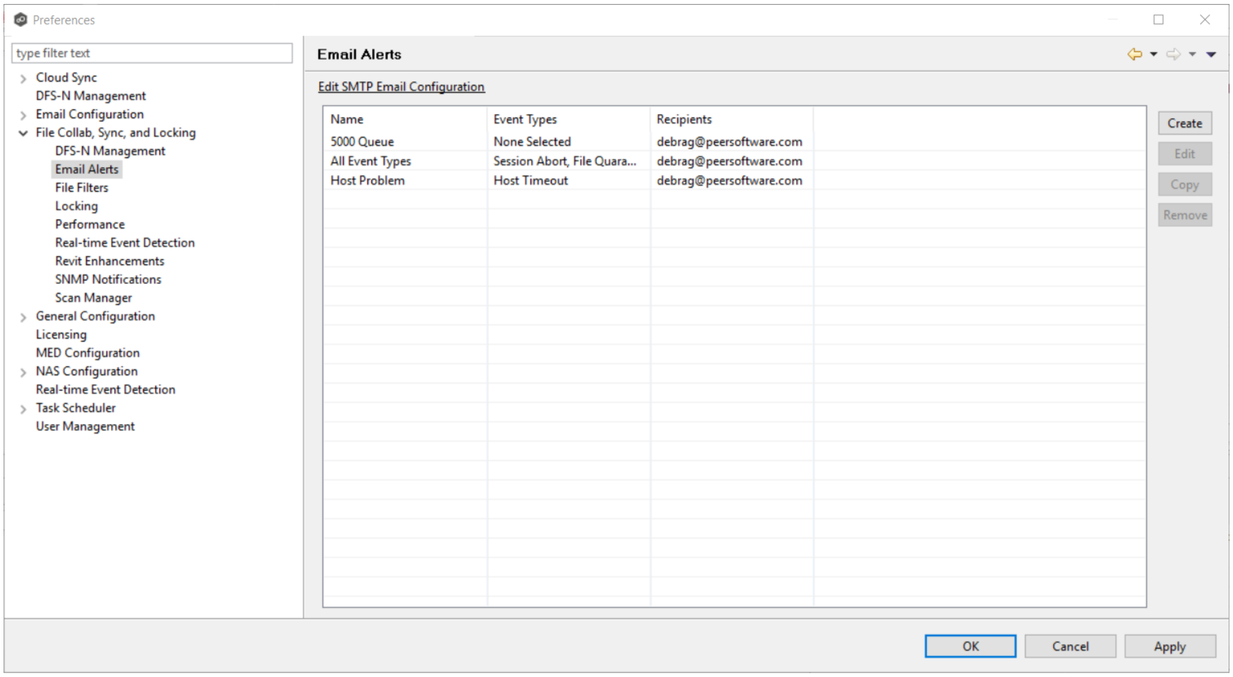

Name	Event Types	Recipients
5000 Queue	None Selected	debrag@peersoftware.com
All Event Types	Session Abort, File Quara...	debrag@peersoftware.com
Host Problem	Host Timeout	debrag@peersoftware.com

3. Click **Create**.

   The **Create Email Alert** dialog appears.

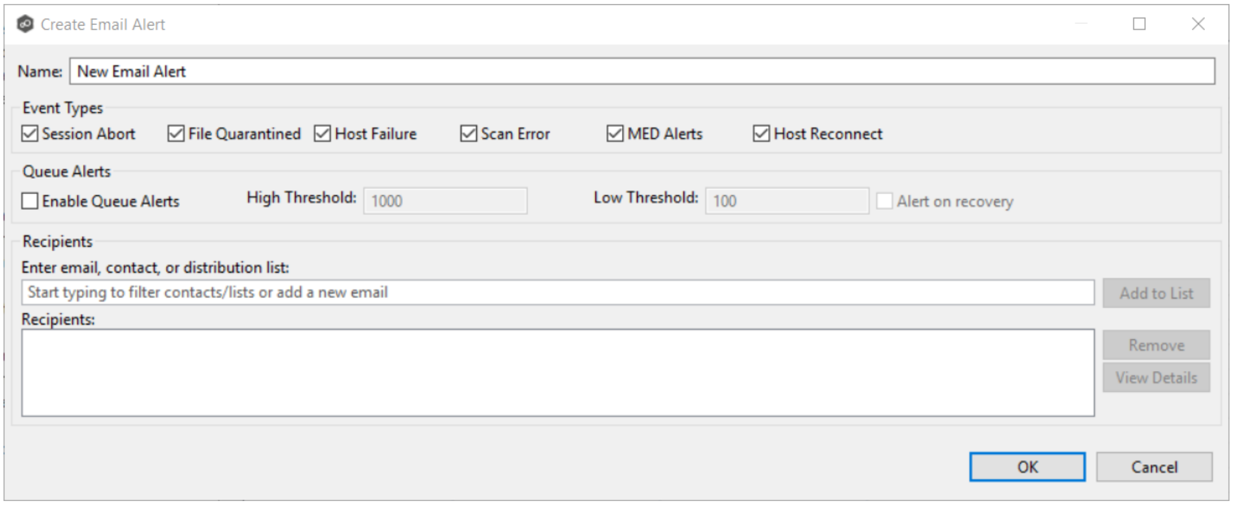

Create Email Alert

Name: New Email Alert

Event Types

☑ Session Abort ☑ File Quarantined ☑ Host Failure ☑ Scan Error ☑ MED Alerts ☑ Host Reconnect

Queue Alerts

□ Enable Queue Alerts High Threshold: 1000 Low Threshold: 100 □ Alert on recovery

Recipients

Enter email, contact, or distribution list:

Start typing to filter contacts/lists or add a new email

Recipients:

4. Enter a name for the alert.

5. Select the type of events for which you want alerts sent.

Session Abort	Enables sending an alert when a session is aborted because of lack of quorum due to one or more failed hosts.
File Quarantined	Enables sending an alert when a file is marked as quarantined because a file conflict was not able to be resolved.
Host Timeout	Enables sending an alert when a host timeout occurs and the host is taken out of session.
Scan Error	Enables sending an alert when an error occurs during the initial synchronization process.
MED Alerts	Enables sending an alert when Peer MED detects potentially malicious activity. For more information, see MED Configuration.
Host Reconnect	Alerts when host comes back online.

6. If you want queue alerts sent, select **Enable Queue Alerts** and enter threshold values.

Enable Queue Alerts	When enabled, sends email alerts when the **Queued Items** counter in the Collaboration Summary view exceeds the configured **High Threshold** value. This counter is the combination of the **Real-time** and **File Sync** queues as they are displayed in the user interface for the job. This counter is checked every 20 seconds and if it exceeds the configured **High Threshold**, an email alert is sent. Another alert will not be sent until the counter has dropped below the configured **Low Threshold** value and then exceeds the **High Threshold** value again.
High Threshold	The high value of the **Queued Items** counter on the Collaboration Summary view. When this value is exceeded, an email will be sent
Low Threshold	Once an email has been sent, no additional emails will be sent until the configured **Low Threshold** value is met and then the **High Threshold** value is met again.
Alert on recovery	The **Alert on recovery** option controls whether or not an email will be sent indicating that the counter has recovered to the **Low Threshold** value after an alert had been previously sent.

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 and can now be applied to jobs.

File Filters

A file filter enables you to specify which files (and folders) should be included and/or excluded from a job's watch set. For more information about file filters, see File and Folder Filters in the Basic Concepts section.

To create a file filter:

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

2. Expand **File Collab, Sync, and Locking** in the navigation tree, and then select **File Filters**.

   Any existing file filters are listed in the **File Filters** table.

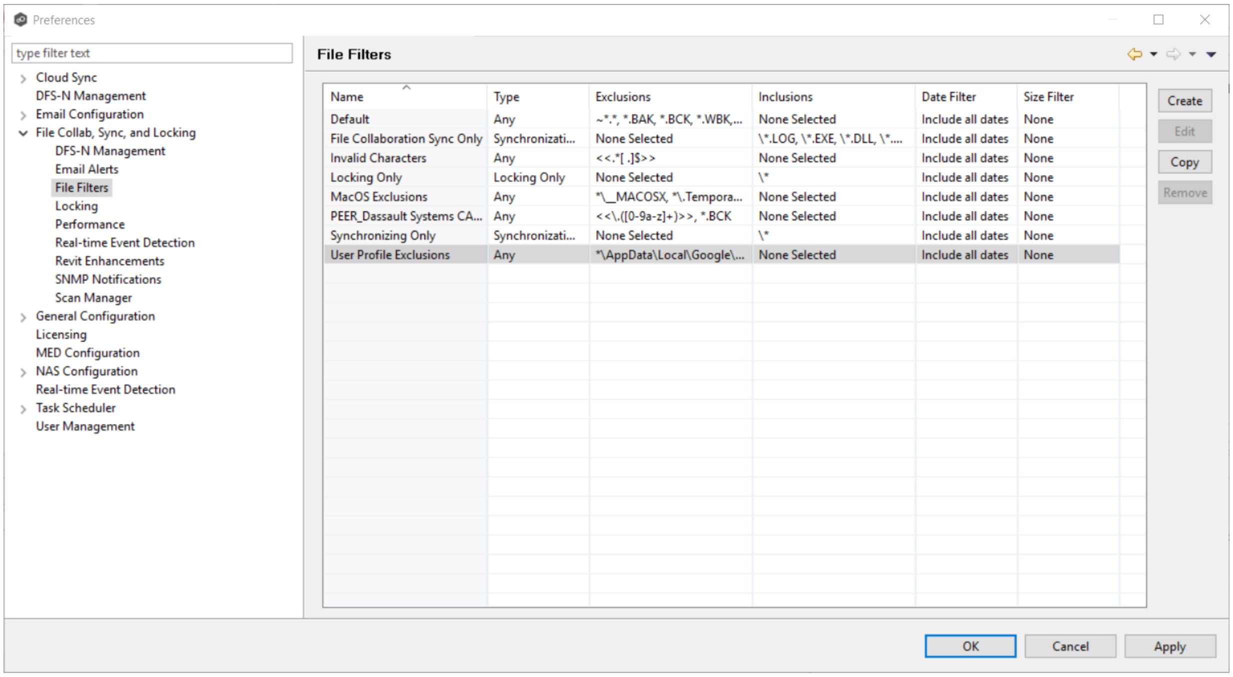

Name	Type	Exclusions	Inclusions	Date Filter	Size Filter
Default	Any	~*.*, *.BAK, *.BCK, *.WBK,...	None Selected	Include all dates	None
File Collaboration Sync Only	Synchronizati...	None Selected	\*.LOG, \*.EXE, \*.DLL, \*....	Include all dates	None
Invalid Characters	Any	<<.*[.]$>>	None Selected	Include all dates	None
Locking Only	Locking Only	None Selected	\*	Include all dates	None
MacOS Exclusions	Any	*\__MACOSX, *\\.Tempora...	None Selected	Include all dates	None
PEER_Dassault Systems CA...	Any	<<\\.([0-9a-z]+)>>, *.BCK	None Selected	Include all dates	None
Synchronizing Only	Synchronizati...	None Selected	\*	Include all dates	None
User Profile Exclusions	Any	*\\AppData\\Local\\Google\\...	None Selected	Include all dates	None

3. Click **Create**.

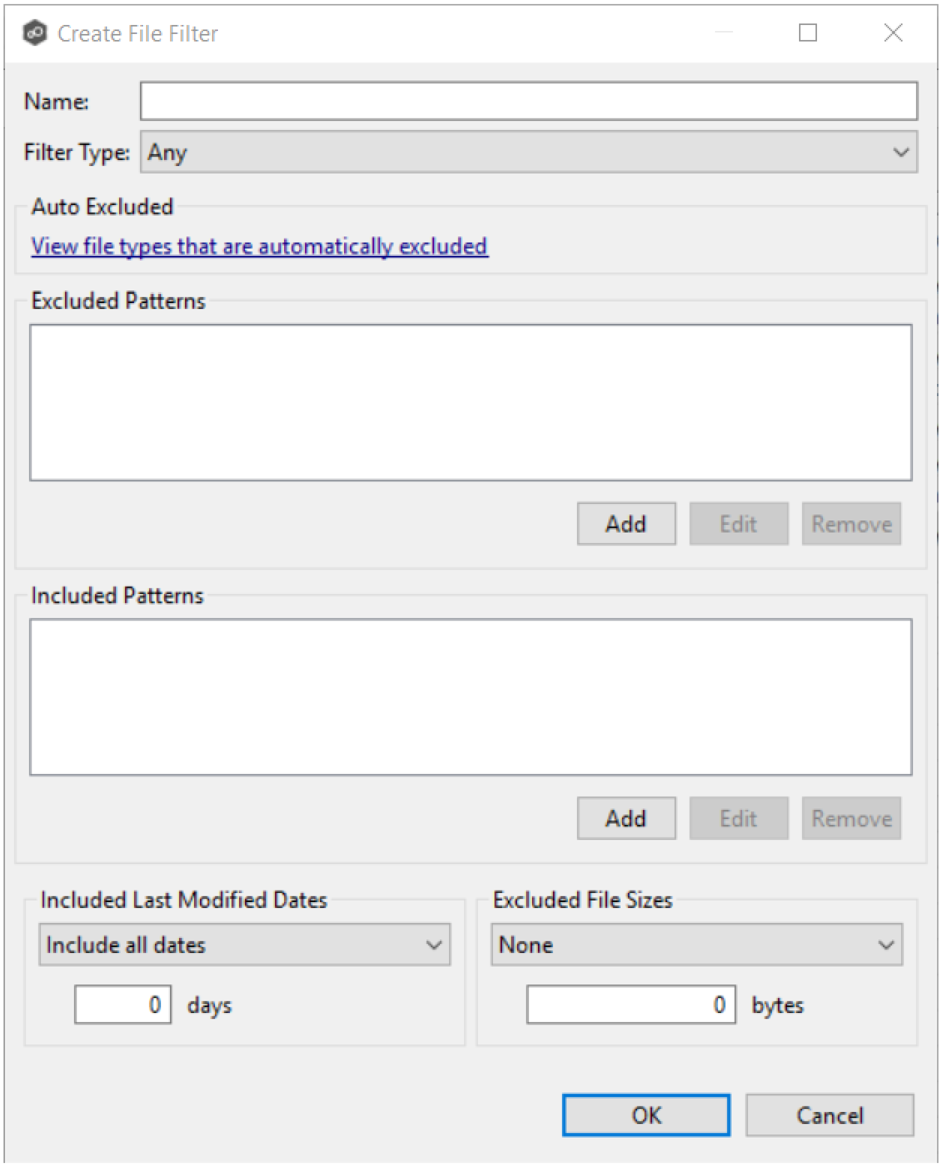

4. Enter a name for the filter.

5. Select the filter type.

6. (Optional) Enter a filter pattern in **Excluded Patterns**:

   a. Click **Add**.

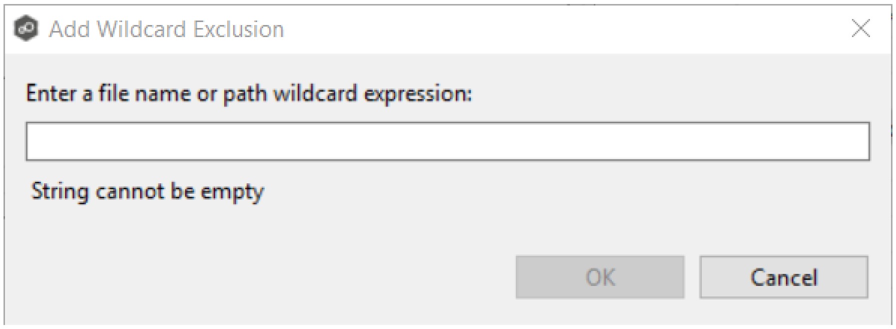

Add Wildcard Exclusion

Enter a file name or path wildcard expression:

String cannot be empty

b. Enter a filter pattern.

c. Click **OK**.

d. Repeat to add more filter patterns.

7. (Optional) Enter a filter pattern in **Included Patterns**..

a. Click **Add**.

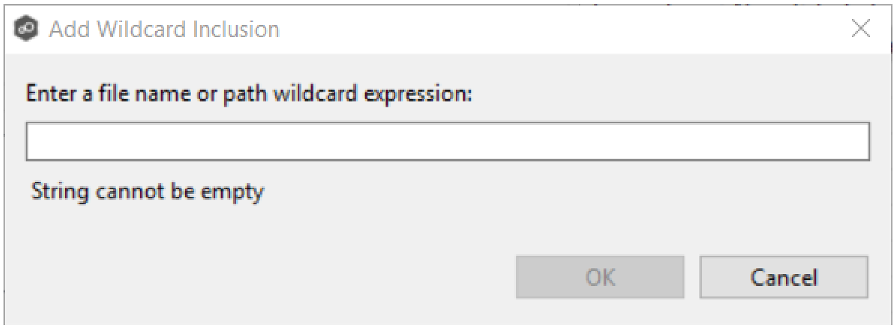

Add Wildcard Inclusion

Enter a file name or path wildcard expression:

String cannot be empty

b. Enter a string.

c. Click **OK**.

d. Repeat to add more filter patterns.

8. (Optional) Select **Included Last Modified Dates**.

9. (Optional) Select **Excluded File Sizes**.

Note that you cannot create a filter that uses **Excluded File Sizes** and filter patterns.

10. Click **OK** or **Apply**.

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

File Filter Example

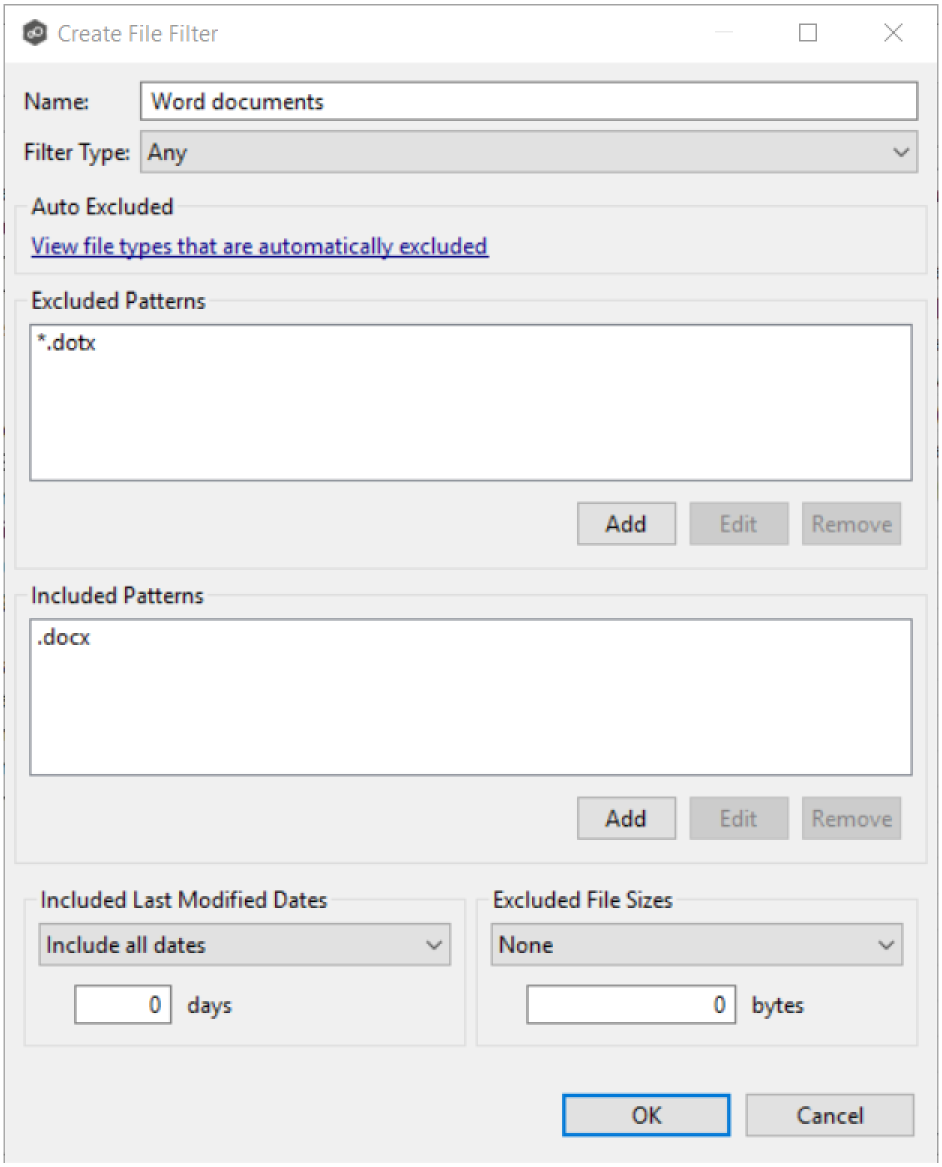

Create File Filter

Name: Word documents

Filter Type: Any

Auto Excluded

View file types that are automatically excluded

Excluded Patterns

*.dotx

Included Patterns

.docx

Included Last Modified Dates: Include all dates

0 days

Excluded File Sizes: None

0 bytes

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 are configured on a global level for all File Collaboration jobs and are critical for certain file types so that the file collaboration solution is able to correctly read any part of these files, ensuring any managed database type files are synchronized in a consistent and usable state.

To modify the locking settings:

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

2. Expand **File Collab, Sync, and Locking** in the navigation tree, and then select **Locking**.

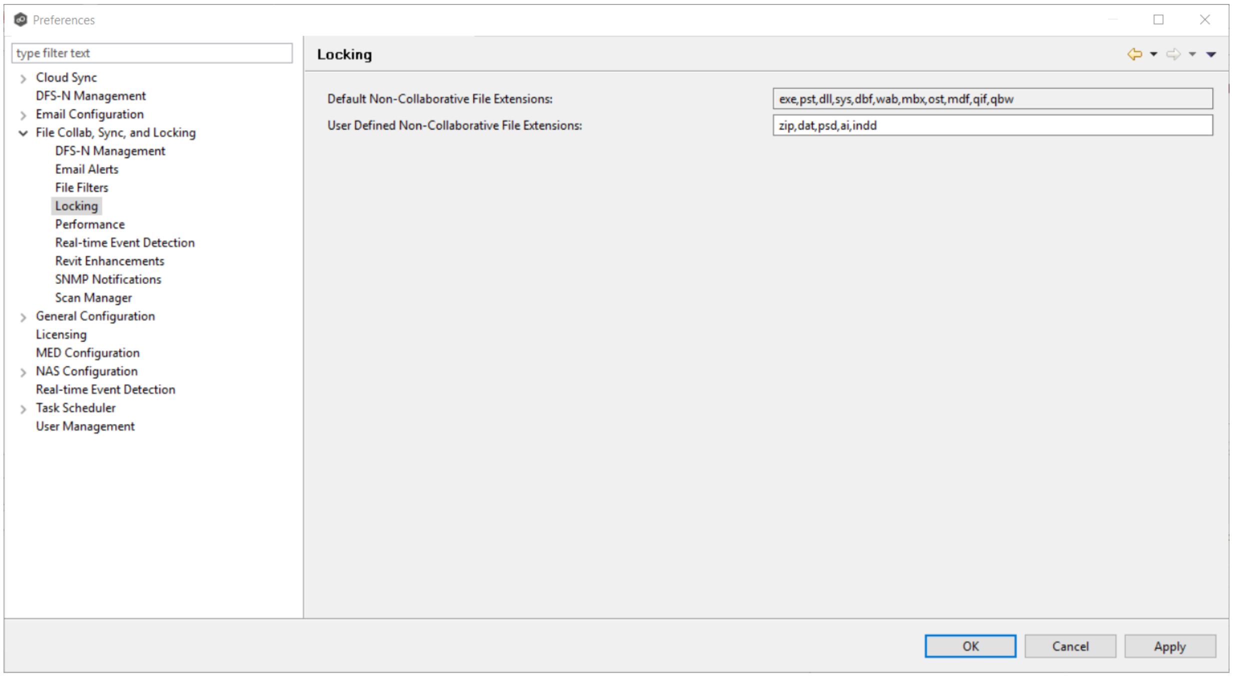

Default Non-Collaborative File Extensions: exe,pst,dll,sys,dbf,wab,mbx,ost,mdf,qif,qbw

User Defined Non-Collaborative File Extensions: zip,dat,psd,ai,indd

3. Modify the settings as needed.

Default Non-Collaborative File Extensions	The default, non-editable, comma separated list of file extensions of non-collaborative file types (e.g. database files, etc.). Write access to source files of these types will be denied while the files are being synchronized.
User Defined Non-Collaborative File Extensions	An editable, comma separated list of file extensions of non-collaborative file types (e.g. database files). Write access to source files of these types will be denied while the files are being synchronized.

4. Click **OK** or **Apply**.

Performance

To customize the File Collaboration performance settings:

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

2. Expand **File Collab, Sync, and Locking** in the navigation tree, and then select **Performance**.

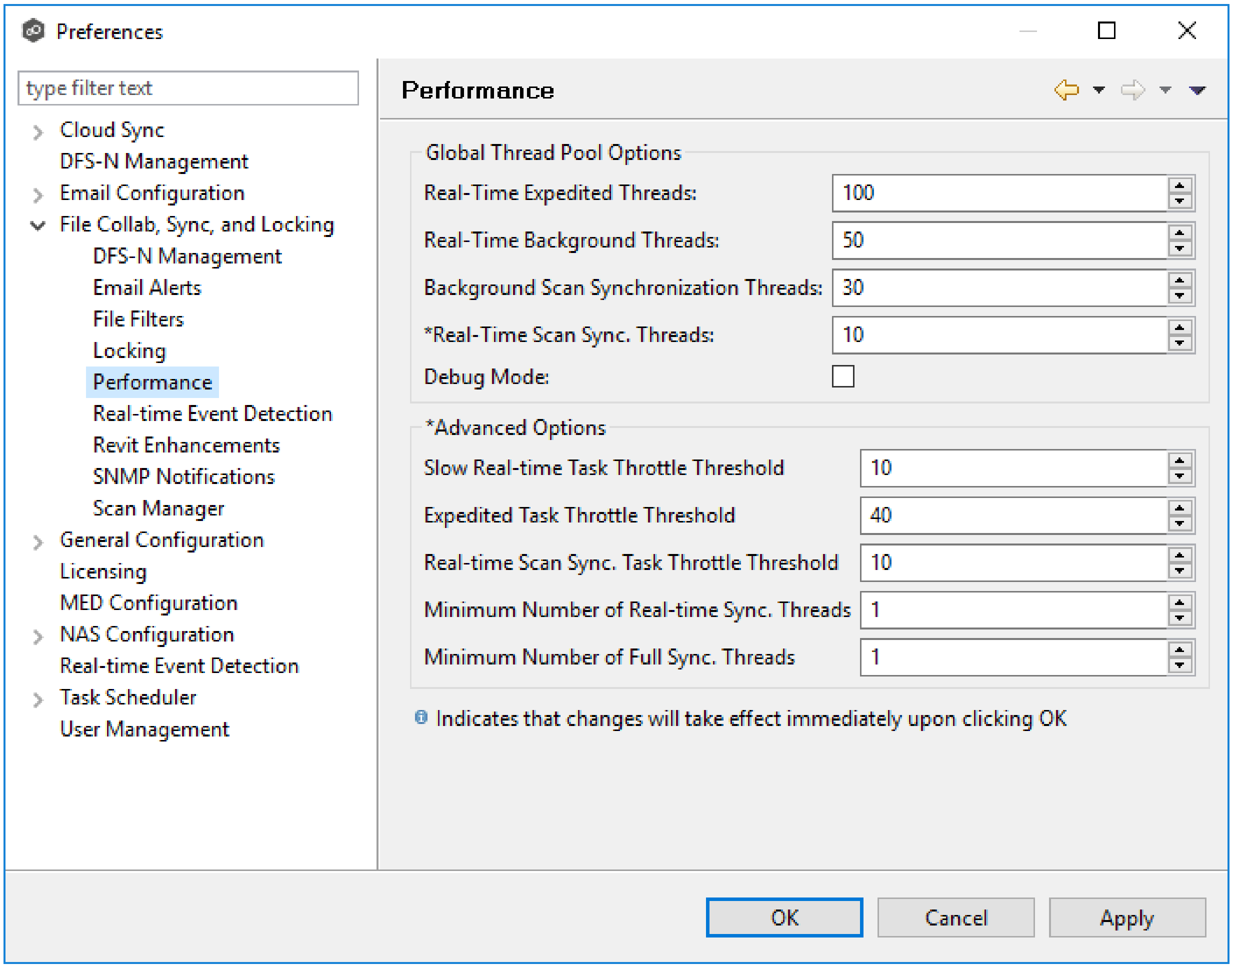

Performance

Global Thread Pool Options

Real-Time Expedited Threads: 100

Real-Time Background Threads: 50

Background Scan Synchronization Threads: 30

*Real-Time Scan Sync. Threads: 10

Debug Mode: □

*Advanced Options

Slow Real-time Task Throttle Threshold 10

Expedited Task Throttle Threshold 40

Real-time Scan Sync. Task Throttle Threshold 10

Minimum Number of Real-time Sync. Threads 1

Minimum Number of Full Sync. Threads 1

Indicates that changes will take effect immediately upon clicking OK

3. Modify the settings as needed. Do NOT modify the settings under **Advanced Options** unless directed by Peer Software Support.

Real-Time Expedited Threads	The threads that control file locking and renames.
Real-Time Background Threads	The threads that control the replication of file modifications.
Background Scan Synchronization Threads	The threads that scan directories for differences.
Real-Time Scan Sync. Threads	The threads that handle scans of newly added directories.
Debug Mode	Turns on debug for the various types of threads.

4. Click **OK** or **Apply**.

Real-time Event Detection

To modify the File Collaboration real-time detection settings:

1. Select **Preferences** from the Window menu.
2. Expand **File Collab, Sync, and Locking** in the navigation tree, and then select **Real-time Event Detection**.

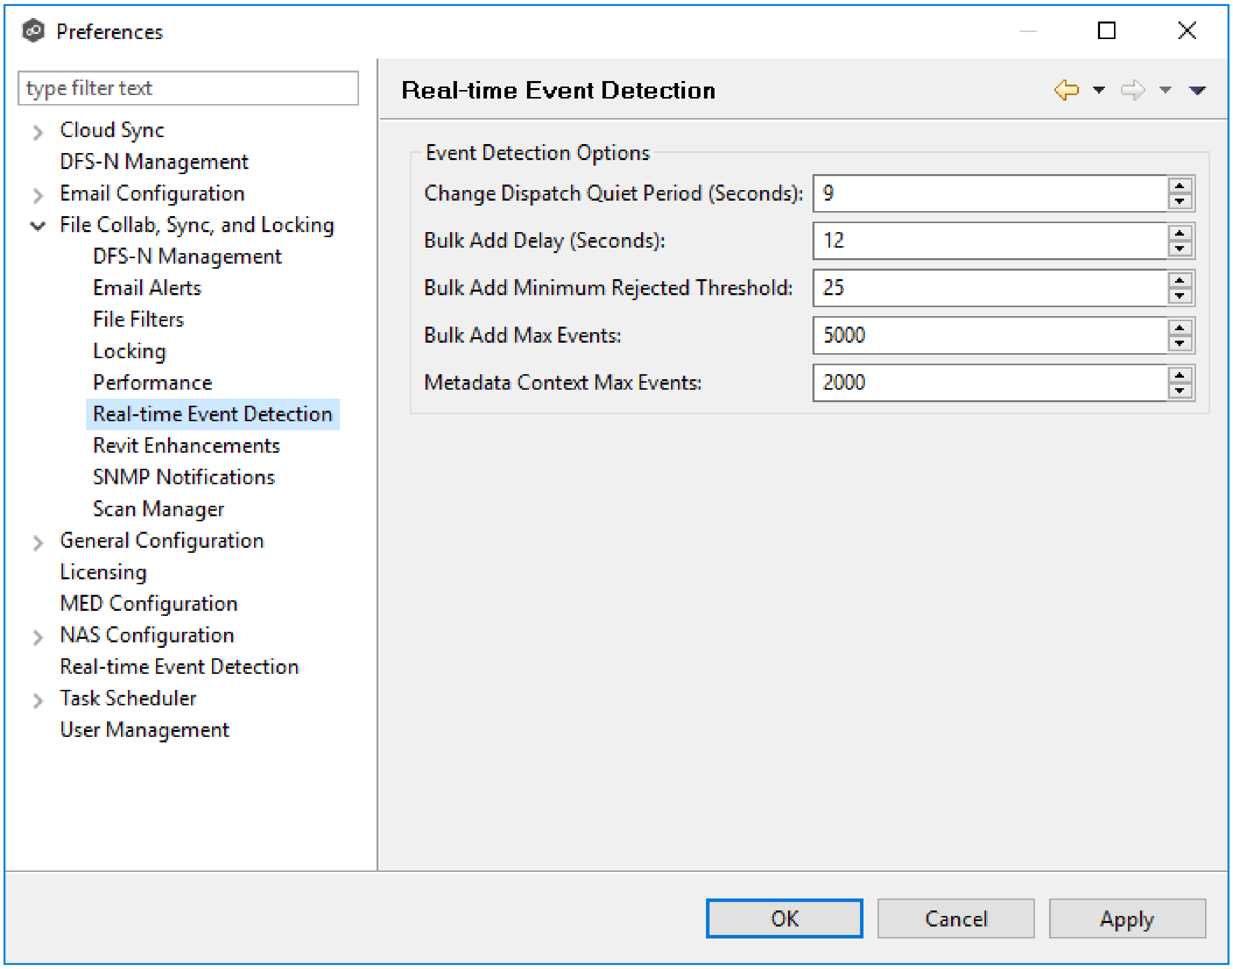

Real-time Event Detection

Event Detection Options

Change Dispatch Quiet Period (Seconds): 9

Bulk Add Delay (Seconds): 12

Bulk Add Minimum Rejected Threshold: 25

Bulk Add Max Events: 5000

Metadata Context Max Events: 2000

3. Modify the settings as needed.

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 de-prioritize mass copying or adding of files to a directory.
Bulk Add Minimu	The minimum number of file adds that must occur within the Bulk Add Delay for bulk add logic to be triggered.

m Rejected Threshold	
Bulk Add Max Events	The maximum number of file adds to lump together in one batch.
Metadata Context Max Events	The maximum number of metadata-related (security) events to batch at once. This reduces overhead when making mass changes to file and folder permissions.

4. Click **OK** 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 **File Collab, Sync, and Locking** in the navigation tree, and then select **Revit Enhancements**.

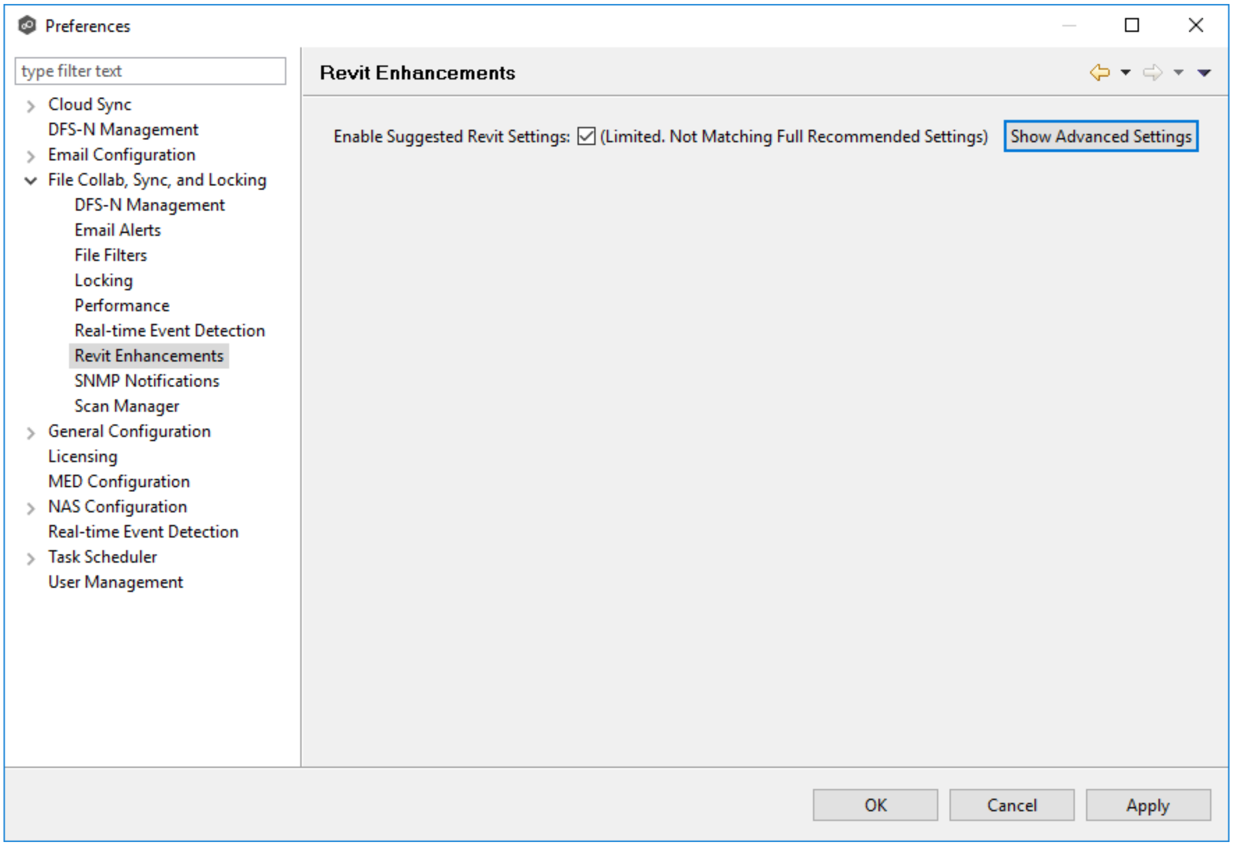

Enable Suggested Revit Settings: ☑ (Limited. Not Matching Full Recommended Settings)

3. Click **Show Advanced Settings**.

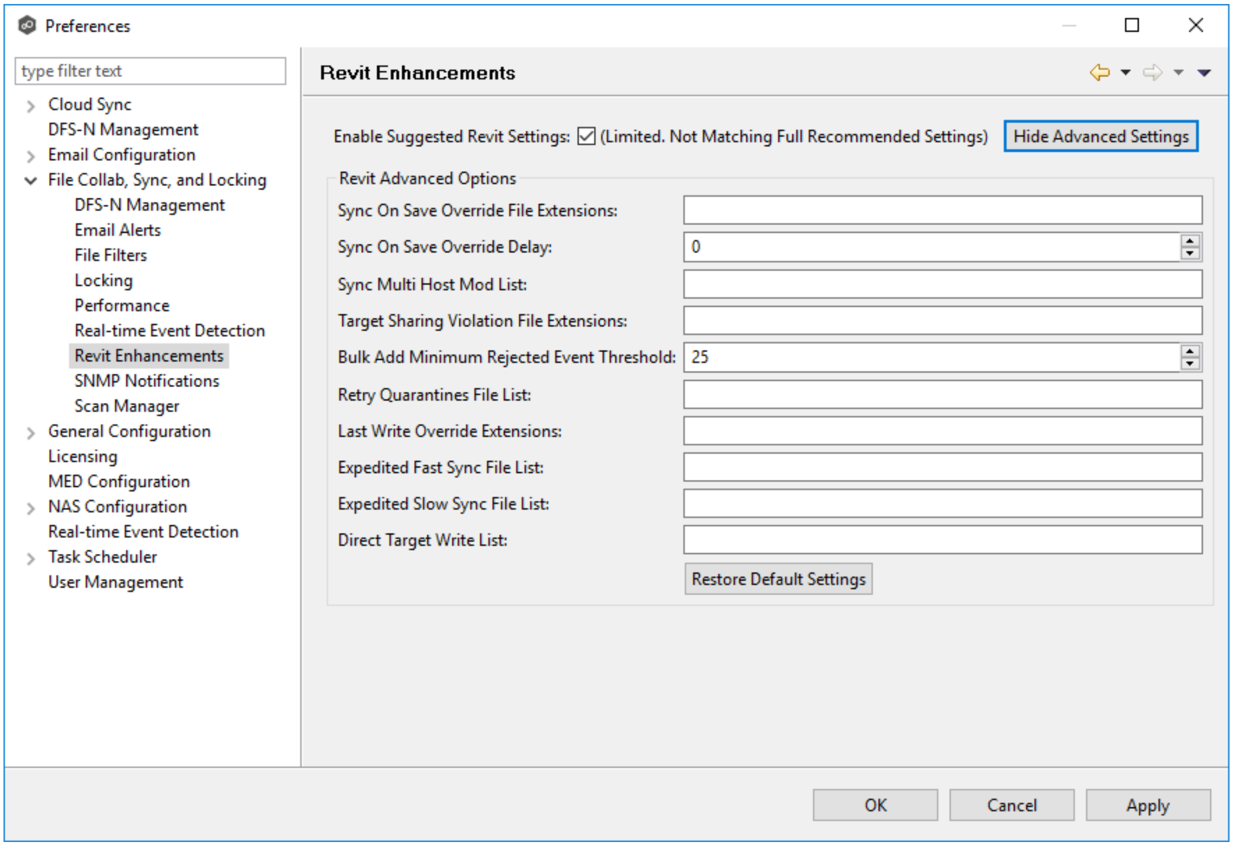

Revit Enhancements

Enable Suggested Revit Settings: ☑ (Limited. Not Matching Full Recommended Settings)

Revit Advanced Options

Sync On Save Override File Extensions:

Sync On Save Override Delay: 0

Sync Multi Host Mod List:

Target Sharing Violation File Extensions:

Bulk Add Minimum Rejected Event Threshold: 25

Retry Quarantines File List:

Last Write Override Extensions:

Expedited Fast Sync File List:

Expedited Slow Sync File List:

Direct Target Write List:

4. Modify the settings as needed.

Sync On Save Override File Extension	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.

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 **OK** or **Apply**.

SNMP Notifications

The Peer Management Center provides basic support for SNMP messaging. SNMP notifications are set through the concept of SNMP Notification, where a single notification (consisting of a unique name, a selection of notification types along with a trap prefix and destination) can be applied to multiple File Collaboration jobs without requiring repeat entry for each job. When an SNMP notification is applied to a job, a SNMP trap will be sent to the destination IP address or hostname whenever a selected notification type is triggered by the job.

When you create a job, you can select an existing SNMP configuration 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 when it is applied to a job.

To create a SMNP notification:

1. From the **Window** menu, select **Preferences**.

2. Expand **File Collab, Sync, and Locking** in the navigation tree, and then select **SNMP Notifications**.

   The existing SNMP notifications are listed in the **SNMP Notifications** table.

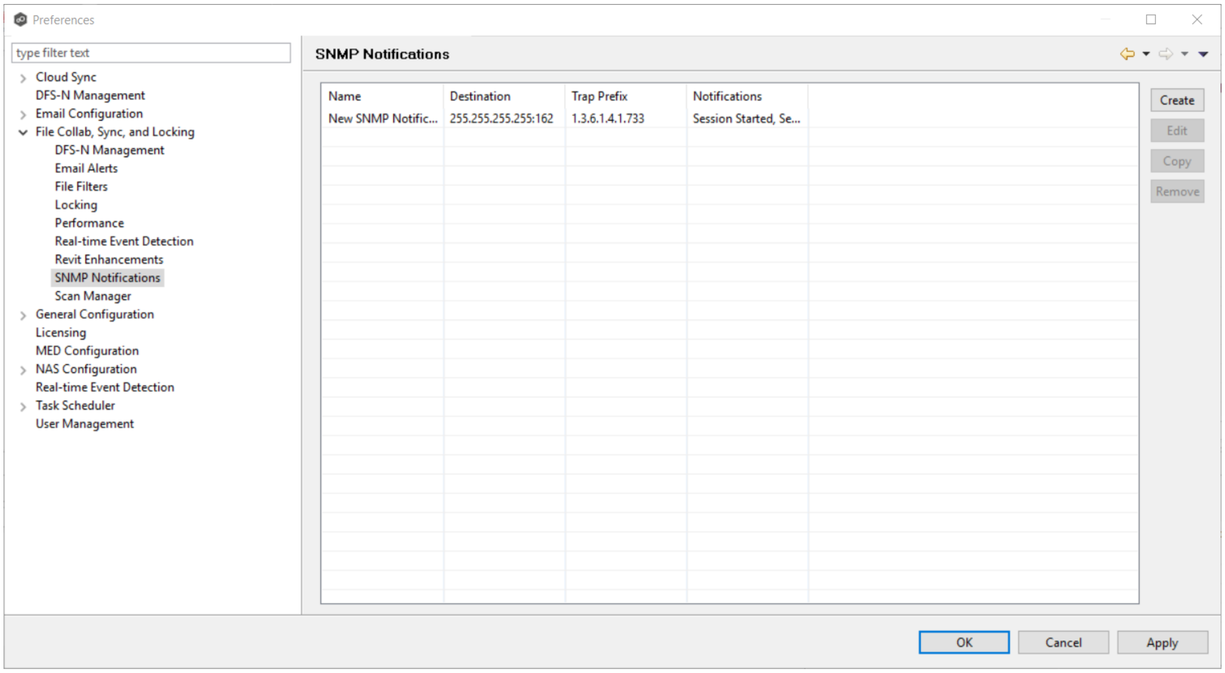

SNMP Notifications

Name	Destination	Trap Prefix	Notifications
New SNMP Notific...	255.255.255.255:162	1.3.6.1.4.1.733	Session Started, Se...

3. Click the **Create** button.

   The **Add SNMP Notification** dialog appears. Within this dialog, you can select specific triggers on which an SNMP trap will be generated, configure the source IP address over which the trap will be sent, set the destination host name, IP address, or broadcast address, set the prefix that is attached to every message (helping to identify messages coming from specific instances of the Peer Management Center or jobs across a network), and test the aforementioned settings.

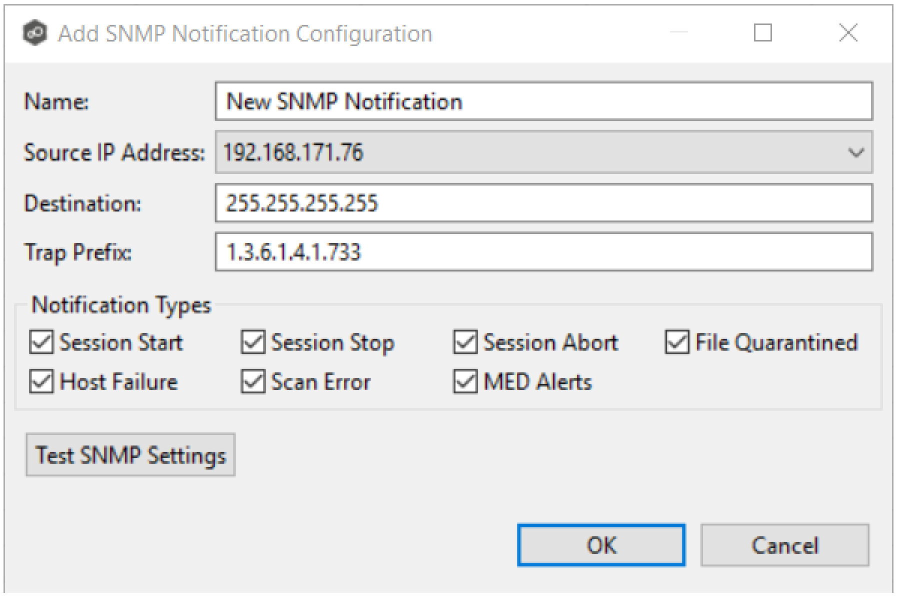

Name: New SNMP Notification

Source IP Address: 192.168.171.76

Destination: 255.255.255.255

Trap Prefix: 1.3.6.1.4.1.733

Notification Types

☑ Session Start ☑ Session Stop ☑ Session Abort ☑ File Quarantined

☑ Host Failure ☑ Scan Error ☑ MED Alerts

4. Enter the values.

The notification types are:

Session Start	Enables sending a notification when a session is started.
Session Stop	Enables sending a notification when a session is stopped.
Session Abort	Enables sending a notification when a session is aborted because of lack of quorum due to a failed host(s).
File Quarantined	Enables sending a notification when a file is marked as quarantined because a file conflict was not able to be resolved.
Host Timeout	Enables sending a notification when a host timeout occurs and the host is taken out of session.
Scan Error	Enables sending a notification when an error occurs during the initial synchronization process.
MED Alerts	Enables sending a notification when Peer MED detects potentially malicious activity. For more information, see MED Configuration.

5. Click **Test SNMP Settings**, and then click **OK** in the **Test Connection** dialog.

6. Click **OK** or **Apply**.

   The new notification is listed in the **SNMP Notifications** table and can now be applied to jobs.

Scan Manager

A number of options are available to tune the way scans are performed for File Collaboration, File Locking, File Replication, and File Synchronization jobs.

To modify the Scan Manager settings for File Collaboration, File Locking, File Replication, and File Synchronization jobs:

1. From the **Window** menu, select **Preferences**.

2. Expand **File Collab, Sync, and Locking** in the navigation tree, and then select **Scan Manager**.

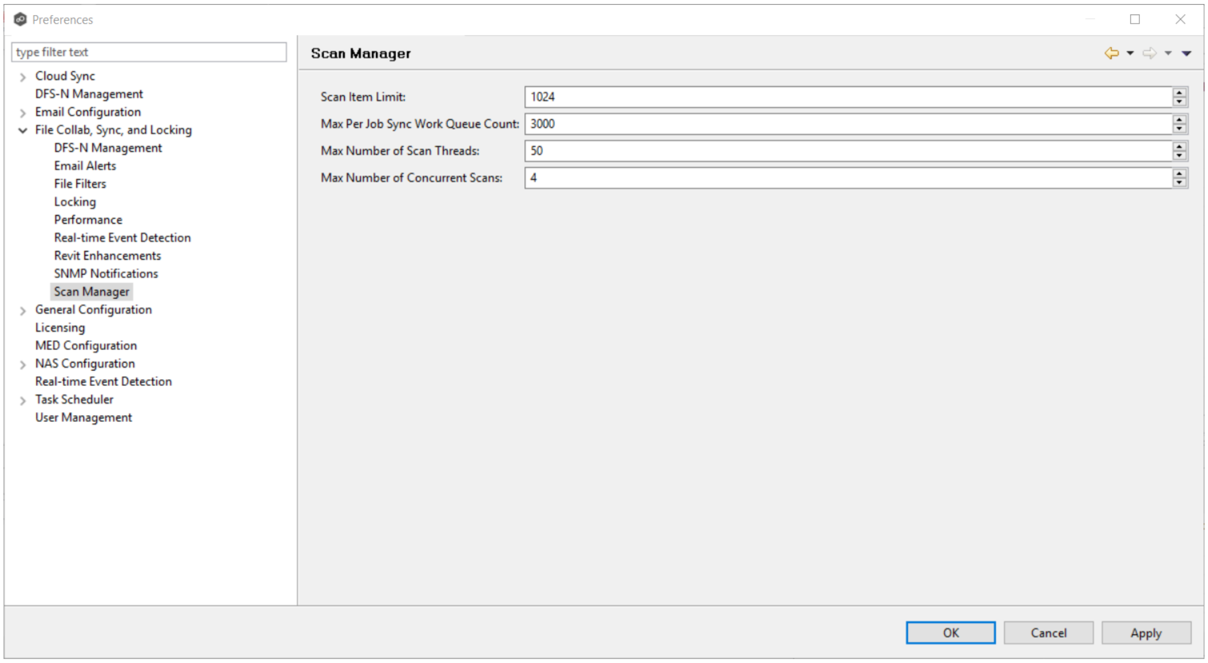

Scan Manager

Scan Item Limit: 1024

Max Per Job Sync Work Queue Count: 3000

Max Number of Scan Threads: 50

Max Number of Concurrent Scans: 4

3. Modify the settings as needed.

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 sessions.
Max Sync Work Queue Count	The maximum number of pending file transfers (as a result of the initial scan) that are queued in memory before pausing the current scan. This value only has an effect on sessions that require a massive amount of initial synchronization.
Max Number of Scan Threads	The maximum number of threads that can be created for use when scanning 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.

4. Click **OK** or **Apply**.

Licensing

Peer Management Center is licensed by the number of unique participating hosts 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**.

   The **Preferences** dialog appears.

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

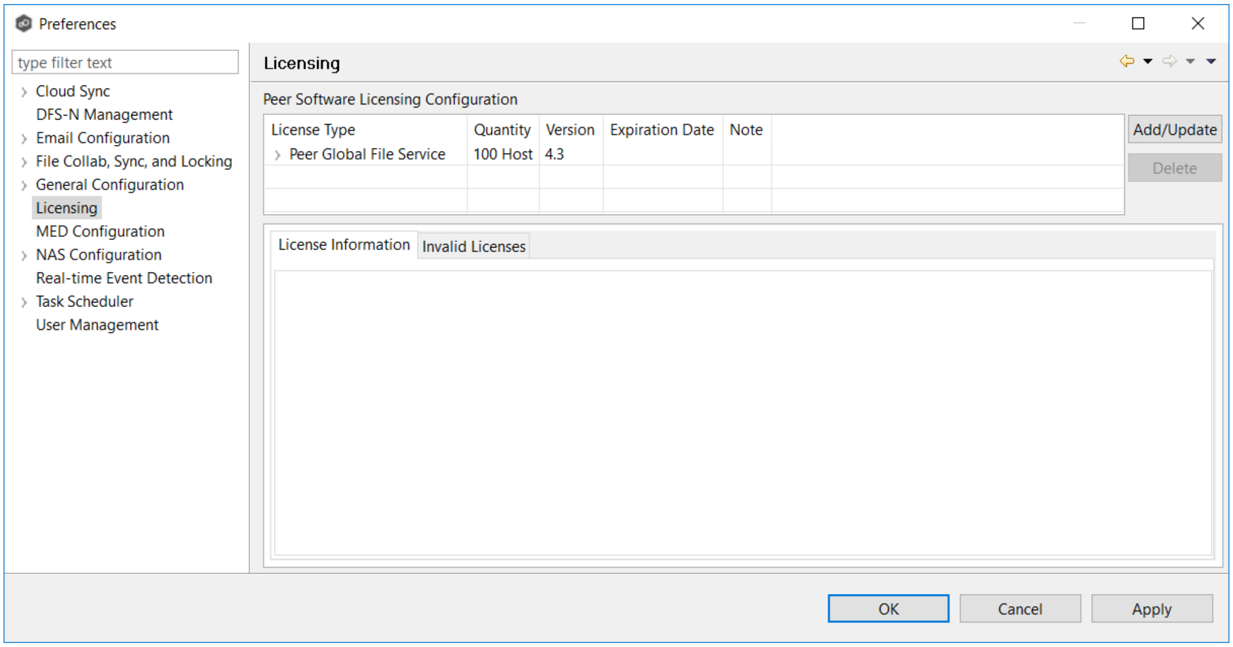

License Type	Quantity	Version	Expiration Date	Note
Peer Global File Service	100 Host	4.3		

3. Click the **Add/Update** button to browse for and install the license file.

4. Select the license file, and then click **Open**.

   If a license already exists for the same type, then the existing license will be ovewritten with the new license. After successful installation of the license file, the license is displayed in the **License Configuration** table, along with licensed quantity and an expiration date (if applicable). You can now create, configure, and run jobs using the new license type.

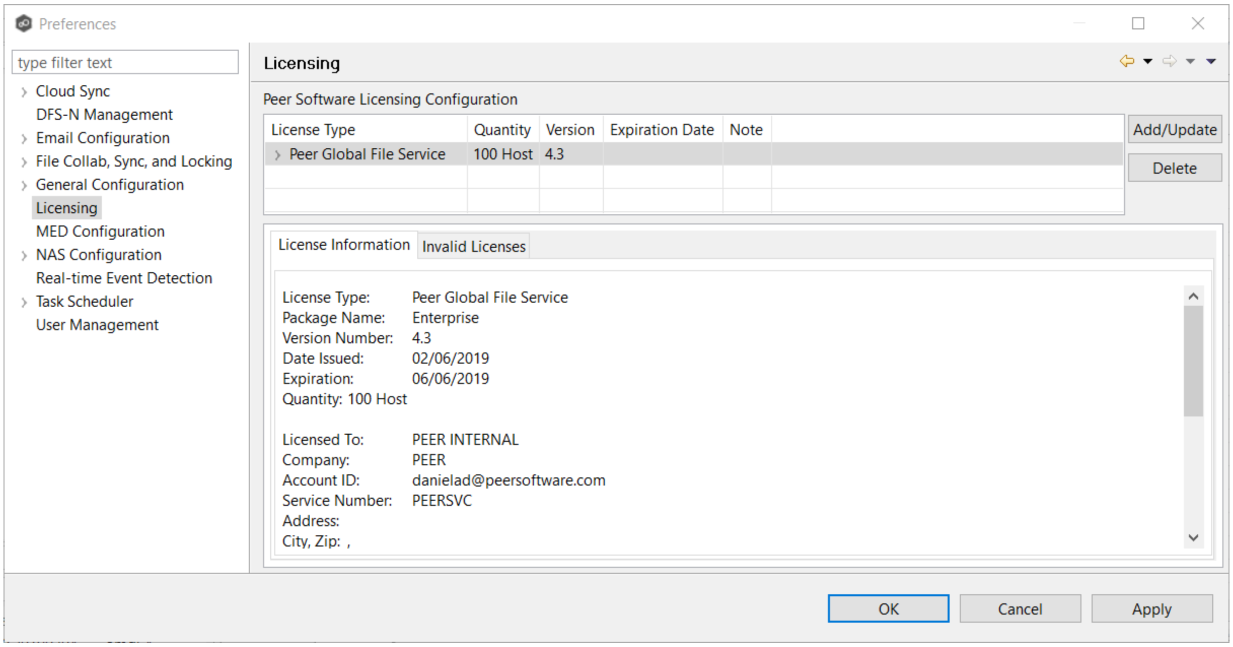

Peer Software Licensing Configuration

License Type	Quantity	Version	Expiration Date	Note
Peer Global File Service	100 Host	4.3		

License Type: Peer Global File Service
Package Name: Enterprise
Version Number: 4.3
Date Issued: 02/06/2019
Expiration: 06/06/2019
Quantity: 100 Host

Licensed To: PEER INTERNAL
Company: PEER
Account ID: danielad@peersoftware.com
Service Number: PEERSVC
Address:
City, Zip: ,

5. Click **OK** or **Apply**.

Deleting a License File

Click the **Delete** button to permanently remove all licenses of the selected type (both valid and invalid licenses). Any job types enabled by that license will be hidden from the Peer Management Center.

Expired licenses will be listed in the **Invalid Licenses** tab.

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 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 view and modify these settings,

1. From the **Window** menu, select **Preferences**.
2. Select **MED Configuration** in the navigation tree.

   The following page is displayed.

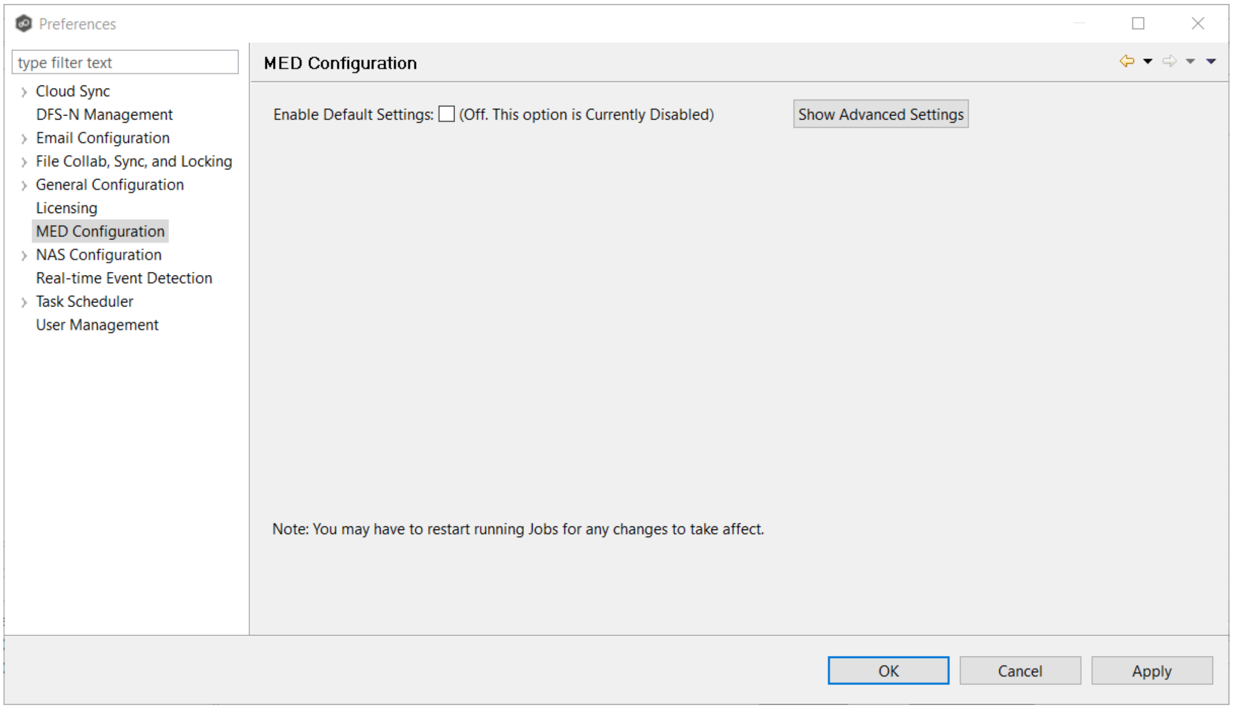

Preferences

MED Configuration

Enable Default Settings: □ (Off. This option is Currently Disabled)

Note: You may have to restart running Jobs for any changes to take affect.

3. Select the **Enable Default Settings** or click **Show Advanced Settings**.

   If you selected Show Advanced Settings, the following is displayed.

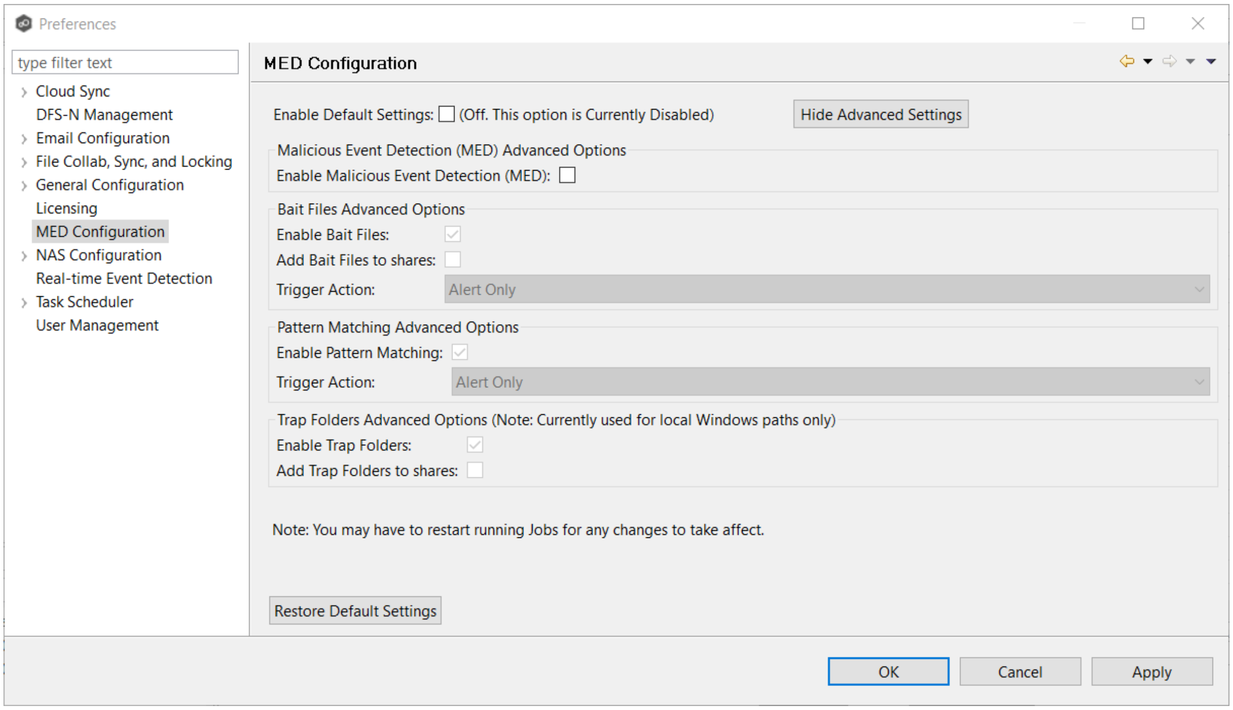

MED Configuration

Enable Default Settings: □ (Off. This option is Currently Disabled)

Malicious Event Detection (MED) Advanced Options

Enable Malicious Event Detection (MED): □

Bait Files Advanced Options

Enable Bait Files: ☑

Add Bait Files to shares: □

Trigger Action: Alert Only

Pattern Matching Advanced Options

Enable Pattern Matching: ☑

Trigger Action: Alert Only

Trap Folders Advanced Options (Note: Currently used for local Windows paths only)

Enable Trap Folders: ☑

Add Trap Folders to shares: □

Note: You may have to restart running Jobs for any changes to take affect.

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:

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	The master on/off switch for MED. If unchecked, all MED mechanisms will be disabled.

Detection (MED)	
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:

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 configured watch directory. To see the watch directory for a job, review Host Participants and Directories.
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:

Alert Only	Triggers an alert in the Peer Management Center. If SMTP email alerts are configured for MED Alerts and enabled for a job, an email will also be sent. For details on SMTP email alerts, see Global Email options. If SNMP traps are configured for MED Alerts and enabled for a job, an SNMP trap will also be sent. For more details on SNMP, see Global SNMP options.
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 the Peer Management Center is as follows:

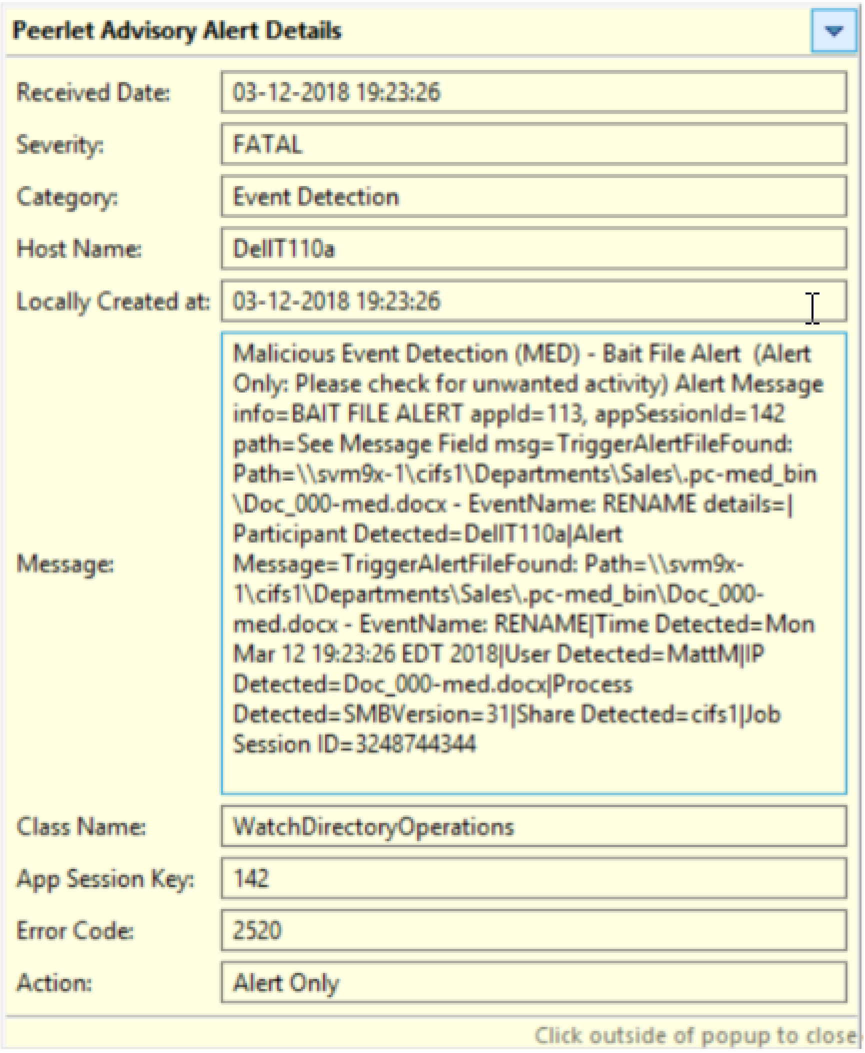

Peerlet Advisory Alert Details

Received Date:	03-12-2018 19:23:26
Severity:	FATAL
Category:	Event Detection
Host Name:	DellT110a
Locally Created at:	03-12-2018 19:23:26
Message:	Malicious Event Detection (MED) - Bait File Alert (Alert Only: Please check for unwanted activity) Alert Message info=BAIT FILE ALERT appId=113, appSessionId=142 path=See Message Field msg=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\.pc-med_bin \Doc_000-med.docx - EventName: RENAME details=|Participant Detected=DellT110a|Alert Message=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\.pc-med_bin\Doc_000-med.docx - EventName: RENAME|Time Detected=Mon Mar 12 19:23:26 EDT 2018|User Detected=MattM|IP Detected=Doc_000-med.docx|Process Detected=SMBVersion=31|Share Detected=cifs1|Job Session ID=3248744344
Class Name:	WatchDirectoryOperations
App Session Key:	142
Error Code:	2520
Action:	Alert Only

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:

Enable Trap Folders	Enables/disables the creation and monitoring of trap folders.
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 detailed information about configuring your NAS:

- EMC Configuration
- NetApp Configuration
- Nutanix Configuration

EMC Configuration

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more EMC storage devices within the majority of available job types. These EMC devices can be running Isilon, Unity, or VNX.

For more detailed information about prerequisites and configuration, see EMC Prerequisites and Configuration.

Prerequisites

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

- For EMC Isilon environments: https://kb.peersoftware.com/tb/emc-isilon-prerequisites
- For EMC Unity environments: https://kb.peersoftware.com/tb/emc-unity-prerequisites
- For EMC VNX/Celerra environments: https://kb.peersoftware.com/tb/emc-vnx-celerra-prerequisites

CEE Server Configuration Guide

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

- EMC Isilon CEE Server Configuration Guide: https://kb.peersoftware.com/tb/emc-isilon-configuration-guide
- EMC Unity CEE Server Configuration Guide: https://kb.peersoftware.com/tb/emc-unity-configuration-guide
- EMC VNX/Celerra CEE Server Configuration Guide: https://kb.peersoftware.com/tb/emc-vnx-celerra-configuration-guide

Configuration

The creation of a new jobs in the Peer Management Center will automatically walk you through entering the most important settings for use in EMC environments. If you need to modify existing settings or tune advanced settings, you can do so by following these steps:

1. Right-click on a File Collaboration, File Synchronization, File Locking, or File Replication job and click **Edit Job(s)**.
2. Click on the **Participants** page on the left.
3. Select the Peer Agent that is managing the appropriate EMC storage device then click the **Edit Detector Settings** button.

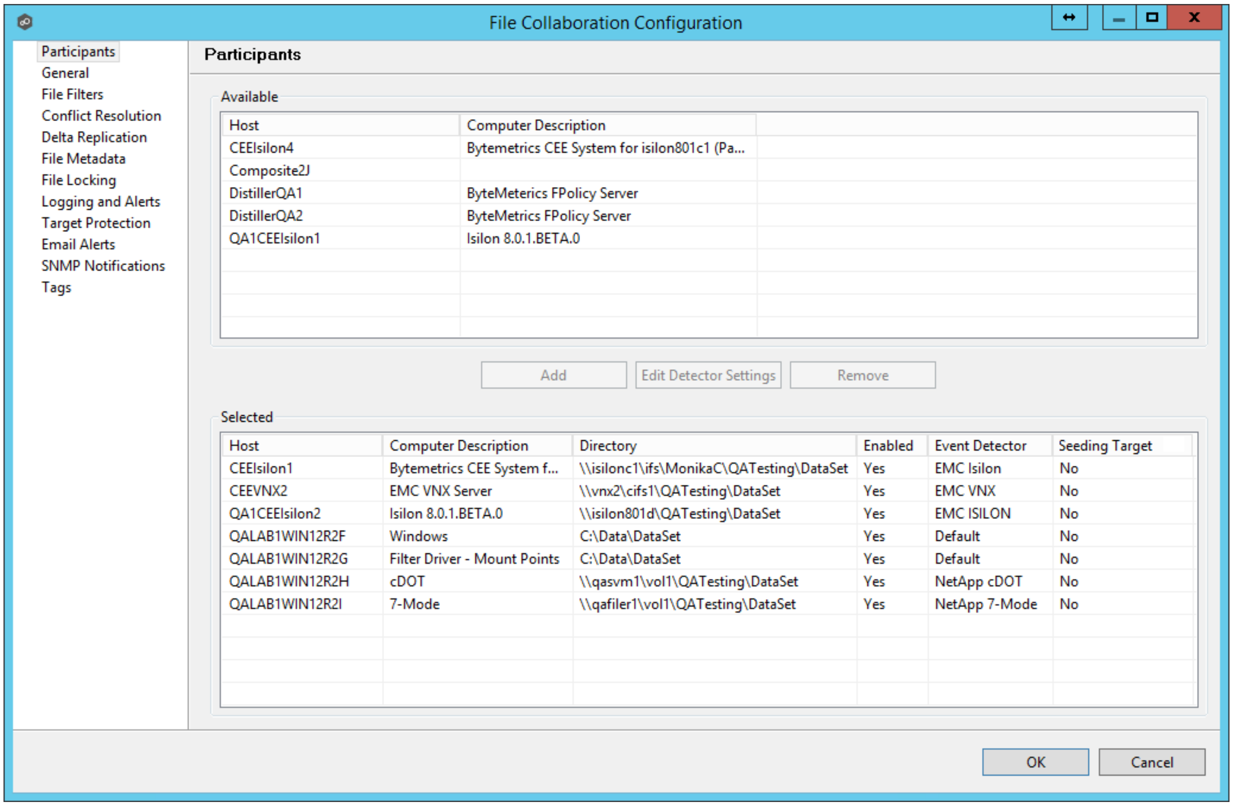

Participants

Available

Host	Computer Description
CEEIsilon4	Bytemetrics CEE System for isilon801c1 (Pa...
Composite2J	
DistillerQA1	ByteMeterics FPolicy Server
DistillerQA2	ByteMetrics FPolicy Server
QA1CEEIsilon1	Isilon 8.0.1.BETA.0

Selected

Host	Computer Description	Directory	Enabled	Event Detector	Seeding Target
CEEIsilon1	Bytemetrics CEE System f...	\\isilonc1\ifs\MonikaC\QATesting\DataSet	Yes	EMC Isilon	No
CEEVNX2	EMC VNX Server	\\vnx2\cifs1\QATesting\DataSet	Yes	EMC VNX	No
QA1CEEIsilon2	Isilon 8.0.1.BETA.0	\\isilon801d\QATesting\DataSet	Yes	EMC ISILON	No
QALAB1WIN12R2F	Windows	C:\Data\DataSet	Yes	Default	No
QALAB1WIN12R2G	Filter Driver - Mount Points	C:\Data\DataSet	Yes	Default	No
QALAB1WIN12R2H	cDOT	\\qasvm1\vol1\QATesting\DataSet	Yes	NetApp cDOT	No
QALAB1WIN12R2I	7-Mode	\\qafiler1\vol1\QATesting\DataSet	Yes	NetApp 7-Mode	No

4. A configuration dialog will be displayed showing various settings used to tune how the Agent will communicate with its associated EMC storage device.

   With EMC Isilon devices, the following configuration options are available:

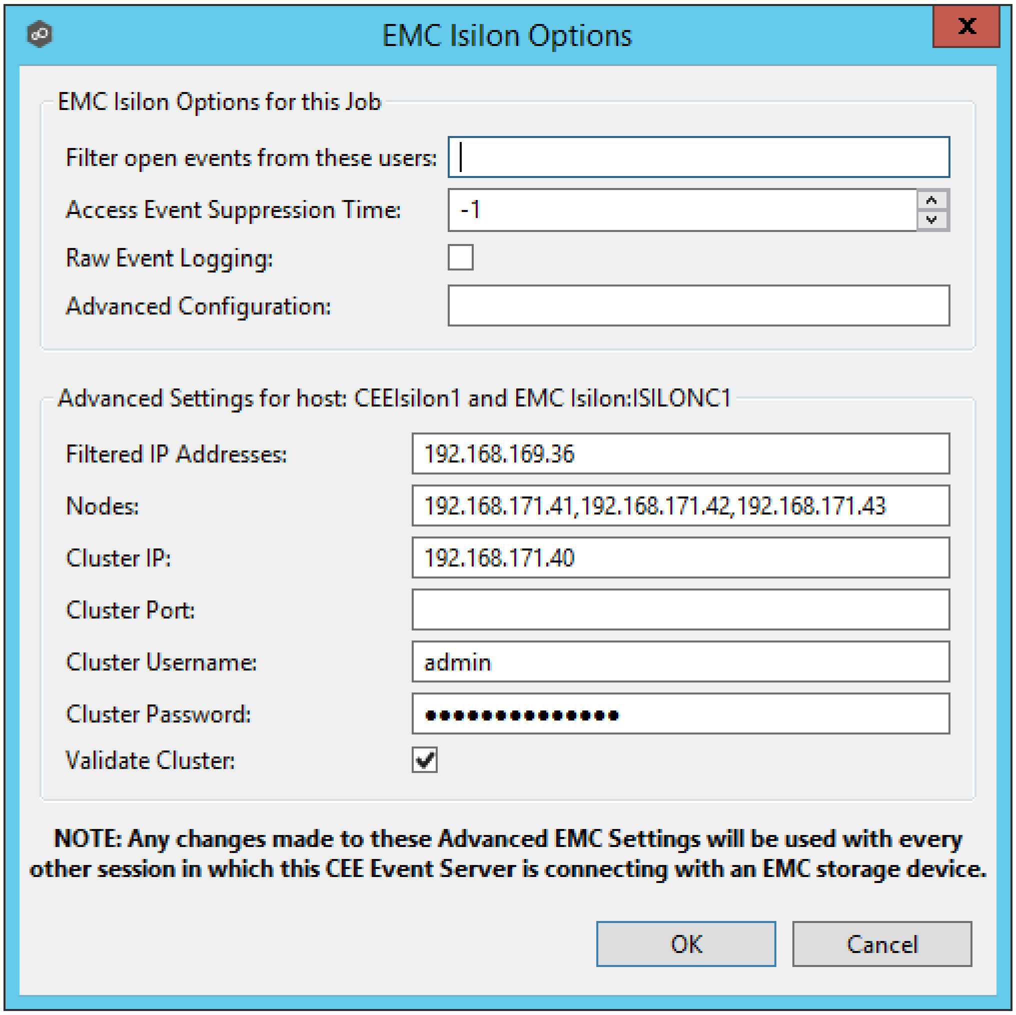

EMC Isilon Options for this Job

Filter open events from these users:

Access Event Suppression Time: -1

Raw Event Logging: □

Advanced Configuration:

Advanced Settings for host: CEEIsilon1 and EMC Isilon:ISILONC1

Filtered IP Addresses: 192.168.169.36

Nodes: 192.168.171.41,192.168.171.42,192.168.171.43

Cluster IP: 192.168.171.40

Cluster Port:

Cluster Username: admin

Cluster Password: ●●●●●●●●●●●●●●

Validate Cluster: ☑

NOTE: Any changes made to these Advanced EMC Settings will be used with every other session in which this CEE Event Server is connecting with an EMC storage device.

EMC Isilon Options

Filter open 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.
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 7 clients when mousing over files in Windows Explorer. The default value is -1, which will use a global set value. 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.
Raw Event Logging	Enables raw event logging for event detection debugging. Technical support may ask you to enable this feature if you are

	experiencing certain issues.
Advanced Configuration	Advanced settings for Event Detection and logging that will override the defaults. Technical support will provide you with a value to put in this field if you are experiencing certain issues.

Advanced Settings for EMC Isilon

Filtered IP Addresses	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	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.
Cluster IP	The cluster IP address of the Isilon system.
Custer Port	The cluster port number of the Isilon system. Default value is 8080.
Cluster Username	Username used to sign into the Isilon cluster.
Cluster Password	Password used to sign into the Isilon cluster.
Validate Cluster	If enabled, the Isilon cluster will be validated both on registration and periodically by a maintenance thread.

With EMC Unity devices, the following configuration options are available:

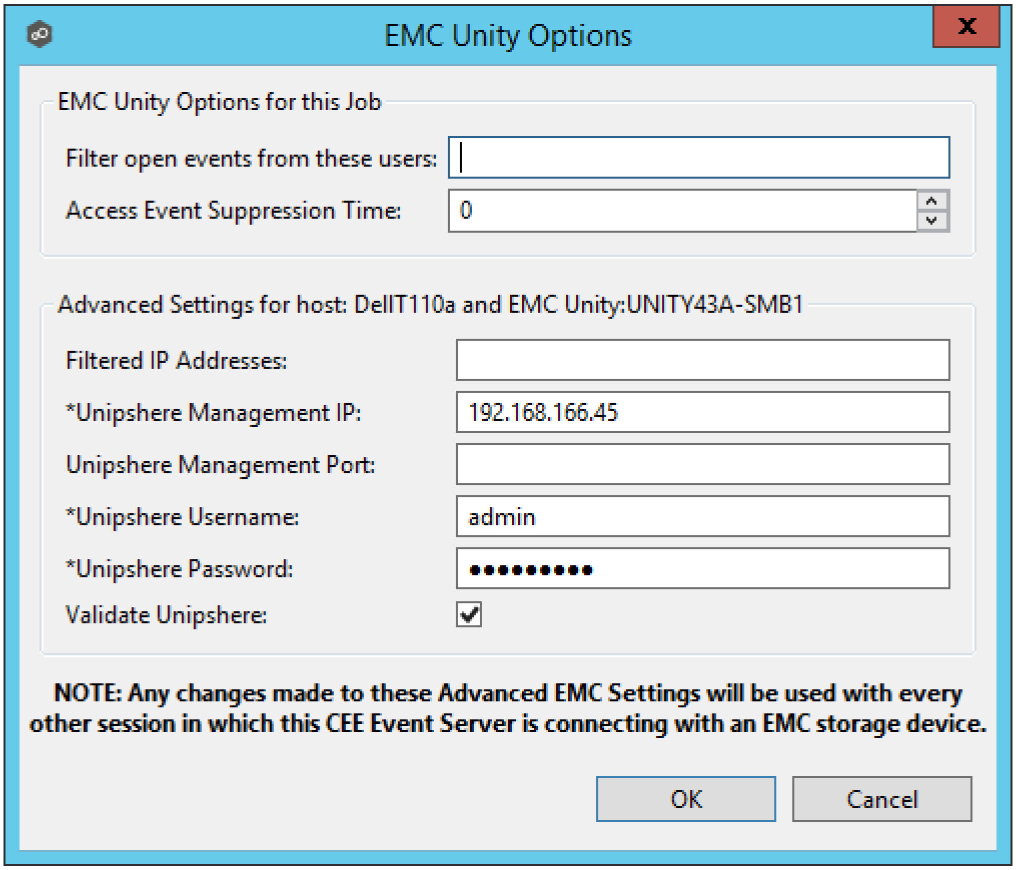

EMC Unity Options

EMC Unity Options for this Job

Filter open events from these users:

Access Event Suppression Time: 0

Advanced Settings for host: DellT110a and EMC Unity:UNITY43A-SMB1

Filtered IP Addresses:

*Unipshere Management IP: 192.168.166.45

Unipshere Management Port:

*Unipshere Username: admin

*Unipshere Password: ••••••••••

Validate Unipshere: ☑

NOTE: Any changes made to these Advanced EMC Settings will be used with every other session in which this CEE Event Server is connecting with an EMC storage device.

EMC Unity Options

Filter open 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.
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 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 that an open event will be delayed based on the load of the system.

Advanced Settings for EMC Unity

Filtered IP Addresses	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	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 connect to Unisphere.
Validate Unisphere	If enabled, Unisphere settings will be validated both on registration and periodically by a maintenance thread.

With EMC VNX devices, the following configuration options are available:

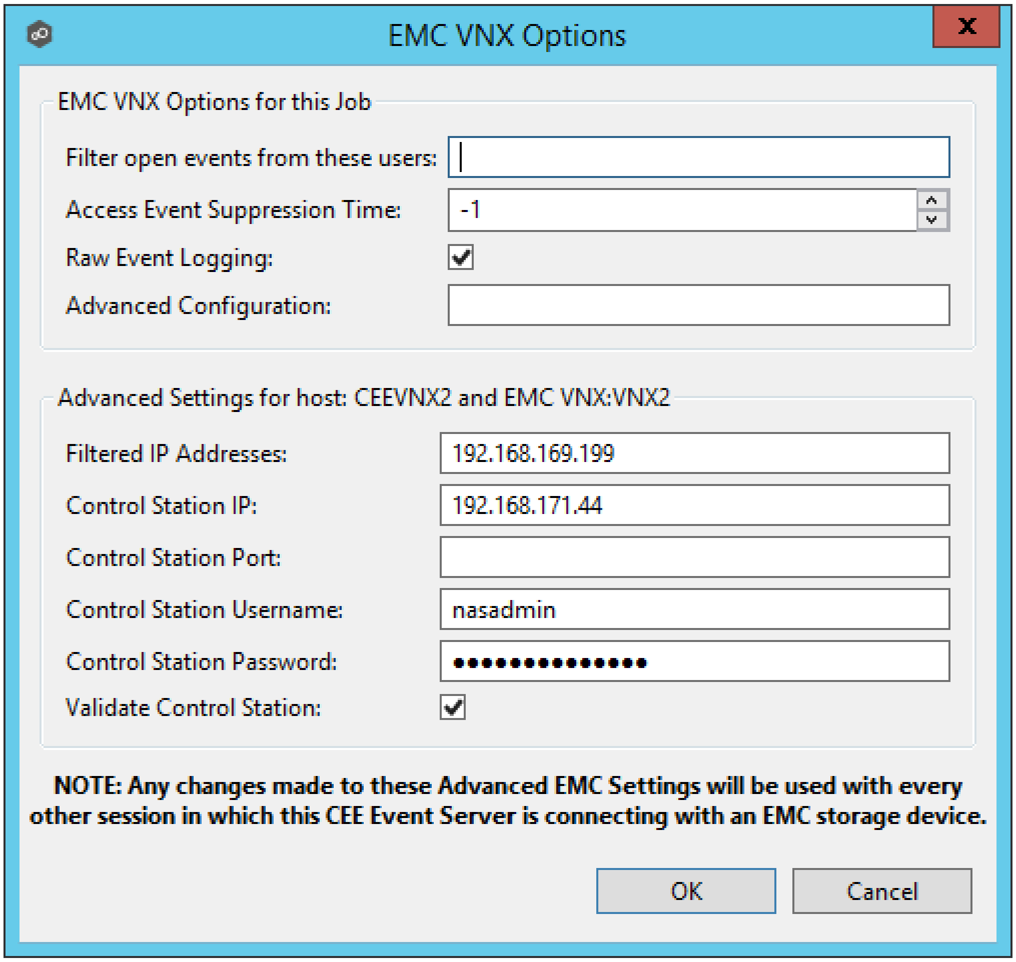

EMC VNX Options

EMC VNX Options for this Job

Filter open events from these users:

Access Event Suppression Time: -1

Raw Event Logging: ☑

Advanced Configuration:

Advanced Settings for host: CEEVNX2 and EMC VNX:VNX2

Filtered IP Addresses: 192.168.169.199

Control Station IP: 192.168.171.44

Control Station Port:

Control Station Username: nasadmin

Control Station Password: ••••••••••••••

Validate Control Station: ☑

NOTE: Any changes made to these Advanced EMC Settings will be used with every other session in which this CEE Event Server is connecting with an EMC storage device.

EMC VNX Options

Filter open 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.
Access Event Suppression Time	Represents the 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 global set value. 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.
Raw Event Logging	Enables raw event logging for event detection debugging. Technical support may ask you to enable this feature if you are experiencing certain issues.

Advanced Configuration	Advanced settings for Event Detection and logging that will override the defaults. Technical support will provide you with a value to put in this field if you are experiencing certain issues.

Advanced Settings for EMC VNX

Filtered IP Addresses	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 system.
Control Station Port	The Control Station Port number of the VNX system. Default value is 443.
Control Station Username	Username used to sign into the VNX Control Station.
Control Station Password	Password used to sign into the VNX Control Station.
Validate Control Station	If enabled, the VNX Control Station will be validated both on registration and periodically by a maintenance thread.

5. After making the necessary changes, click **OK** twice to save them. If the selected job is already running, you will need to restart it for the changes to take effect.

NetApp Configuration

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more NetApp storage devices within the majority of 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 more detailed information about prerequisites and configuration, see NetApp Prerequisites and Configuration.

Prerequisites

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

- For NetApp 7-Mode environments: https://kb.peersoftware.com/tb/netapp-7-mode-prerequisites
- For NetApp cDOT environments: https://kb.peersoftware.com/tb/netapp-cdot-prerequisites

Configuration

The creation of a new jobs in the Peer Management Center will automatically walk you through entering the most important settings for use in NetApp environments. If you need to modify existing settings or tune advanced settings, you can do so by following these steps:

1. Right-click on a File Collaboration, File Synchronization, File Locking, or File Replication job and click **Edit Job(s)**.
2. Click on the **Participants** page on the left.
3. Select the Peer Agent that is managing the appropriate NetApp storage device then click the **Edit Detector Settings** button.

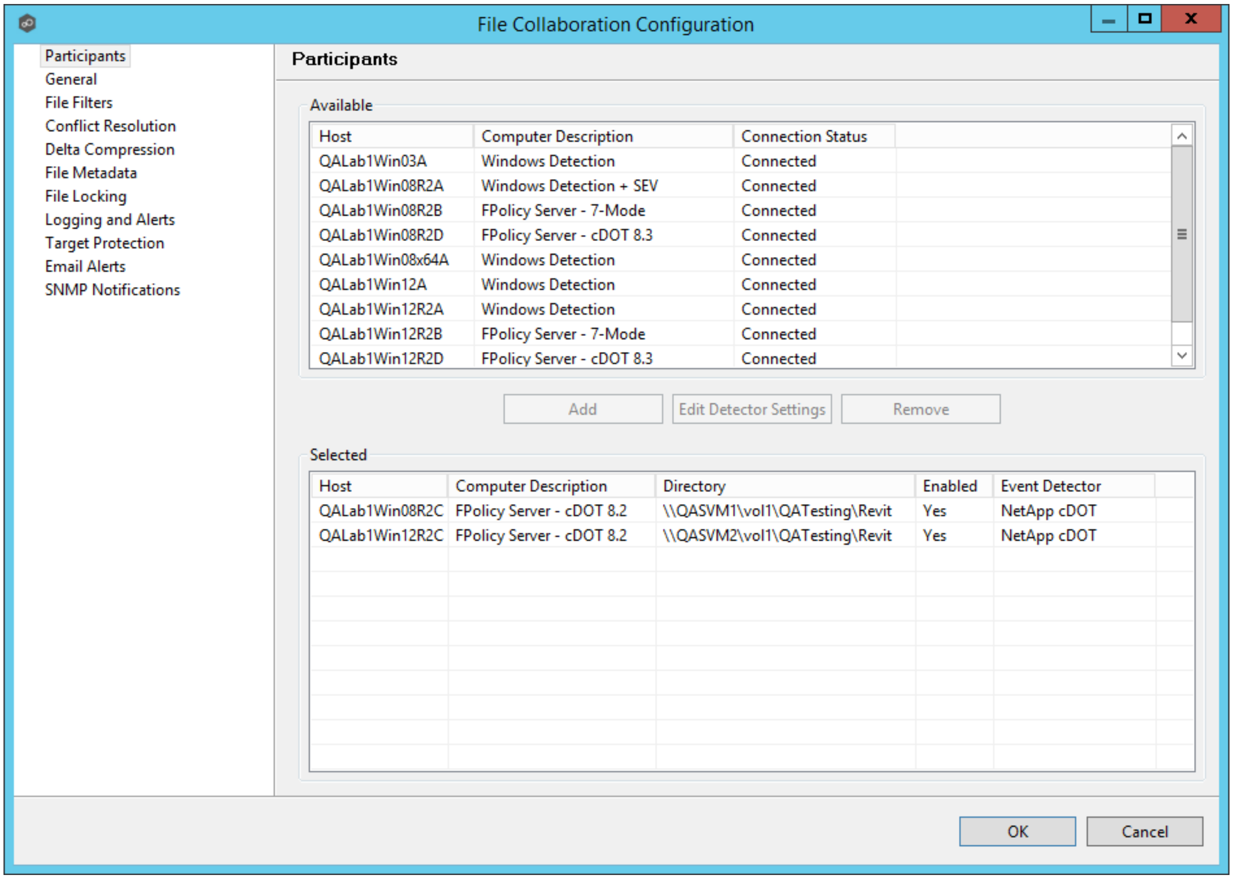

Participants

Available

Host	Computer Description	Connection Status
QALab1Win03A	Windows Detection	Connected
QALab1Win08R2A	Windows Detection + SEV	Connected
QALab1Win08R2B	FPolicy Server - 7-Mode	Connected
QALab1Win08R2D	FPolicy Server - cDOT 8.3	Connected
QALab1Win08x64A	Windows Detection	Connected
QALab1Win12A	Windows Detection	Connected
QALab1Win12R2A	Windows Detection	Connected
QALab1Win12R2B	FPolicy Server - 7-Mode	Connected
QALab1Win12R2D	FPolicy Server - cDOT 8.3	Connected

Selected

Host	Computer Description	Directory	Enabled	Event Detector
QALab1Win08R2C	FPolicy Server - cDOT 8.2	\\QASVM1\vol1\QATesting\Revit	Yes	NetApp cDOT
QALab1Win12R2C	FPolicy Server - cDOT 8.2	\\QASVM2\vol1\QATesting\Revit	Yes	NetApp cDOT

4. A configuration dialog will be displayed showing various settings used to tune how the Agent will communicate with its associated NetApp storage device.

   With NetApp 7-Mode devices, you will see the following:

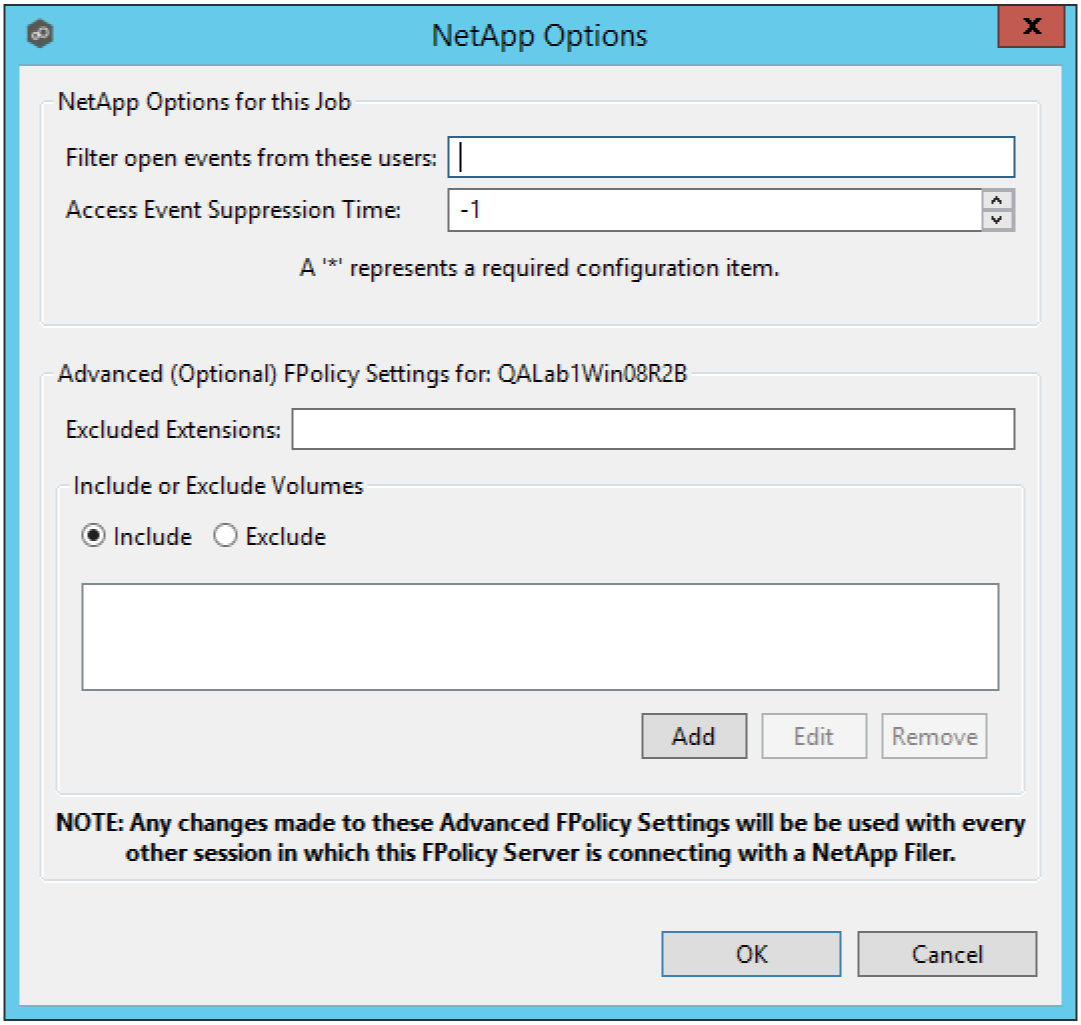

NetApp Options

NetApp Options for this Job

Filter open events from these users:

Access Event Suppression Time: -1

A '*' represents a required configuration item.

Advanced (Optional) FPolicy Settings for: QALab1Win08R2B

Excluded Extensions:

Include or Exclude Volumes

● Include ○ Exclude

NOTE: Any changes made to these Advanced FPolicy Settings will be be used with every other session in which this FPolicy Server is connecting with a NetApp Filer.

Some of the advanced optional settings for 7-Mode devices are as follows:

Filter open 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.
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 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 that an open event will be delayed based on the load of the system.
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.

With NetApp cDOT devices, you will see the following:

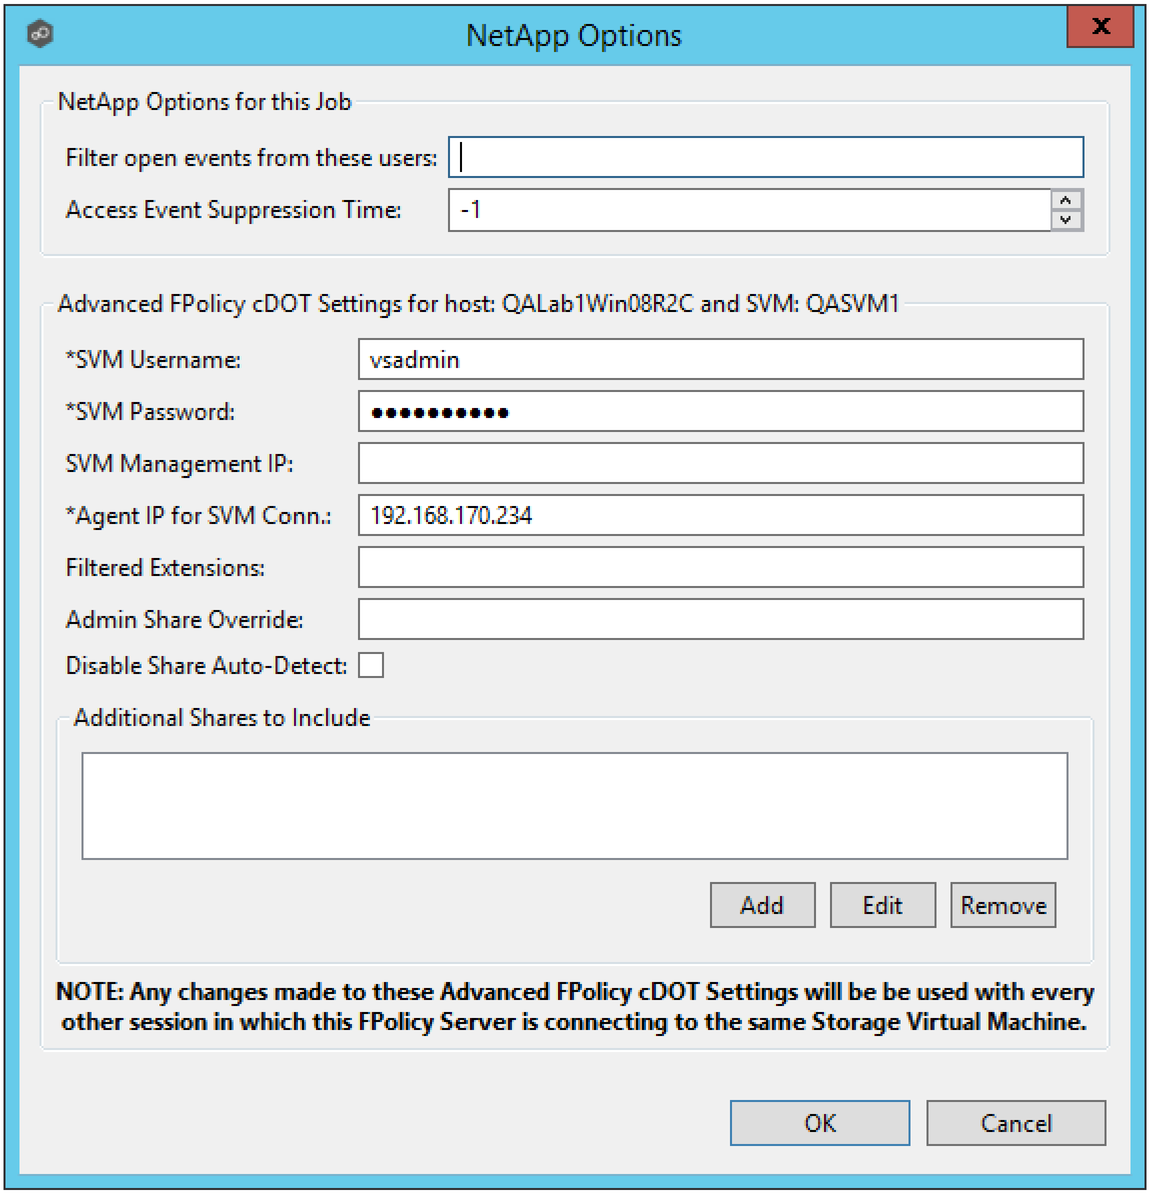

NetApp Options

NetApp Options for this Job

Filter open events from these users:

Access Event Suppression Time: -1

Advanced FPolicy cDOT Settings for host: QALab1Win08R2C and SVM: QASVM1

*SVM Username: vsadmin

*SVM Password: ●●●●●●●●●●●

SVM Management IP:

*Agent IP for SVM Conn.: 192.168.170.234

Filtered Extensions:

Admin Share Override:

Disable Share Auto-Detect: □

Additional Shares to Include

NOTE: Any changes made to these Advanced FPolicy cDOT Settings will be be used with every other session in which this FPolicy Server is connecting to the same Storage Virtual Machine.

Filter open 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.
Access Event Suppres	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 7 clients when mousing over files in Windows Explorer. The default vault is -1, which will use a

sion Time	globally set value. 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.
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	A comma separated list of file extensions to exclude (without a leading asterisk (*).
Admin Share Override	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.
Disable Share Auto-Detect	Disable the option to auto-detect hares and only use the shares defined in the Participants screen and the Additional Shares to Include option below.
Additional Shares to Include	Specify the shares on the SVM that users can use to access the data that Peer Management Center will be collaborating with. For example, if Peer Management Center is collaborating on data that resides under the Departments share with a local namespace path /departments, but users access data via shares to individual sub folders under the Departments folder (such as Marketing with a local namespace path of /departments/marketing and Sales with a

	local namespace path of /departments/sales). In this example, the list of shares would be Departments, Marketing, and Sales.

After making the necessary changes, click **OK** twice to save them. If the selected job is already running, you will need to restart it for the changes to take effect.

Nutanix Configuration

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 the majority of available job types.

For more detailed information about prerequisites and configuration, see Nutanix Prerequisites and Configuration.

Prerequisites

In addition to the standard Peer Global File Service Environmental Requirements, the following prerequisites must be met: https://kb.peersoftware.com/tb/nutanix-files-prerequisites

Configuration

The creation of a new jobs in the Peer Management Center will automatically walk you through entering the most important settings for use in NetApp environments. If you need to modify existing settings or tune advanced settings, you can do so by following these steps:

1. Right-click on a File Collaboration, File Synchronization, File Locking, or File Replication job and click **Edit Job(s)**.

2. Click on the **Participants** page on the left.

3. Select the Peer Agent that is managing the appropriate NetApp storage device then click the **Edit Detector Settings** button.

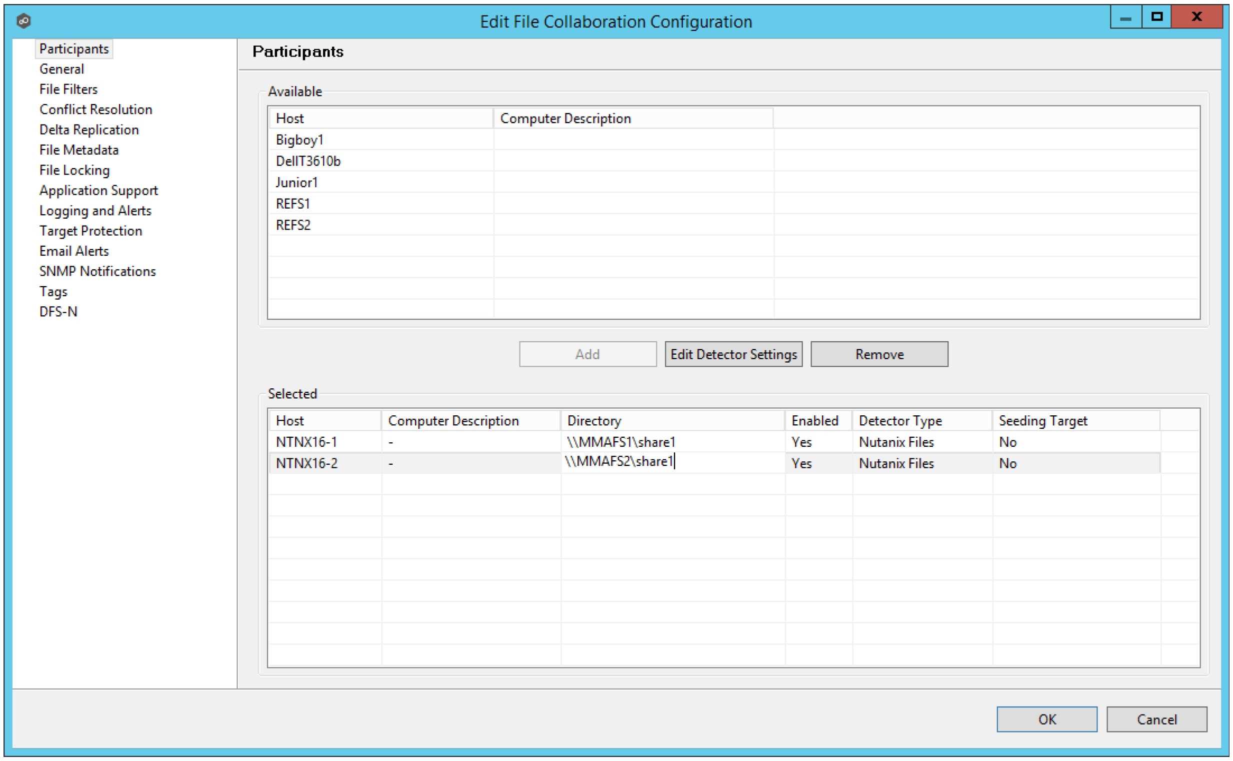

Participants

Available

Host	Computer Description
Bigboy1	
DellT3610b	
Junior1	
REFS1	
REFS2	

Selected

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
NTNX16-1	-	\\MMAFS1\share1	Yes	Nutanix Files	No
NTNX16-2	-	\\MMAFS2\share1	Yes	Nutanix Files	No

4. A configuration dialog will be displayed showing various settings used to tune how the Agent will communicate with its associated Nutanix storage device.

   With Nutanix Files clusters, the following configuration options are available:

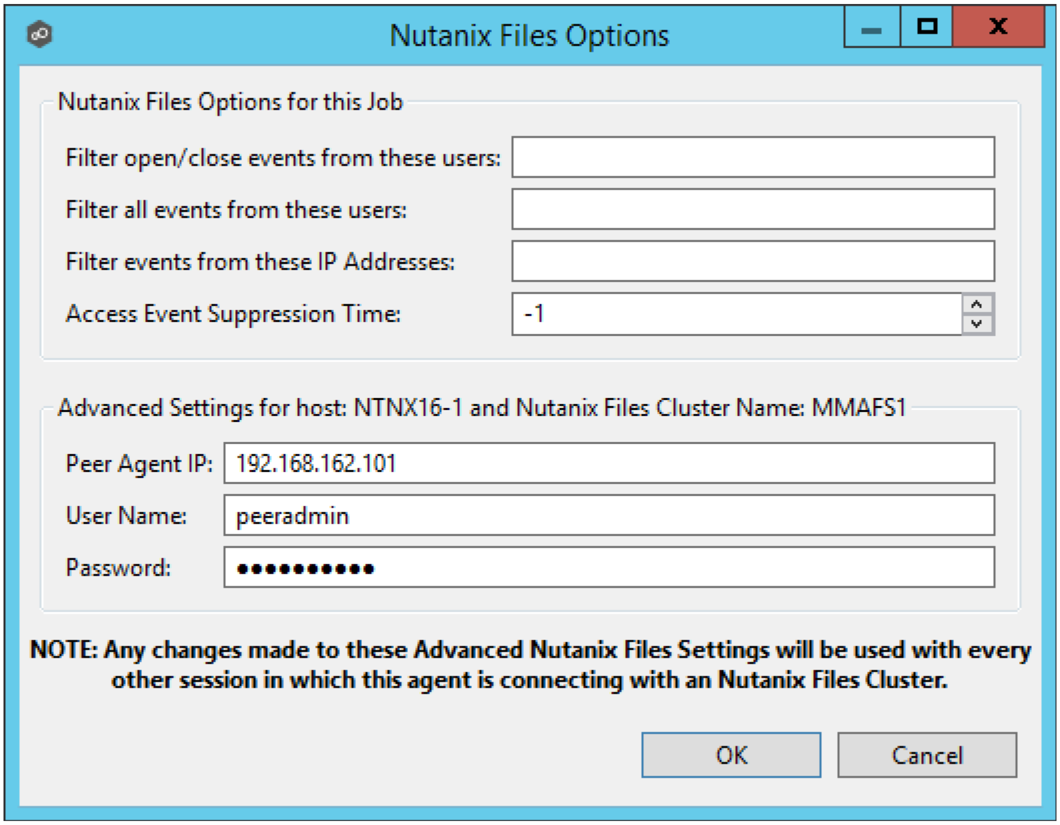

Nutanix Files Options

Nutanix Files Options for this Job

Filter open/close events from these users:

Filter all events from these users:

Filter events from these IP Addresses:

Access Event Suppression Time: -1

Advanced Settings for host: NTNX16-1 and Nutanix Files Cluster Name: MMAFS1

Peer Agent IP: 192.168.162.101

User Name: peeradmin

Password: ●●●●●●●●●●●

NOTE: Any changes made to these Advanced Nutanix Files Settings will be used with every other session in which this agent is connecting with an Nutanix Files Cluster.

Nutanix Files Options

Filter open 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.
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 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 that an open event will be delayed based on the load of the system.

Advanced Settings for Nutanix Files

Partner Server 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 configured Files cluster.

Password	Password used to access the APIs on the configured Files cluster.

5. After making the necessary changes, click **OK** twice to save them. If the selected job is already running, you will need to restart it for the changes to take effect.

Real-time Event Detection

A number of 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 the Peer Management Center. See Real-time Event Detection in the File Collab, Sync, and Locking Preferences topic for more information.

To view and modify real-time event detection settings for all job types:

1. From the **Window** menu, and select **Preferences**.

2. Select **Real-time Detection** in the navigation tree.

   The following page is displayed.

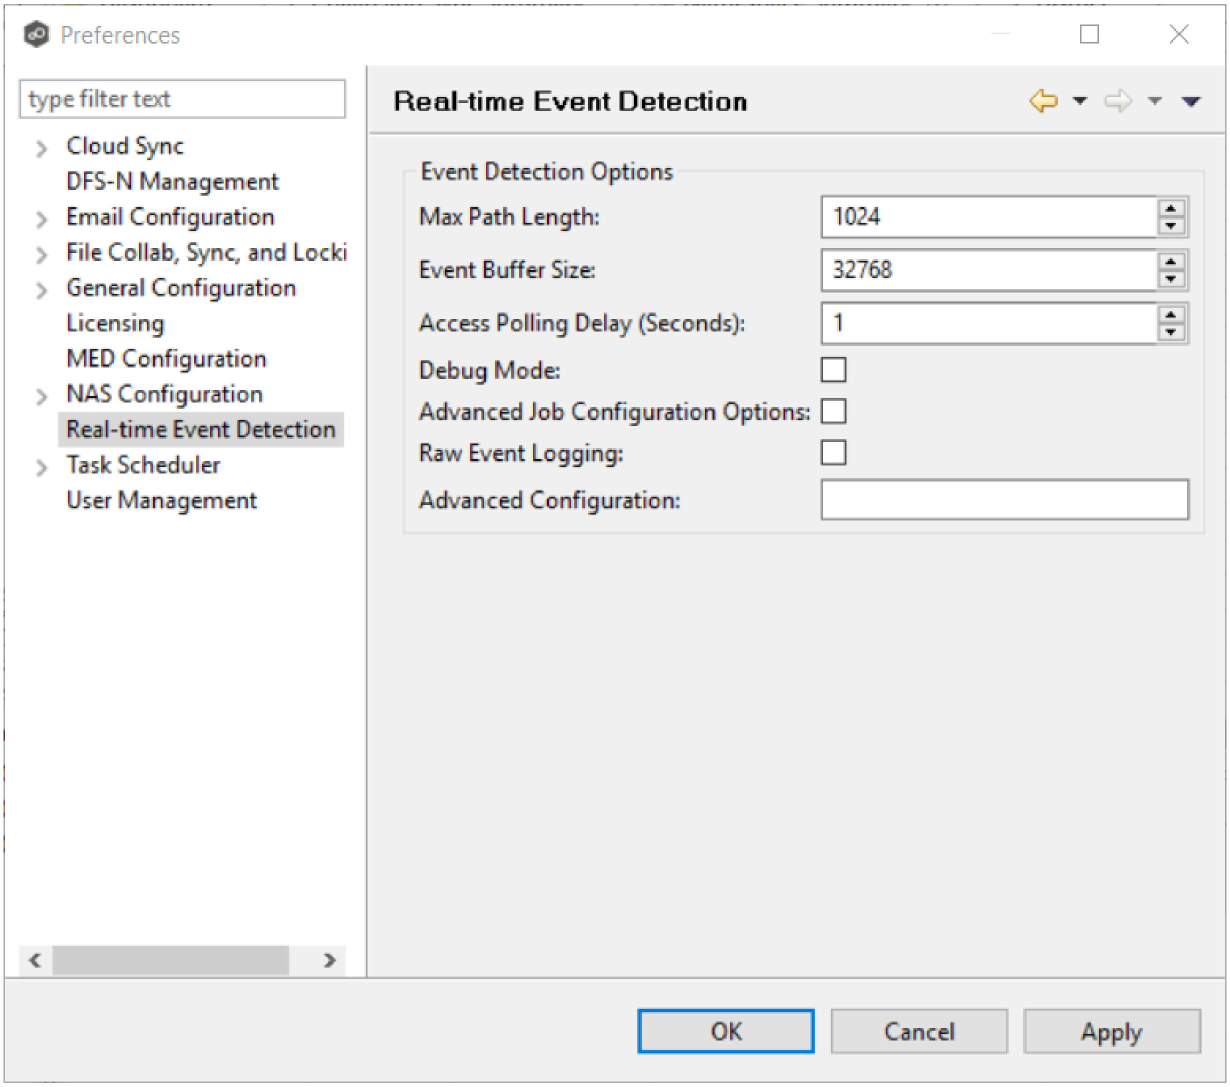

Real-time Event Detection

Event Detection Options

Max Path Length: 1024

Event Buffer Size: 32768

Access Polling Delay (Seconds): 1

Debug Mode: □

Advanced Job Configuration Options: □

Raw Event Logging: □

Advanced Configuration:

3. Modify values as needed:

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**.

User Management

Management of users with access to Peer Management Center's web interface can be performed through either the Peer Management Center's rich client or through an **admin** account logged into the web interface.

From the User Management page, you can add, edit, and remove internal user accounts, roles, and active directory users and groups.

To access the User Management page:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.

   The following page is displayed:

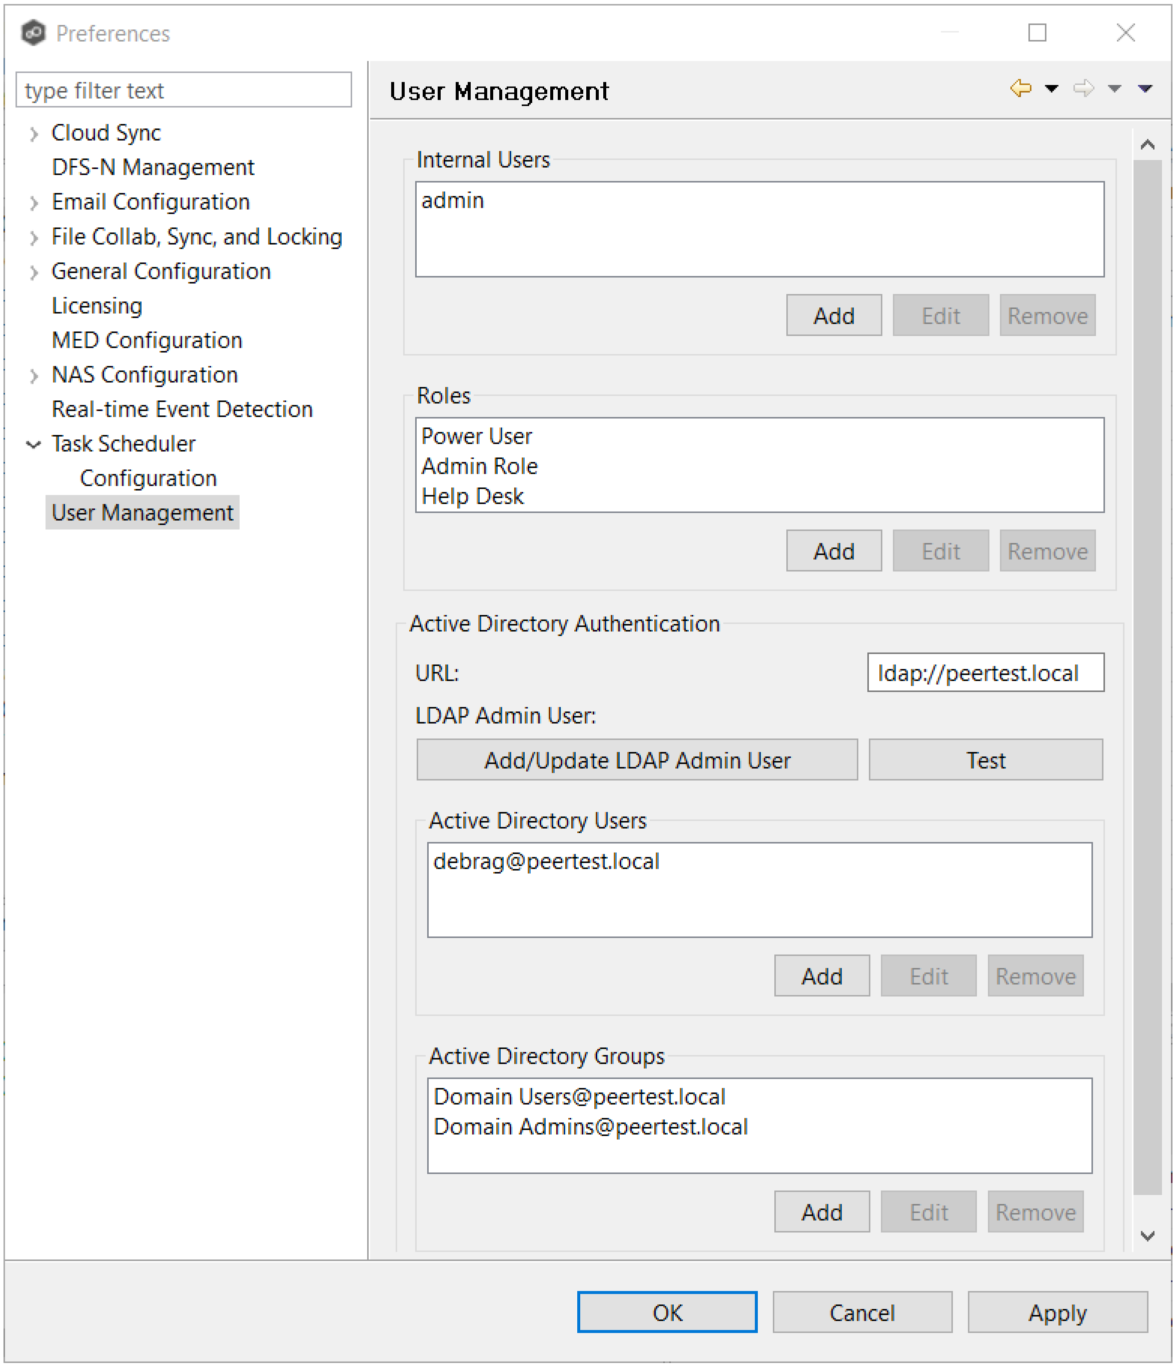

Internal Users: admin

Roles: Power User, Admin Role, Help Desk

URL: ldap://peertest.local

LDAP Admin User:

Active Directory Users: debrag@peertest.local

Active Directory Groups: Domain Users@peertest.local, Domain Admins@peertest.local

3. From this screen, you can add, edit, and remove internal user accounts, roles, and active directory users and groups.

Web Interface Roles

Roles define user permissions to access and edit resources in the web interface.

There are three predefined roles with specific set of permissions: **Power User**, **Admin Role**, and **Help Desk**.

You can create custom roles, edit roles, delete roles, and assign roles:

- To create a custom role or to edit a role, click select **User Management** > **Add** or **Edit** button in the **Roles** section. Adding a custom role requires a name, display name, description, and base role.

- To delete a role, select the role from the **User Management** > **Roles** section, and then click **Remove**.

- Roles can also be assigned existing tags to define the resources users with that role can view and edit.

The following table outlines the Hub resources that each role can edit and view.

	Power User	**Admin Role**	**Help Desk**
Tag Resources Dialog		Edit	
PeerSync Summary View	Edit	Edit	Edit
PeerSync Job Stats View Part	Edit	Edit	
Memory Dump Action	Edit	Edit	
Advisory Alert View	Edit	Edit	
Runtime Summary Interface	Edit	Edit	View Only
Permission Mode	Edit	Edit	
Status Agent Tree View	View Only	Edit	
Session View	Edit	Edit	View Only

Peerlet View	Edit	Edit	View Only
Preferences		Edit	
Broker Statistics Action	Edit	Edit	
Hub Alert View	Edit	Edit	
PeerSync Configuration Interface	Edit	Edit	View Only
PeerSync Job Stats View	Edit	Edit	Edit
Event Analyzer Configuration Interface	Edit	Edit	
Collaboration Summary View	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
Folder Analyzer View	Edit	Edit	View Only
Hub Save All	Edit	Edit	
PeerSync Participant View	Edit	Edit	
New Peerlet Action		Edit	
File Conflict View	Edit	Edit	Edit

Configuration Interface	Edit	Edit	
Hub View Progress	Edit	Edit	
PeerSync Advisory Alert View	Edit	Edit	
Event Analyzer Validation View	Edit	Edit	View Only
Expression Info Dialog	Edit	Edit	
Hub Refresh Perspective	Edit	Edit	
Event Analyzer Runtime Summary Interface	Edit	Edit	View Only
Peerlet Alert View	Edit	Edit	
Log Dump Action	Edit	Edit	
Event Log View	Edit	Edit	
PeerSync Messages Log View	Edit	Edit	Edit
Event Analyzer Participant view	Edit	Edit	
Event Analyzer Log View	Edit	Edit	View Only
PeerSync Delete Log View	Edit	Edit	Edit
Thread Dump Action	Edit	Edit	
Participant View	Edit	Edit	
Hub Download Agent	Edit	Edit	

PeerSync Event Log View	Edit	Edit	
File Sync Advisory Alert View	Edit	Edit	
Expression List Dialog	Edit	Edit	

Active Directory Authentication

In addition to internal users, the Peer Management Center also provides Active Directory user and group authentication.

To configure Active Directory authentication:

1. Provide the URL of the LDAP server on the network in the one of the following formats:

   ldap://MYDOMAIN.LOCAL

   or

   ldaps://MYDOMAIN.LOCAL

2. Add an Active Directory user or group by clicking **User Management > Active Directory Users/Active Directory Groups > Add** button.

3. Enter the Domain, Username or Group, and the Role in the configuration dialog.

To delete an Active Directory user or group:

1. Click the **User Management > Active Directory Users/Active Directory Groups > Remove** button.

Note: Active Directory users and groups are saved in the following format:

Username@MYDOMAIN.LOCAL

Use this format to log into the Peer Management Center's web interface:

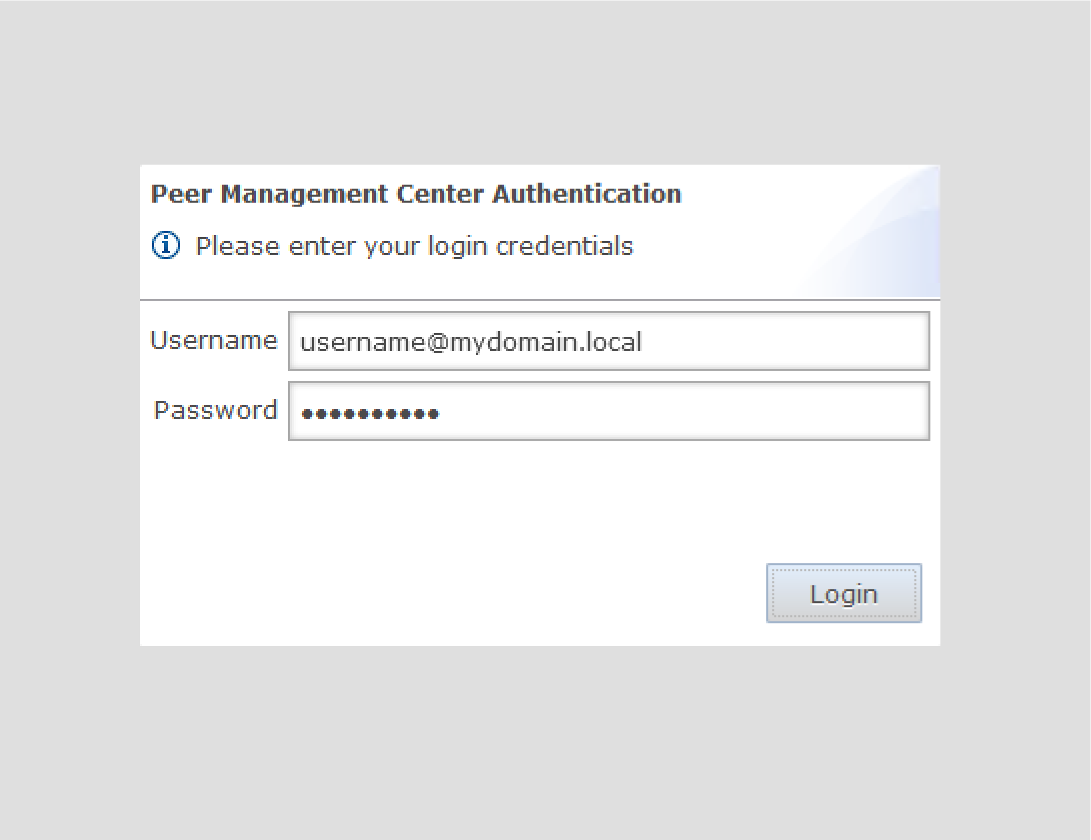

Peer Management Center Authentication

ⓘ Please enter your login credentials

Username: username@mydomain.local

Password: ●●●●●●●●●●

Internal Users

Adding an internal account requires a username, a password, an email address, and a selected role. For more details on the available roles, see Web Interface. Once an account has been created, its username, password, email address, and role can all be changed. The default **admin** user account password is **password**.

Notes:

- The default **admin** user cannot be renamed, nor can its role be changed.
- These user accounts have no impact on access to the rich client.

Cloud Sync Jobs

This section provides information about creating, running, and managing a Cloud Sync job:

- Overview
- Before You Create Your First Cloud Sync Job
- Creating a Cloud Sync Job
- Running a Cloud Sync Job
- Monitoring Your Cloud Sync Jobs
- Deleting a Cloud Sync Job
- Recovering Data from the Cloud

Overview

Cloud Sync 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 Sync efficiently pushes data into Microsoft Azure Blob storage in an open format that is immediately consumable by other applications and services.

Use cases for Cloud Sync include pushing exact replicas of on-premises data sets into object storage for use with burstable compute and cloud-borne services and for tape replacement-style backup to object with point-in-time recovery capability.

Before You Create Your First Cloud Sync Job

We strongly recommend that you configure the Cloud Sync settings as well as other global settings such as SMTP configuration and email alerts before configuring your first Cloud Sync job. See Preferences for details on what and how to configure these settings.

Creating a Cloud Sync Job

The **Create Job Wizard** walks you through the process of creating a Cloud Sync 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: Storage Information

Step 5: Source Paths

Step 6: File Filters

Step 7: Destination

Step 8: Cloud Platform Credentials

Step 9: Miscellaneous Options

Step 10: Sync and Retention Policy

Step 11: Sync Schedule

Step 12: Retention

Step 13: Email Alerts

Step 14: Confirmation

Step 1: Job Type and Name

1. Open the Peer Management Center.

2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **New Job** wizard displays a list of job types you can create.

3. Click **Cloud Sync**, and then click **Create**.

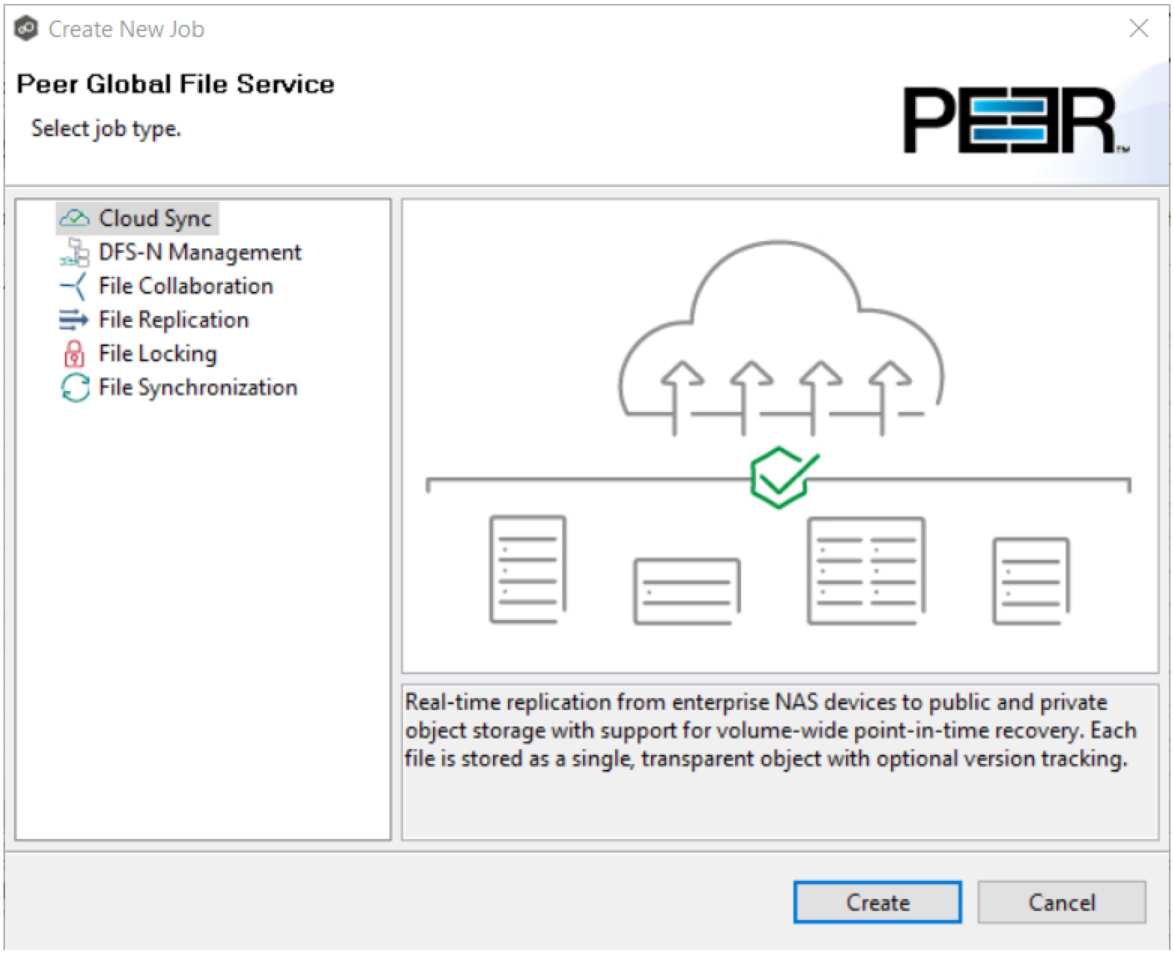

Real-time replication from enterprise NAS devices to public and private object storage with support for volume-wide point-in-time recovery. Each file is stored as a single, transparent object with optional version tracking.

4. Enter a name for the job in the dialog that appears.

   The job name must be unique.

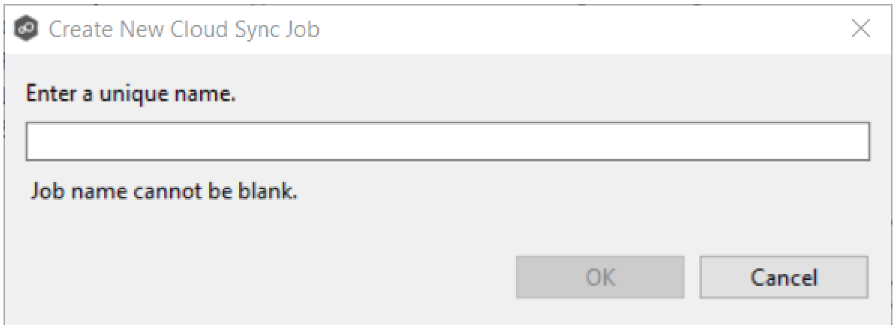

Enter a unique name.

Job name cannot be blank.

5. Click **OK**.

Step 2: Source Storage Platform

The **Source Storage Platform** page lists the types of source storage platforms that Cloud Sync supports. The source storage device hosts the data you want to replicate.

1. Select the type of storage platform you want to replicate.

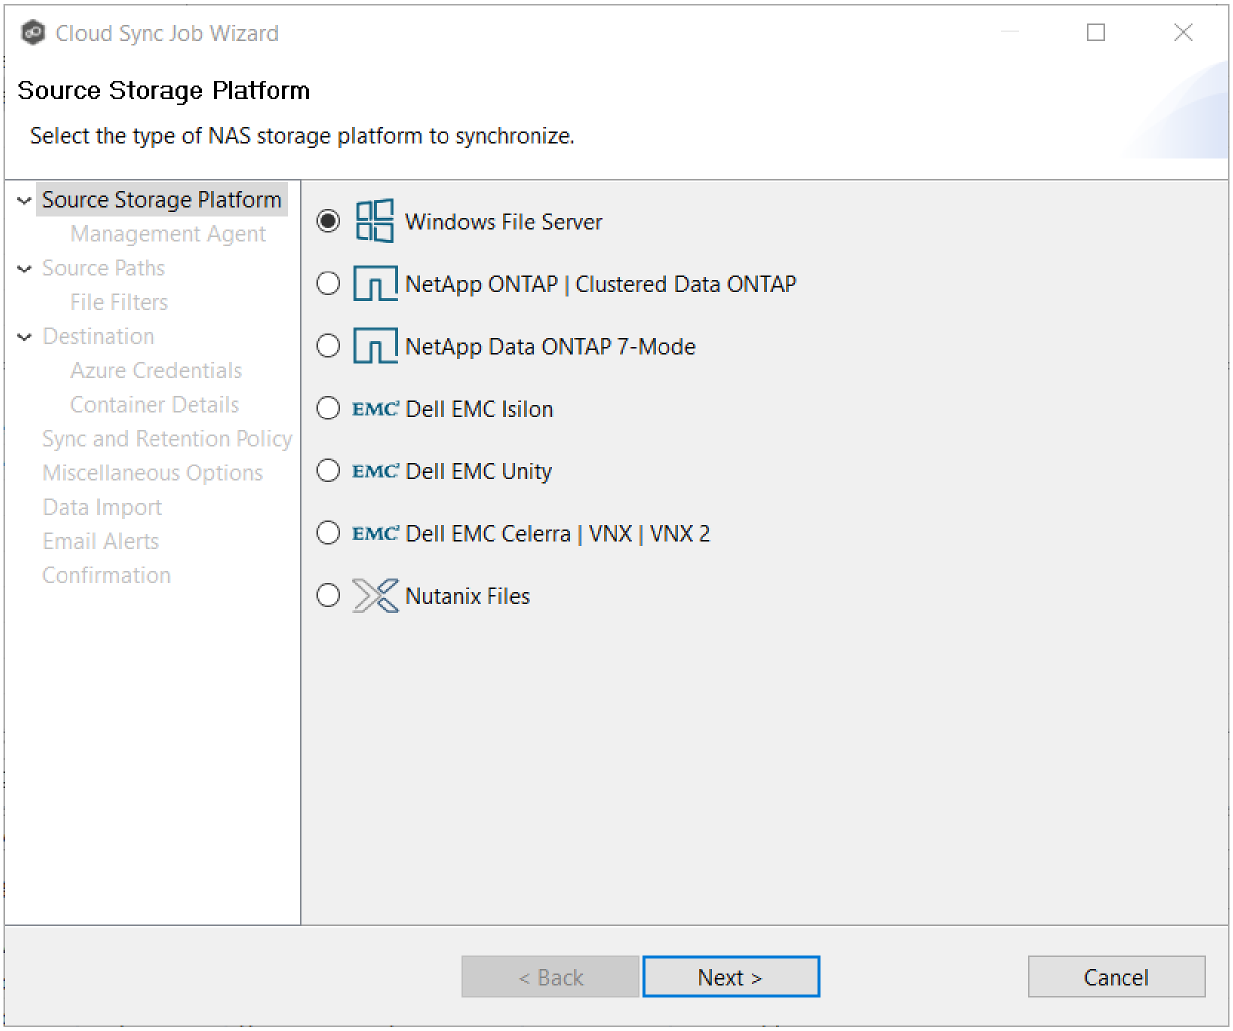

Cloud Sync Job Wizard

Source Storage Platform

Select the type of NAS storage platform to synchronize.

- ● Windows File Server
- ○ NetApp ONTAP | Clustered Data ONTAP
- ○ NetApp Data ONTAP 7-Mode
- ○ EMC² Dell EMC Isilon
- ○ EMC² Dell EMC Unity
- ○ EMC² Dell EMC Celerra | VNX | VNX 2
- ○ Nutanix Files

2. Click **Next**.

Step 3: Management Agent

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. You should select the Management 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.

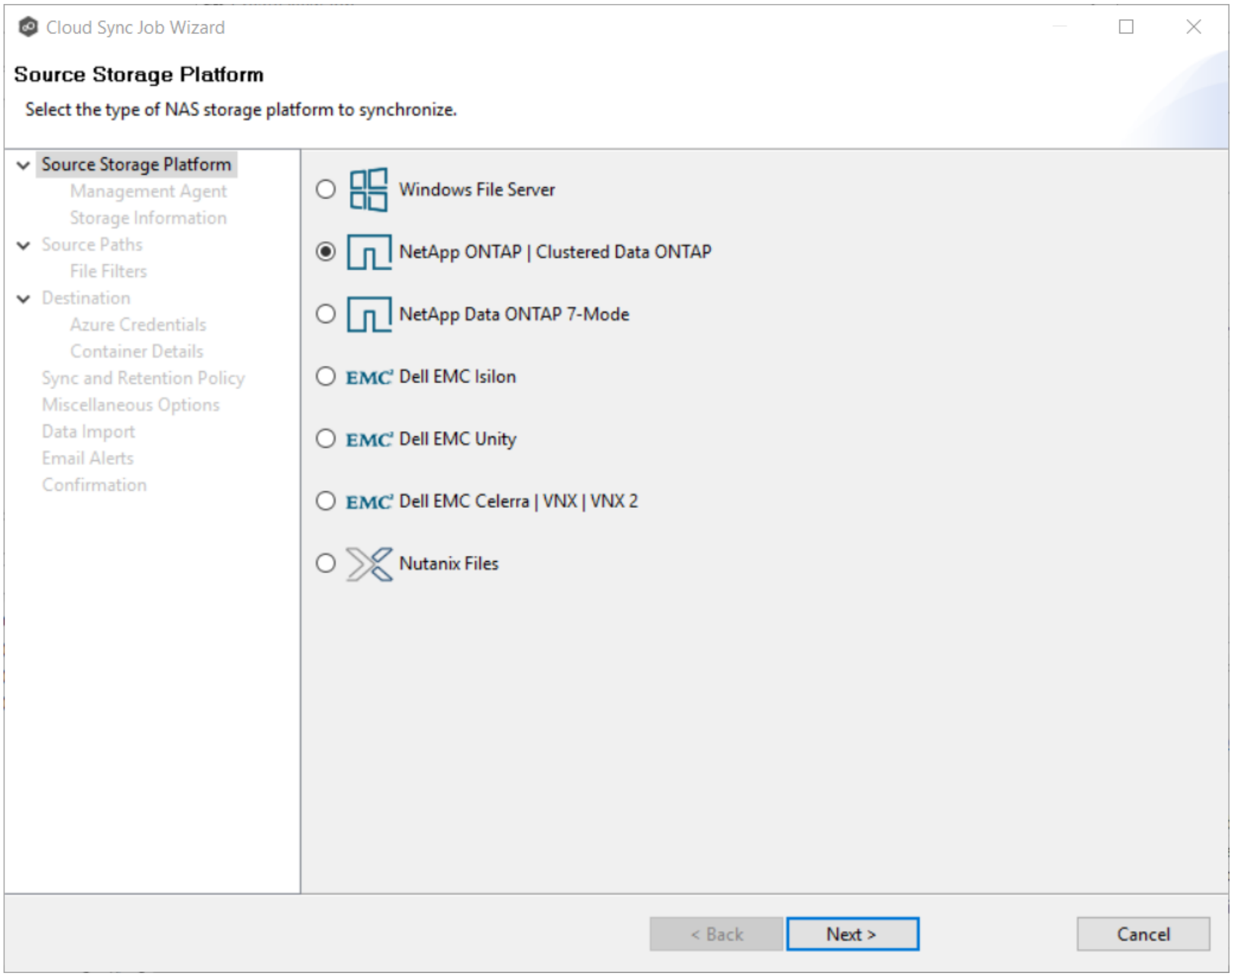

Source Storage Platform

Select the type of NAS storage platform to synchronize.

○ Windows File Server
● NetApp ONTAP | Clustered Data ONTAP
○ NetApp Data ONTAP 7-Mode
○ EMC Dell EMC Isilon
○ EMC Dell EMC Unity
○ EMC Dell EMC Celerra | VNX | VNX 2
○ Nutanix Files

2. Click **Next**.

   If you get a message that the database connection is not found, click **OK** and then configure the database connection for the selected Management Agent. See Database Connections for information about configuring a database connection. After configuring the database connection, continue with the Storage Information page.

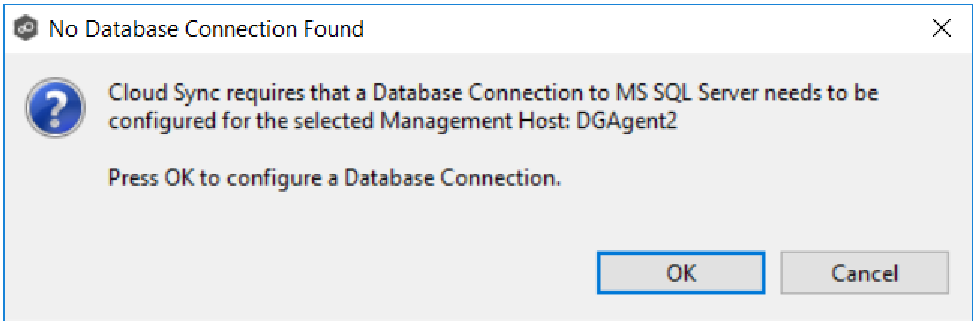

Cloud Sync requires that a Database Connection to MS SQL Server needs to be configured for the selected Management Host: DGAgent2

Press OK to configure a Database Connection.

Step 4: Storage Information

If you selected any storage platform other than Windows File Server in Step 2, the **Storage Information** page appears. It requests the credentials necessary to connect to the storage device you want to replicate.

If you selected **Windows File Server**, skip to Step 5: Source Paths.

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 Paths.

   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**.

1. Enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

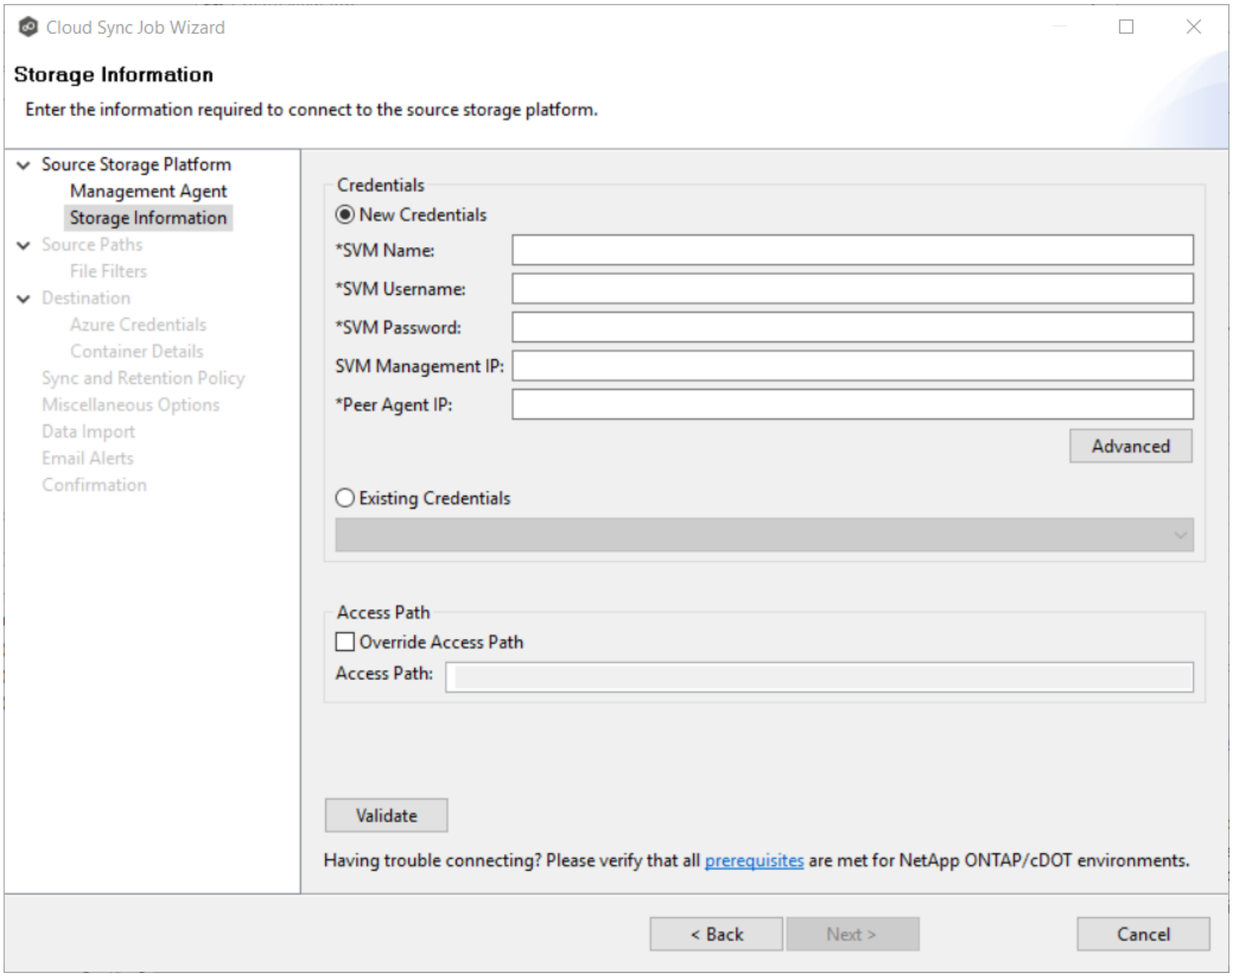

Cloud Sync Job Wizard

Storage Information

Enter the information required to connect to the source storage platform.

Credentials

● New Credentials

*SVM Name:

*SVM Username:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

○ Existing Credentials

Access Path

□ Override Access Path

Access Path:

Having trouble connecting? Please verify that all prerequisites are met for NetApp ONTAP/cDOT environments.

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	Enter the IP address of the server hosting the Agent that manages the Storage Virtual Machine.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

1. Enter the credentials to connect to the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

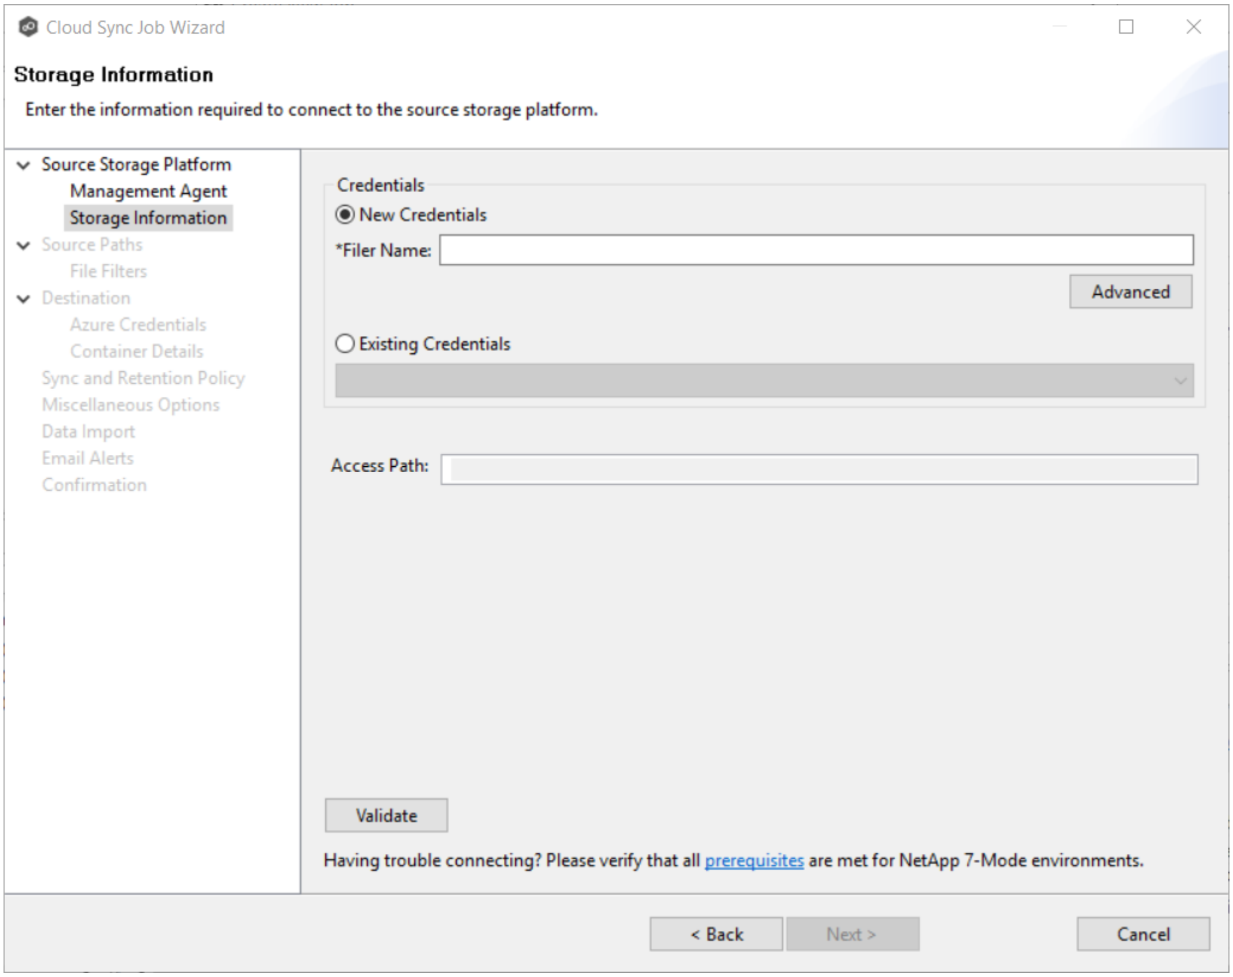

Cloud Sync Job Wizard

Storage Information

Enter the information required to connect to the source storage platform.

Credentials

● New Credentials

*Filer Name:

○ Existing Credentials

Access Path:

Having trouble connecting? Please verify that all prerequisites are met for NetApp 7-Mode environments.

Filer Name	Enter the name of the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

1. Enter the credentials to connect to the EMC Isilon cluster hosting the data to be replicated.

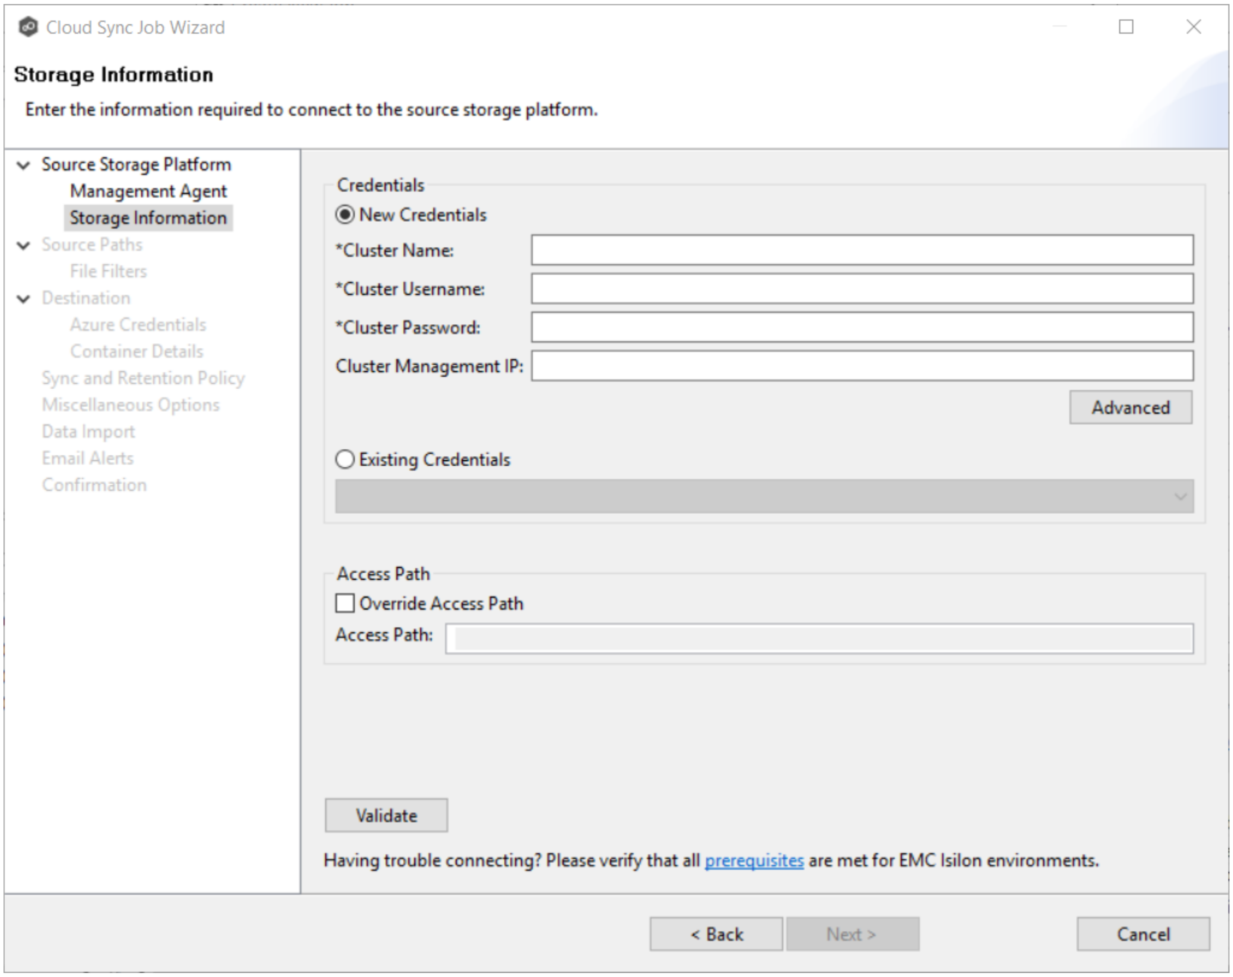

Cloud Sync Job Wizard

Storage Information

Enter the information required to connect to the source storage platform.

● New Credentials

○ Existing Credentials

□ Override Access Path

Having trouble connecting? Please verify that all prerequisites are met for EMC Isilon environments.

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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

1. Enter the credentials to connect to the NAS Server hosting the data to be replicated.

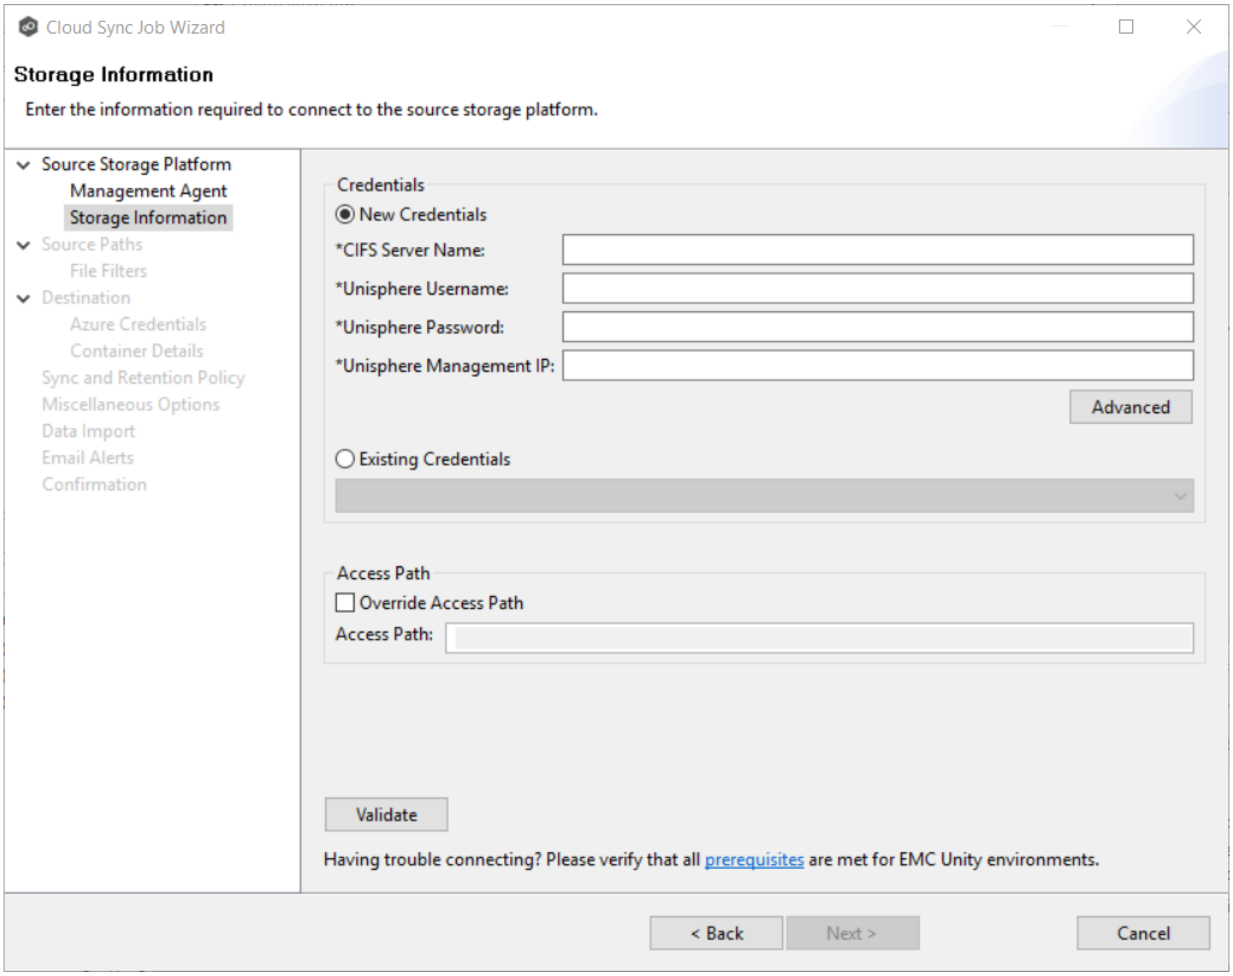

Storage Information

Enter the information required to connect to the source storage platform.

Credentials

● New Credentials

*CIFS Server Name:

*Unisphere Username:

*Unisphere Password:

*Unisphere Management IP:

○ Existing Credentials

Access Path

□ Override Access Path

Access Path:

Having trouble connecting? Please verify that all prerequisites are met for EMC Unity environments.

NAS 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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

1. Enter the credentials to connect to the CIFS Server hosting the data to be replicated.

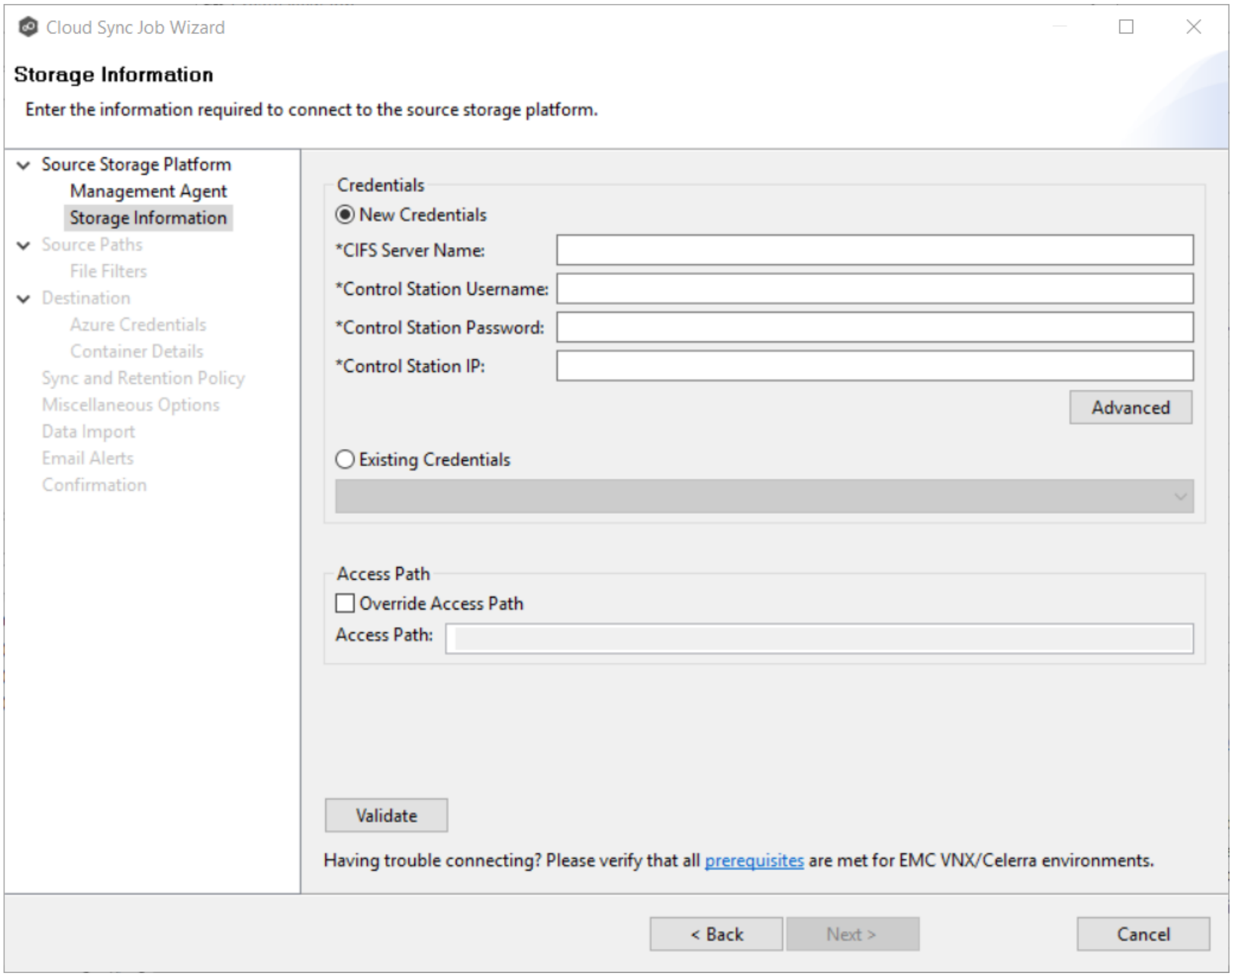

Cloud Sync Job Wizard

Storage Information

Enter the information required to connect to the source storage platform.

Credentials

● New Credentials

*CIFS Server Name:

*Control Station Username:

*Control Station Password:

*Control Station IP:

○ Existing Credentials

Access Path

□ Override Access Path

Access Path:

Having trouble connecting? Please verify that all prerequisites are met for EMC VNX/Celerra environments.

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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

1. Enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated.

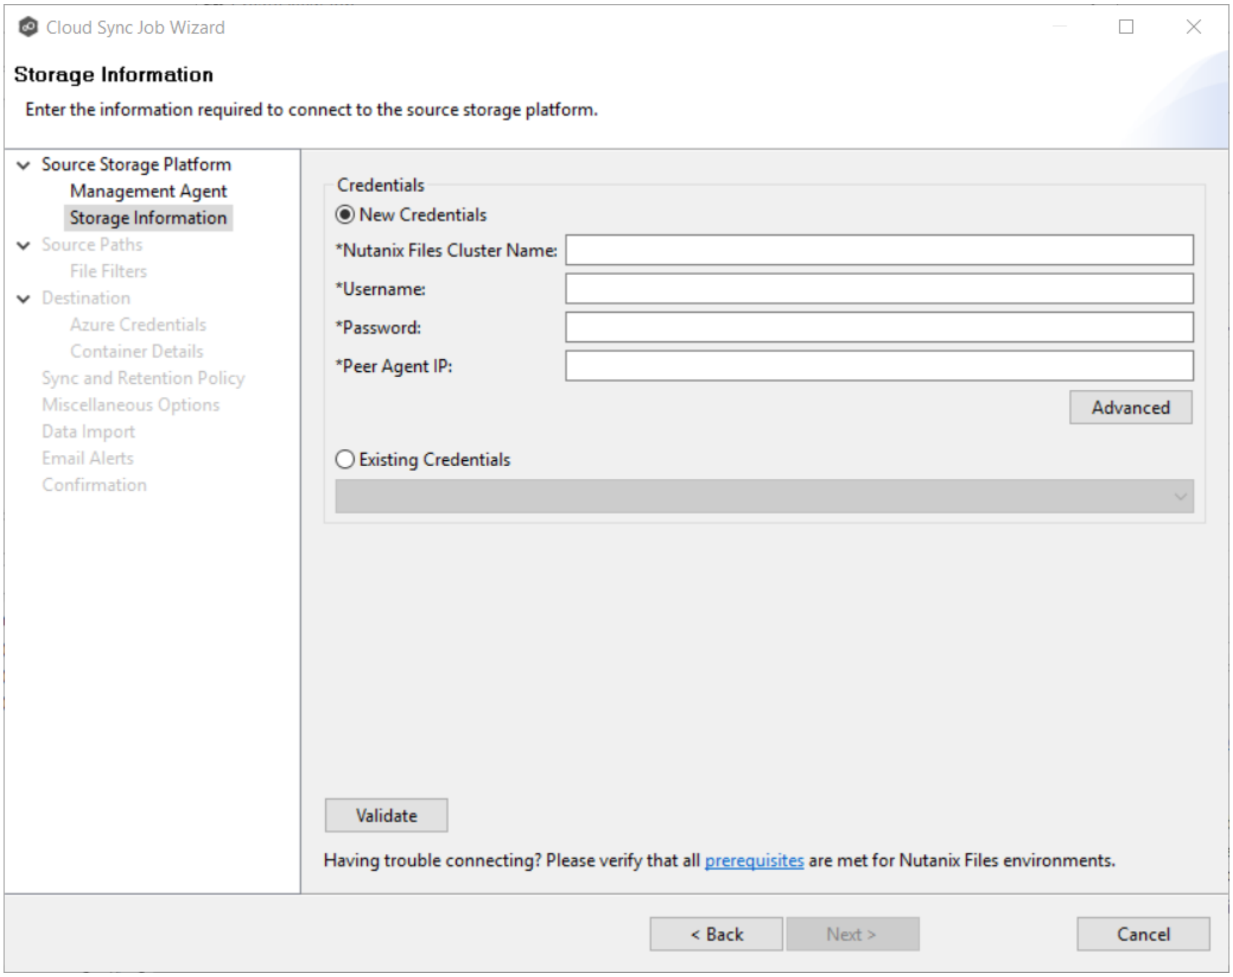

Nutanix Files Cluster Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the user name for the account managing the AFS cluster via its management APIs.
Password	Enter the password for the account managing the AFS cluster via its management APIs.
Peer Agent IP	Enter the IP address of the Agent server used to manage the storage platform. This should not point to the AFS cluster itself.

Step 5: 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 and folders selected for replication are referred to as the watch set.

1. Select the paths to the files and folders you want to replicate.

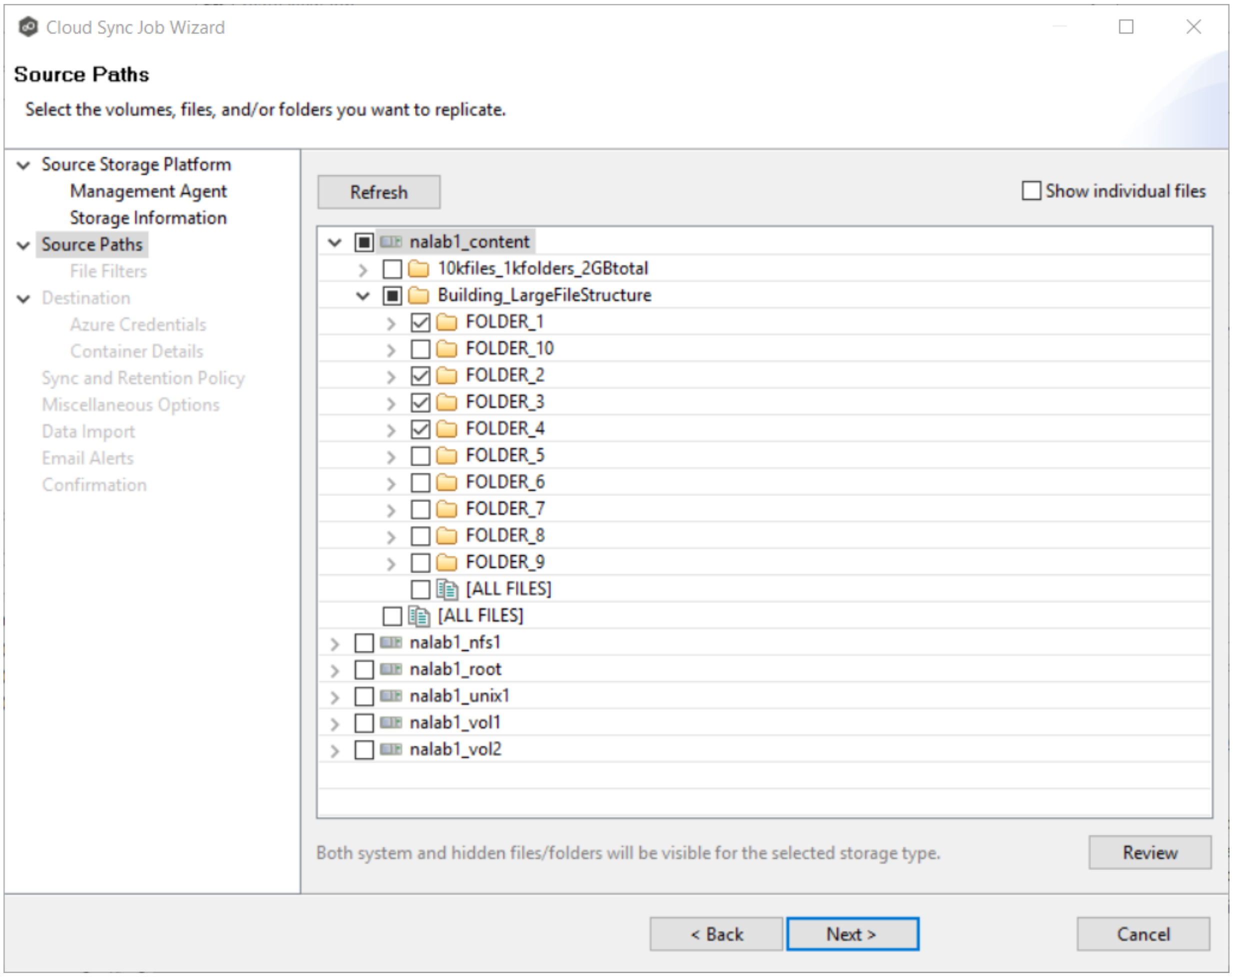

Cloud Sync Job Wizard

Source Paths

Select the volumes, files, and/or folders you want to replicate.

□ Show individual files

- ■ nalab1_content
  - □ 10kfiles_1kfolders_2GBtotal
  - ■ Building_LargeFileStructure
    - ☑ FOLDER_1
    - □ FOLDER_10
    - ☑ FOLDER_2
    - ☑ FOLDER_3
    - ☑ FOLDER_4
    - □ FOLDER_5
    - □ FOLDER_6
    - □ FOLDER_7
    - □ FOLDER_8
    - □ FOLDER_9
    - □ [ALL FILES]
  - □ [ALL FILES]
- □ nalab1_nfs1
- □ nalab1_root
- □ nalab1_unix1
- □ nalab1_vol1
- □ nalab1_vol2

Both system and hidden files/folders will be visible for the selected storage type.

To replicate:

The entire volume (all files and folders, including subfolders and their files)	Select the volume check box.
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 check box.
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** check box, 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**.

Step 6: File Filters

The **File Filters** page displays a list of file filters. A file 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 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.

For general information about file filters, see File Filters in the Concepts section for information about creating file filters for a Cloud Sync job, see File Filters in the Cloud Sync Preferences section.

1. Select the file filters you want to apply to the job.

   Click **Edit File Filters** if you want to create a new file filter or modify an existing one. See Cloud Sync File Filters for more details.

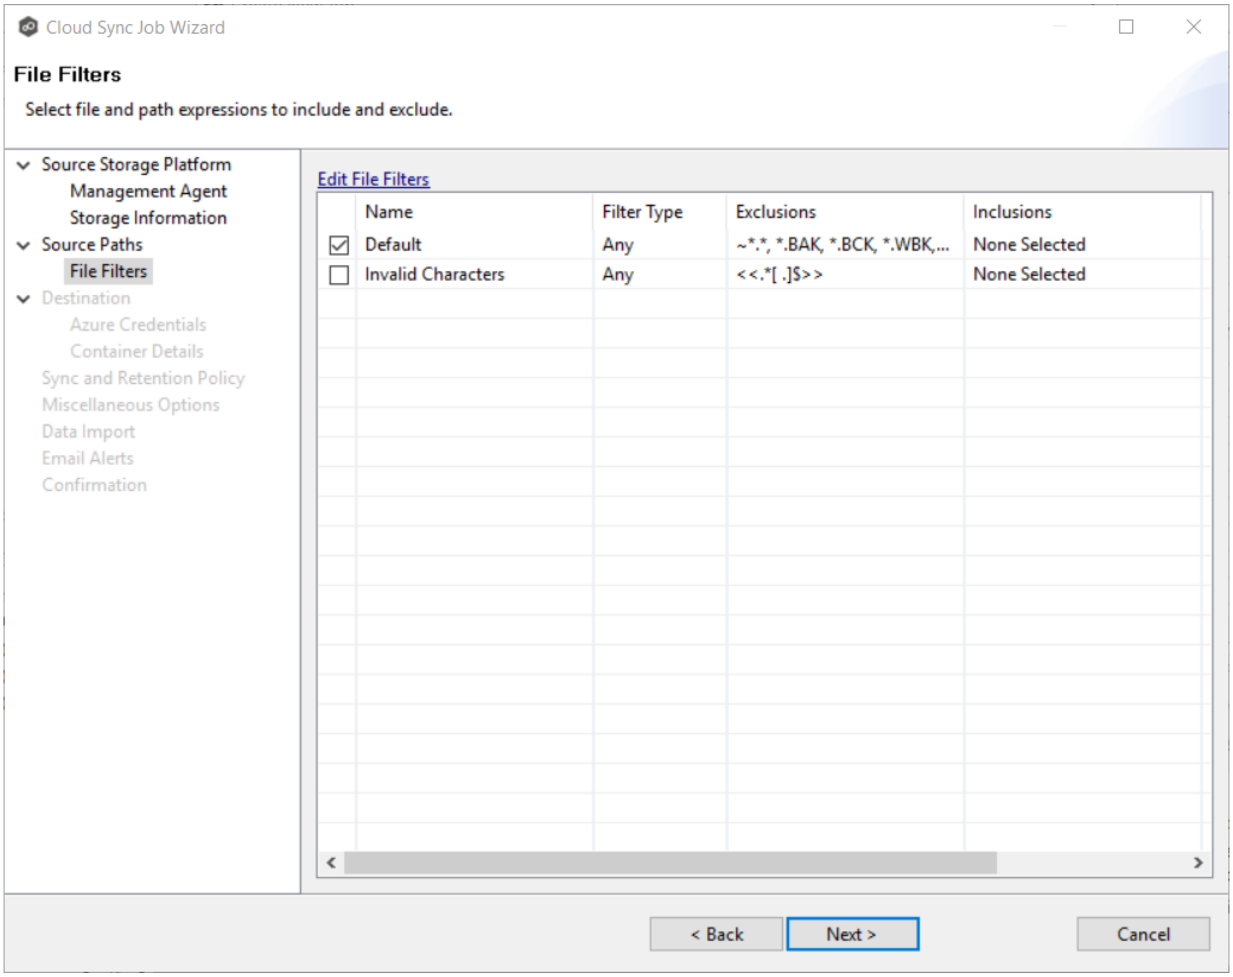

File Filters

Select file and path expressions to include and exclude.

	Name	Filter Type	Exclusions	Inclusions
☑	Default	Any	~*.*, *.BAK, *.BCK, *.WBK,...	None Selected
☐	Invalid Characters	Any	<<.*[.]$>>	None Selected

2. Click **Next**.

Step 7: Destination

The **Destination** page displays a list of the currently available storage platforms to which Cloud Sync can replicate.

1. Select the destination storage platform.

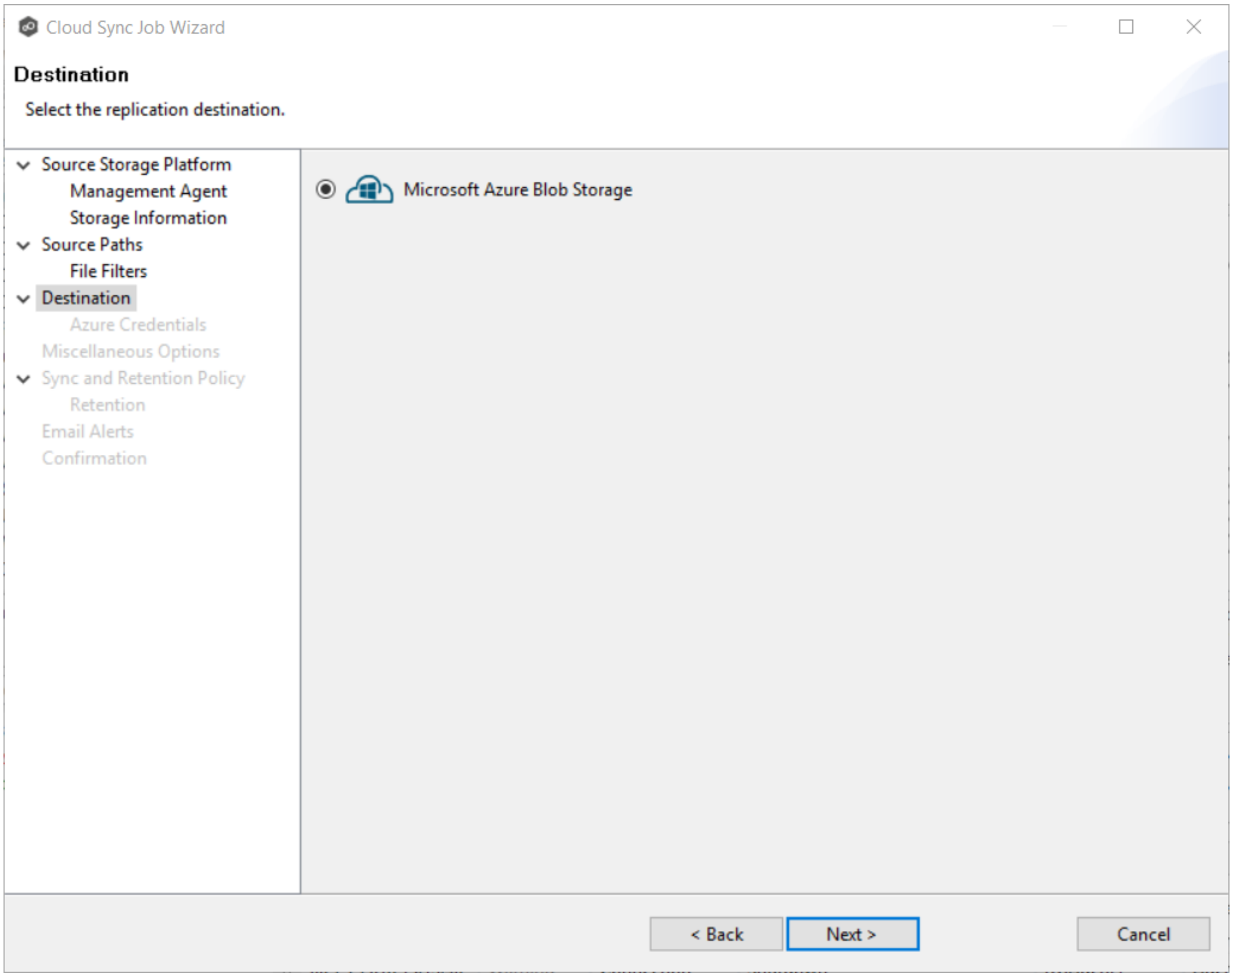
Cloud Sync Job Wizard

Destination

Select the replication destination.

● Microsoft Azure Blob Storage

2. Click **Next**.

Step 8: Cloud Platform Credentials

The **Credentials** page requests the credentials necessary to connect to the destination storage device.

1. Select **New Credentials** to enter a new set of credentials for the destination cloud storage platform or select **Existing Credentials**.

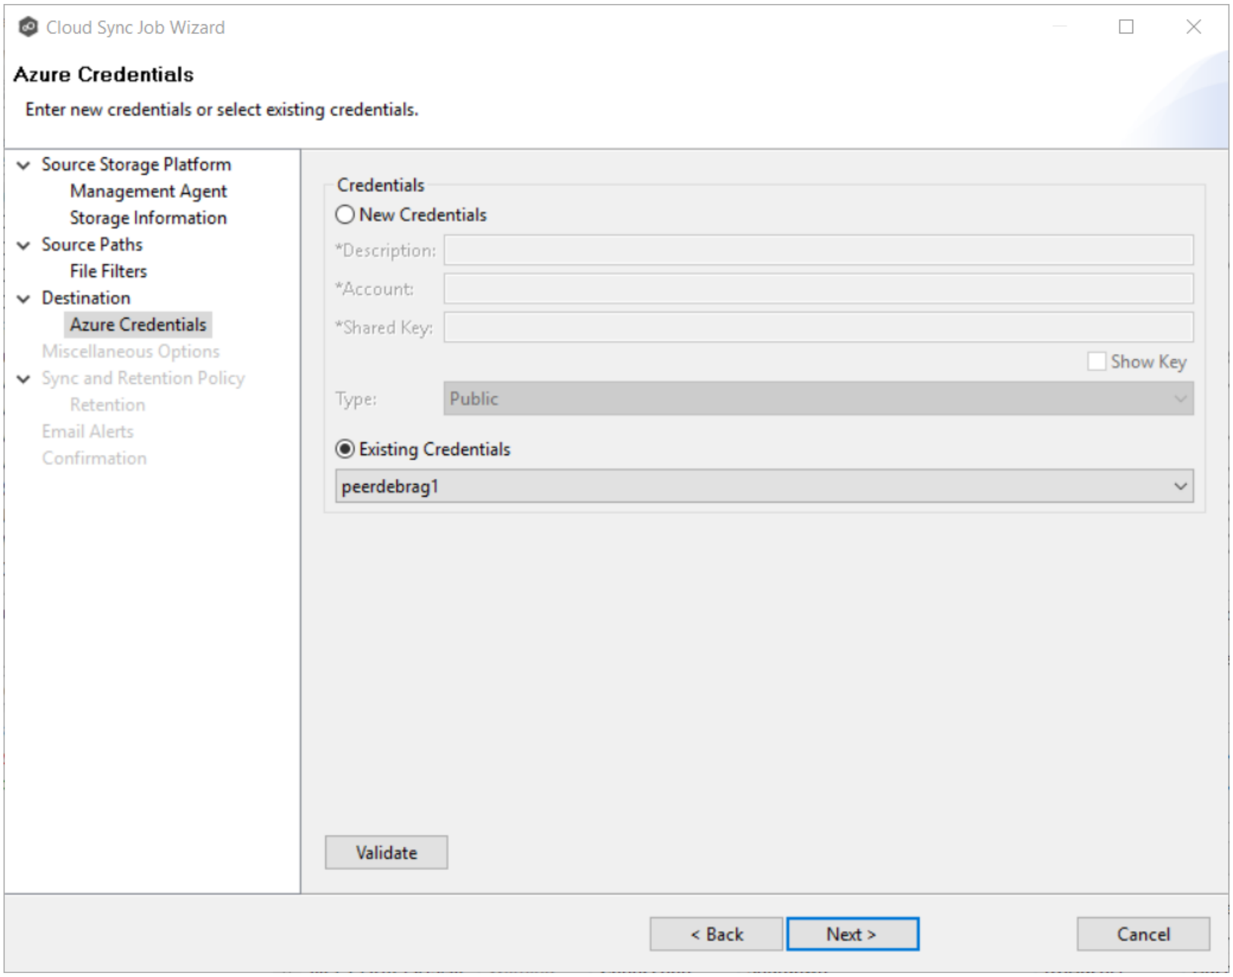
Cloud Sync Job Wizard

Azure Credentials

Enter new credentials or select existing credentials.

Credentials

○ New Credentials

*Description:

*Account:

*Shared Key:

□ Show Key

Type: Public

● Existing Credentials

peerdebrag1

2. If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next**. Continue with Step 9. Miscellaneous Options.

   If you selected **New Credentials**, enter the credentials for connecting to the destination cloud storage platform.

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 of the Azure Storage account, which can be found in the Azure Portal.

3. Click **Validate** to test the connection.

4. Click **Next**.

Step 9: Miscellaneous Options

The **Miscellaneous Options** page displays options for file metadata replication, delta-level replication, and cloud tiering.

1. Select the options to apply to this job.

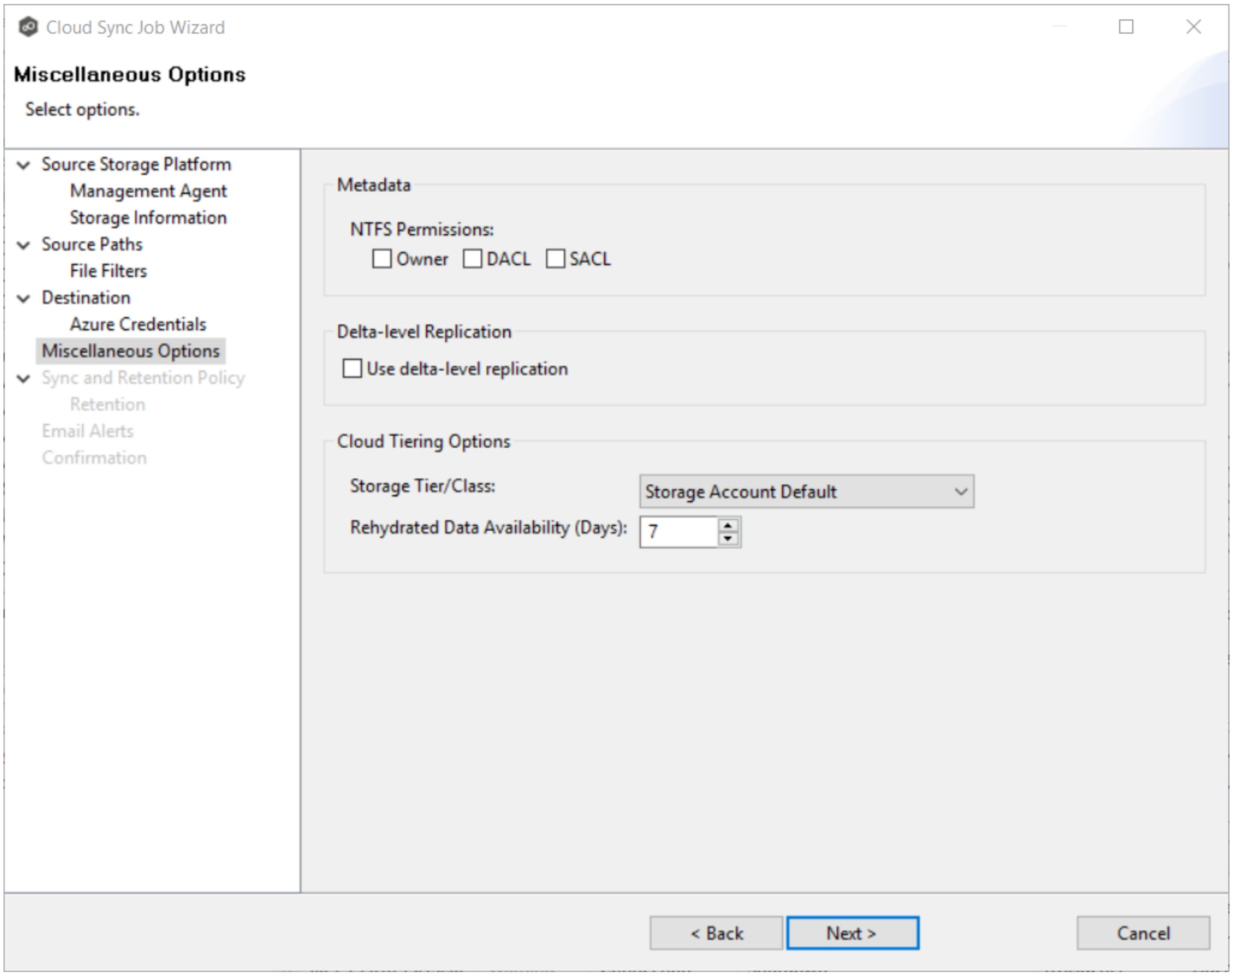

Cloud Sync Job Wizard

Miscellaneous Options

Select options.

Metadata

NTFS Permissions:

□ Owner □ DACL □ SACL

Delta-level Replication

□ Use delta-level replication

Cloud Tiering Options

Storage Tier/Class: Storage Account Default

Rehydrated Data Availability (Days): 7

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.
Delta-level Replication	If you want delta-level replication, select this check box. Delta-level replication enables Cloud Sync 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 over a slow link.
Storage Tier/Class	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 Sync 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 Sync automatically returns it to archive storage.
Rehydrated Data Availability (Days)	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**.

Step 10: Sync and Retention Policy

When Cloud Sync scans the source storage device or is notified of user activity on the source storage device, it replicates changed files in the watch set to the destination storage device. Cloud Sync can also take a snapshot of the watch set when replicating. A **snapshot** captures the state of a file system at a point in time. Snapshots are useful for backing up data at different intervals, which allows information to be recovered from different periods of time. (For more information about recovering data, see Recovering Data from the Cloud.)

Each Cloud Sync job must have a Sync and Replication policy. A Sync and Replication 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 create snapshots of the replicated data.
- How long you want to retain the snapshots.

The **Sync and Replication Policy** page enables you to create a new Sync and Replication policy or choose an existing policy.

1. Select **New Policy** or **Existing Policy**.

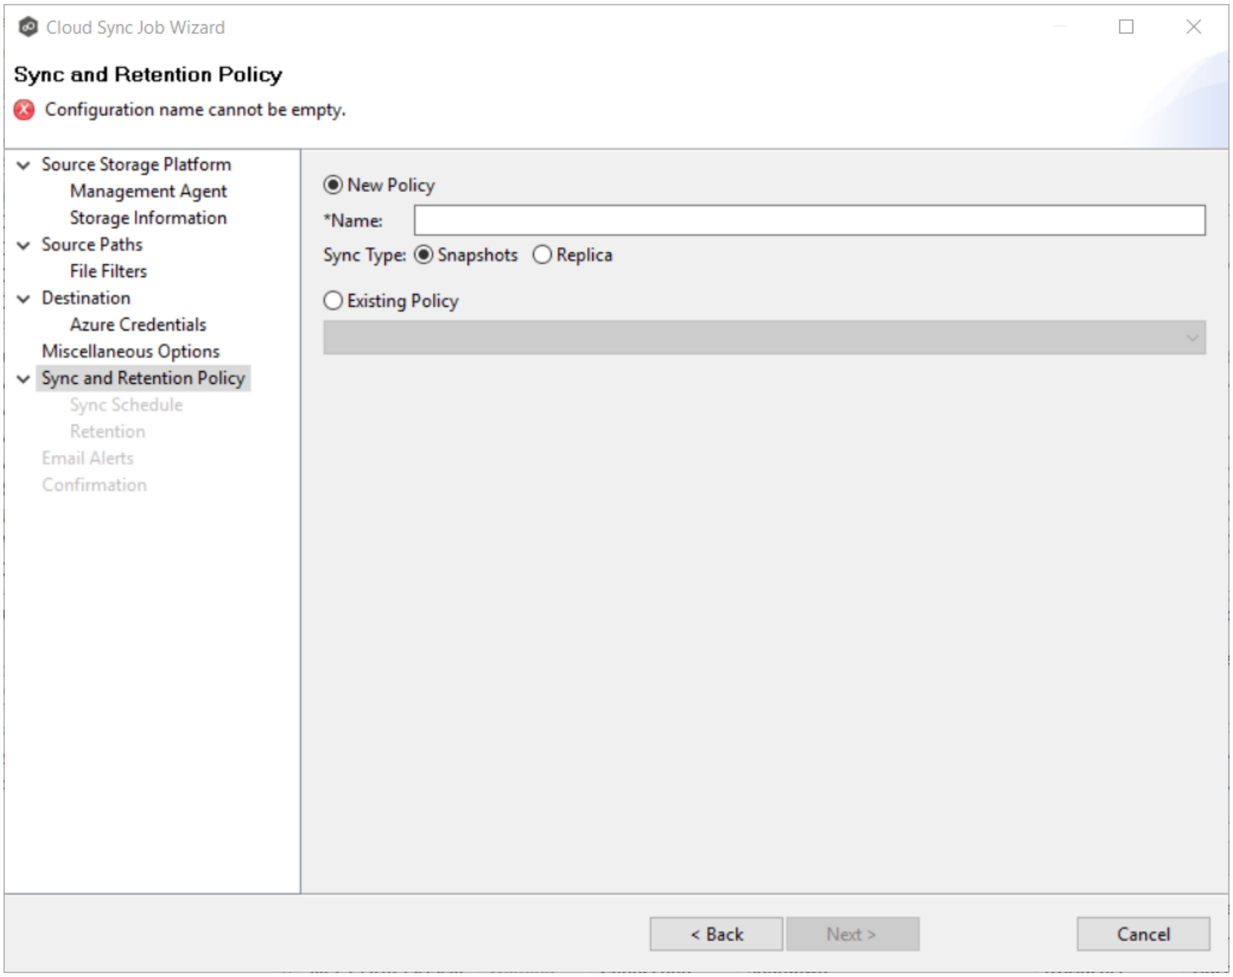
Cloud Sync Job Wizard

Sync and Retention Policy

Configuration name cannot be empty.

● New Policy

*Name:

Sync Type: ● Snapshots ○ Replica

○ Existing Policy

2. If you selected **Existing Policy**, select a policy from the drop-down list, and then click **Next**. Continue with Step 14. Email Alerts.

   If you selected **New Policy**, enter a name for the policy in the **Name** field.

3. Select the sync type:

   - Select **Snapshots** if you want to replicate what is on premises to the destination cloud storage platform, and, in addition to replication, you want to keep versions of changed files and take snapshots of the watch set at specific points in time.

   - Select **Replica** if you want to replicate what is on premises to the destination cloud storage platform and do not want to save versions of changed files or take snapshots.

Step 11: Sync Schedule

The **Sync Schedule** page enables you to select the frequency of the replication and when snapshots should be performed (if you selected **Snapshots** as your sync option). You can choose replication to be performed on a scheduled basis or a continuous, real-time basis.

1. Select the frequency of the replication:

   - Scheduled Scans – Replication is scheduled on a daily or weekly basis.

   - Continuous Data Protection – Replication occurs in real-time: whenever a file changes, the change is replicated.

If you selected **Scheduled Scans** for the replication frequency:

1. Select **Scan at Start** 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.

  Then, if you chose **Snapshots** as the sync type, choose when snapshots are taken (you must take at least one snapshot). If you chose **Replica** as the sync type, the **Trigger Snapshot** options will not appear.

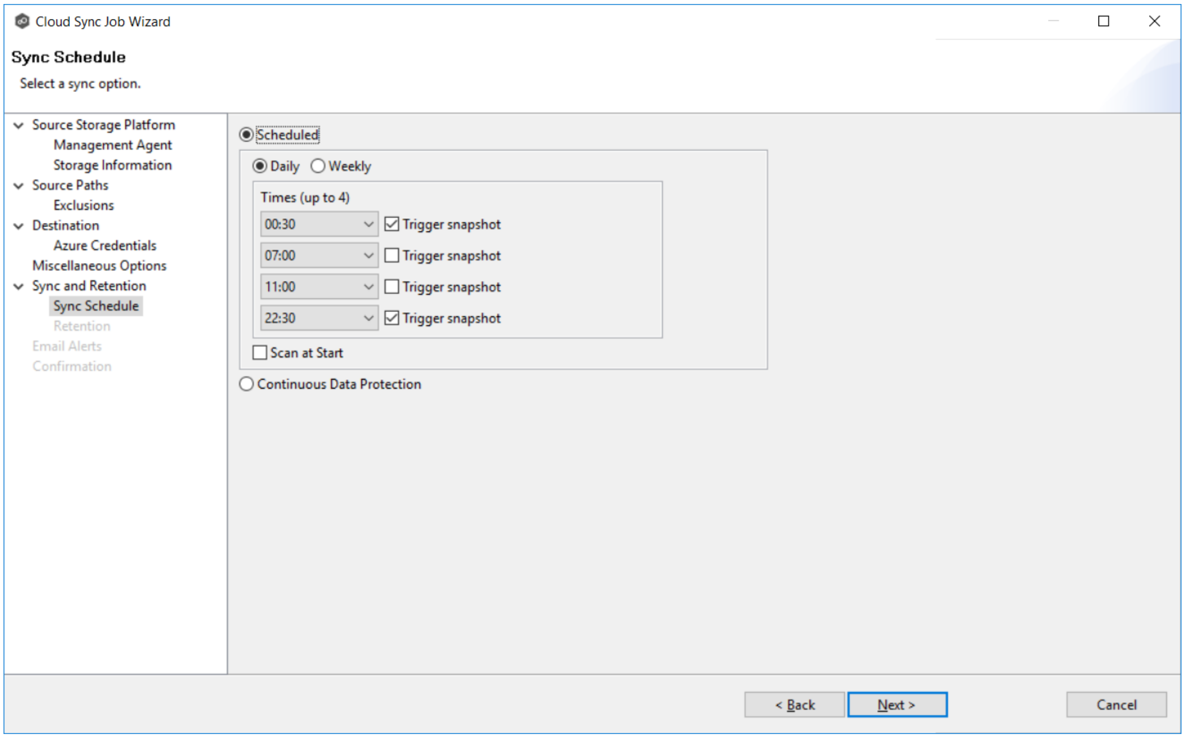

Sync Schedule

Select a sync option.

● Scheduled

● Daily ○ Weekly

Times (up to 4)

00:30	☑ Trigger snapshot
07:00	□ Trigger snapshot
11:00	□ Trigger snapshot
22:30	☑ Trigger snapshot

□ Scan at Start

○ Continuous Data Protection

- Select **Weekly** if you want to select specific days for replication. However, you can schedule only one scan per day.

  Then, if you chose **Snapshots** as the sync type, click the **Trigger** snapshot check box. A snapshot will be taken at the scan time. If you chose **Replica** as the sync type, the **Trigger Snapshot** option will not appear.

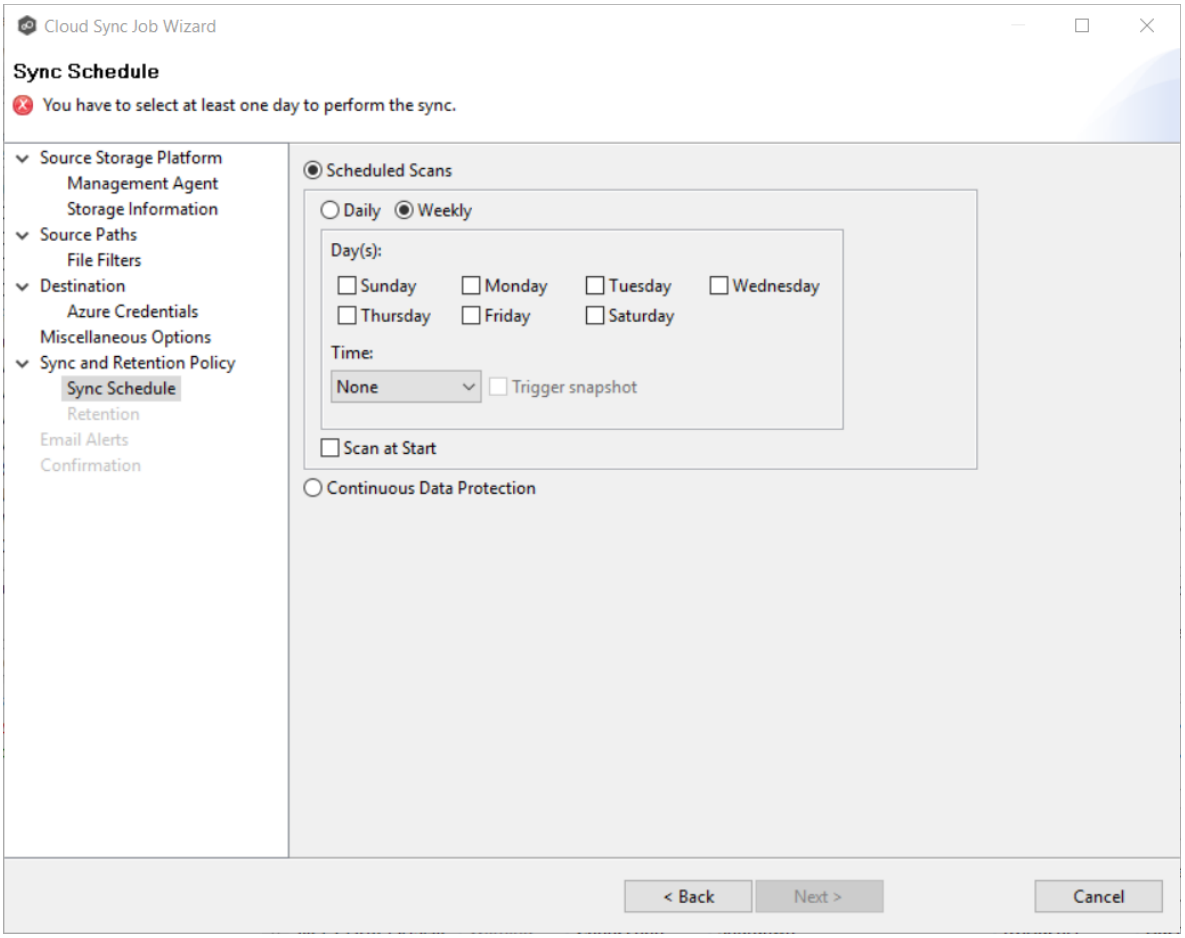

Sync Schedule

You have to select at least one day to perform the sync.

● Scheduled Scans

○ Daily ● Weekly

Day(s):

□ Sunday □ Monday □ Tuesday □ Wednesday
□ Thursday □ Friday □ Saturday

Time: None □ Trigger snapshot

□ Scan at Start

○ Continuous Data Protection

3. Click **Next** and continue with Step 12.

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.

2. If you chose **Snapshots** as the sync type, choose when snapshots are taken (you must take at least one snapshot). If you chose **Replica** as the sync type, the **Trigger Snapshot** options will not appear.

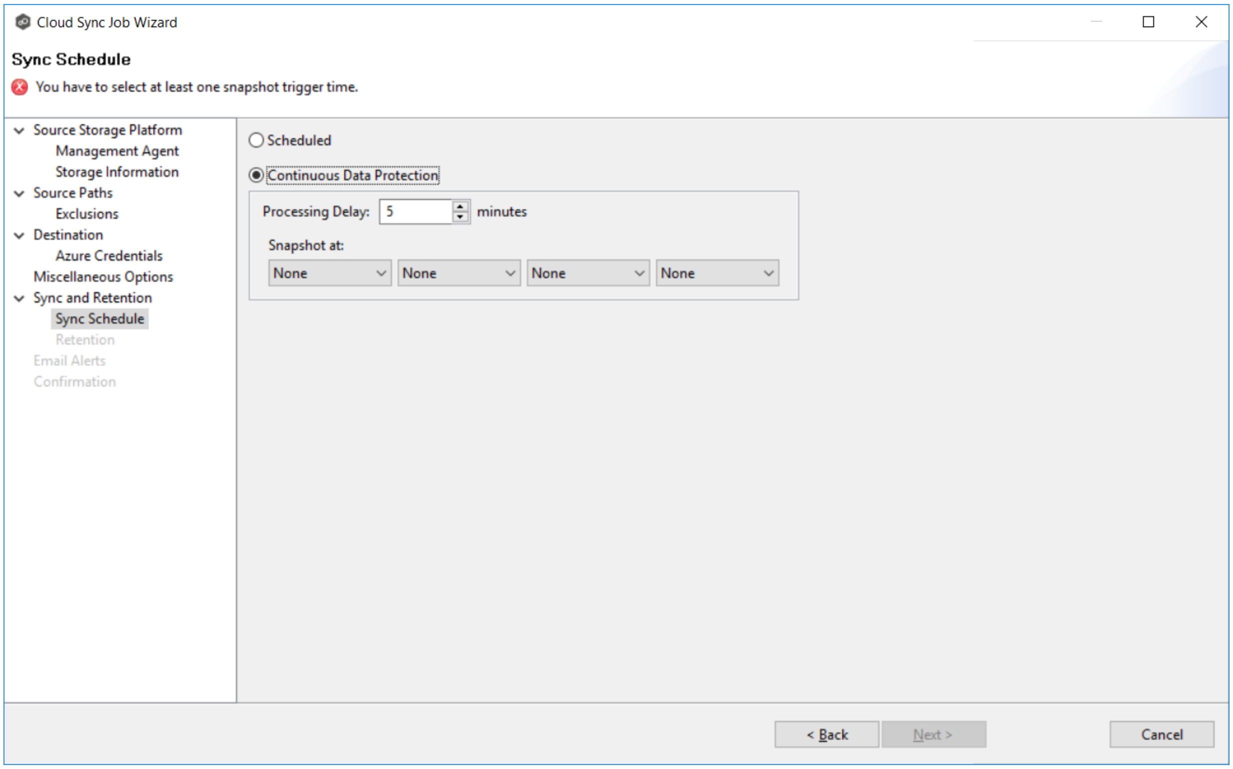

Cloud Sync Job Wizard

Sync Schedule

You have to select at least one snapshot trigger time.

○ Scheduled

● Continuous Data Protection

Processing Delay: 5 minutes

Snapshot at: None None None None

3. Click **Next** and continue with Step 12.

Step 12: Retention

The **Retention** page enables you to define how long you want to retain snapshots. You have the option to retain snapshots on a daily, weekly, monthly, and yearly basis. If you selected **Replica** for the sync type, the **Retention** page will not appear.

1. Select the **Purge all versions between snapshots** check box if you do not want to indefinitely retain all versions.

2. Select the retention options.

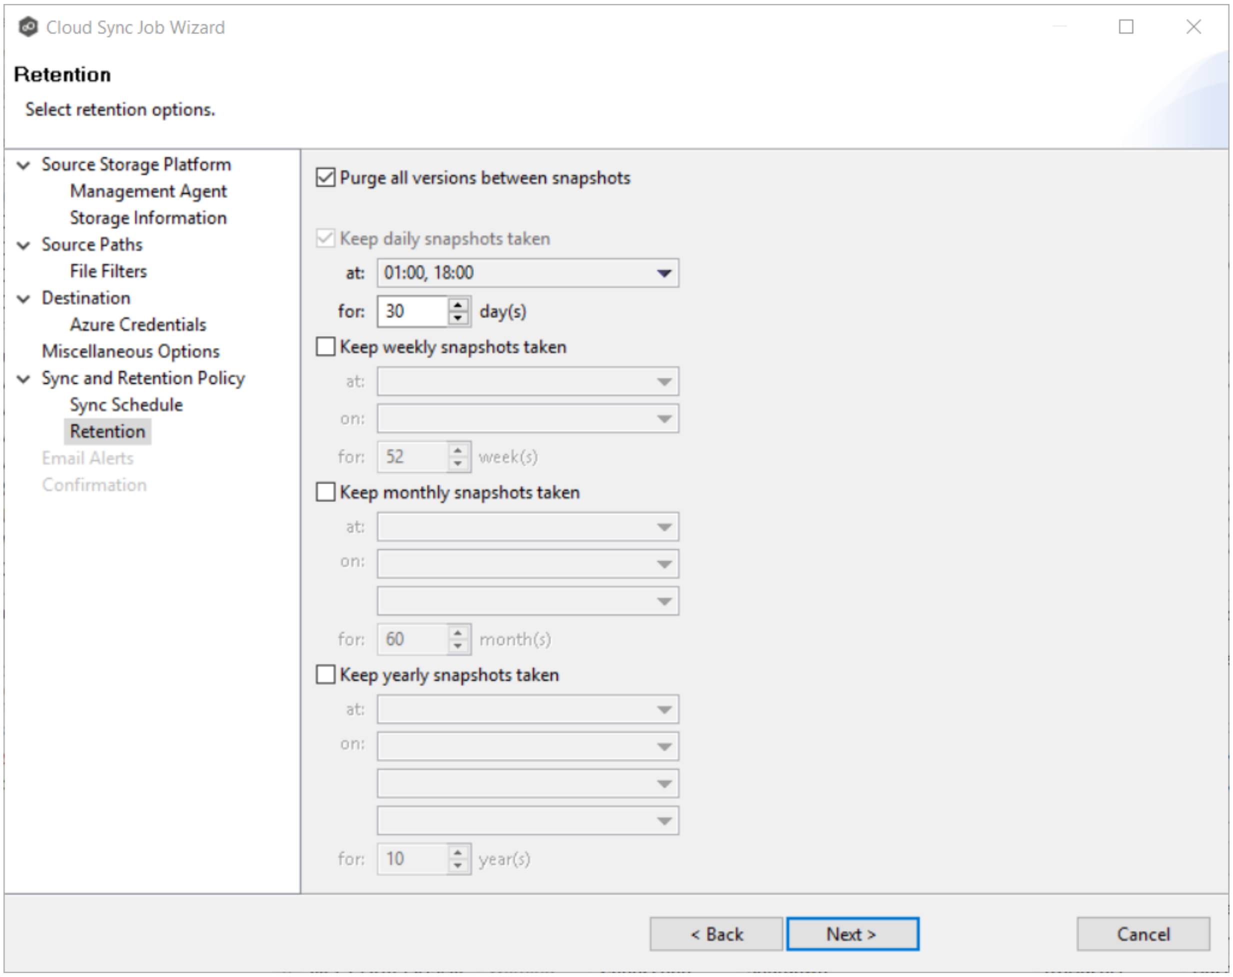

Retention

Select retention options.

☑ Purge all versions between snapshots

☑ Keep daily snapshots taken

at: 01:00, 18:00

for: 30 day(s)

□ Keep weekly snapshots taken

at:

on:

for: 52 week(s)

□ Keep monthly snapshots taken

at:

on:

for: 60 month(s)

□ Keep yearly snapshots taken

at:

on:

for: 10 year(s)

3. Click **Next**.

Step 13: Email Alerts

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 apply an existing email alert to the job:

1. Click the **Select** button.

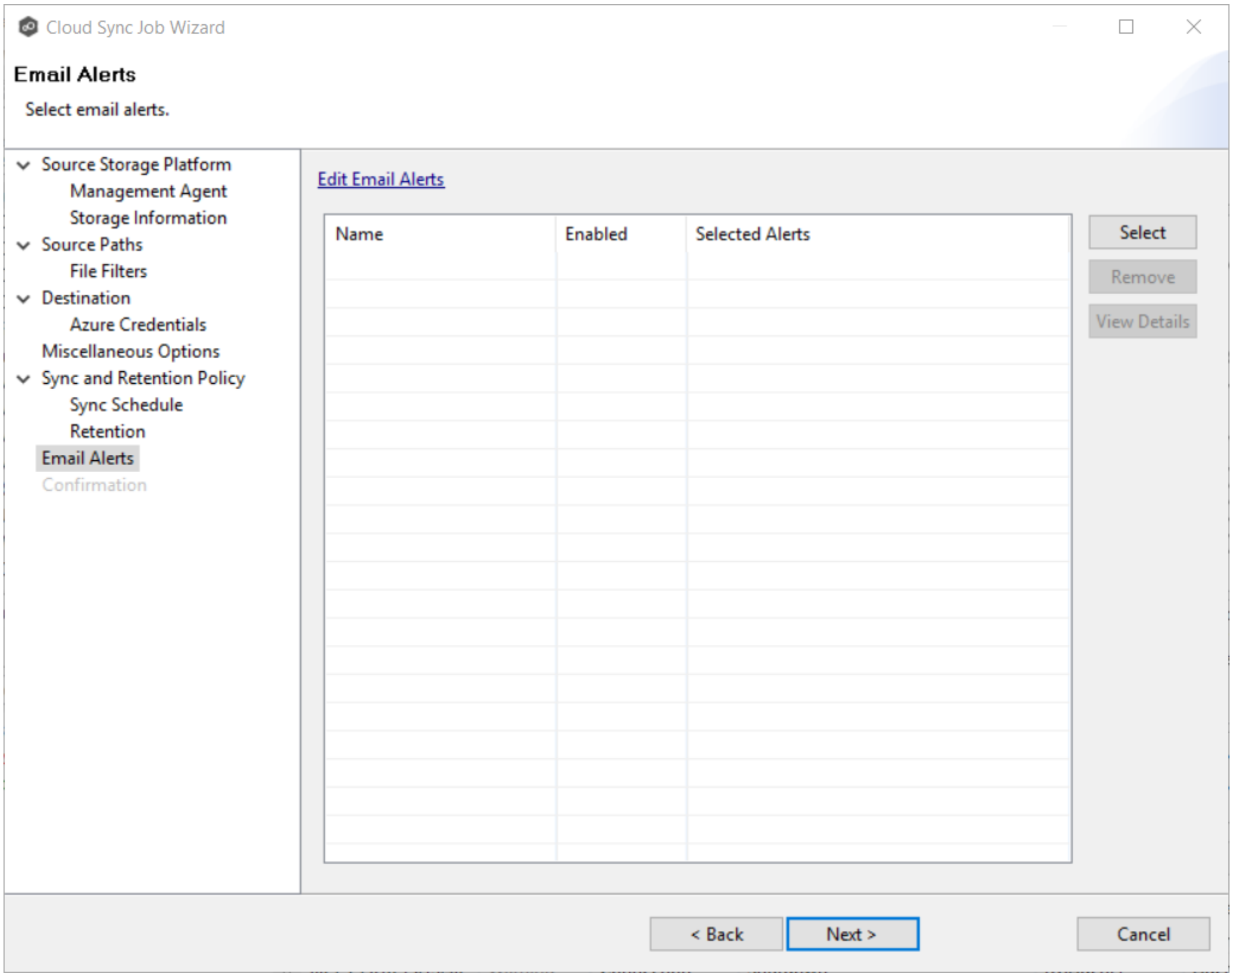

Email Alerts

Select email alerts.

Name	Enabled	Selected Alerts

The **Select Alert Configuration** dialog appears.

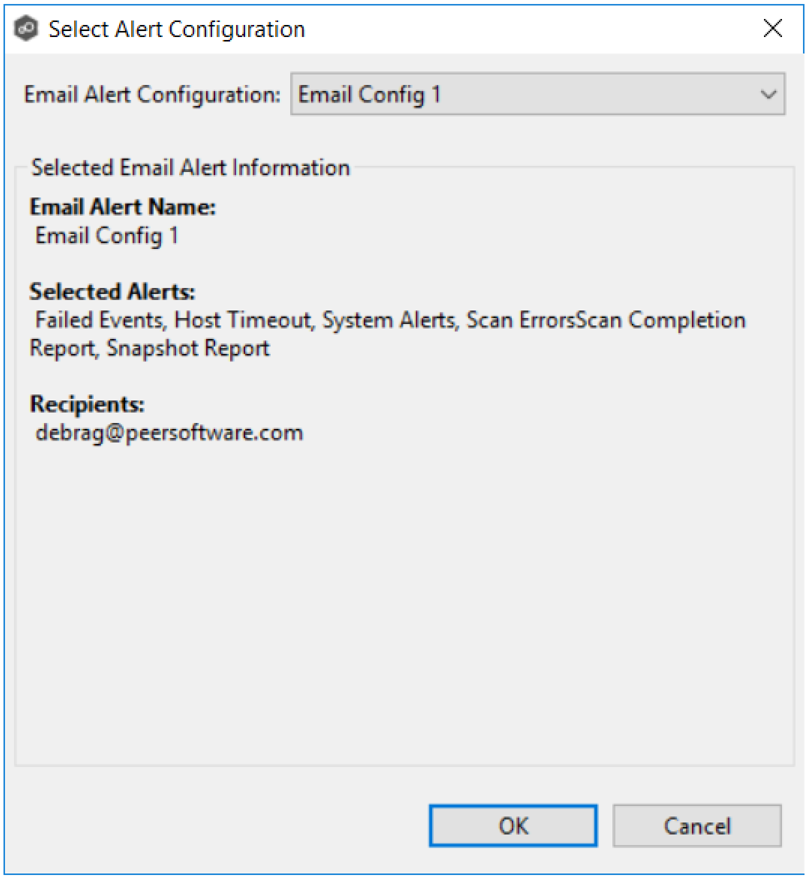

Email Alert Configuration: Email Config 1

Selected Email Alert Information

Email Alert Name:
Email Config 1

Selected Alerts:
Failed Events, Host Timeout, System Alerts, Scan ErrorsScan Completion Report, Snapshot Report

Recipients:
debrag@peersoftware.com

2. From the **Email Alert Configuration** drop-down list, select the email alert to apply to the job.

3. Click **OK**.

4. Repeat to apply additional alerts.

Step 14: 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.

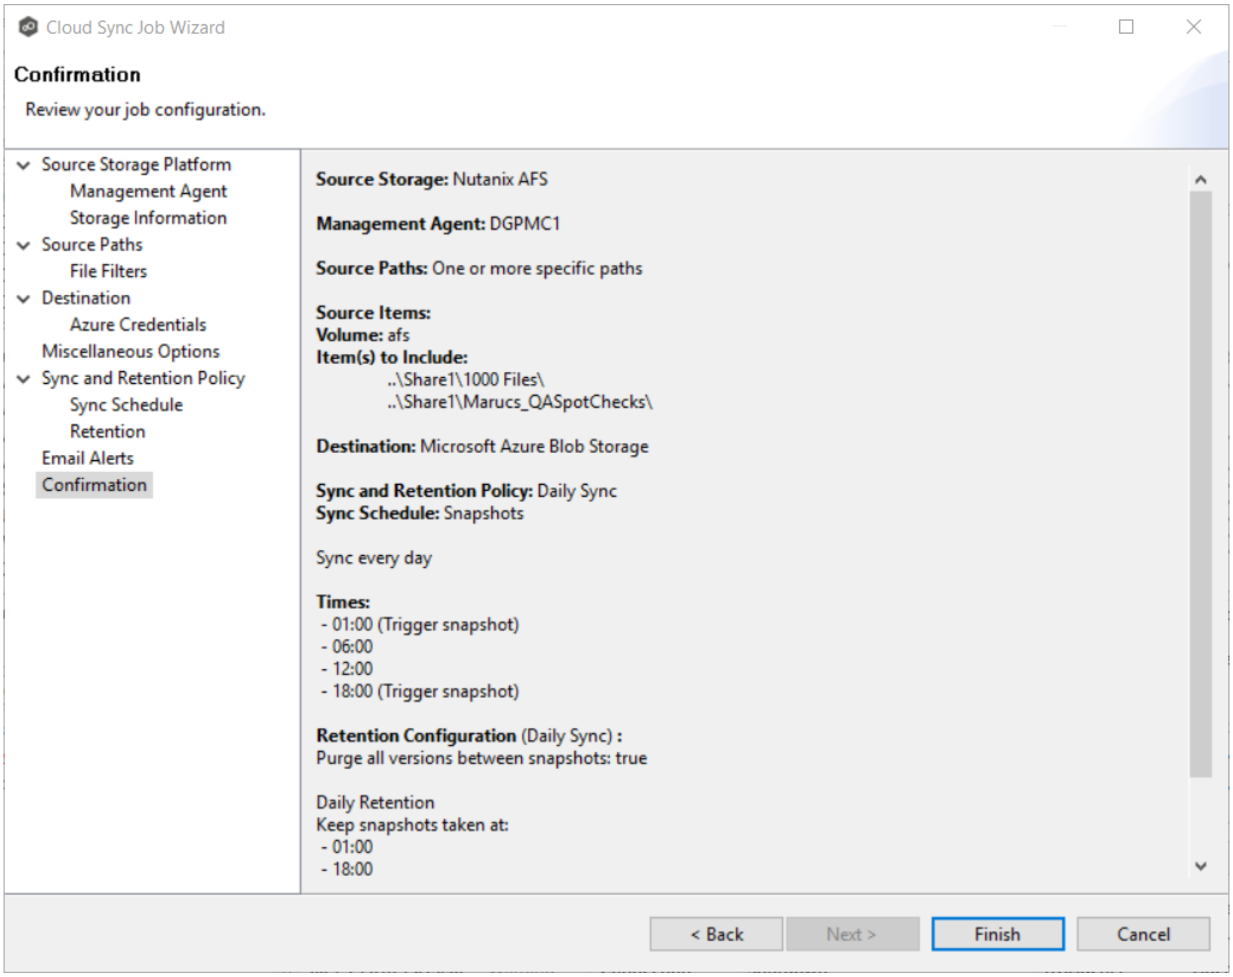

Source Storage: Nutanix AFS

Management Agent: DGPMC1

Source Paths: One or more specific paths

Source Items:
Volume: afs
Item(s) to Include:
..\Share1\1000 Files\
..\Share1\Marucs_QASpotChecks\

Destination: Microsoft Azure Blob Storage

Sync and Retention Policy: Daily Sync
Sync Schedule: Snapshots

Sync every day

Times:
- 01:00 (Trigger snapshot)
- 06:00
- 12:00
- 18:00 (Trigger snapshot)

Retention Configuration (Daily Sync) **:**
Purge all versions between snapshots: true

Daily Retention
Keep snapshots taken at:
- 01:00
- 18:00

3. Once satisfied, click the **Finish** button.

   The **Summary** tab in the **Cloud Sync Job** view is displayed.

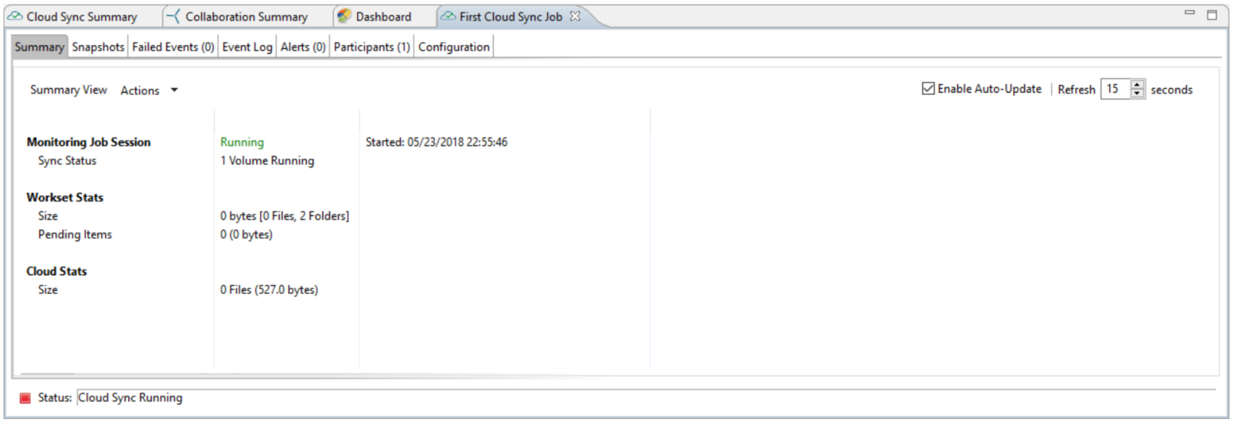

Monitoring Job Session	Running	Started: 05/23/2018 22:55:46
Sync Status	1 Volume Running	
Workset Stats		
Size	0 bytes [0 Files, 2 Folders]	
Pending Items	0 (0 bytes)	
Cloud Stats		
Size	0 Files (527.0 bytes)	

Status: Cloud Sync Running

Congratulations! Now you are ready to start running the job. See Starting a Cloud Sync Job for details.

Running a Cloud Sync Job

This section describes:

- Starting a Cloud Sync Job
- Stopping a Cloud Sync Job

Starting a Cloud Sync Job

When running a Cloud Sync 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 **Cloud Sync Job Summary** view, and then choose **Start** from the pop-up menu.
   - Open the job and then click the **Start/Stop** button in the bottom left corner of the job's **Summary** tab (shown below).

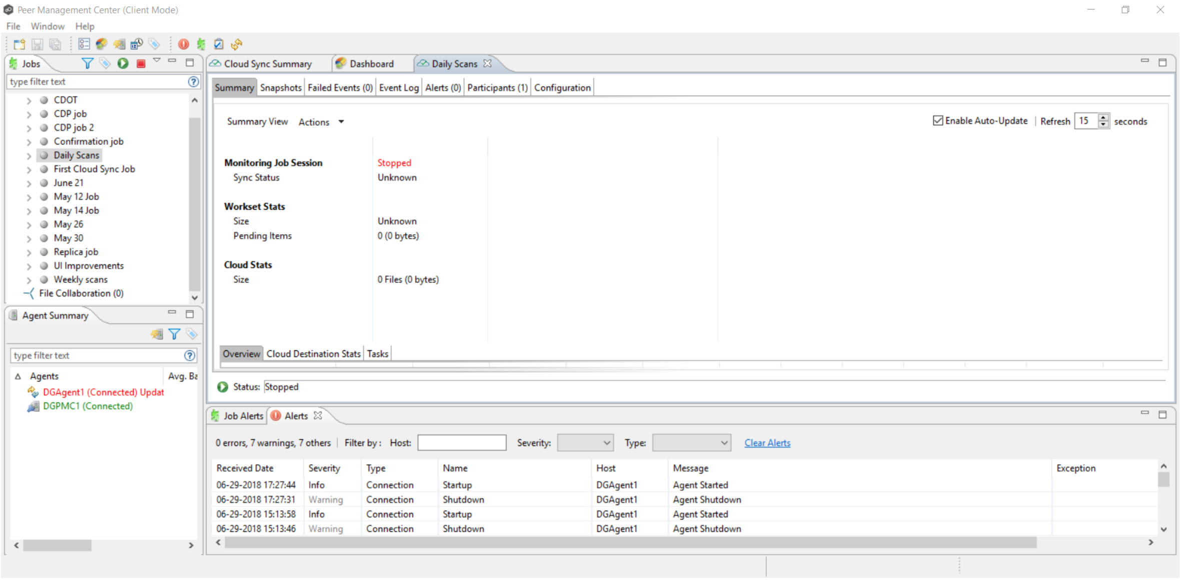

☑ Enable Auto-Update | Refresh 15 seconds

Monitoring Job Session Stopped
Sync Status Unknown

Workset Stats
Size Unknown
Pending Items 0 (0 bytes)

Cloud Stats
Size 0 Files (0 bytes)

Status: Stopped

0 errors, 7 warnings, 7 others | Filter by : Host: Severity: Type:

Received Date	Severity	Type	Name	Host	Message	Exception
06-29-2018 17:27:44	Info	Connection	Startup	DGAgent1	Agent Started	
06-29-2018 17:27:31	Warning	Connection	Shutdown	DGAgent1	Agent Shutdown	
06-29-2018 15:13:58	Info	Connection	Startup	DGAgent1	Agent Started	
06-29-2018 15:13:46	Warning	Connection	Shutdown	DGAgent1	Agent Shutdown	

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.

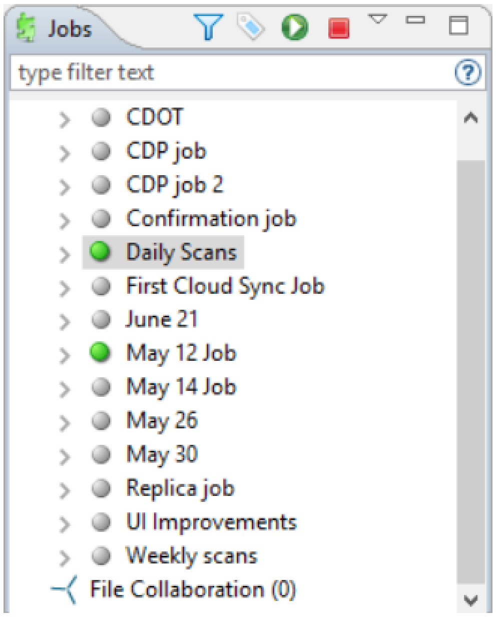

Stopping a Cloud Sync Job

You can stop a Cloud Sync job at any time.

To stop a Cloud Sync job:

1. Right-click the job name in the **Jobs** view or in the **Cloud Sync Job Summary** view, and then choose **Stop** from the pop-up menu.

   Or, open the job and click the **Start/Stop** button in the bottom left corner of the job's **Summary** tab.

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 Sync Jobs

Monitoring your Cloud Sync 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 Sync provides several views to help you monitor the health and performance of your Cloud Sync 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 Sync** in the **Jobs** view to display information about all Cloud Sync jobs. The **Cloud Sync Volume Summary** tab of the **Cloud Sync Summary** view is displayed.

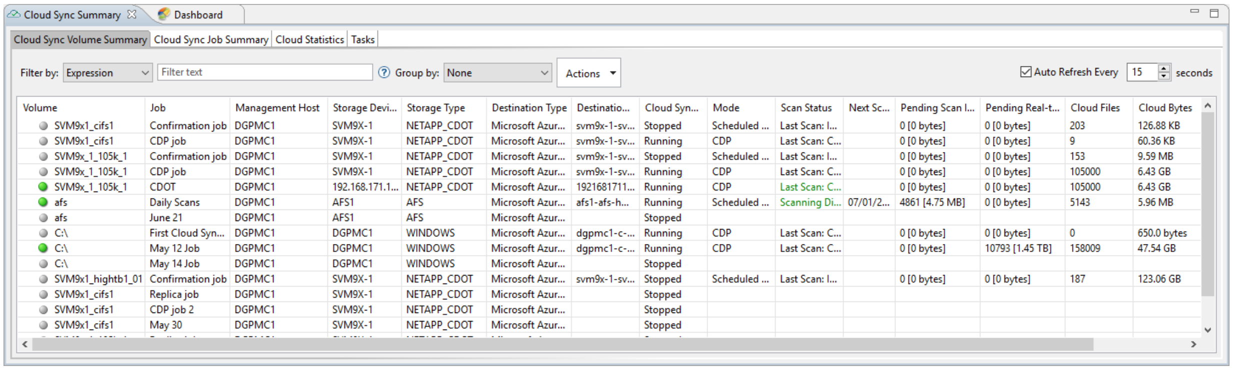

Filter by: Expression | Filter text | Group by: None | ☑ Auto Refresh Every 15 seconds

Volume	Job	Management Host	Storage Devi...	Storage Type	Destination Type	Destinatio...	Cloud Syn...	Mode	Scan Status	Next Sc...	Pending Scan I...	Pending Real-t...	Cloud Files	Cloud Bytes
SVM9x1_cifs1	Confirmation job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Stopped	Scheduled ...	Last Scan: I...		0 [0 bytes]	0 [0 bytes]	203	126.88 KB
SVM9x1_cifs1	CDP job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Running	CDP	Last Scan: C...		0 [0 bytes]	0 [0 bytes]	9	60.36 KB
SVM9x_1_105k_1	Confirmation job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Stopped	Scheduled ...	Last Scan: I...		0 [0 bytes]	0 [0 bytes]	153	9.59 MB
SVM9x_1_105k_1	CDP job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Running	CDP	Last Scan: C...		0 [0 bytes]	0 [0 bytes]	105000	6.43 GB
SVM9x_1_105k_1	CDOT	DGPMC1	192.168.171.1...	NETAPP_CDOT	Microsoft Azur...	1921681711...	Running	CDP	Last Scan: C...		0 [0 bytes]	0 [0 bytes]	105000	6.43 GB
afs	Daily Scans	DGPMC1	AFS1	AFS	Microsoft Azur...	afs1-afs-h...	Running	Scheduled ...	Scanning Di...	07/01/2...	4861 [4.75 MB]	0 [0 bytes]	5143	5.96 MB
afs	June 21	DGPMC1	AFS1	AFS	Microsoft Azur...		Stopped							
C:\	First Cloud Syn...	DGPMC1	DGPMC1	WINDOWS	Microsoft Azur...	dgpmc1-c-...	Running	CDP	Last Scan: C...		0 [0 bytes]	0 [0 bytes]	0	650.0 bytes
C:\	May 12 Job	DGPMC1	DGPMC1	WINDOWS	Microsoft Azur...	dgpmc1-c-...	Running	CDP	Last Scan: C...		0 [0 bytes]	10793 [1.45 TB]	158009	47.54 GB
C:\	May 14 Job	DGPMC1	DGPMC1	WINDOWS	Microsoft Azur...		Stopped							
SVM9x1_hightb1_01	Confirmation job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Stopped	Scheduled ...	Last Scan: I...		0 [0 bytes]	0 [0 bytes]	187	123.06 GB
SVM9x1_cifs1	Replica job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...		Stopped							
SVM9x1_cifs1	CDP job 2	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...		Stopped							
SVM9x1_cifs1	May 30	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...		Stopped							

- Double-click a job name in the **Jobs** view to display the views associated with a job. The **Summary** tab of the **Cloud Sync Job** view is displayed.

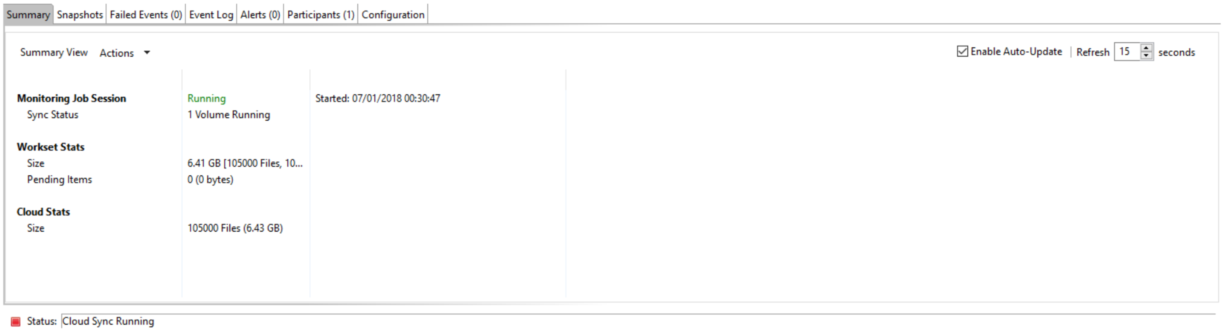

Refresh 15 seconds

Monitoring Job Session Running Started: 07/01/2018 00:30:47
Sync Status 1 Volume Running

Workset Stats
Size 6.41 GB [105000 Files, 10...
Pending Items 0 (0 bytes)

Cloud Stats
Size 105000 Files (6.43 GB)

Status: Cloud Sync Running

Deleting a Cloud Sync Job

To delete a Cloud Sync job:

1. Right-click on the job name in the **Jobs** view, and then choose **Delete** from the menu. A confirmation dialog appears.

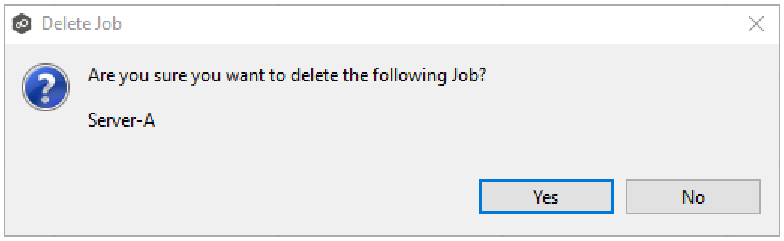

Are you sure you want to delete the following Job?

Server-A

2. Click **OK** in the confirmation dialog.

   Another dialog appears, prompting you to choose whether to delete data associated with the job.

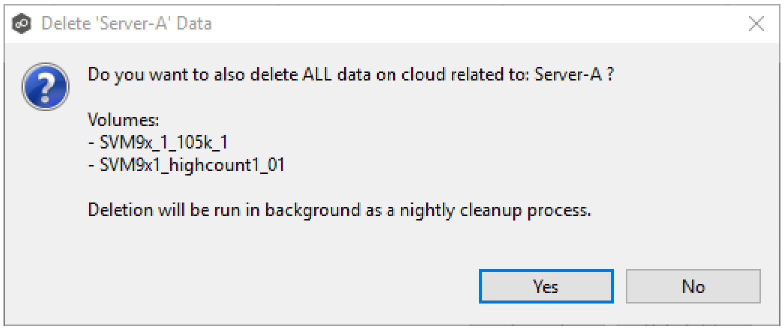

Do you want to also delete ALL data on cloud related to: Server-A ?

Volumes:
- SVM9x_1_105k_1
- SVM9x1_highcount1_01

Deletion will be run in background as a nightly cleanup process.

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 from the Cloud

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 Sync job that has been replicating that data.

Note: You can recover data from a running job—unless you plan to restore the data to the original location. If so, you should stop the job first.

To recover data:

1. Open the Peer Management Center.

2. In the **Jobs** view, identify the Cloud Sync 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 platform.

4. Select the volume that was the source of the replicated data from the **Volume** drop-down list.

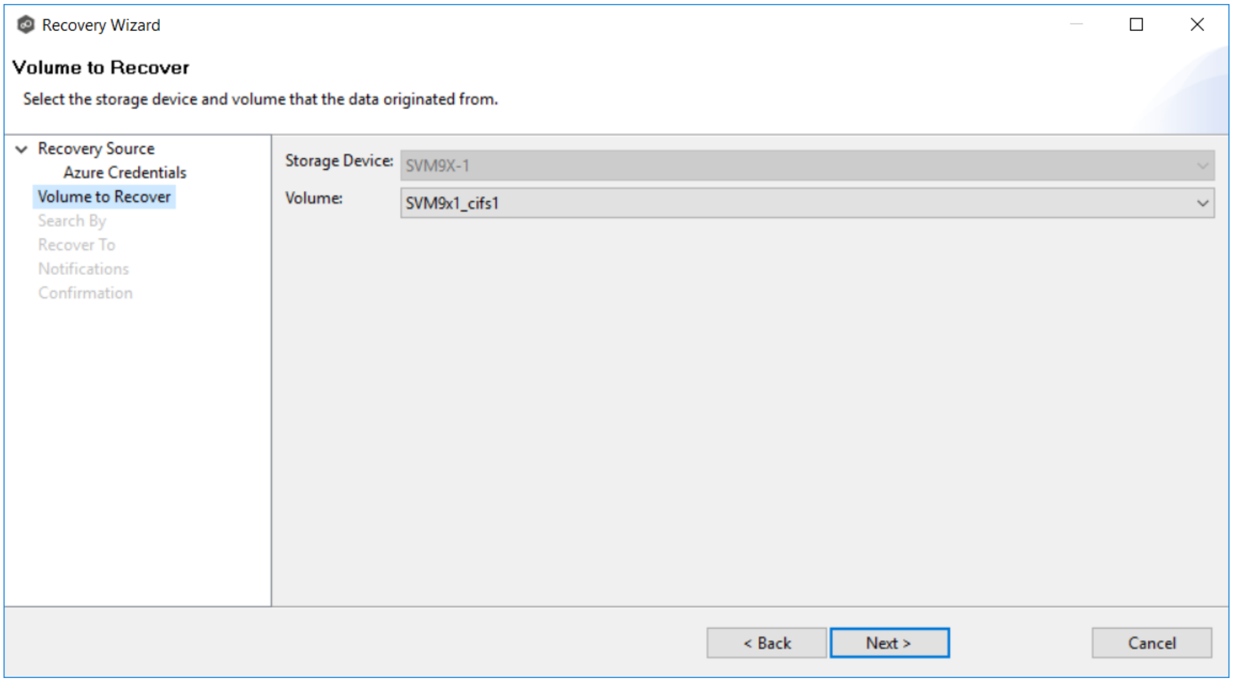

Recovery Wizard

Volume to Recover

Select the storage device and volume that the data originated from.

Storage Device: SVM9X-1

Volume: SVM9x1_cifs1

5. Click **Next**.

   The **Search By** page is displayed. It presents four search options:

   [Name]

   [Snapshot]

   [Point in Time]

   [Latest Replication]

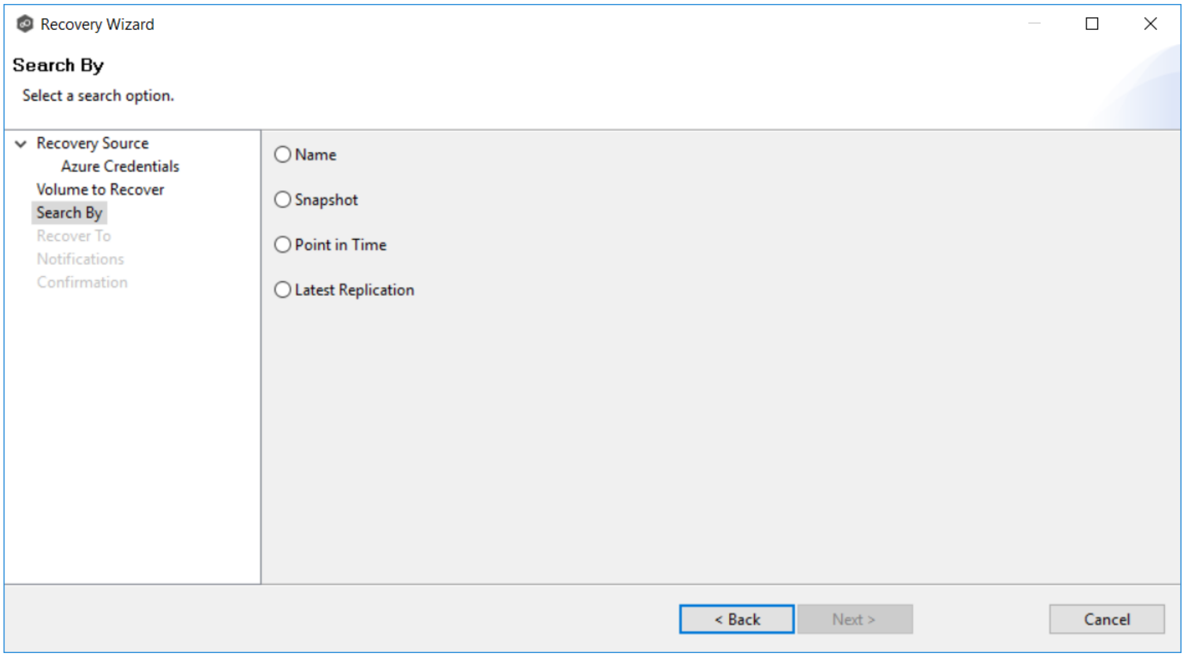

6. Select one of the search options.

7. Click **Next**.

   The search pages vary, depending on the search option you selected.

Search Options

The four 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.

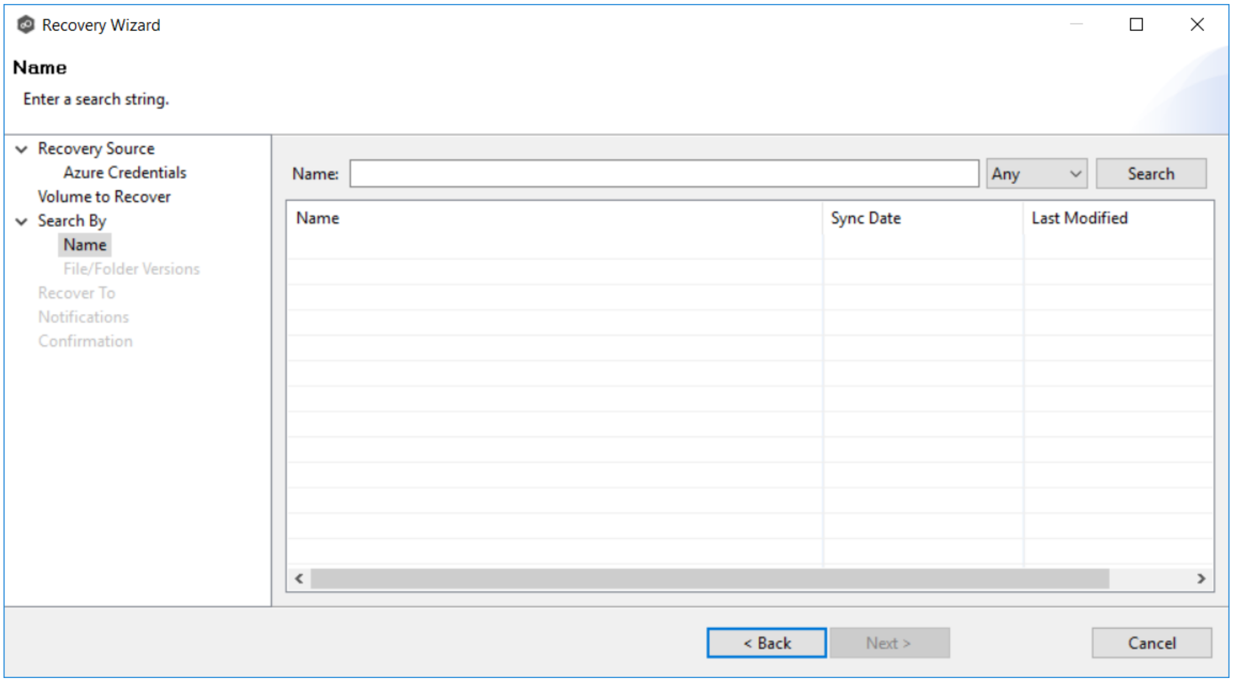

2. Select **File** or **Folder** from the **Any** drop-down list; if you want to search for both files and folders, select **Any**.

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.

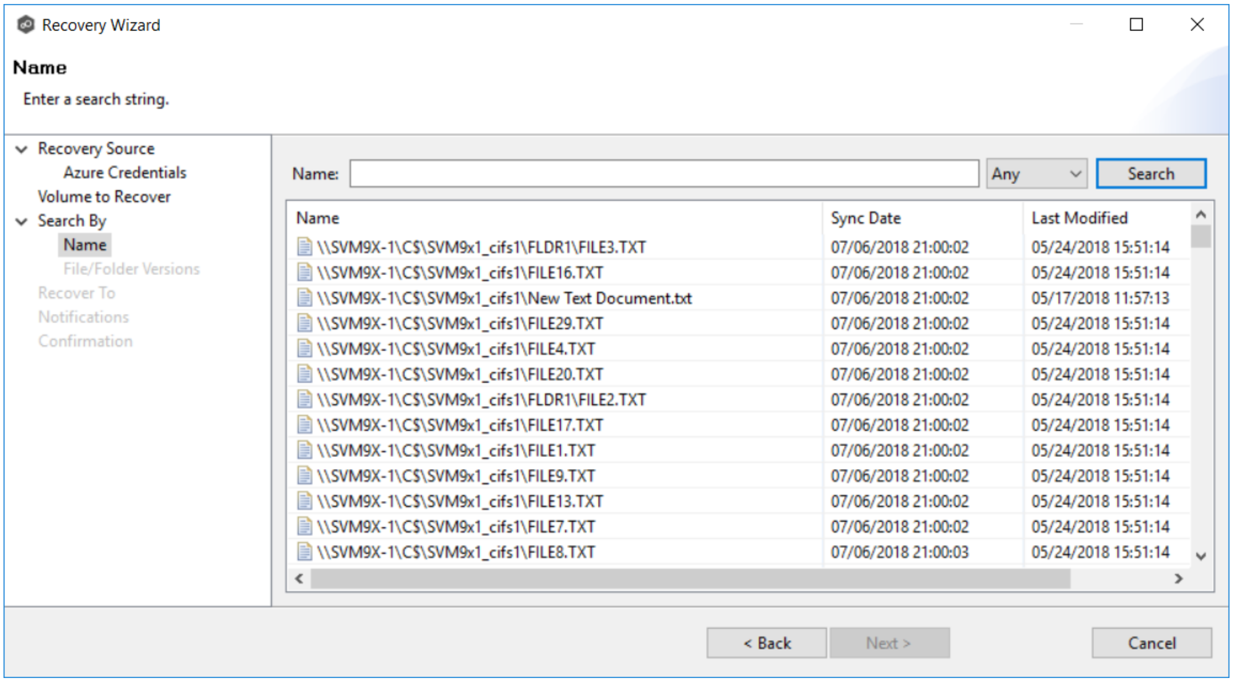

Name

Enter a search string.

Name:

Name	Sync Date	Last Modified
\\SVM9X-1\C$\SVM9x1_cifs1\FLDR1\FILE3.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FILE16.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\New Text Document.txt	07/06/2018 21:00:02	05/17/2018 11:57:13
\\SVM9X-1\C$\SVM9x1_cifs1\FILE29.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FILE4.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FILE20.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FLDR1\FILE2.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FILE17.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FILE1.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FILE9.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FILE13.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FILE7.TXT	07/06/2018 21:00:02	05/24/2018 15:51:14
\\SVM9X-1\C$\SVM9x1_cifs1\FILE8.TXT	07/06/2018 21:00:03	05/24/2018 15:51:14

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.

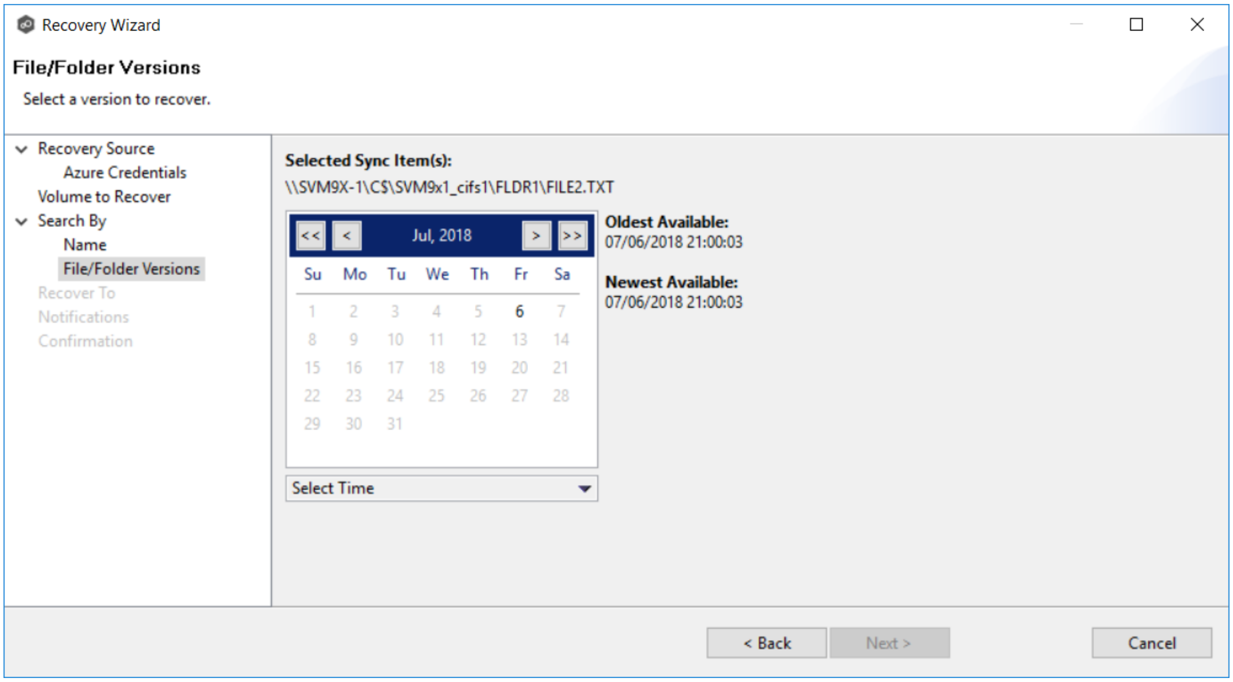

File/Folder Versions

Select a version to recover.

Selected Sync Item(s):
\\SVM9X-1\C$\SVM9x1_cifs1\FLDR1\FILE2.TXT

Oldest Available:
07/06/2018 21:00:03

Newest Available:
07/06/2018 21:00:03

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.

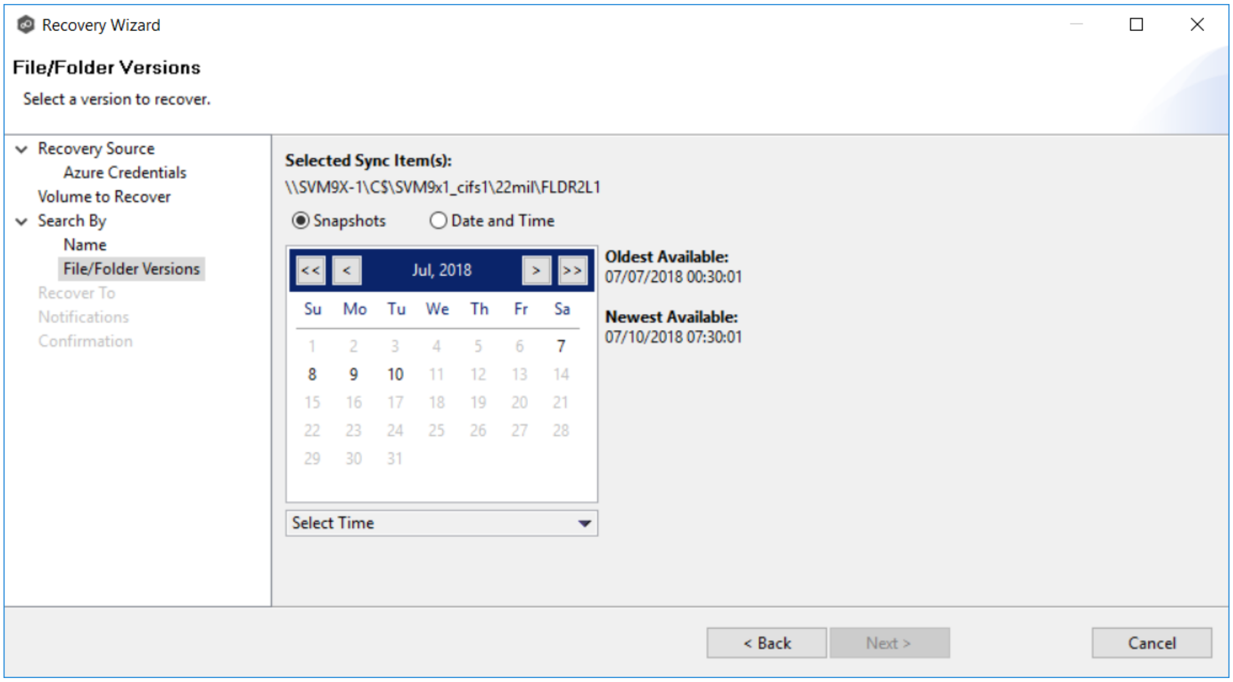

File/Folder Versions

Select a version to recover.

Selected Sync Item(s):
\\SVM9X-1\C$\SVM9x1_cifs1\22mil\FLDR2L1

● Snapshots ○ Date and Time

Oldest Available:
07/07/2018 00:30:01

Newest Available:
07/10/2018 07:30:01

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 snapshot. All available snapshots will be represented in the calendar widget below.

To search by snapshot:

1. Select the date of the snapshot.

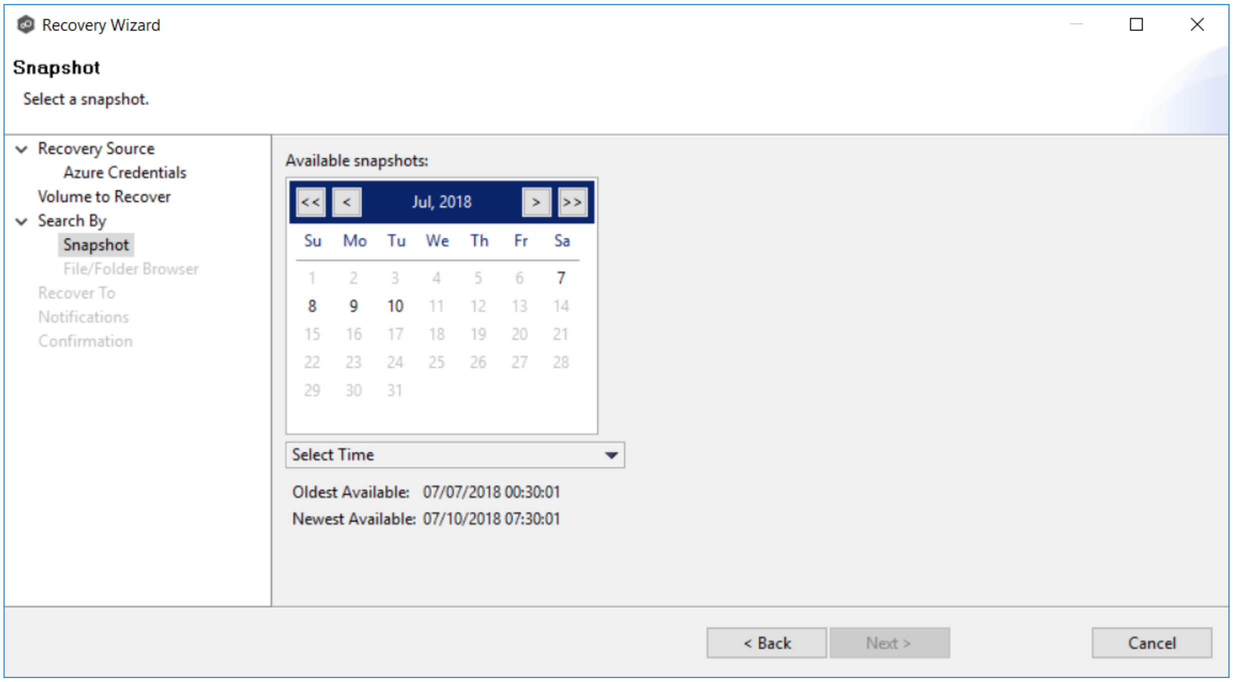

Snapshot

Select a snapshot.

Available snapshots:

Jul, 2018

Su	Mo	Tu	We	Th	Fr	Sa
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31				

Oldest Available: 07/07/2018 00:30:01

Newest Available: 07/10/2018 07:30:01

2. Select the time of the snapshot, and then click elsewhere in the page.

3. Click **Next**.

   The **File/Folder Browser** page appears.

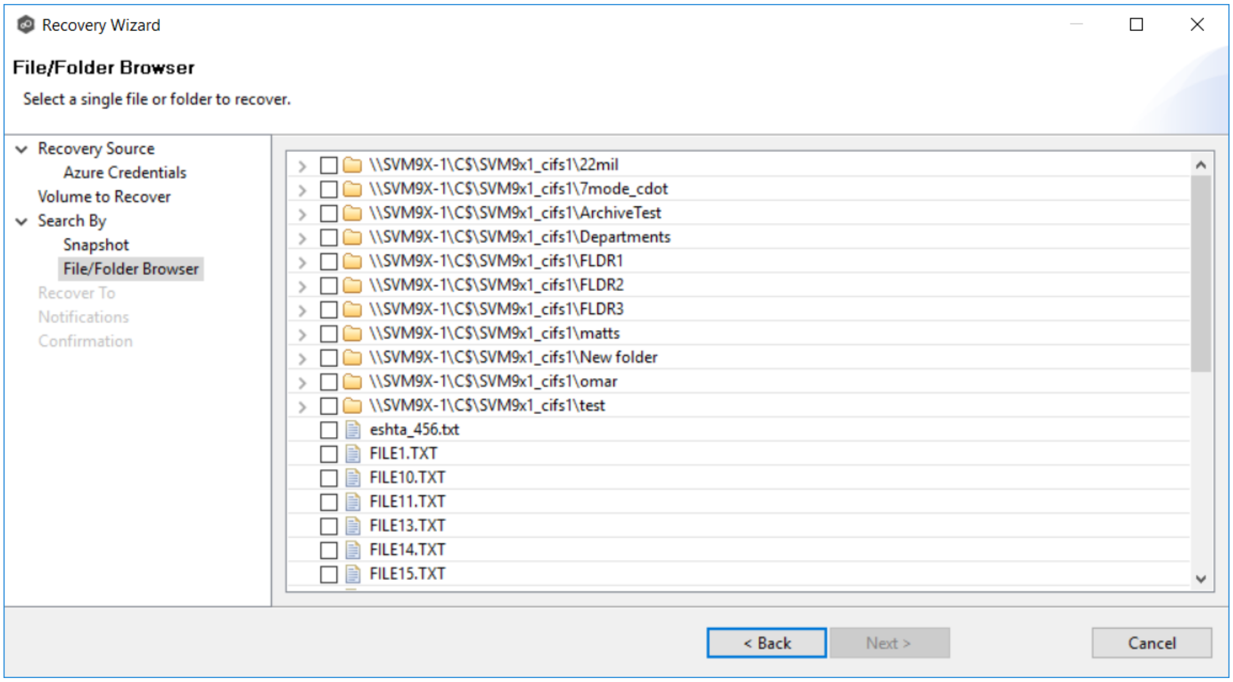

Recovery Wizard

File/Folder Browser

Select a single file or folder to recover.

- □ \\SVM9X-1\C$\SVM9x1_cifs1\22mil
- □ \\SVM9X-1\C$\SVM9x1_cifs1\7mode_cdot
- □ \\SVM9X-1\C$\SVM9x1_cifs1\ArchiveTest
- □ \\SVM9X-1\C$\SVM9x1_cifs1\Departments
- □ \\SVM9X-1\C$\SVM9x1_cifs1\FLDR1
- □ \\SVM9X-1\C$\SVM9x1_cifs1\FLDR2
- □ \\SVM9X-1\C$\SVM9x1_cifs1\FLDR3
- □ \\SVM9X-1\C$\SVM9x1_cifs1\matts
- □ \\SVM9X-1\C$\SVM9x1_cifs1\New folder
- □ \\SVM9X-1\C$\SVM9x1_cifs1\omar
- □ \\SVM9X-1\C$\SVM9x1_cifs1\test
- □ eshta_456.txt
- □ FILE1.TXT
- □ FILE10.TXT
- □ FILE11.TXT
- □ FILE13.TXT
- □ FILE14.TXT
- □ FILE15.TXT

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.

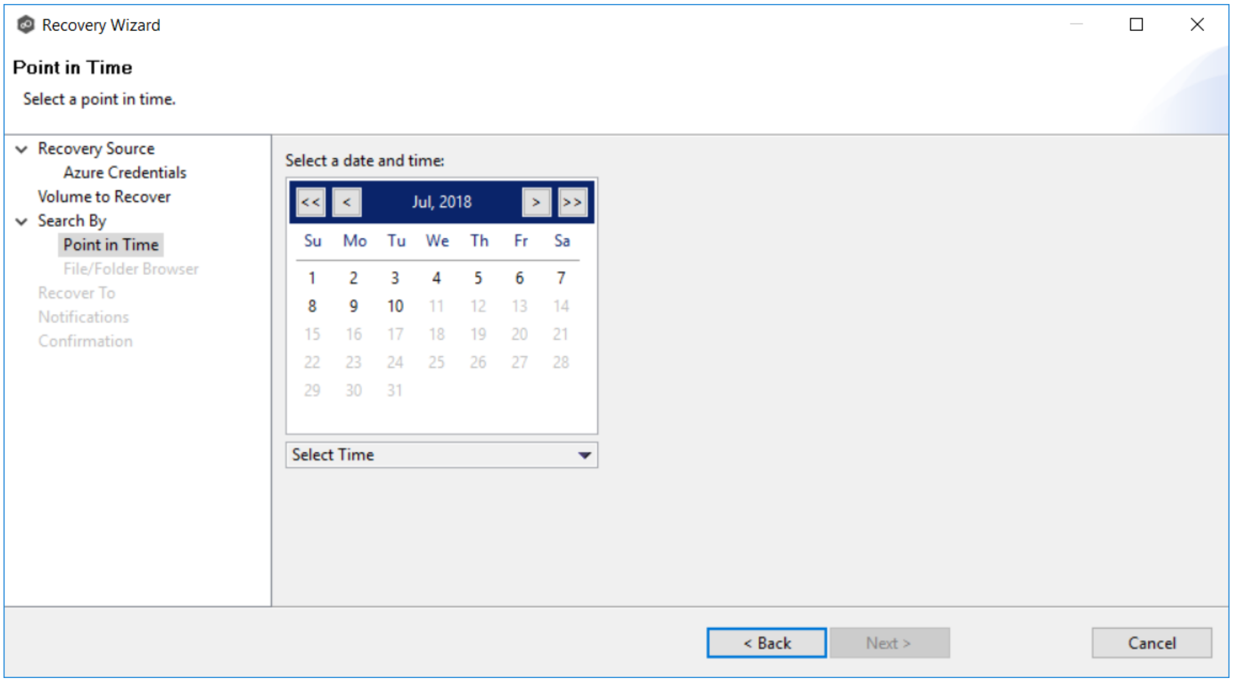

2. Select a date and time, and then click elsewhere in the page.

3. Click **Next**.

   The **File/Folder Browser** page appears.

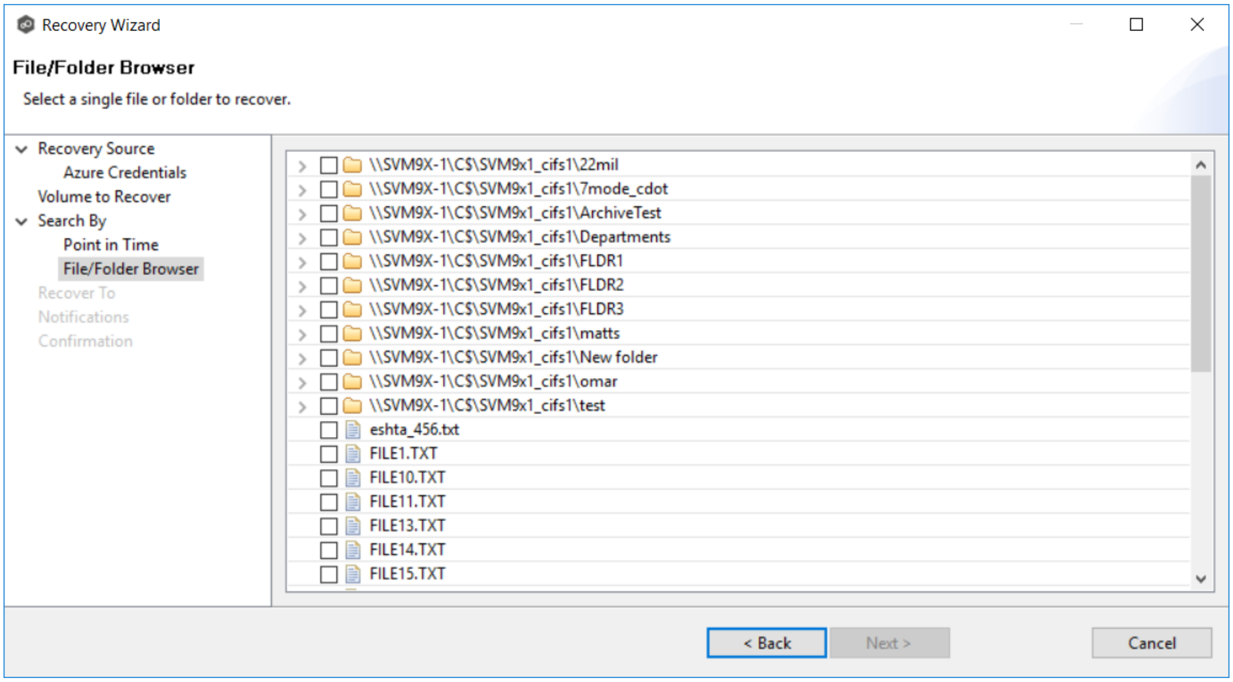
Recovery Wizard
File/Folder Browser
Select a single file or folder to recover.
Recovery Source
Azure Credentials
Volume to Recover
Search By
Point in Time
File/Folder Browser
Recover To
Notifications
Confirmation
\\SVM9X-1\C$\SVM9x1_cifs1\22mil
\\SVM9X-1\C$\SVM9x1_cifs1\7mode_cdot
\\SVM9X-1\C$\SVM9x1_cifs1\ArchiveTest
\\SVM9X-1\C$\SVM9x1_cifs1\Departments
\\SVM9X-1\C$\SVM9x1_cifs1\FLDR1
\\SVM9X-1\C$\SVM9x1_cifs1\FLDR2
\\SVM9X-1\C$\SVM9x1_cifs1\FLDR3
\\SVM9X-1\C$\SVM9x1_cifs1\matts
\\SVM9X-1\C$\SVM9x1_cifs1\New folder
\\SVM9X-1\C$\SVM9x1_cifs1\omar
\\SVM9X-1\C$\SVM9x1_cifs1\test
eshta_456.txt
FILE1.TXT
FILE10.TXT
FILE11.TXT
FILE13.TXT
FILE14.TXT
FILE15.TXT
< Back
Next >
Cancel

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.

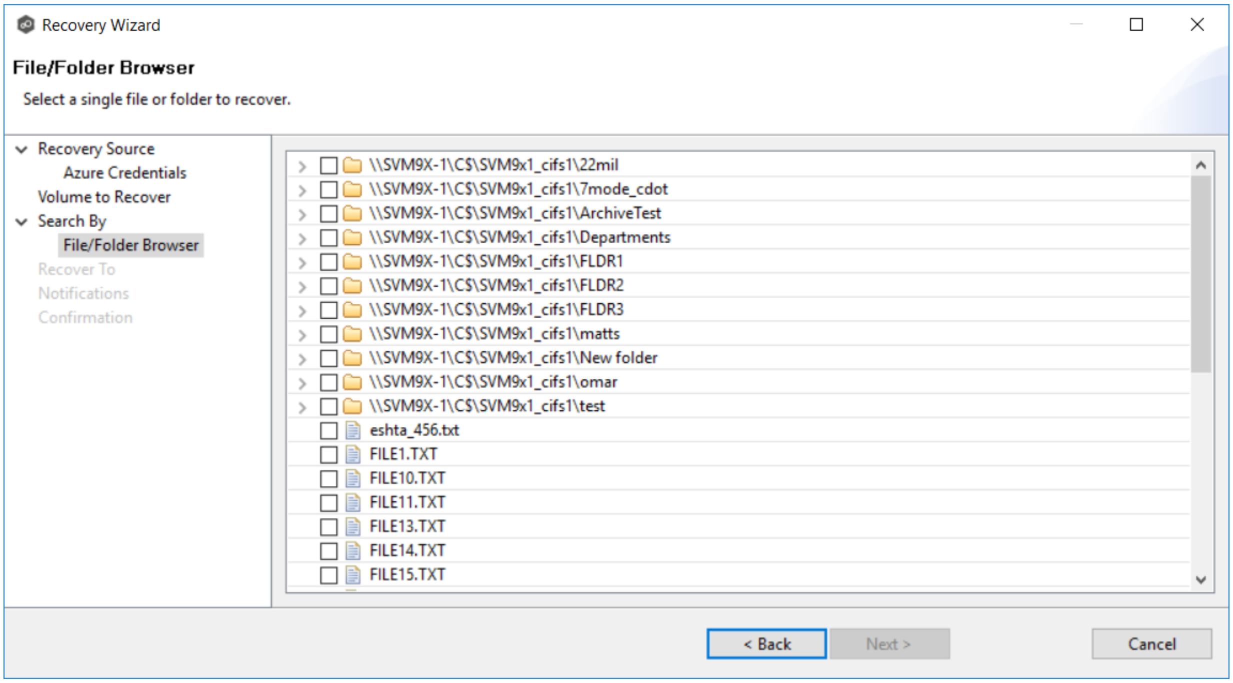

File/Folder Browser

Select a single file or folder to recover.

- □ \\SVM9X-1\C$\SVM9x1_cifs1\22mil
- □ \\SVM9X-1\C$\SVM9x1_cifs1\7mode_cdot
- □ \\SVM9X-1\C$\SVM9x1_cifs1\ArchiveTest
- □ \\SVM9X-1\C$\SVM9x1_cifs1\Departments
- □ \\SVM9X-1\C$\SVM9x1_cifs1\FLDR1
- □ \\SVM9X-1\C$\SVM9x1_cifs1\FLDR2
- □ \\SVM9X-1\C$\SVM9x1_cifs1\FLDR3
- □ \\SVM9X-1\C$\SVM9x1_cifs1\matts
- □ \\SVM9X-1\C$\SVM9x1_cifs1\New folder
- □ \\SVM9X-1\C$\SVM9x1_cifs1\omar
- □ \\SVM9X-1\C$\SVM9x1_cifs1\test
- □ eshta_456.txt
- □ FILE1.TXT
- □ FILE10.TXT
- □ FILE11.TXT
- □ FILE13.TXT
- □ FILE14.TXT
- □ FILE15.TXT

2. Click **Next** and continue with Recovery Options.

Recovery Options

After you find 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.

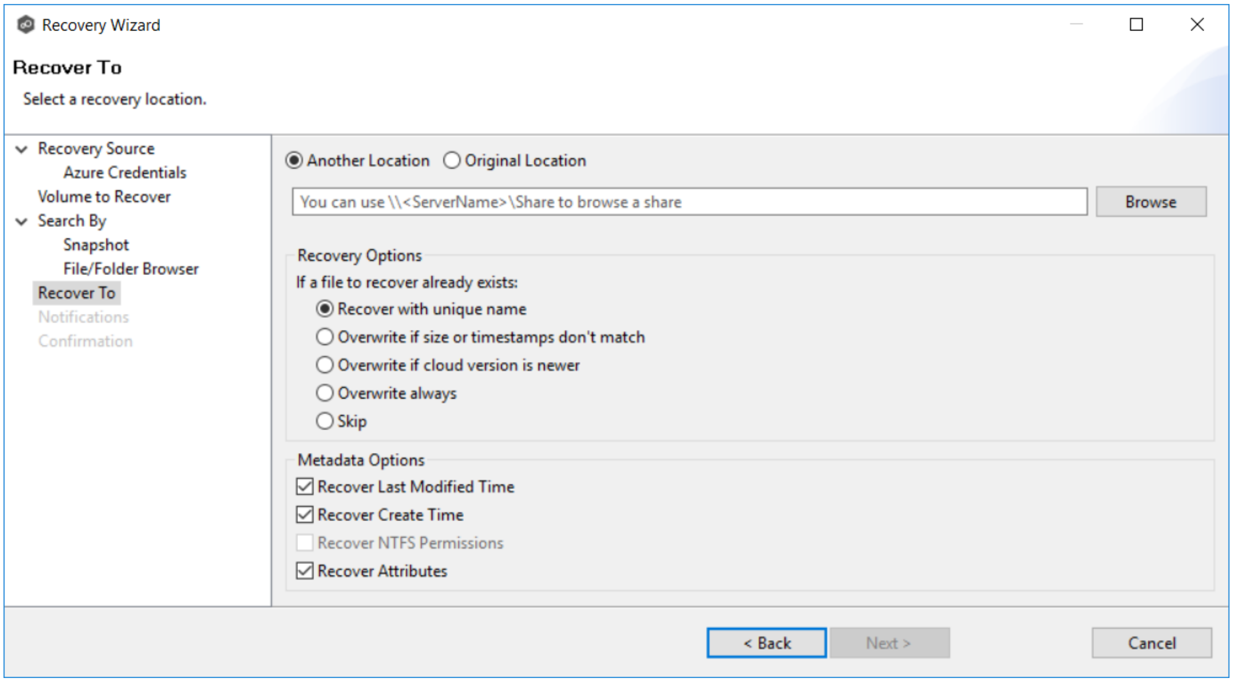
Recovery Wizard

Recover To

Select a recovery location.

● Another Location ○ Original Location

You can use \\<ServerName>\Share to browse a share

Recovery Options

If a file to recover already exists:

● Recover with unique name

○ Overwrite if size or timestamps don't match

○ Overwrite if cloud version is newer

○ Overwrite always

○ Skip

Metadata Options

☑ Recover Last Modified Time

☑ Recover Create Time

□ Recover NTFS Permissions

☑ Recover Attributes

2. Select the recovery options for if 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.

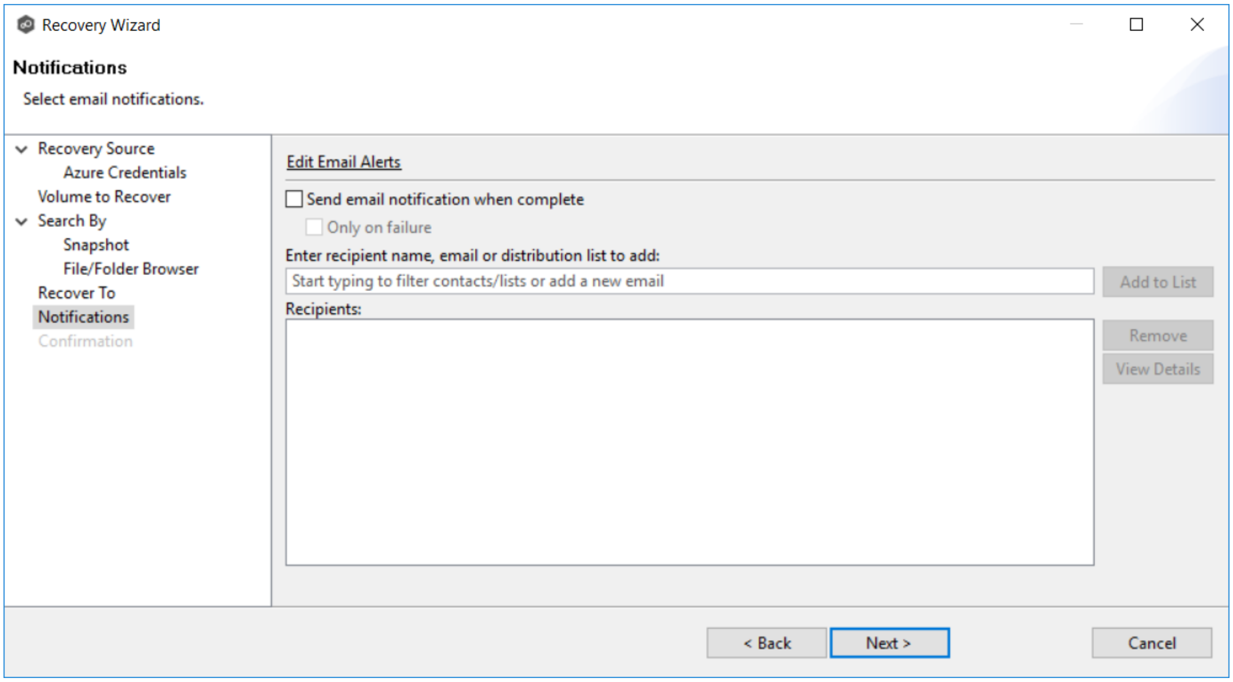

6. (Optional) Select the **Send email notification when complete** check box if you want notifications sent when the recovery process is complete. Select **Only on failure** if you want notifications sent only if the recovery does not successfully complete.

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.

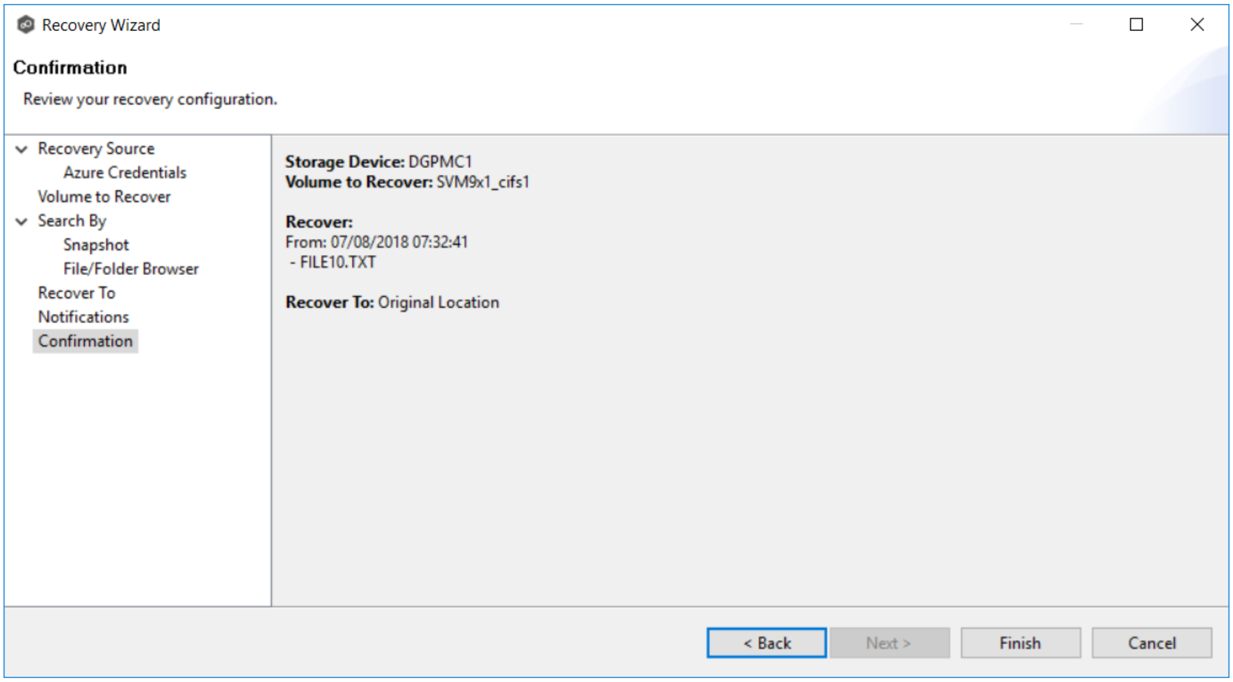

Storage Device: DGPMC1
Volume to Recover: SVM9x1_cifs1

Recover:
From: 07/08/2018 07:32:41
 - FILE10.TXT

Recover To: Original Location

10. Click **Finish**.

DFS-N Management Jobs

This section provides information about creating, editing, running, and managing a DFS-N Management job:

- Overview
- Namespace Elements
- Getting Started With DFS Namespaces
- Creating a DFS_N Management Job
- Running a DFS-N Management Job
- Managing DFS Namespaces
  - o Adding an Existing Namespace
  - o Adding a Namespace Folder

o Adding a Namespace Folder Target

o Connecting DFS Namespaces with File Collaboration and File Synchronization Jobs

Overview

The purpose of creating a DFS Namespace Management job is to allow you to manage various activities related to DFS namespaces, such as creating a namespace, creating namespace folders, and adding folder targets. A DFS namespace enables you to group shared folders located on different servers into one or more logically structured namespaces. DFS namespace activities can be performed using a Microsoft tool; however, the benefits of creating and configuring namespace within the Peer Management Center are:

- **Ease of managing a namespace** - You can create and manage a namespace within the same interface that manages our synchronization and collaboration 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 Synchronization and Collaboration** - When combined with PMC's file synchronization technology, DFS namespaces can provide redundancy to file shares across file servers and locations.
- **Automating failover and failback** - If a file server goes offline, the PMC will disable the associated folder target in DFS namespace. This automatically redirects users to another available file server. When the original file server comes back, the PMC 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.

Namespace Elements

Namespace Elements

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.

- **Folder** - Folders with folder targets provide users with actual content. When users browse a folder that has folder targets in the namespace, 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 is stored. For example, if a user navigates to **\\Domain.local\MyNamespace\MyFolder**, they are transparently redirected to **\\NYC-FS.Domain.local\MyFolder** or **\\LA-FS.Domain.local\MyFolder**, depending on which site the user is currently located in.

Getting Started With DFS Namespaces

If you need to create a namespace, begin by creating a DFS-N Management job. You may want to configure your DFS preferences before you create the job.

If you already have a namespace that you want to import, see Adding an Existing Namespace.

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

Once you have configured your namespace, you can link it to a File Collaboration or File Synchronization job.

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: Review](#)

[Step 9: Results](#)

Step 1: Job Type

1. Open the Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **New Job** wizard displays a list of job types you can create.
3. Click **DFS-N**, and then click **Create**.

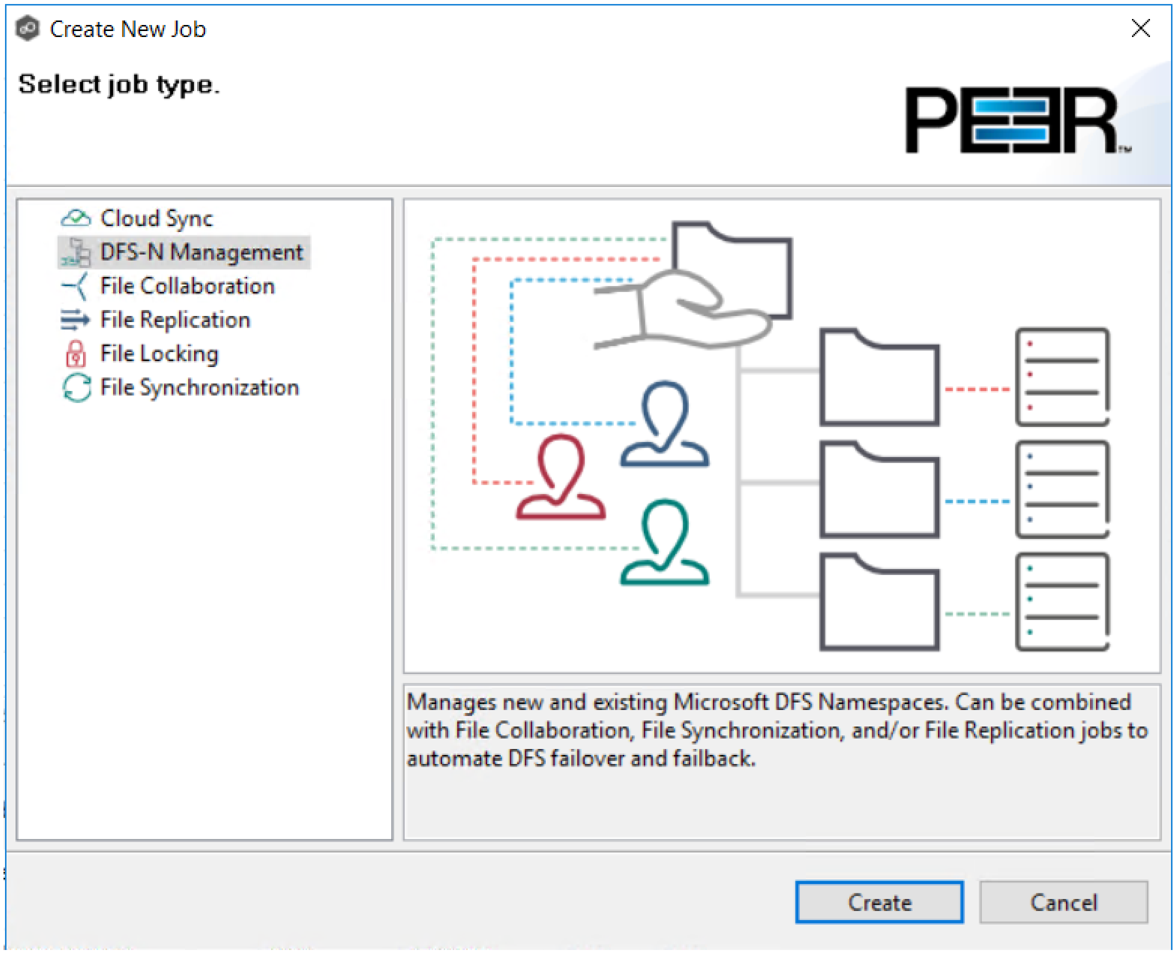

Manages new and existing Microsoft DFS Namespaces. Can be combined with File Collaboration, File Synchronization, and/or File Replication jobs to automate DFS failover and failback.

The Management Agent page appears.

Step 2: Management Agent

The **Management Agent** page presents a list of servers that have a Peer Agent installed.

1. Select an Agent that is in the domain where the DFS namespace or where you want to createe 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 the next step.

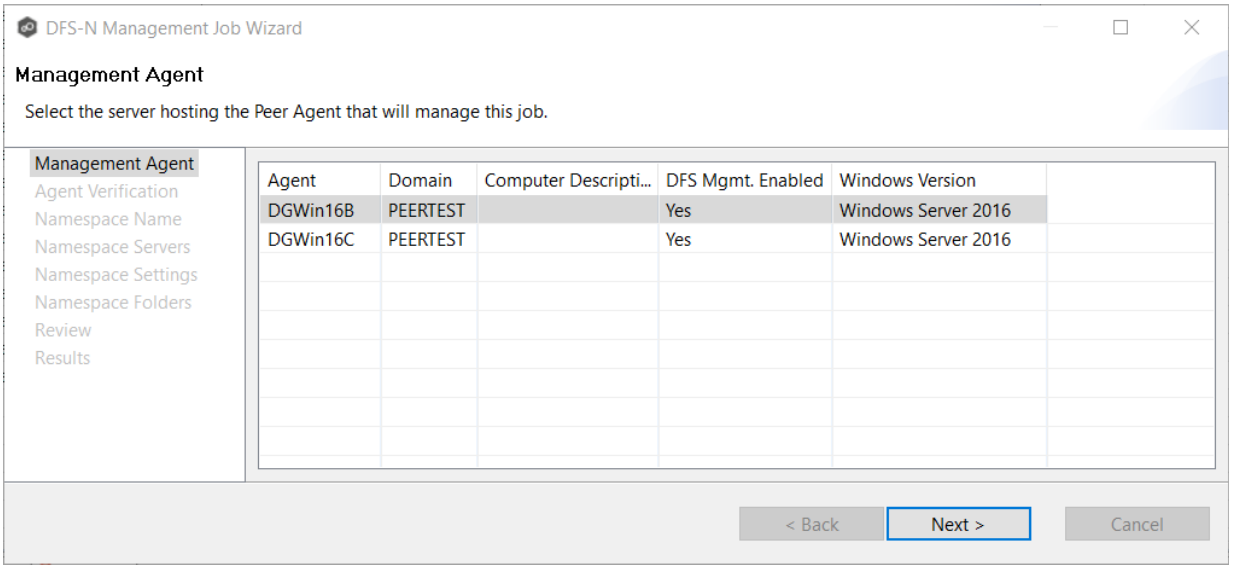

Agent	Domain	Computer Descripti...	DFS Mgmt. Enabled	Windows Version
DGWin16B	PEERTEST		Yes	Windows Server 2016
DGWin16C	PEERTEST		Yes	Windows Server 2016

2. Click **Next**.

   The Agent Verification page appears.

Step 3: Agent Verification

The **Agent Verification** page presents a list of steps that are performed to verify that the Microsoft DFS Powershell Management toolkit is installed on the same system as the Agent and configured correctly.

Note: The verification does not include checking whether DFS Services is running because the DFS service doesn't have to run on the agent server itself, it typically runs on a domain controller.

1. Click **Start Verification**.

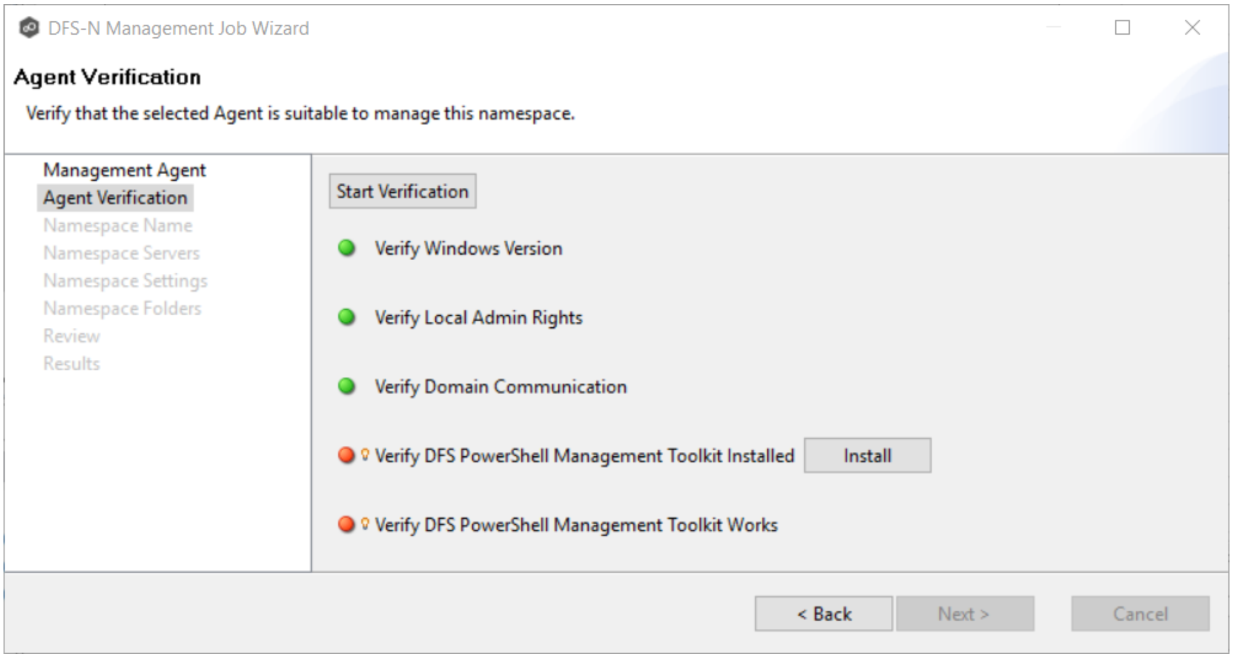

DFS-N Management Job Wizard

Agent Verification

Verify that the selected Agent is suitable to manage this namespace.

2. If the DFS PowerShell Management toolkit is not installed, click the **Install** button, and then after the toolkit is installed, click the **Start Verification** button again.

3. After the verification has successfully completed, click **Next**.

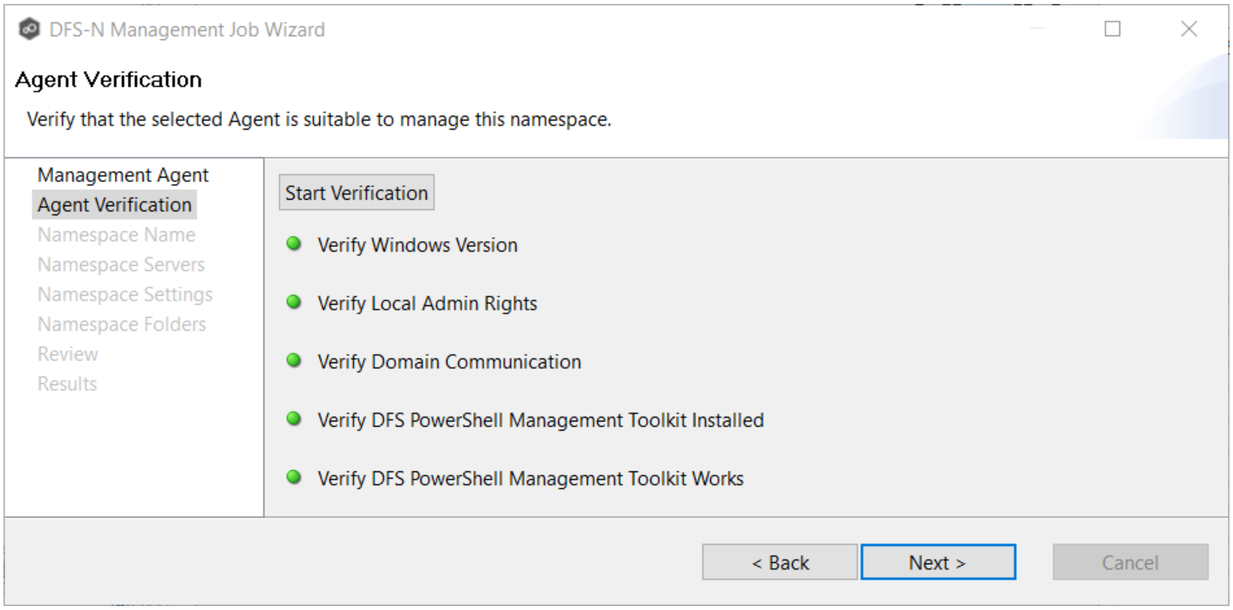

DFS-N Management Job Wizard

Agent Verification

Verify that the selected Agent is suitable to manage this namespace.

The Namespace Name page appears.

Step 4: Namespace Name

The name of the namespace will also be the name of the DFS-N Management job.

1. Enter the name of the namespace.

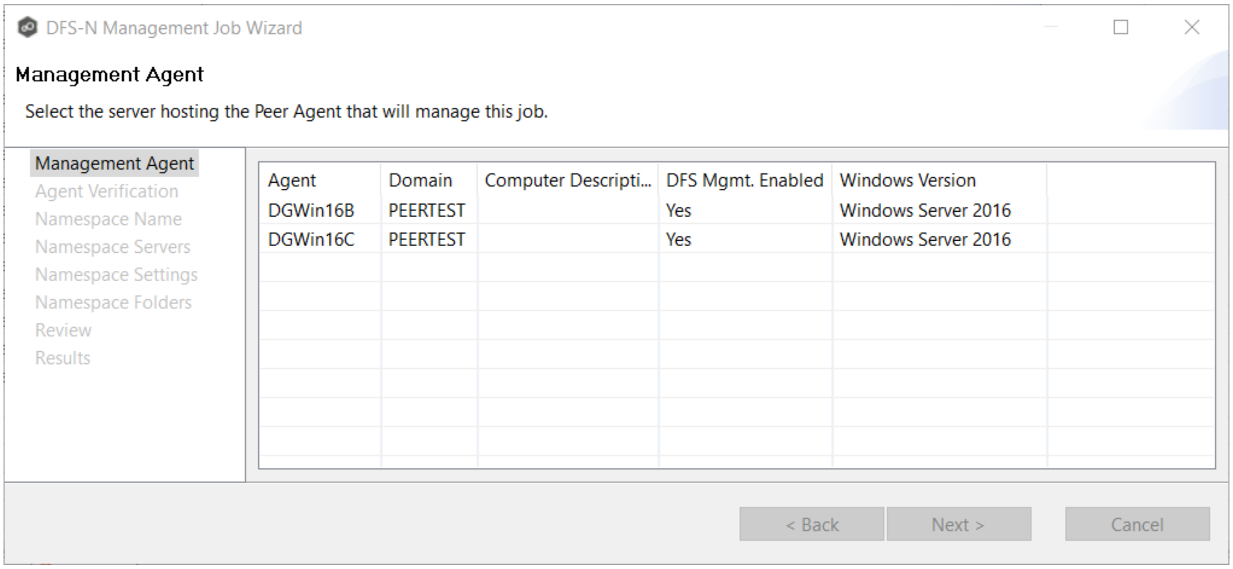

Management Agent

Select the server hosting the Peer Agent that will manage this job.

Agent	Domain	Computer Descripti...	DFS Mgmt. Enabled	Windows Version
DGWin16B	PEERTEST		Yes	Windows Server 2016
DGWin16C	PEERTEST		Yes	Windows Server 2016

2. Click **Next**.

   The Namespace Servers page appears.

Step 5: Namespace Servers

A server that you want to host a namespace is called a namespace server. It does not have to host the data. However, a namespace server must be running the Microsoft DFS Namespace Service. In most cases, a namespace server should be a domain controller.

1. Enter the fully qualified path of a file server in the **Server Name** field, and then click **Add**.

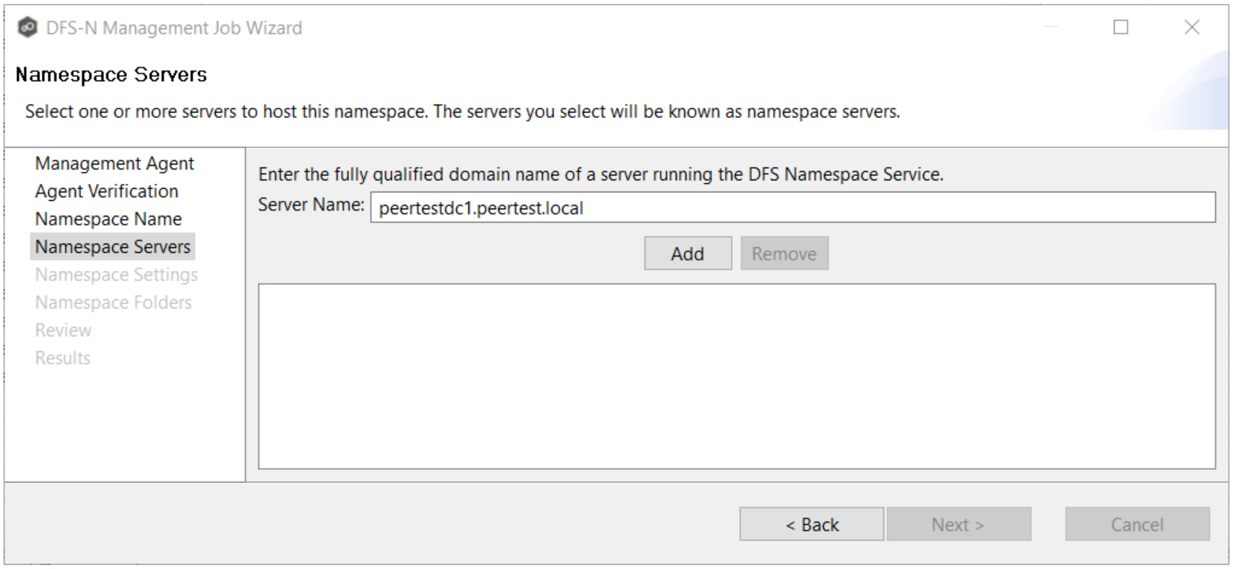

DFS-N Management Job Wizard

Namespace Servers

Select one or more servers to host this namespace. The servers you select will be known as namespace servers.

Enter the fully qualified domain name of a server running the DFS Namespace Service.

Server Name: peertestdc1.peertest.local

The server path is listed in the area below.

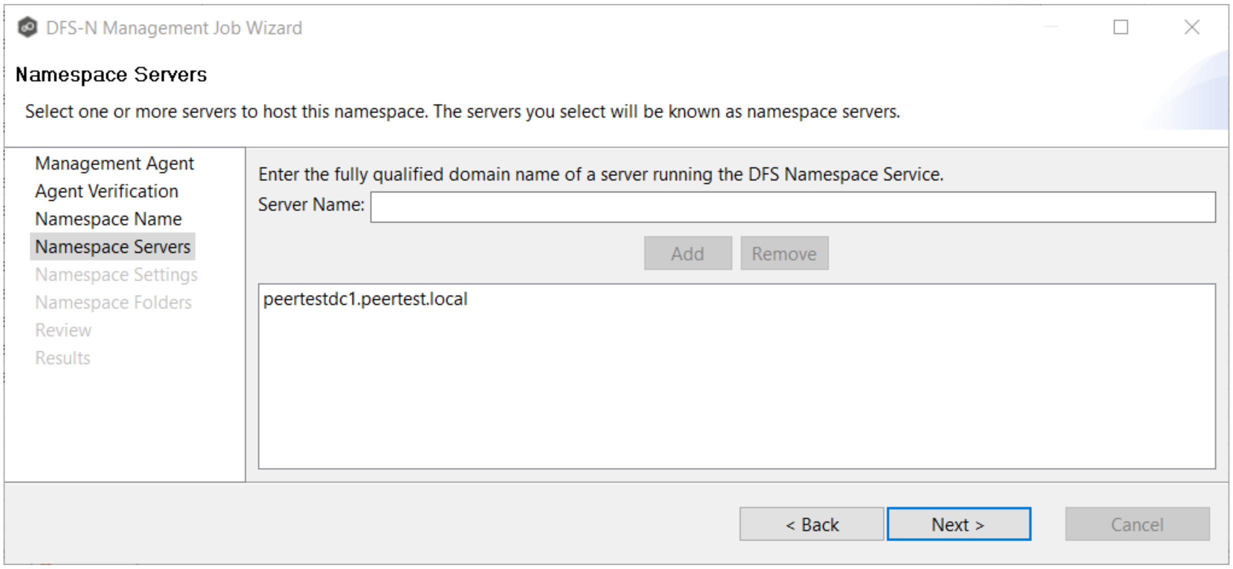

DFS-N Management Job Wizard

Namespace Servers

Select one or more servers to host this namespace. The servers you select will be known as namespace servers.

Enter the fully qualified domain name of a server running the DFS Namespace Service.

Server Name:

peertestdc1.peertest.local

2. Add additional servers if desired.

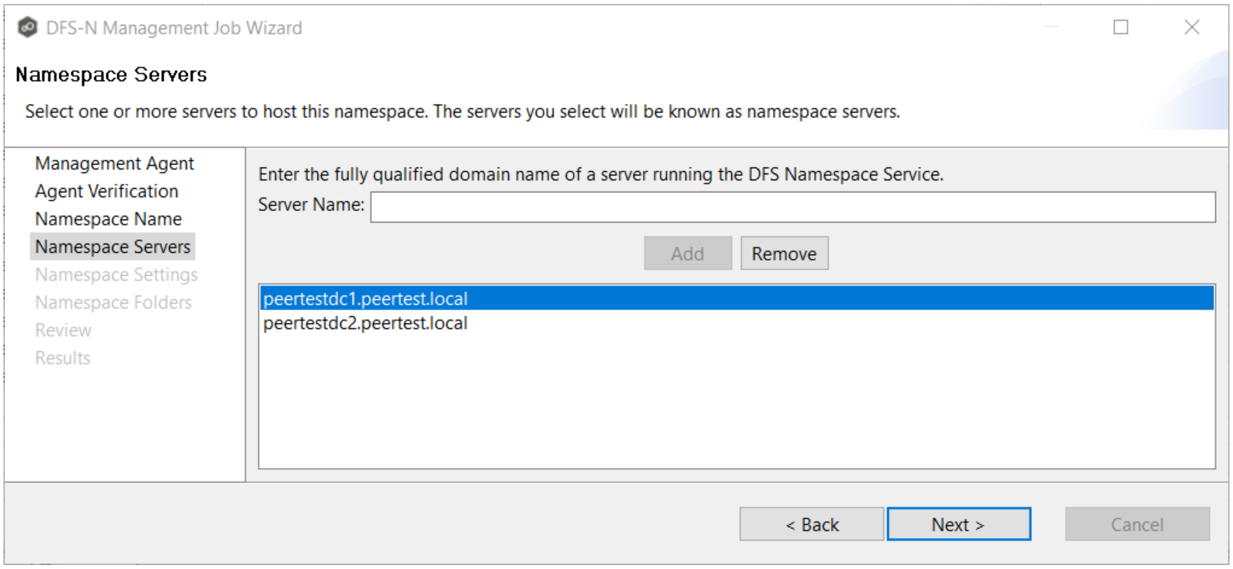

DFS-N Management Job Wizard

Namespace Servers

Select one or more servers to host this namespace. The servers you select will be known as namespace servers.

Enter the fully qualified domain name of a server running the DFS Namespace Service.

Server Name:

peertestdc1.peertest.local
peertestdc2.peertest.local

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 a server's local path and access permissions.

To edit a server's settings:

1. In the **Shared Folder Path** column for the server, type the new path.

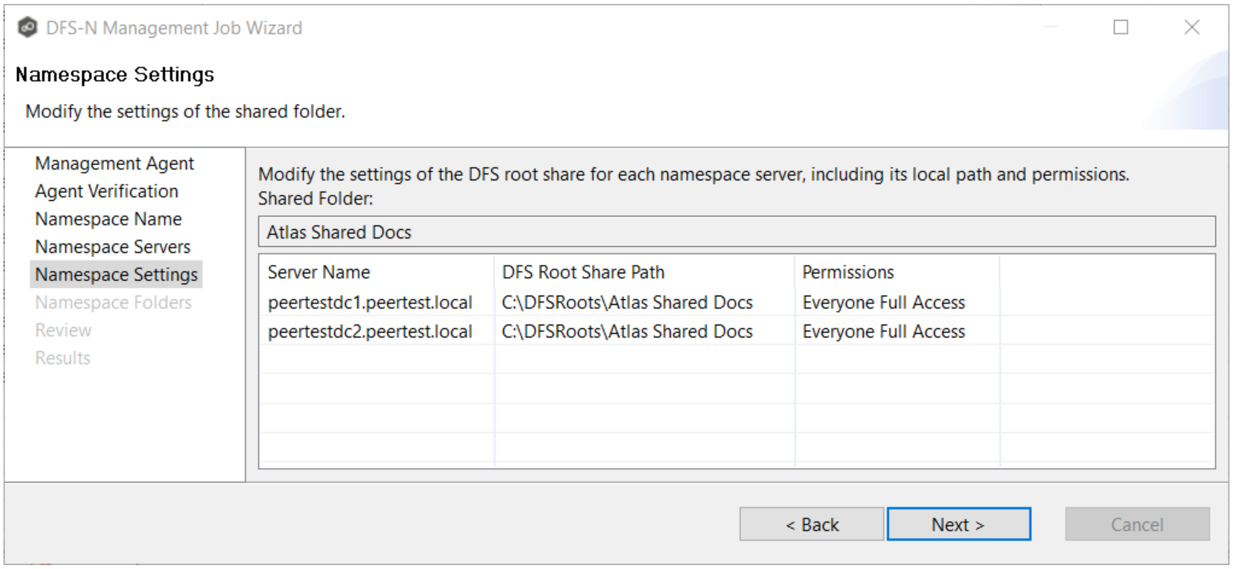

Modify the settings of the DFS root share for each namespace server, including its local path and permissions.

Shared Folder:

Atlas Shared Docs

Server Name	DFS Root Share Path	Permissions
peertestdc1.peertest.local	C:\DFSRoots\Atlas Shared Docs	Everyone Full Access
peertestdc2.peertest.local	C:\DFSRoots\Atlas Shared Docs	Everyone Full Access

2. In the **Permissions** column, select the desired access level.

3. Click **Next**.

   The Namespace Folders page appears.

Step 7: Namespace Folders

A namespace folder contains folder targets, which provide users with actual content. A folder target is the Universal Naming Convention (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 is stored. Adding multiple folder targets increases the availability of the folder in the namespace.

The **Namespace Folders** page lists the namespace folders and folder targets.

1. Click the **Add** button.

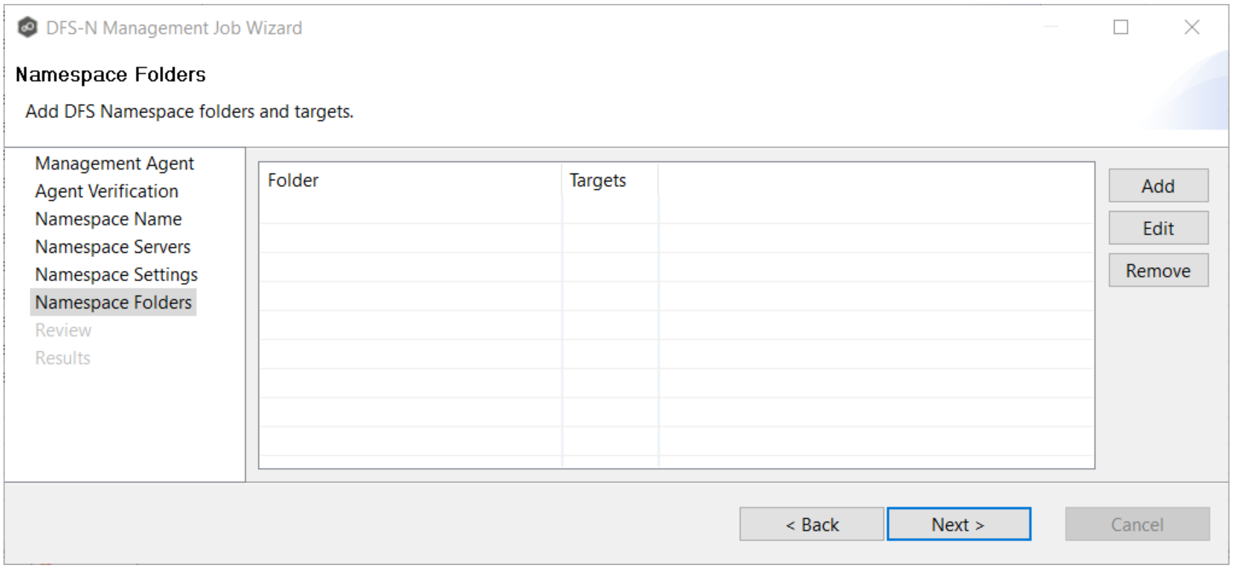

Namespace Folders

Add DFS Namespace folders and targets.

The **Folder Name** dialog appears.

2. Enter a name for the namespace folder in the **Folder Name** field.

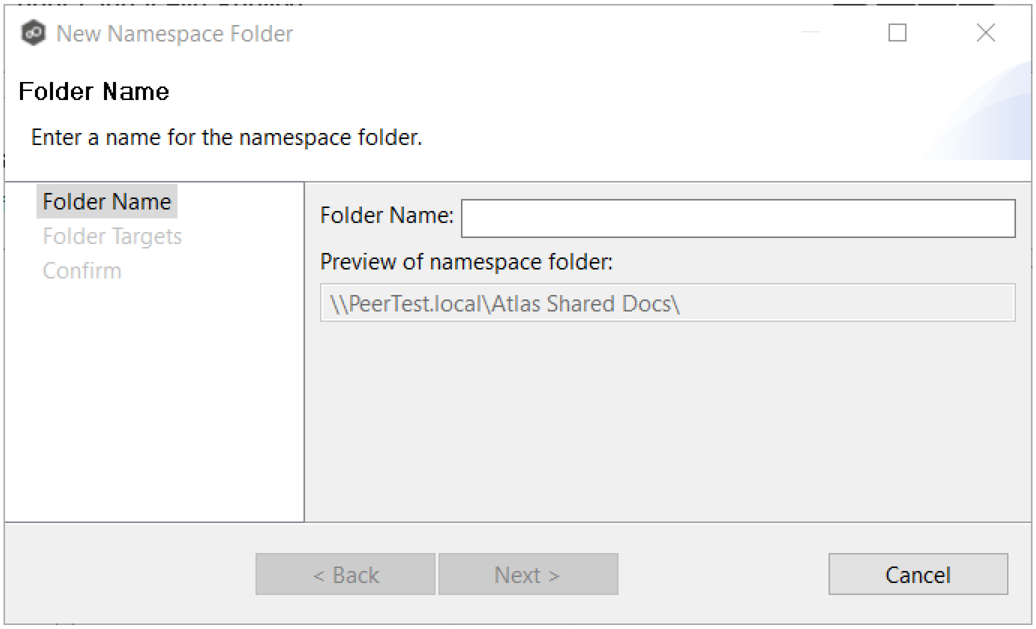

Folder Name

Enter a name for the namespace folder.

Folder Name:

Preview of namespace folder: \\PeerTest.local\Atlas Shared Docs\

After you enter the folder name, a preview of the folder and path name appears below the **Folder Name** field.

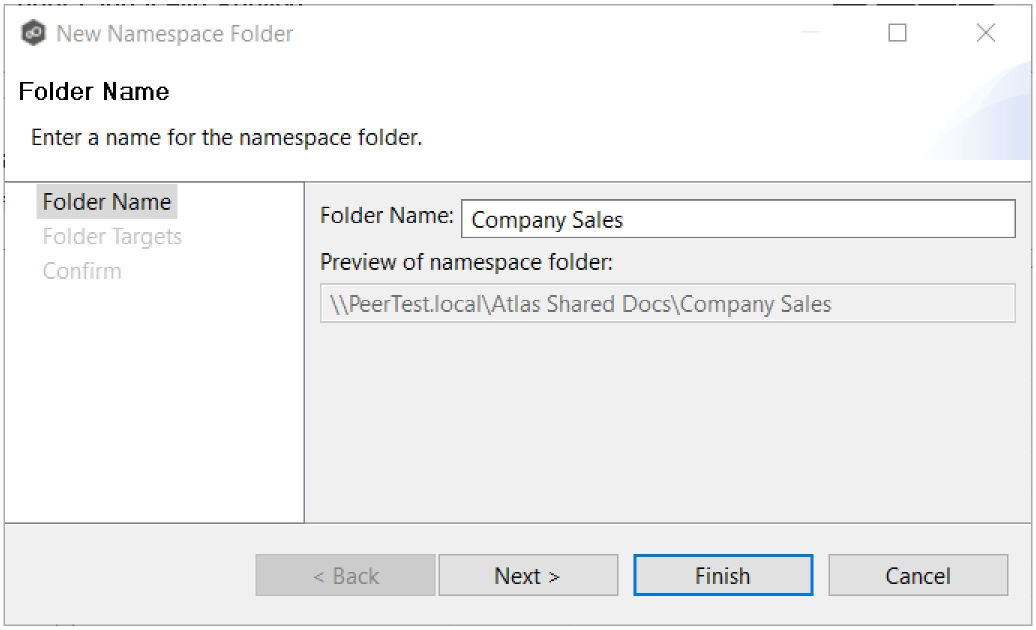
New Namespace Folder

Folder Name

Enter a name for the namespace folder.

Folder Name: Company Sales

Preview of namespace folder:

\\PeerTest.local\Atlas Shared Docs\Company Sales

3. If you want to add folder targets for the namespace folder, click **Next**.

   The **Folder Targets** dialog appears.

4. Enter the UNC path to a shared folder, and then click **Add**.

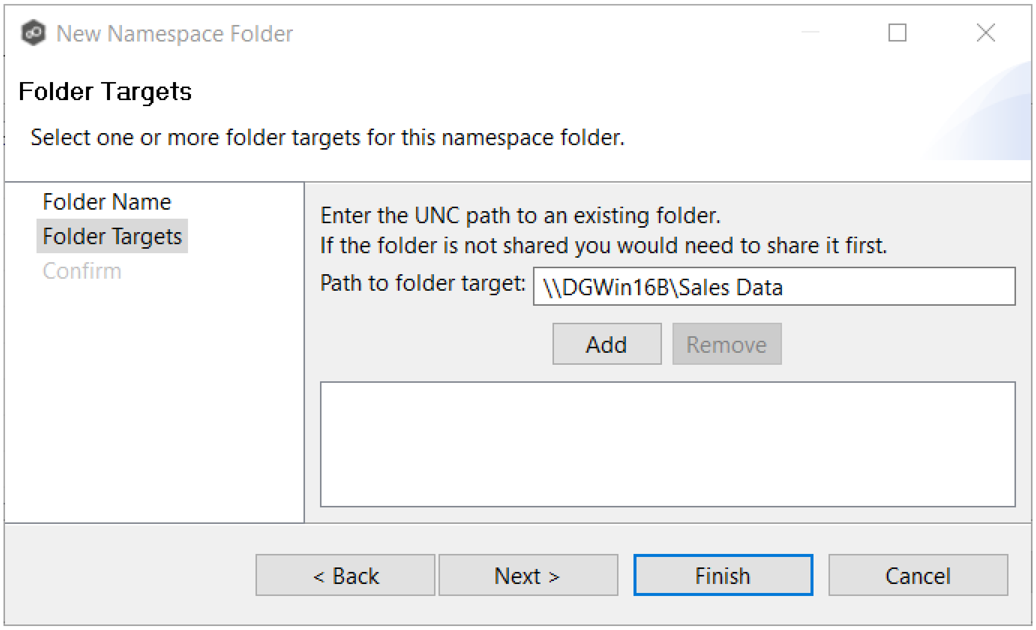
New Namespace Folder

Folder Targets

Select one or more folder targets for this namespace folder.

Enter the UNC path to an existing folder.
If the folder is not shared you would need to share it first.

Path to folder target: \\DGWin16B\Sales Data

The folder target path is listed in the field below.

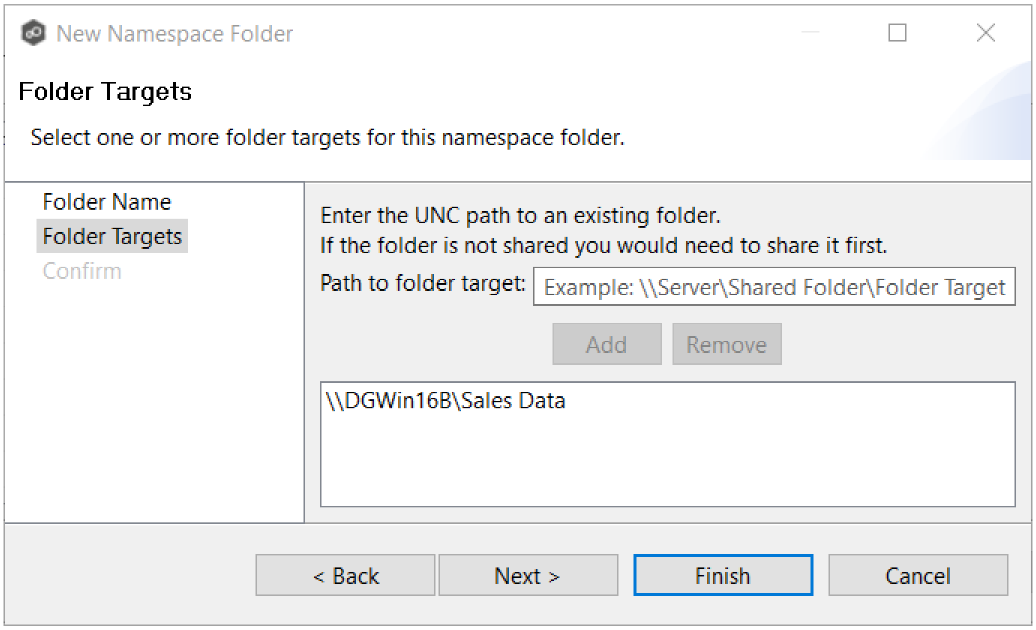

New Namespace Folder

Folder Targets

Select one or more folder targets for this namespace folder.

Enter the UNC path to an existing folder.
If the folder is not shared you would need to share it first.

Path to folder target: Example: \\Server\Shared Folder\Folder Target

\\DGWin16B\Sales Data

5. (Optional) Add additional folder targets.

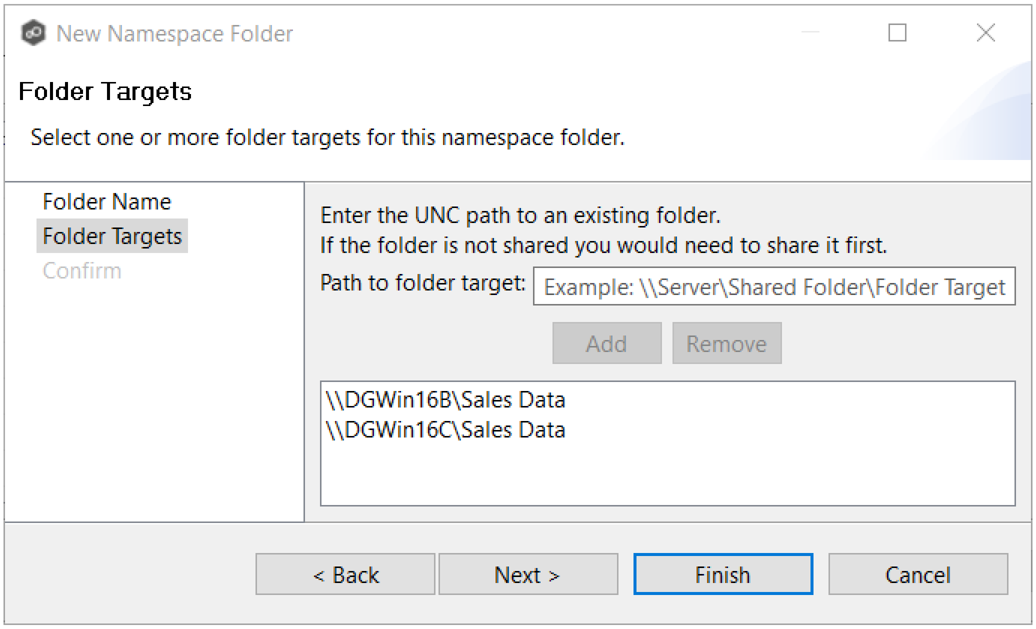

New Namespace Folder

Folder Targets

Select one or more folder targets for this namespace folder.

Enter the UNC path to an existing folder.
If the folder is not shared you would need to share it first.

Path to folder target: Example: \\Server\Shared Folder\Folder Target

\\DGWin16B\Sales Data
\\DGWin16C\Sales Data

6. Click **Next**.

   The **Confirm** dialog appears.

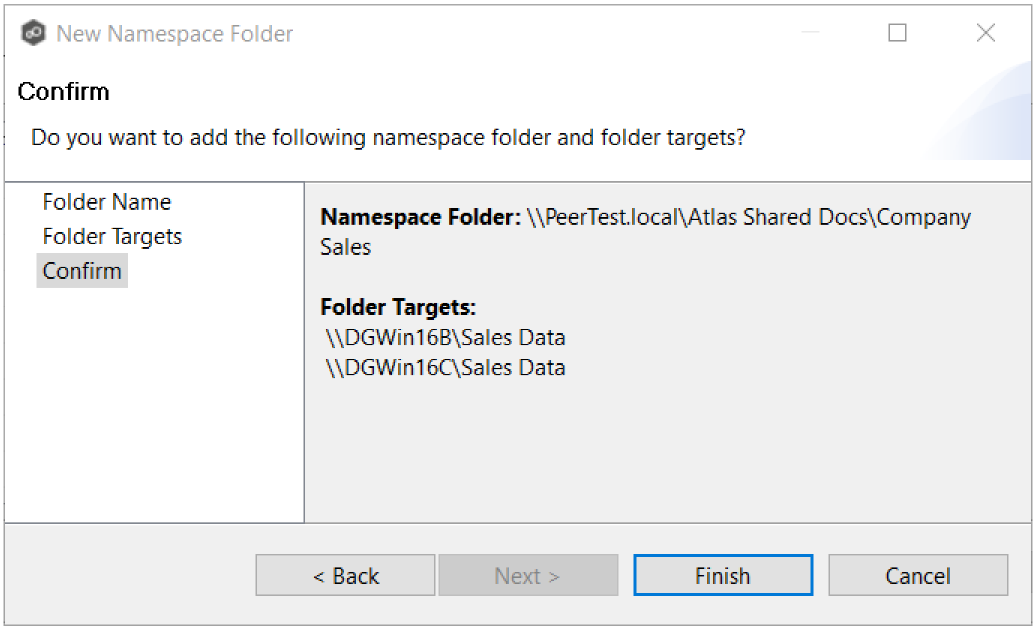

Namespace Folder: \\PeerTest.local\Atlas Shared Docs\Company Sales

Folder Targets:
\\DGWin16B\Sales Data
\\DGWin16C\Sales Data

7. Review the folders and folder targets, and then click **Back** to add more folder and folder targets; otherwise, click **Finish**.

   The **Namespace Folders** page reappears; it lists the folder you added and the number of its targets.

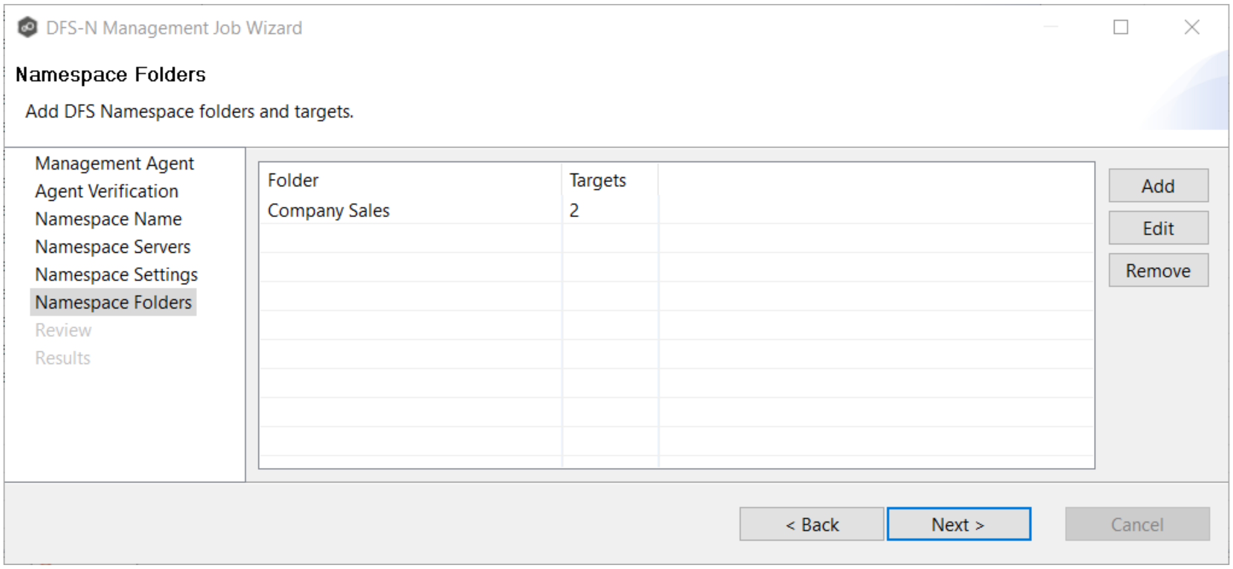

Folder	Targets
Company Sales	2

8. Click **Next**.

   The Review page appears.

Step 8: Review

The **Review** page allows you to review the configuration before it is actually created.

1. Review the namespace configuration.

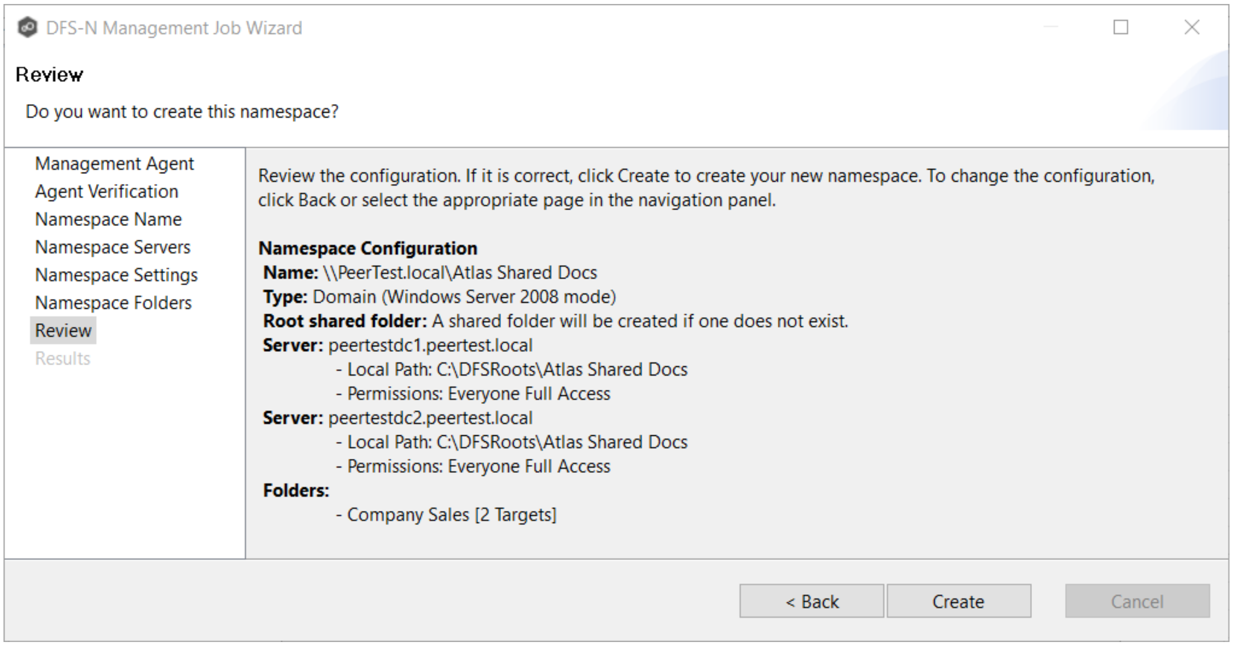

Review the configuration. If it is correct, click Create to create your new namespace. To change the configuration, click Back or select the appropriate page in the navigation panel.

Namespace Configuration
Name: \\PeerTest.local\Atlas Shared Docs
Type: Domain (Windows Server 2008 mode)
Root shared folder: A shared folder will be created if one does not exist.
Server: peertestdc1.peertest.local
- Local Path: C:\DFSRoots\Atlas Shared Docs
- Permissions: Everyone Full Access

Server: peertestdc2.peertest.local
- Local Path: C:\DFSRoots\Atlas Shared Docs
- Permissions: Everyone Full Access

Folders:
- Company Sales [2 Targets]

2. Click **Create** if the configuration is correct; otherwise, click **Back** and correct the configuration.

   After you click **Create**, the Results page appears.

Step 9: Results

The **Results** page has two tabs: **Tasks** and **Errors**.

1. Review the results in the **Tasks** and **Errors** tabs.

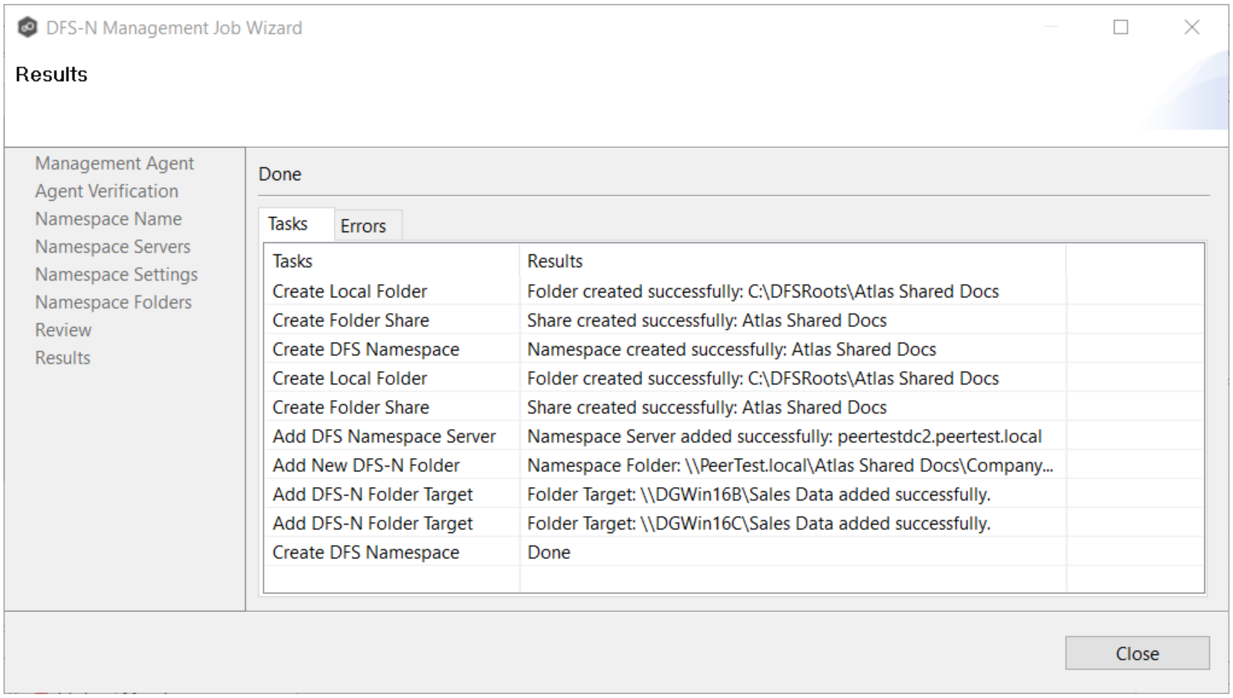

Tasks	Results
Create Local Folder	Folder created successfully: C:\DFSRoots\Atlas Shared Docs
Create Folder Share	Share created successfully: Atlas Shared Docs
Create DFS Namespace	Namespace created successfully: Atlas Shared Docs
Create Local Folder	Folder created successfully: C:\DFSRoots\Atlas Shared Docs
Create Folder Share	Share created successfully: Atlas Shared Docs
Add DFS Namespace Server	Namespace Server added successfully: peertestdc2.peertest.local
Add New DFS-N Folder	Namespace Folder: \\PeerTest.local\Atlas Shared Docs\Company...
Add DFS-N Folder Target	Folder Target: \\DGWin16B\Sales Data added successfully.
Add DFS-N Folder Target	Folder Target: \\DGWin16C\Sales Data added successfully.
Create DFS Namespace	Done

2. Click **Close**.

The **Namespace Summary** view is displayed.

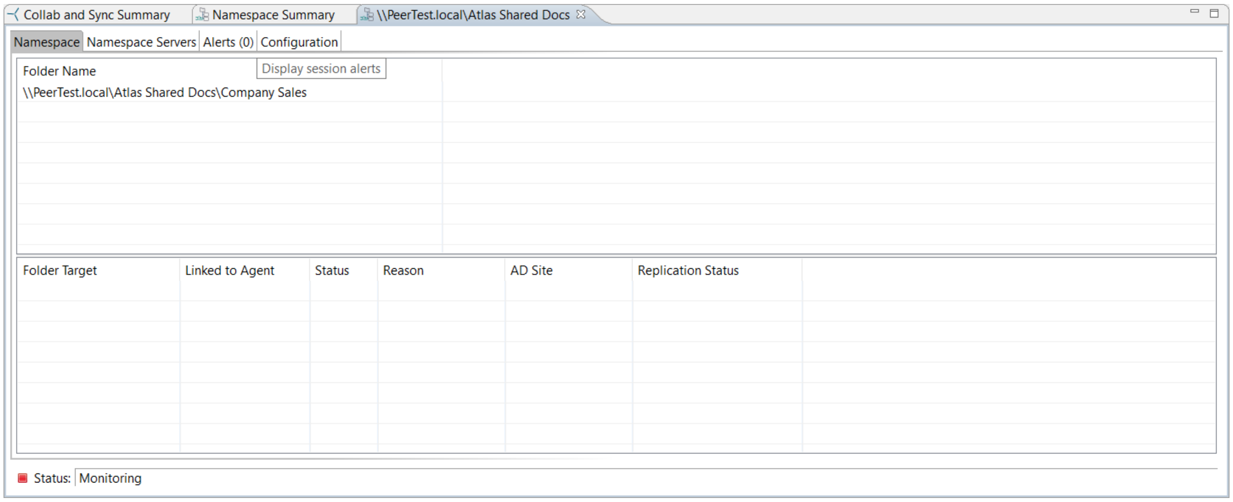

Folder Name
\\PeerTest.local\Atlas Shared Docs\Company Sales

Folder Target	Linked to Agent	Status	Reason	AD Site	Replication Status

Status: Monitoring

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

When running a DFS-N Management 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.

To manually start a job:

1. Choose one of these options:

   - Right-click the job name in the **Jobs** view.
   - Open the job and then click the **Start/Stop** button in the bottom left corner of the job's **Namespace** tab in the run-time view (shown below).

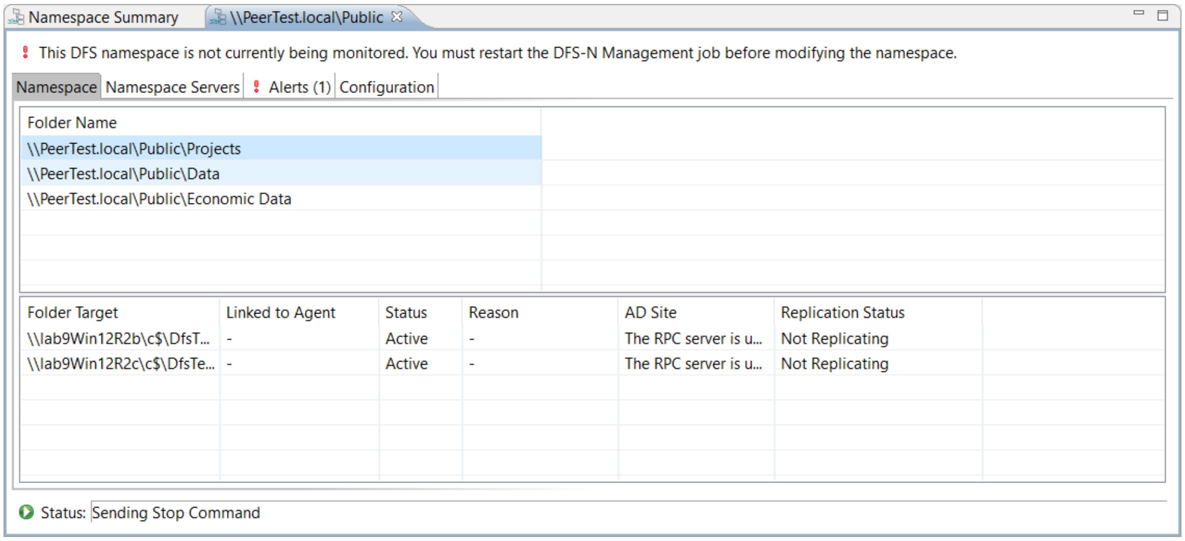

This DFS namespace is not currently being monitored. You must restart the DFS-N Management job before modifying the namespace.

Folder Name
\\PeerTest.local\Public\Projects
\\PeerTest.local\Public\Data
\\PeerTest.local\Public\Economic Data

Folder Target	Linked to Agent	Status	Reason	AD Site	Replication Status
\\lab9Win12R2b\c$\DfsT...	-	Active	-	The RPC server is u...	Not Replicating
\\lab9Win12R2c\c$\DfsTe...	-	Active	-	The RPC server is u...	Not Replicating

Status: Sending Stop Command

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.

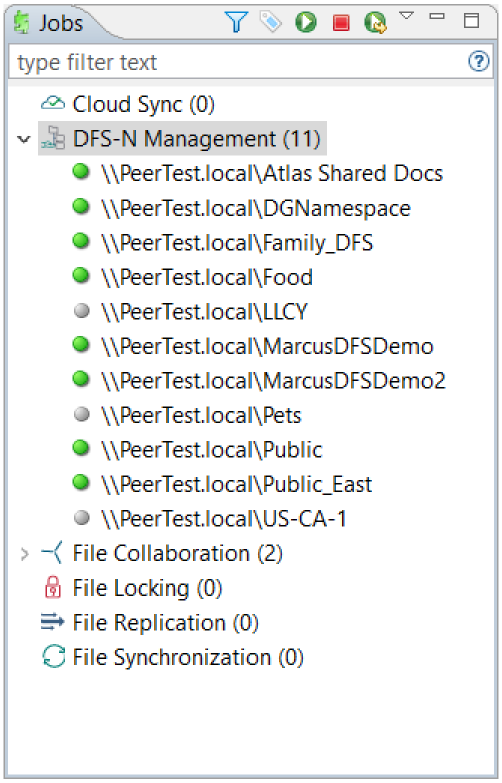

Cloud Sync (0)
DFS-N Management (11)
- \\PeerTest.local\Atlas Shared Docs
- \\PeerTest.local\DGNamespace
- \\PeerTest.local\Family_DFS
- \\PeerTest.local\Food
- \\PeerTest.local\LLCY
- \\PeerTest.local\MarcusDFSDemo
- \\PeerTest.local\MarcusDFSDemo2
- \\PeerTest.local\Pets
- \\PeerTest.local\Public
- \\PeerTest.local\Public_East
- \\PeerTest.local\US-CA-1

File Collaboration (2)
File Locking (0)
File Replication (0)
File Synchronization (0)

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 pop-up menu.

   Or, open the job and click the **Start/Stop** button in the bottom left corner of the job's **Namespace** tab in the run-time 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 an Existing Namespace
- Adding a Namespace Server
- Adding a Namespace Folder
- Adding a Namespace Folder Target

Adding 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. You can then either link the namespace to an existing File Collaboration or 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 **Namespace Summary** view, and then select **Add Existing Namespaces**.

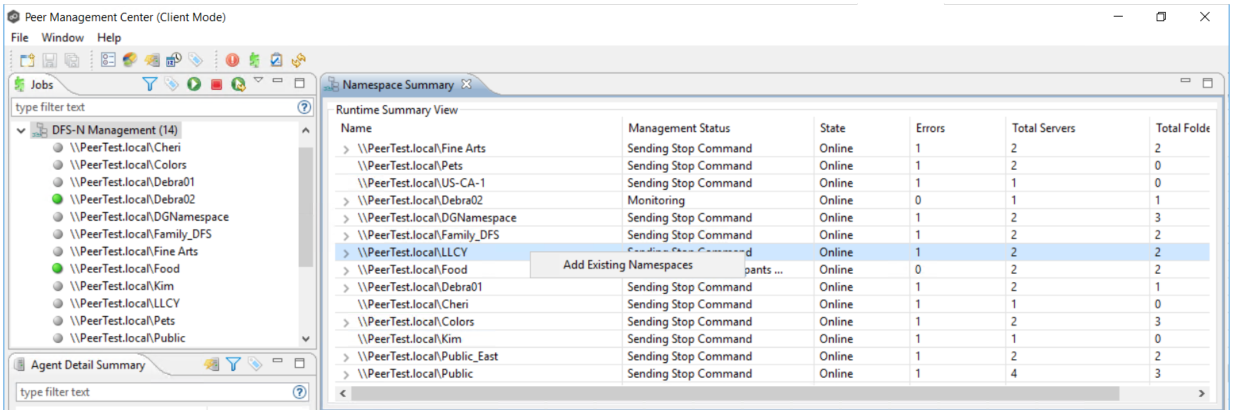

Name	Management Status	State	Errors	Total Servers	Total Folde
\\PeerTest.local\Fine Arts	Sending Stop Command	Online	1	2	2
\\PeerTest.local\Pets	Sending Stop Command	Online	1	2	0
\\PeerTest.local\US-CA-1	Sending Stop Command	Online	1	1	0
\\PeerTest.local\Debra02	Monitoring	Online	0	1	1
\\PeerTest.local\DGNamespace	Sending Stop Command	Online	1	2	3
\\PeerTest.local\Family_DFS	Sending Stop Command	Online	1	2	2
\\PeerTest.local\LLCY	Sending Stop Command	Online	1	2	2
\\PeerTest.local\Food	pants ...	Online	0	2	2
\\PeerTest.local\Debra01	Sending Stop Command	Online	1	2	1
\\PeerTest.local\Cheri	Sending Stop Command	Online	1	1	0
\\PeerTest.local\Colors	Sending Stop Command	Online	1	2	3
\\PeerTest.local\Kim	Sending Stop Command	Online	1	1	0
\\PeerTest.local\Public_East	Sending Stop Command	Online	1	2	2
\\PeerTest.local\Public	Sending Stop Command	Online	1	4	3

   The **Add Existing Namespace** wizard appears.

2. Select a management agent.

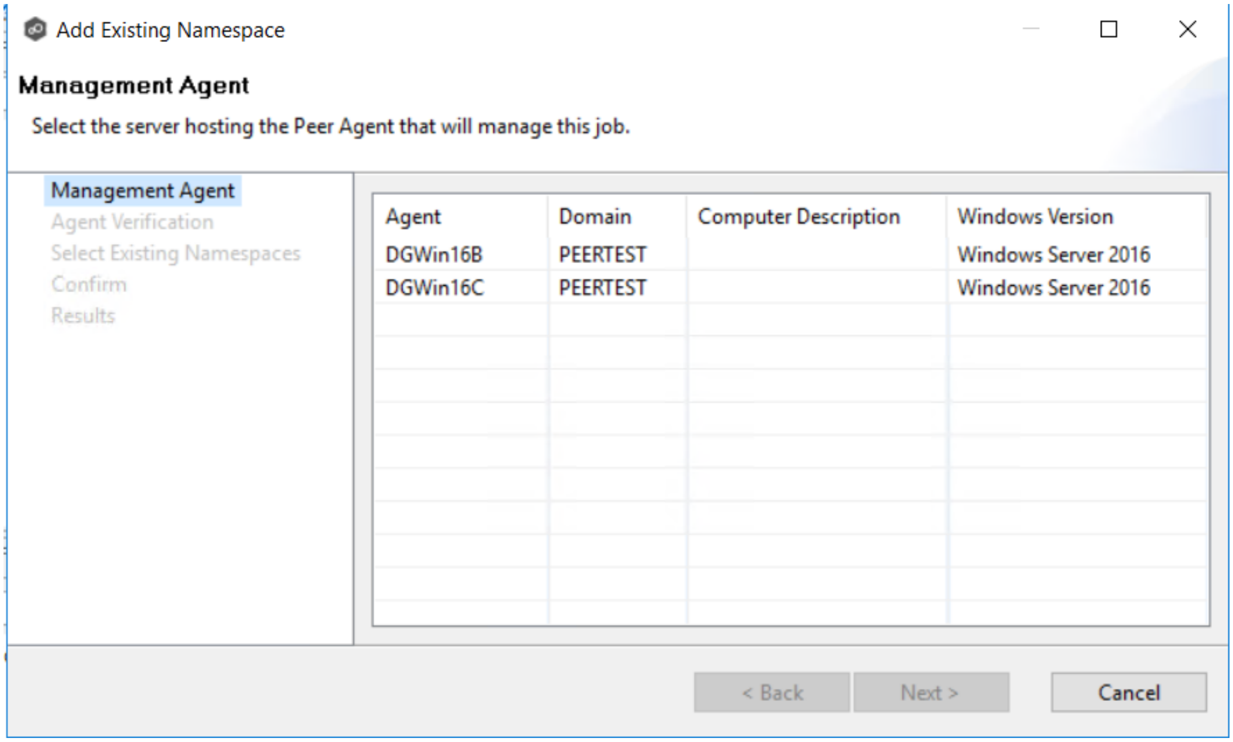

Management Agent

Select the server hosting the Peer Agent that will manage this job.

Agent	Domain	Computer Description	Windows Version
DGWin16B	PEERTEST		Windows Server 2016
DGWin16C	PEERTEST		Windows Server 2016

3. Verify the Agent environment.

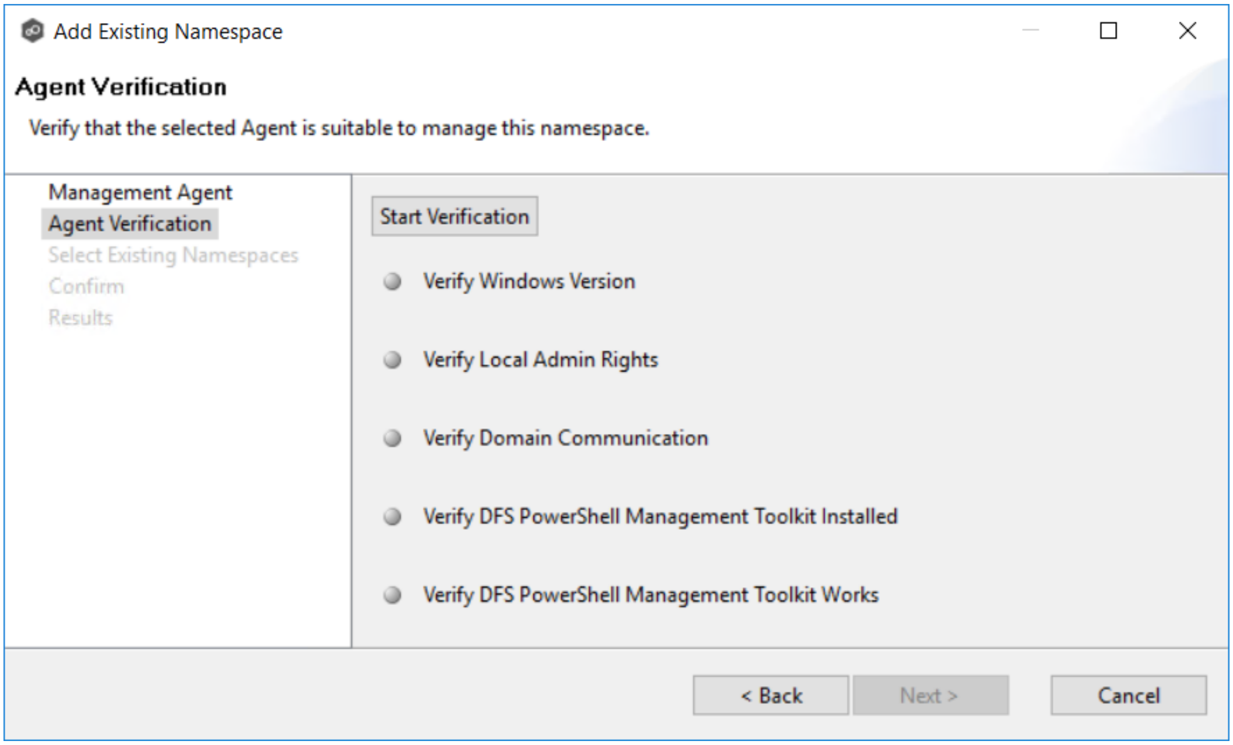

Agent Verification

Verify that the selected Agent is suitable to manage this namespace.

- Verify Windows Version
- Verify Local Admin Rights
- Verify Domain Communication
- Verify DFS PowerShell Management Toolkit Installed
- Verify DFS PowerShell Management Toolkit Works

4. Select one or more existing namespaces.

It may take a few minutes for the existing namespaces to appear in the table.

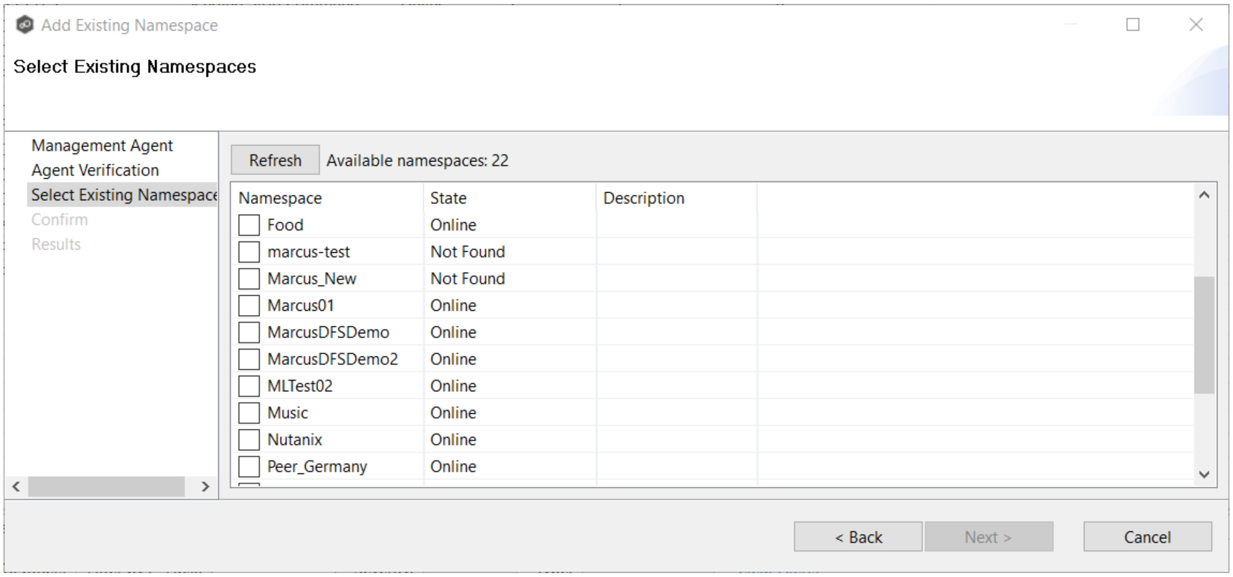

Add Existing Namespace

Select Existing Namespaces

Management Agent
Agent Verification
Select Existing Namespace
Confirm
Results

Available namespaces: 22

Namespace	State	Description
☐ Food	Online	
☐ marcus-test	Not Found	
☐ Marcus_New	Not Found	
☐ Marcus01	Online	
☐ MarcusDFSDemo	Online	
☐ MarcusDFSDemo2	Online	
☐ MLTest02	Online	
☐ Music	Online	
☐ Nutanix	Online	
☐ Peer_Germany	Online	

The **Confirm** dialog appears.

5. Review the configuration, and then click **Next**.

   The **Results** dialog appears.

6. Review the results, and then click **Close**.

   The namespace is displayed in the **Jobs** view.

Adding a Namespace Server

You can add a namespace server to a namespace.

To add a namespace server to a namespace:

1. Double-click the job name in the **Jobs** view or the **Namespace Summary** view to open the run-time view for the namespace.

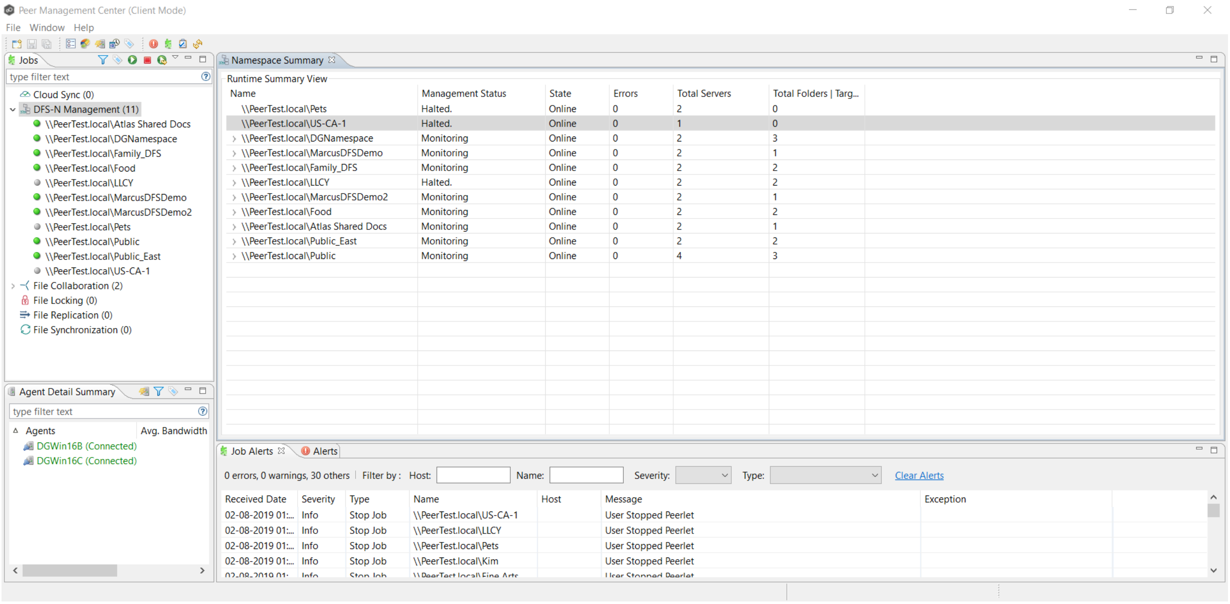

Runtime Summary View

Name	Management Status	State	Errors	Total Servers	Total Folders | Targ...
\\PeerTest.local\Pets	Halted.	Online	0	2	0
\\PeerTest.local\US-CA-1	Halted.	Online	0	1	0
\\PeerTest.local\DGNamespace	Monitoring	Online	0	2	3
\\PeerTest.local\MarcusDFSDemo	Monitoring	Online	0	2	1
\\PeerTest.local\Family_DFS	Monitoring	Online	0	2	2
\\PeerTest.local\LLCY	Halted.	Online	0	2	2
\\PeerTest.local\MarcusDFSDemo2	Monitoring	Online	0	2	1
\\PeerTest.local\Food	Monitoring	Online	0	2	2
\\PeerTest.local\Atlas Shared Docs	Monitoring	Online	0	2	1
\\PeerTest.local\Public_East	Monitoring	Online	0	2	2
\\PeerTest.local\Public	Monitoring	Online	0	4	3

0 errors, 0 warnings, 30 others | Filter by : Host: Name: Severity: Type:

Received Date	Severity	Type	Name	Host	Message	Exception
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\US-CA-1		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\LLCY		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Pets		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Kim		User Stopped Peerlet	
02-08-2019 01:	Info	Stop Job	\\PeerTest.local\Fine Arts		User Stopped Peerlet	

The job's run-time view is displayed.

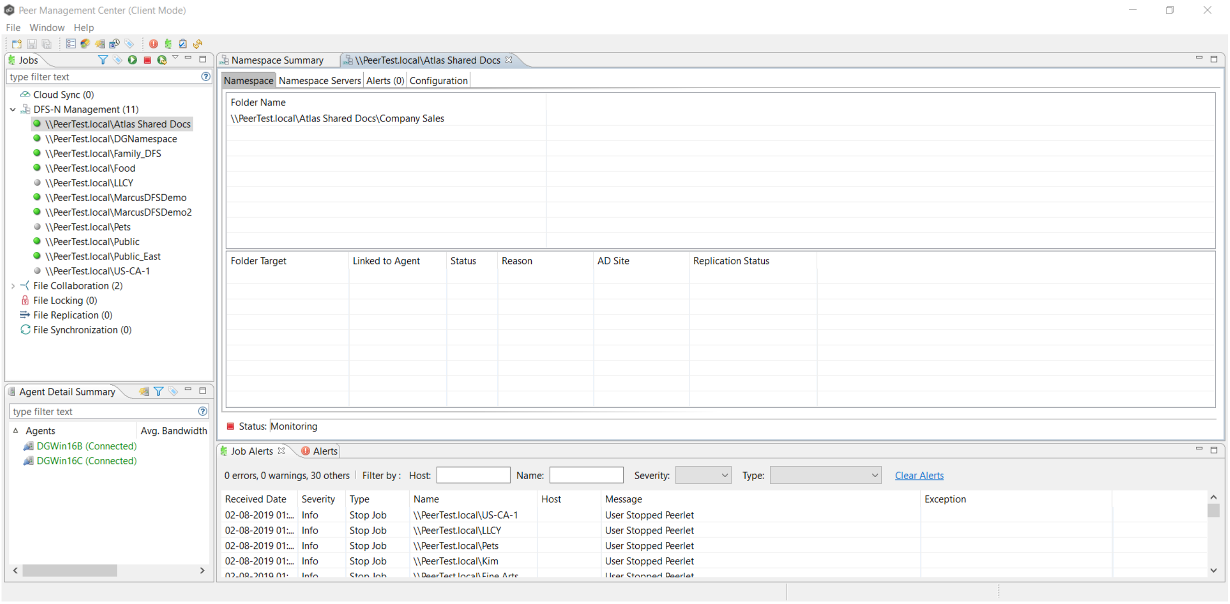

Folder Name
\\PeerTest.local\Atlas Shared Docs\Company Sales

Folder Target	Linked to Agent	Status	Reason	AD Site	Replication Status

Status: Monitoring

0 errors, 0 warnings, 30 others | Filter by : Host: Name: Severity: Type:

Received Date	Severity	Type	Name	Host	Message	Exception
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\US-CA-1		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\LLCY		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Pets		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Kim		User Stopped Peerlet	
02-08-2019 01:	Info	Stop Job	\\PeerTest.local\Fine Arts		User Stopped Peerlet	

2. Click the **Namespace Servers** tab.

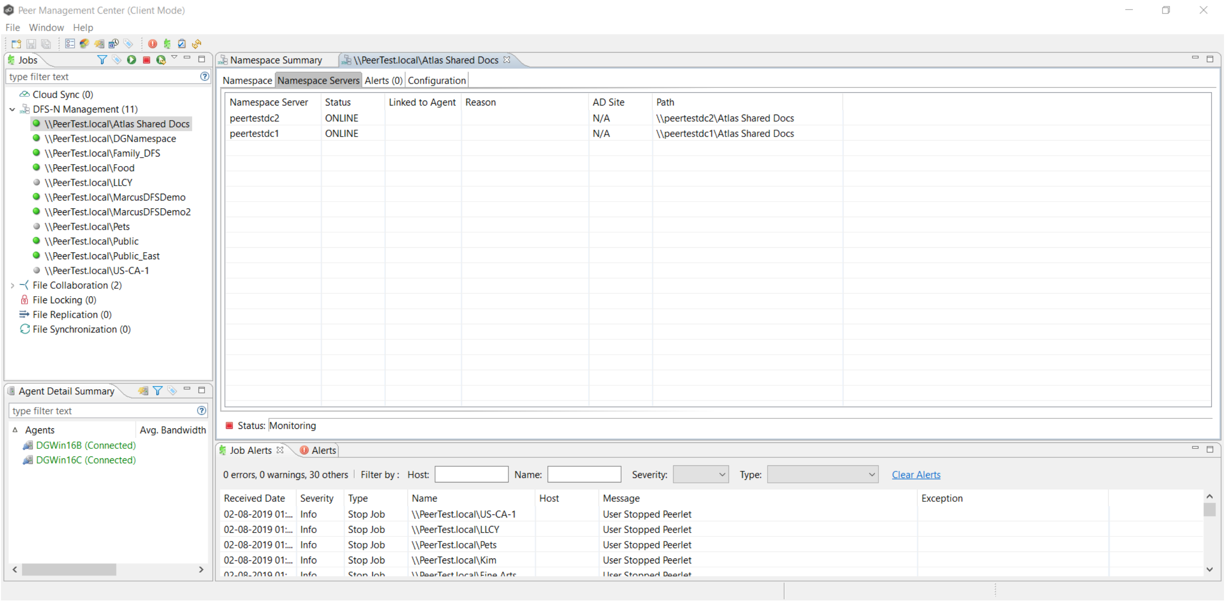

Namespace Server	Status	Linked to Agent	Reason	AD Site	Path
peertestdc2	ONLINE			N/A	\\peertestdc2\Atlas Shared Docs
peertestdc1	ONLINE			N/A	\\peertestdc1\Atlas Shared Docs

Status: Monitoring

0 errors, 0 warnings, 30 others | Filter by : Host: Name: Severity: Type:

Received Date	Severity	Type	Name	Host	Message	Exception
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\US-CA-1		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\LLCY		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Pets		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Kim		User Stopped Peerlet	
02-08-2019 01:	Info	Stop Job	\\PeerTest.local\Fine Arts		User Stopped Peerlet	

3. Right-click anywhere in the **Namespace Servers** tab, and then select **Add Servers**.

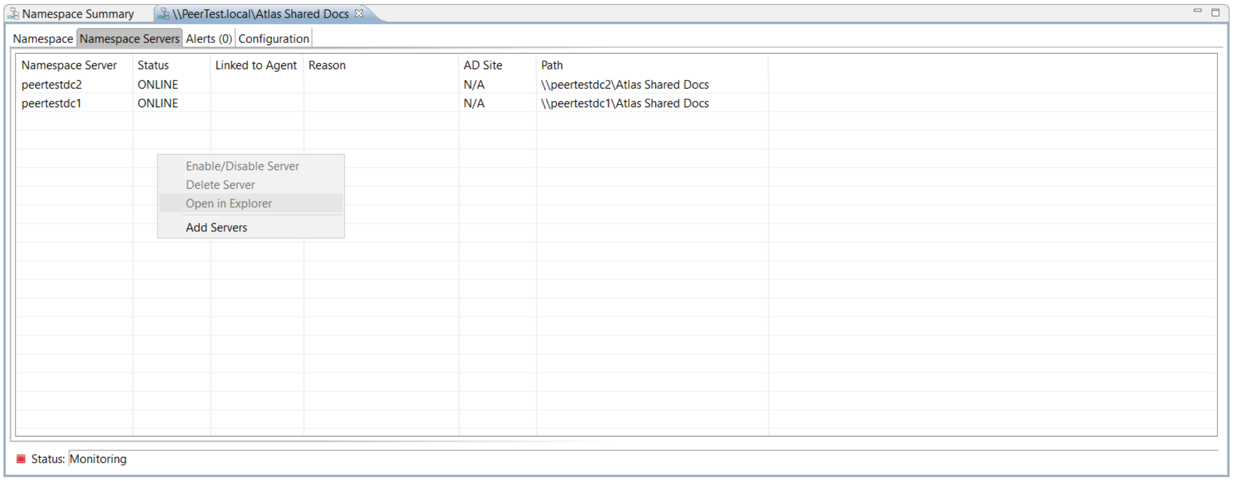

Namespace Server	Status	Linked to Agent	Reason	AD Site	Path
peertestdc2	ONLINE			N/A	\\peertestdc2\Atlas Shared Docs
peertestdc1	ONLINE			N/A	\\peertestdc1\Atlas Shared Docs

Status: Monitoring

The **Add DFS Namespace Server** wizard appears.

4. Enter the fully qualified path of a file server in the **Server Name** field, and then click **Add**.

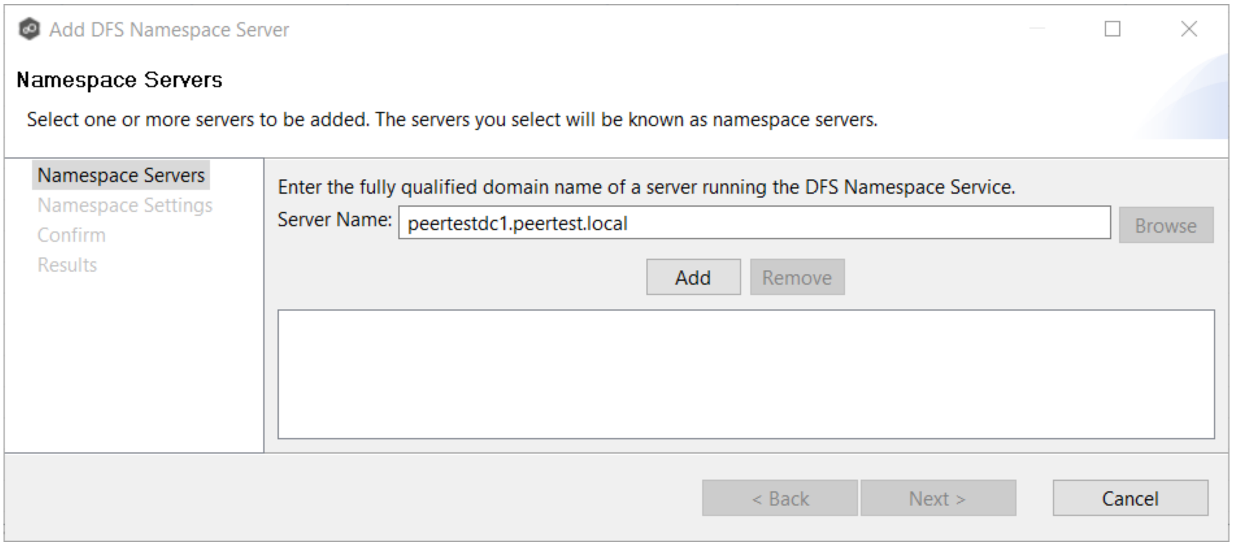

Add DFS Namespace Server

Namespace Servers

Select one or more servers to be added. The servers you select will be known as namespace servers.

Enter the fully qualified domain name of a server running the DFS Namespace Service.

Server Name: peertestdc1.peertest.local

The server path is listed in the area below.

5. Add additional servers if desired.

6. Click **Next**.

   The **Namespace Settings** page is displayed.

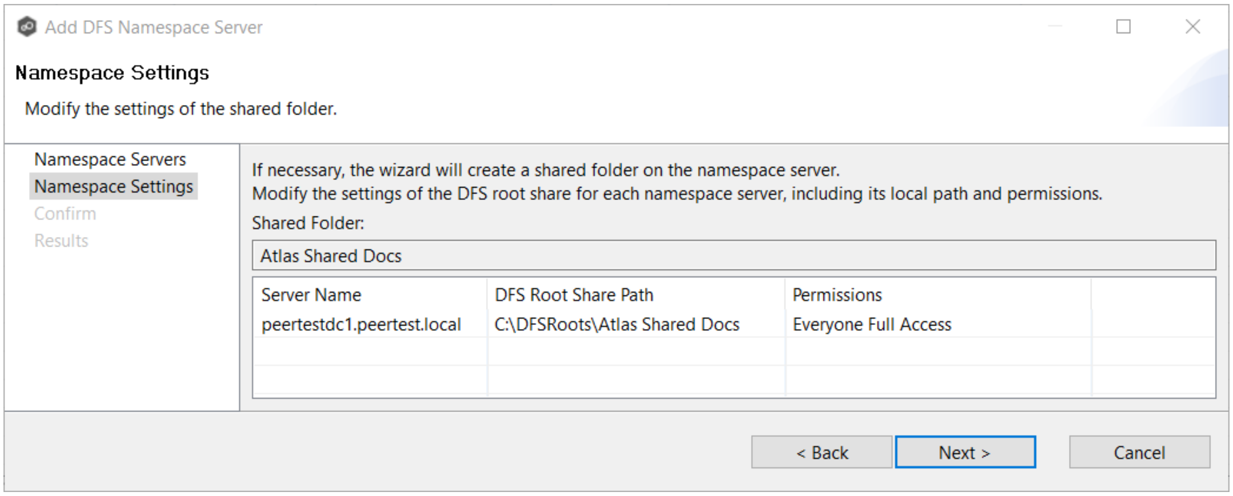

Add DFS Namespace Server

Namespace Settings

Modify the settings of the shared folder.

If necessary, the wizard will create a shared folder on the namespace server.
Modify the settings of the DFS root share for each namespace server, including its local path and permissions.

Shared Folder: Atlas Shared Docs

Server Name	DFS Root Share Path	Permissions
peertestdc1.peertest.local	C:\DFSRoots\Atlas Shared Docs	Everyone Full Access

7. (Optional) Edit the namespace server settings: **DFS Root Share Path** and **Permissions**.

8. Click **Next**.

   The **Confirm** page is displayed.

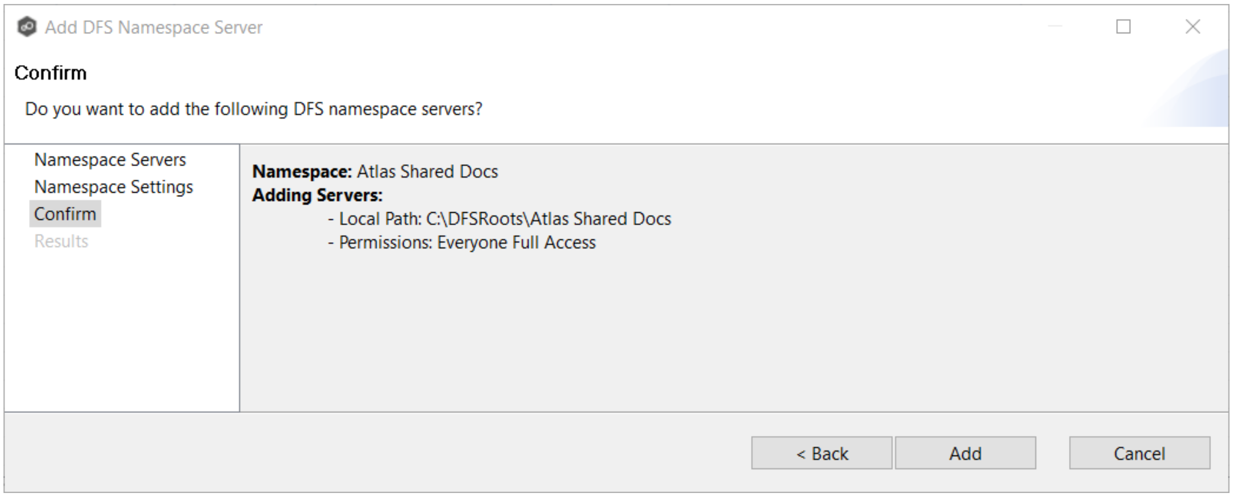

Do you want to add the following DFS namespace servers?

Namespace: Atlas Shared Docs
Adding Servers:
- Local Path: C:\DFSRoots\Atlas Shared Docs
- Permissions: Everyone Full Access

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.

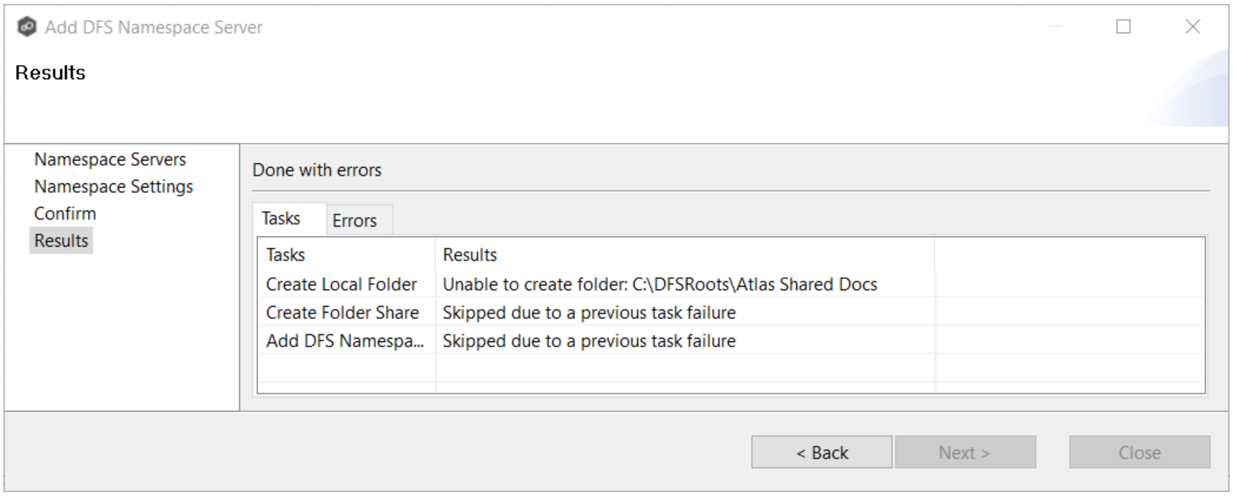

Done with errors

Tasks	Results
Create Local Folder	Unable to create folder: C:\DFSRoots\Atlas Shared Docs
Create Folder Share	Skipped due to a previous task failure
Add DFS Namespa...	Skipped due to a previous task failure

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.

To add a namespace folder to a namespace:

1. Double-click the job name in the **Jobs** view or the **Namespace Summary** view to open the run-time view for the namespace.

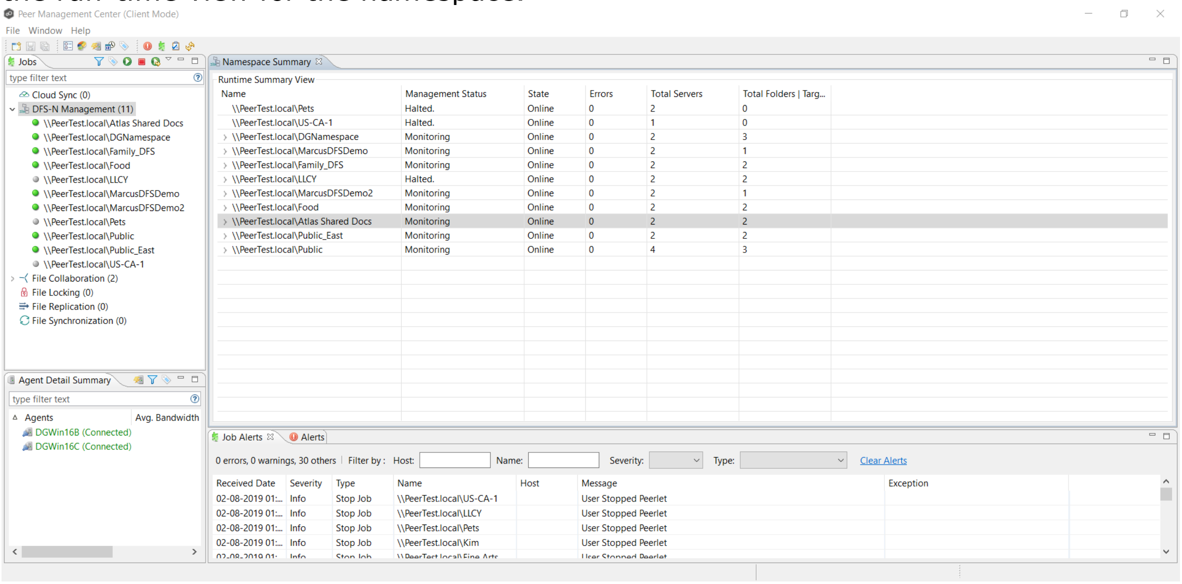

Runtime Summary View

Name	Management Status	State	Errors	Total Servers	Total Folders | Targ...
\\PeerTest.local\Pets	Halted.	Online	0	2	0
\\PeerTest.local\US-CA-1	Halted.	Online	0	1	0
\\PeerTest.local\DGNamespace	Monitoring	Online	0	2	3
\\PeerTest.local\MarcusDFSDemo	Monitoring	Online	0	2	1
\\PeerTest.local\Family_DFS	Monitoring	Online	0	2	2
\\PeerTest.local\LLCY	Halted.	Online	0	2	2
\\PeerTest.local\MarcusDFSDemo2	Monitoring	Online	0	2	1
\\PeerTest.local\Food	Monitoring	Online	0	2	2
\\PeerTest.local\Atlas Shared Docs	Monitoring	Online	0	2	2
\\PeerTest.local\Public_East	Monitoring	Online	0	2	2
\\PeerTest.local\Public	Monitoring	Online	0	4	3

0 errors, 0 warnings, 30 others | Filter by : Host: Name: Severity: Type:

Received Date	Severity	Type	Name	Host	Message	Exception
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\US-CA-1		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\LLCY		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Pets		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Kim		User Stopped Peerlet	

The job's run-time view is displayed.

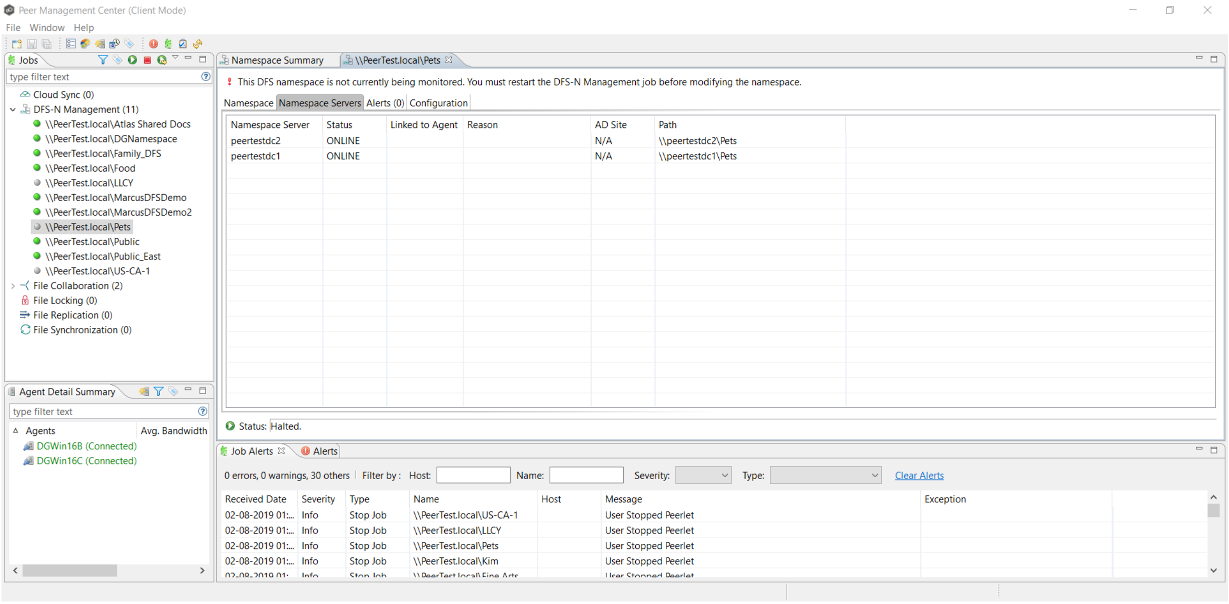

This DFS namespace is not currently being monitored. You must restart the DFS-N Management job before modifying the namespace.

Namespace Server	Status	Linked to Agent	Reason	AD Site	Path
peertestdc2	ONLINE			N/A	\\peertestdc2\Pets
peertestdc1	ONLINE			N/A	\\peertestdc1\Pets

Status: Halted.

0 errors, 0 warnings, 30 others | Filter by : Host: Name: Severity: Type:

Received Date	Severity	Type	Name	Host	Message	Exception
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\US-CA-1		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\LLCY		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Pets		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Kim		User Stopped Peerlet	

2. Right-click anywhere in the **Namespace** tab, and then select **Add Folder**.

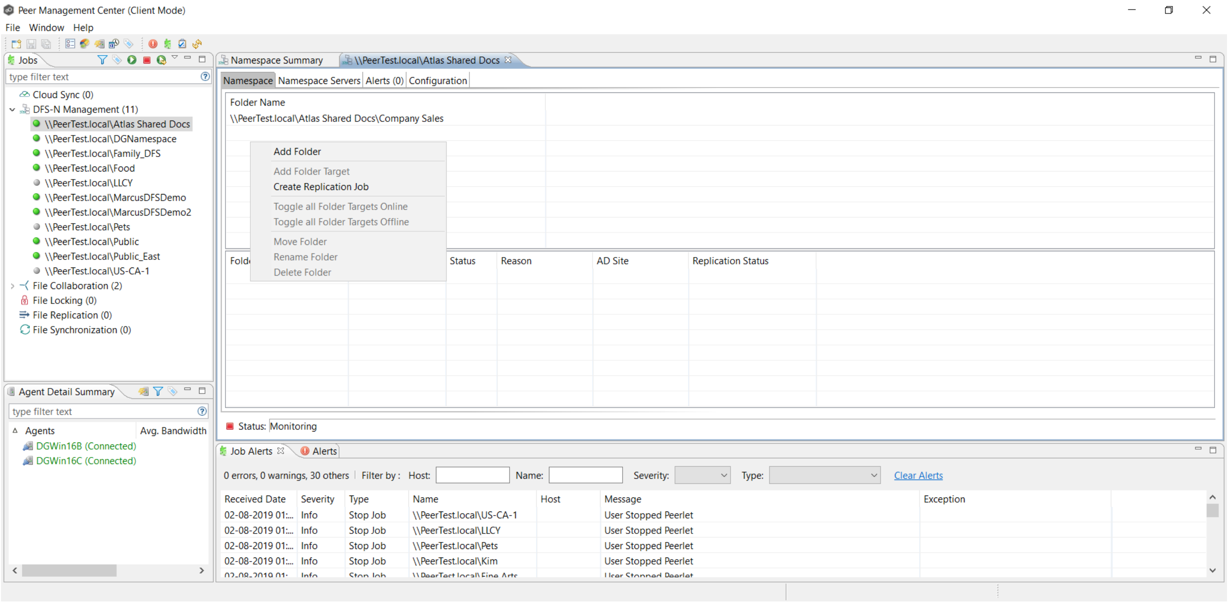

Folder Name
\\PeerTest.local\Atlas Shared Docs\Company Sales

Status: Monitoring

0 errors, 0 warnings, 30 others

Received Date	Severity	Type	Name	Host	Message	Exception
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\US-CA-1		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\LLCY		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Pets		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Kim		User Stopped Peerlet	
02-08-2019 01:	Info	Stop Job	\\PeerTest.local\Fine Arts		User Stopped Peerlet	

The **New Namespace Folder** wizard appears.

3. Enter a name for the namespace folder in the **Folder Name** field.

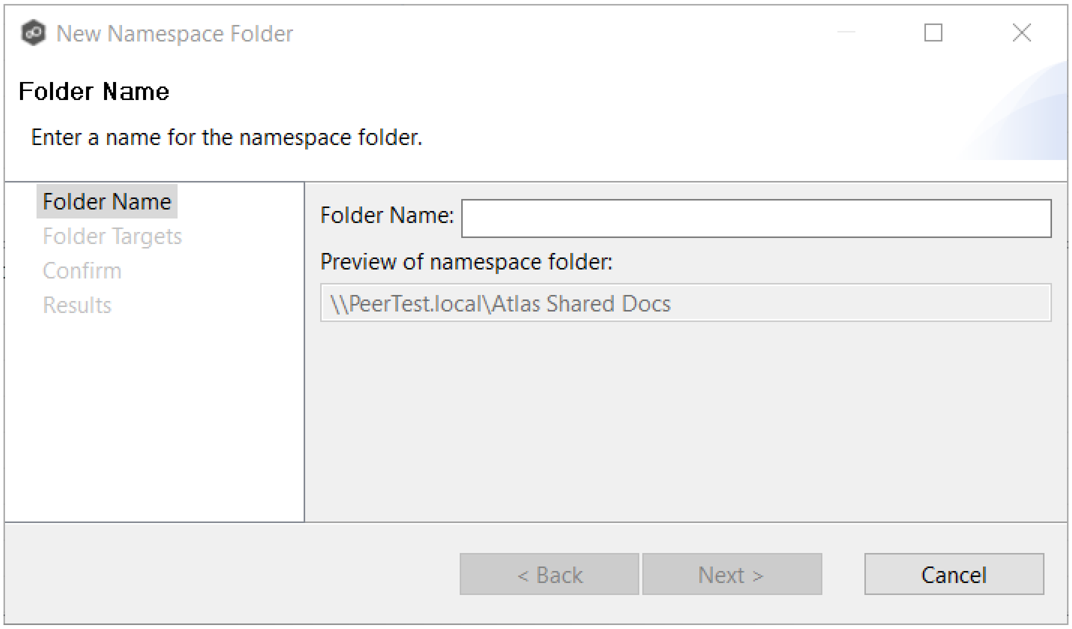

Folder Name

Enter a name for the namespace folder.

Folder Name:

Preview of namespace folder:

\\PeerTest.local\Atlas Shared Docs

After you enter the folder name, a preview of the folder and path name appears below the **Folder Name** field.

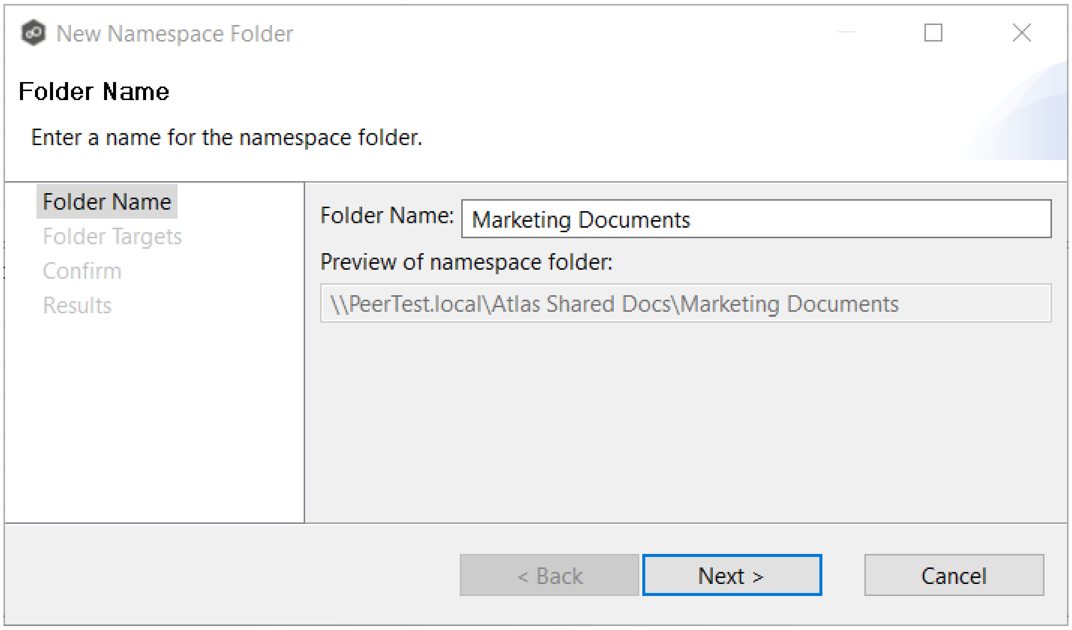

New Namespace Folder

Folder Name

Enter a name for the namespace folder.

Folder Name: Marketing Documents

Preview of namespace folder:

\\PeerTest.local\Atlas Shared Docs\Marketing Documents

4. Click **Next**.

   The **Folder Targets** page is displayed.

5. Enter the UNC path to a shared folder, and then click **Add**.

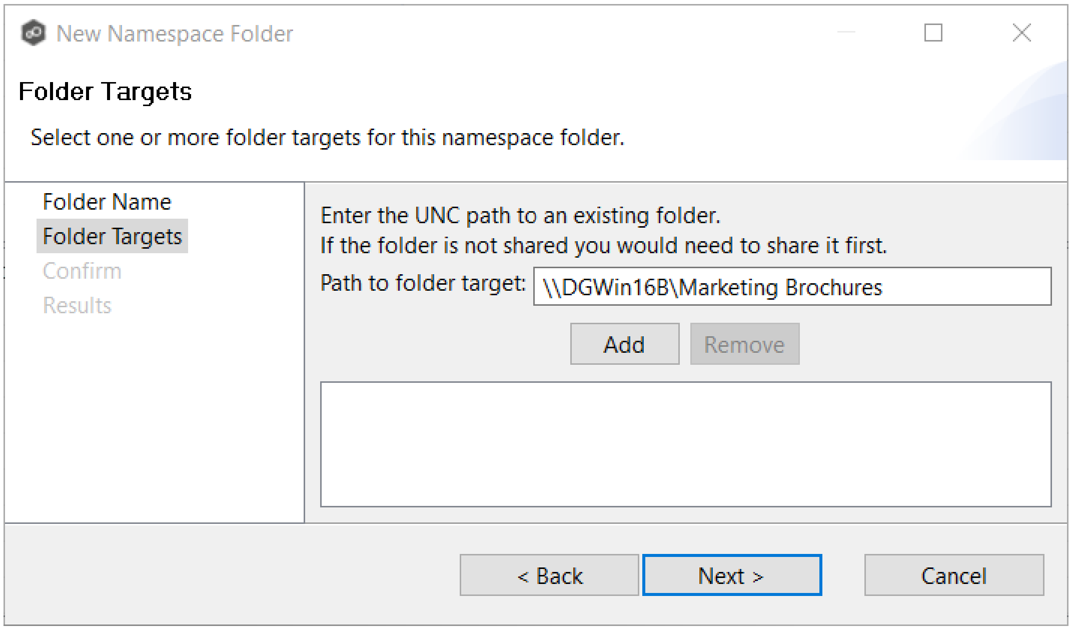

New Namespace Folder

Folder Targets

Select one or more folder targets for this namespace folder.

Enter the UNC path to an existing folder.
If the folder is not shared you would need to share it first.

Path to folder target: \\DGWin16B\Marketing Brochures

The folder target path is listed in the field below.

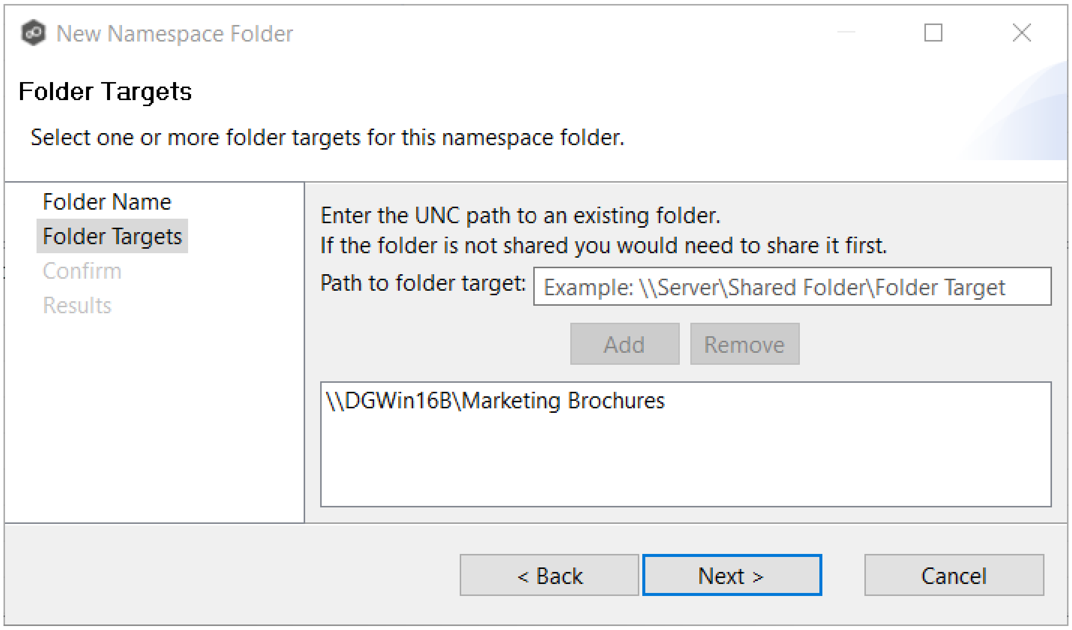
New Namespace Folder

Folder Targets

Select one or more folder targets for this namespace folder.

Enter the UNC path to an existing folder.
If the folder is not shared you would need to share it first.

Path to folder target: Example: \\Server\Shared Folder\Folder Target

\\DGWin16B\Marketing Brochures

6. (Optional) Add additional folder targets.

7. Click **Next**.

   The **Confirm** page is displayed.

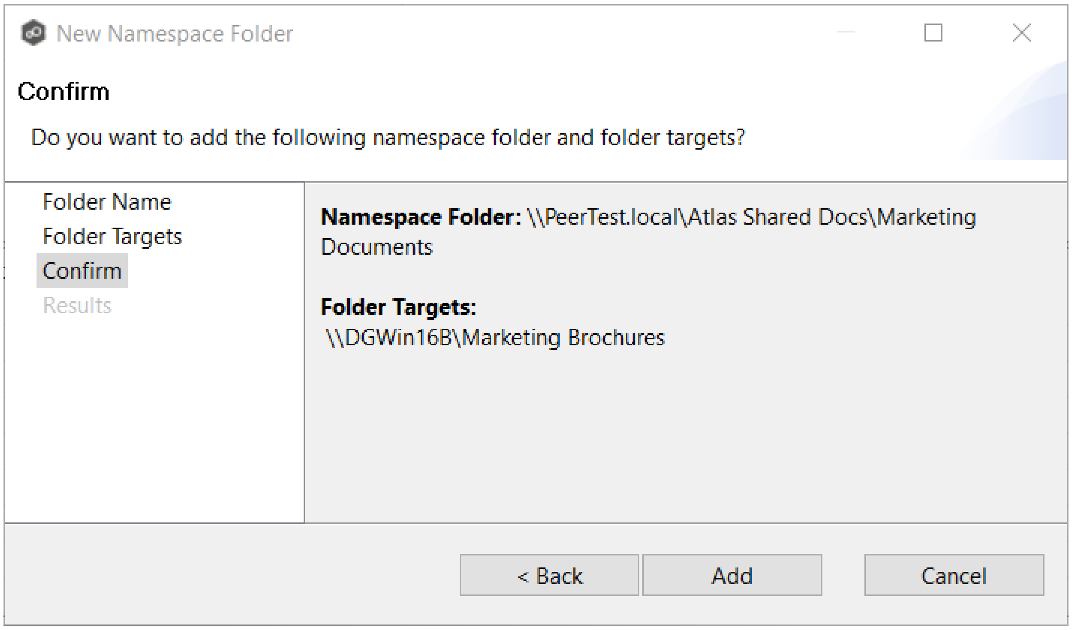
New Namespace Folder

Confirm

Do you want to add the following namespace folder and folder targets?

Namespace Folder: \\PeerTest.local\Atlas Shared Docs\Marketing Documents

Folder Targets:
\\DGWin16B\Marketing Brochures

8. Review the folders and folder targets.

9. Click **Add** if the configuration is correct; otherwise, click **Back** and correct the configuration.

The **Results** page is displayed.

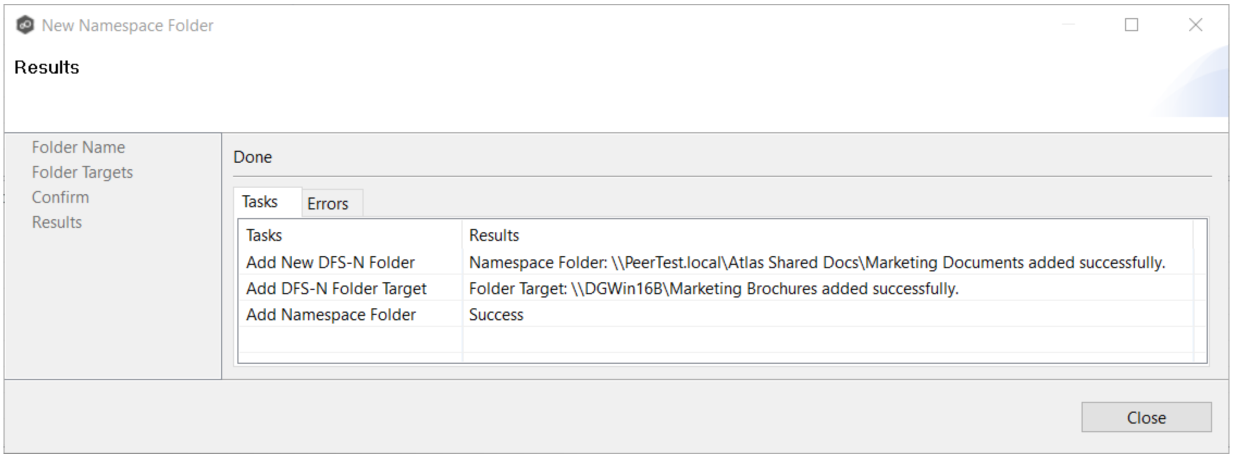

Done

Tasks	Results
Add New DFS-N Folder	Namespace Folder: \\PeerTest.local\Atlas Shared Docs\Marketing Documents added successfully.
Add DFS-N Folder Target	Folder Target: \\DGWin16B\Marketing Brochures added successfully.
Add Namespace Folder	Success

10. Click **Close**.

    The newly added folder and folder targets are listed in the **Namespace** tab.

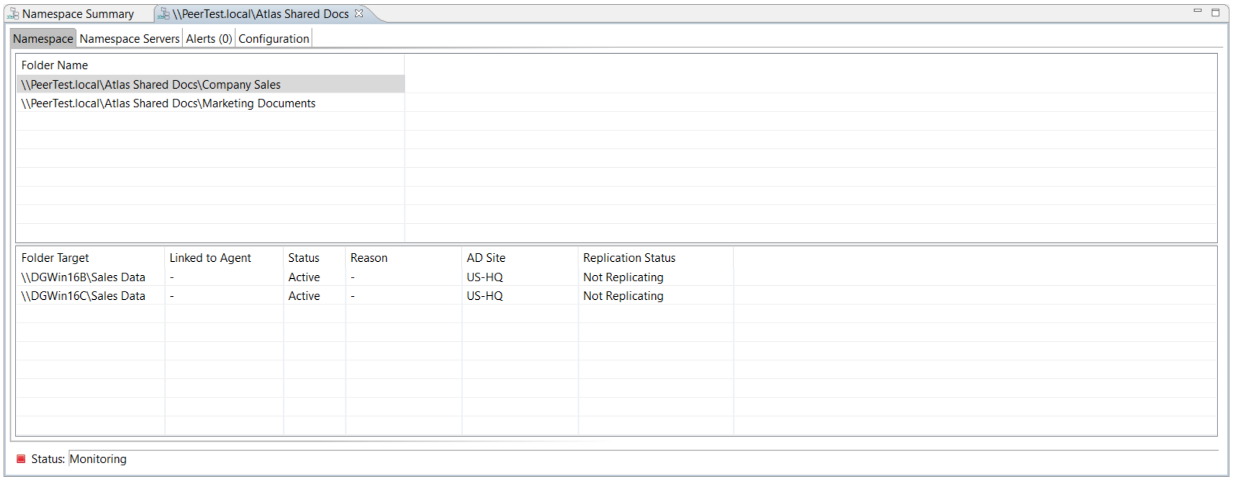

Folder Name
\\PeerTest.local\Atlas Shared Docs\Company Sales
\\PeerTest.local\Atlas Shared Docs\Marketing Documents

Folder Target	Linked to Agent	Status	Reason	AD Site	Replication Status
\\DGWin16B\Sales Data	-	Active	-	US-HQ	Not Replicating
\\DGWin16C\Sales Data	-	Active	-	US-HQ	Not Replicating

Status: Monitoring

Adding a Namespace Folder Target

You can add a folder target to a namespace.

To add a folder target to a namespace:

1. Double-click the job name in the **Jobs** view or the **Namespace Summary** view to open the run-time view for the namespace.

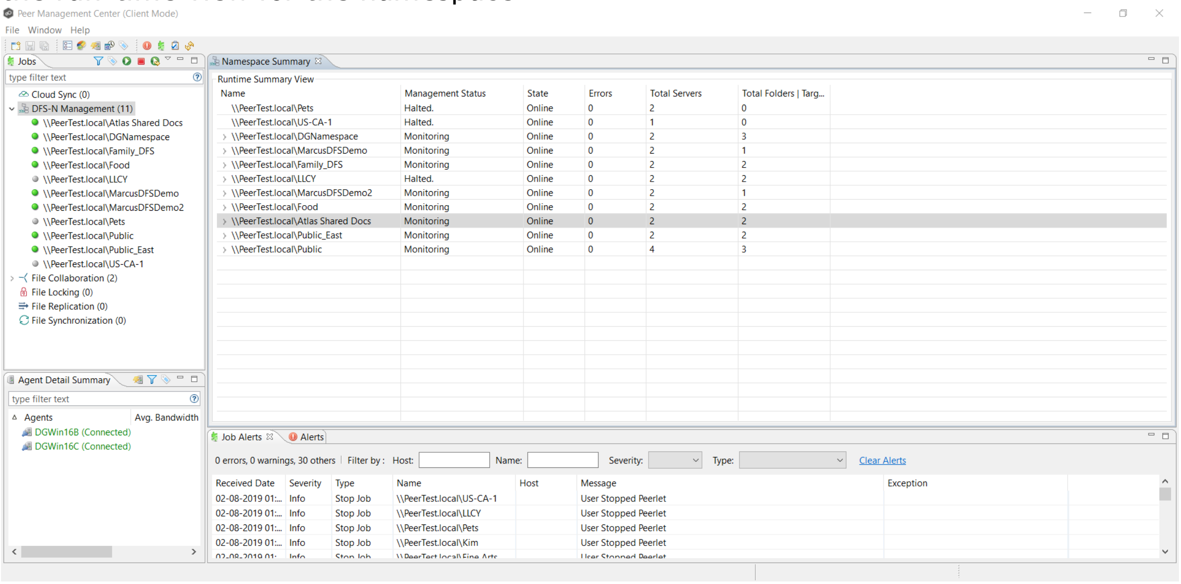

Runtime Summary View

Name	Management Status	State	Errors	Total Servers	Total Folders | Targ...
\\PeerTest.local\Pets	Halted.	Online	0	2	0
\\PeerTest.local\US-CA-1	Halted.	Online	0	1	0
\\PeerTest.local\DGNamespace	Monitoring	Online	0	2	3
\\PeerTest.local\MarcusDFSDemo	Monitoring	Online	0	2	1
\\PeerTest.local\Family_DFS	Monitoring	Online	0	2	2
\\PeerTest.local\LLCY	Halted.	Online	0	2	2
\\PeerTest.local\MarcusDFSDemo2	Monitoring	Online	0	2	1
\\PeerTest.local\Food	Monitoring	Online	0	2	2
\\PeerTest.local\Atlas Shared Docs	Monitoring	Online	0	2	2
\\PeerTest.local\Public_East	Monitoring	Online	0	2	2
\\PeerTest.local\Public	Monitoring	Online	0	4	3

Job Alerts | Alerts

0 errors, 0 warnings, 30 others | Filter by : Host: Name: Severity: Type: Clear Alerts

Received Date	Severity	Type	Name	Host	Message	Exception
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\US-CA-1		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\LLCY		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Pets		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Kim		User Stopped Peerlet	
02-08-2019 01:	Info	Stop Job	\\PeerTest.local\Fine Arts		User Stopped Peerlet	

The job's run-time view is displayed.

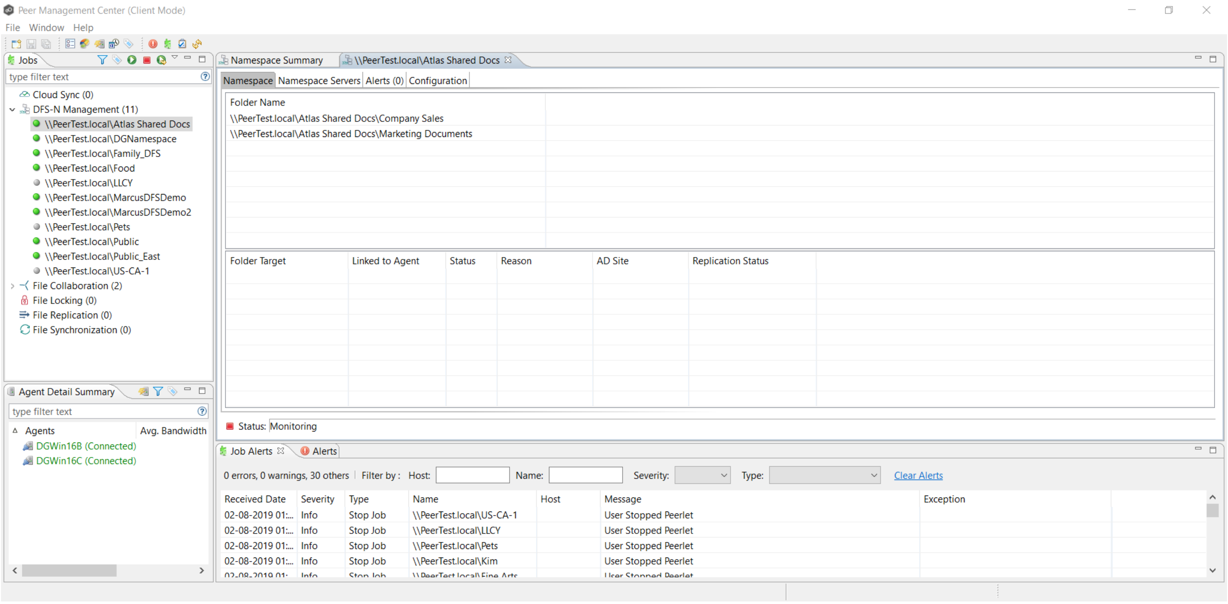

Folder Name
\\PeerTest.local\Atlas Shared Docs\Company Sales
\\PeerTest.local\Atlas Shared Docs\Marketing Documents

Folder Target	Linked to Agent	Status	Reason	AD Site	Replication Status

Status: Monitoring

Job Alerts | Alerts

0 errors, 0 warnings, 30 others | Filter by : Host: Name: Severity: Type: Clear Alerts

Received Date	Severity	Type	Name	Host	Message	Exception
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\US-CA-1		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\LLCY		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Pets		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Kim		User Stopped Peerlet	
02-08-2019 01:	Info	Stop Job	\\PeerTest.local\Fine Arts		User Stopped Peerlet	

2. Right-click the folder you want to add a folder target to, and then select **Add Folder Target**.

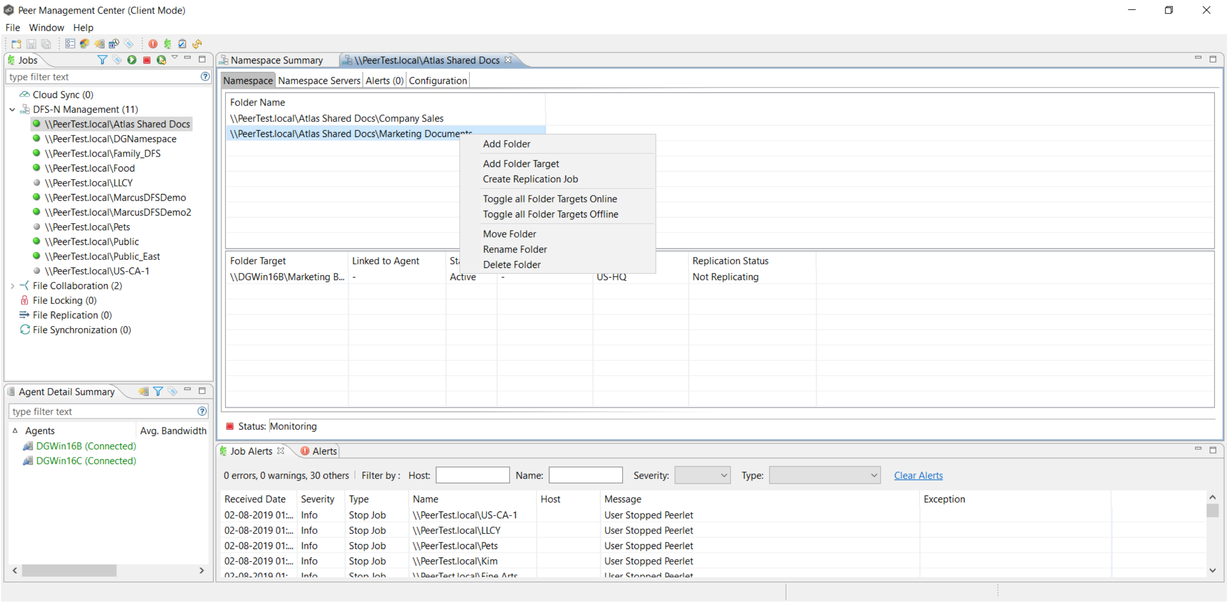

Folder Name
\\PeerTest.local\Atlas Shared Docs\Company Sales
\\PeerTest.local\Atlas Shared Docs\Marketing Docume...

Folder Target	Linked to Agent	St			Replication Status
\\DGWin16B\Marketing B...	-	Active	-	US-HQ	Not Replicating

0 errors, 0 warnings, 30 others

Received Date	Severity	Type	Name	Host	Message	Exception
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\US-CA-1		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\LLCY		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Pets		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Kim		User Stopped Peerlet	
02-08-2019 01:...	Info	Stop Job	\\PeerTest.local\Fine Arts		User Stopped Peerlet	

The **New Folder Target** wizard appears.

3. Enter the UNC path to a shared folder, and then click **Add**.

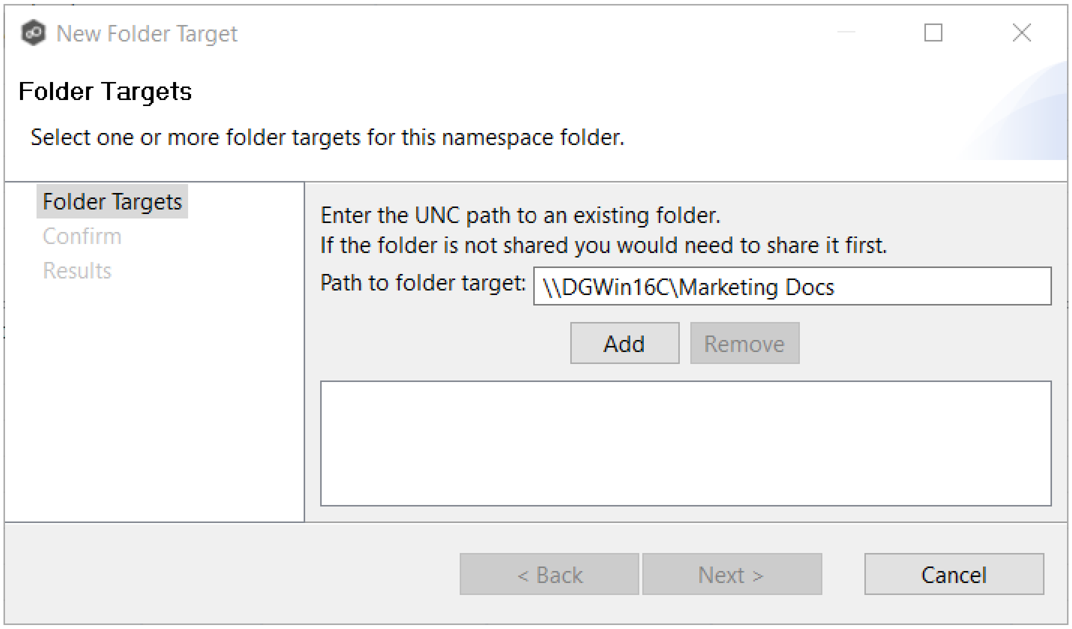

Folder Targets

Select one or more folder targets for this namespace folder.

Enter the UNC path to an existing folder.
If the folder is not shared you would need to share it first.

Path to folder target: \\DGWin16C\Marketing Docs

The folder target path is listed in the field below.

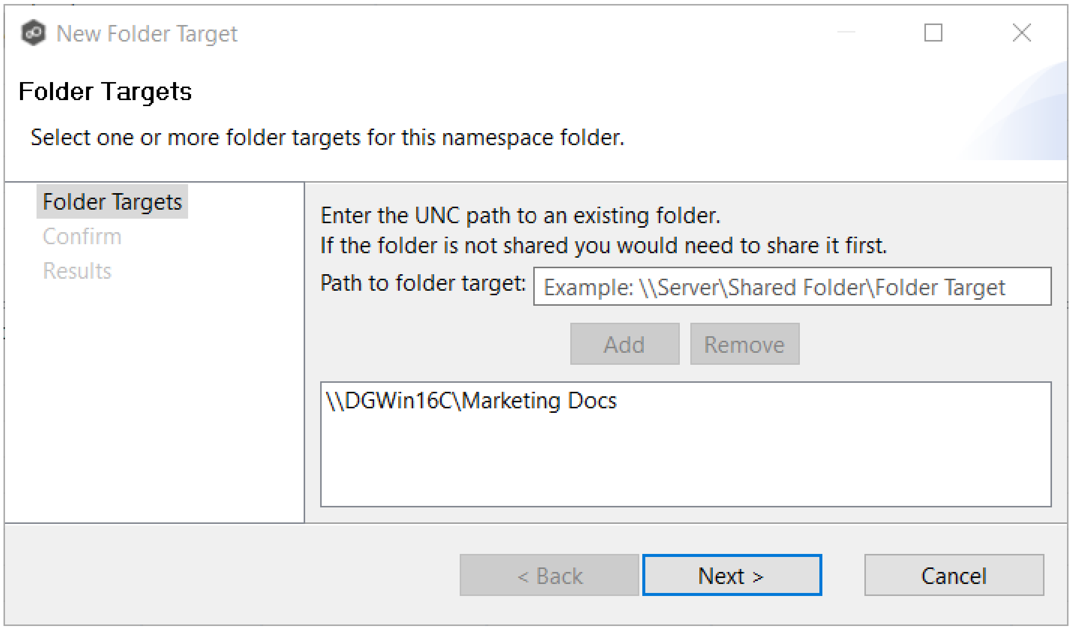

New Folder Target

Folder Targets

Select one or more folder targets for this namespace folder.

Enter the UNC path to an existing folder.
If the folder is not shared you would need to share it first.

Path to folder target: Example: \\Server\Shared Folder\Folder Target

\\DGWin16C\Marketing Docs

4. (Optional) Add additional folder targets.

5. Click **Next**.

   The **Confirm** page is displayed.

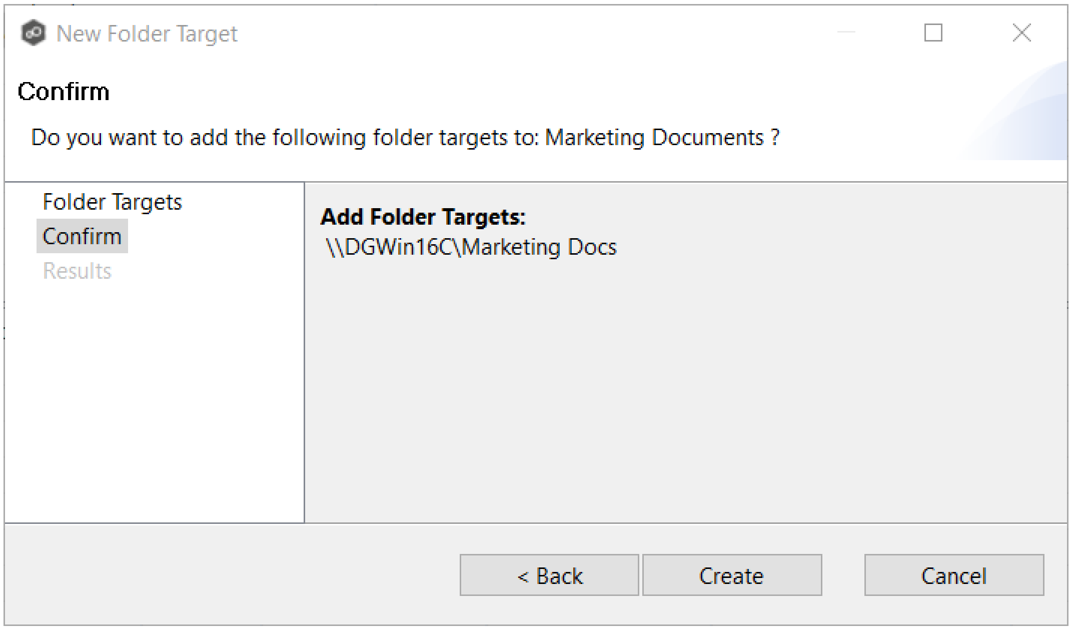

New Folder Target

Confirm

Do you want to add the following folder targets to: Marketing Documents ?

Add Folder Targets:
\\DGWin16C\Marketing Docs

6. Review the folder targets.

7. Click **Create** if the configuration is correct; otherwise, click **Back** and correct the configuration.

The **Results** page is displayed.

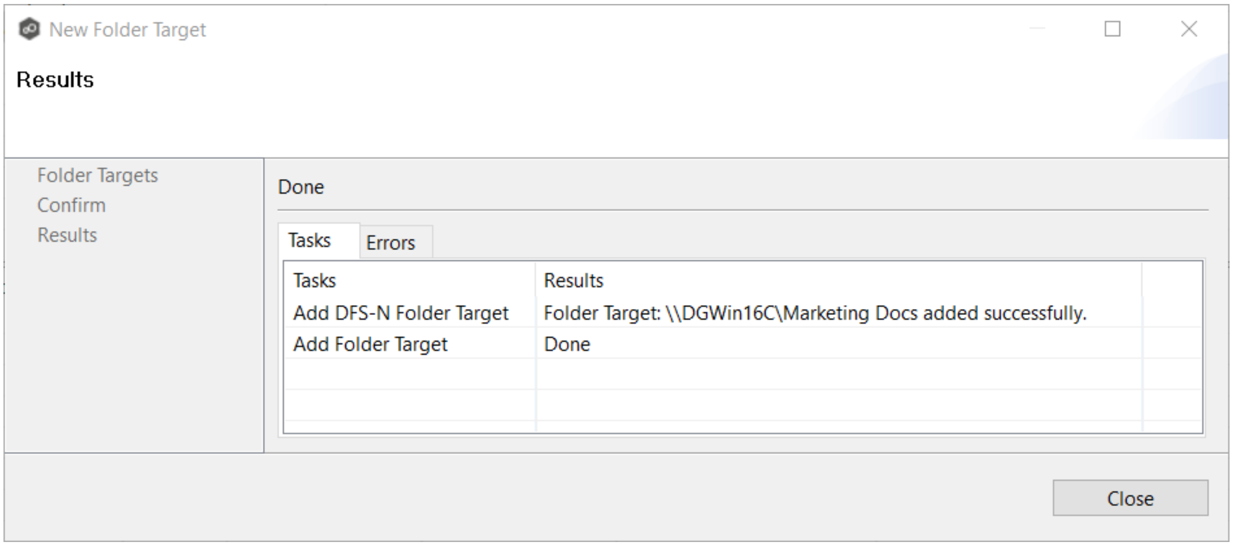

Done

Tasks | Errors

Tasks	Results
Add DFS-N Folder Target	Folder Target: \\DGWin16C\Marketing Docs added successfully.
Add Folder Target	Done

8. Click **Close**.

   The newly added folder targets are listed in the **Namespace** tab.

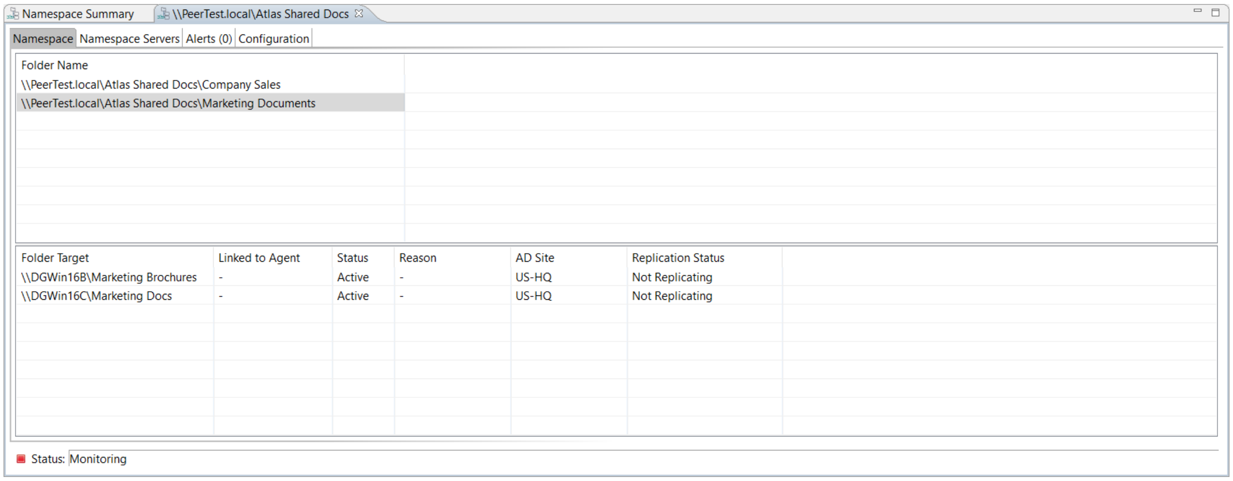

Folder Name
\\PeerTest.local\Atlas Shared Docs\Company Sales
\\PeerTest.local\Atlas Shared Docs\Marketing Documents

Folder Target	Linked to Agent	Status	Reason	AD Site	Replication Status
\\DGWin16B\Marketing Brochures	-	Active	-	US-HQ	Not Replicating
\\DGWin16C\Marketing Docs	-	Active	-	US-HQ	Not Replicating

Status: Monitoring

Connecting DFS Namespaces with File Collaboration and File Synchronization Jobs

In order to allow the PMC synchronization engine to automate the state of folder targets, a File Collaboration or file synchronization job must be linked to a job that manages the appropriate DFS namespace.

The two main ways to create this link are:

- If the File Collaboration or file synchronization job already exists, edit the job and use the DFS-N settings page to link the collaboration or synchronization to the DFS-N Management job. See Link a Namespace with an Existing File Collaboration or Synchronization Job for step-by-step instructions.

- If the File Collaboration or synchronization job does not yet exist, create one from the DFS namespace folder. See Create a File Collaboration or Synchronization Job from a DFS Namespace Folder for step-by-step instructions.

Note: Currently, only two job types, File Collaboration and File Synchronization, can be linked to a DFS-N Management job.

You can link a DFS namespace 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 the PMC.

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 Configuration** wizard appears.

3. Select **DFS-N** in the navigation tree.

   The following page is displayed:

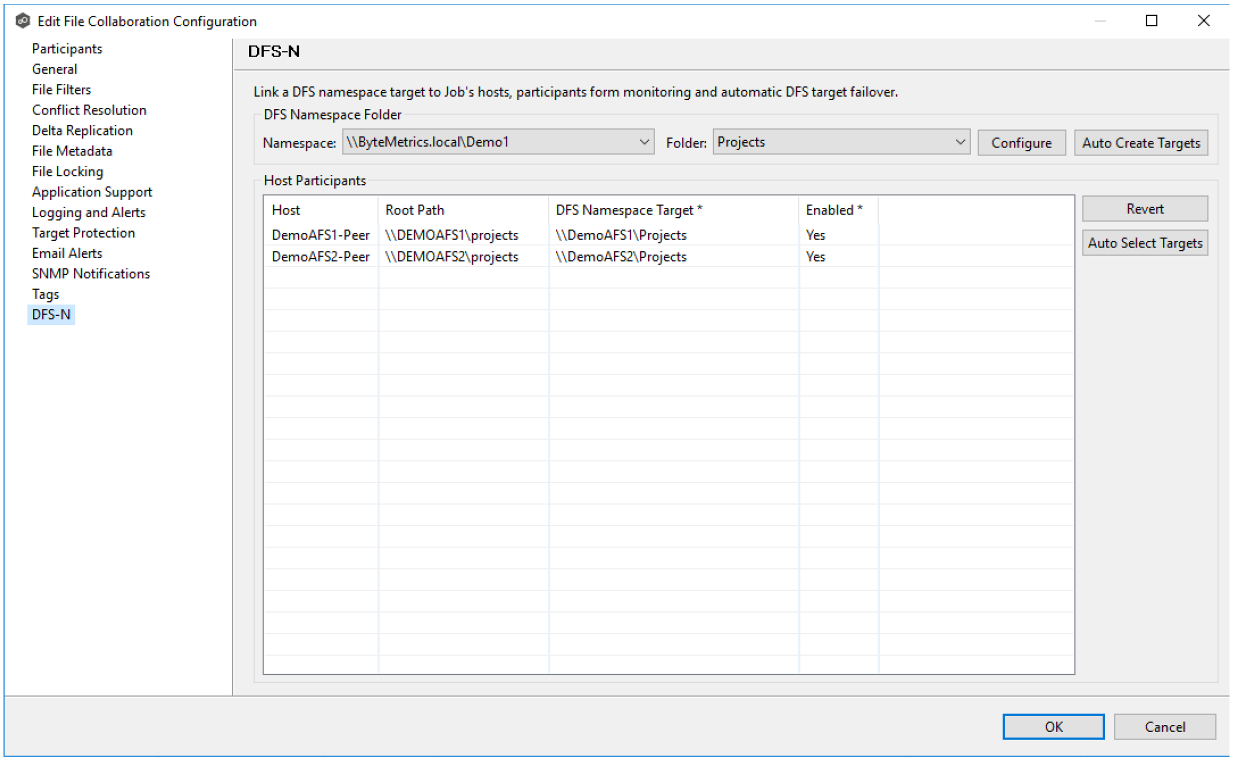

DFS-N

Link a DFS namespace target to Job's hosts, participants form monitoring and automatic DFS target failover.

DFS Namespace Folder

Namespace: \\ByteMetrics.local\Demo1 Folder: Projects

Host Participants

Host	Root Path	DFS Namespace Target *	Enabled *
DemoAFS1-Peer	\\DEMOAFS1\projects	\\DemoAFS1\Projects	Yes
DemoAFS2-Peer	\\DEMOAFS2\projects	\\DemoAFS2\Projects	Yes

4. In the **DFS Namespace Folder** area, select the namespace you want to link to from the first drop down list.

   Note: If your desired namespace does not exist, you either need to create it (via the Create DFS-N Management Job wizard) or you can import an existing namespace into the PMC.

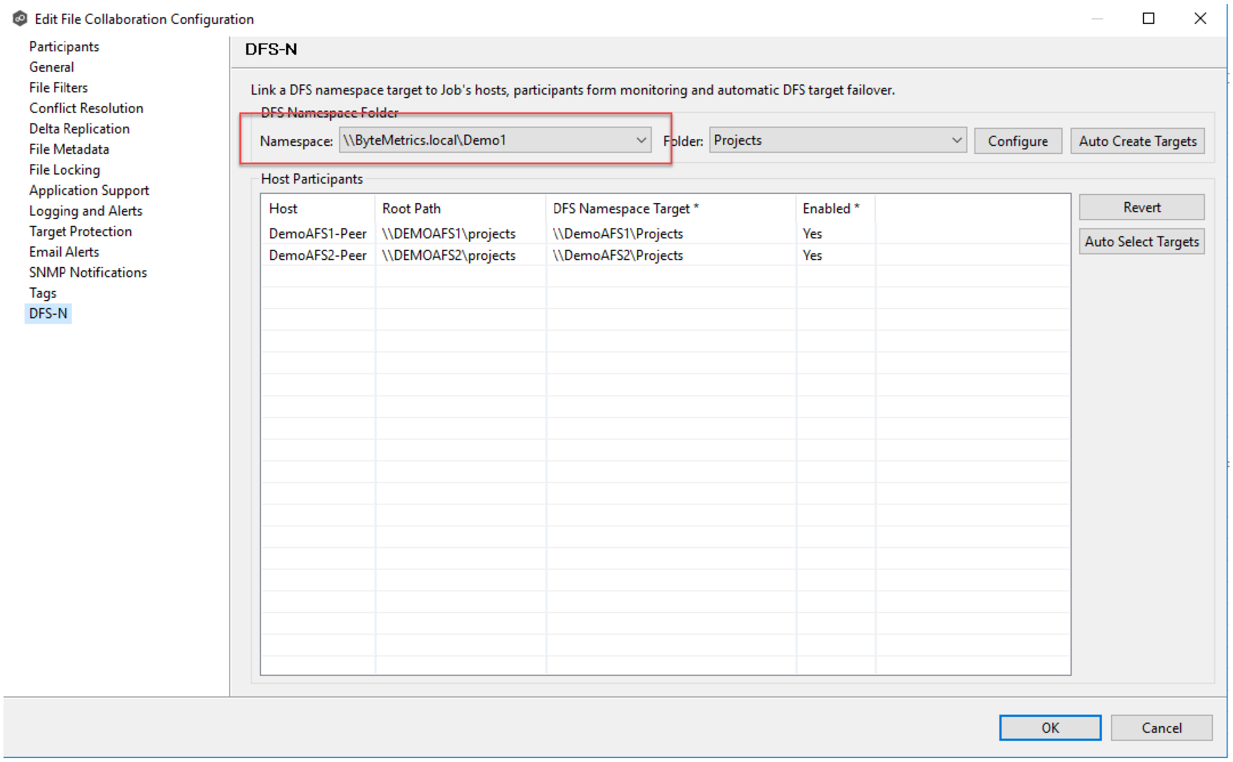

DFS-N

Link a DFS namespace target to Job's hosts, participants form monitoring and automatic DFS target failover.

DFS Namespace Folder

Namespace: \\ByteMetrics.local\Demo1

Folder: Projects

Host Participants

Host	Root Path	DFS Namespace Target *	Enabled *
DemoAFS1-Peer	\\DEMOAFS1\projects	\\DemoAFS1\Projects	Yes
DemoAFS2-Peer	\\DEMOAFS2\projects	\\DemoAFS2\Projects	Yes

Once you’ve selected a namespace, a list of namespace folders are available in the second drop-down list.

5. Select the namespace folder.

   If your desired folder does not show up in this list, click the **Configure** button to open a wizard to make changes to the selected namespace (including the ability to add folders and folder targets).

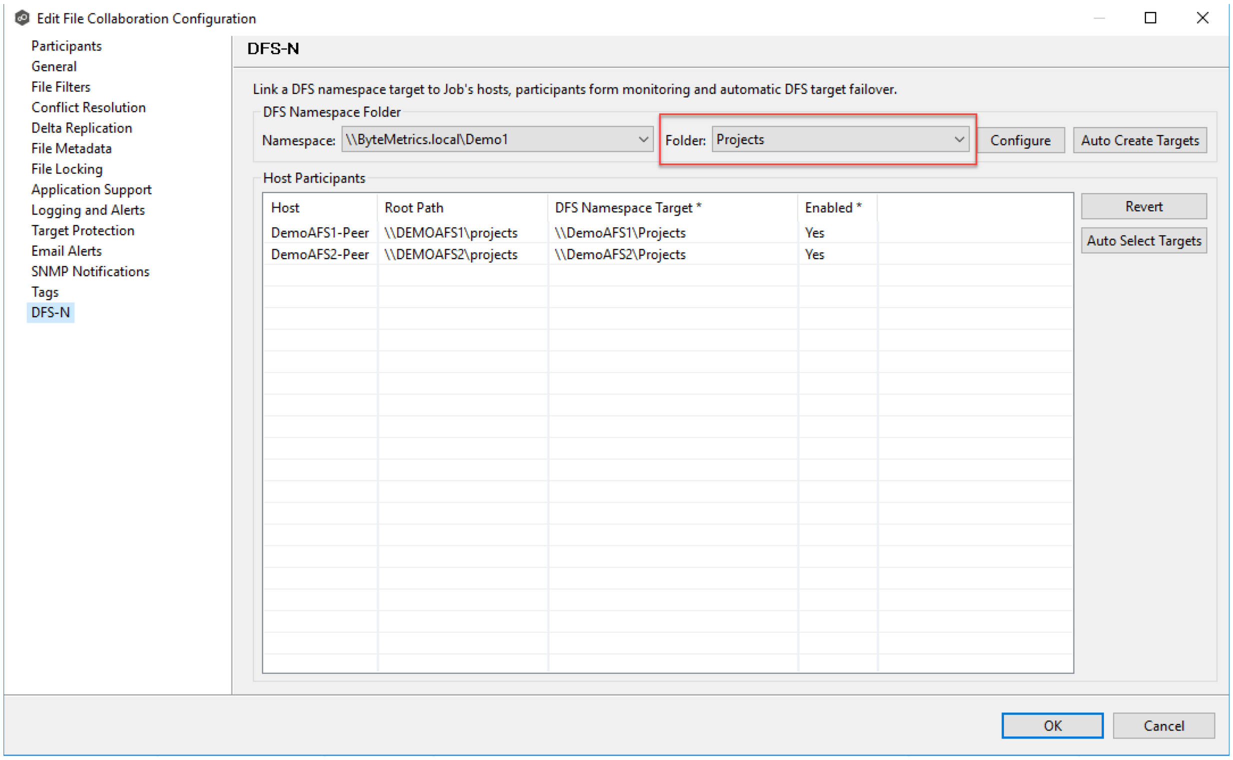

DFS-N

Link a DFS namespace target to Job's hosts, participants form monitoring and automatic DFS target failover.

DFS Namespace Folder

Namespace: \\ByteMetrics.local\Demo1

Folder: Projects

Host Participants

Host	Root Path	DFS Namespace Target *	Enabled *
DemoAFS1-Peer	\\DEMOAFS1\projects	\\DemoAFS1\Projects	Yes
DemoAFS2-Peer	\\DEMOAFS2\projects	\\DemoAFS2\Projects	Yes

If your desired folder does not have the appropriate folder targets, click the **Auto Create Targets** button. This 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:

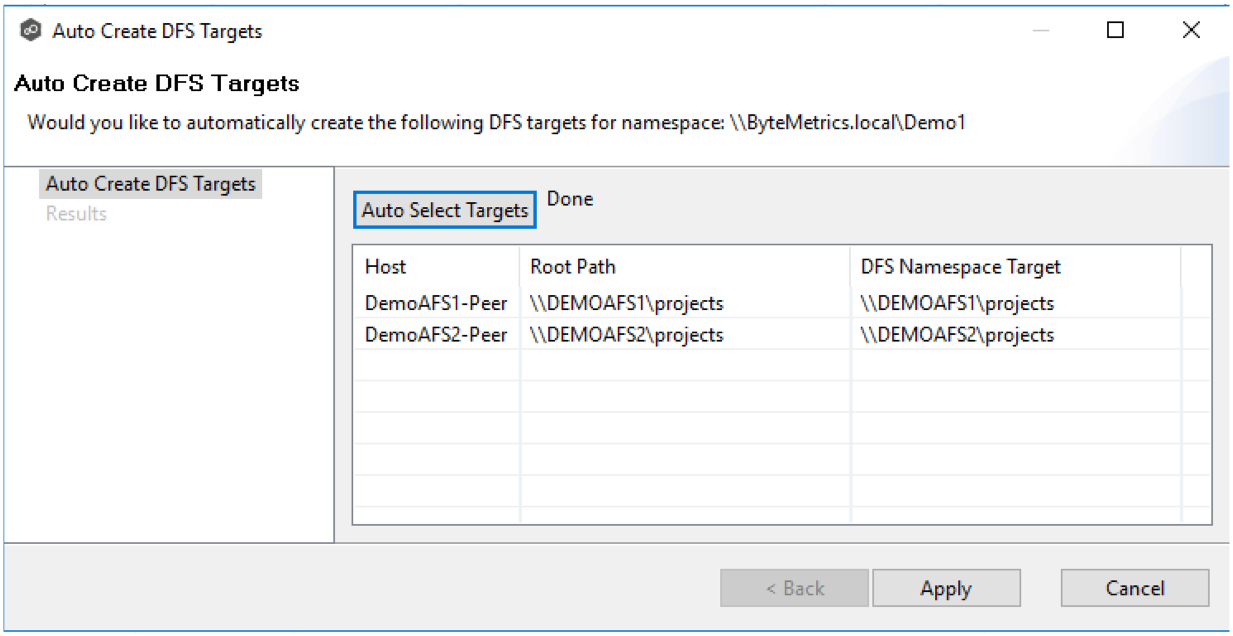

Auto Create DFS Targets

Would you like to automatically create the following DFS targets for namespace: \\ByteMetrics.local\Demo1

Done

Host	Root Path	DFS Namespace Target
DemoAFS1-Peer	\\DEMOAFS1\projects	\\DEMOAFS1\projects
DemoAFS2-Peer	\\DEMOAFS2\projects	\\DEMOAFS2\projects

6. Once you've selected a namespace and a folder, you will need to assign a target to each participant in the collaboration or synchronization job. In most scenarios, clicking the

Auto Select Targets button will be able to automatically link a folder target with the appropriate participant.

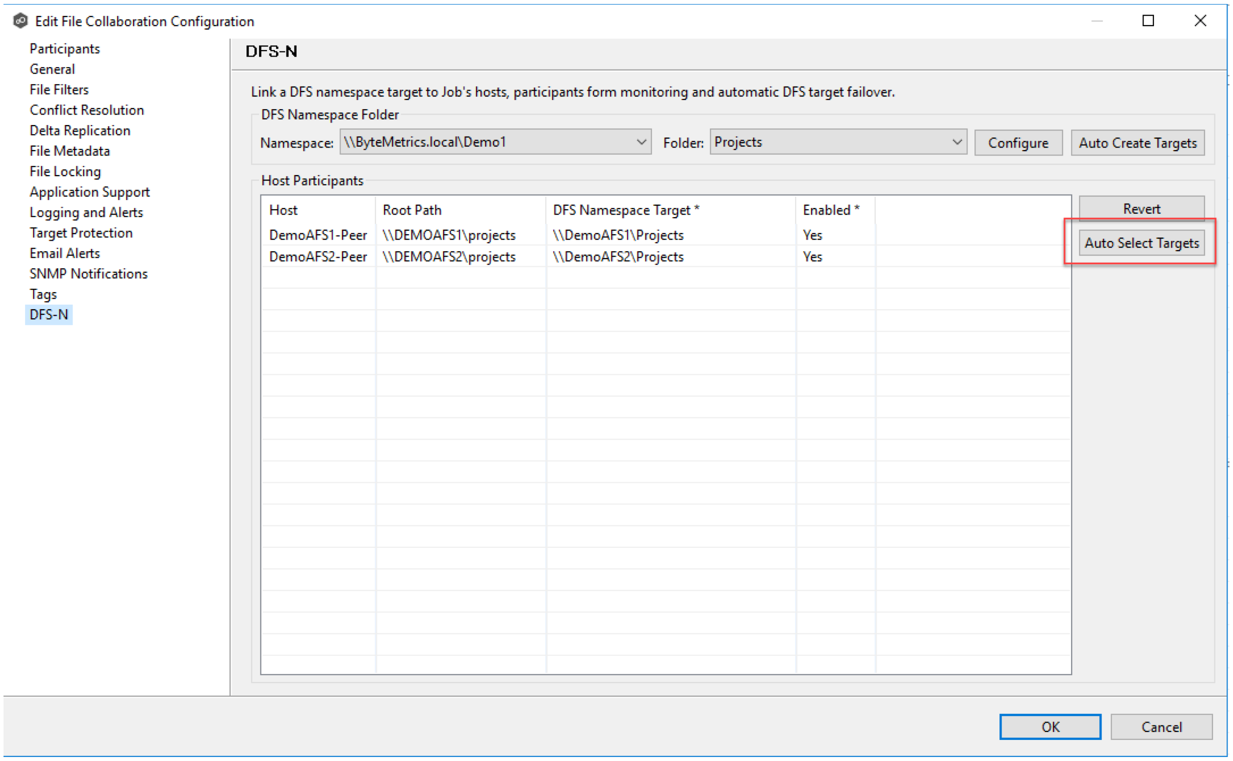

DFS-N

Link a DFS namespace target to Job's hosts, participants form monitoring and automatic DFS target failover.

Namespace: \\ByteMetrics.local\Demo1 Folder: Projects

Host	Root Path	DFS Namespace Target *	Enabled *
DemoAFS1-Peer	\\DEMOAFS1\projects	\\DemoAFS1\Projects	Yes
DemoAFS2-Peer	\\DEMOAFS2\projects	\\DemoAFS2\Projects	Yes

7. Once all participants are linked to the appropriate folder targets, click **OK** to save your changes.

   From this point forward, if this collaboration or synchronization job is running along with its paired Namespace job, the PMC will automatically failover and failback folder targets.

You can create a File Collaboration or File Synchronization job from a DFS namespace folder. These steps require that a DFS namespace has been already created and is being managed by the PMC. Collaboration and synchronization jobs can be created only with a namespace folder that has at least two folder targets.

1. From the **Jobs** view, open the DFS-N Management job managing the namespace.

   The job's namespace management view appears.

2. In the **Namespace** tab, right-click the desired namespace folder and select **Create Replication Job**.

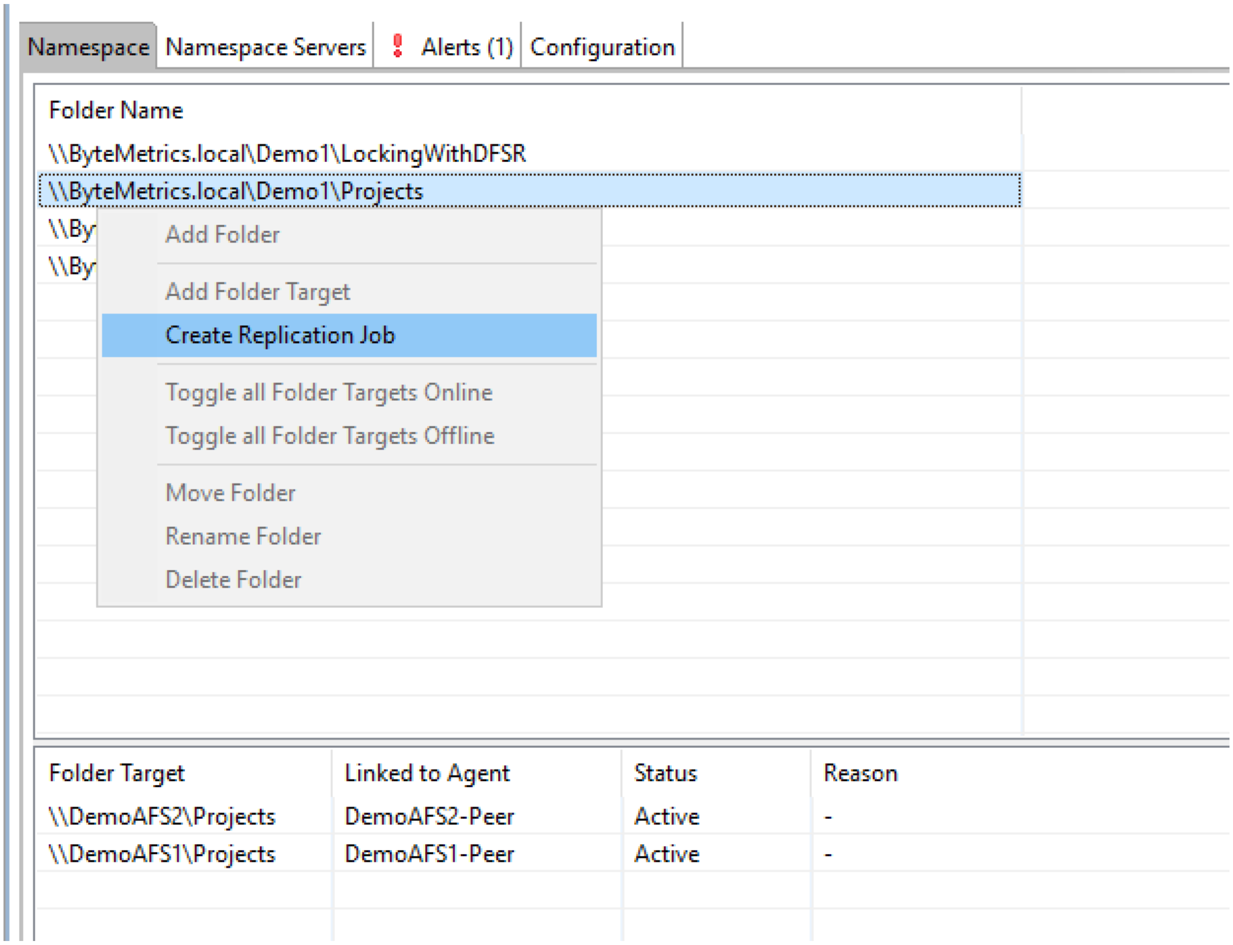

Folder Name
\\ByteMetrics.local\Demo1\LockingWithDFSR
\\ByteMetrics.local\Demo1\Projects
\\By
\\By

Folder Target	Linked to Agent	Status	Reason
\\DemoAFS2\Projects	DemoAFS2-Peer	Active	-
\\DemoAFS1\Projects	DemoAFS1-Peer	Active	-

The **Create New Job** wizard appears. The two supported job types are listed: File Collaboration and File Synchronization. All other job types are not supported for use with DFS namespace management.

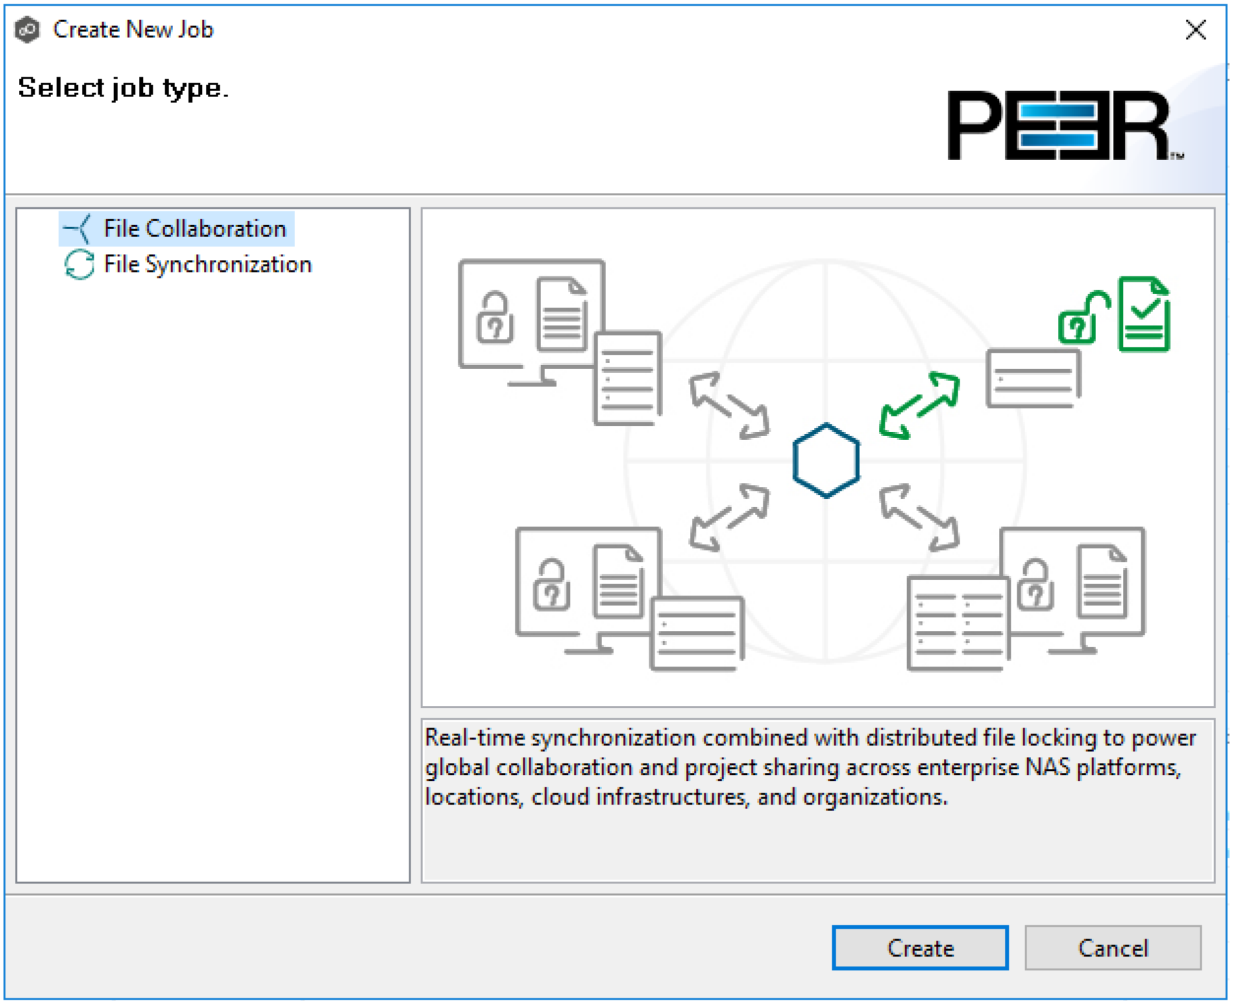

Real-time synchronization combined with distributed file locking to power global collaboration and project sharing across enterprise NAS platforms, locations, cloud infrastructures, and organizations.

3. Select a job type:

   - Select **File Collaboration** if locking is required in additional to synchronization (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).

   The **New Job** wizard appears. The **Participants** page is slightly different than the standard **New Job** wizard.

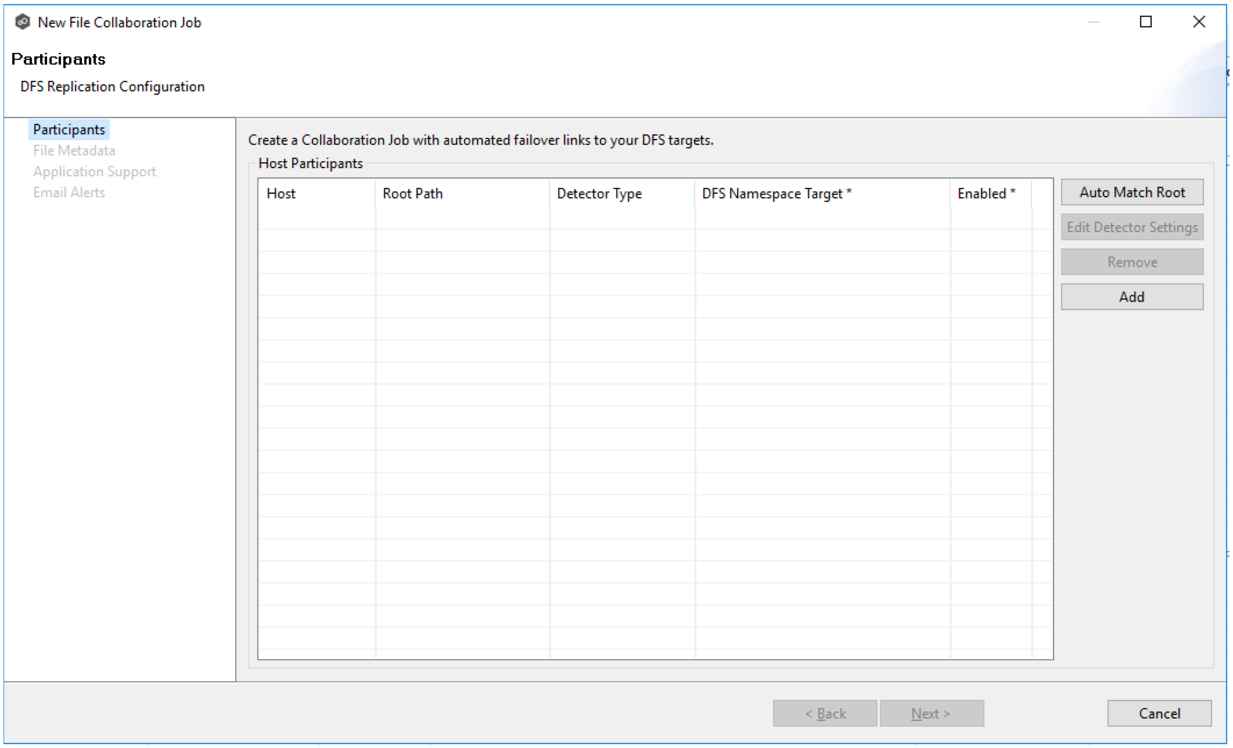

Create a Collaboration Job with automated failover links to your DFS targets.

Host	Root Path	Detector Type	DFS Namespace Target *	Enabled *

4. Click the **Add** button.

The **Host Selection** dialog appears.

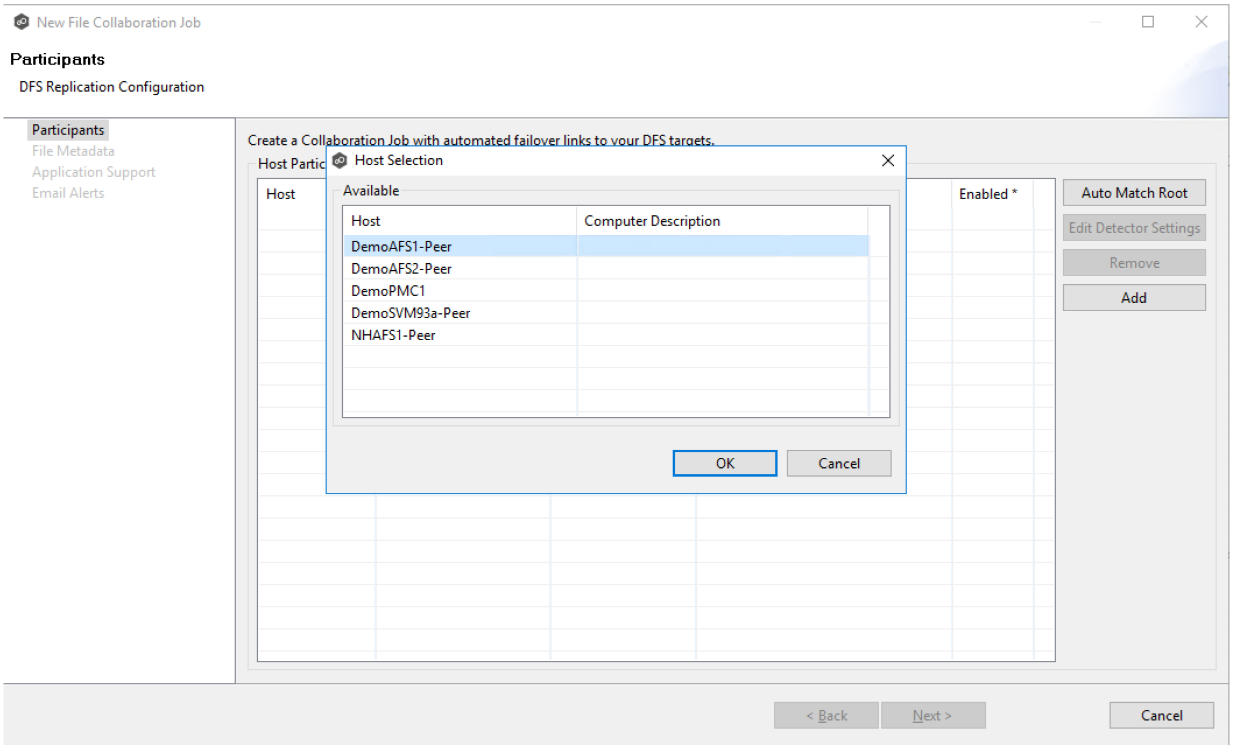

Create a Collaboration Job with automated failover links to your DFS targets.

Host	Computer Description
DemoAFS1-Peer	
DemoAFS2-Peer	
DemoPMC1	
DemoSVM93a-Peer	
NHAFS1-Peer	

5. Select two or more available Agents to add to this collaboration or synchronization job.

6. Click **OK**.

   The selected Agents are added to the **Host Participants** table.

7. Click **Auto Match Root** to automatically match a participant with the appropriate namespace folder target.

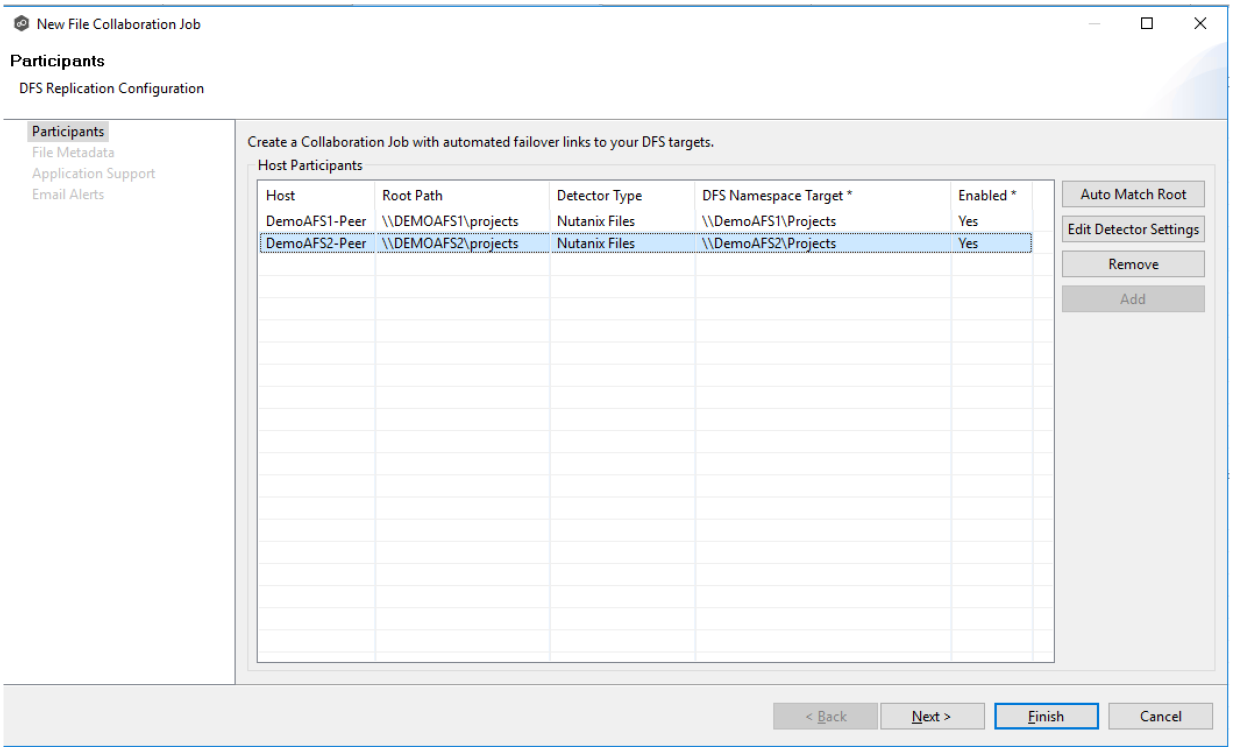

Create a Collaboration Job with automated failover links to your DFS targets.

Host Participants

Host	Root Path	Detector Type	DFS Namespace Target *	Enabled *
DemoAFS1-Peer	\\DEMOAFS1\projects	Nutanix Files	\\DemoAFS1\Projects	Yes
DemoAFS2-Peer	\\DEMOAFS2\projects	Nutanix Files	\\DemoAFS2\Projects	Yes

8. If an Agent has talked to that file server before, the **Detector Type** column will be auto-populated.

9. If this is the first collaboration or synchronization job created with these Agents, you may need to populate the **Root Path** column manually.

   - If using a Windows file server, the Root Path should be the local path on that file server that corresponds to the share path of the folder target. You must also select **Windows** in the **Detector Type** column.

   - If using a non-Windows NAS device, this path should match the Namespace Folder Target. You will also need to select the appropriate storage platform from the Detector Type column. After selecting the storage platform, you will need to press the Edit Detector Settings to enter the required settings for your selected platform.

10. Once all participants are added and associated with folder targets, click **Next**.

11. (Optional) In the **File Metadata** page, enable file metadata replication, and then click **Next**.

12. (Optional) In the **Application Support** page, select the applications you want optimized, and then click **Next**

13. (Optional) In the **Email Alerts** page, select the emails alerts to apply to the job.

14. Click **Finish** to complete the creation of this job.

15. From this point forward, if this collaboration or synchronization job is running along with its paired Namespace job, the PMC will automatically failover and failback 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 the Peer Management Center.

2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **New Job** wizard displays a list of job types you can create.

3. Click **File Collaboration**, and then click **Create**.

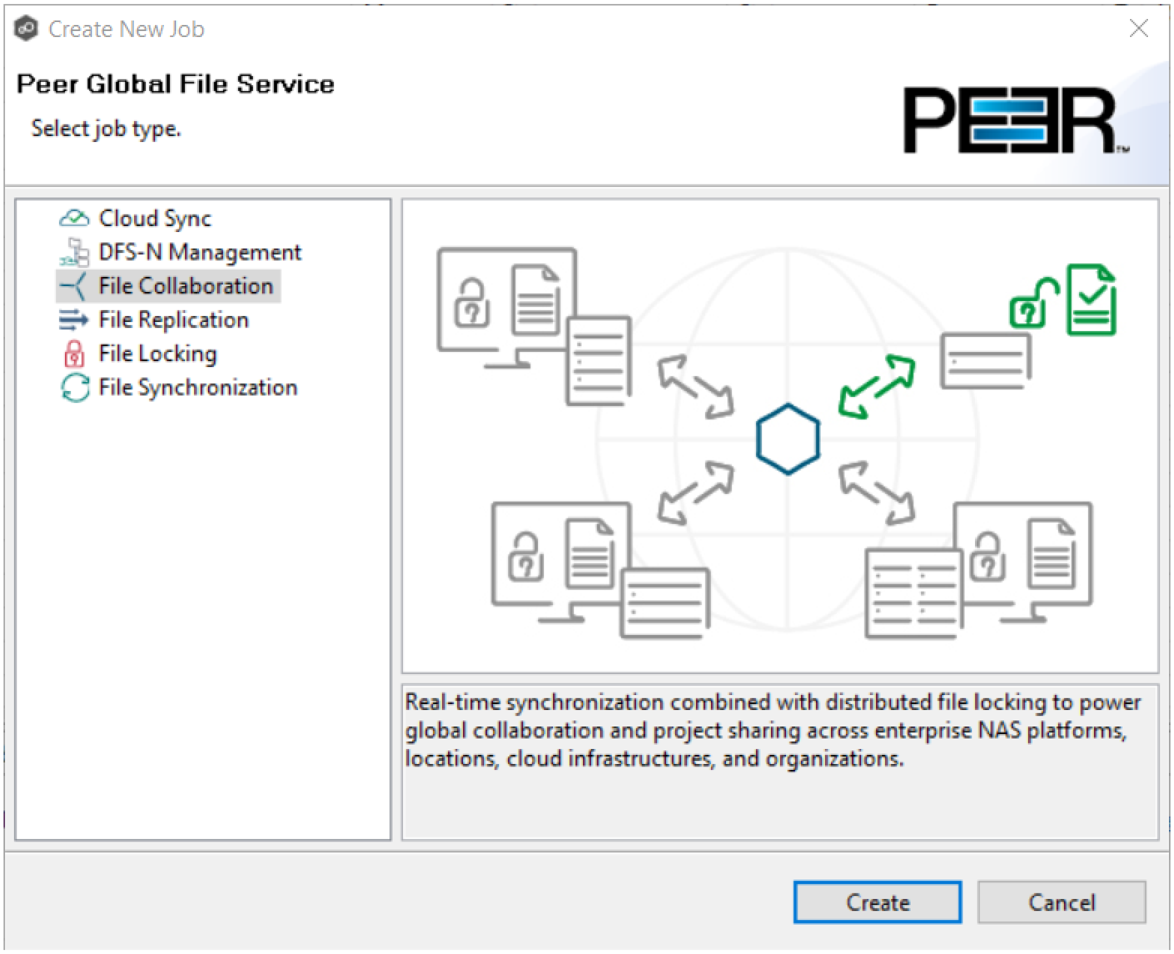

Real-time synchronization combined with distributed file locking to power global collaboration and project sharing across enterprise NAS platforms, locations, cloud infrastructures, and organizations.

4. Enter a name for the job.

   The job name must be unique.

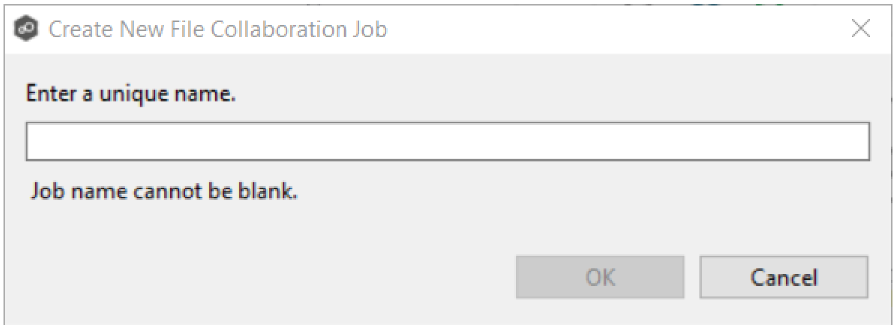

Job name cannot be blank.

5. Click **OK**.

   The Participants page appears.

Step 2: Participants

A File Collaboration job must have two or more participants. A participant consists of an Agent and the volume/share/folder to be replicated. The server that the Agent is installed upon is called the host (or host participant). A File Collaboration job replicates the files of participants in real-time.

1. Complete the five substeps:

   Participants

   Storage Platform

   Management Agent

   Storage Information

   Path

   After you add a participant, it appears in the **Participants** table.

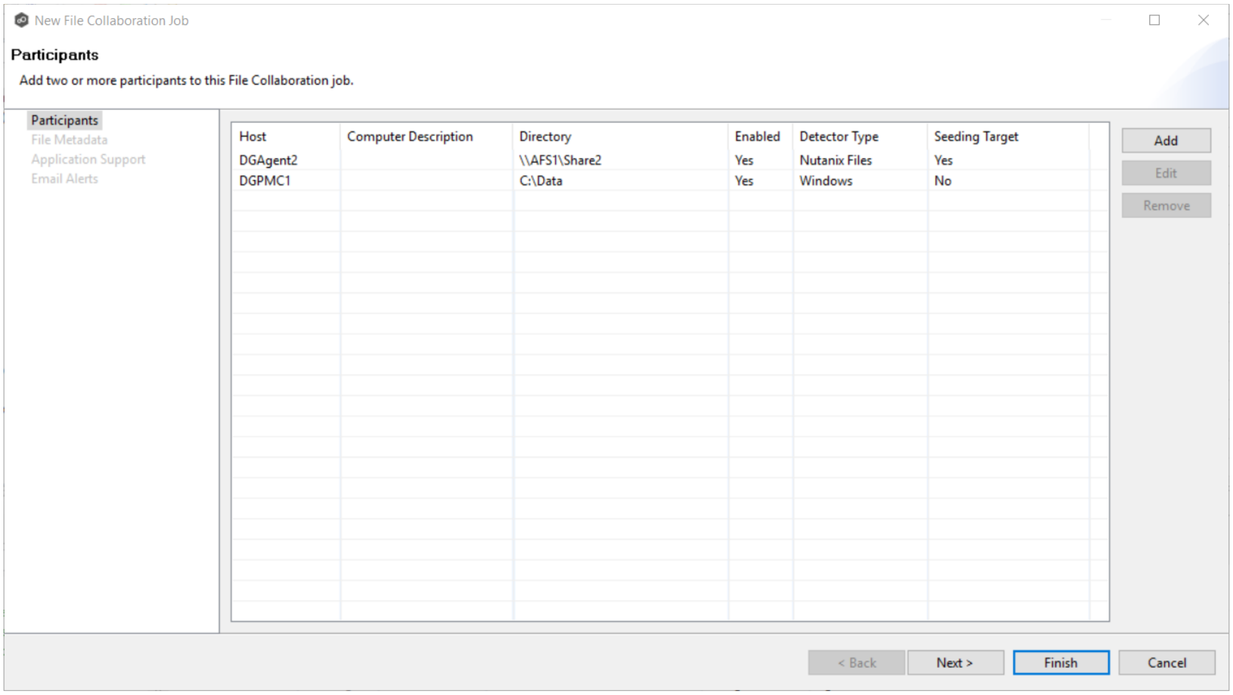

New File Collaboration Job

Participants

Add two or more participants to this File Collaboration job.

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent2		\\AFS1\Share2	Yes	Nutanix Files	Yes
DGPMC1		C:\Data	Yes	Windows	No

2. Repeat the five substeps for each participant you want to add to the job.

3. Once you have added all the participants, click **Next** to specify file metadata for the job. (Don't click **Finish**.)

The **Participants** page is where you select and configure which hosts will be participating in this job. The **Participants** page is empty until you finish the process of adding your first participant. Once you have added the participants, they are listed on the **Participants** page.

To begin the process of adding a participant:

1. Click the **Add** button.

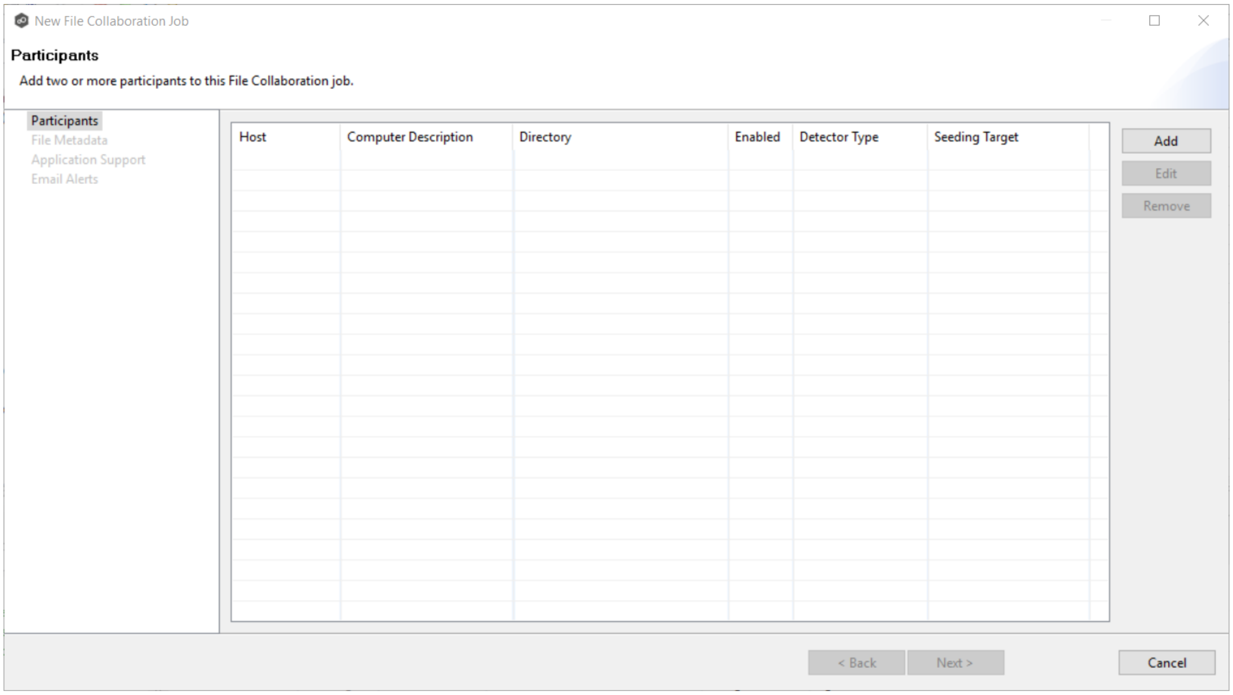

New File Collaboration Job

Participants

Add two or more participants to this File Collaboration job.

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target

Another wizard opens to guide you through the process of adding a participant to the job. The first step in the process involves selecting the storage platform.

The **Storage Platform** page lists the types of storage platforms that File Collaboration supports. A storage device hosts data you want to replicate. It is often referred to as the host or host participant.

1. Select the type of storage platform that hosts the data you want to replicate/for the job participant.

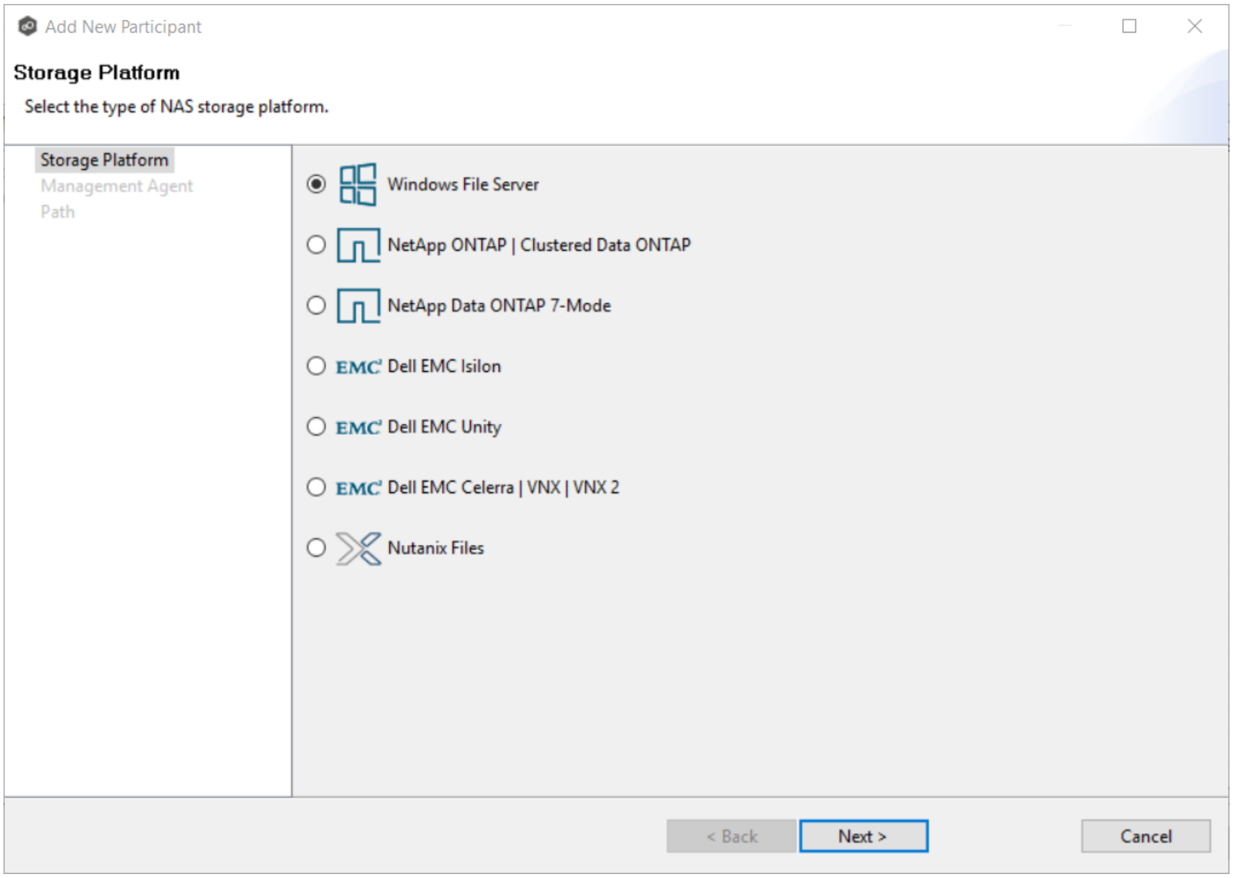

Add New Participant

Storage Platform

Select the type of NAS storage platform.

● Windows File Server

○ NetApp ONTAP | Clustered Data ONTAP

○ NetApp Data ONTAP 7-Mode

○ EMC Dell EMC Isilon

○ EMC Dell EMC Unity

○ EMC Dell EMC Celerra | VNX | VNX 2

○ Nutanix Files

2. Click **Next**.

   The Management Agent page is displayed.

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 the Management Agent that manages the volumes/shares/folders you want to replicate in this job.

1. Select the Agent that manages the host.

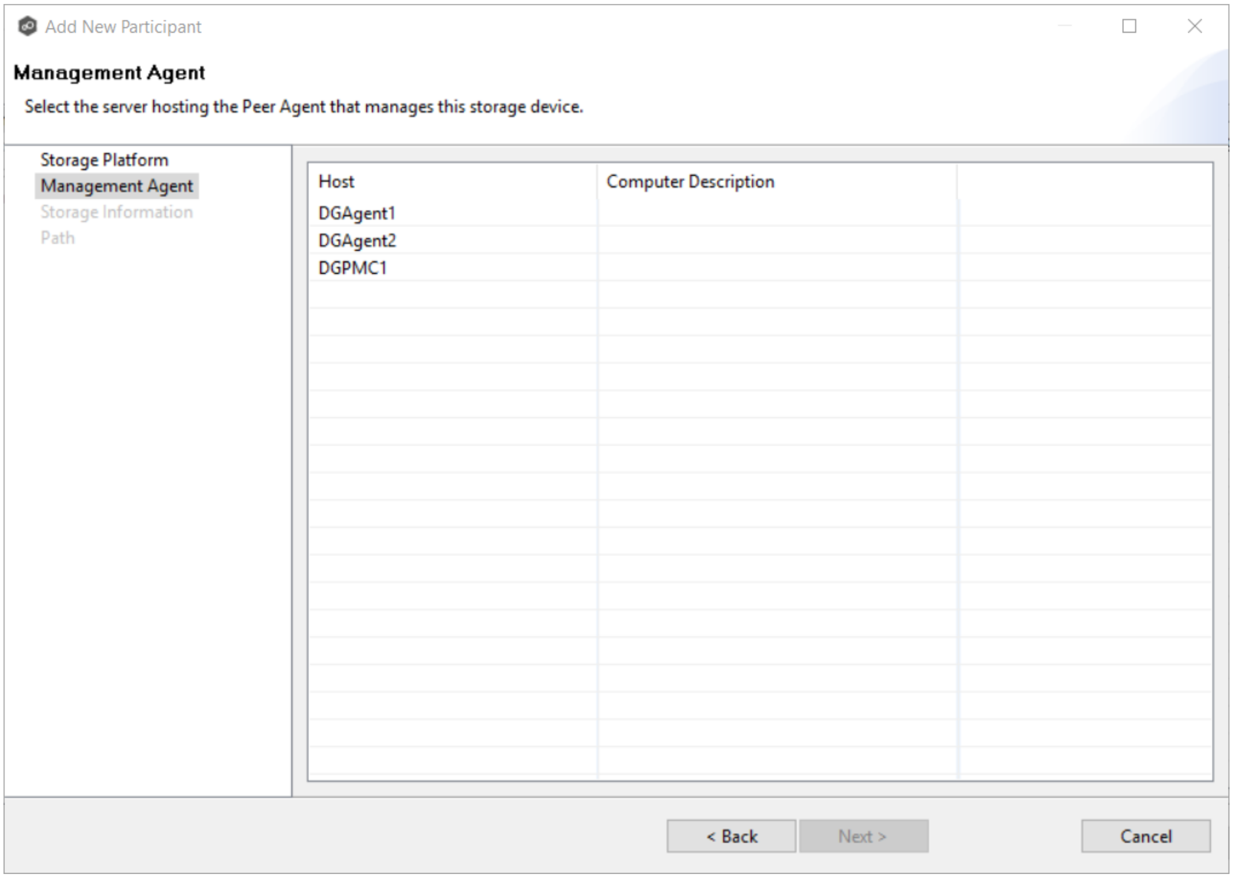

Management Agent

Select the server hosting the Peer Agent that manages this storage device.

Host	Computer Description
DGAgent1	
DGAgent2	
DGPMC1	

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 Storage Information page is displayed.

If you selected any storage platform type other than Windows File Server in the previous wizard page, the **Storage Information** page appears. It requests the credentials necessary to connect to the storage device you want to replicate. If you selected Windows Files Server in the previous wizard page, skip to Step 3: File Metadata.

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.

NetApp ONTAP | Clustered Data ONTAP

1. Enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

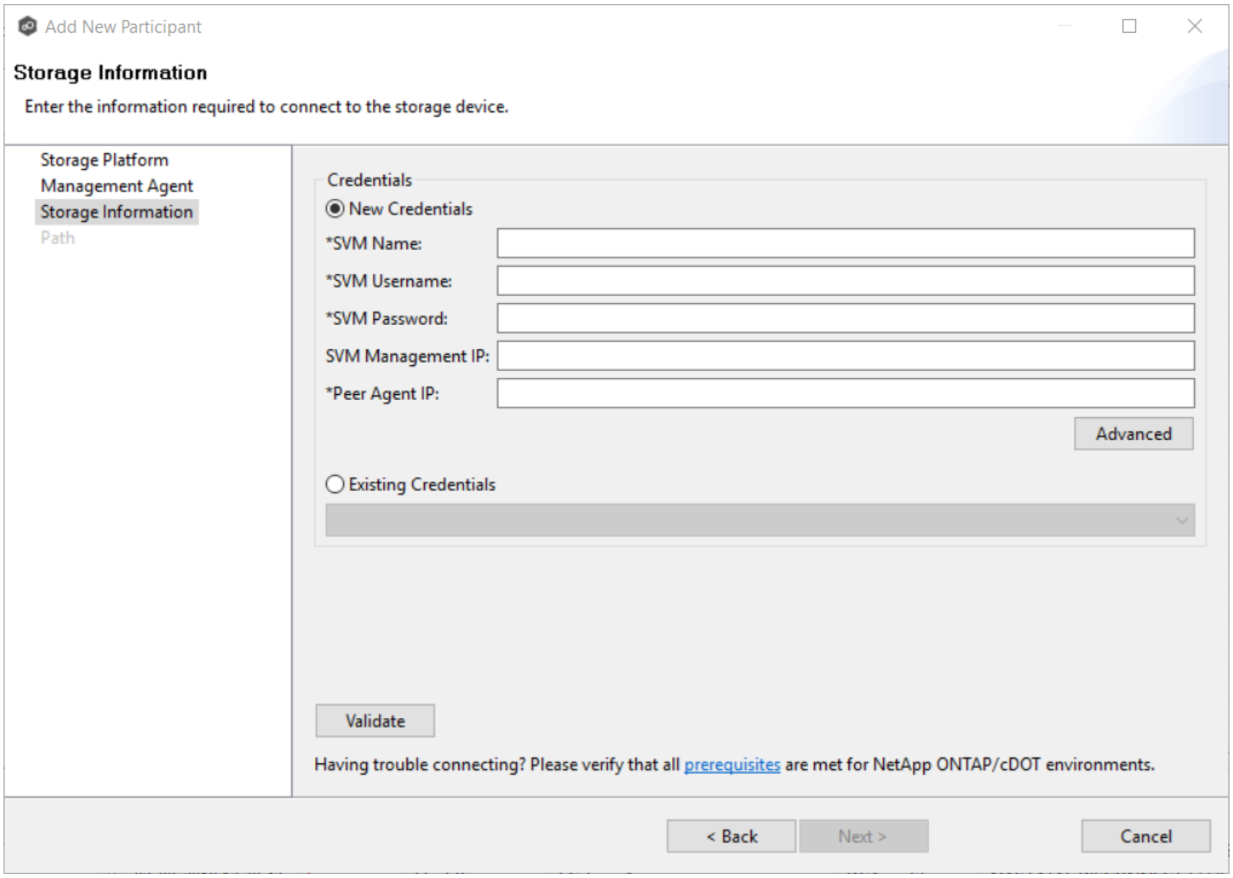

Add New Participant

Storage Information

Enter the information required to connect to the storage device.

Credentials

● New Credentials

*SVM Name:

*SVM Username:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

○ Existing Credentials

Having trouble connecting? Please verify that all prerequisites are met for NetApp ONTAP/cDOT environments.

2. Click **Next**.

   The Path page is displayed.

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	Enter the IP address of the server hosting the Agent that manages the Storage Virtual Machine.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

NetApp Data ONTAP 7-Mode

1. Enter the credentials to connect to the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

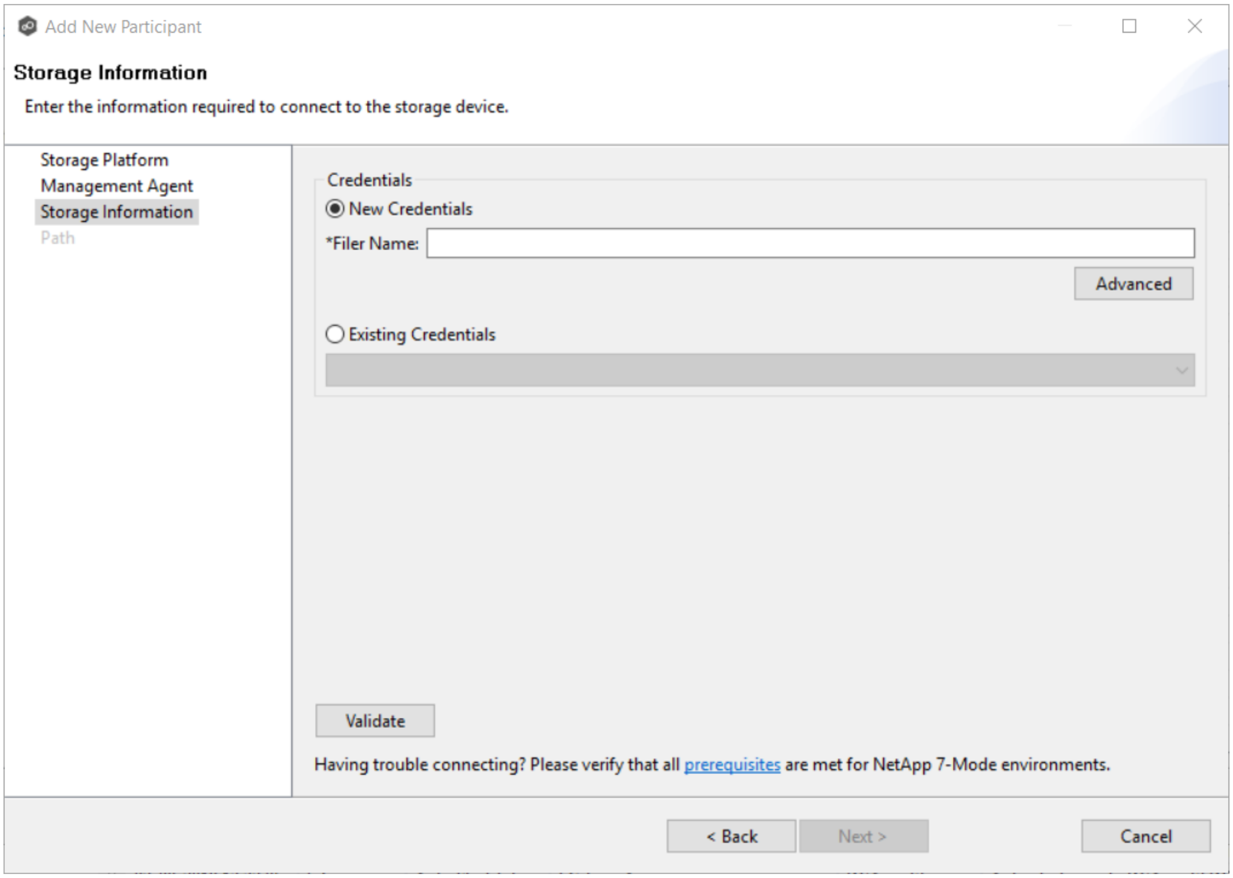

2. Click **Next**.

   The Path page is displayed.

Filer Name	Enter the name of the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

Dell EMC Isilon

1. Enter the credentials to connect the EMC Isilon cluster hosting the data to be replicated.

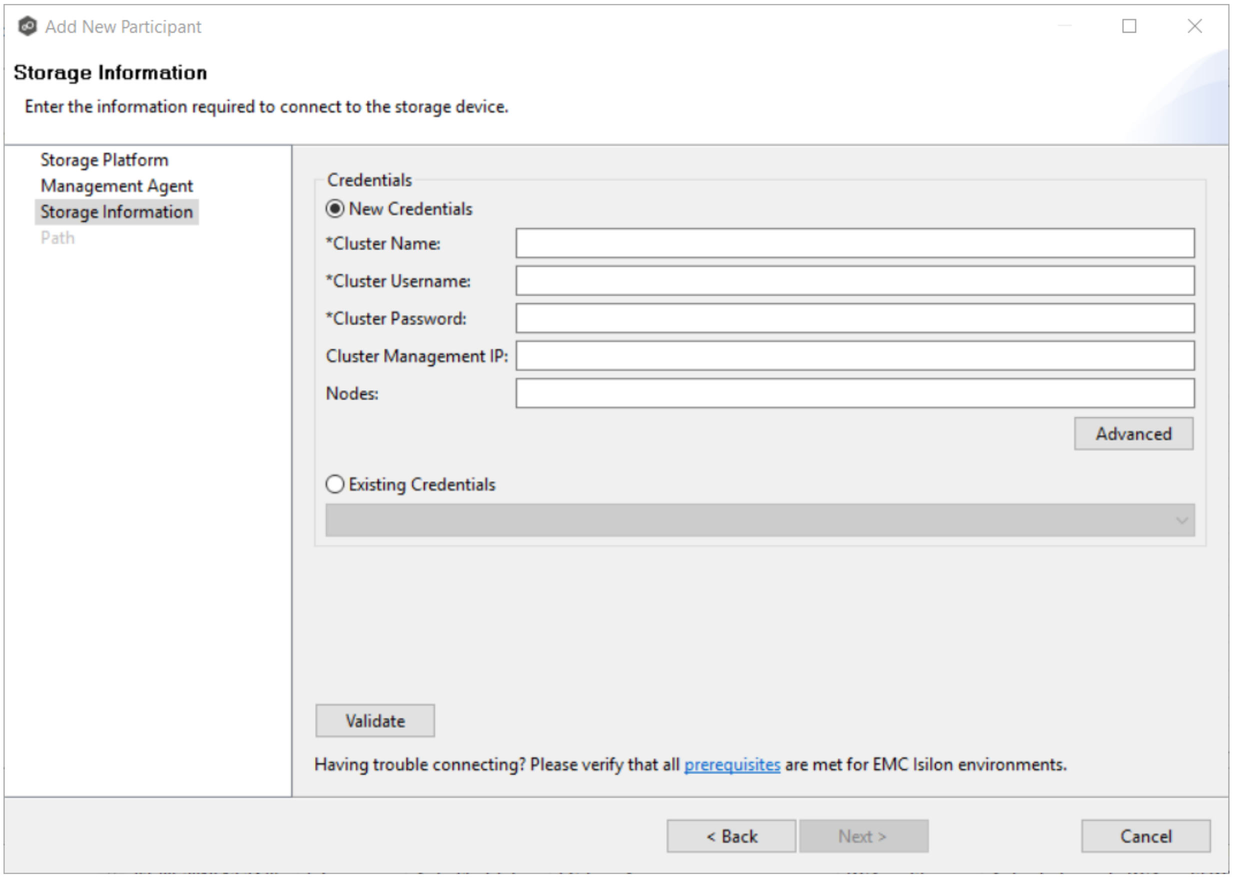

Add New Participant

Storage Information

Enter the information required to connect to the storage device.

Storage Platform
Management Agent
Storage Information
Path

Credentials

● New Credentials

*Cluster Name:

*Cluster Username:

*Cluster Password:

Cluster Management IP:

Nodes:

Advanced

○ Existing Credentials

Validate

Having trouble connecting? Please verify that all prerequisites are met for EMC Isilon environments.

< Back | Next > | Cancel

2. Click **Next**.

   The Path page is displayed.

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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

Dell EMC Unity

1. Enter the credentials to connect to the NAS Server hosting the data to be replicated.

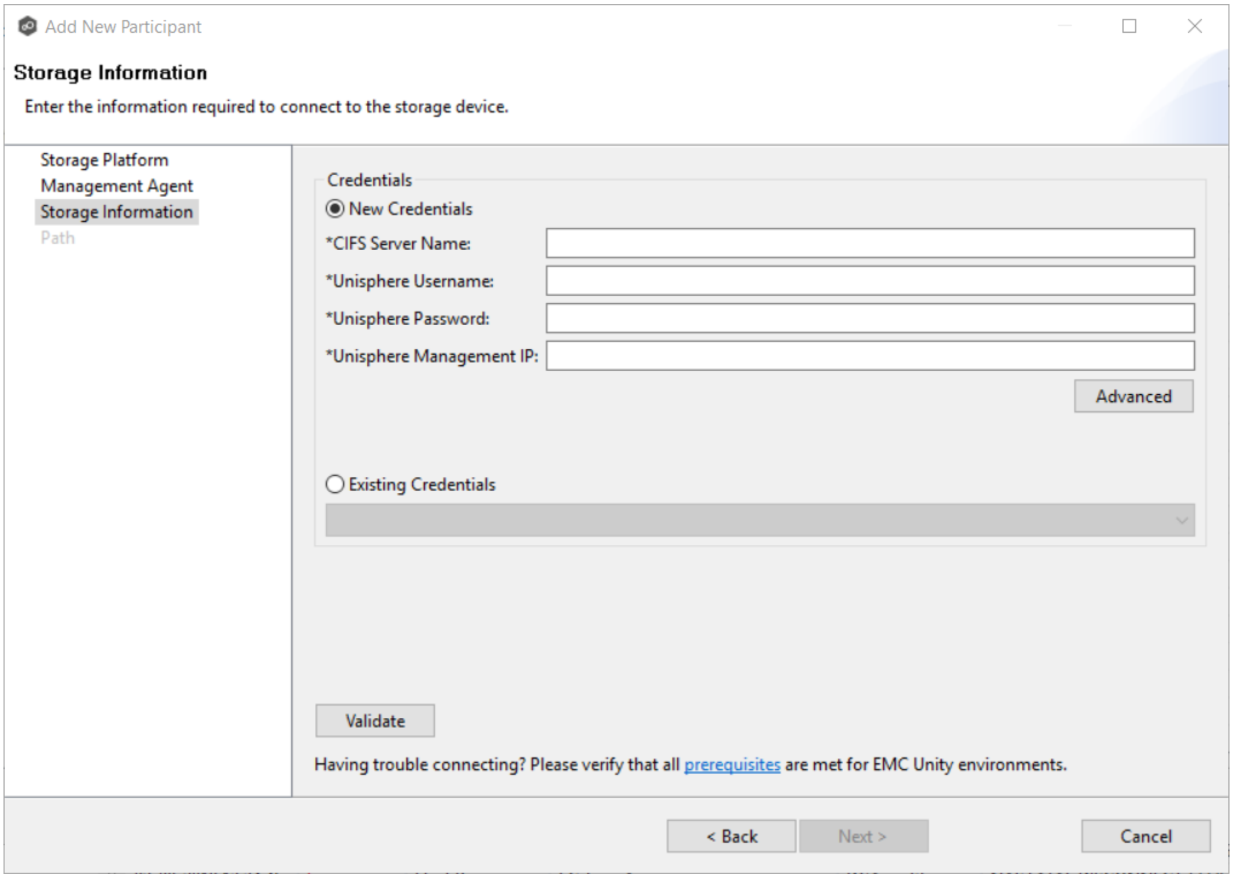

Add New Participant

Storage Information

Enter the information required to connect to the storage device.

Credentials

● New Credentials

*CIFS Server Name:

*Unisphere Username:

*Unisphere Password:

*Unisphere Management IP:

○ Existing Credentials

Having trouble connecting? Please verify that all prerequisites are met for EMC Unity environments.

2. Click **Next**.

   The Path page is displayed.

NAS 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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

Dell EMC Celerra | VNX | VNX 2

1. Enter the credentials to connect to the CIFS Server hosting the data to be replicated.

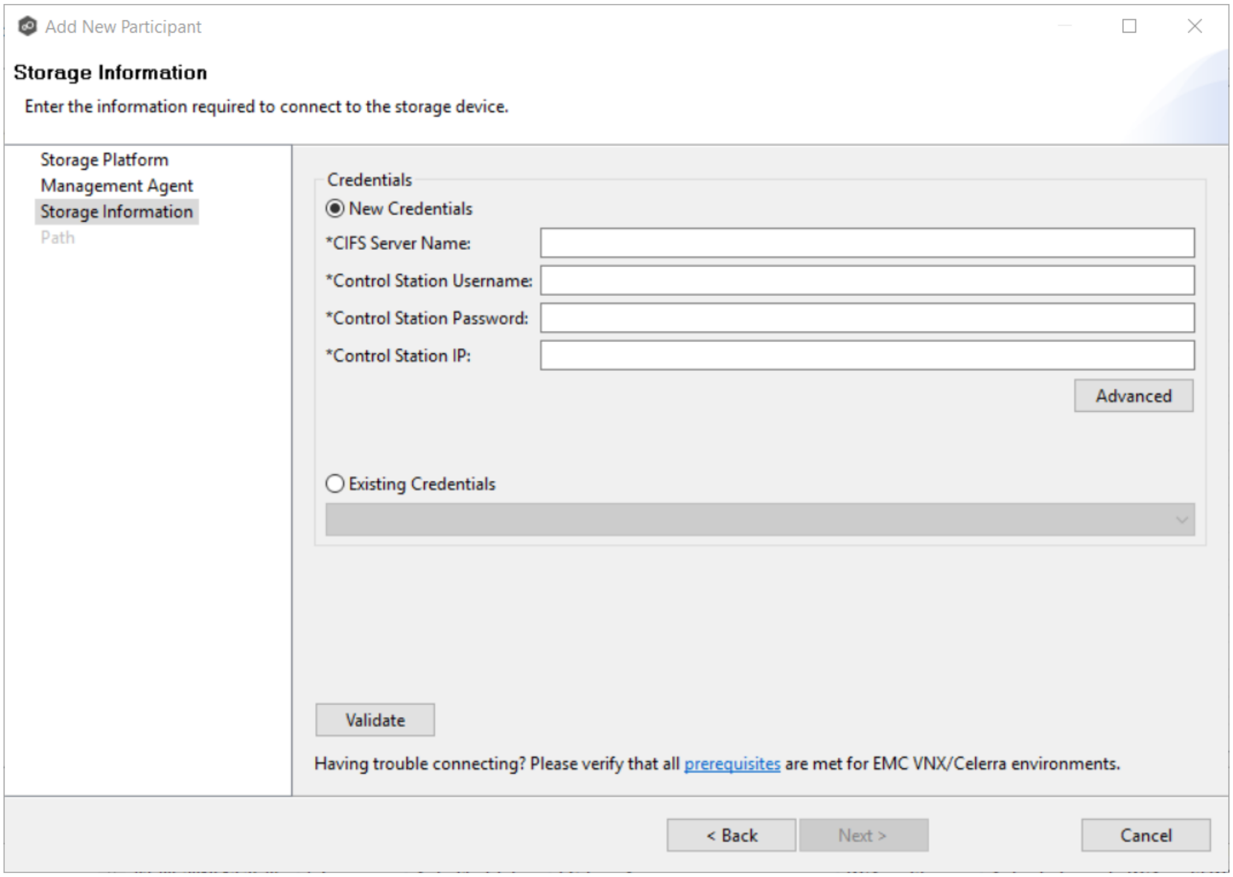

Add New Participant
Storage Information
Enter the information required to connect to the storage device.
Storage Platform
Management Agent
Storage Information
Path
Credentials
● New Credentials
*CIFS Server Name:
*Control Station Username:
*Control Station Password:
*Control Station IP:
Advanced
○ Existing Credentials
Validate
Having trouble connecting? Please verify that all prerequisites are met for EMC VNX/Celerra environments.
< Back
Next >
Cancel

2. Click **Next**.

   The Path page is displayed.

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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

Nutanix Files

1. Enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated.

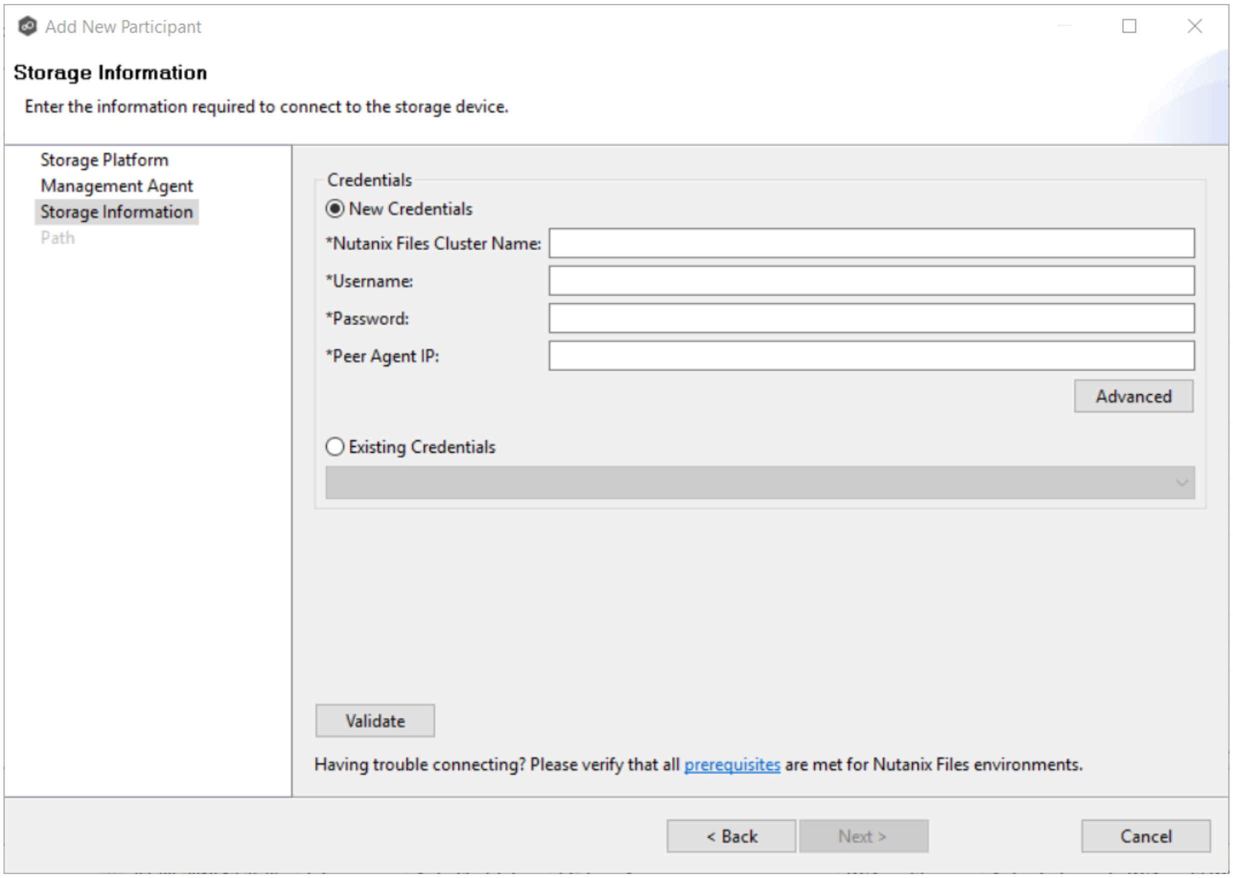

Add New Participant

Storage Information

Enter the information required to connect to the storage device.

Credentials

● New Credentials

*Nutanix Files Cluster Name:

*Username:

*Password:

*Peer Agent IP:

○ Existing Credentials

Having trouble connecting? Please verify that all prerequisites are met for Nutanix Files environments.

2. Click Next.

   The Path page is displayed.

Nutanix Files Cluster 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 AFS cluster via its management APIs.
Password	Enter the password for the account managing the AFS cluster via its management APIs.
Peer Agent IP	Enter the IP address of the Agent server used to manage the storage platform. This should not point to the AFS cluster itself.

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.

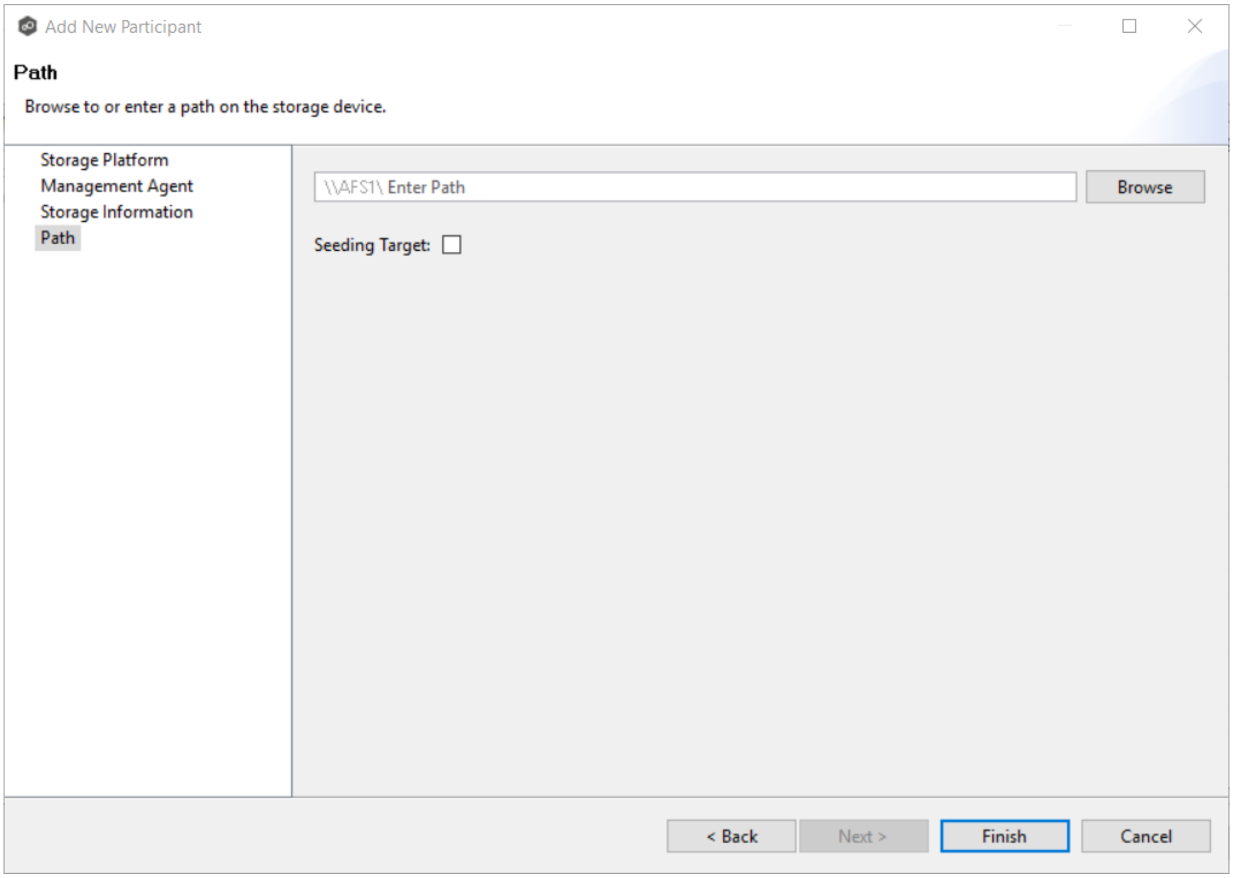

Add New Participant

Path

Browse to or enter a path on the storage device.

\\AFS1\ Enter Path

Seeding Target: □

If you selected **Browse**, the **Folder Browser** dialog appears:

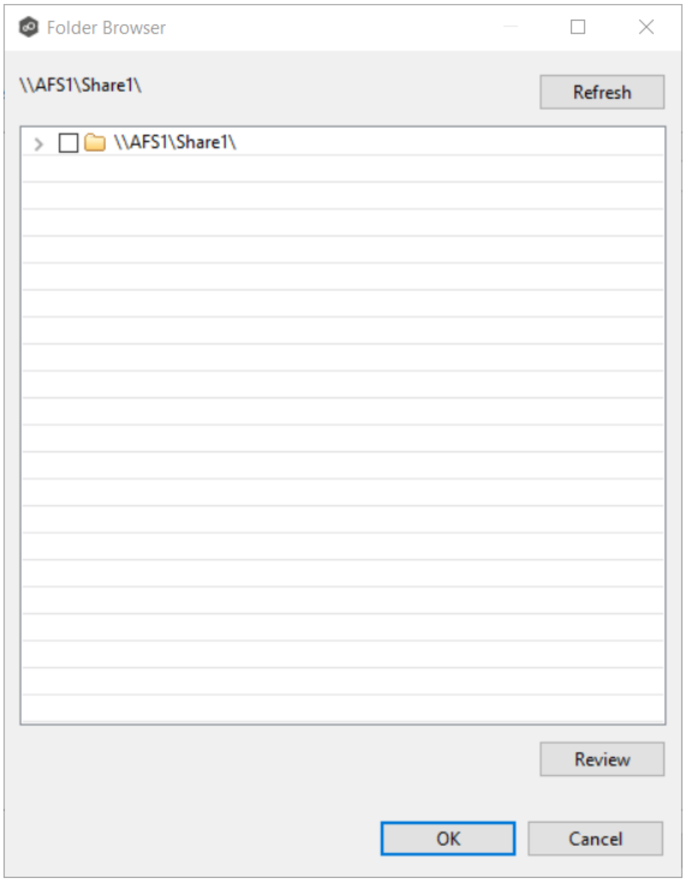

a. Expand the folder tree.

b. Select the appropriate volume/share/folder/

c. (Optional) Click the **Review** button to see your selection.

d. Click **OK**.

2. (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. For more information about smart data seeding, see Smart Data Seeding or contact support@peersoftware.com.

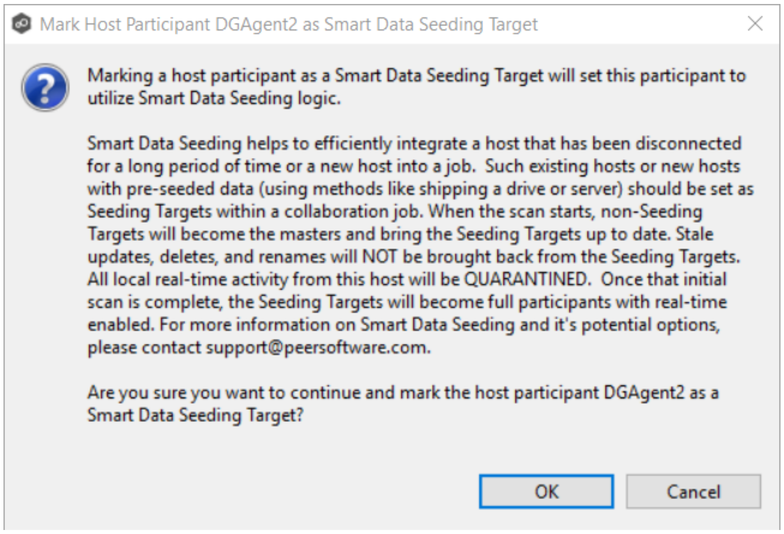

Mark Host Participant DGAgent2 as Smart Data Seeding Target

Marking a host participant as a Smart Data Seeding Target will set this participant to utilize Smart Data Seeding logic.

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 job. Such existing hosts or new hosts with pre-seeded data (using 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 from this host 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 it's potential options, please contact support@peersoftware.com.

Are you sure you want to continue and mark the host participant DGAgent2 as a Smart Data Seeding Target?

Note: Multiple participants in a File Collaboration job can use Smart Data Seeding logic; however, at least one participant should not use Smart Data Seeding logic.

3. Click **Finish** to complete the wizard for this participant.

4. Return to Step 2: Participants to add more participants, if applicable. A File Collaboration job must have at least two participants. If you have added all of the participants, continue with Step 3: File Metadata.

Step 3: 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 of the 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 (DACL, SACL, Owner.) 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 synchronized during the initial scan. If enabled, changes to the selected security descriptor components (DACL, SACL, Owner) will be synchronized during the initial scan.

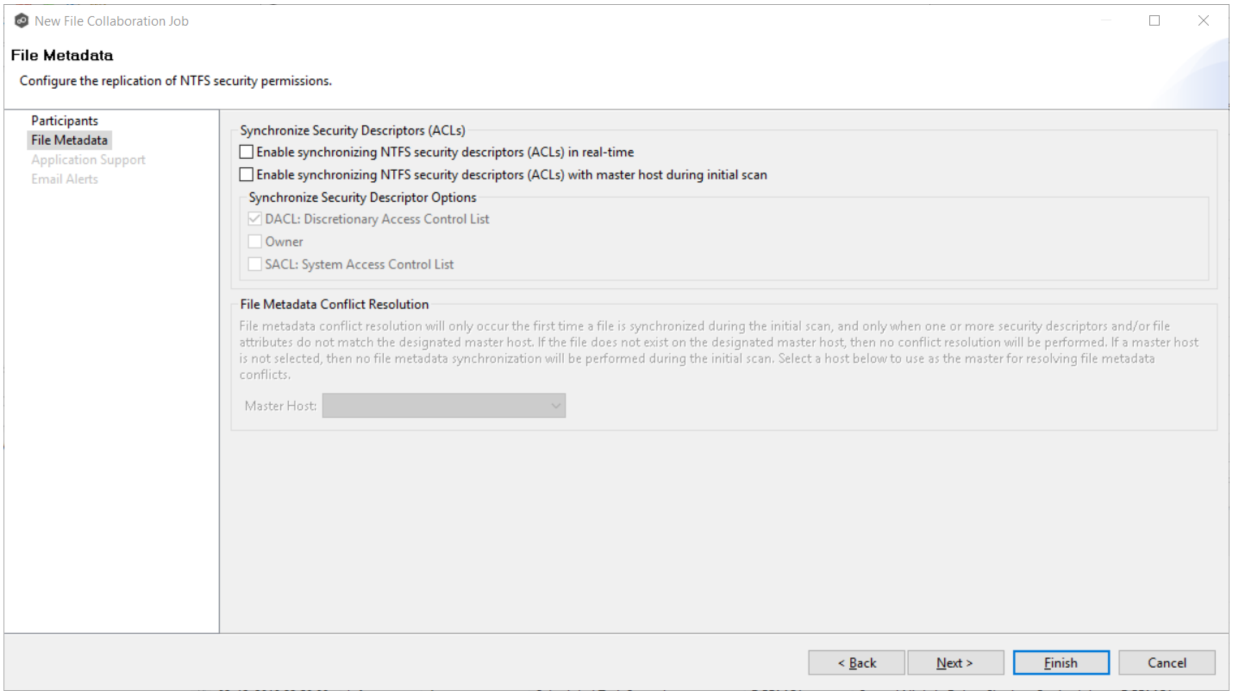

New File Collaboration Job

File Metadata

Configure the replication of NTFS security permissions.

Synchronize Security Descriptors (ACLs)

☐ Enable synchronizing NTFS security descriptors (ACLs) in real-time

☐ Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan

Synchronize Security Descriptor Options

☑ DACL: Discretionary Access Control List

☐ Owner

☐ SACL: System Access Control List

File Metadata Conflict Resolution

File metadata conflict resolution will only occur the first time a file is synchronized during the initial scan, and only when one or more security descriptors and/or file attributes do not match the designated master host. If the file does not exist on the designated master host, then no conflict resolution will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan. Select a host below to use as the master for resolving file metadata conflicts.

Master Host:

2. Click **OK** in the informational dialog that appears after selecting a metadata option.

3. Select which security descriptor components (DACL, SACL and Owner) are synchronized.

   In general, you will usually need to synchronize only DACLs. If you need to synchronize SACLs or Owner, then the user that a Peer Agent service is run under on each participating host must have permission to read and write SACLs and Owner.

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.

   Conflict resolution for file metadata 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 will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan.

5. Click **Next**.

The Application Support page is displayed.

Step 4: 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.

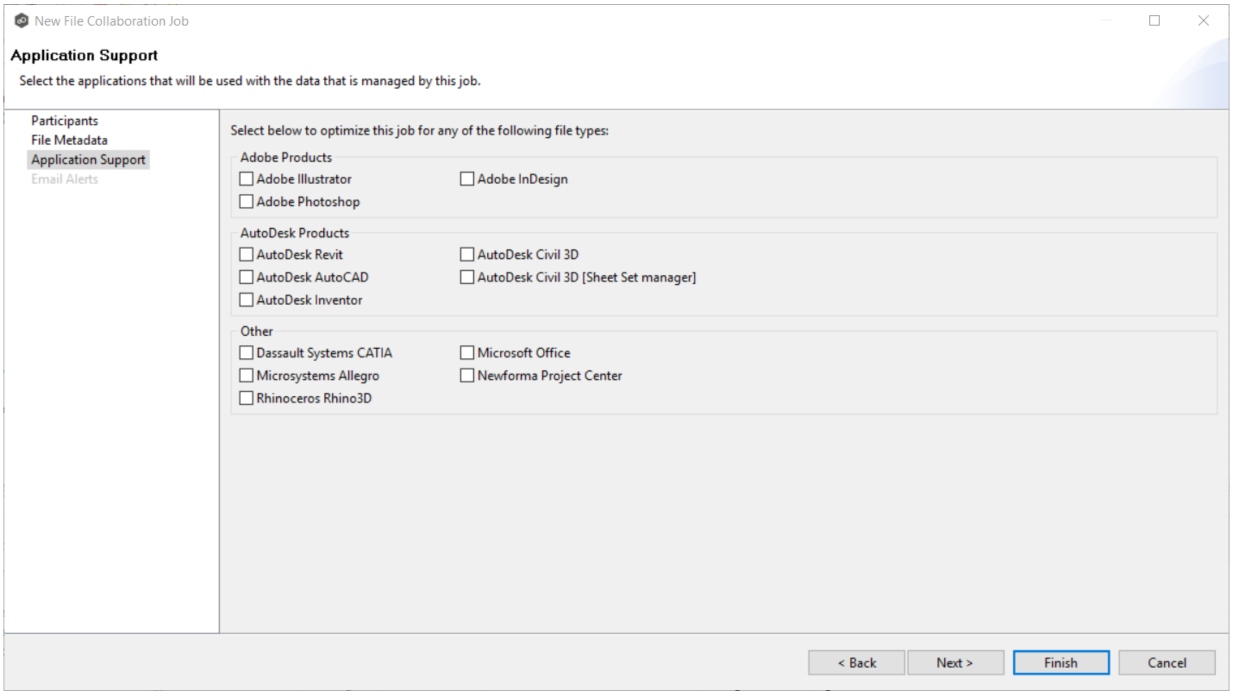

Application Support

Select the applications that will be used with the data that is managed by this job.

Select below to optimize this job for any of the following file types:

Adobe Products

- □ Adobe Illustrator
- □ Adobe InDesign
- □ Adobe Photoshop

AutoDesk Products

- □ AutoDesk Revit
- □ AutoDesk Civil 3D
- □ AutoDesk AutoCAD
- □ AutoDesk Civil 3D [Sheet Set manager]
- □ AutoDesk Inventor

Other

- □ Dassault Systems CATIA
- □ Microsoft Office
- □ Microsystems Allegro
- □ Newforma Project Center
- □ Rhinoceros Rhino3D

2. 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, 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.

Apply an Existing Email Alert

To apply an existing email alert to the job.

1. Click the **Select** button.

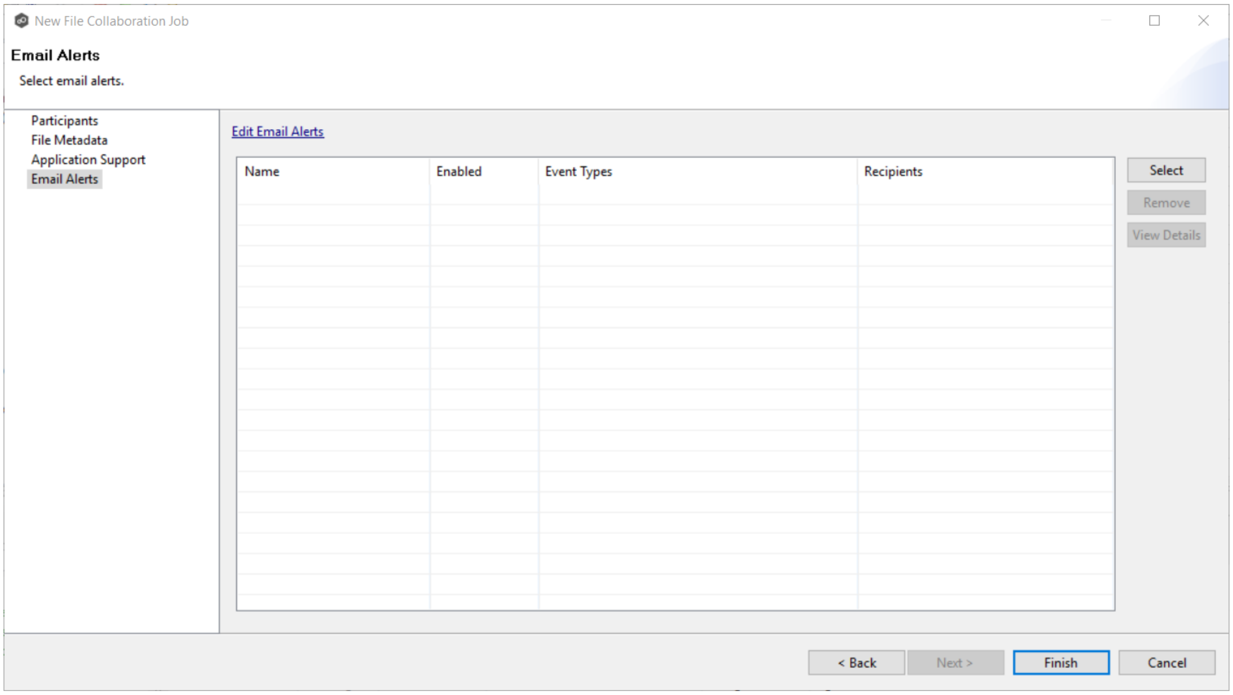

The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.

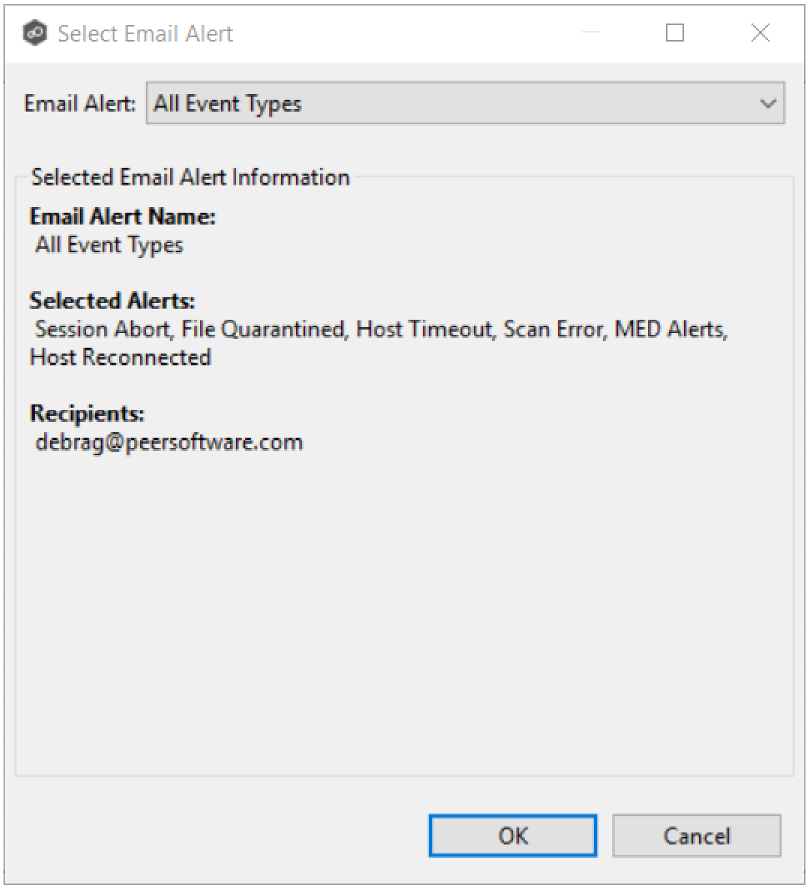

Email Alert: All Event Types

Selected Email Alert Information

Email Alert Name:
All Event Types

Selected Alerts:
Session Abort, File Quarantined, Host Timeout, Scan Error, MED Alerts, Host Reconnected

Recipients:
debrag@peersoftware.com

3. Click **OK**.

   The alert is listed on the **Email Alerts** page.

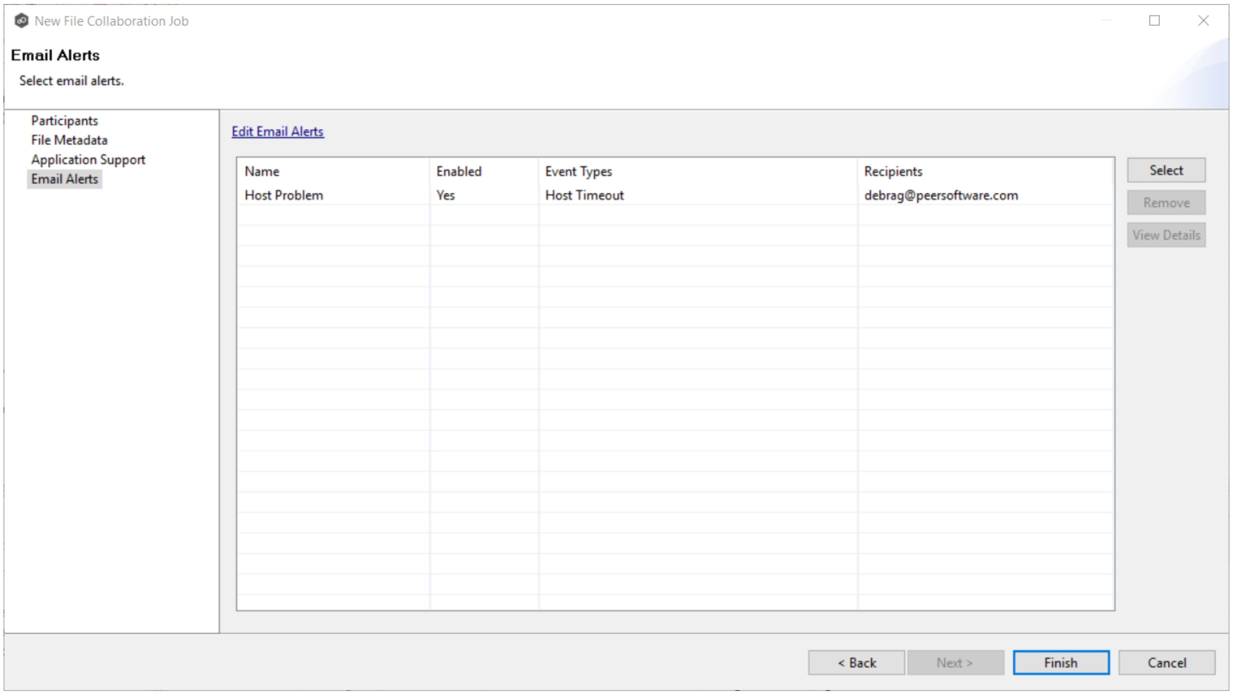

Name	Enabled	Event Types	Recipients
Host Problem	Yes	Host Timeout	debrag@peersoftware.com

4. (Optional) Repeat steps 1-3 to apply additional alerts.

5. After applying email alerts, continue to Step 6: Save Job.

Create an Email Alert

To create an alert or edit an existing alert from the **Email Alerts** page of the **Create Job** wizard:

1. Click the **Edit Email Alerts** link.

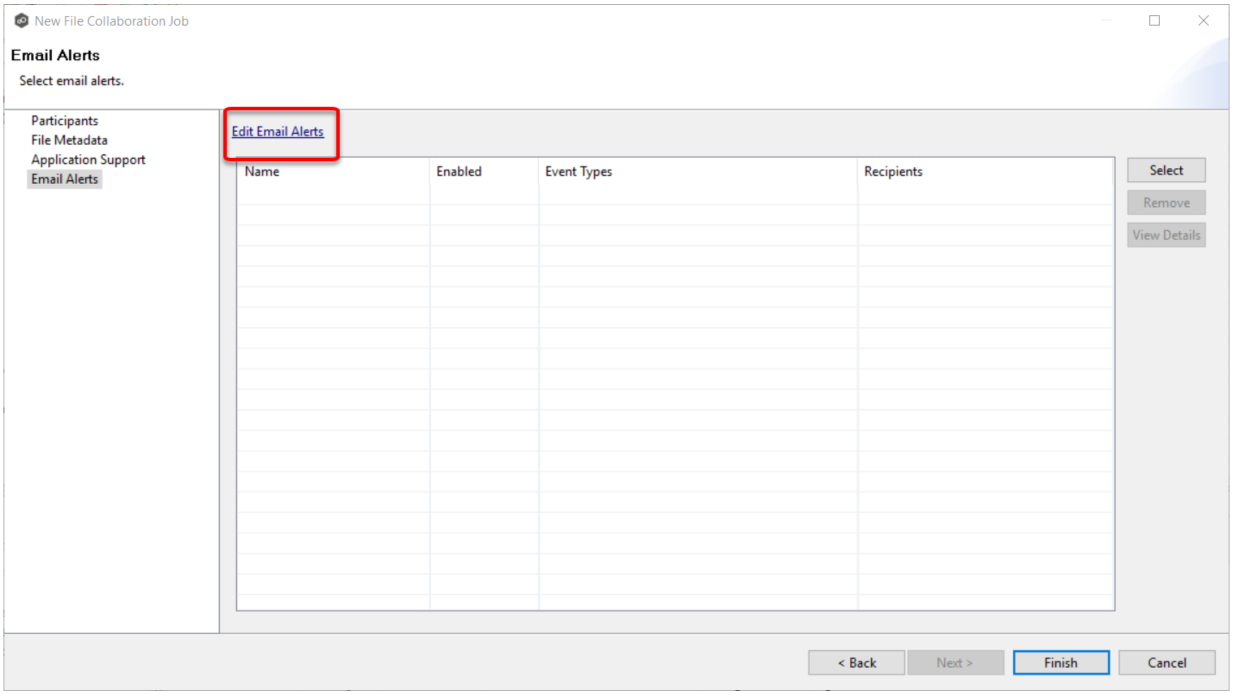

Name	Enabled	Event Types	Recipients

The **Collaboration Email Alerts** dialog appears.

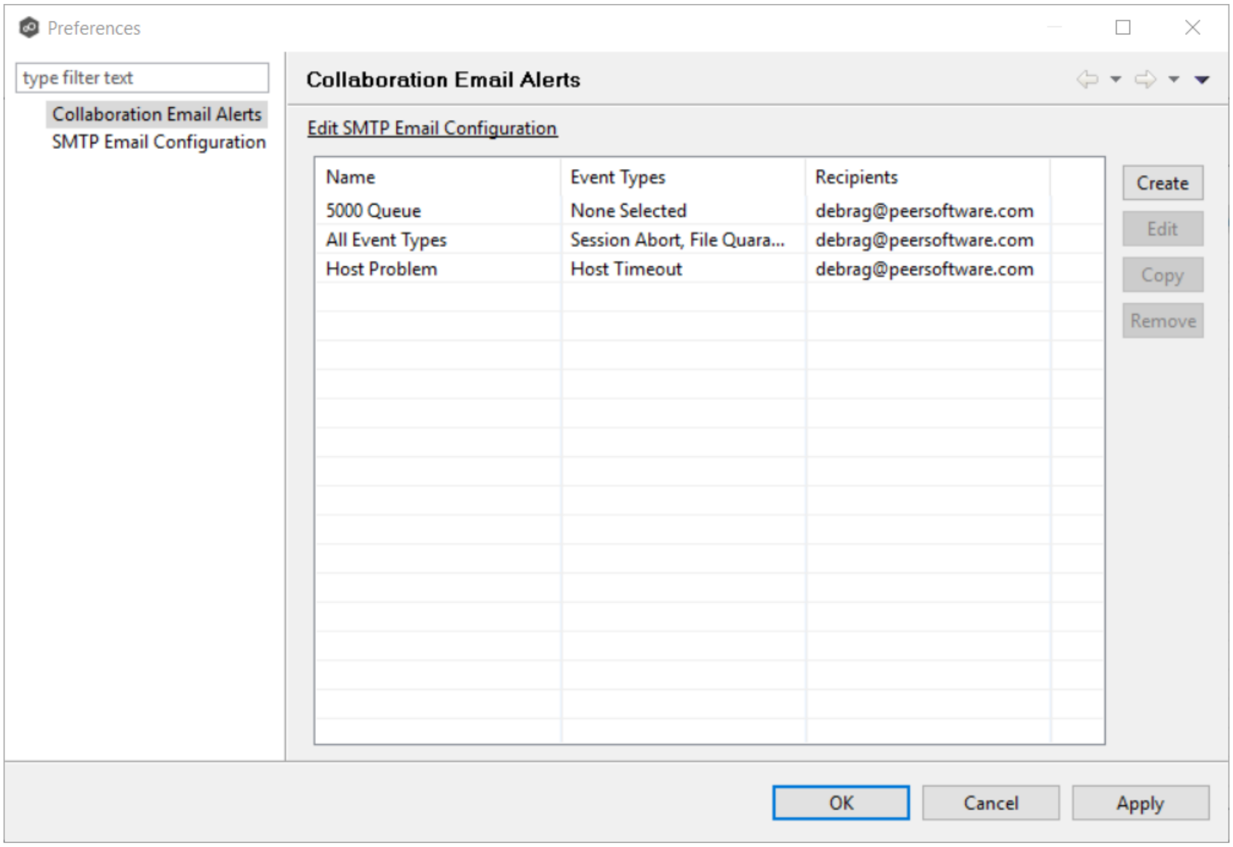

Name	Event Types	Recipients
5000 Queue	None Selected	debrag@peersoftware.com
All Event Types	Session Abort, File Quara...	debrag@peersoftware.com
Host Problem	Host Timeout	debrag@peersoftware.com

2. Click **Create**.

   The **Create Email Alert** dialog appears.

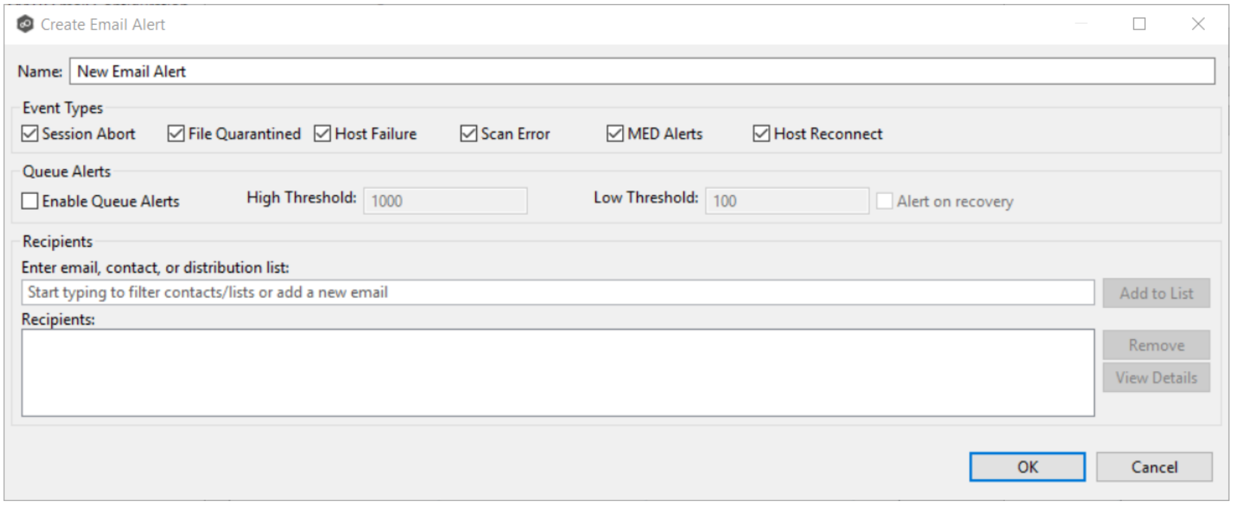

Name: New Email Alert

Event Types

☑ Session Abort ☑ File Quarantined ☑ Host Failure ☑ Scan Error ☑ MED Alerts ☑ Host Reconnect

Queue Alerts

□ Enable Queue Alerts High Threshold: 1000 Low Threshold: 100 □ Alert on recovery

Recipients

Enter email, contact, or distribution list: Start typing to filter contacts/lists or add a new email

Recipients:

3. Continue with Step 4 of the instructions in Email Alerts in the **Preferences** section.

Step 6: Save Job

Now that you have completed the first five 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 Collaboration job. It is now listed in the **Jobs** view under **File Collaboration** and a job run-time view appears in the **Summaries** area. You can start the job from either place. See Running and Managing a File Collaboration job for more information.

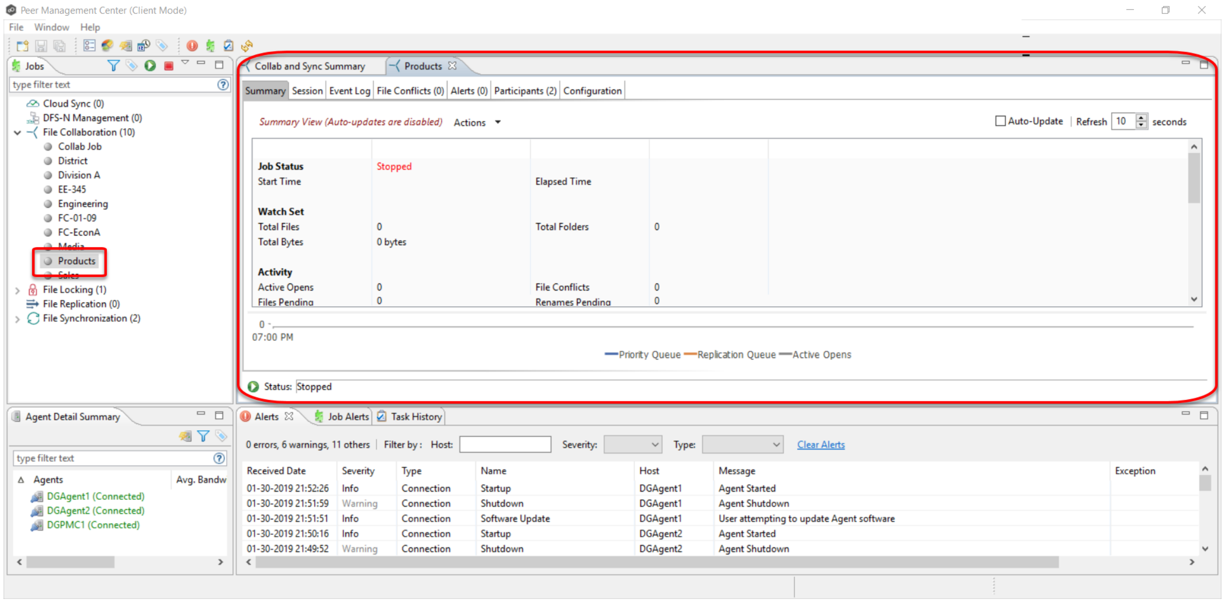

Summary View (Auto-updates are disabled)

□ Auto-Update | Refresh 10 seconds

Job Status Stopped
Start Time | Elapsed Time

Watch Set
Total Files 0 | Total Folders 0
Total Bytes 0 bytes

Activity
Active Opens 0 | File Conflicts 0
Files Pending 0 | Renames Pending 0

Status: Stopped

0 errors, 6 warnings, 11 others | Filter by : Host: Severity: Type:

Received Date	Severity	Type	Name	Host	Message	Exception
01-30-2019 21:52:26	Info	Connection	Startup	DGAgent1	Agent Started	
01-30-2019 21:51:59	Warning	Connection	Shutdown	DGAgent1	Agent Shutdown	
01-30-2019 21:51:51	Info	Connection	Software Update	DGAgent1	User attempting to update Agent software	
01-30-2019 21:50:16	Info	Connection	Startup	DGAgent2	Agent Started	
01-30-2019 21:49:52	Warning	Connection	Shutdown	DGAgent2	Agent Shutdown	

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
- DFS-N
- File Filters
- File Locking
- General
- Logging and Alerts
- SNMP Notifications

- [Target Protection]
- [Tags]

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 Configuration** dialog appears.

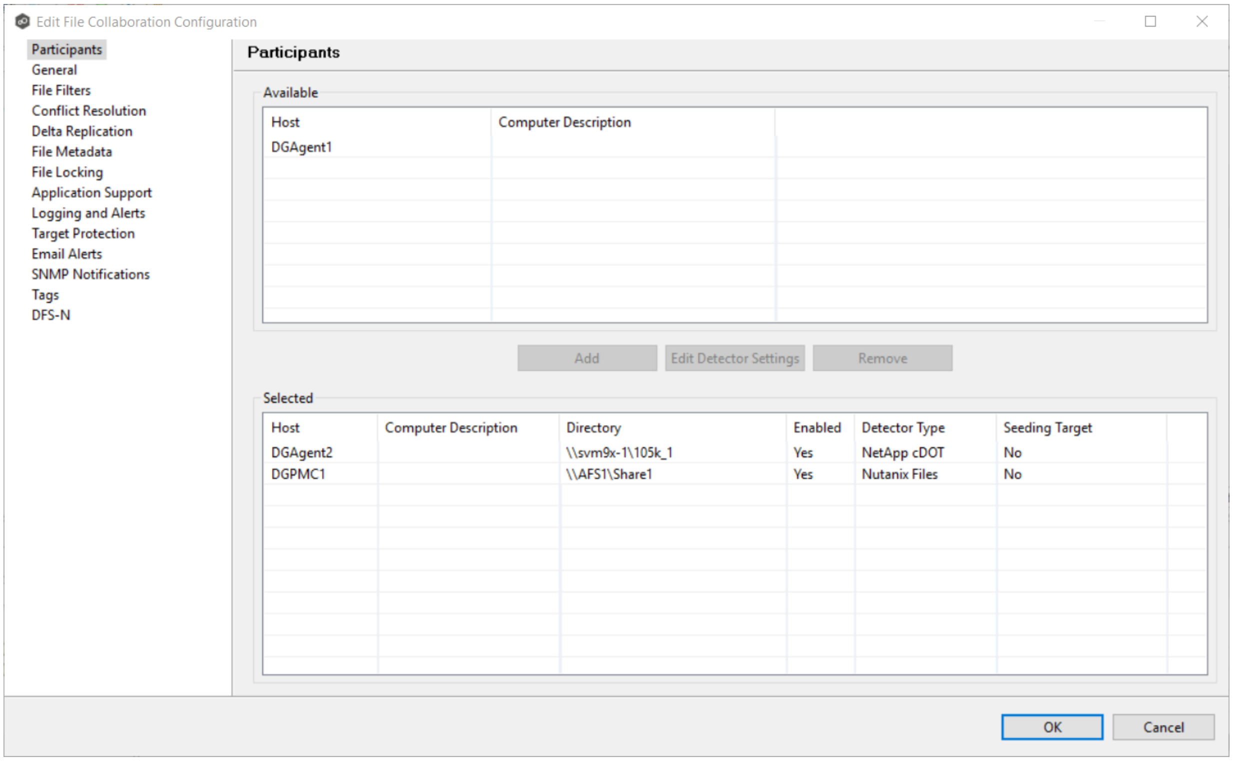

Participants

Available

Host	Computer Description
DGAgent1	

Selected

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent2		\\svm9x-1\105k_1	Yes	NetApp cDOT	No
DGPMC1		\\AFS1\Share1	Yes	Nutanix Files	No

3. Select a configuration item in the navigation tree and make the desired changes:

   - Participants
   - General
   - File Filters

- 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 Collaboration Configuration** wizard allows you to:

- Add and remove participants from a job.
- Modify a participant's attributes.
- Modify a participant's detector settings.

The **Participants** page in the **Edit File Collaboration Configuration** wizard has two tables: the **Available** table and the **Selected** table. The **Available** table lists the available hosts and the **Selected** table lists hosts that have already been added to the job. The **Computer Description** field displays the name of the server that the Peer Agent is running on.

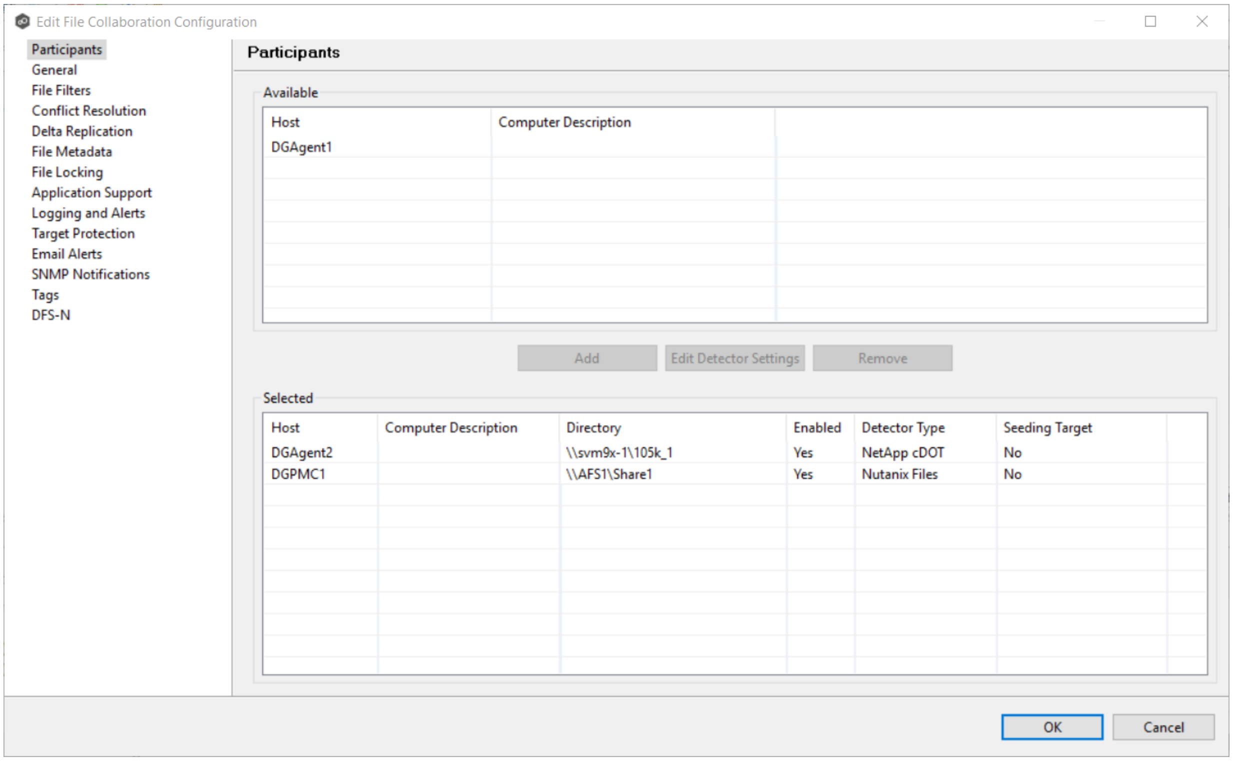

Participants

Available

Host	Computer Description
DGAgent1	

Selected

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent2		\\svm9x-1\105k_1	Yes	NetApp cDOT	No
DGPMC1		\\AFS1\Share1	Yes	Nutanix Files	No

This topic describes adding and removing participants in a File Collaboration job.

Adding a Participant

To add a participant:

1. Click the participant in the **Available** table.

   To be available, a host must have Peer Agent installed and successfully connect to the Peer Management Broker. If a particular host is not displayed in the list, try restarting the Peer Agent Windows Service on that host, and if it successfully connects to the Peer Management Center Broker, then the list will be updated with the computer name of that host.

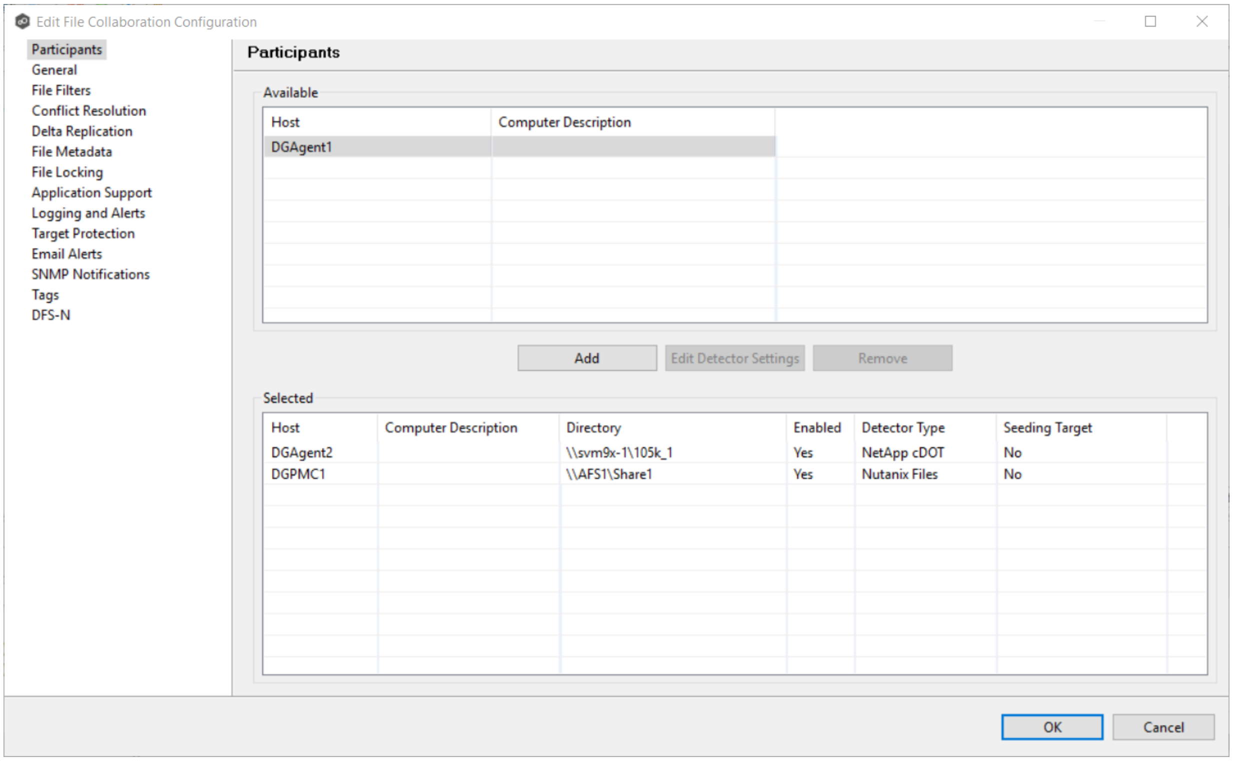

Participants

Available

Host	Computer Description
DGAgent1	

Selected

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent2		\\svm9x-1\105k_1	Yes	NetApp cDOT	No
DGPMC1		\\AFS1\Share1	Yes	Nutanix Files	No

2. Click the **Add** button.

   The participant is moved to the **Selected** table.

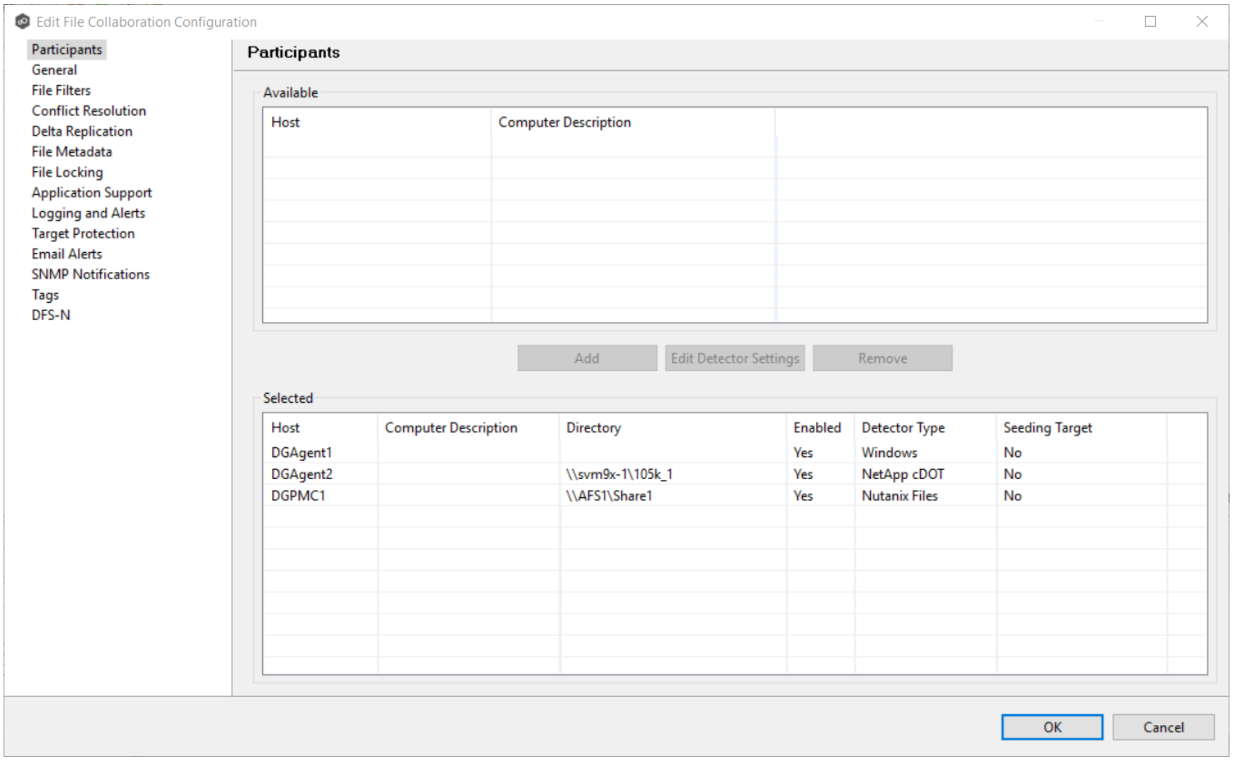

Participants

Available

Host	Computer Description

Selected

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent1			Yes	Windows	No
DGAgent2		\\svm9x-1\105k_1	Yes	NetApp cDOT	No
DGPMC1		\\AFS1\Share1	Yes	Nutanix Files	No

3. (Optional) Enter the computer's name in the **Computer Description** column.

4. Enter the path to the folder to be watched in the **Directory** column.

5. (Optional) Modify whether the participant is a seeding target.

6. (Optional) Modify the participant's detector settings.

7. Click **OK**.

Removing a Participant

To remove a participant:

1. Click the participant in the **Selected** table.

2. Click the **Remove** button.

   The participant is moved to the **Available** table.

   Note: A File Collaboration job must have at least two participants, so if after removing a participant, there is only a single participant, you must add another participant to the job.

3. Click **OK**.

For more information on smart data seeding, see Smart Data Seeding in Advanced Topics or contact support@peersoftware.com.

To set a host as a smart data seeding target:

1. Select the host in the **Selected** table.

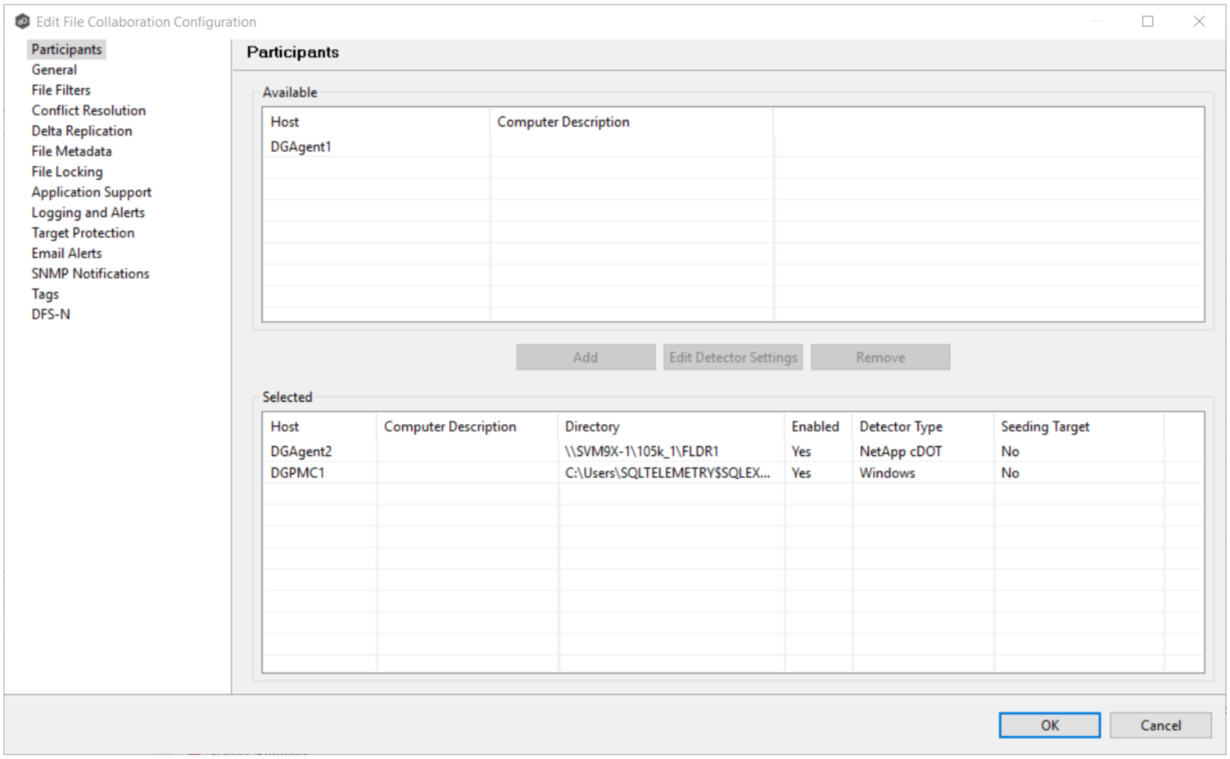

Host	Computer Description
DGAgent1	

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent2		\\SVM9X-1\105k_1\FLDR1	Yes	NetApp cDOT	No
DGPMC1		C:\Users\SQLTELEMETRY$SQLEX...	Yes	Windows	No

2. Select **Yes** in the **Seeding Target** column.

3. Review the information in the message dialog that appears:

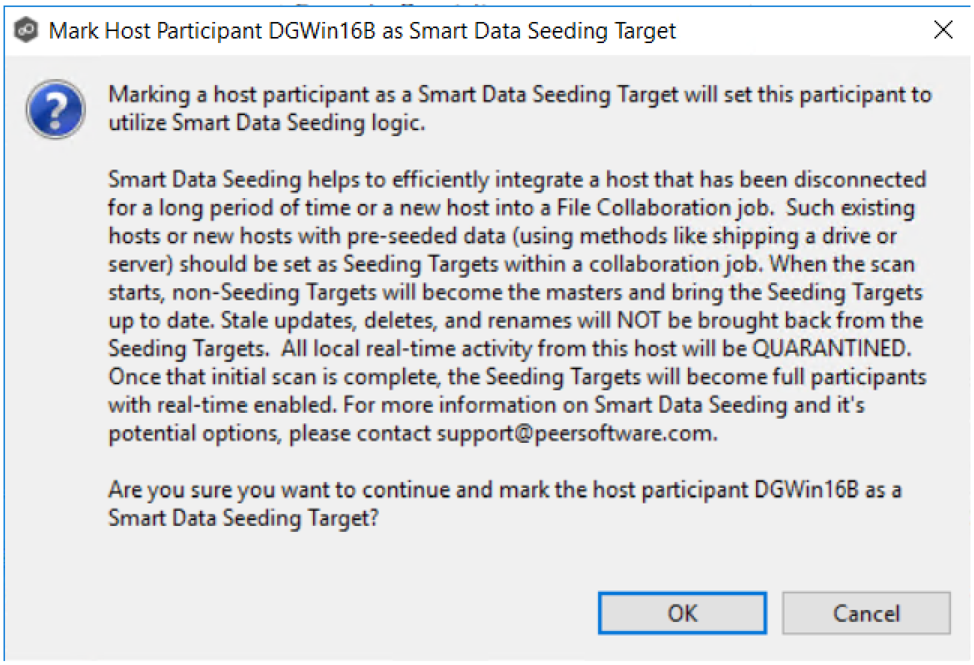

Marking a host participant as a Smart Data Seeding Target will set this participant to utilize Smart Data Seeding logic.

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 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 from this host 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 it's potential options, please contact support@peersoftware.com.

Are you sure you want to continue and mark the host participant DGWin16B as a Smart Data Seeding Target?

4. Click **OK**.

The value in the **Seeding Target** column is updated.

In addition to global real-time detection options that apply to all jobs, you can set additional detection-related options for a specific File Collaboration job. For example, you can exclude real-time events by certain users. This is helpful if you are trying to prevent events generated from backup and/or archival tools from triggering activity.

To modify the detector settings for a host:

1. Select the host in the **Selected** table.

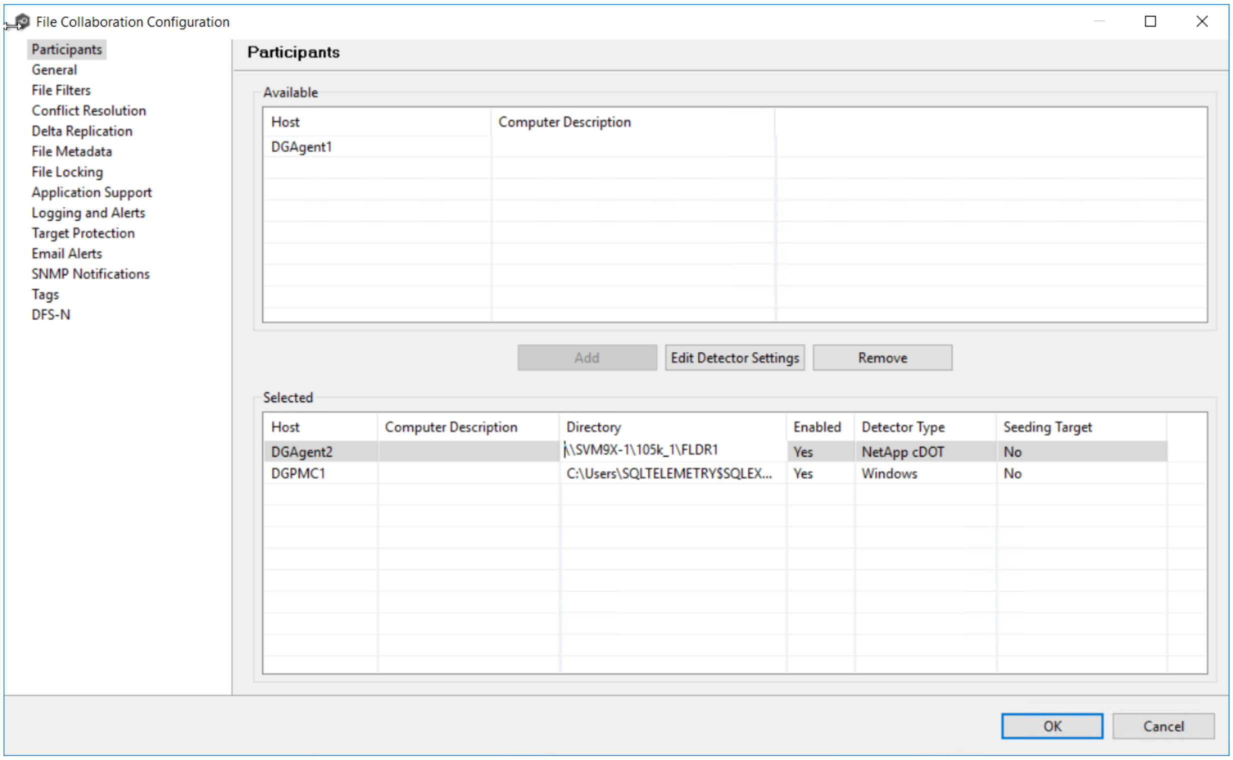

Host	Computer Description
DGAgent1	

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent2		\\SVM9X-1\105k_1\FLDR1	Yes	NetApp cDOT	No
DGPMC1		C:\Users\SQLTELEMETRY$SQLEX...	Yes	Windows	No

2. Click **Edit Detector Settings**.

   The information you are prompted to enter varies, depending on the type of storage platform.

3. Modify the values as needed.

4. Click **OK**.

General

The **General** page in the **Edit File Collaboration Configuration** wizard presents miscellaneous settings pertaining to a File Collaboration job. You may want to consult with Peer Software's support team before modifying these values.

To modify these settings:

1. Enter the values recommended by Peer Software Support.

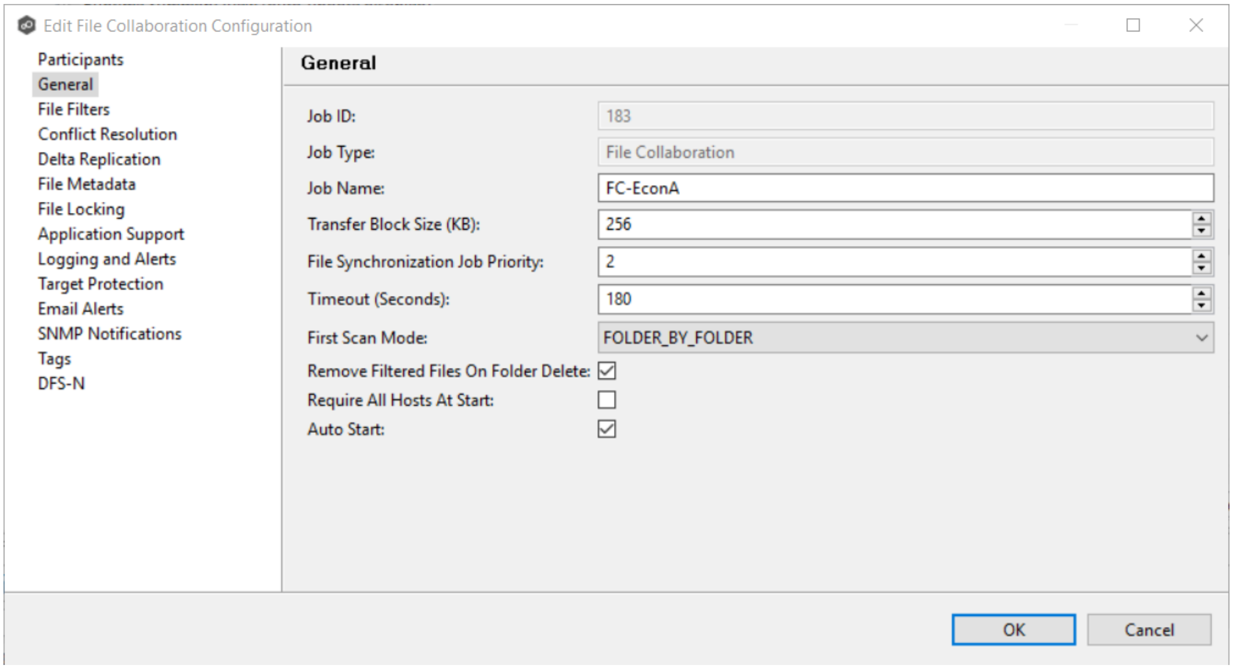

General

Job ID: 183
Job Type: File Collaboration
Job Name: FC-EconA
Transfer Block Size (KB): 256
File Synchronization Job Priority: 2
Timeout (Seconds): 180
First Scan Mode: FOLDER_BY_FOLDER
Remove Filtered Files On Folder Delete: ☑
Require All Hosts At Start: □
Auto Start: ☑

2. Click **OK**.

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.

Job ID	Unique, system-generated job identifier that cannot be edited.
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.
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 results 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.

File Filters

This topic describes how to apply a file filter to a File Collaboration job. File filters are defined global level in Preferences and can then be applied to individual jobs.

- For a description of how to create a file filter, see File Filters in the **Preferences** section for File Collaboration jobs.

- For a discussion of how file filters work, see File and Folder Filters in the **Basic Concepts** section.

The **File Filters** page in the **Edit File Collaboration Configuration** wizard allows you to select which file filters to apply to a File Collaboration job. When the job is run, each selected filter is combined into one large filter (by combining all exclusions and inclusions together). In general, you should have at least one default global file filter that is applied to all jobs and possibly other file filters that apply to specific jobs. However, for most environments, only a single default global file filter is necessary.

To modify which file filters are applied to a File Collaboration job:

1. Select the checkbox next to the filters you want applied to the job.

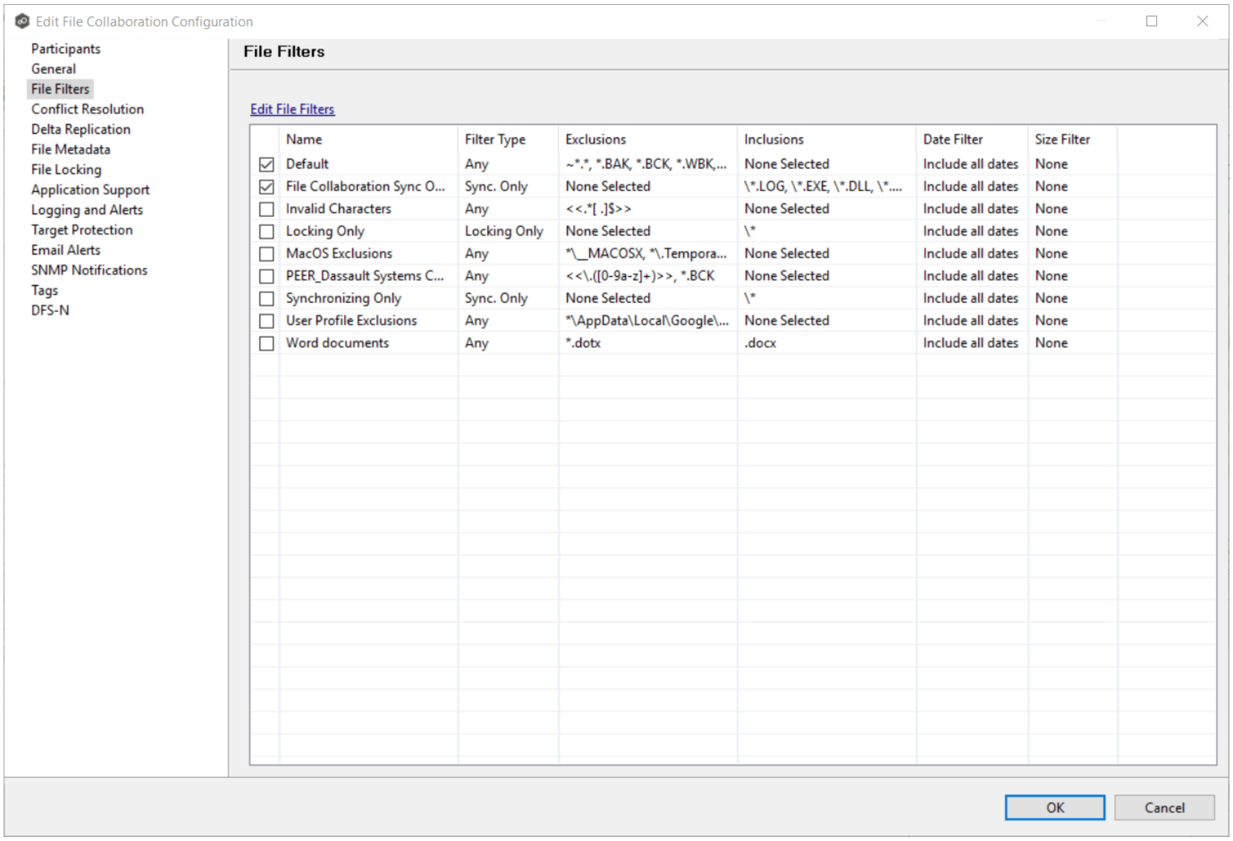

	Name	Filter Type	Exclusions	Inclusions	Date Filter	Size Filter
☑	Default	Any	~*.*, *.BAK, *.BCK, *.WBK,...	None Selected	Include all dates	None
☑	File Collaboration Sync O...	Sync. Only	None Selected	*.LOG, *.EXE, *.DLL, *....	Include all dates	None
☐	Invalid Characters	Any	<<.*[.]$>>	None Selected	Include all dates	None
☐	Locking Only	Locking Only	None Selected	*	Include all dates	None
☐	MacOS Exclusions	Any	*__MACOSX, *\.Tempora...	None Selected	Include all dates	None
☐	PEER_Dassault Systems C...	Any	<<\.([0-9a-z]+)>>, *.BCK	None Selected	Include all dates	None
☐	Synchronizing Only	Sync. Only	None Selected	*	Include all dates	None
☐	User Profile Exclusions	Any	*\AppData\Local\Google\...	None Selected	Include all dates	None
☐	Word documents	Any	*.dotx	.docx	Include all dates	None

2. Click **OK**.

Conflict Resolution

The **Conflict Resolution** page in the **Edit File Collaboration Configuration** wizard allows you to specify the file conflict resolution options to use during the initial scan when a file conflict exists for a file between two or more hosts.

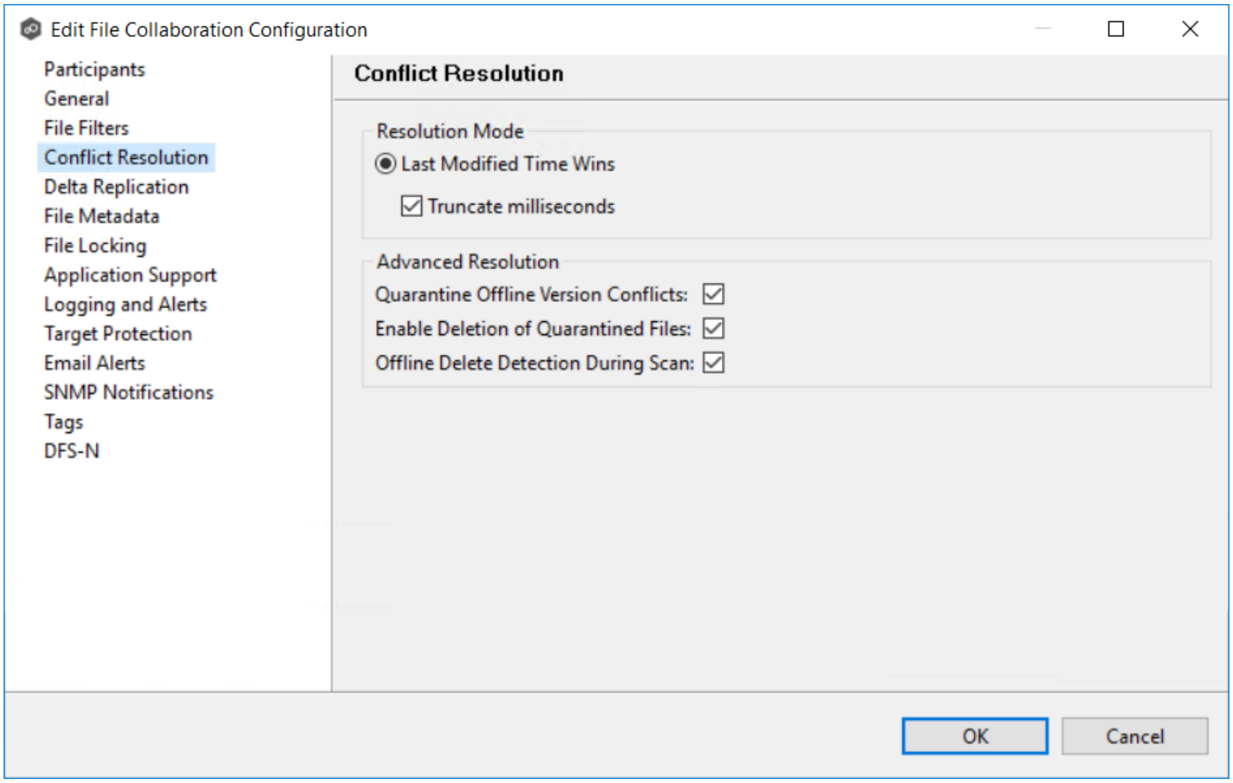

Conflict Resolution

Resolution Mode

● Last Modified Time Wins

☑ Truncate milliseconds

Advanced Resolution

Quarantine Offline Version Conflicts: ☑

Enable Deletion of Quarantined Files: ☑

Offline Delete Detection During Scan: ☑

To modify conflict resolution settings for the File Collaboration job:

1. Select the resolution mode:

Last Modified Time Wins	A file's modification time will be used to designate an instance as a resolution candidate. The later the modification time, the greater the likelihood for a file's selection. Option: **Truncate milliseconds:** When comparing the time stamps of a file on two or more hosts, truncate the millisecond value from each time stamp.
None (Manual Resolution)	This is an advanced option. Contact Peer Software to enable. When selected, any file conflicts that are encountered during the initial synchronization process will result in quarantines that are added to the File Conflict List. These file conflicts must be resolved manually by selecting the host with the correct version of the file from the conflict list.

All the types listed above have the potential for producing multiple resolution candidates. A collaboration session can be configured with any one of the available conflict options. If a option produces more than one candidate for a conflicted file, a winner will be selected arbitrarily.

2. Select the **Advanced Resolution** options you want applied.

Quarantine Offline Version Conflicts	Enable 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.
Enable Deletion of Quarantine d Files	If a file that is quarantined is deleted, Peer Management Center will process the delete event and remove the quarantine when this option is enabled.
Offline Delete Detection During Scan	If this option is enabled and target protection is enabled, and it can be determined that a file or folder has been deleted since the session was stopped, then the file or folder will be deleted from all hosts. If this option is not enabled then the deleted file or folder will be brought back to any host where it was removed.

3. Click **OK**.

Delta Replication

The **Delta Replication** page in the **Edit File Collaboration Configuration** wizard 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 the use of 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.

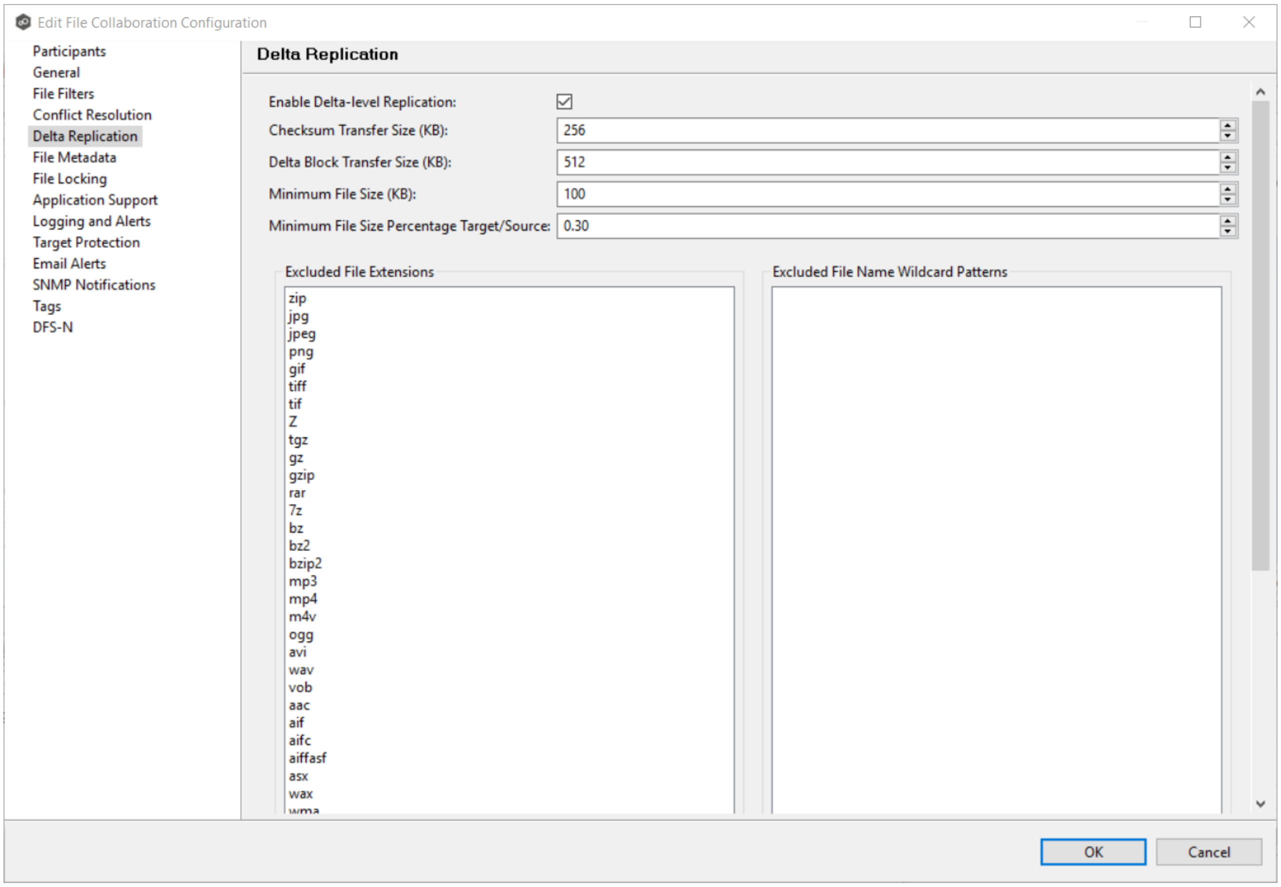

Delta Replication

Enable Delta-level Replication: ☑

Checksum Transfer Size (KB): 256

Delta Block Transfer Size (KB): 512

Minimum File Size (KB): 100

Minimum File Size Percentage Target/Source: 0.30

Excluded File Extensions

zip
jpg
jpeg
png
gif
tiff
tif
Z
tgz
gz
gzip
rar
7z
bz
bz2
bzip2
mp3
mp4
m4v
ogg
avi
wav
vob
aac
aif
aifc
aiffasf
asx
wax
wma

Excluded File Name Wildcard Patterns

To modify delta-level replication options:

1. Modify the following the fields as necessary.

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 wildcard 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 Wildcard Patterns	Enter a list of file name wildcard patterns to be excluded from delta encoding. If a file name matches any wildcard 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**.

File Metadata

The **File Metadata** page in the **Edit File Collaboration Configuration** wizard allows you to modify your file metadata synchronization settings and presents additional options for metadata replication. See File Metadata Replication in Advanced Topics for more information about file metadata replication.

To enable file metadata replication:

1. Select when you want the metadata synchronized (you can select one or both of the 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 (DACL, SACL, Owner.) 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 (DACL, SACL, Owner) will be synchronized during the initial scan.

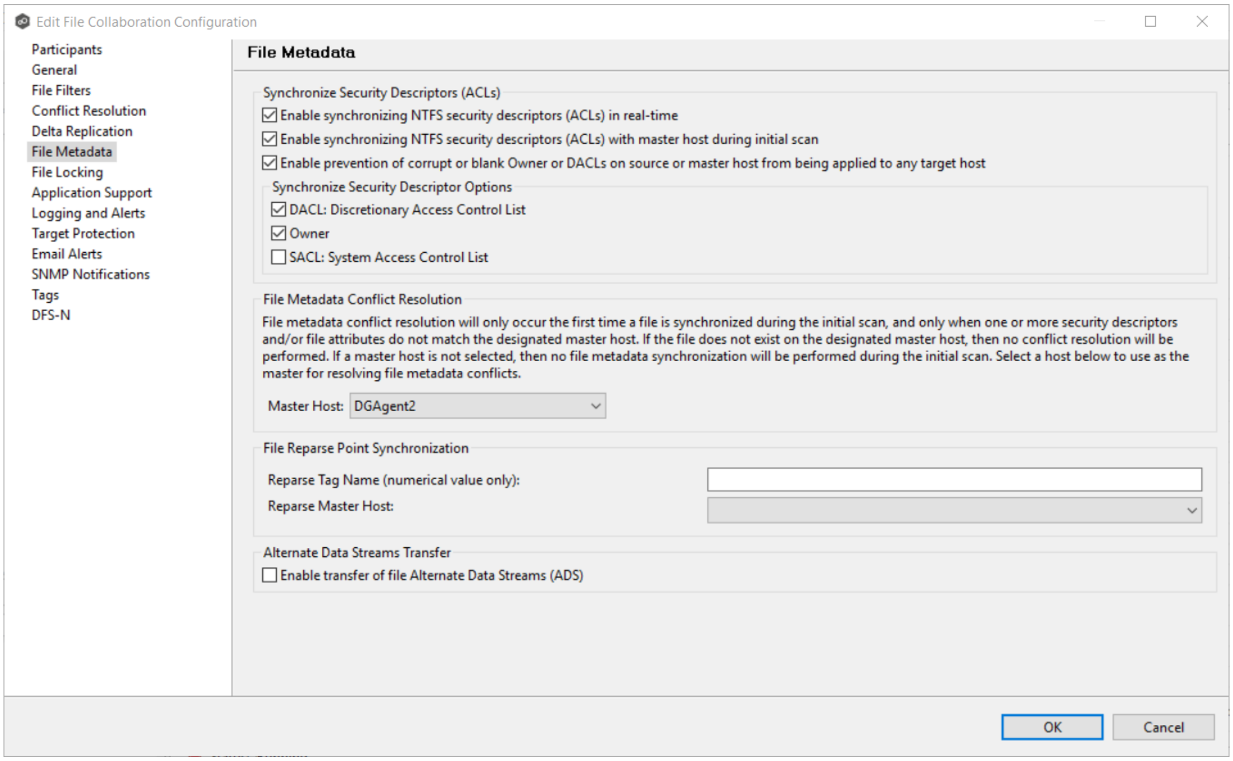

File Metadata

Synchronize Security Descriptors (ACLs)

☑ Enable synchronizing NTFS security descriptors (ACLs) in real-time

☑ Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan

☑ Enable prevention of corrupt or blank Owner or DACLs on source or master host from being applied to any target host

Synchronize Security Descriptor Options

☑ DACL: Discretionary Access Control List

☑ Owner

☐ SACL: System Access Control List

File Metadata Conflict Resolution

File metadata conflict resolution will only occur the first time a file is synchronized during the initial scan, and only when one or more security descriptors and/or file attributes do not match the designated master host. If the file does not exist on the designated master host, then no conflict resolution will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan. Select a host below to use as the master for resolving file metadata conflicts.

Master Host: DGAgent2

File Reparse Point Synchronization

Reparse Tag Name (numerical value only):

Reparse Master Host:

Alternate Data Streams Transfer

☐ Enable transfer of file Alternate Data Streams (ADS)

2. Click **OK** in the informational dialog that appears after selecting a metadata option.

3. Select which security descriptor components (DACL, SACL and Owner) are synchronized.

   In general, you will usually need to synchronize only DACLs. If you need to synchronize SACLs or Owner, then the user that a Peer Agent service is run under on each participating host must have permission to read and write SACLs and Owner.

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.

   Conflict resolution for file metadata 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 will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan.

5. (Optional) Enter values for one or both of the 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/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 verse 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.

6. (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.

7. Click **OK**.

File Locking

The **File Locking** page in the **Edit File Collaboration Configuration** wizard presents options pertaining to how source and target files are locked by Peer Management Center.

To modify file locking options:

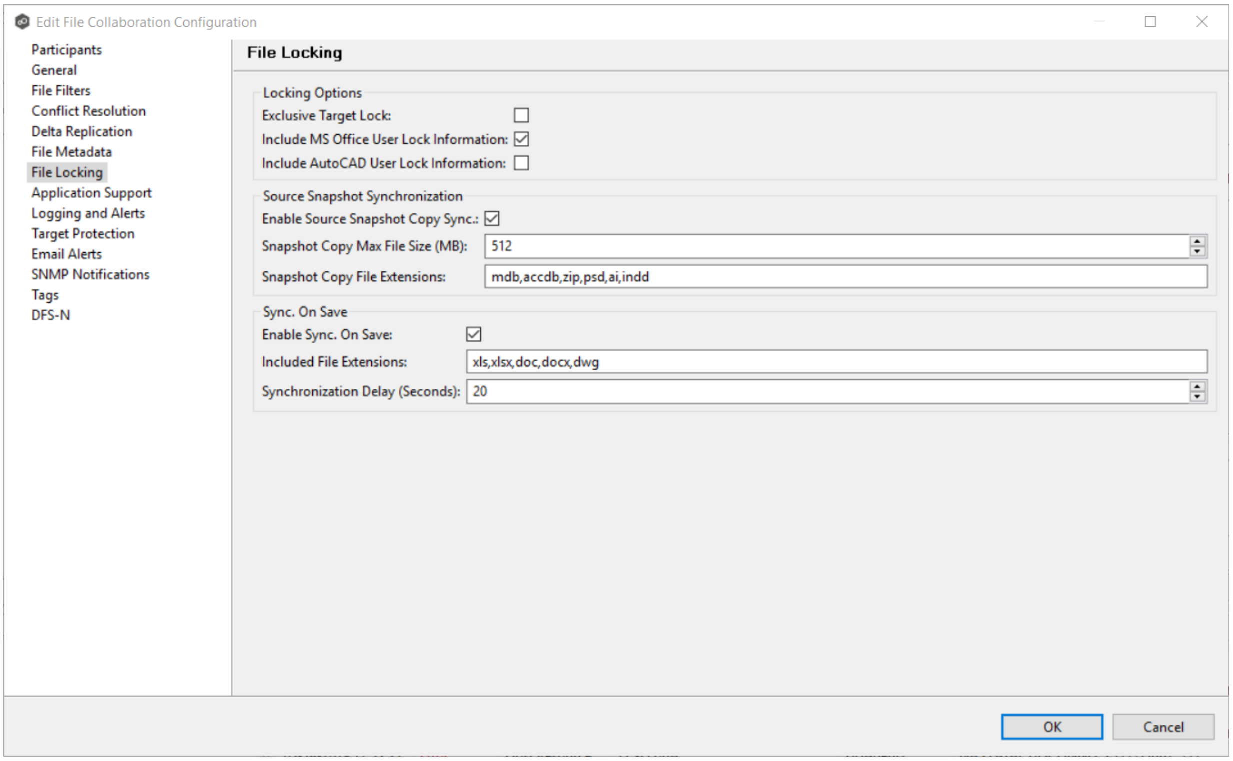

File Locking

Locking Options

Exclusive Target Lock: □

Include MS Office User Lock Information: ☑

Include AutoCAD User Lock Information: □

Source Snapshot Synchronization

Enable Source Snapshot Copy Sync.: ☑

Snapshot Copy Max File Size (MB): 512

Snapshot Copy File Extensions: mdb,accdb,zip,psd,ai,indd

Sync. On Save

Enable Sync. On Save: ☑

Included File Extensions: xls,xlsx,doc,docx,dwg

Synchronization Delay (Seconds): 20

Modify these fields as needed:

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 that are automatically optimized. You can modify your selections when editing the job.

To modify which applications are optimized:

1. Select the applications to be optimized.

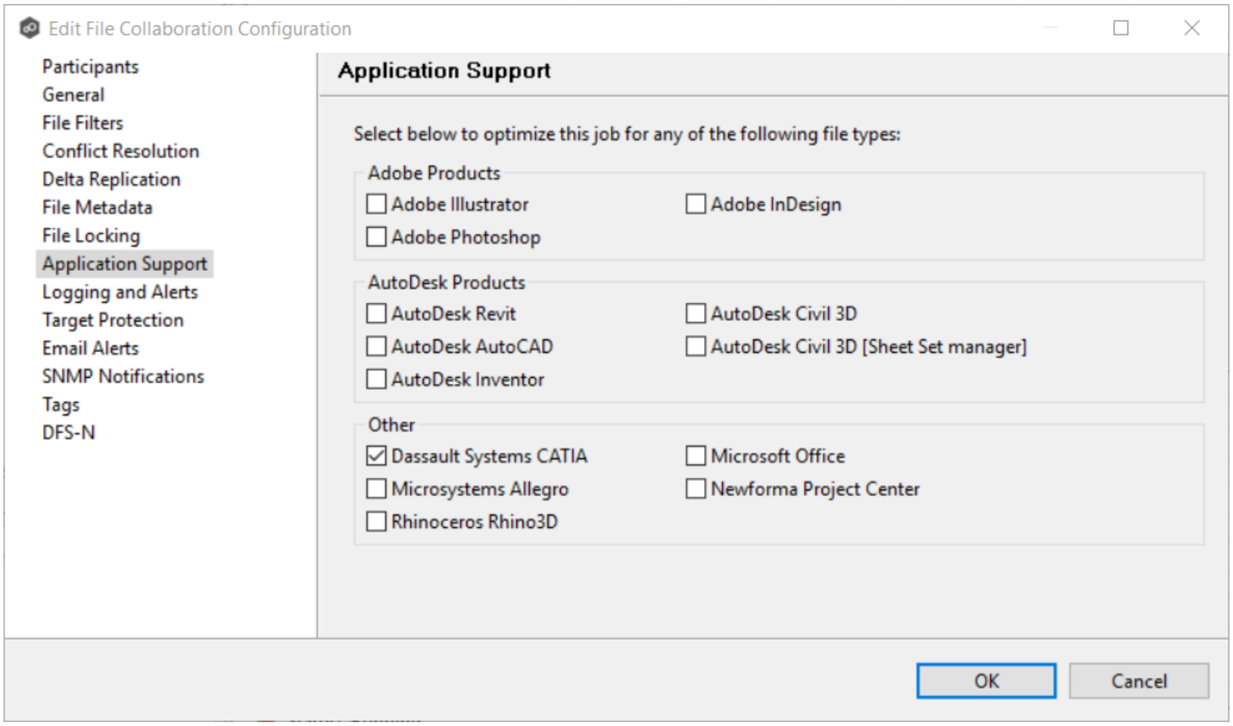

Edit File Collaboration Configuration

Application Support

Select below to optimize this job for any of the following file types:

Adobe Products

□ Adobe Illustrator □ Adobe InDesign

□ Adobe Photoshop

AutoDesk Products

□ AutoDesk Revit □ AutoDesk Civil 3D

□ AutoDesk AutoCAD □ AutoDesk Civil 3D [Sheet Set manager]

□ AutoDesk Inventor

Other

☑ Dassault Systems CATIA □ Microsoft Office

□ Microsystems Allegro □ Newforma Project Center

□ Rhinoceros Rhino3D

2. Click **OK**.

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.

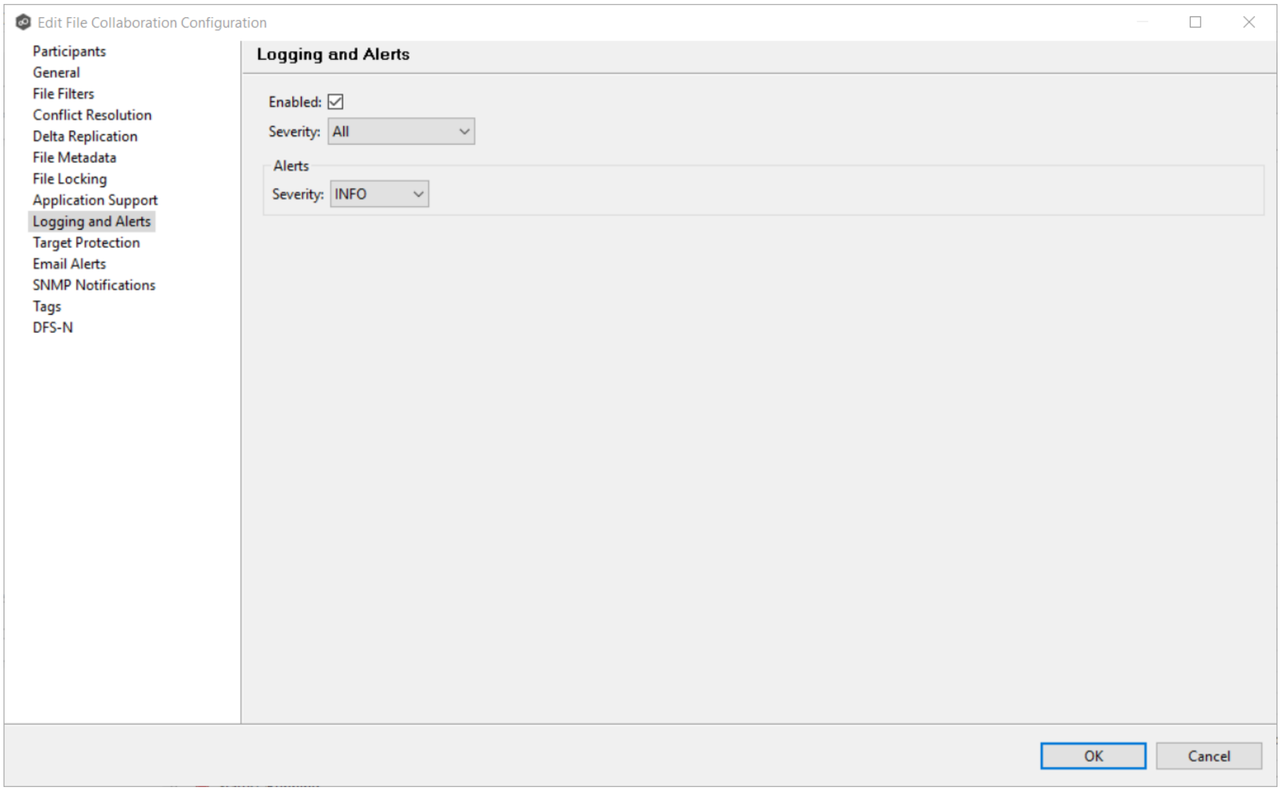

Logging and Alerts

Enabled: ☑

Severity: All

Alerts

Severity: INFO

Below is a list of logging fields and their descriptions:

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 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

Configured in the screen shown above, various types of alerts will 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.

Target protection configuration is available by selecting **Target Protection** from the tree node within the File Collaboration Configuration dialog.

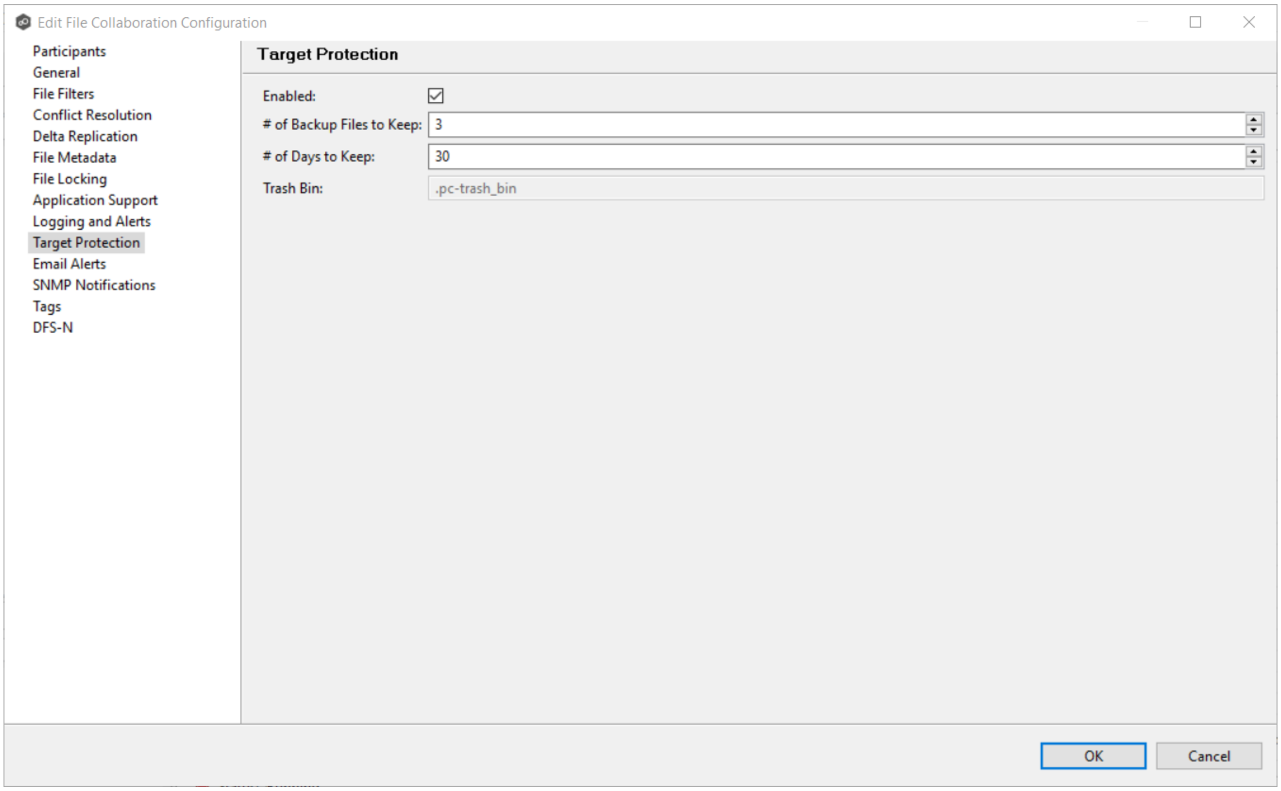

Target Protection

Enabled: ☑

# of Backup Files to Keep: 3

# of Days to Keep: 30

Trash Bin: .pc-trash_bin

Modify the fields as needed:

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 Configuration** wizard 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.

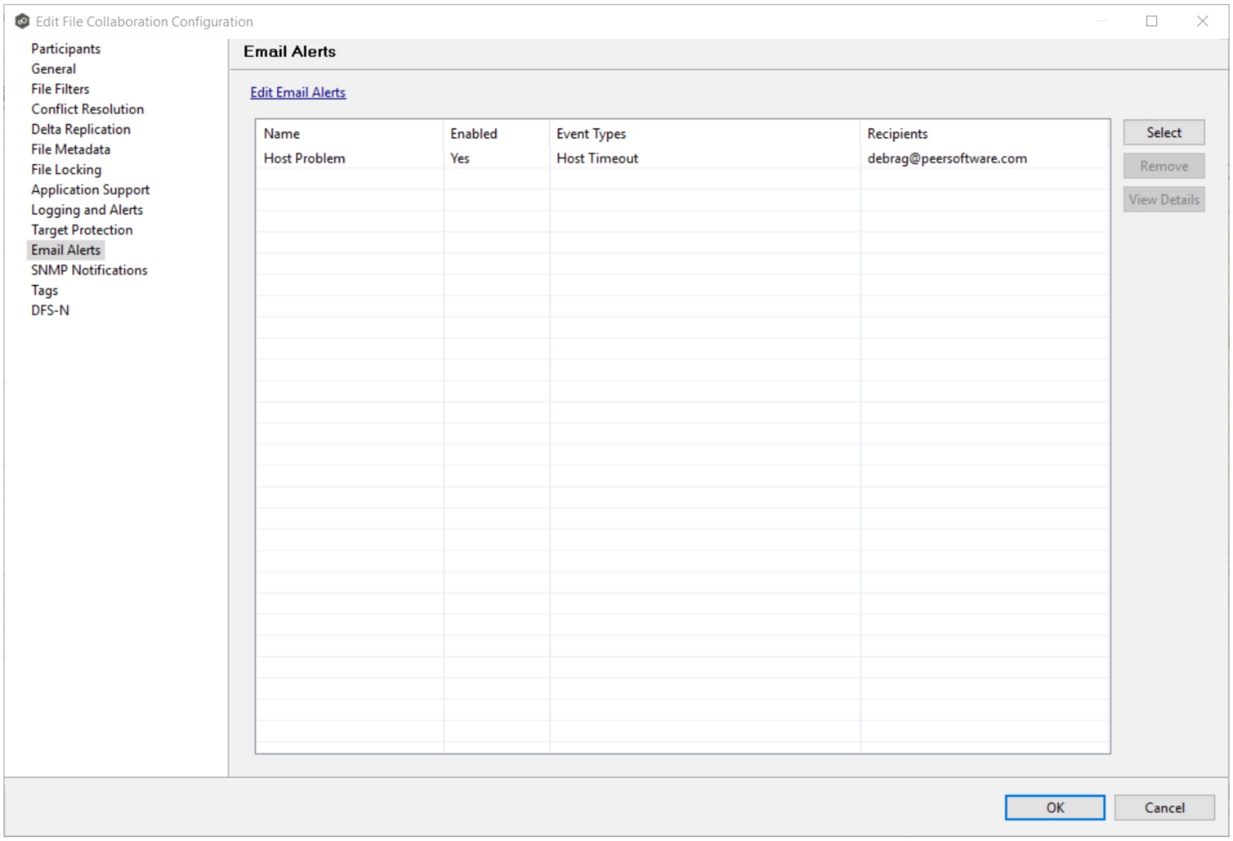

Name	Enabled	Event Types	Recipients
Host Problem	Yes	Host Timeout	debrag@peersoftware.com

The **Select Email Alert** dialog opens.

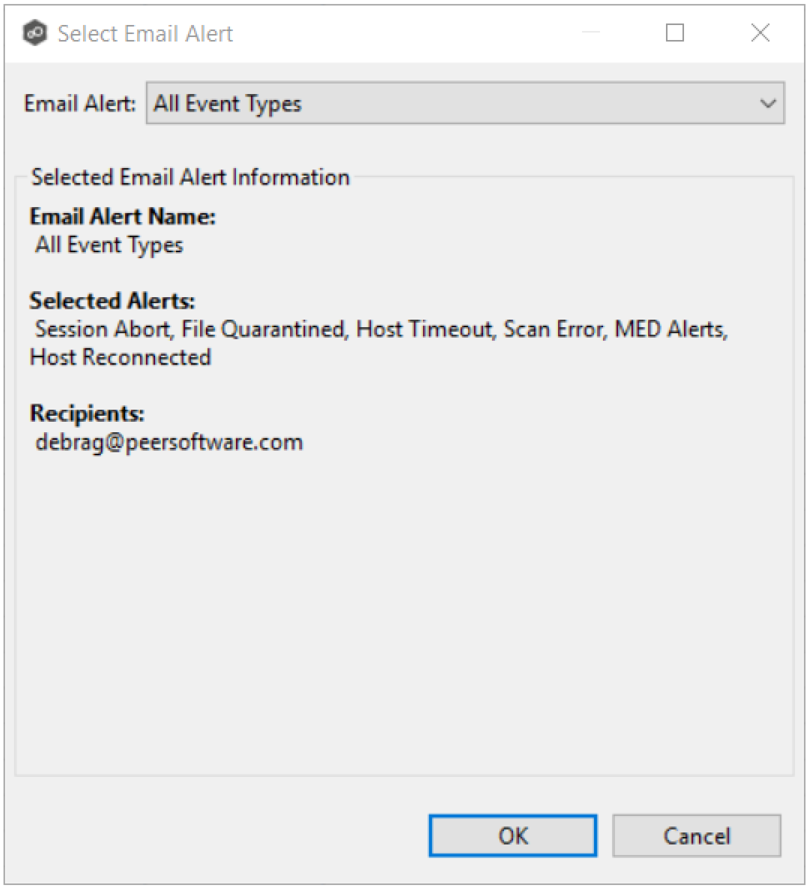
Select Email Alert
Email Alert: All Event Types
Selected Email Alert Information

Email Alert Name:
All Event Types

Selected Alerts:
Session Abort, File Quarantined, Host Timeout, Scan Error, MED Alerts, Host Reconnected

Recipients:
debrag@peersoftware.com

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**.

SNMP Notifications

The **SNMP Notifications** page in the **Edit File Collaboration Configuration** wizard allows you to select which SNMP notifications to apply 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**.

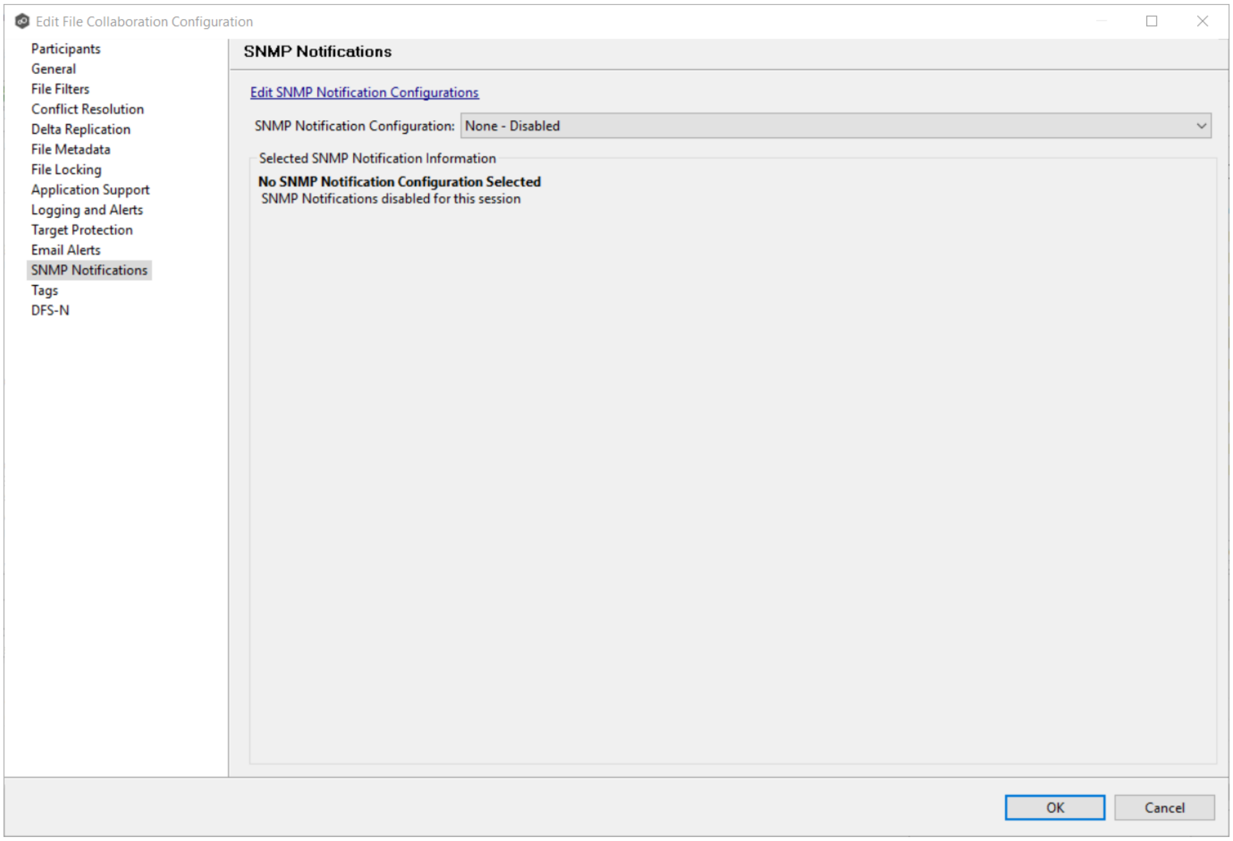

SNMP Notification Configuration: None - Disabled

Selected SNMP Notification Information

No SNMP Notification Configuration Selected

SNMP Notifications disabled for this session

2. Click **OK**.

Tags

The **Tags** page in the **Edit File Collaboration Configuration** wizard allows you to 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.

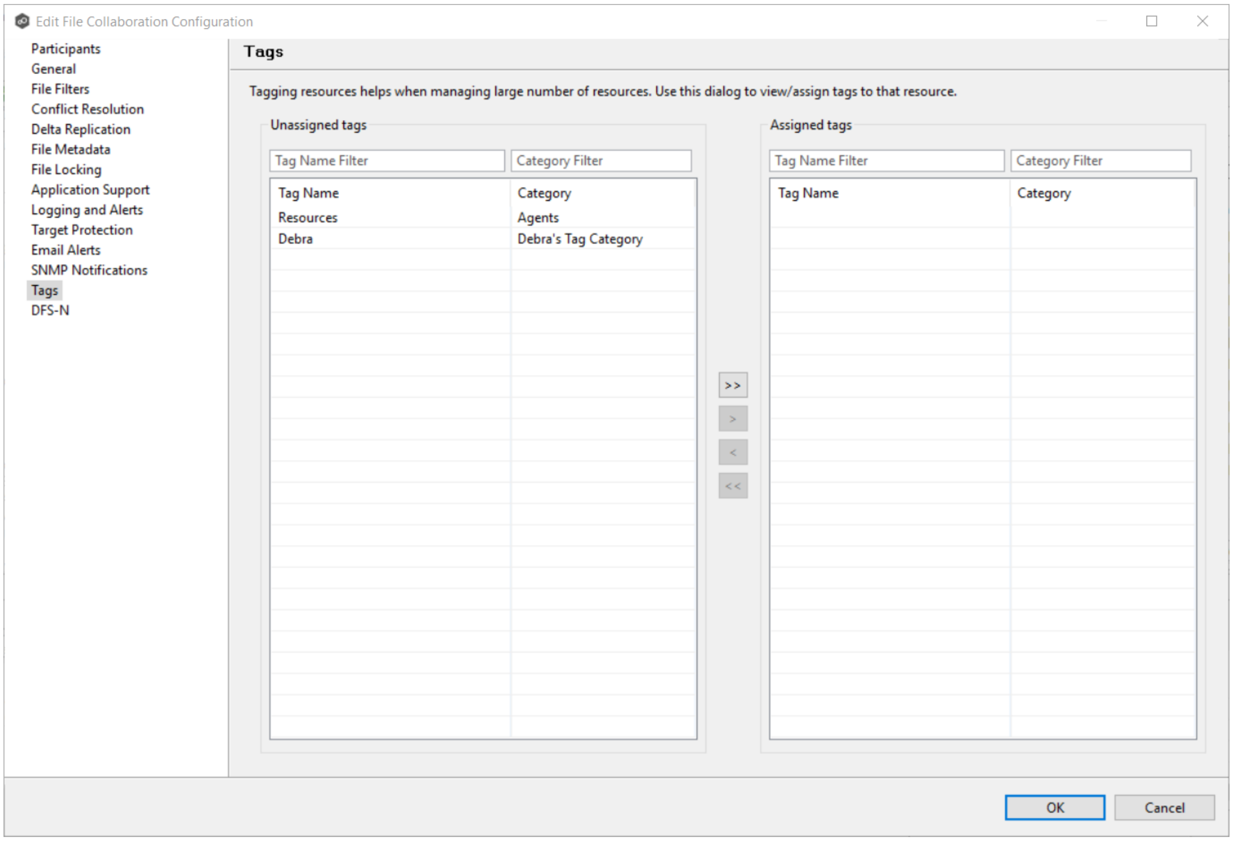

Tagging resources helps when managing large number of resources. Use this dialog to view/assign tags to that resource.

Unassigned tags

Tag Name	Category
Resources	Agents
Debra	Debra's Tag Category

Assigned tags

Tag Name	Category

DFS-N

The **DFS-N** page in the **Edit File Collaboration Configuration** wizard presents options for linking a DFS namespace folder to this job. See Link a Namespace Folder with an Existing File Collaboration or Synchronization Job for more information.

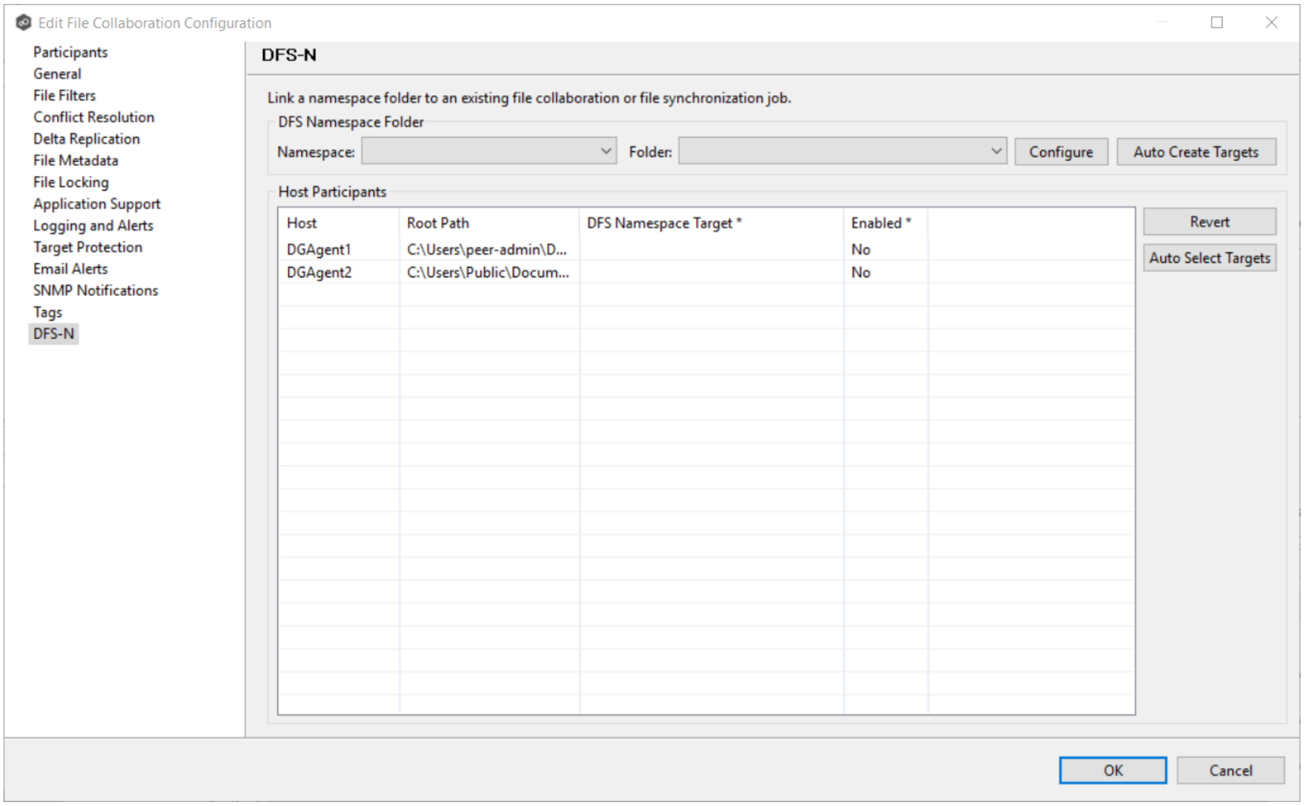

DFS-N

Link a namespace folder to an existing file collaboration or file synchronization job.

DFS Namespace Folder

Namespace:

Folder:

Host Participants

Host	Root Path	DFS Namespace Target *	Enabled *
DGAgent1	C:\Users\peer-admin\D...		No
DGAgent2	C:\Users\Public\Docum...		No

Editing Multiple Jobs

The 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 the 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:

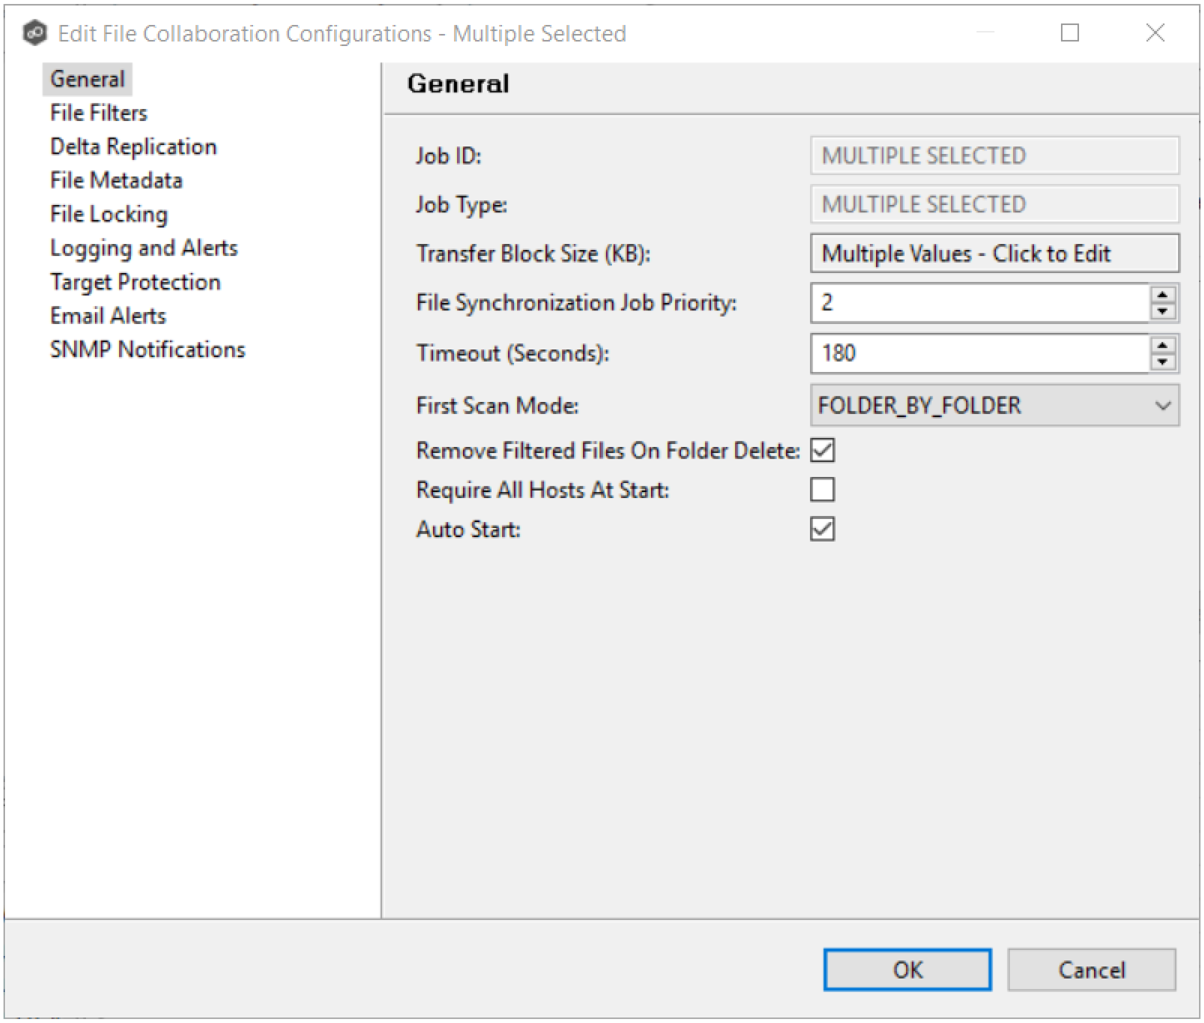

Edit File Collaboration Configurations - Multiple Selected

General

Job ID: MULTIPLE SELECTED

Job Type: MULTIPLE SELECTED

Transfer Block Size (KB): Multiple Values - Click to Edit

File Synchronization Job Priority: 2

Timeout (Seconds): 180

First Scan Mode: FOLDER_BY_FOLDER

Remove Filtered Files On Folder Delete: ☑

Require All Hosts At Start: □

Auto Start: ☑

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:

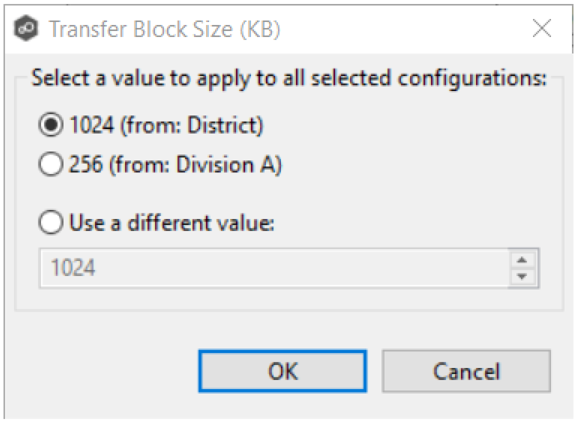

Transfer Block Size (KB)

Select a value to apply to all selected configurations:

● 1024 (from: District)

○ 256 (from: Division A)

○ Use a different value:

1024

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 check box 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 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
- Managing File Conflicts

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 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 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 File 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 **File Collaboration Job Summary** view, and then choose **Start** from the pop-up menu.

- Open a job and then click the **Start/Stop** button in the bottom left corner of the job's **Summary** tab (shown below).

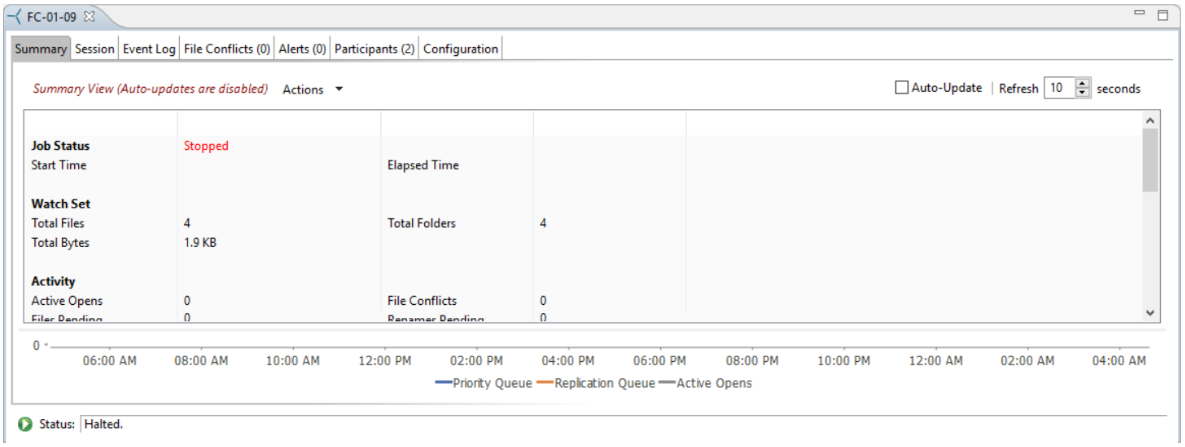

Summary View (Auto-updates are disabled)

☐ Auto-Update | Refresh 10 seconds

Job Status	Stopped		
Start Time		Elapsed Time	
Watch Set			
Total Files	4	Total Folders	4
Total Bytes	1.9 KB		
Activity			
Active Opens	0	File Conflicts	0
Filer Pending	0	Renamer Pending	0

Status: Halted.

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.

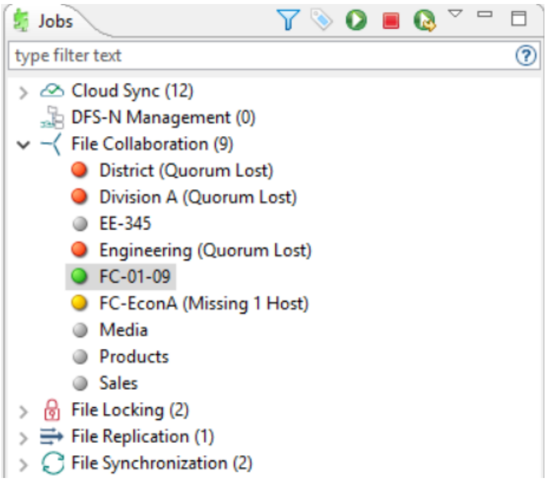

- Cloud Sync (12)
- DFS-N Management (0)
- File Collaboration (9)
  - District (Quorum Lost)
  - Division A (Quorum Lost)
  - EE-345
  - Engineering (Quorum Lost)
  - FC-01-09
  - FC-EconA (Missing 1 Host)
  - Media
  - Products
  - Sales
- File Locking (2)
- File Replication (1)
- File Synchronization (2)

Stopping a File Collaboration Job

You can stop a File Collaboration job at any time by 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 host instance. For more information on disabling auto-restart at the job level, see Participant View.

To enable all File Collaboration jobs to auto-restart:

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

2. Select **File Collab, Sync, and Locking** in the navigation tree.

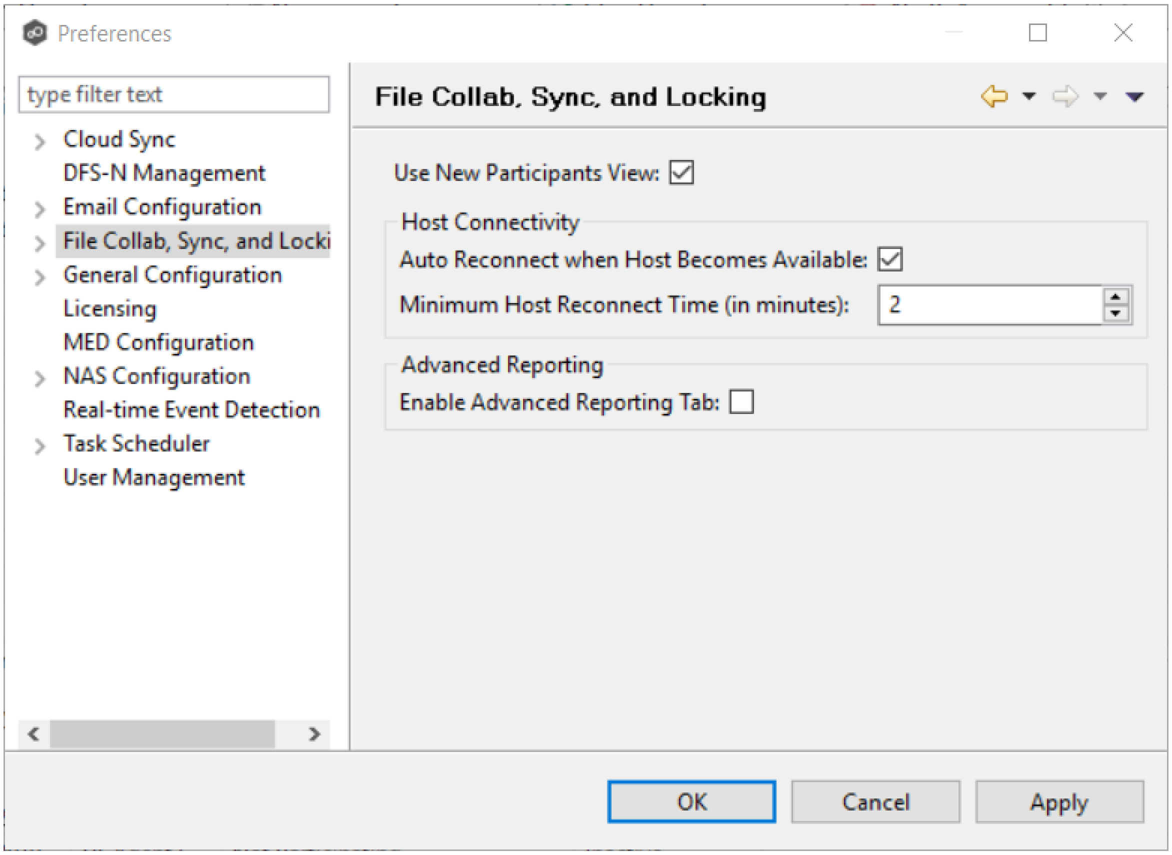

File Collab, Sync, and Locking

Use New Participants View: ☑

Host Connectivity

Auto Reconnect when Host Becomes Available: ☑

Minimum Host Reconnect Time (in minutes): 2

Advanced Reporting

Enable Advanced Reporting Tab: □

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.

Unavailable Hosts

If a host becomes unavailable while a File Collaboration job is running, and is unreachable within the configured timeout period (specified within 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, then the host will be pinged, and if still no response, the host will be taken out of the running session, a FATAL event will be logged, and the

Participants view for the job will be 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) will be sent.

If auto-restart not enabled, you will need to stop and start the File Collaboration job in order 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 jobs.

Quorum

In order for a File Collaboration 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 collaboration session will automatically be terminated. If email alerts and/or SNMP notifications are configured and enabled for **Session Aborts**, then the appropriate message(s) will be sent.

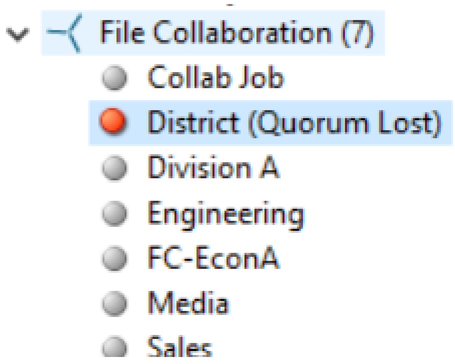

Managing File Conflicts

Files conflicts can occur for the following reasons:

- A file is already opened by a user when a File Collaboration job is started and the file size and timestamp does not match the other target hosts.
- A file is already opened by two or more users when a File Collaboration job is started.
- Two or more users open a file at the same time before all files can be locked down by the running File Collaboration job.
- A file was modified on two or more hosts between job restarts or network outages.

- A general I/O failure occurs on the source host after the file has been modified, but before the file is synchronized to all target hosts. In this case, the file is automatically quarantined.

When a file conflict is detected, the file is placed in the File Conflicts list and assigned a conflict status, which determines how the conflict is resolved. The File Conflicts list is displayed in the File Conflicts view.

For an elaboration on conflict scenarios, see file conflict scenarios.

- A job is started and Initial Scan Logic is performed on a file:

  If the file has never been synchronized by Peer Management Center and if file sizes and last modified times do not match on all participating hosts, or if the file does not exist on one or more hosts, then the file will be synchronized based on the file conflict resolution strategy, which is typically the most recent last modified time.

  Files that have previously been synchronized by Peer Management Center where just a single file's last modified timestamp is newer than the last recorded timestamp, then that file will be synchronized to all other hosts. However, if two or more files have a more recent last modified timestamp than was last recorded timestamp, then the file will be quarantined (this is the default behavior and can be disabled by deselecting the file conflict resolution strategy **Quarantine Offline Version Conflicts** option).

- A single user has a file opened before starting a collaboration job:

  A file conflict will be created with a status of **Pending Initial Synchronization**. After the user closes the file, if all file sizes and timestamps match, then the file conflict is removed and no synchronization is performed. However, if any file last modified times or file sizes do not match, the file will be synchronized or quarantined based on the file conflict resolution strategy and according to the initial scan logic detailed above. Once the file is synchronized, the file conflict will be removed.

- Two or more users have a file open before starting a collaboration job

  A file conflict will be created with a status of **Pending Conflict Resolution**. After the users close all files, the conflict will be removed if the last modified timestamp matches on all files. Otherwise, if the file has never been synchronized by Peer Management Center, then the file conflict status will be updated to Quarantined. However, if the file has previously been synchronized by Peer Management Center, then the file will be synchronized or quarantined based on the file conflict resolution strategy and according to the initial scan logic detailed above.

- Two or more users open a file at the same time

In the rare situation when two users open a file at the same time, or in-and-around the same time and Peer Management Center is unable to obtain corresponding locks on target hosts before this happens (this is dependent on WAN latency and other factors), then a file conflict will be created with a status of **Pending Conflict Resolution**. After all users close the files, file lock conflict resolution will be performed as follows:

- If all files last modified timestamps and file sizes match, then the file conflict will be removed.
- If only a single file has been modified, then the file that changed is synchronized or quarantined based on the configured file conflict resolution strategy and according to the initial scan logic detailed above.
- If two or more files have been modified since it was opened, then the file conflict status will be updated to quarantined.

When a file conflict occurs, the status is set to one of the the following statuses:

- **Pending Conflict Resolution** if the file has already been verified or synchronized by the initial synchronization process.
- **Pending Initial Synchronization** if the file hasn't been verified or synchronized.
- **Quarantined** if the conflict is a result of a fatal I/O error on the source.

Note: If a File Collaboration job is stopped before a file conflict with a status of **Pending Conflict Resolution** is resolved, then that file is automatically quarantined the next time the File Collaboration job is started.

The resolution strategies for the three conflict statuses are as follows:

Conflict Status	Resolution Strategy
Pending Conflict Resolution This status is assigned to files that have already been verified or synchronized by the	When all files in use are closed by users on the source hosts, the files will be analyzed to determine if a file conflict has occurred as follows: • If more than one file has been modified, then the file will be quarantined by updating the file conflict status to Quarantined.

session via the initial synchronization process.	• If only one file as been modified, then that file will be used as the source, synchronized with all other participating hosts, and removed from the File Conflicts list. • If no files have been modified, then no action will be taken and the file will be removed from the File Conflict list.
Pending Initial Synchronization This status is assigned to files that have not been verified or synchronized by session via the initial synchronization process.	When all files in use are closed by users on the source hosts, then standard file conflict resolution will be performed based on the configured File Conflict Resolvers. However, if the **Quarantine Offline Multi-Edits** option is enabled, then if a file is modified on two or more hosts while the collaboration session is not running, and the last modified timestamps are all newer then the last timestamp recorded by the collaboration session, then the file will be quarantined.
Quarantined	A file will be quarantined when a file conflict with **Pending Conflict Resolution** status cannot be resolved or a fatal I/O error occurs. Quarantined files need to be explicitly removed from the File Conflicts list.

.

Once a file is marked as **Quarantined**, the file no longer participates in collaboration, and thus changes to any version of the file will not be propagated to other hosts. However, subsequent file activity on a quarantined file will be logged in the Event Log as a warning so that you can determine who modified the file while it was quarantined.

Quarantined files are saved to disk and will survive session restarts. The File Conflicts list displays the time and date of the quarantine along with an error message indicating the reason for the quarantine.

A Quarantined File event is also logged in the Event Log and you can obtain a more detailed reason for the quarantine by analyzing the Event Log file(s). In addition, if email alerts and/or SNMP notifications are configured and enabled for **File Quarantines**, then the appropriate message(s) will be sent.

Removing a File from Quarantine

You must explicitly remove a file from quarantine in order to have it participate in the collaboration session once again.

You may also chose to perform no action, in which case, the file is removed from the File Conflict 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 removing the file conflict, then the Status field in the File Conflict table is updated to reflect the error.

To remove a file (or multiple files) from quarantine:

1. Select the file(s) in the File Conflicts 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 anybody. 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 Collaboration Views

This section describes File Collaboration views. The views can be grouped into two types:

- Summary
- Runtime

File Collaboration Summary Views

The three File Collaboration Summary views are:

- Collab and Sync Summary View
- Runtime Reports View
- Dashboard Summary View

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

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

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 file conflicts from stopped jobs.

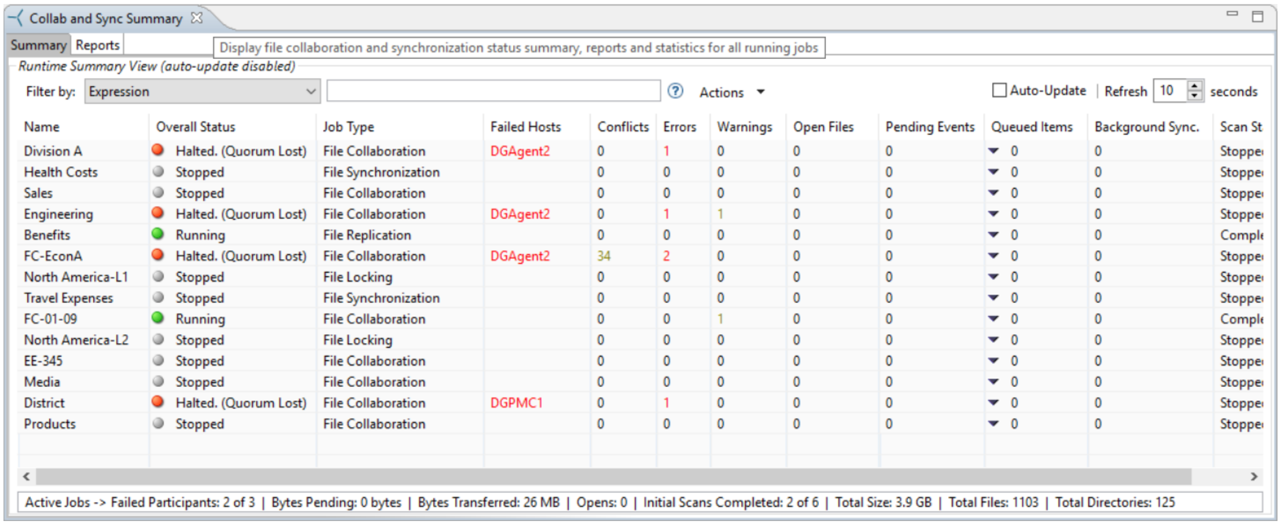

Display file collaboration and synchronization status summary, reports and statistics for all running jobs

Runtime Summary View (auto-update disabled)

Filter by: Expression

□ Auto-Update | Refresh 10 seconds

Name	Overall Status	Job Type	Failed Hosts	Conflicts	Errors	Warnings	Open Files	Pending Events	Queued Items	Background Sync.	Scan St
Division A	● Halted. (Quorum Lost)	File Collaboration	DGAgent2	0	1	0	0	0	▼ 0	0	Stoppe
Health Costs	● Stopped	File Synchronization		0	0	0	0	0	▼ 0	0	Stoppe
Sales	● Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stoppe
Engineering	● Halted. (Quorum Lost)	File Collaboration	DGAgent2	0	1	1	0	0	▼ 0	0	Stoppe
Benefits	● Running	File Replication		0	0	0	0	0	▼ 0	0	Comple
FC-EconA	● Halted. (Quorum Lost)	File Collaboration	DGAgent2	34	2	0	0	0	▼ 0	0	Stoppe
North America-L1	● Stopped	File Locking		0	0	0	0	0	▼ 0	0	Stoppe
Travel Expenses	● Stopped	File Synchronization		0	0	0	0	0	▼ 0	0	Stoppe
FC-01-09	● Running	File Collaboration		0	0	1	0	0	▼ 0	0	Comple
North America-L2	● Stopped	File Locking		0	0	0	0	0	▼ 0	0	Stoppe
EE-345	● Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stoppe
Media	● Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stoppe
District	● Halted. (Quorum Lost)	File Collaboration	DGPMC1	0	1	0	0	0	▼ 0	0	Stoppe
Products	● Stopped	File Collaboration		0	0	0	0	0	▼ 0	0	Stoppe

Active Jobs -> Failed Participants: 2 of 3 | Bytes Pending: 0 bytes | Bytes Transferred: 26 MB | Opens: 0 | Initial Scans Completed: 2 of 6 | Total Size: 3.9 GB | Total Files: 1103 | Total Directories: 125

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.

- **Queue 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.

Automatic Updates

For performance reasons, this view is not updated in real-time. However, the table 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.

Table Context Menu

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. Double-clicking any item in the table will automatically open the selected job in a tab within the a view, allowing you to drill down and view specific information about that single job.

The actions that are unique to this table are as follows:

Purge All Conflicts	Purges all file conflicts from the selected jobs. This can only be performed on jobs that are not running.
Clear Alerts	Clears all alerts for the selected jobs. This can be performed while a job is running.
Trash-Bin Cleanup	The automatic trash-bin cleanup process runs once daily at 11 PM. Select this option to execute the trash-bin cleanup process on demand.
Show Details	Choose this option to display all the statistics for the selected job in the **Runtime Summary Details** dialog.
Copy Details	Choose this option to copy detailed information to the system clipboard for the job(s) selected in the table. This information can then be pasted into a document editor.

Filtering Jobs in the Table

You can change which items are displayed in the table by filtering the list or by it 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 drop0down list.

Actions Menu

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

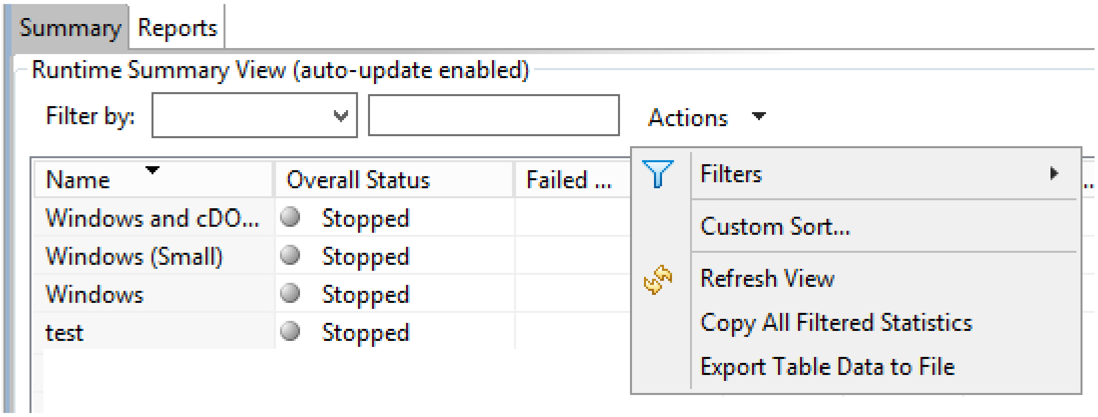
Summary
Reports
Runtime Summary View (auto-update enabled)
Filter by:
Actions
Filters
Custom Sort...
Refresh View
Copy All Filtered Statistics
Export Table Data to File

Name	Overall Status	Failed ...
Windows and cDO...	Stopped	
Windows (Small)	Stopped	
Windows	Stopped	
test	Stopped	

Filters	Allows for the selection of built-in or user-defined filters and to save/manage list filters. Default job filters include Failed Jobs, Jobs with Backlog, and Running Scans. For example, filter "Running Scans."
Custom Sort...	Use the Custom Sort option to configure and save how you want the Collaboration Summary view table to be sorted and keep 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	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 document editor.
Export Entire Table to File	Dump the entire contents of the table to a text file that can be viewed in any document editor.

The **Runtime Reports** tab aggregates critical statistical information from all File Collaboration, File Locking, File Replication, and File Synchronization jobs in a single table view. The **Reports** tab is visible when the global **Enable Advanced Reporting Tab** option in Preferences checked. This tab is especially useful to see the number of files that are in the queue waiting to be synchronized (File Sync Queue).

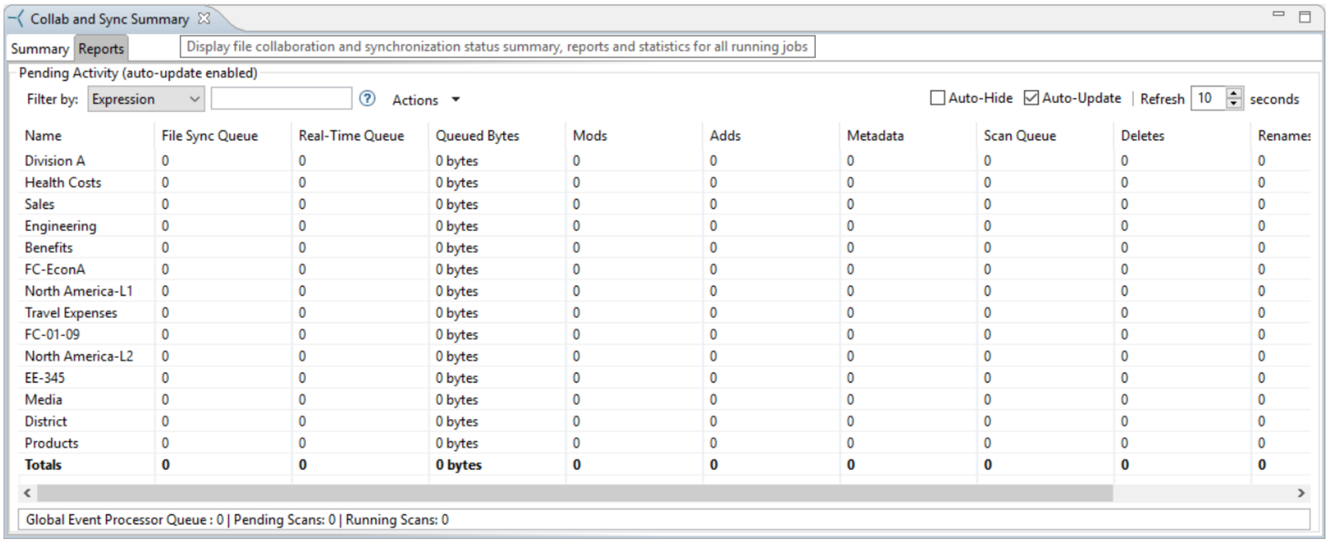

Collab and Sync Summary

Summary | Reports | Display file collaboration and synchronization status summary, reports and statistics for all running jobs

Pending Activity (auto-update enabled)

Filter by: Expression | Actions | □ Auto-Hide ☑ Auto-Update | Refresh 10 seconds

Name	File Sync Queue	Real-Time Queue	Queued Bytes	Mods	Adds	Metadata	Scan Queue	Deletes	Renames
Division A	0	0	0 bytes	0	0	0	0	0	0
Health Costs	0	0	0 bytes	0	0	0	0	0	0
Sales	0	0	0 bytes	0	0	0	0	0	0
Engineering	0	0	0 bytes	0	0	0	0	0	0
Benefits	0	0	0 bytes	0	0	0	0	0	0
FC-EconA	0	0	0 bytes	0	0	0	0	0	0
North America-L1	0	0	0 bytes	0	0	0	0	0	0
Travel Expenses	0	0	0 bytes	0	0	0	0	0	0
FC-01-09	0	0	0 bytes	0	0	0	0	0	0
North America-L2	0	0	0 bytes	0	0	0	0	0	0
EE-345	0	0	0 bytes	0	0	0	0	0	0
Media	0	0	0 bytes	0	0	0	0	0	0
District	0	0	0 bytes	0	0	0	0	0	0
Products	0	0	0 bytes	0	0	0	0	0	0
Totals	**0**	**0**	**0 bytes**	**0**	**0**	**0**	**0**	**0**	**0**

Global Event Processor Queue : 0 | Pending Scans: 0 | Running Scans: 0

Automatic Updates

For performance reasons, this view is not updated in real-time. However, the table 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.

Column Descriptions

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

Name	The name of the configured 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 global **Performance Real-Time Background Threads** option.
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 global **Performance Real-Time Expedited Threads** option.

Queued Bytes	The number of bytes that are in queue waiting to be processed.
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.

Items in the table can be filtered by a filter expression, job name, Participant, Session Status, or 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.

Actions Menu

Clicking on the **Actions** table menu provides the following options:

Filters	Allows for the selection of built-in or user-defined filters and to save/manage filter expressions. Default job filters include Failed Jobs, Jobs with Backlog, and Running Scans. For example, filter:"Running Scans".
Custom Sort...	Use the Custom Sort option to configure and save how you want the Runtime Reports table to be sorted and keep 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	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 document editor.
Export Entire Table to File	Dump the entire contents of the table to a text file that can be viewed in any document 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.

The **Dashboard Summary View** displays metrics and key performance indicators from all running File Collaboration, File Locking, File Replication, and File Synchronization jobs. It is automatically displayed when the Peer Management Center client is started and can be opened at any other time by selecting **View Dashboard** from the **Window** menu or by clicking the **View Dashboard** icon in the Peer Management Center toolbar.

Entries in the first column of the **File Collaboration** and **Agents** categories can be double-clicked, which will take the user to a filtered runtime view of the selected item for additional details.

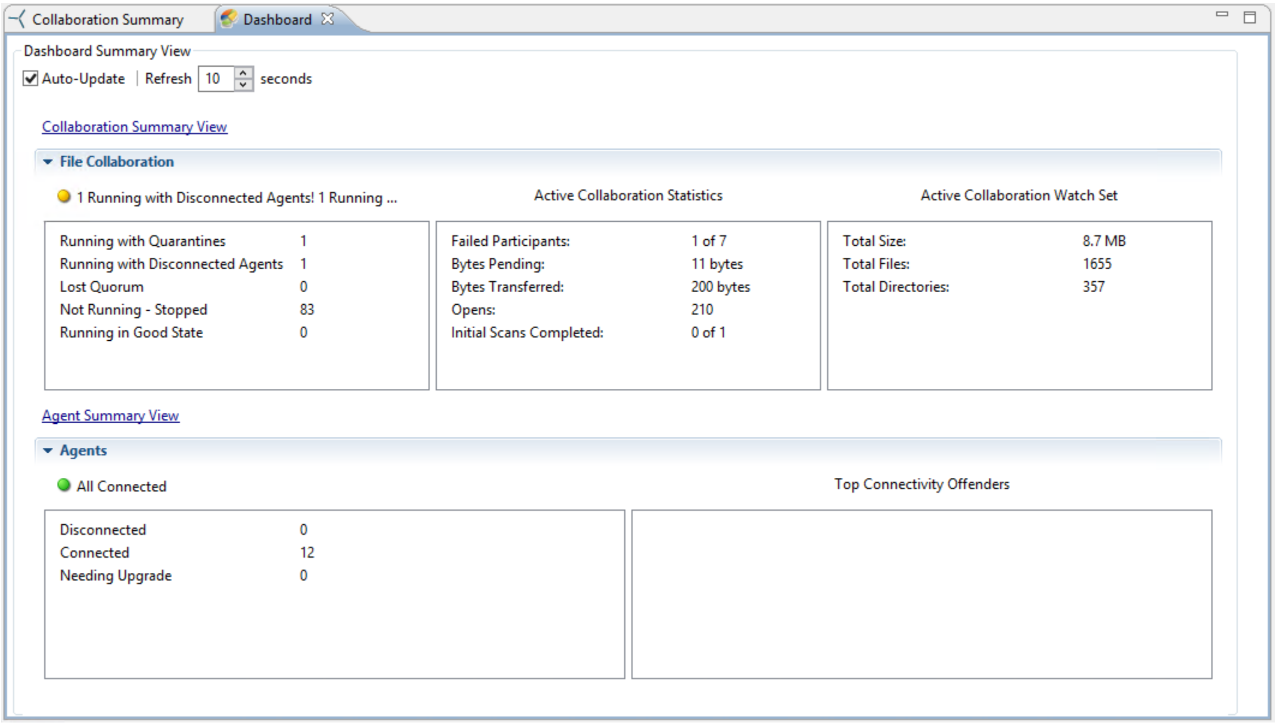

Collaboration Summary
Dashboard
Dashboard Summary View
☑ Auto-Update | Refresh 10 seconds
Collaboration Summary View
File Collaboration
1 Running with Disconnected Agents! 1 Running ...
Active Collaboration Statistics
Active Collaboration Watch Set

Running with Quarantines	1
Running with Disconnected Agents	1
Lost Quorum	0
Not Running - Stopped	83
Running in Good State	0

Failed Participants:	1 of 7
Bytes Pending:	11 bytes
Bytes Transferred:	200 bytes
Opens:	210
Initial Scans Completed:	0 of 1

Total Size:	8.7 MB
Total Files:	1655
Total Directories:	357

Agent Summary View
Agents
All Connected
Top Connectivity Offenders

Disconnected	0
Connected	12
Needing Upgrade	0

Automatic Updates

For performance reasons, this view is not updated in real-time. However, the table 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.

File Collaboration Runtime Job Views

Each File Collaboration job has seven primary runtime views used for viewing a combination of real-time file I/O activity, history, and configuration. The File Collaboration Runtime View is located in the large center section of the Peer Management Center. It is composed of various tabs (or editors) representing individual File Collaboration jobs and/or cross-job summary information:

- Summary (or Status) view - Shows overall statistics for the selected job.
- Session view - Shows active open files and files that are currently in transit between participating hosts.
- Event Log view - Shows a list of all runtime activity that has occurred within the selected job.
- File Conflicts view - Shows a list of all files that are quarantined for the session or are in conflict between two or more participating hosts.
- Alerts view - Shows a list of all job alerts specifically tied to the selected job.
- Participants view - Shows a list of all hosts participating in the selected job.
- Configuration view - Shows a summary of all configurable options for the selected job.

These views also provide the ability to manage specific collaboration runtime functionality.

Runtime View Sub Tabs	Displays runtime statistics.
Job Start/Stop	The button allows you to start and stop the File Collaboration job.
Job Status Display	Displays status-related messages when the job is running.

The **Summary** runtime view allows you to view current and cumulative file collaboration and synchronization statistics, as well background synchronization status.

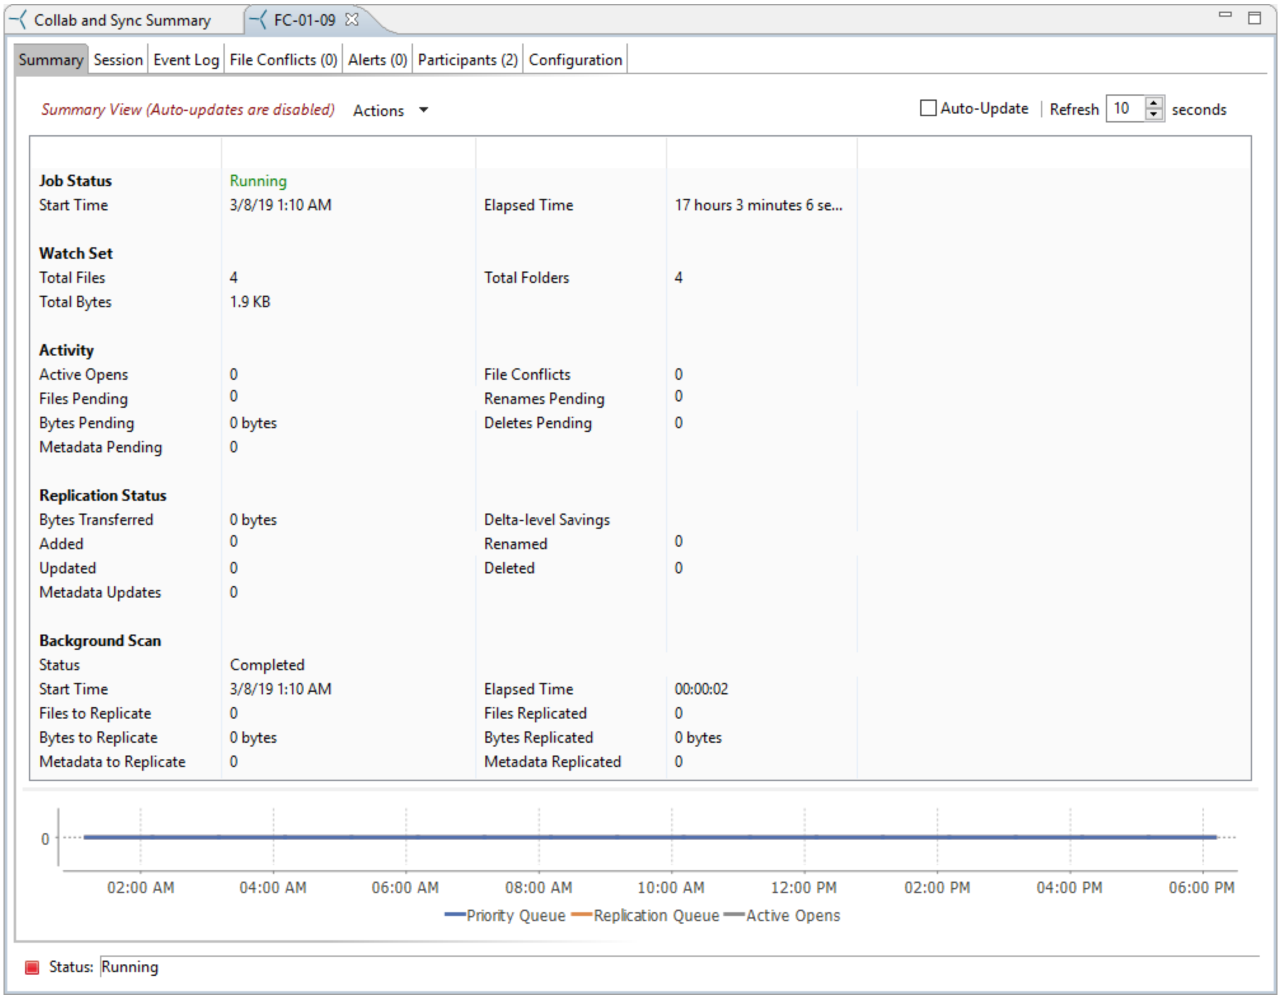

Summary View (Auto-updates are disabled)

☐ Auto-Update | Refresh 10 seconds

Job Status	Running		
Start Time	3/8/19 1:10 AM	Elapsed Time	17 hours 3 minutes 6 se...
Watch Set			
Total Files	4	Total Folders	4
Total Bytes	1.9 KB		
Activity			
Active Opens	0	File Conflicts	0
Files Pending	0	Renames Pending	0
Bytes Pending	0 bytes	Deletes Pending	0
Metadata Pending	0		
Replication Status			
Bytes Transferred	0 bytes	Delta-level Savings	
Added	0	Renamed	0
Updated	0	Deleted	0
Metadata Updates	0		
Background Scan			
Status	Completed		
Start Time	3/8/19 1:10 AM	Elapsed Time	00:00:02
Files to Replicate	0	Files Replicated	0
Bytes to Replicate	0 bytes	Bytes Replicated	0 bytes
Metadata to Replicate	0	Metadata Replicated	0

Status: Running

Key statistics in this view are presented in the Activity, Replication Status, and Background Scan sections.

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 and Sync 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 off iles 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.

The **Session** view 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.

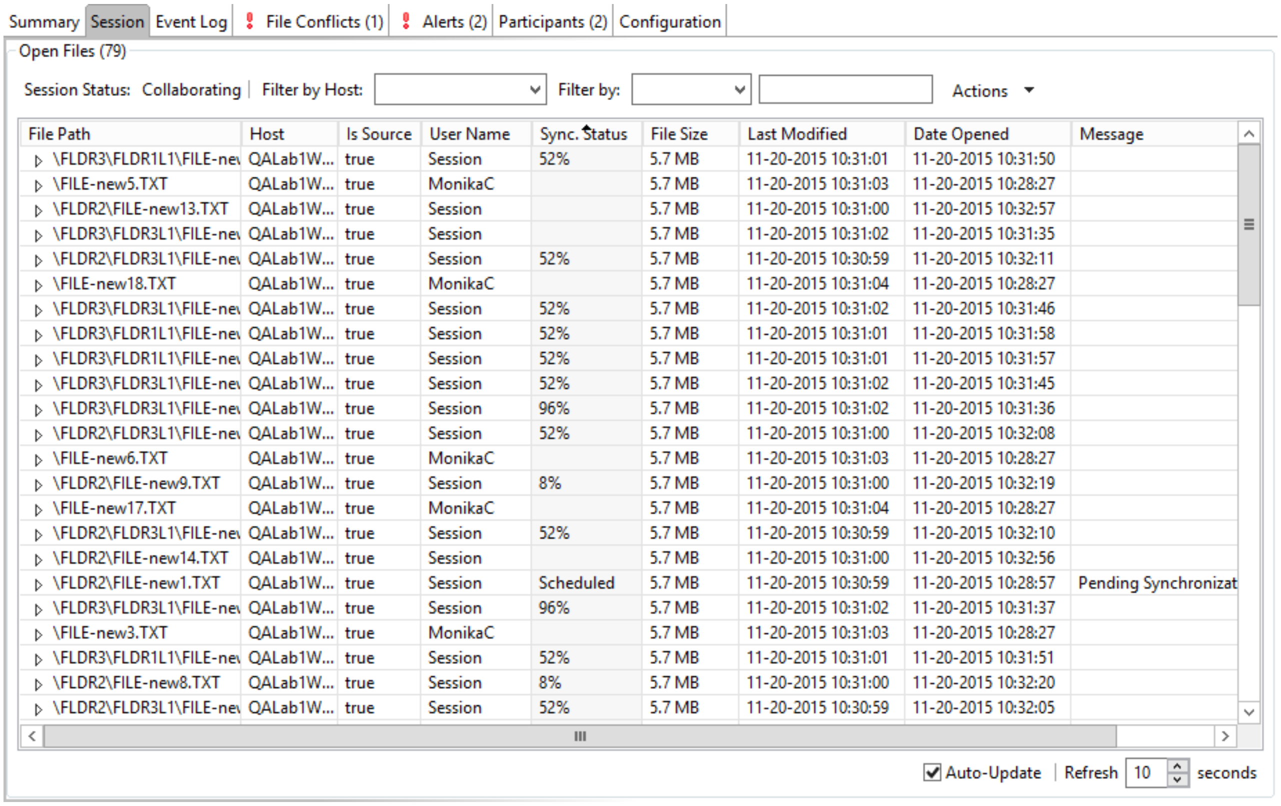

Session Status: Collaborating

File Path	Host	Is Source	User Name	Sync. Status	File Size	Last Modified	Date Opened	Message
\FLDR3\FLDR1L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:01	11-20-2015 10:31:50	
\FILE-new5.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:03	11-20-2015 10:28:27	
\FLDR2\FILE-new13.TXT	QALab1W...	true	Session		5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:57	
\FLDR3\FLDR3L1\FILE-nev	QALab1W...	true	Session		5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:35	
\FLDR2\FLDR3L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:30:59	11-20-2015 10:32:11	
\FILE-new18.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:04	11-20-2015 10:28:27	
\FLDR3\FLDR3L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:46	
\FLDR3\FLDR1L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:01	11-20-2015 10:31:58	
\FLDR3\FLDR1L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:01	11-20-2015 10:31:57	
\FLDR3\FLDR3L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:45	
\FLDR3\FLDR3L1\FILE-nev	QALab1W...	true	Session	96%	5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:36	
\FLDR2\FLDR3L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:08	
\FILE-new6.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:03	11-20-2015 10:28:27	
\FLDR2\FILE-new9.TXT	QALab1W...	true	Session	8%	5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:19	
\FILE-new17.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:04	11-20-2015 10:28:27	
\FLDR2\FLDR3L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:30:59	11-20-2015 10:32:10	
\FLDR2\FILE-new14.TXT	QALab1W...	true	Session		5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:56	
\FLDR2\FILE-new1.TXT	QALab1W...	true	Session	Scheduled	5.7 MB	11-20-2015 10:30:59	11-20-2015 10:28:57	Pending Synchronizat
\FLDR3\FLDR3L1\FILE-nev	QALab1W...	true	Session	96%	5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:37	
\FILE-new3.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:03	11-20-2015 10:28:27	
\FLDR3\FLDR1L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:01	11-20-2015 10:31:51	
\FLDR2\FILE-new8.TXT	QALab1W...	true	Session	8%	5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:20	
\FLDR2\FLDR3L1\FILE-nev	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:30:59	11-20-2015 10:32:05	

☑ Auto-Update | Refresh 10 seconds

The **Session** view is made up of the following components:

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 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	A drop-down list of participating hosts to filter on. Selecting a specific host will filter the Open Files to just show files on that host.
Filter By Combo	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.

The **Event Log** view 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.

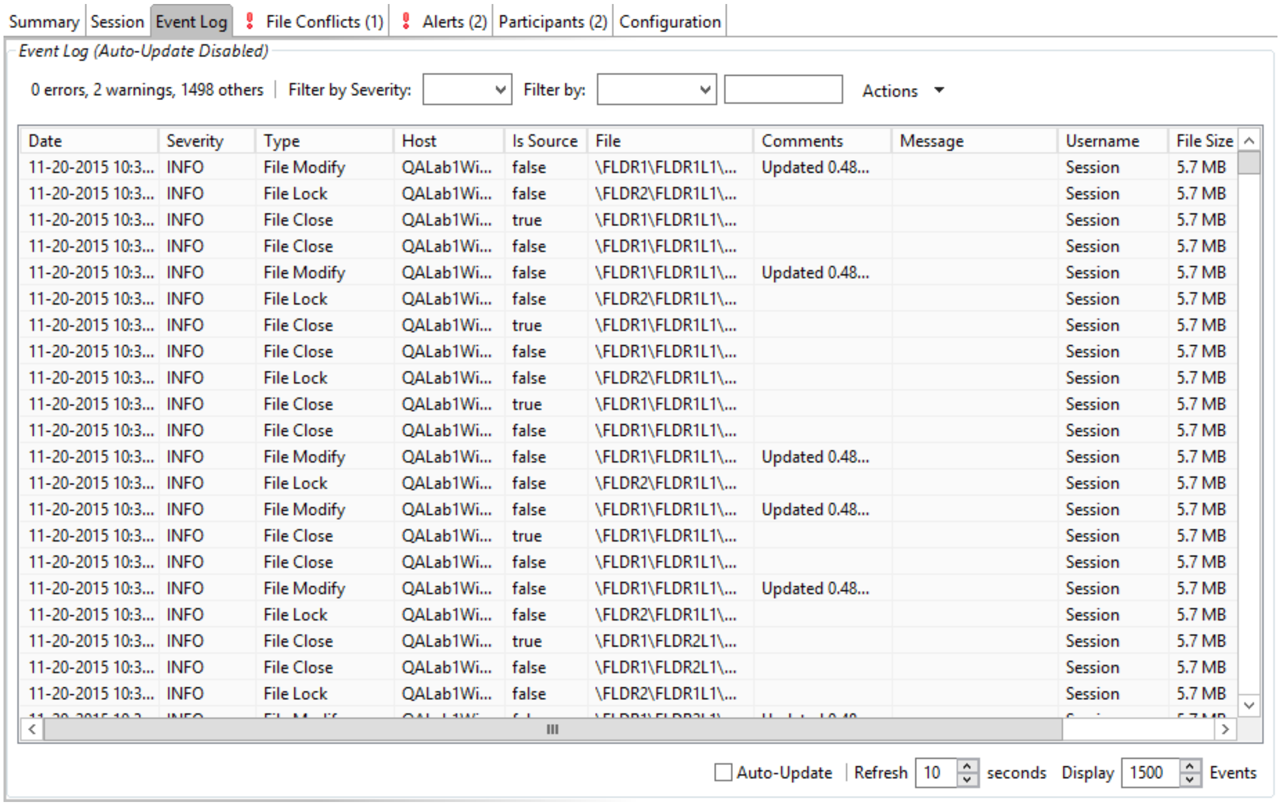

0 errors, 2 warnings, 1498 others

☐ Auto-Update | Refresh 10 seconds Display 1500 Events

Date	Severity	Type	Host	Is Source	File	Comments	Message	Username	File Size
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\FLDR1\FLDR1L1\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\FLDR2\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\FLDR1\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\FLDR1\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\FLDR1\FLDR1L1\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\FLDR2\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\FLDR1\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\FLDR1\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\FLDR2\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\FLDR1\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\FLDR1\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\FLDR1\FLDR1L1\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\FLDR2\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\FLDR1\FLDR1L1\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\FLDR1\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\FLDR1\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\FLDR1\FLDR1L1\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\FLDR2\FLDR1L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\FLDR1\FLDR2L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\FLDR1\FLDR2L1\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\FLDR2\FLDR1L1\...			Session	5.7 MB

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.

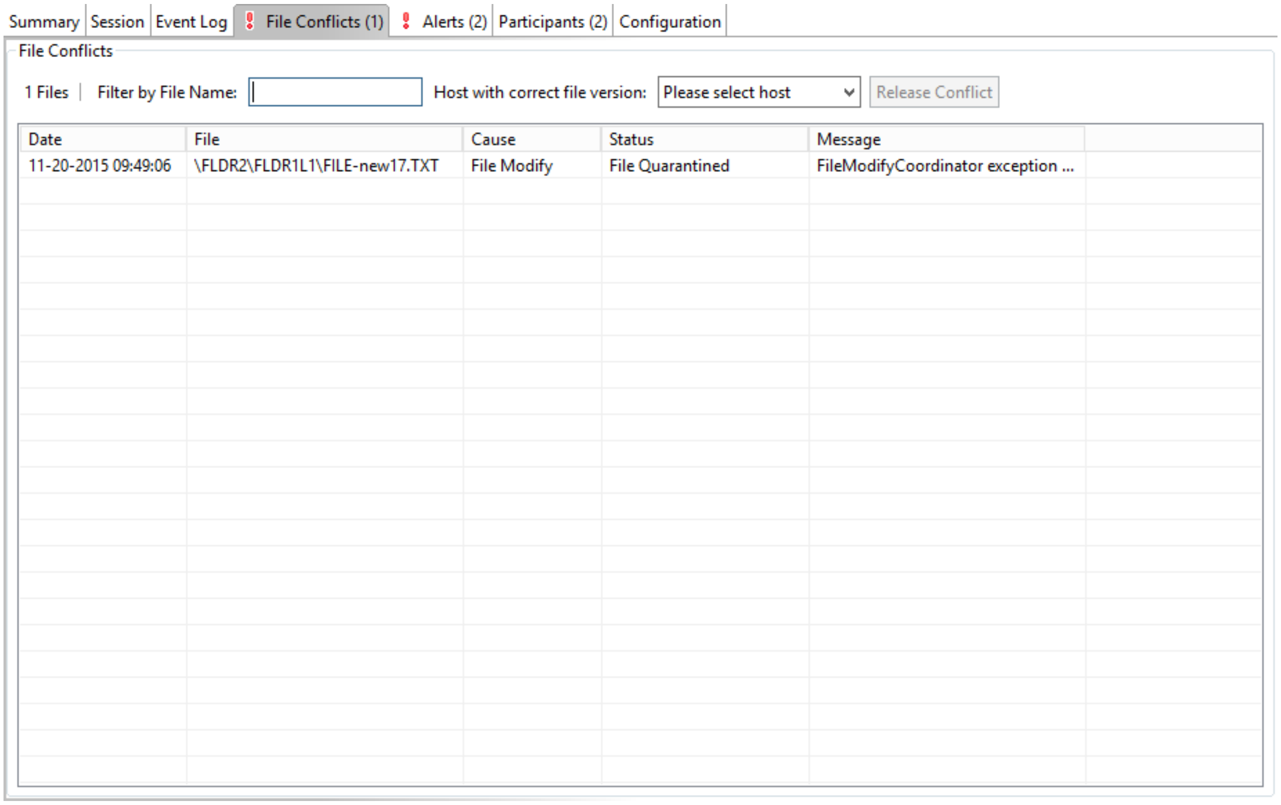

Date	File	Cause	Status	Message
11-20-2015 09:49:06	\FLDR2\FLDR1L1\FILE-new17.TXT	File Modify	File Quarantined	FileModifyCoordinator exception ...

The right-click context menu for the table contains the following actions that are unique to this particular view:

Refresh View	Refresh all information provided in the table.
Clear Alerts	Clears all alerts for the selected job. This can be performed while a job is running.

The **Alerts** view 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 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, 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.

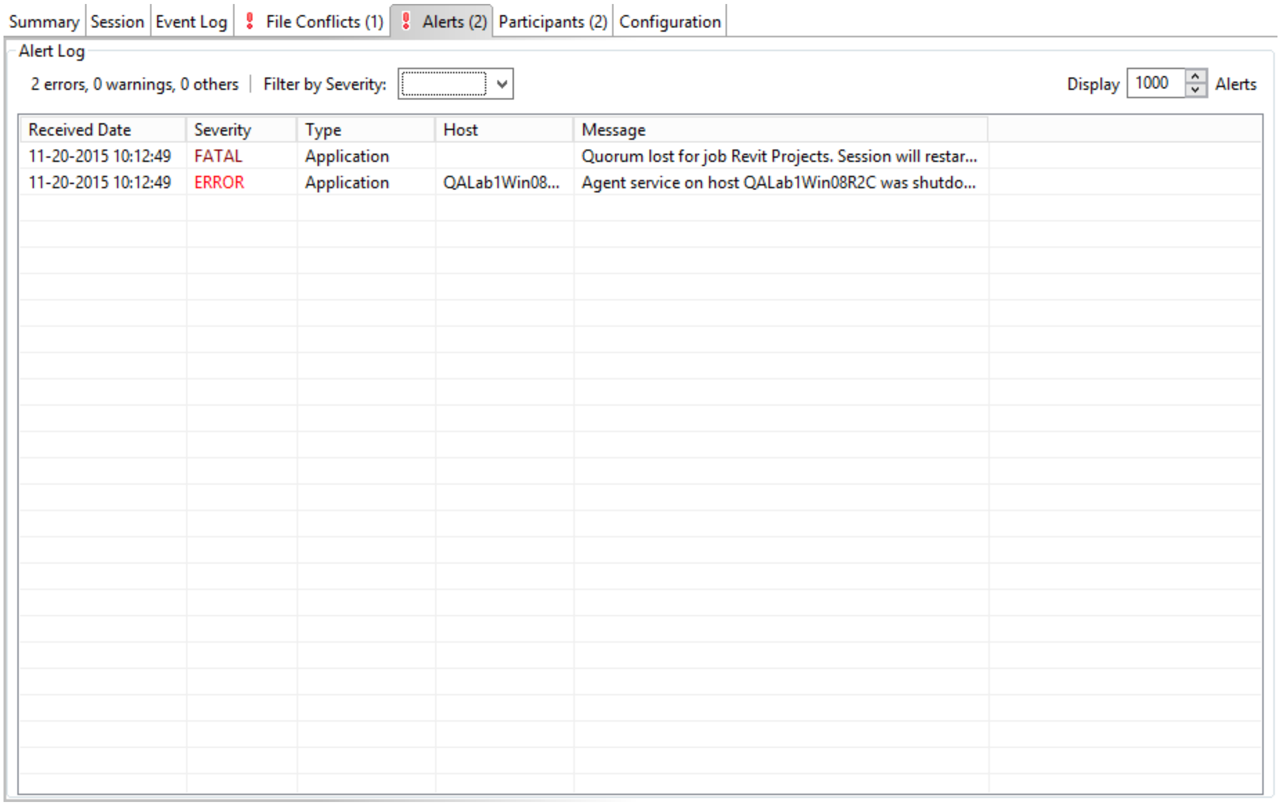

2 errors, 0 warnings, 0 others

Display 1000 Alerts

Received Date	Severity	Type	Host	Message
11-20-2015 10:12:49	FATAL	Application		Quorum lost for job Revit Projects. Session will restar...
11-20-2015 10:12:49	ERROR	Application	QALab1Win08...	Agent service on host QALab1Win08R2C was shutdo...

The following right-click menu items are unique to this particular 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 is divided into two sections:

- Host Participants

- [Host Participant State Change Log]

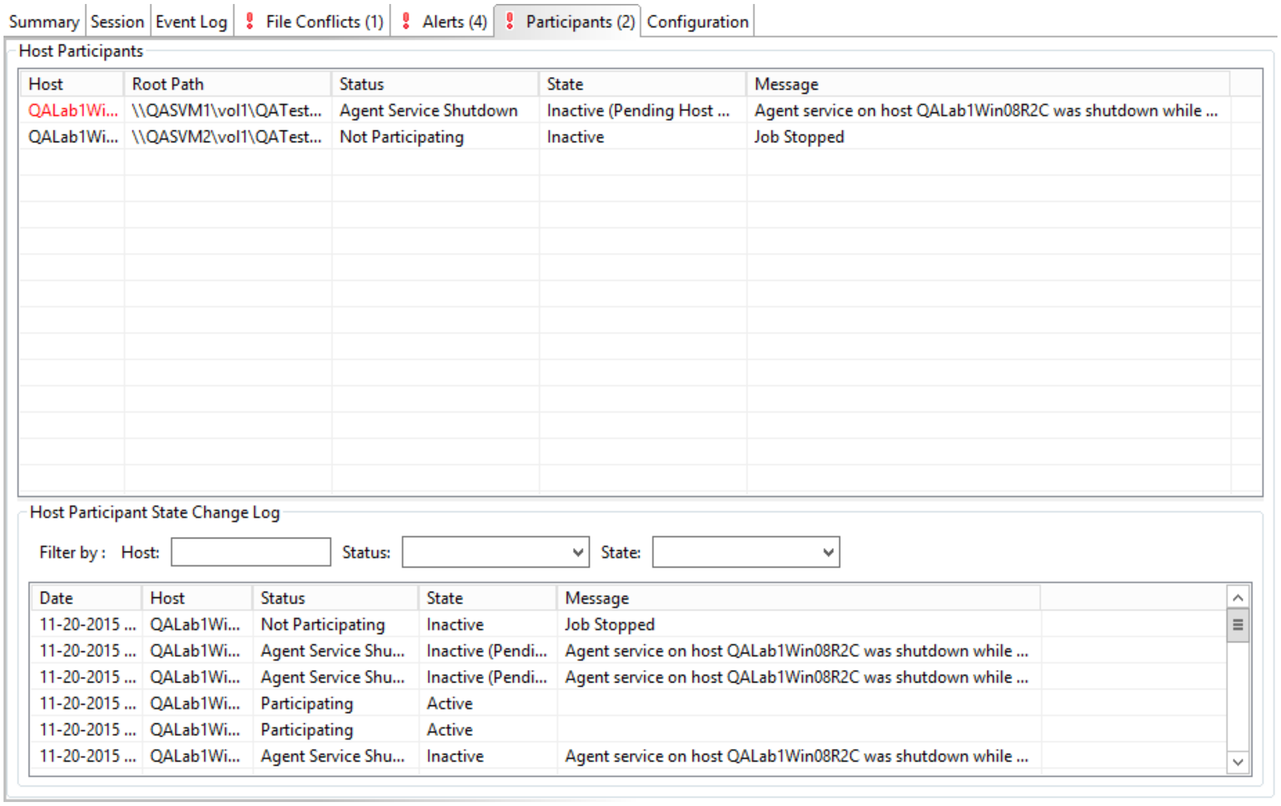

Host Participants

Host	Root Path	Status	State	Message
QALab1Wi...	\\QASVM1\vol1\QATest...	Agent Service Shutdown	Inactive (Pending Host ...	Agent service on host QALab1Win08R2C was shutdown while ...
QALab1Wi...	\\QASVM2\vol1\QATest...	Not Participating	Inactive	Job Stopped

Host Participant State Change Log

Filter by : Host: Status: State:

Date	Host	Status	State	Message
11-20-2015 ...	QALab1Wi...	Not Participating	Inactive	Job Stopped
11-20-2015 ...	QALab1Wi...	Agent Service Shu...	Inactive (Pendi...	Agent service on host QALab1Win08R2C was shutdown while ...
11-20-2015 ...	QALab1Wi...	Agent Service Shu...	Inactive (Pendi...	Agent service on host QALab1Win08R2C was shutdown while ...
11-20-2015 ...	QALab1Wi...	Participating	Active	
11-20-2015 ...	QALab1Wi...	Participating	Active	
11-20-2015 ...	QALab1Wi...	Agent Service Shu...	Inactive	Agent service on host QALab1Win08R2C was shutdown while ...

Host Participants

The **Host Participants** section contains a table that displays all the currently configured [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.

The **Configuration** view 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** will immediately bring you to the File Collaboration Configuration edit wizard where you can edit the current configuration.

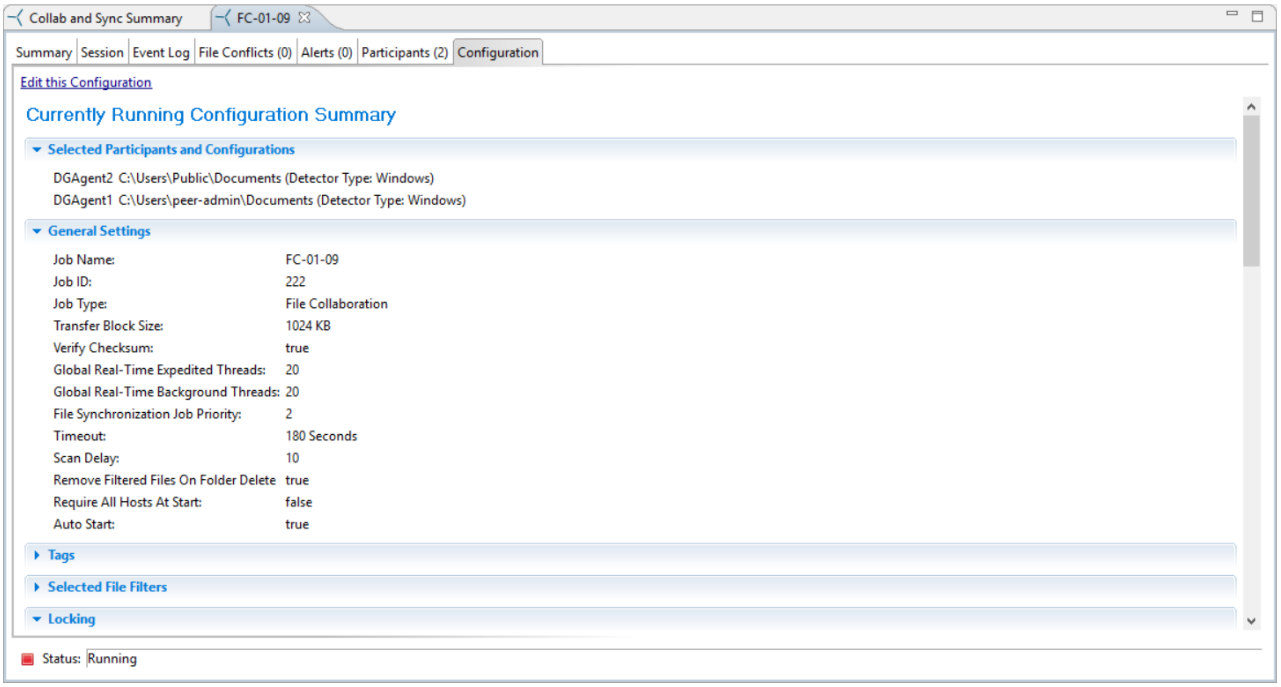

Currently Running Configuration Summary

Selected Participants and Configurations

DGAgent2 C:\Users\Public\Documents (Detector Type: Windows)
DGAgent1 C:\Users\peer-admin\Documents (Detector Type: Windows)

General Settings

Job Name:	FC-01-09
Job ID:	222
Job Type:	File Collaboration
Transfer Block Size:	1024 KB
Verify Checksum:	true
Global Real-Time Expedited Threads:	20
Global Real-Time Background Threads:	20
File Synchronization Job Priority:	2
Timeout:	180 Seconds
Scan Delay:	10
Remove Filtered Files On Folder Delete	true
Require All Hosts At Start:	false
Auto Start:	true

Tags

Selected File Filters

Locking

Status: Running

File Locking Jobs

This section provides information about creating a File Locking job:

- Overview
- Before You Create Your First File Locking Job
- Creating a File Locking Job

Overview

A File Locking job is designed to prevent multiple users from simultaneously accessing the same file across multiple file servers. You can combine it with Microsoft DFS Replication to provide an entry-level collaboration option for teams sharing project files that are distributed across multiple locations. Since it is only available with Microsoft DFS-R, File Locking jobs are only available for use with Windows file servers.

Before You Create Your First File Locking Job

We strongly recommend that you configure global settings such as SMTP configuration and email alerts before configuring your first File Locking job. See Preferences for details on what and how to configure these settings.

Creating a File Locking Job

The **Create Job Wizard** walks you through the process of creating a File Locking job:

Step 1: Job Type and Name

Step 2: Participants

Step 3: Email Alerts

Step 4: Save Job

Step 1: Job Type and Name

1. Open the Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **New Job** wizard displays a list of job types you can create.
3. Click **File Locking**, and then click **Create**.

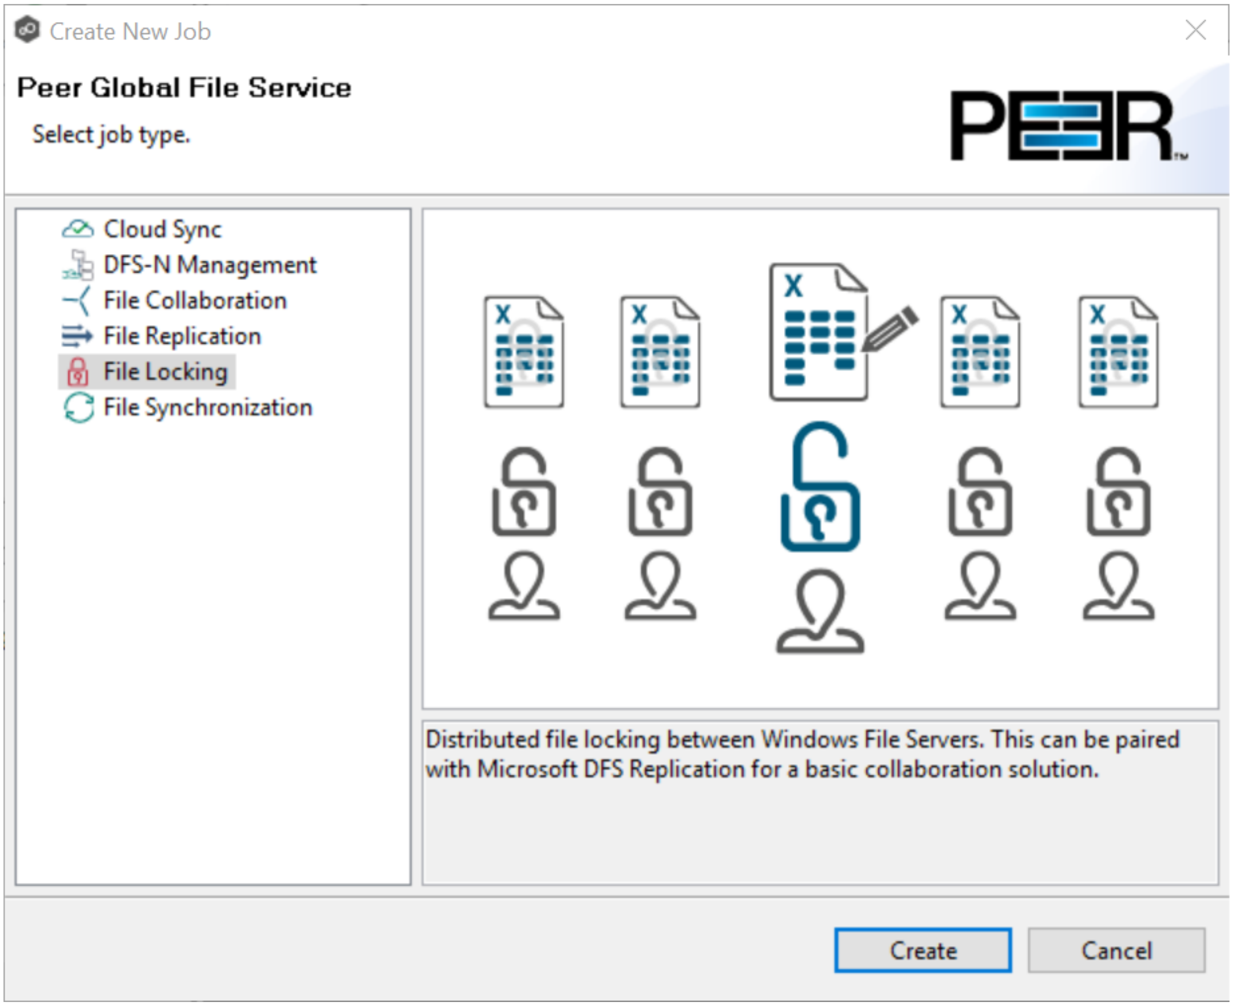

Distributed file locking between Windows File Servers. This can be paired with Microsoft DFS Replication for a basic collaboration solution.

4. Enter a name for the job in the dialog that appears.

   The job name must be unique.

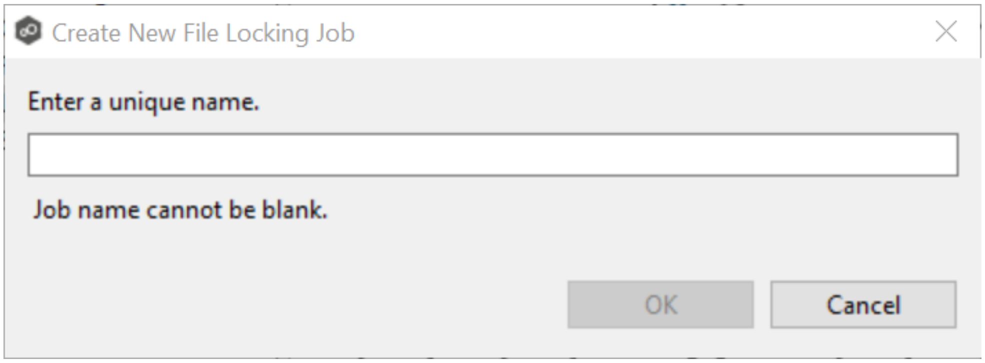

Job name cannot be blank.

5. Click **OK**.

   The Participants page is displayed.

Step 2: Participants

A File Locking job must have two or more participants. A participant consists of an Agent and the volume/share/folder to be protected from version conflicts. The server that the Agent is installed upon is called the host (or host participant). A File Locking job prevents other users from modifying the files of the participants while the files are locked.

1. Complete the four substeps:

   Participants

   Storage Platform

   Management Agent

   Path

   After you add a participant, it appears in the **Participants** table.

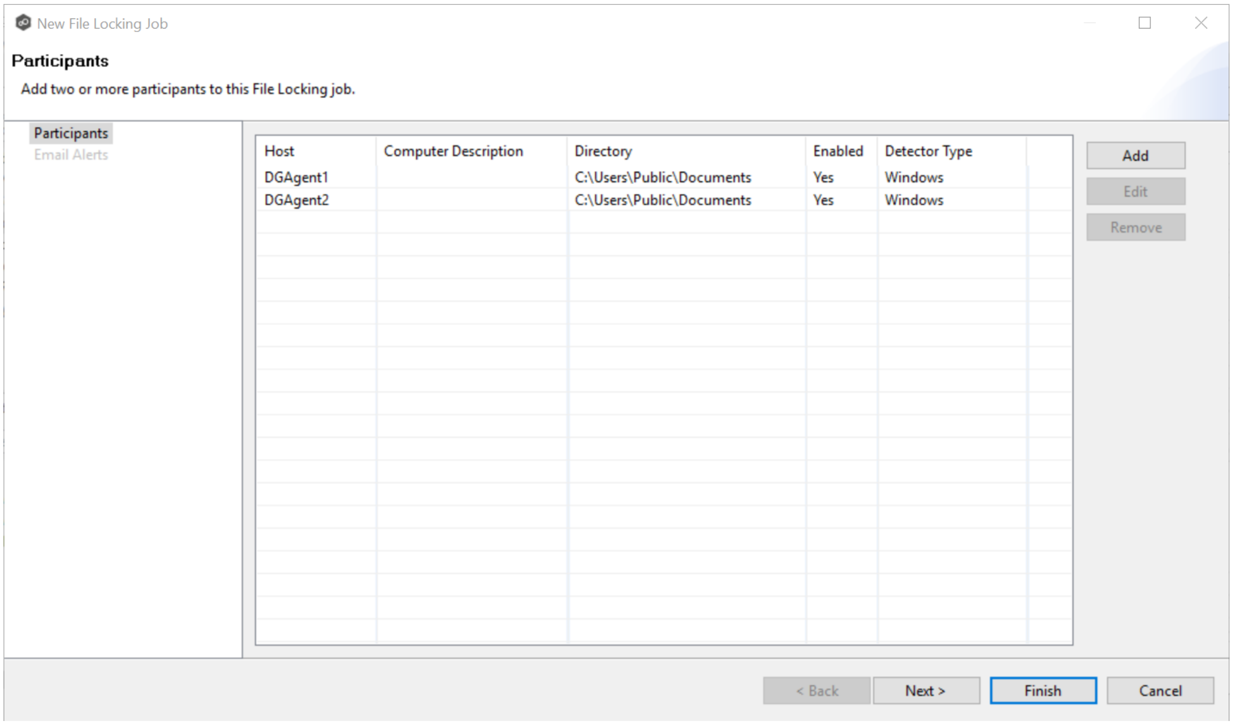

New File Locking Job

Participants

Add two or more participants to this File Locking job.

Host	Computer Description	Directory	Enabled	Detector Type
DGAgent1		C:\Users\Public\Documents	Yes	Windows
DGAgent2		C:\Users\Public\Documents	Yes	Windows

2. Repeat the five substeps for each participant you want to add to the job.

3. Once you have added all the participants, click **Next** to specify email alerts for the job. (Don't click **Finish**.)

The **Participants** page is where you select and configure which hosts will be participating in this job. The **Participants** page is empty until you finish the process of adding your first participant. Once you have added the participants, they are listed on the **Participants** page.

To begin the process of adding a participant:

1. Click the **Add** button.

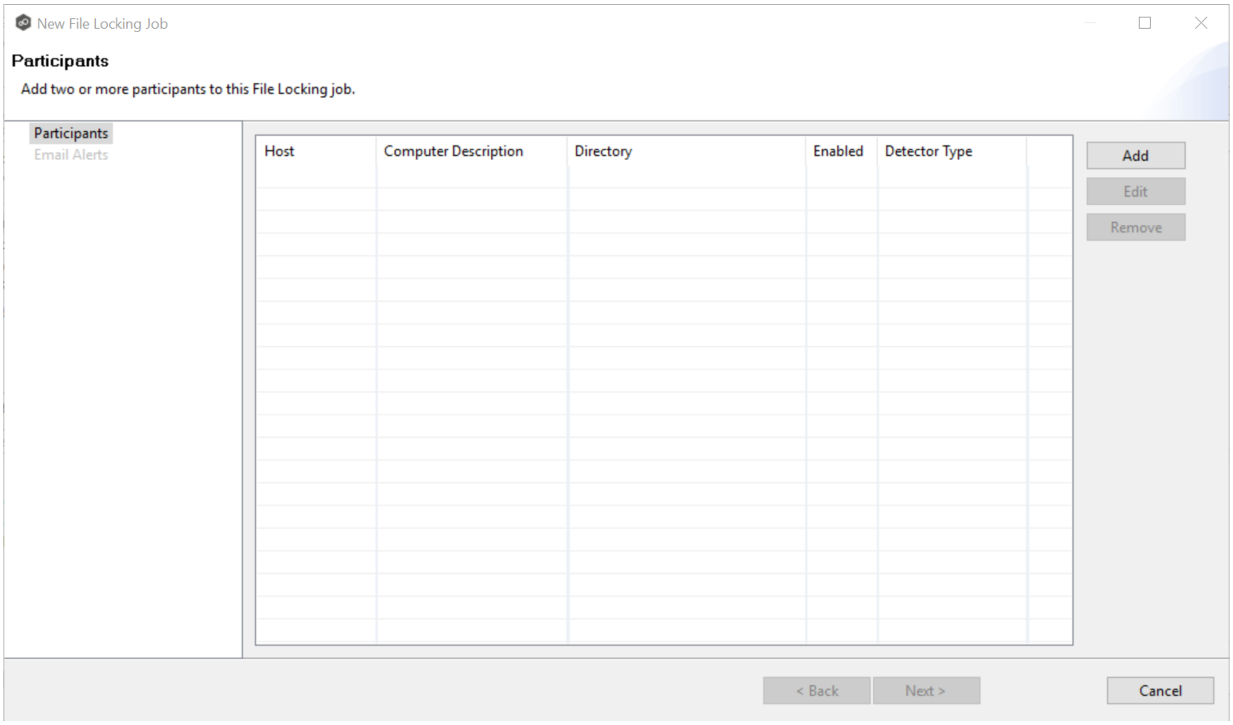

Another wizard opens to guide you through the process of adding a participant to the job. The first step in the process involves selecting the storage platform.

The **Storage Platform** page lists the types of storage platforms that File Collaboration supports. A storage device hosts data you want to replicate. It is often referred to as the host or host participant.

1. Select the type of storage platform that hosts the data you want to replicate.

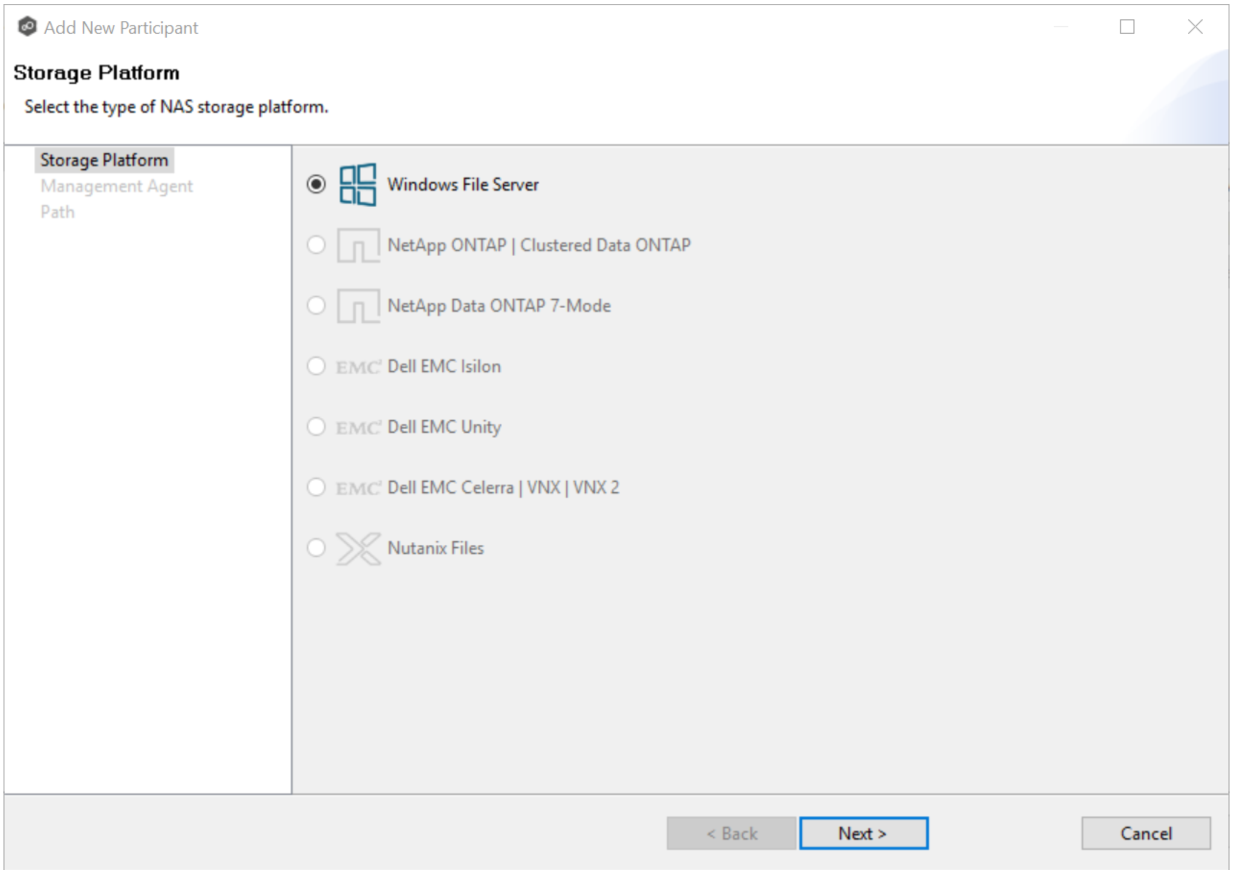

2. Click **Next**.

   The Management Agent page is displayed.

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 Locking job, you should select the Management Agent that manages the volumes/shares/folders you want to lock in this job.

1. Select the Agent that manages the host.

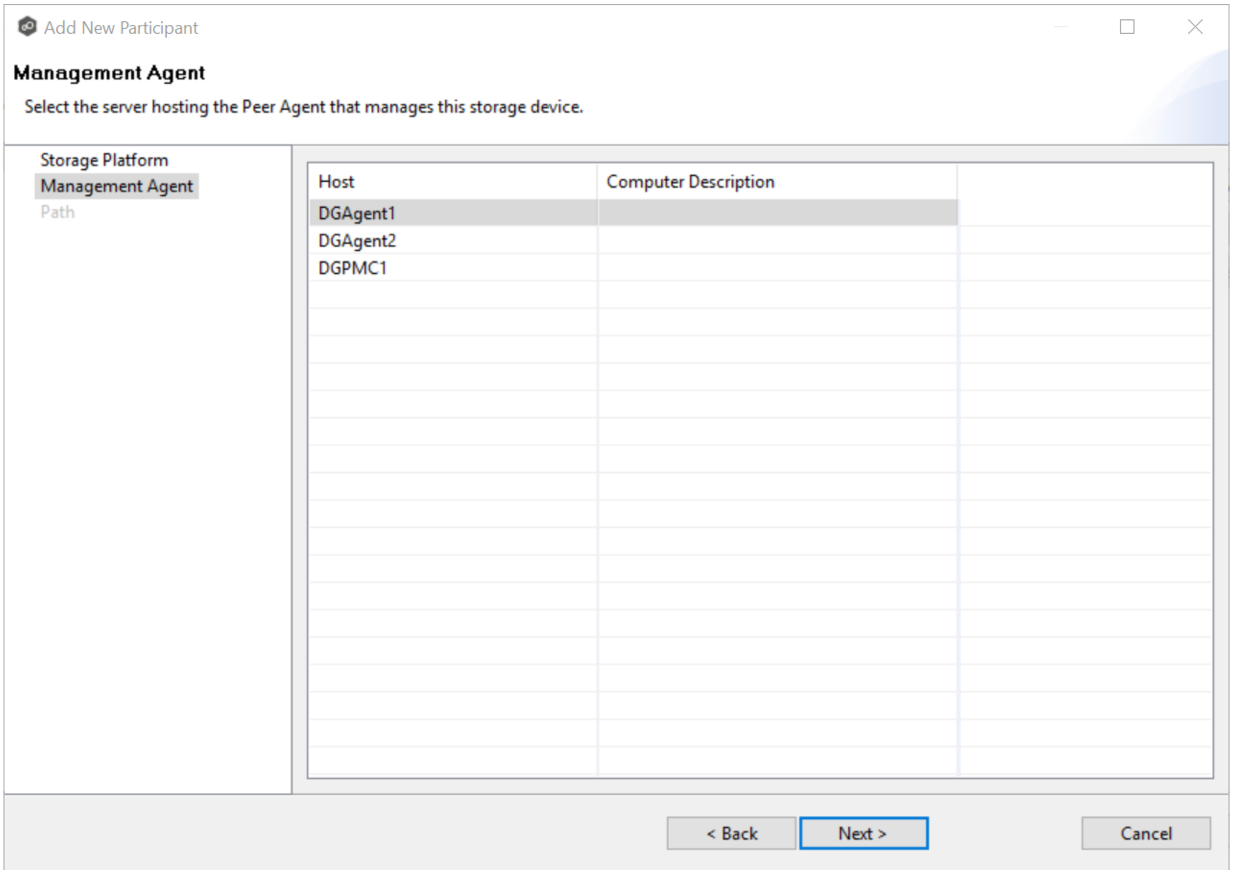

Add New Participant

Management Agent

Select the server hosting the Peer Agent that manages this storage device.

Host	Computer Description
DGAgent1	
DGAgent2	
DGPMC1	

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 Path page is displayed.

The **Path** page is where you specify the path to the volume/share/folder you want to lock. 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 lock 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.

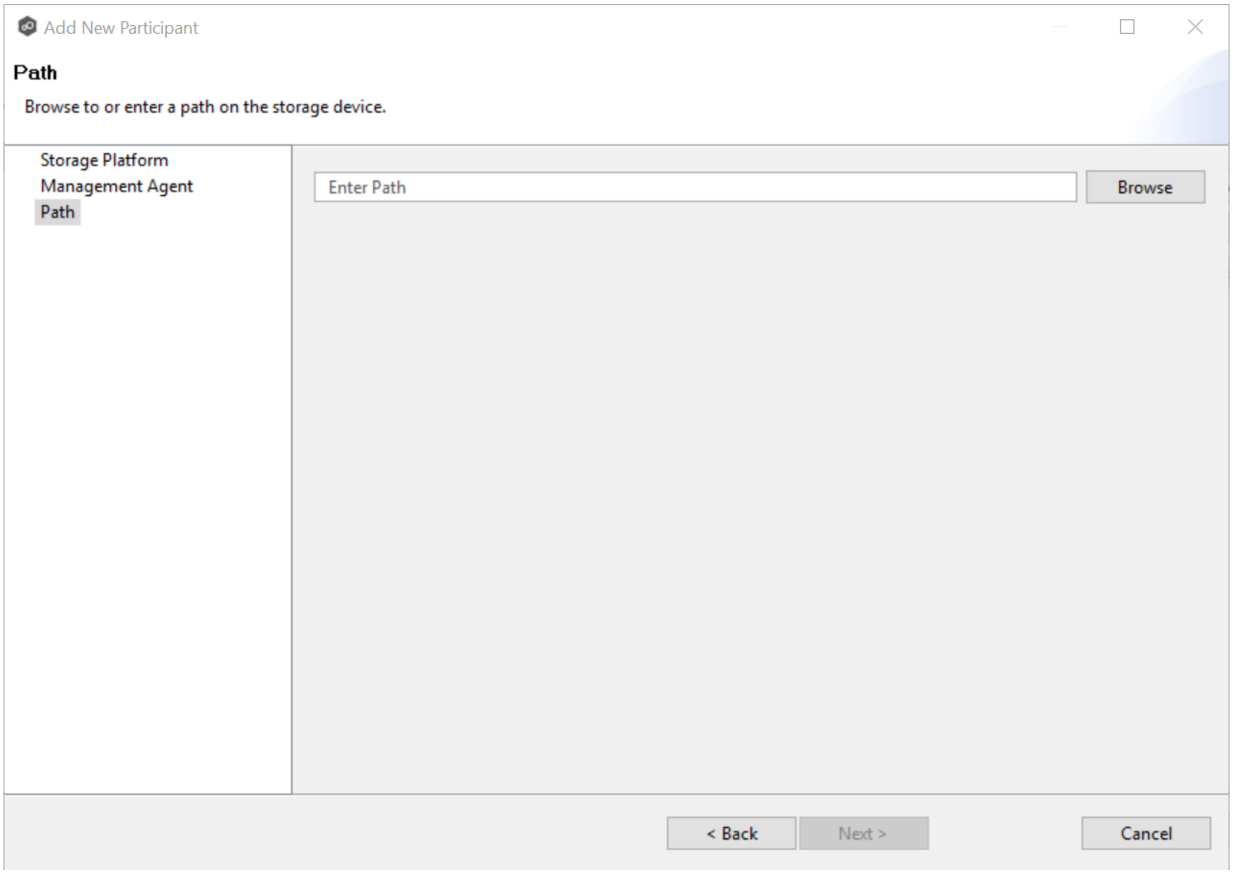

If you selected **Browse**, the **Folder Browser** dialog appears:

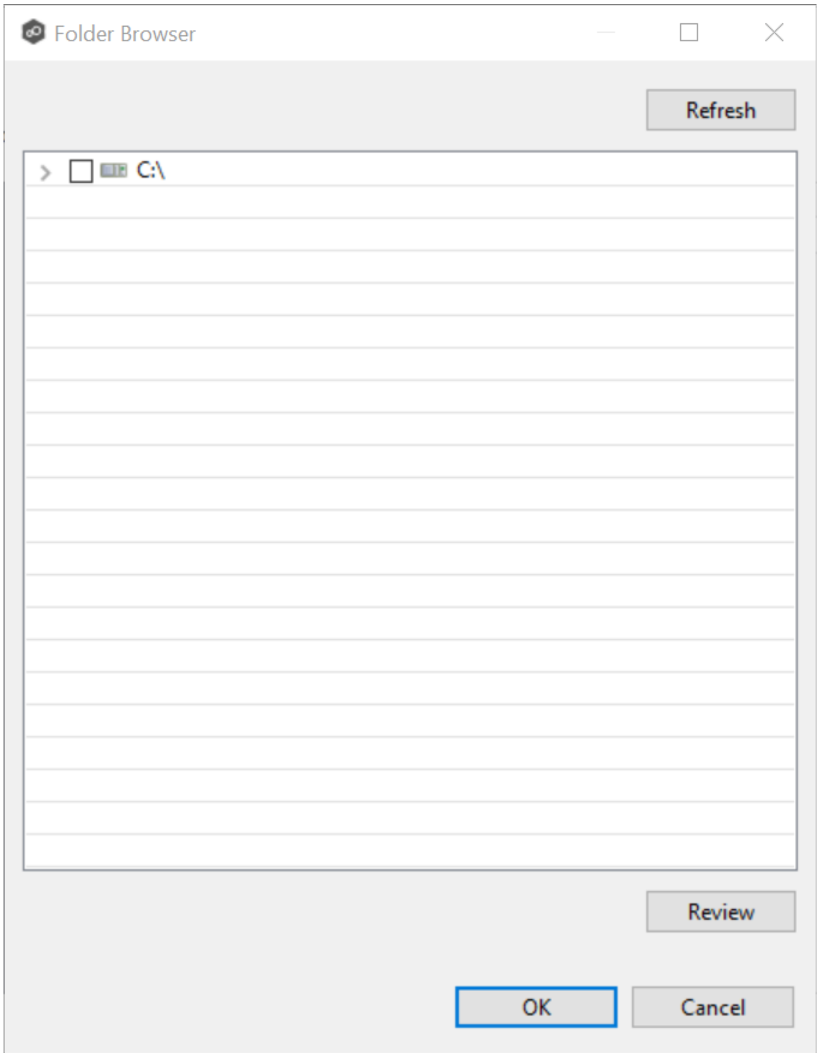

a. Expand the folder tree.

b. Select the appropriate volume/share/folder/.

c. (Optional) Click the **Review** button to see your selection.

d. Click **OK**.

2. Click **Finish** to complete the wizard for this participant.

3. Return to Step 2: Participants to add more participants, if applicable. A File Locking job must have at least two participants. If you have added all of the participants, continue with Step 3: Email Alerts.

Step 3: 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.

Apply an Existing Email Alert

To apply an existing email alert to the job.

1. Click the **Select** button.

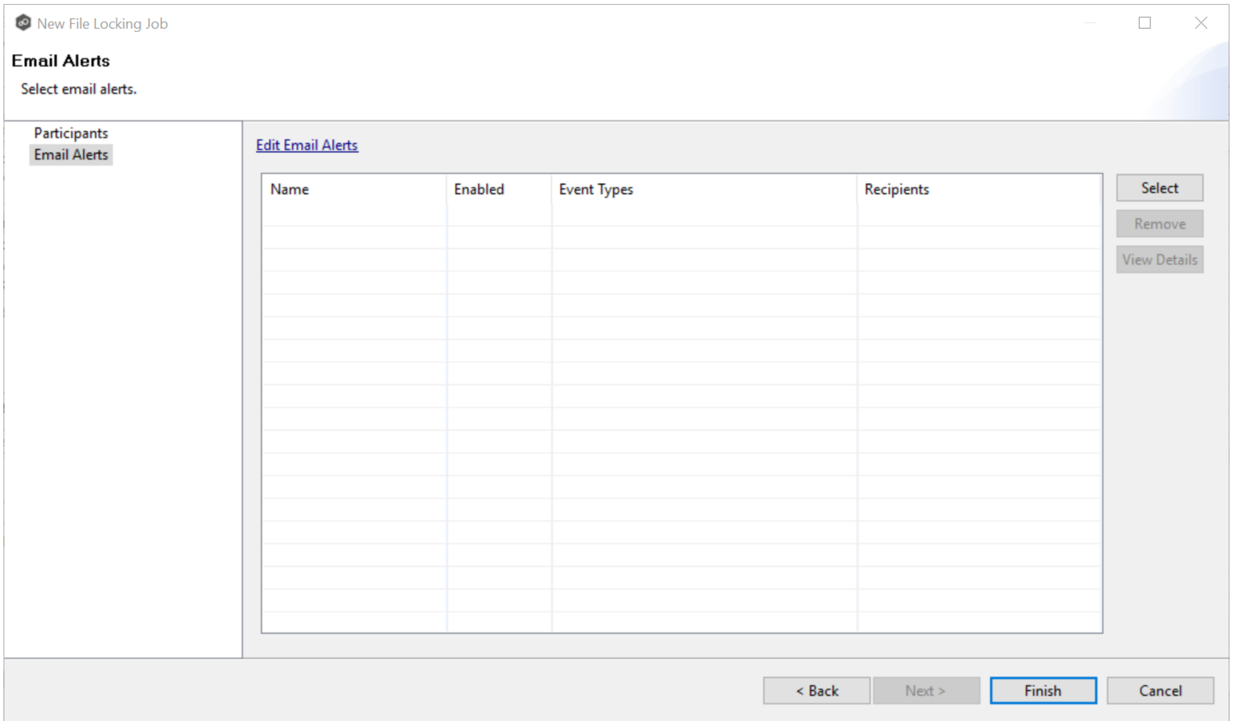

New File Locking Job

Email Alerts

Select email alerts.

Name	Enabled	Event Types	Recipients

The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.

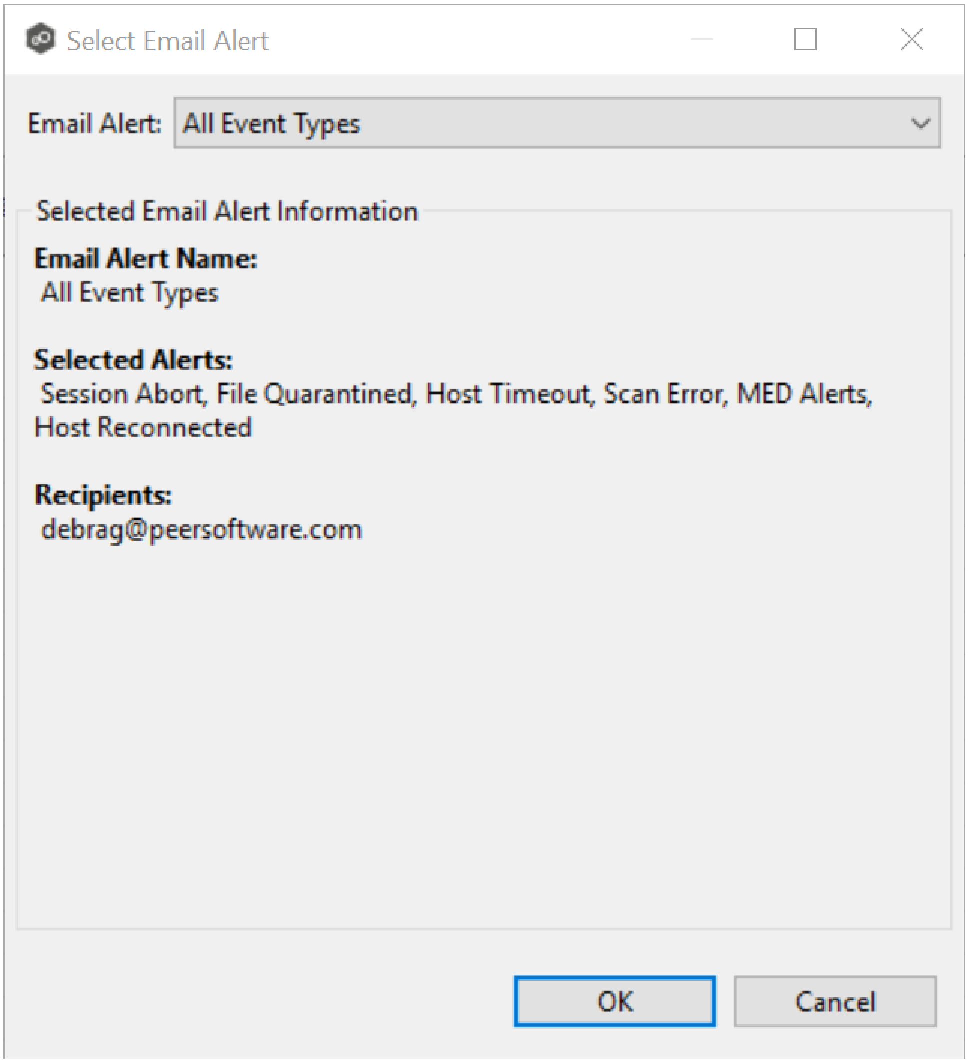

Email Alert: All Event Types

Selected Email Alert Information

Email Alert Name:
All Event Types

Selected Alerts:
Session Abort, File Quarantined, Host Timeout, Scan Error, MED Alerts, Host Reconnected

Recipients:
debrag@peersoftware.com

3. Click **OK**.

   The alert is listed on the **Email Alerts** page.

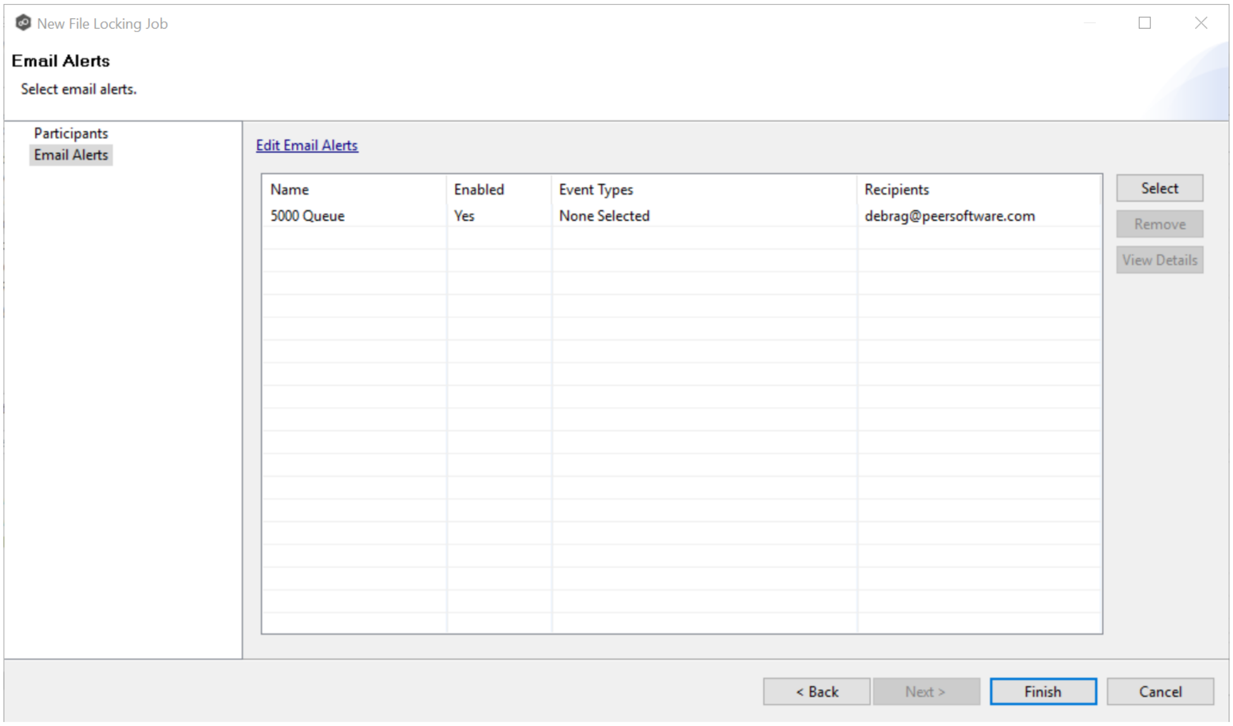

New File Locking Job

Email Alerts

Select email alerts.

Name	Enabled	Event Types	Recipients
5000 Queue	Yes	None Selected	debrag@peersoftware.com

4. (Optional) Repeat steps 1-3 to apply additional alerts.

5. After applying email alerts, continue to Step 4: Save Job.

Create an Email Alert

To create an alert or edit an existing alert from the **Email Alerts** page of the **Create Job** wizard:

1. Click the **Edit Email Alerts** link.

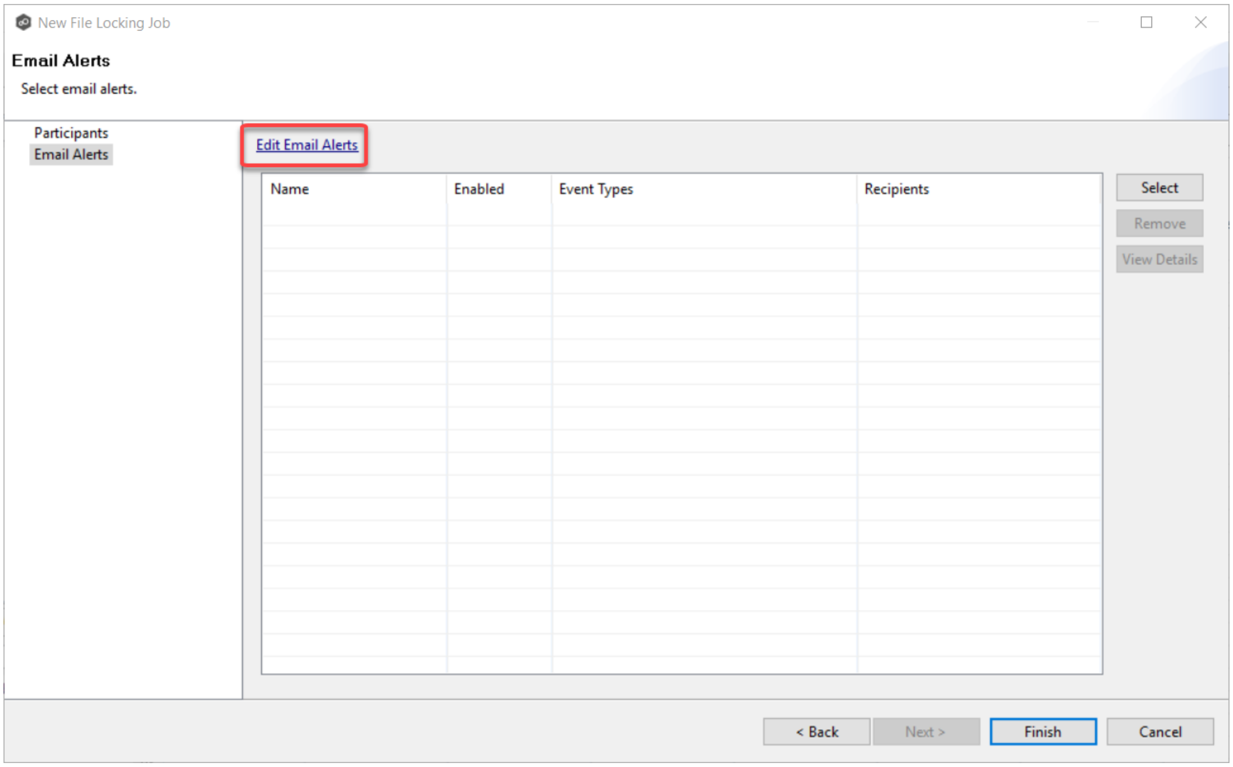

Email Alerts

Select email alerts.

Name	Enabled	Event Types	Recipients

The **Collaboration Email Alerts** dialog appears.

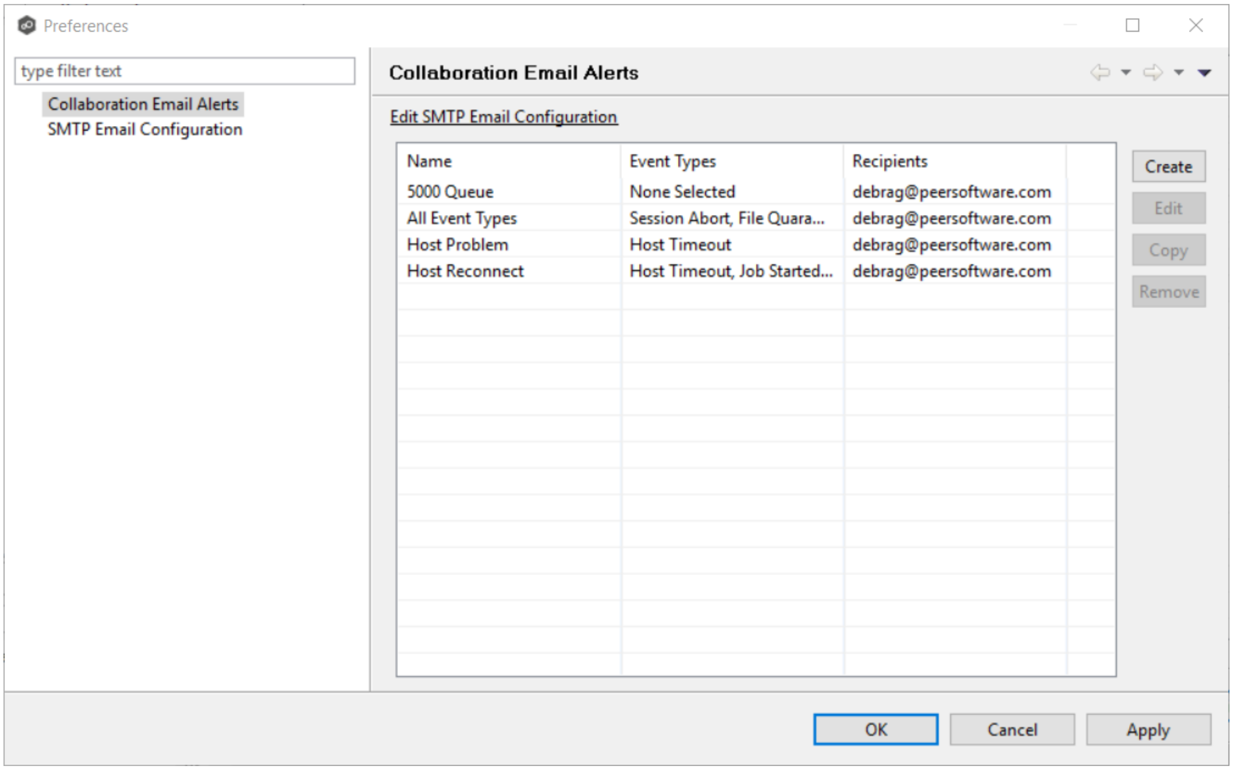

Collaboration Email Alerts

Name	Event Types	Recipients
5000 Queue	None Selected	debrag@peersoftware.com
All Event Types	Session Abort, File Quara...	debrag@peersoftware.com
Host Problem	Host Timeout	debrag@peersoftware.com
Host Reconnect	Host Timeout, Job Started...	debrag@peersoftware.com

2. Click **Create**.

   The **Create Email Alert** dialog appears.

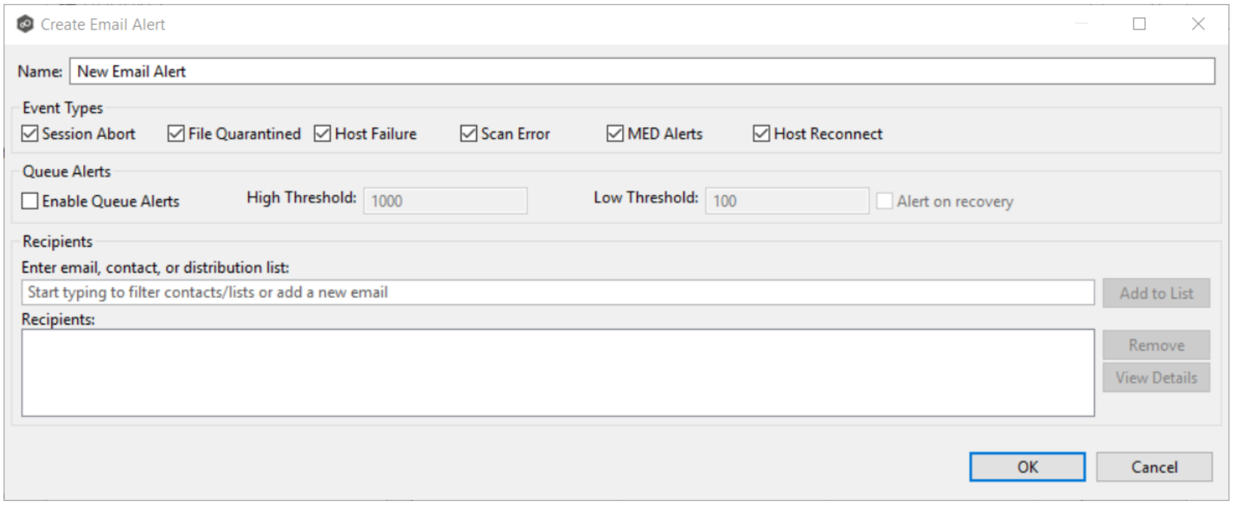

Create Email Alert

Name: New Email Alert

Event Types
☑ Session Abort ☑ File Quarantined ☑ Host Failure ☑ Scan Error ☑ MED Alerts ☑ Host Reconnect

Queue Alerts
☐ Enable Queue Alerts High Threshold: 1000 Low Threshold: 100 ☐ Alert on recovery

Recipients
Enter email, contact, or distribution list:
Start typing to filter contacts/lists or add a new email

Recipients:

3. Continue with Step 4 of the instructions in Email Alerts in the **Preferences** section.

Step 4: Save Job

Now that you have completed the first three 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 locking job. It is now listed in the **Jobs** view under **File Locking** and a job run-time view appears in **Summaries** area. You can start the job from either place.

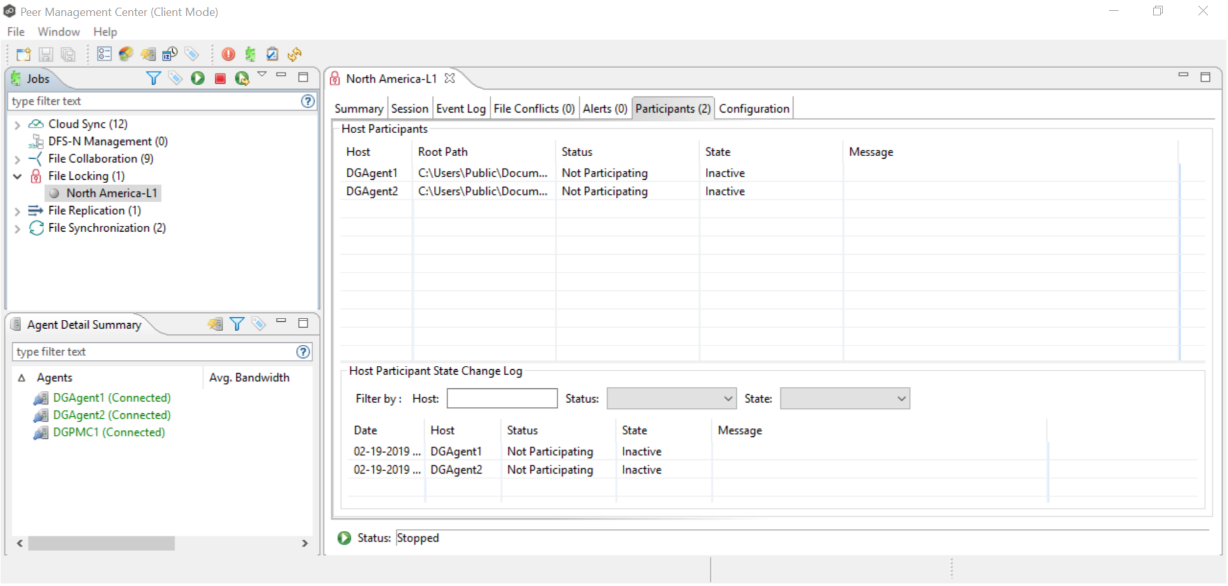

Agents	Avg. Bandwidth
DGAgent1 (Connected)	
DGAgent2 (Connected)	
DGPMC1 (Connected)	

Host Participants

Host	Root Path	Status	State	Message
DGAgent1	C:\Users\Public\Docum...	Not Participating	Inactive	
DGAgent2	C:\Users\Public\Docum...	Not Participating	Inactive	

Host Participant State Change Log

Date	Host	Status	State	Message
02-19-2019 ...	DGAgent1	Not Participating	Inactive	
02-19-2019 ...	DGAgent2	Not Participating	Inactive	

Status: Stopped

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 rely 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 the Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **New Job** wizard displays a list of job types you can create.
3. Click **File Replication**, and then click **Create**.

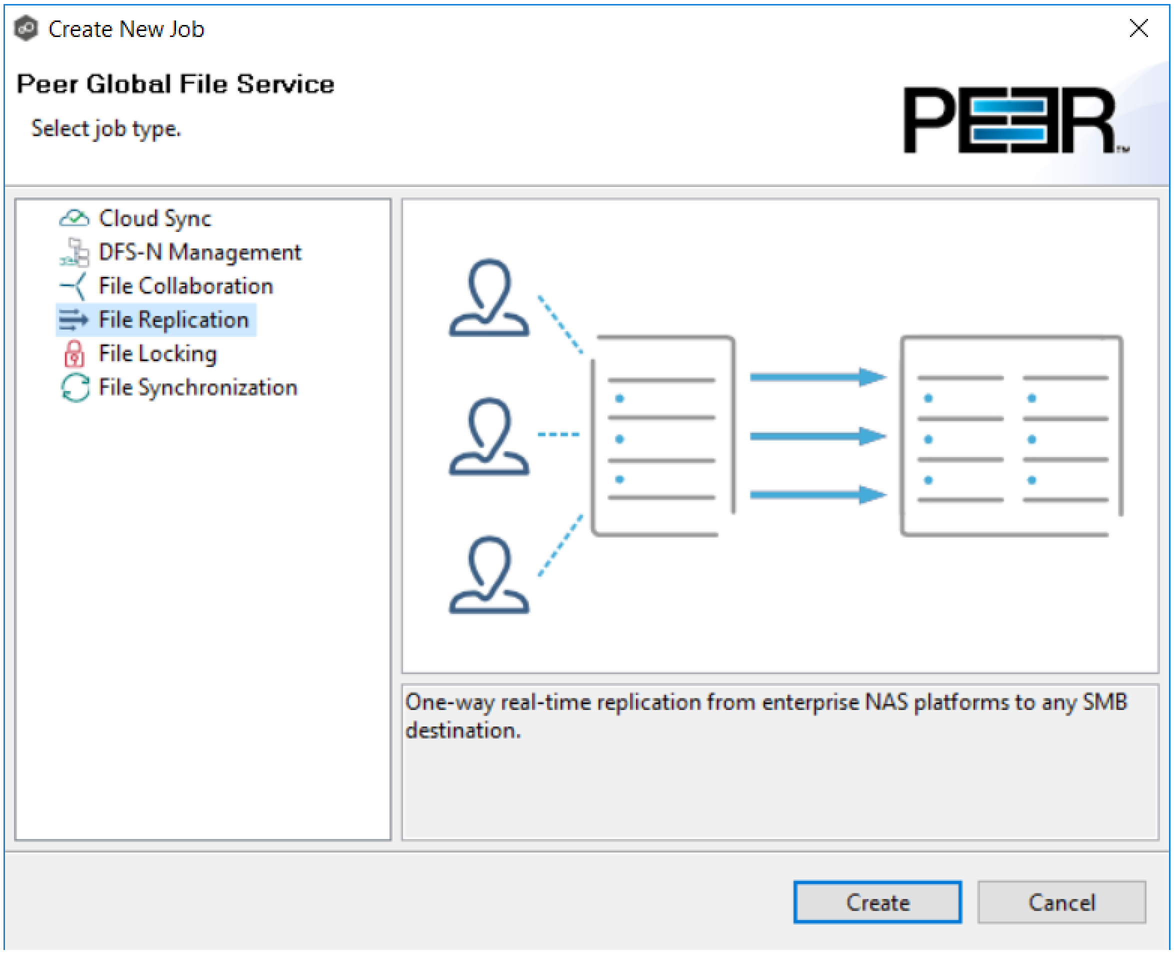
Create New Job

Peer Global File Service

Select job type.

One-way real-time replication from enterprise NAS platforms to any SMB destination.

4. Enter a name for the job in the dialog that appears.

   The job name must be unique.

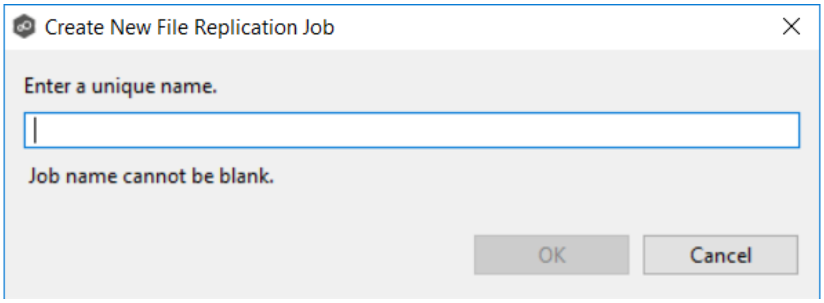
Create New File Replication Job

Enter a unique name.

Job name cannot be blank.

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.

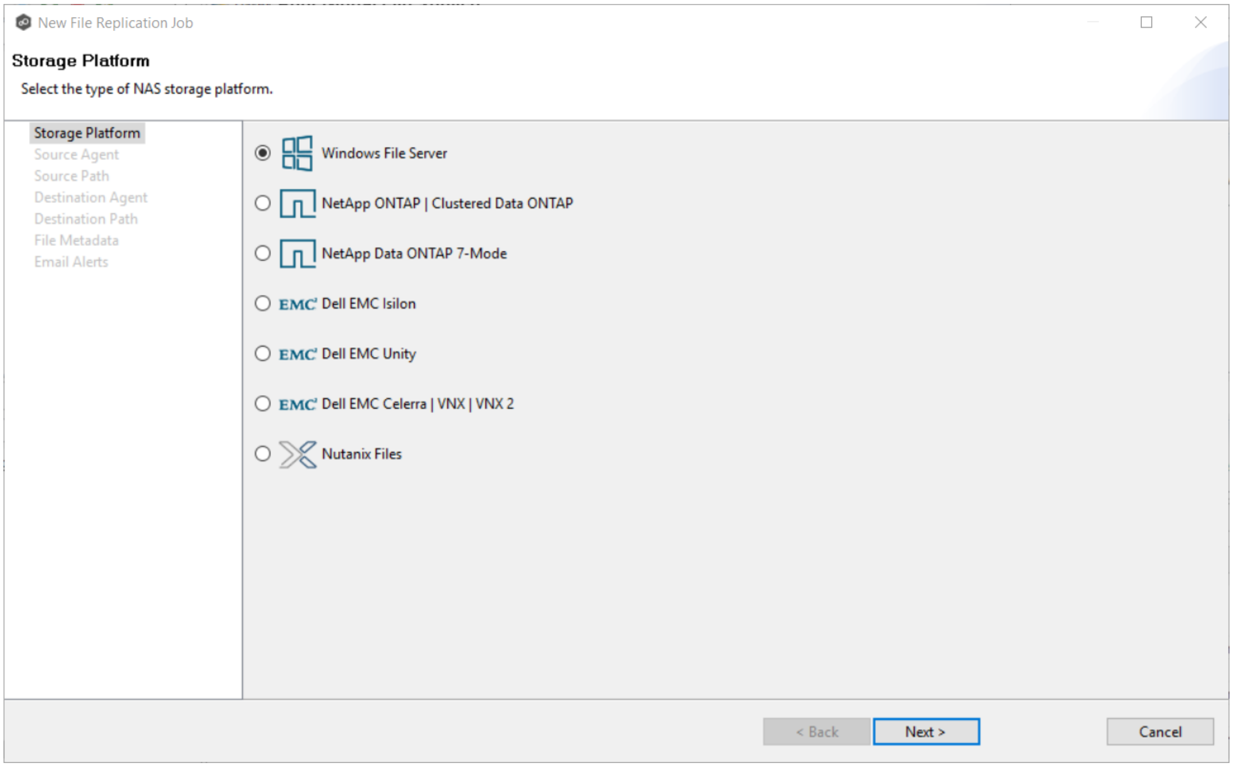

New File Replication Job

Storage Platform

Select the type of NAS storage platform.

- ● 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

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.

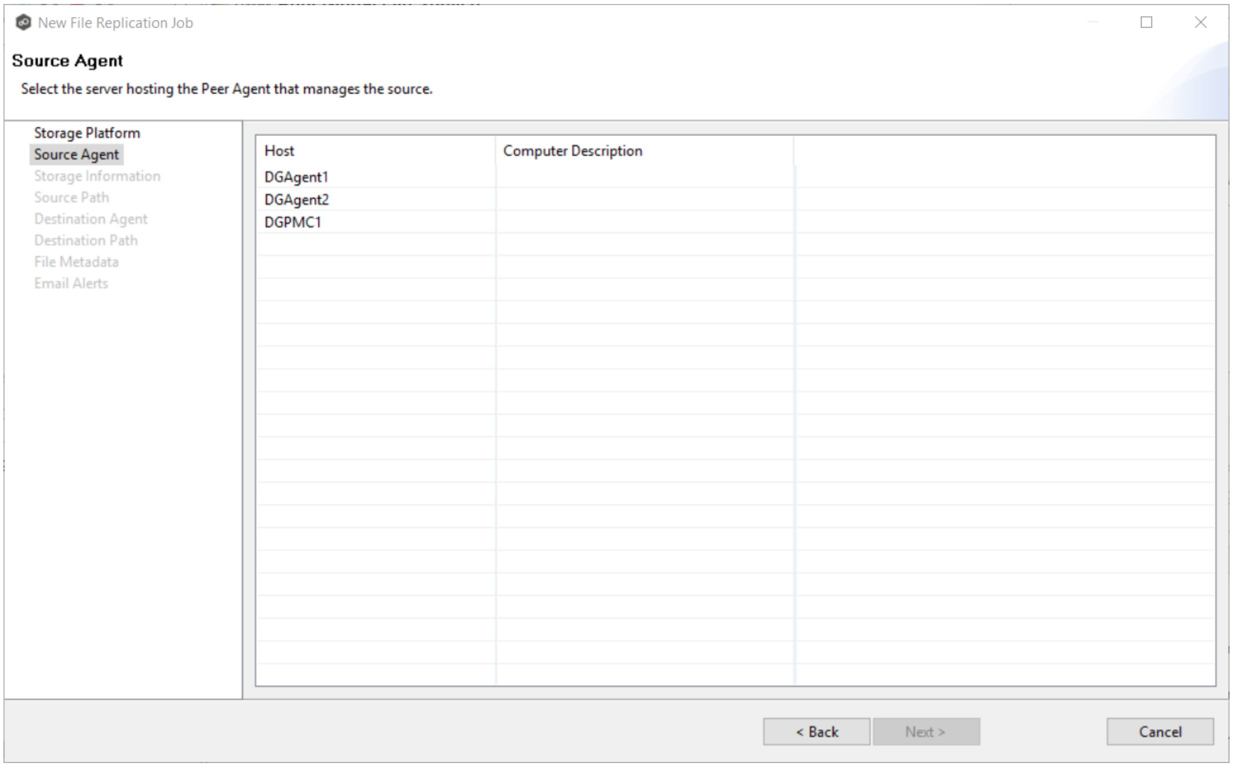

Source Agent

Select the server hosting the Peer Agent that manages the source.

Host	Computer Description
DGAgent1	
DGAgent2	
DGPMC1	

2. Click **Next**.

   The Storage Information page is displayed.

Step 4: Storage Information

If you selected any storage platform other than Windows File Server in Step 2, the **Storage Information** page appears. It requests the credentials necessary to connect to the storage device you want to replicate.

If you selected **Windows File Server**, skip to Step 5: Source Paths.

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 Paths.

   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. Enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

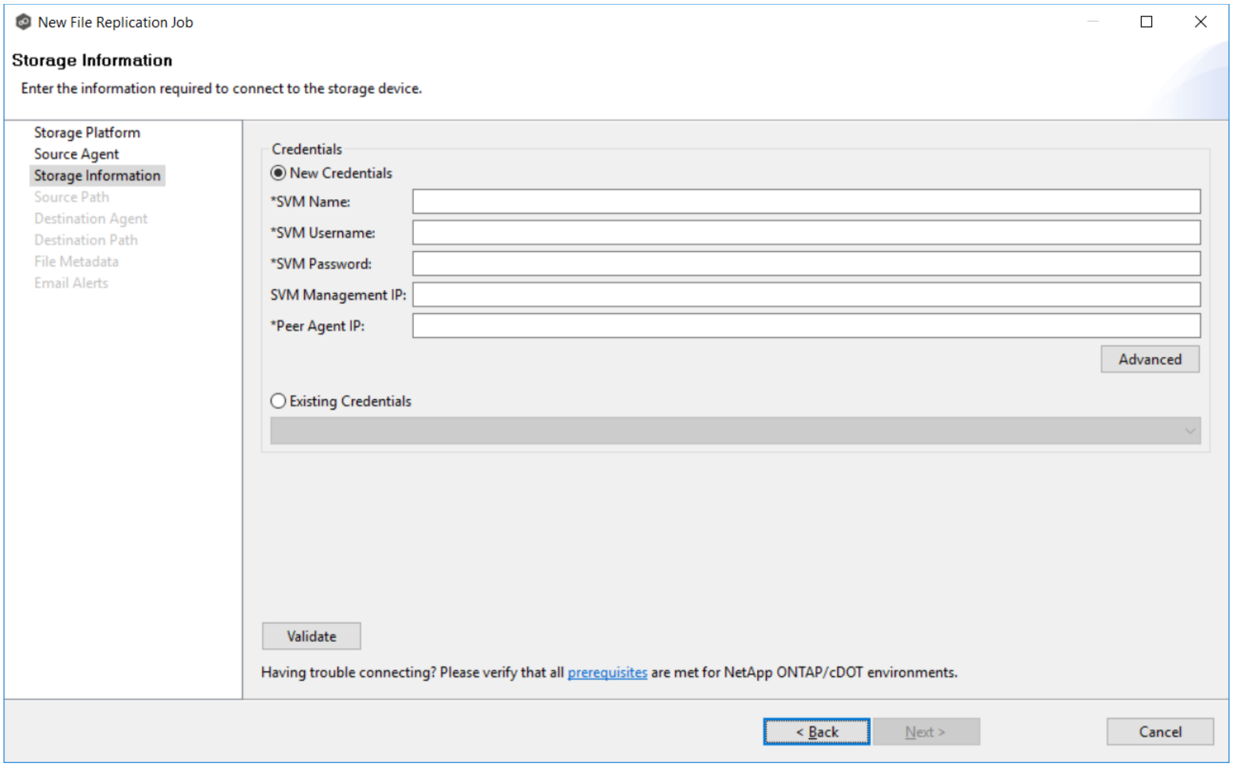
New File Replication Job

Storage Information

Enter the information required to connect to the storage device.

Having trouble connecting? Please verify that all prerequisites are met for NetApp ONTAP/cDOT environments.

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	Enter the IP address of the server hosting the Agent that manages the Storage Virtual Machine.

1. Enter the credentials to connect to the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

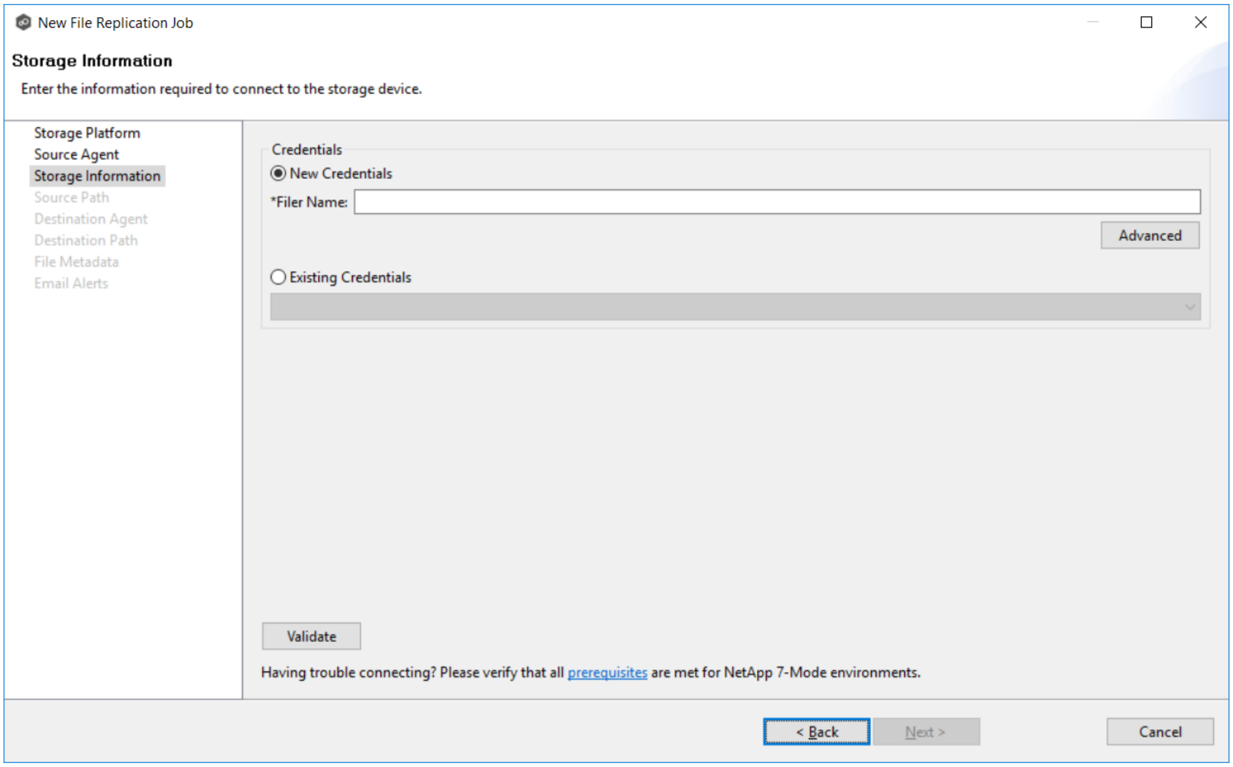

New File Replication Job

Storage Information

Enter the information required to connect to the storage device.

Having trouble connecting? Please verify that all prerequisites are met for NetApp 7-Mode environments.

Filer Name	Enter the name of the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

1. Enter the credentials to connect to the EMC Isilon cluster hosting the data to be replicated.

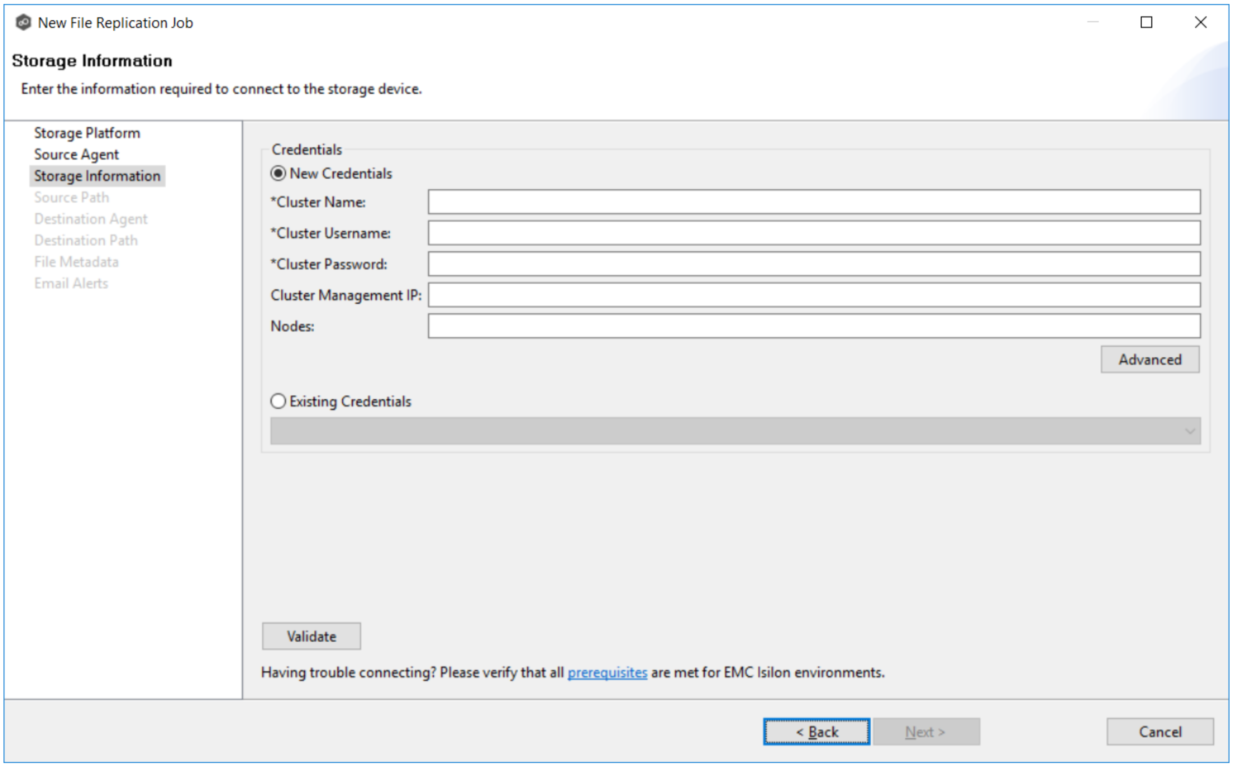

New File Replication Job

Storage Information

Enter the information required to connect to the storage device.

Credentials

● New Credentials

*Cluster Name:

*Cluster Username:

*Cluster Password:

Cluster Management IP:

Nodes:

○ Existing Credentials

Having trouble connecting? Please verify that all prerequisites are met for EMC Isilon environments.

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.
Nodes	For each node in the Isilon cluster, enter an IP address that can be reached by the Agent. Separate multiple values with a comma.

1. Enter the credentials to connect to the NAS Server hosting the data to be replicated.

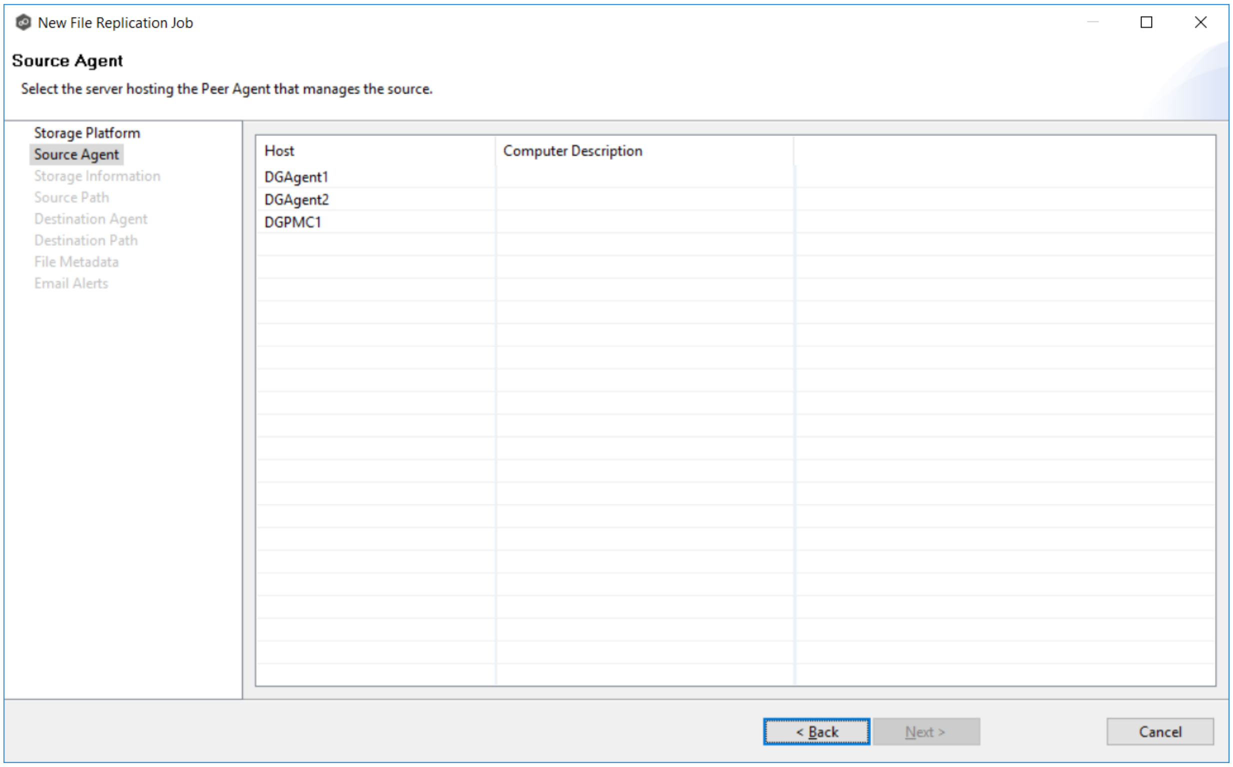
New File Replication Job

Source Agent

Select the server hosting the Peer Agent that manages the source.

Host	Computer Description
DGAgent1	
DGAgent2	
DGPMC1	

NAS 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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

1. Enter the credentials to connect to the CIFS Server hosting the data to be replicated.

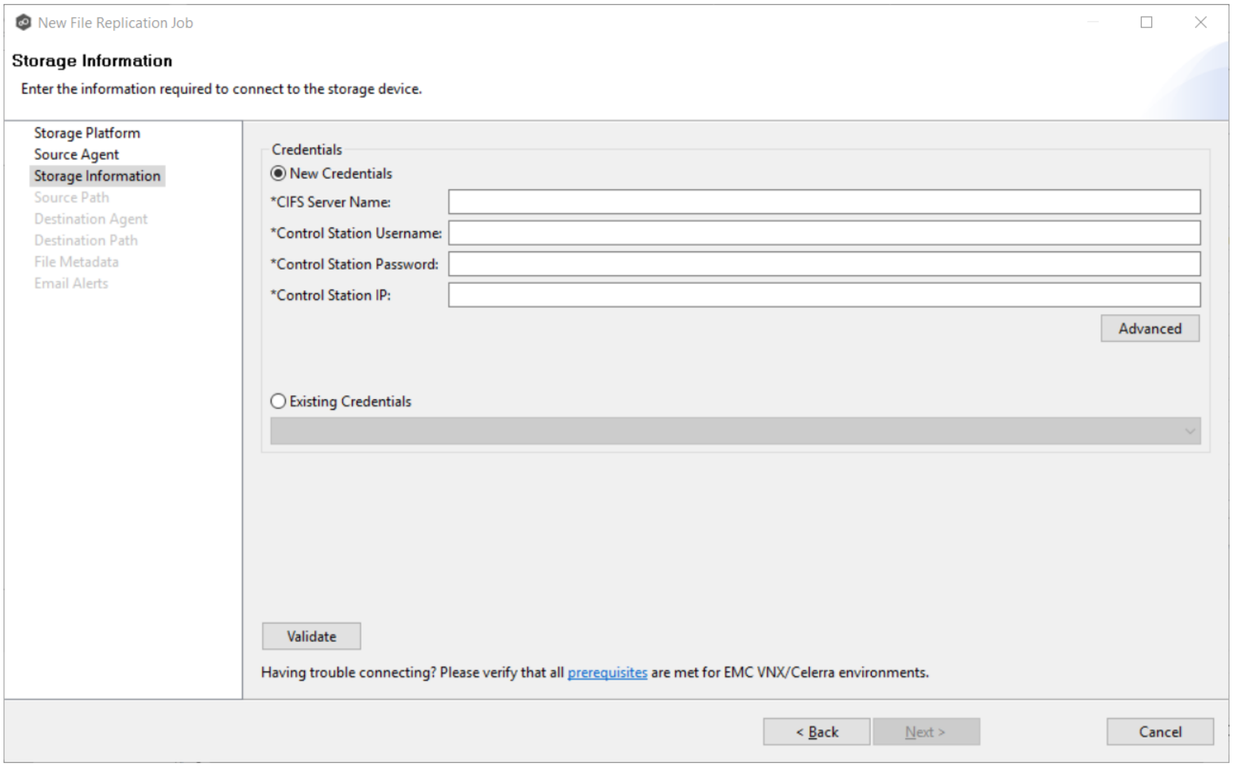

Storage Information

Enter the information required to connect to the storage device.

Credentials

● New Credentials

*CIFS Server Name:

*Control Station Username:

*Control Station Password:

*Control Station IP:

○ Existing Credentials

Having trouble connecting? Please verify that all prerequisites are met for EMC VNX/Celerra environments.

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.

Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

1. Enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated.

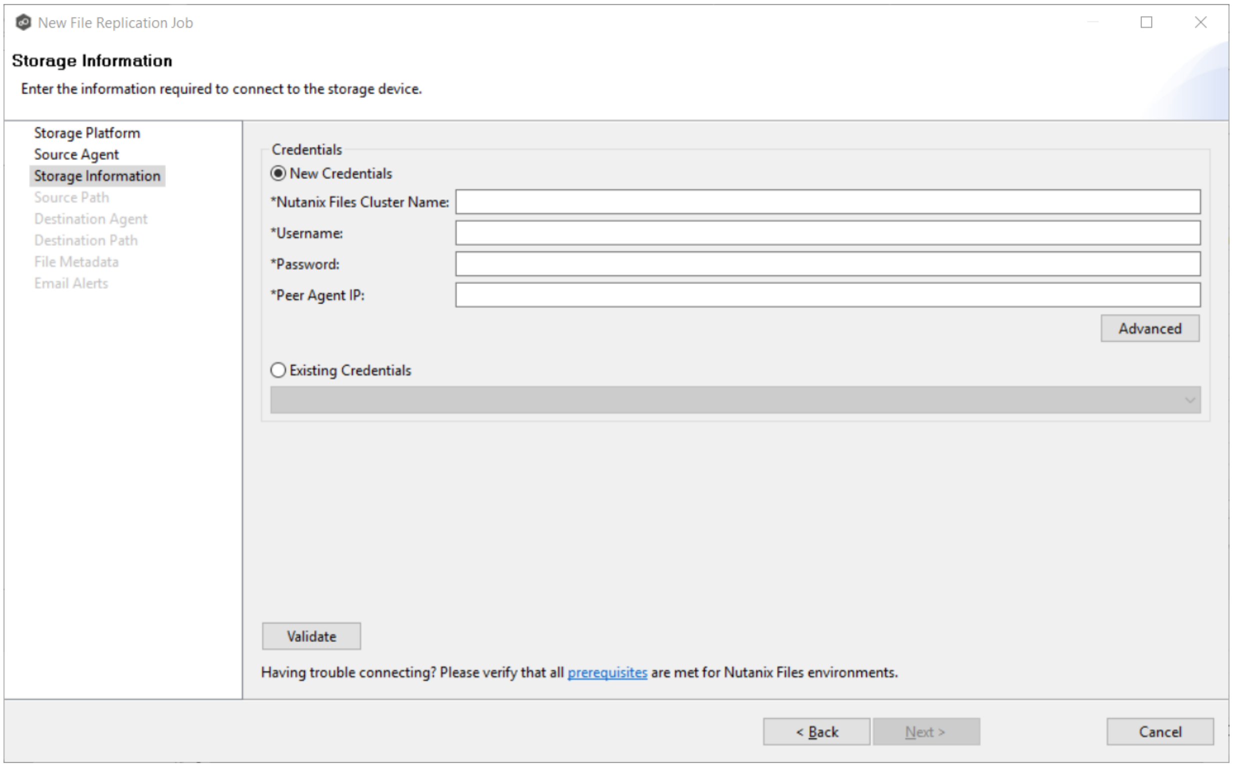

New File Replication Job

Storage Information

Enter the information required to connect to the storage device.

Credentials

● New Credentials

*Nutanix Files Cluster Name:

*Username:

*Password:

*Peer Agent IP:

○ Existing Credentials

Having trouble connecting? Please verify that all prerequisites are met for Nutanix Files environments.

Nutanix Files Cluster 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 AFS cluster via its management APIs.
Password	Enter the password for the account managing the AFS cluster via its management APIs.

Peer Agent IP	Enter the IP address of the Agent server used to manage the storage platform. This should not point to the AFS cluster itself.

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.

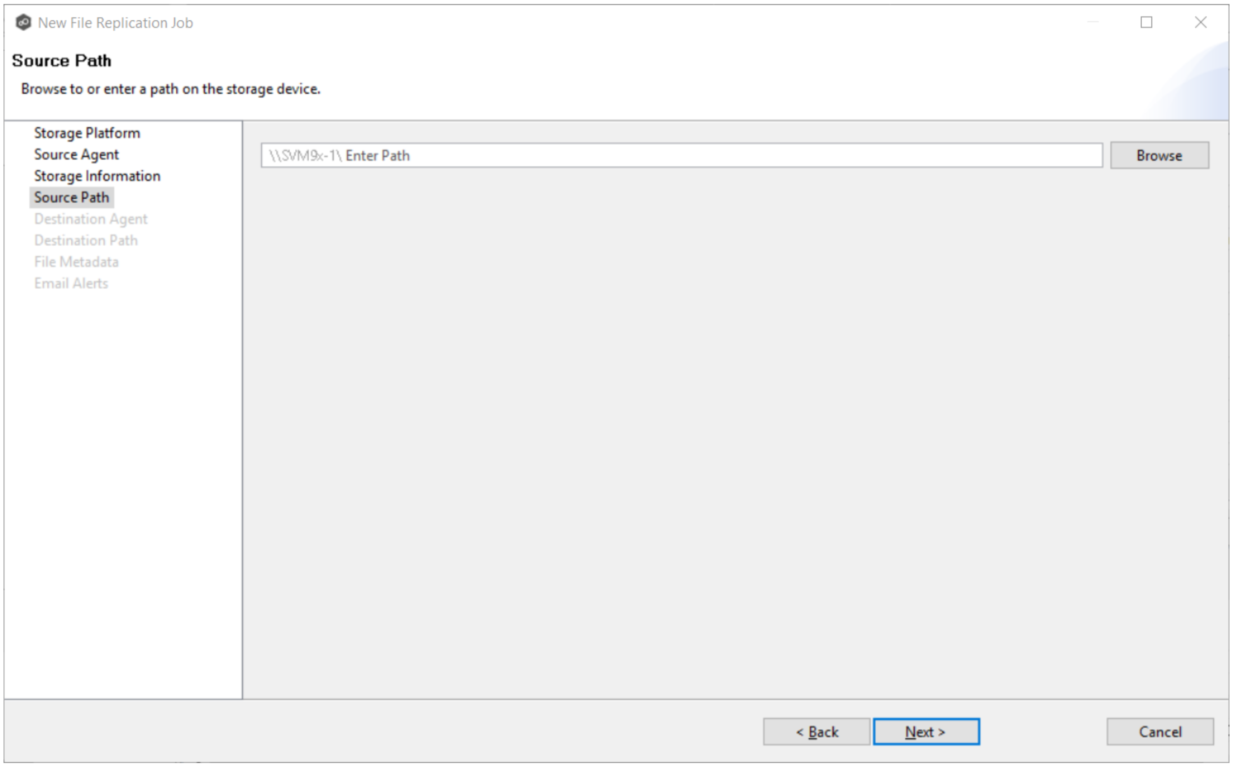

New File Replication Job

Source Path

Browse to or enter a path on the storage device.

If you selected **Browse**, the **Folder Browser** dialog appears:

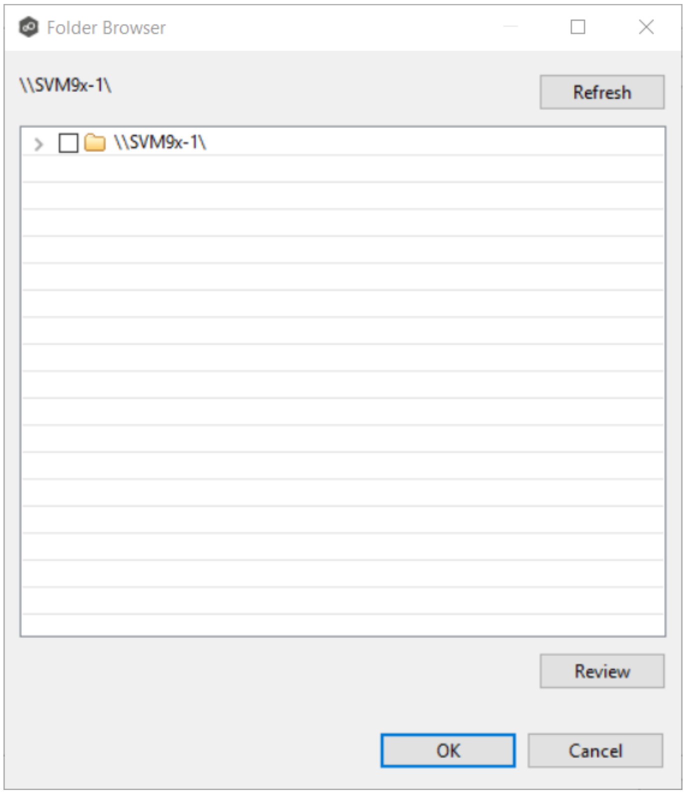

a. Expand the folder tree.

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.

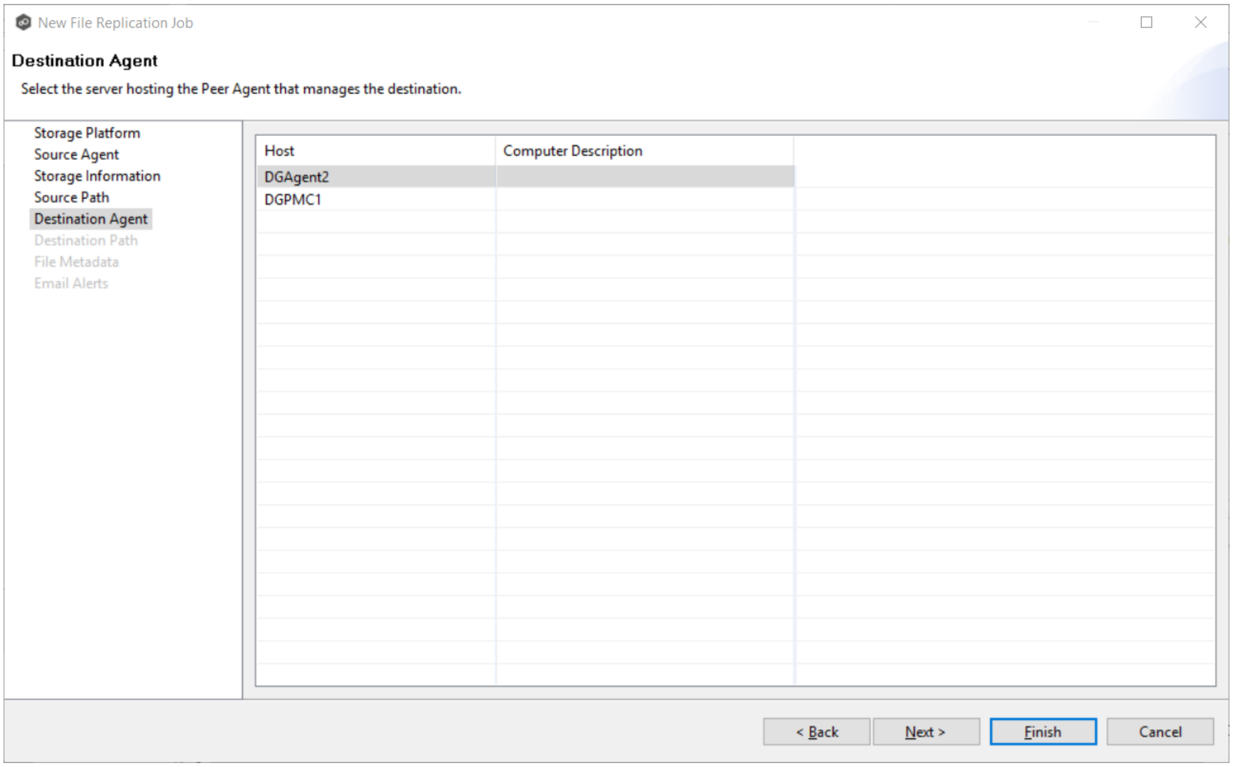

New File Replication Job

Destination Agent

Select the server hosting the Peer Agent that manages the destination.

Host	Computer Description
DGAgent2	
DGPMC1	

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.

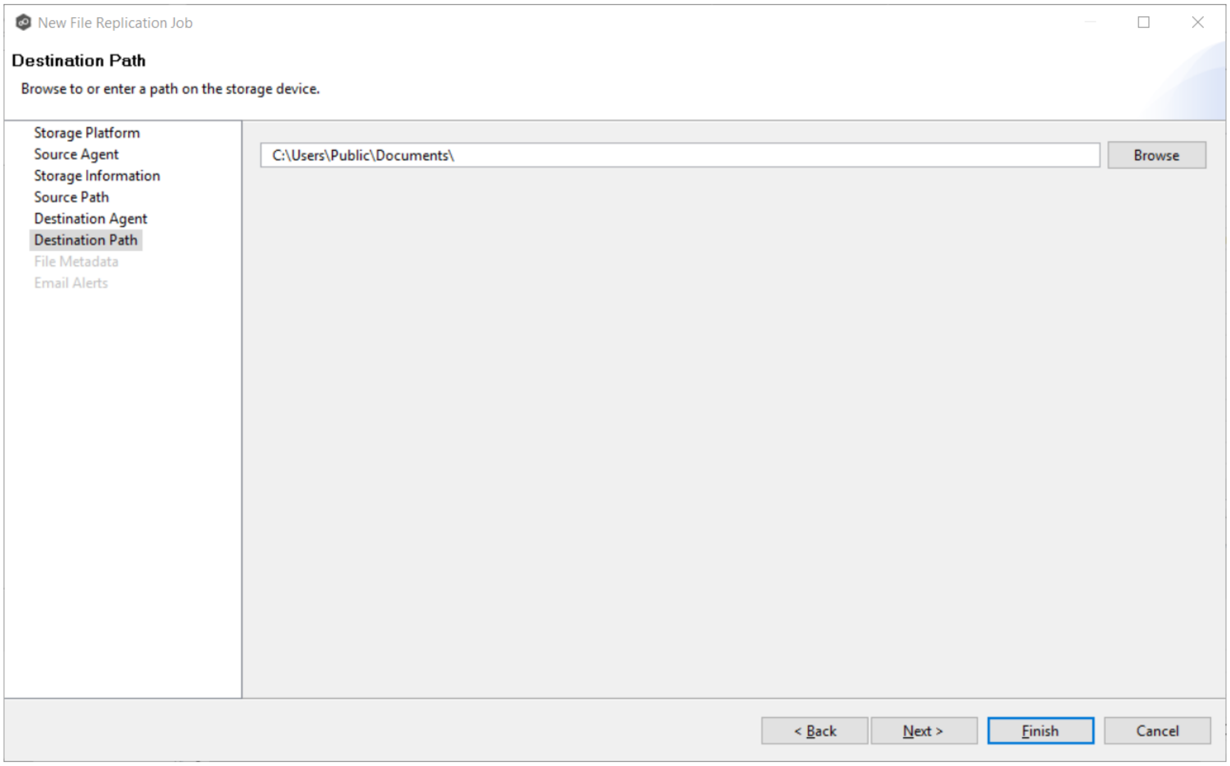
New File Replication Job

Destination Path

Browse to or enter a path on the storage device.

C:\Users\Public\Documents\

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 the Advanced Topic File Metadata Replication.

To enable file metadata replication:

1. Select when you want the metadata replicated (you can select one or both of the 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 (DACL, SACL, Owner.) 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 (DACL, SACL, Owner) will be synchronized during the initial scan.

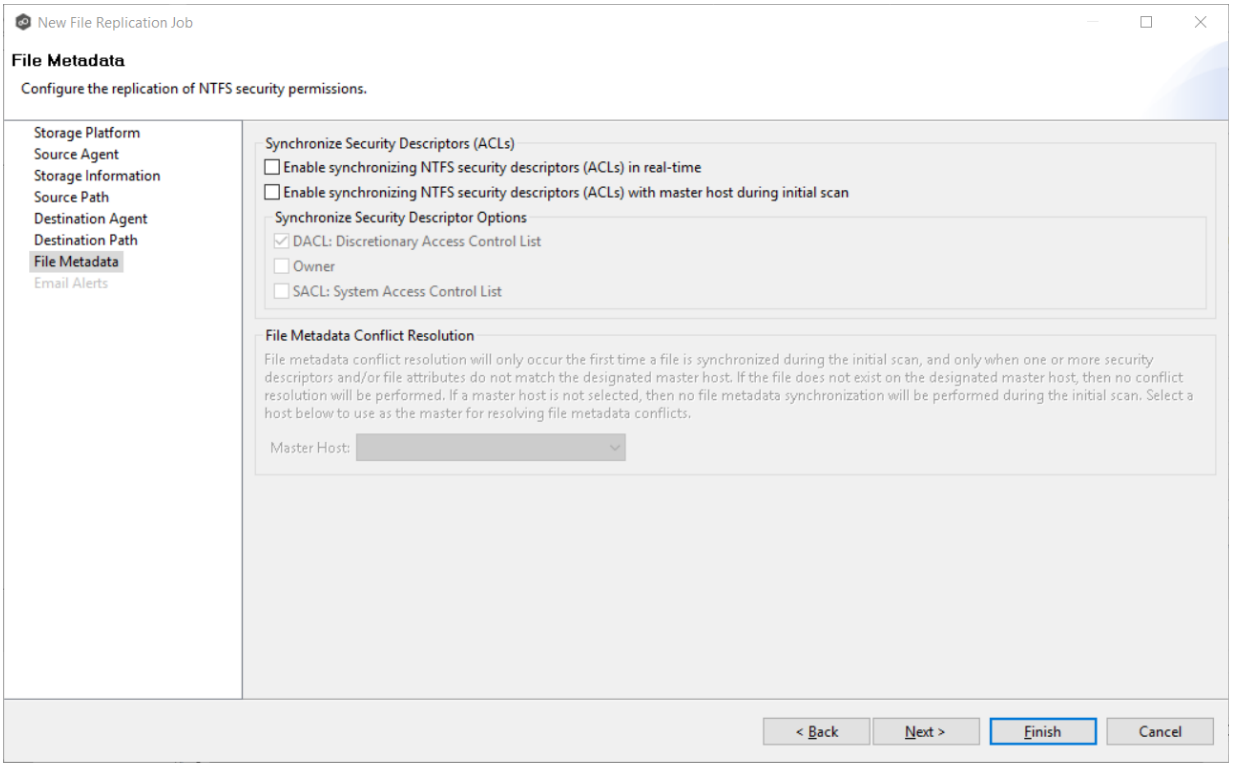

New File Replication Job

File Metadata

Configure the replication of NTFS security permissions.

Synchronize Security Descriptors (ACLs)

□ Enable synchronizing NTFS security descriptors (ACLs) in real-time

□ Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan

Synchronize Security Descriptor Options

☑ DACL: Discretionary Access Control List

□ Owner

□ SACL: System Access Control List

File Metadata Conflict Resolution

File metadata conflict resolution will only occur the first time a file is synchronized during the initial scan, and only when one or more security descriptors and/or file attributes do not match the designated master host. If the file does not exist on the designated master host, then no conflict resolution will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan. Select a host below to use as the master for resolving file metadata conflicts.

Master Host:

2. Click **OK** in the informational dialog that appears after selecting a metadata option.

3. Select which security descriptor components (DACL, SACL and Owner) are synchronized.

   In general, you will usually need to synchronize only DACLs. If you need to synchronize SACLs or Owner, then the user that a Peer Agent service is run under on each participating host must have permission to read and write SACLs and Owner.

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.

Conflict resolution for file metadata 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 will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan.

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.

Apply an Existing Email Alert

To apply an existing email alert to the job.

1. Click the **Select** button.

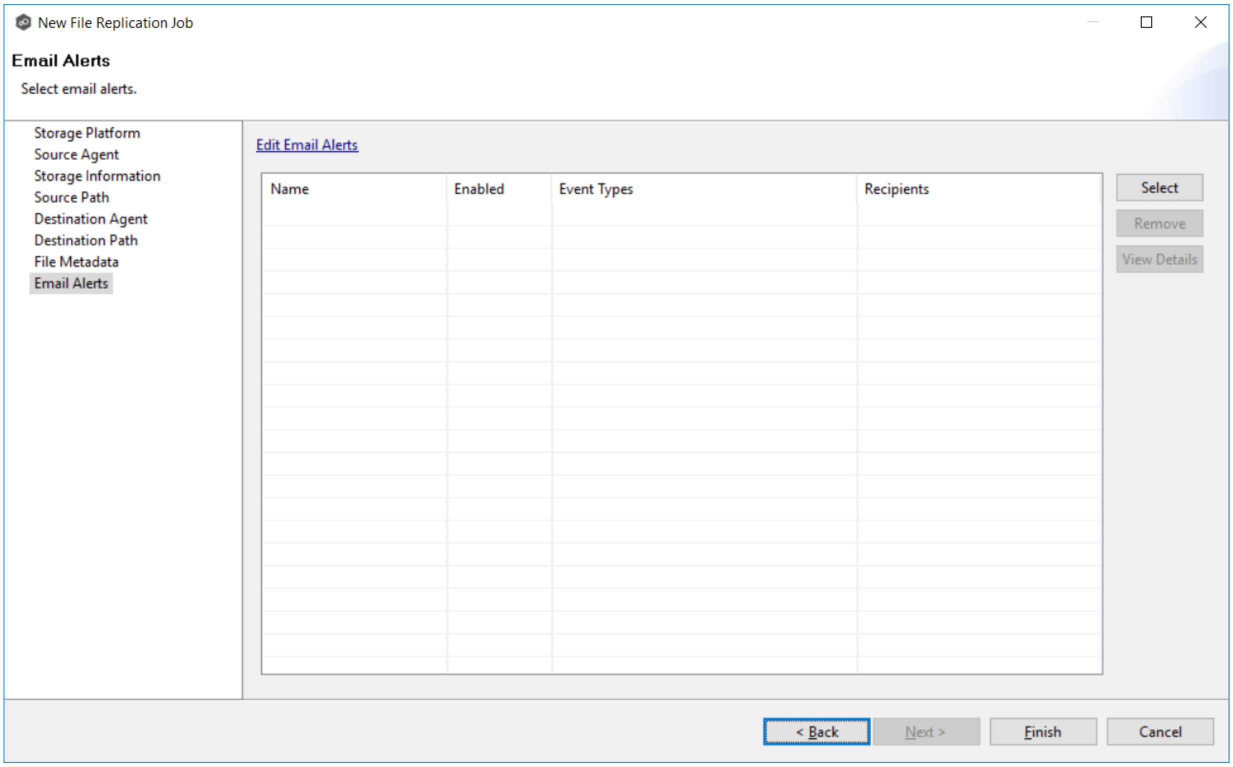

The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.

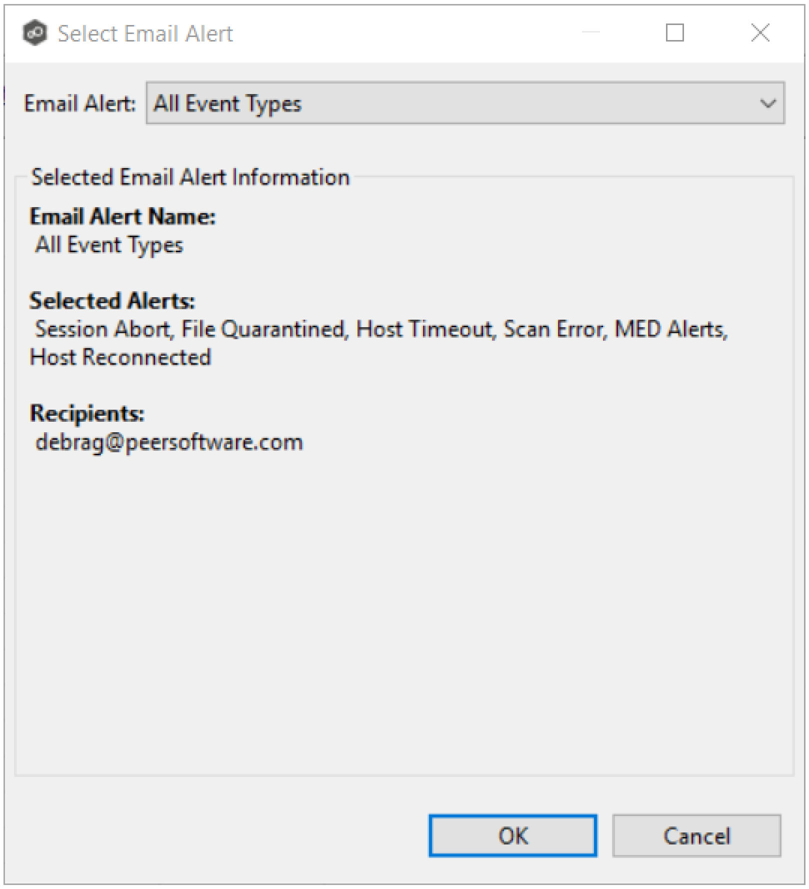

Email Alert: All Event Types

Selected Email Alert Information

Email Alert Name:
All Event Types

Selected Alerts:
Session Abort, File Quarantined, Host Timeout, Scan Error, MED Alerts, Host Reconnected

Recipients:
debrag@peersoftware.com

3. Click **OK**.

   The alert is listed on the **Email Alerts** page.

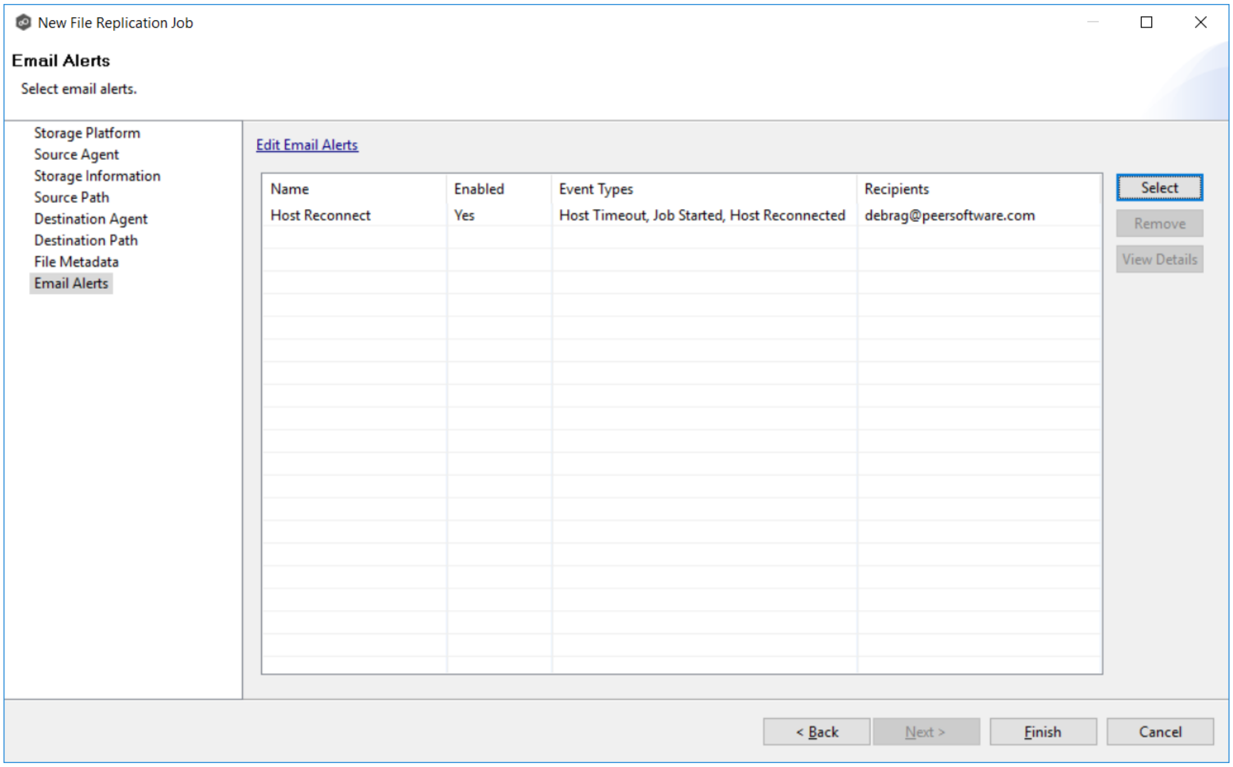

New File Replication Job

Email Alerts

Select email alerts.

Name	Enabled	Event Types	Recipients
Host Reconnect	Yes	Host Timeout, Job Started, Host Reconnected	debrag@peersoftware.com

4. (Optional) Repeat steps 1-3 to apply additional alerts.

5. After applying email alerts, continue to Step 10: Save Job.

Create an Email Alert

To create an alert or edit an existing alert from the **Email Alerts** page of the **Create Job** wizard:

1. Click the **Edit Email Alerts** link.

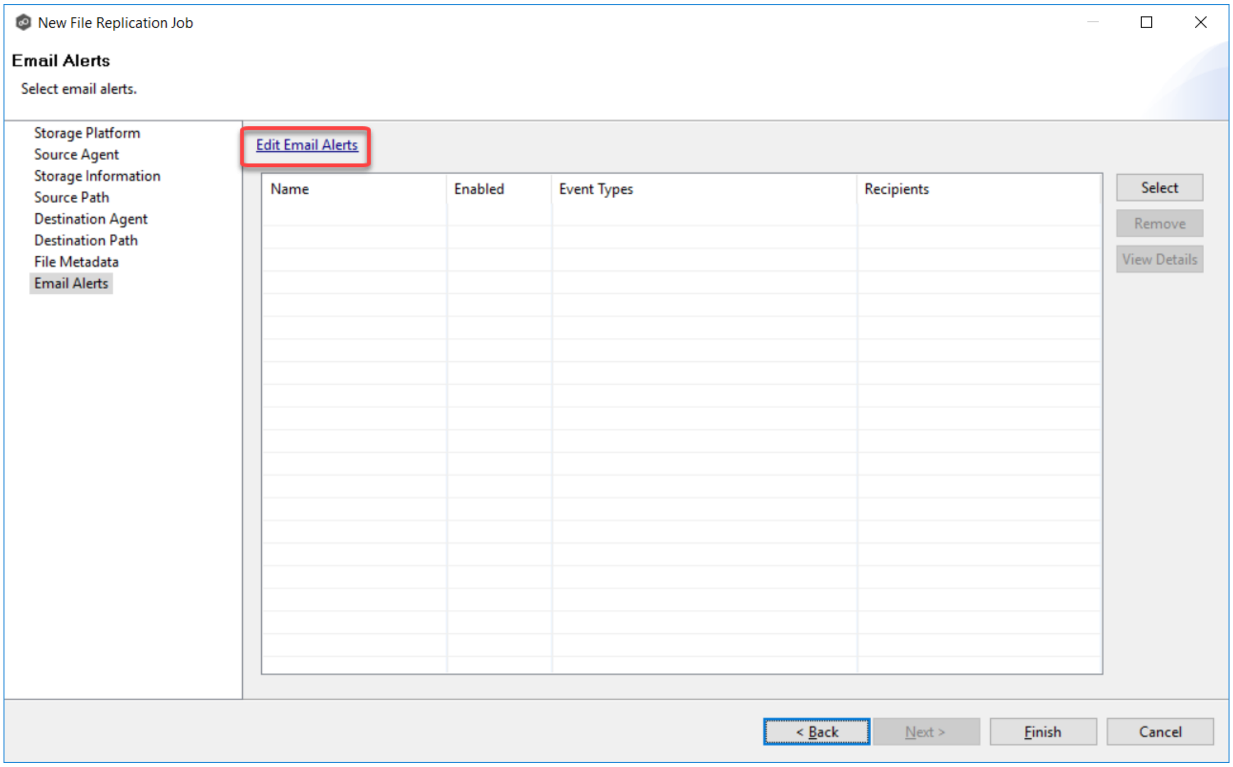

New File Replication Job

Email Alerts

Select email alerts.

Name	Enabled	Event Types	Recipients

The **Collaboration Email Alerts** dialog appears.

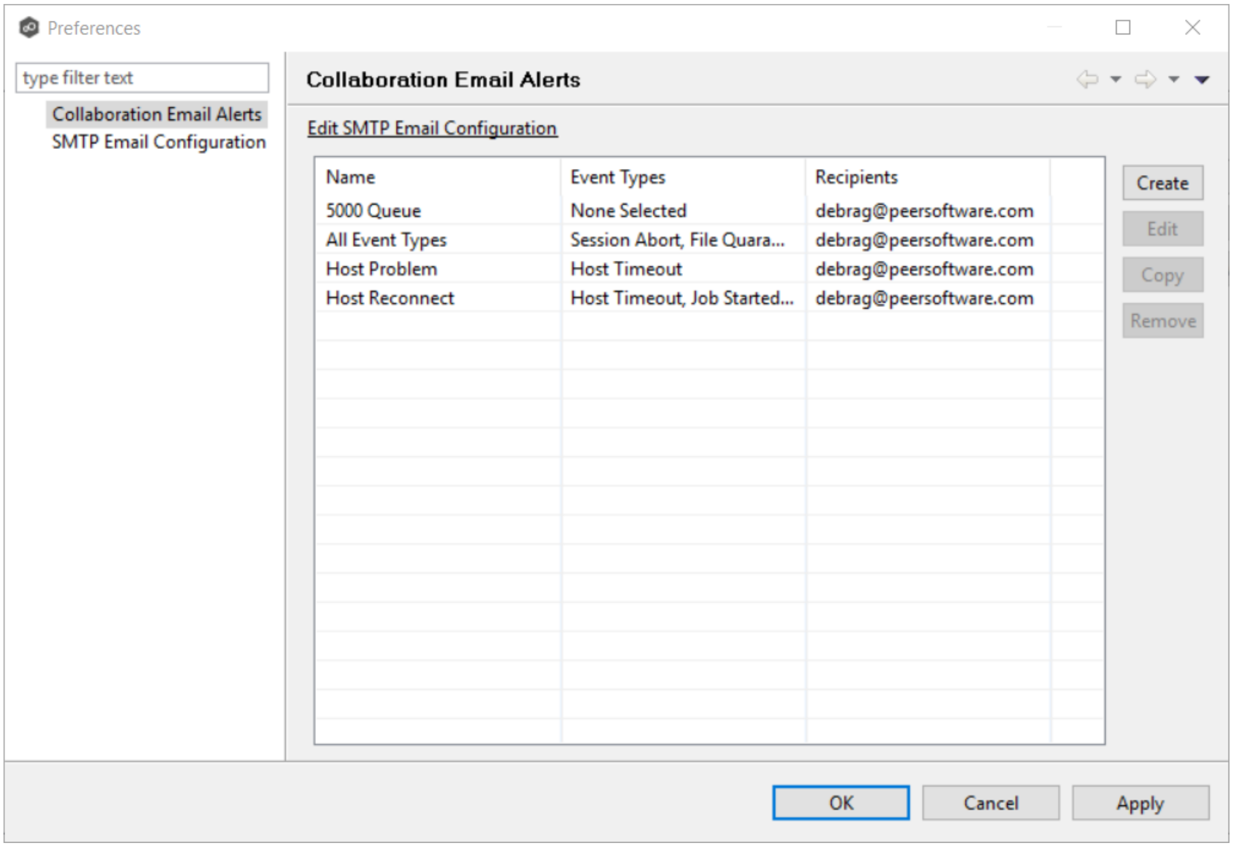

Collaboration Email Alerts

Name	Event Types	Recipients
5000 Queue	None Selected	debrag@peersoftware.com
All Event Types	Session Abort, File Quara...	debrag@peersoftware.com
Host Problem	Host Timeout	debrag@peersoftware.com
Host Reconnect	Host Timeout, Job Started...	debrag@peersoftware.com

2. Click **Create**.

   The **Create Email Alert** dialog appears.

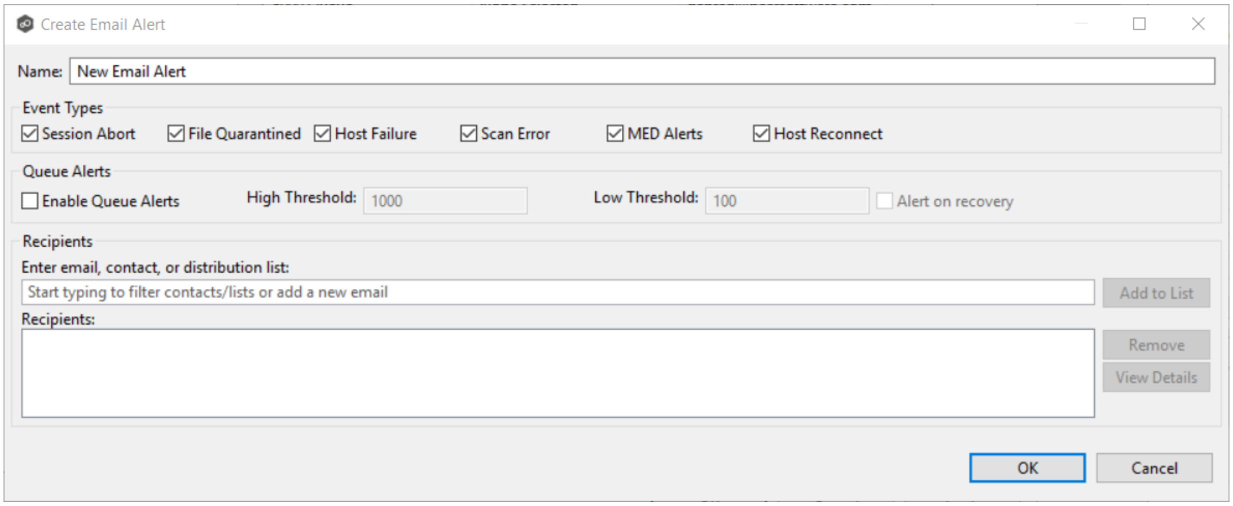

Name: New Email Alert

Event Types

☑ Session Abort ☑ File Quarantined ☑ Host Failure ☑ Scan Error ☑ MED Alerts ☑ Host Reconnect

Queue Alerts

☐ Enable Queue Alerts High Threshold: 1000 Low Threshold: 100 ☐ Alert on recovery

Recipients

Enter email, contact, or distribution list:

Start typing to filter contacts/lists or add a new email

Recipients:

3. Continue with Step 4 of the instructions in Email Alerts in the **Preferences** section.

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 Summaries area. You can start the job from either place.

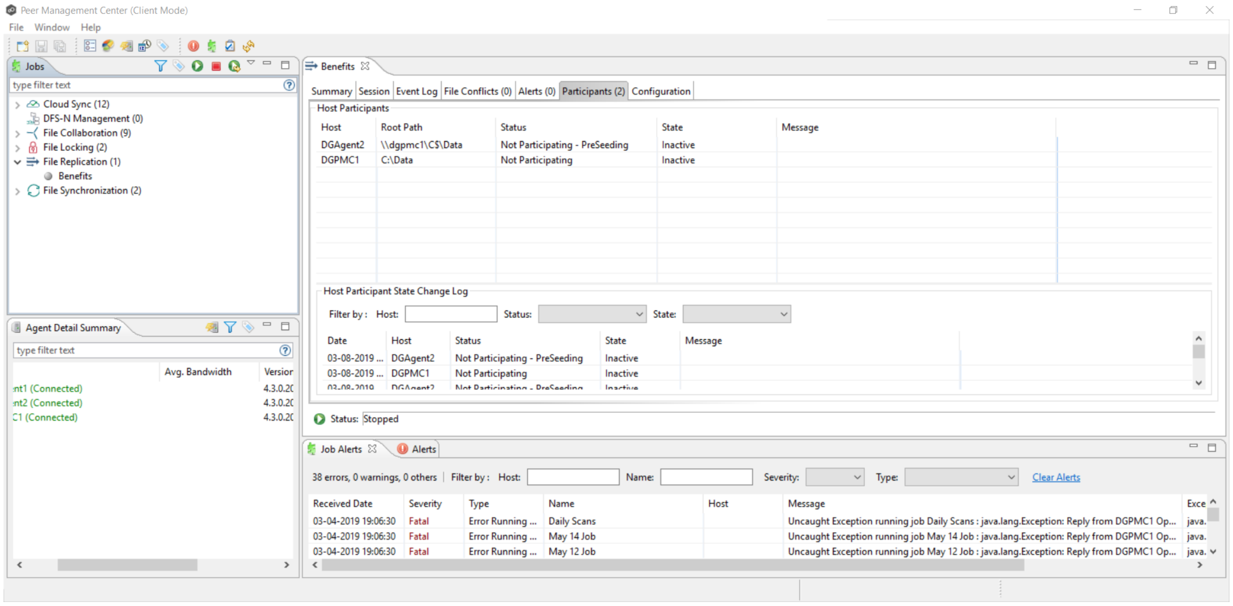

Avg. Bandwidth	Version
nt1 (Connected)	4.3.0.20
nt2 (Connected)	4.3.0.20
C1 (Connected)	4.3.0.20

Host Participants

Host	Root Path	Status	State	Message
DGAgent2	\\dgpmc1\C$\Data	Not Participating - PreSeeding	Inactive	
DGPMC1	C:\Data	Not Participating	Inactive	

Host Participant State Change Log

Filter by : Host: Status: State:

Date	Host	Status	State	Message
03-08-2019 ...	DGAgent2	Not Participating - PreSeeding	Inactive	
03-08-2019 ...	DGPMC1	Not Participating	Inactive	
03-08-2019	DGAgent2	Not Participating - PreSeeding	Inactive	

Status: Stopped

38 errors, 0 warnings, 0 others | Filter by : Host: Name: Severity: Type:

Received Date	Severity	Type	Name	Host	Message	Exce
03-04-2019 19:06:30	Fatal	Error Running ...	Daily Scans		Uncaught Exception running job Daily Scans : java.lang.Exception: Reply from DGPMC1 Op...	java.
03-04-2019 19:06:30	Fatal	Error Running ...	May 14 Job		Uncaught Exception running job May 14 Job : java.lang.Exception: Reply from DGPMC1 Op...	java.
03-04-2019 19:06:30	Fatal	Error Running ...	May 12 Job		Uncaught Exception running job May 12 Job : java.lang.Exception: Reply from DGPMC1 Op...	java.

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 the Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

   The **New Job** wizard displays a list of job types you can create.

3. Click **File Locking**, and then click **Create**.

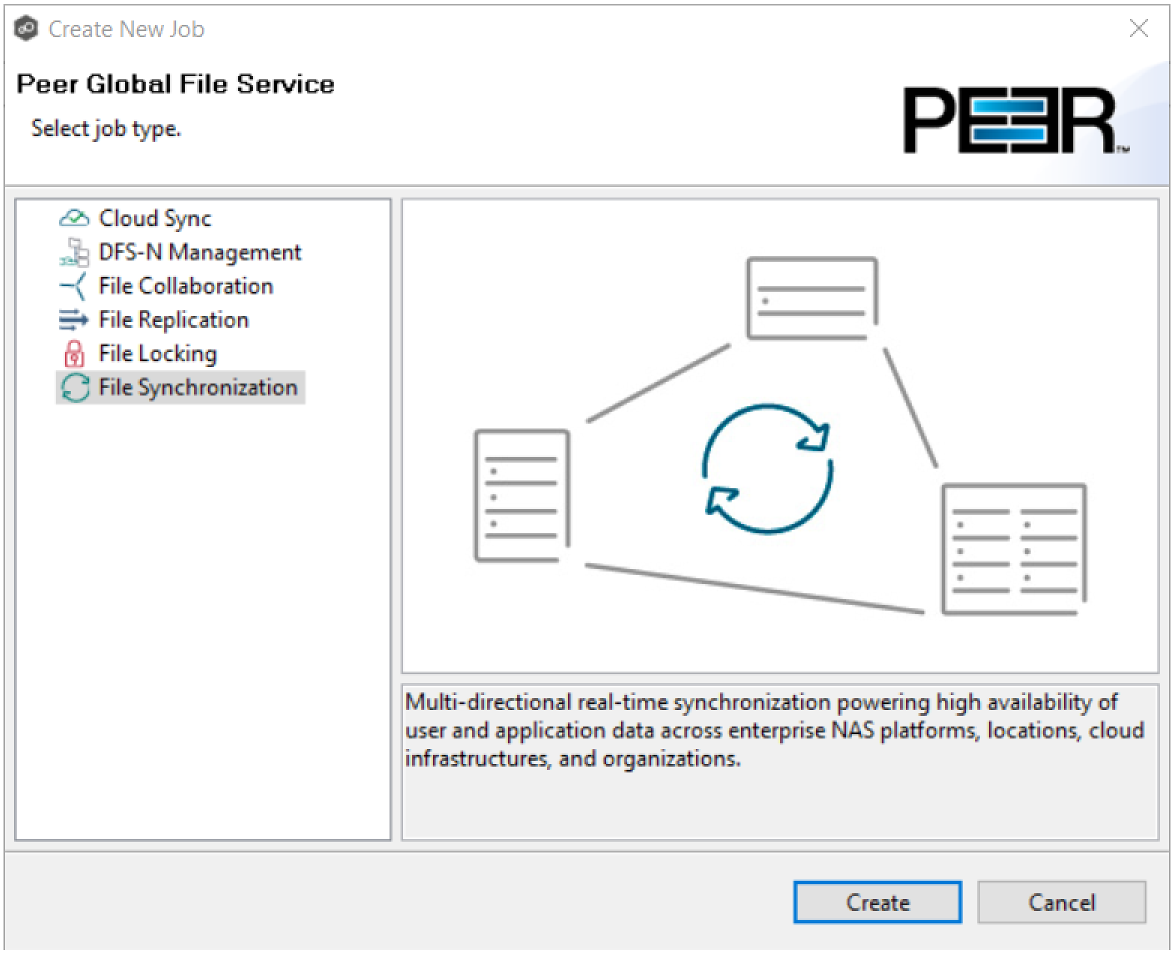

Multi-directional real-time synchronization powering high availability of user and application data across enterprise NAS platforms, locations, cloud infrastructures, and organizations.

4. Enter a name for the job in the dialog that appears.

   The job name must be unique.

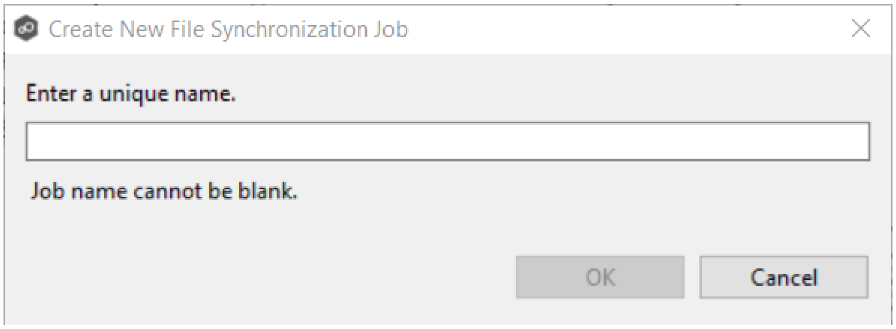

Job name cannot be blank.

5. Click **OK**.

   The Participants page is displayed.

Step 2: Participants

A File Synchronization job must have two or more participants. A participant consists of an Agent and the volume/share/folder to be replicated. The server that the Agent is installed upon is called the host (or host participant). A File Synchronization job synchronizes the files of participants in real-time.

1. Complete the five substeps:

   Participants

   Storage Platform

   Management Agent

   Storage Information

   Path

   After you add a participant, it appears in the **Participants** table.

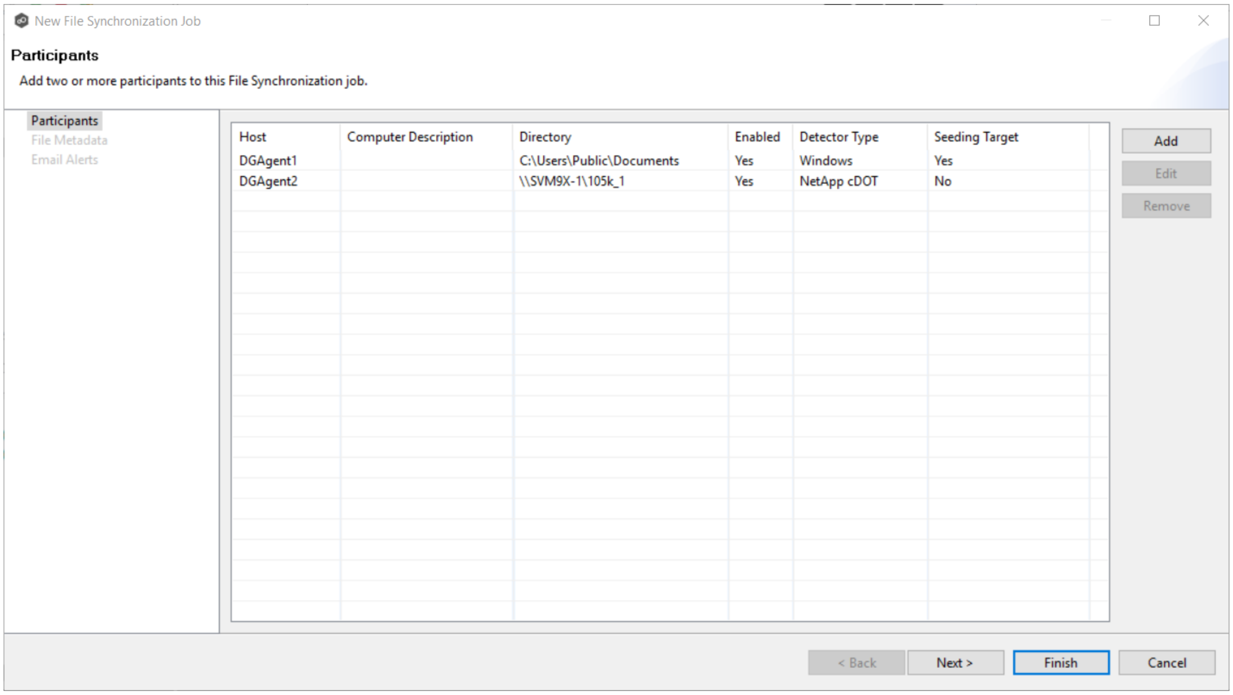

Participants

Add two or more participants to this File Synchronization job.

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent1		C:\Users\Public\Documents	Yes	Windows	Yes
DGAgent2		\\SVM9X-1\105k_1	Yes	NetApp cDOT	No

2. Repeat the five substeps for each participant you want to add to the job.

3. Once you have added all the participants, click **Next** to specify file metadata for the job. (Don't click **Finish**.)

The **Participants** page is where you select and configure which hosts will be participating in this job. The **Participants** page is empty until you finish the process of adding your first participant. Once you have added the participants, they are listed on the **Participants** page.

To begin the process of adding a participant:

1. Click the **Add** button.

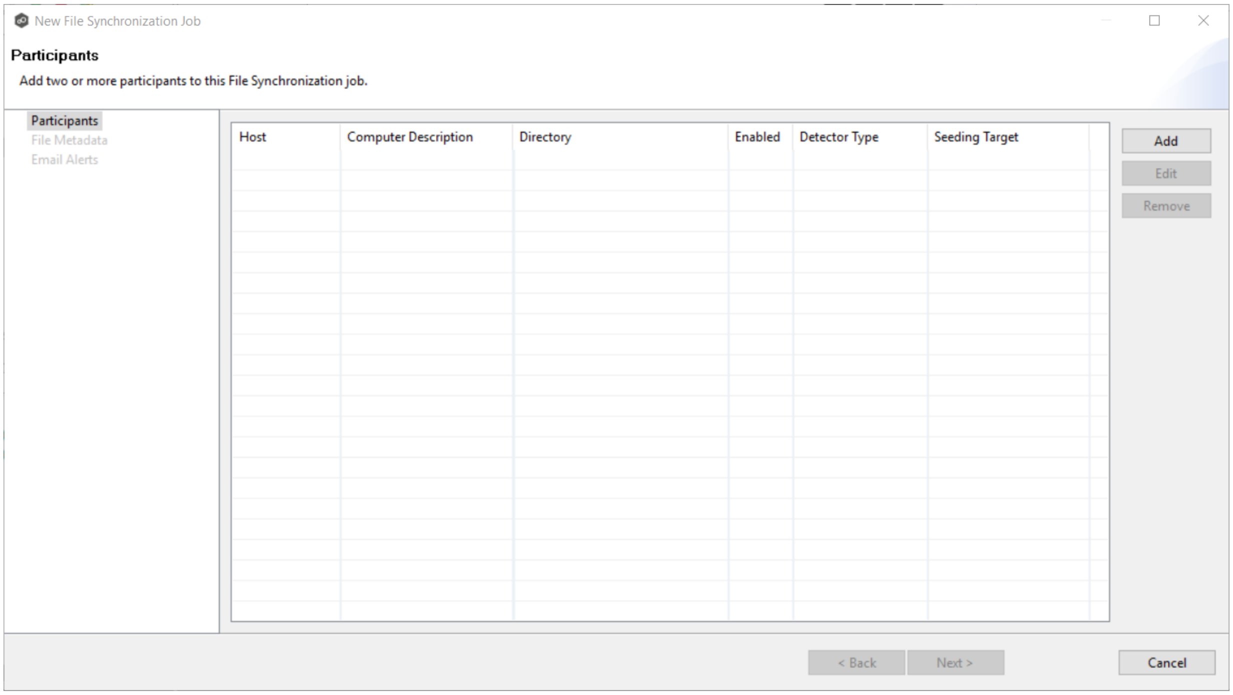

Participants

Add two or more participants to this File Synchronization job.

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target

Another wizard opens to guide you through the process of adding a participant to the job. The first step in the process involves selecting the storage platform.

The **Storage Platform** page lists the types of storage platforms that File Synchronization supports. A storage device hosts data you want to synchronize. It is often referred to as the host or host participant.

1. Select the type of storage platform that hosts the data you want to synchronize.

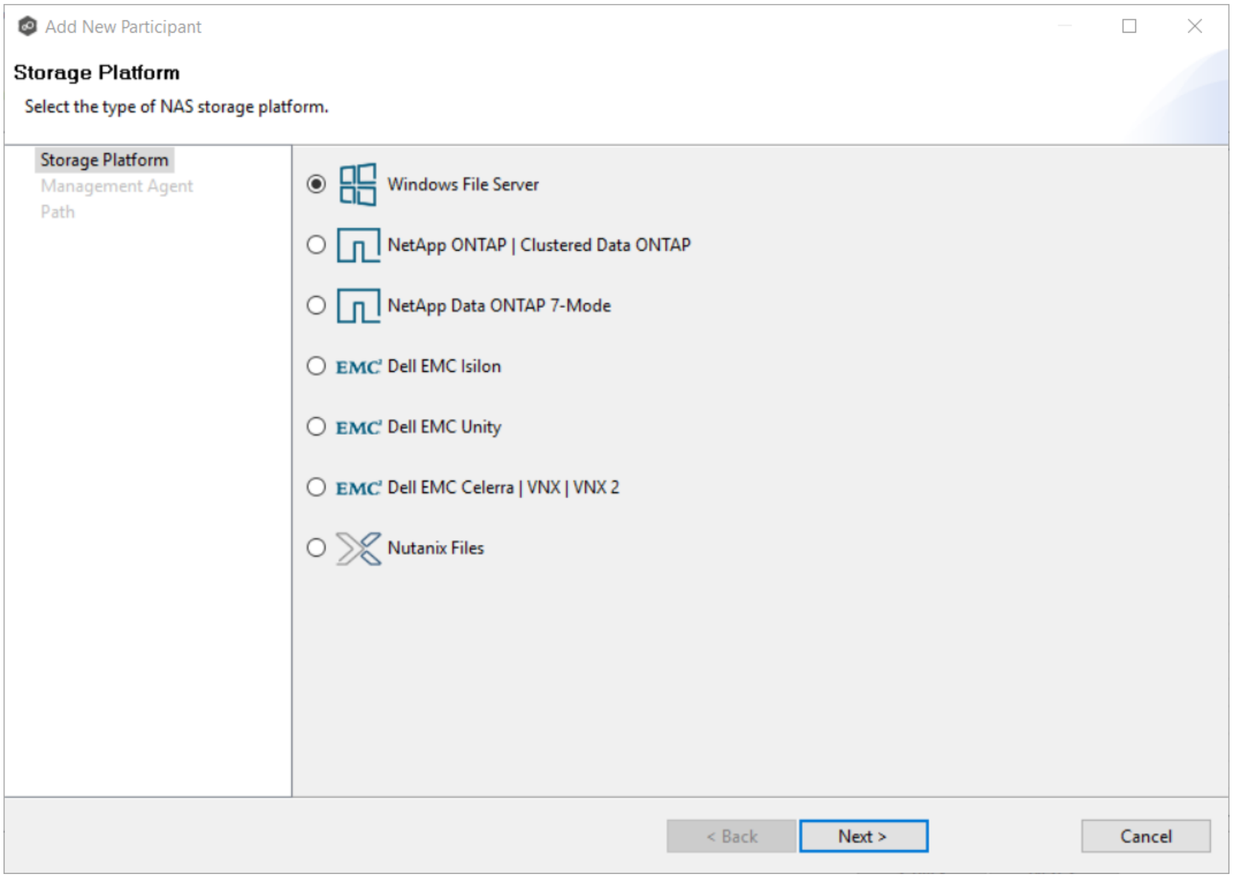

2. Click **Next**.

   The Management Agent page is displayed.

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 the Management Agent that manages the volumes/shares/folders you want to replicate in this job.

1. Select the Agent that manages the host.

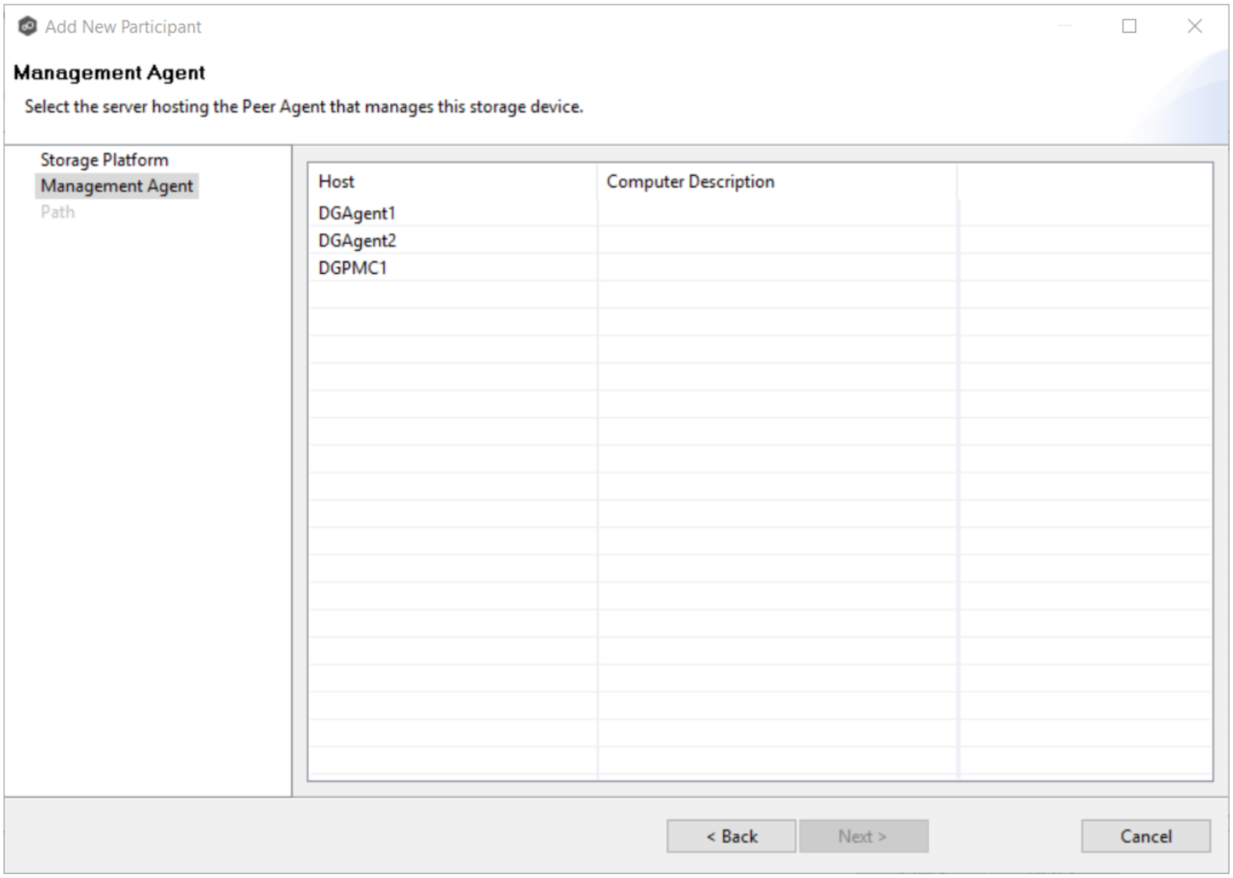

Management Agent

Select the server hosting the Peer Agent that manages this storage device.

Host	Computer Description
DGAgent1	
DGAgent2	
DGPMC1	

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 Storage Information page is displayed.

If you selected any storage platform type other than Windows File Server in the previous wizard page, the **Storage Information** page appears. It requests the credentials necessary to connect to the storage device you want to replicate. If you selected Windows Files Server in the previous wizard page, skip to Step 3: File Metadata.

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.

NetApp ONTAP | Clustered Data ONTAP

1. Enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

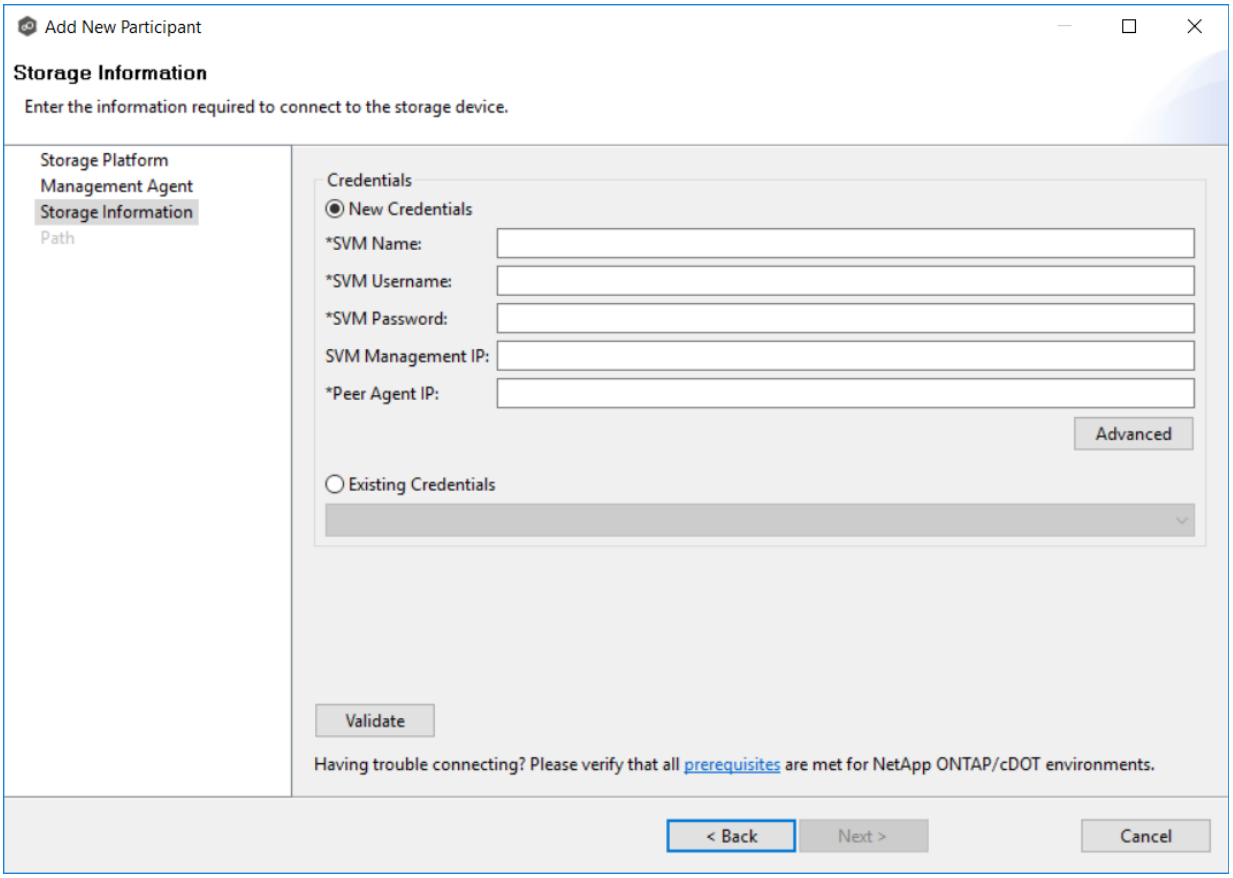

Add New Participant

Storage Information

Enter the information required to connect to the storage device.

Credentials

● New Credentials

*SVM Name:

*SVM Username:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

○ Existing Credentials

Having trouble connecting? Please verify that all prerequisites are met for NetApp ONTAP/cDOT environments.

2. Click **Next**.

   The Path page is displayed.

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	Enter the IP address of the server hosting the Agent that manages the Storage Virtual Machine.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

NetApp Data ONTAP 7-Mode

1. Enter the credentials to connect to the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

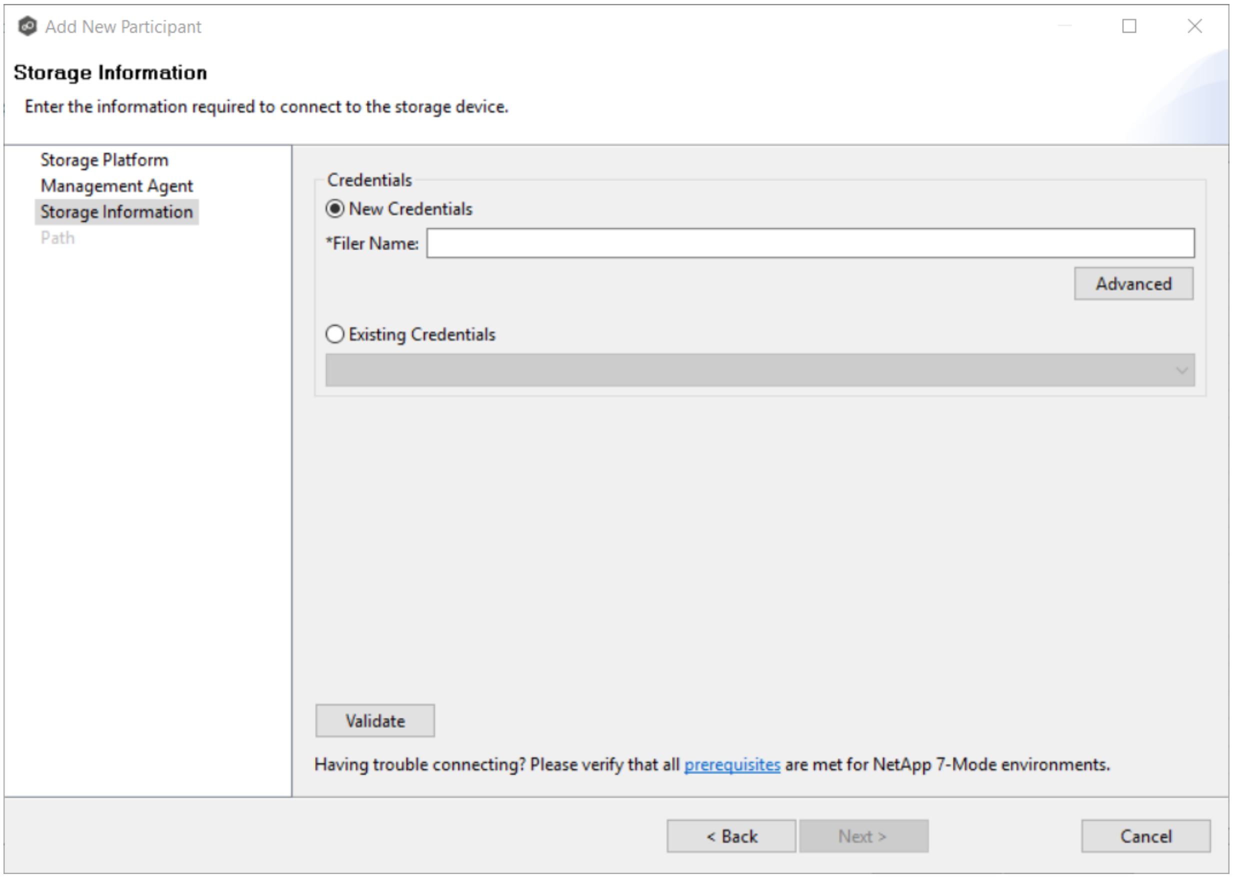

Add New Participant

Storage Information

Enter the information required to connect to the storage device.

Credentials

● New Credentials

*Filer Name:

○ Existing Credentials

Having trouble connecting? Please verify that all prerequisites are met for NetApp 7-Mode environments.

2. Click **Next**.

   The Path page is displayed.

Filer Name	Enter the name of the NetApp 7-Mode filer or vFiler hosting the data to be replicated.

Dell EMC Isilon

1. Enter the credentials to connect the EMC Isilon cluster hosting the data to be replicated.

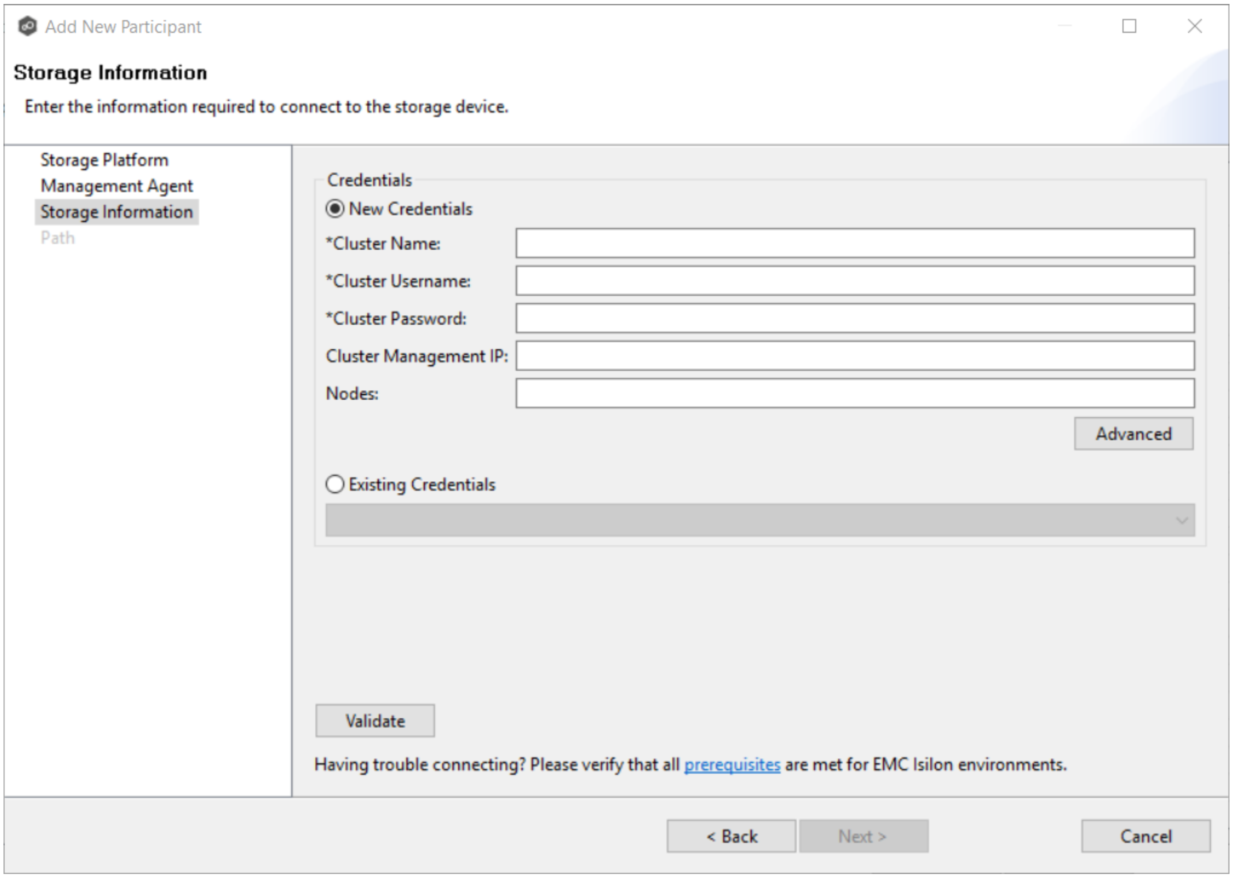

Add New Participant

Storage Information

Enter the information required to connect to the storage device.

Storage Platform
Management Agent
Storage Information
Path

Credentials

● New Credentials

*Cluster Name:

*Cluster Username:

*Cluster Password:

Cluster Management IP:

Nodes:

Advanced

○ Existing Credentials

Validate

Having trouble connecting? Please verify that all prerequisites are met for EMC Isilon environments.

< Back   Next >   Cancel

2. Click **Next**.

   The Path page is displayed.

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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

Dell EMC Unity

1. Enter the credentials to connect to the NAS Server hosting the data to be replicated.

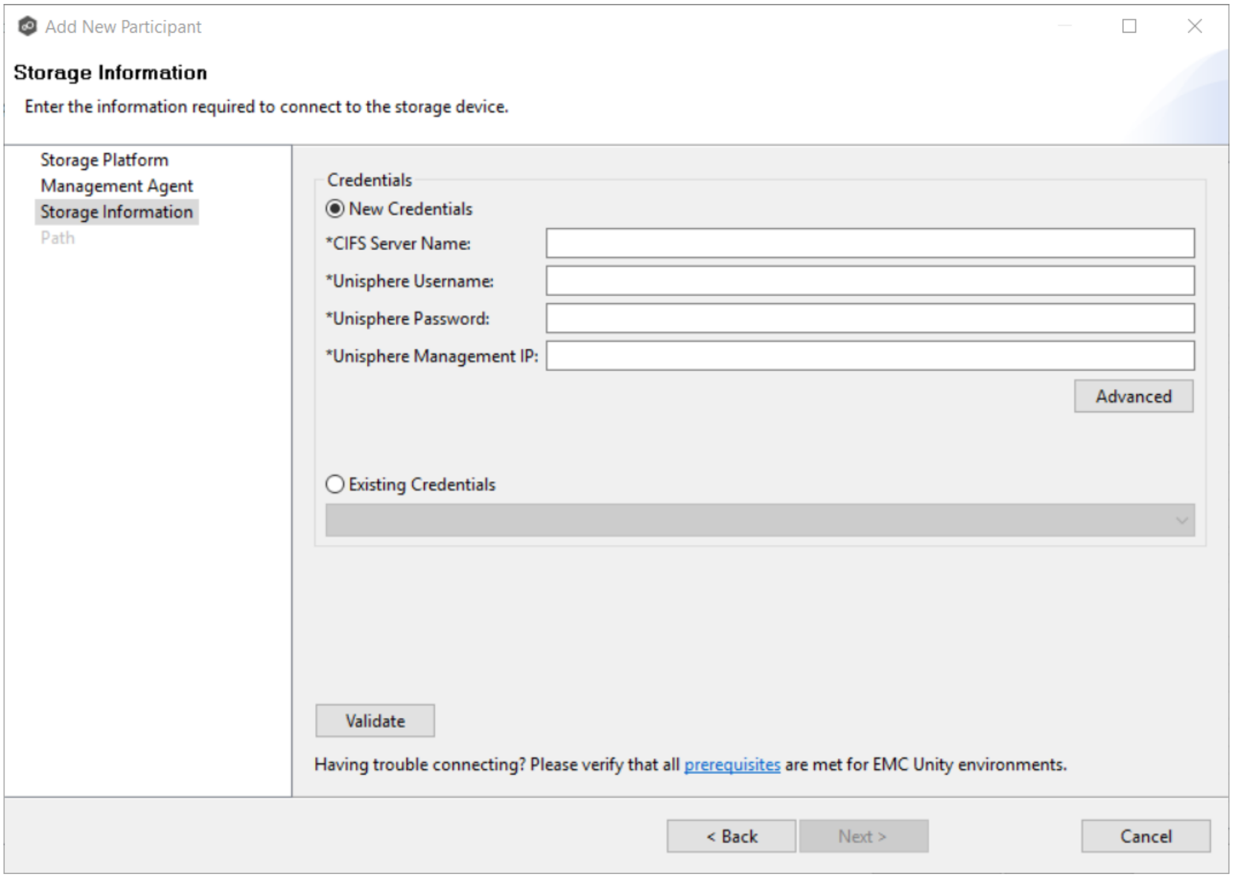

2. Click **Next**.

   The Path page is displayed.

NAS 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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

Dell EMC Celerra | VNX | VNX 2

1. Enter the credentials to connect to the CIFS Server hosting the data to be replicated.

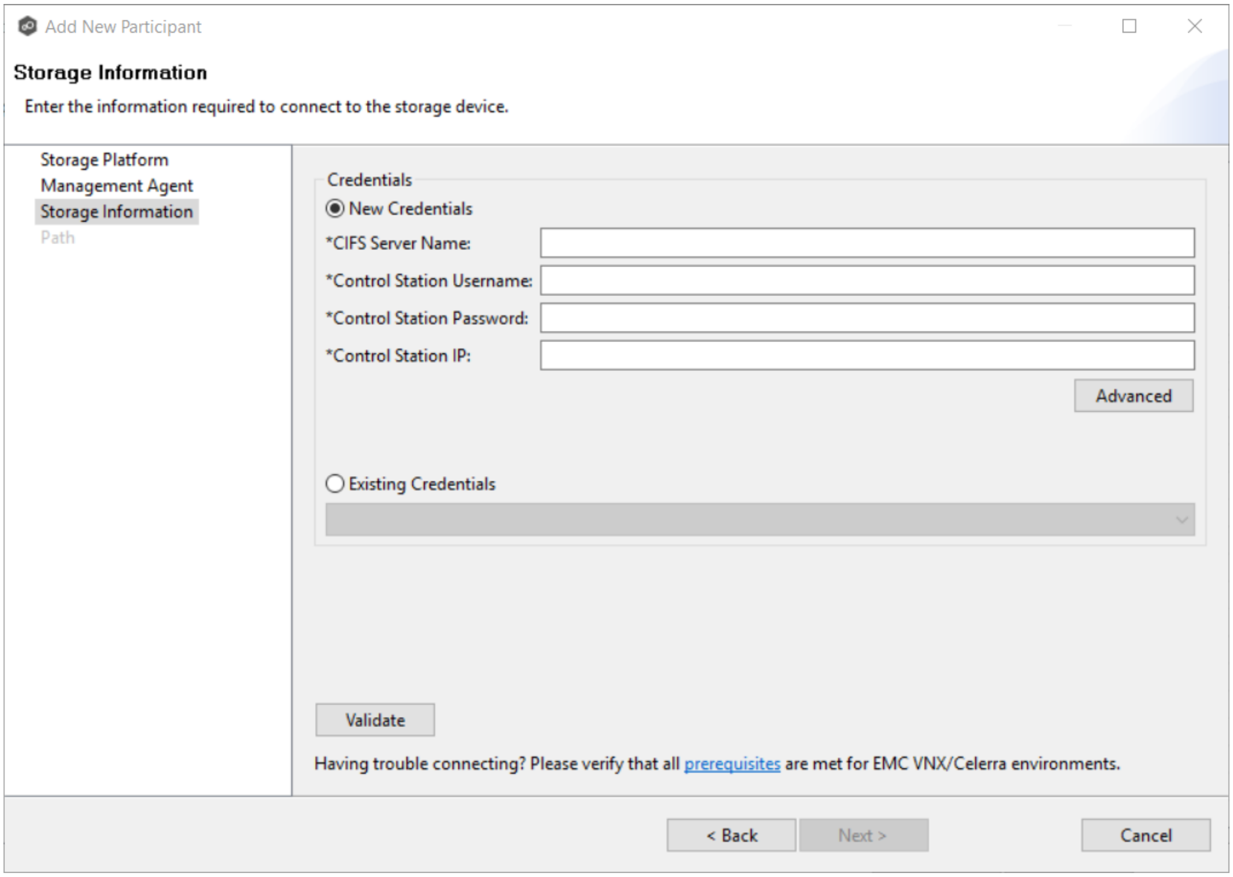

Add New Participant
Storage Information
Enter the information required to connect to the storage device.
Storage Platform
Management Agent
Storage Information
Path
Credentials
● New Credentials
*CIFS Server Name:
*Control Station Username:
*Control Station Password:
*Control Station IP:
Advanced
○ Existing Credentials
Validate
Having trouble connecting? Please verify that all prerequisites are met for EMC VNX/Celerra environments.
< Back
Next >
Cancel

2. Click **Next**.

   The Path page is displayed.

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.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

Nutanix Files

1. Enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated.

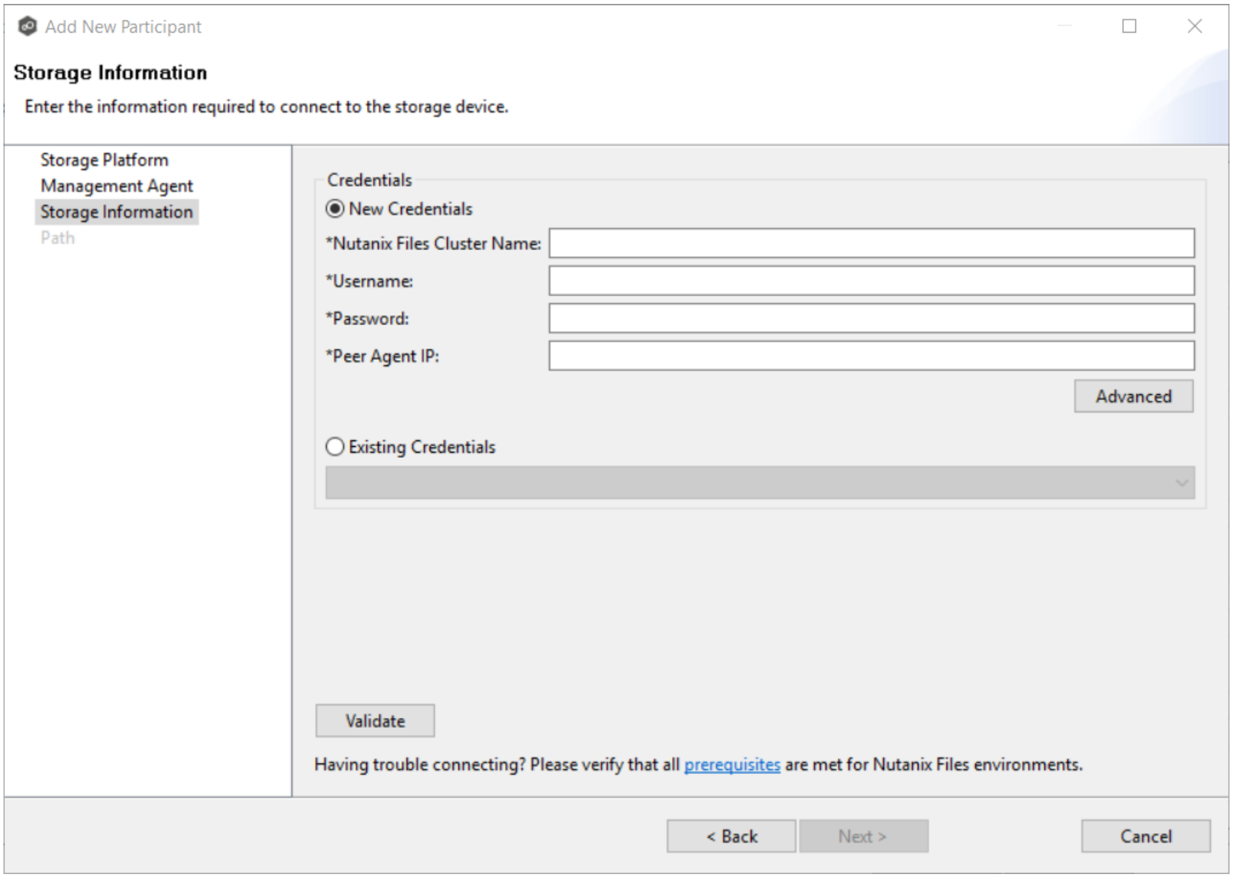

Having trouble connecting? Please verify that all prerequisites are met for Nutanix Files environments.

2. Click **Next**.

   The Path page is displayed.

Nutanix Files Cluster 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 AFS cluster via its management APIs.
Password	Enter the password for the account managing the AFS cluster via its management APIs.
Peer Agent IP	Enter the IP address of the Agent server used to manage the storage platform. This should not point to the AFS cluster itself.

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.

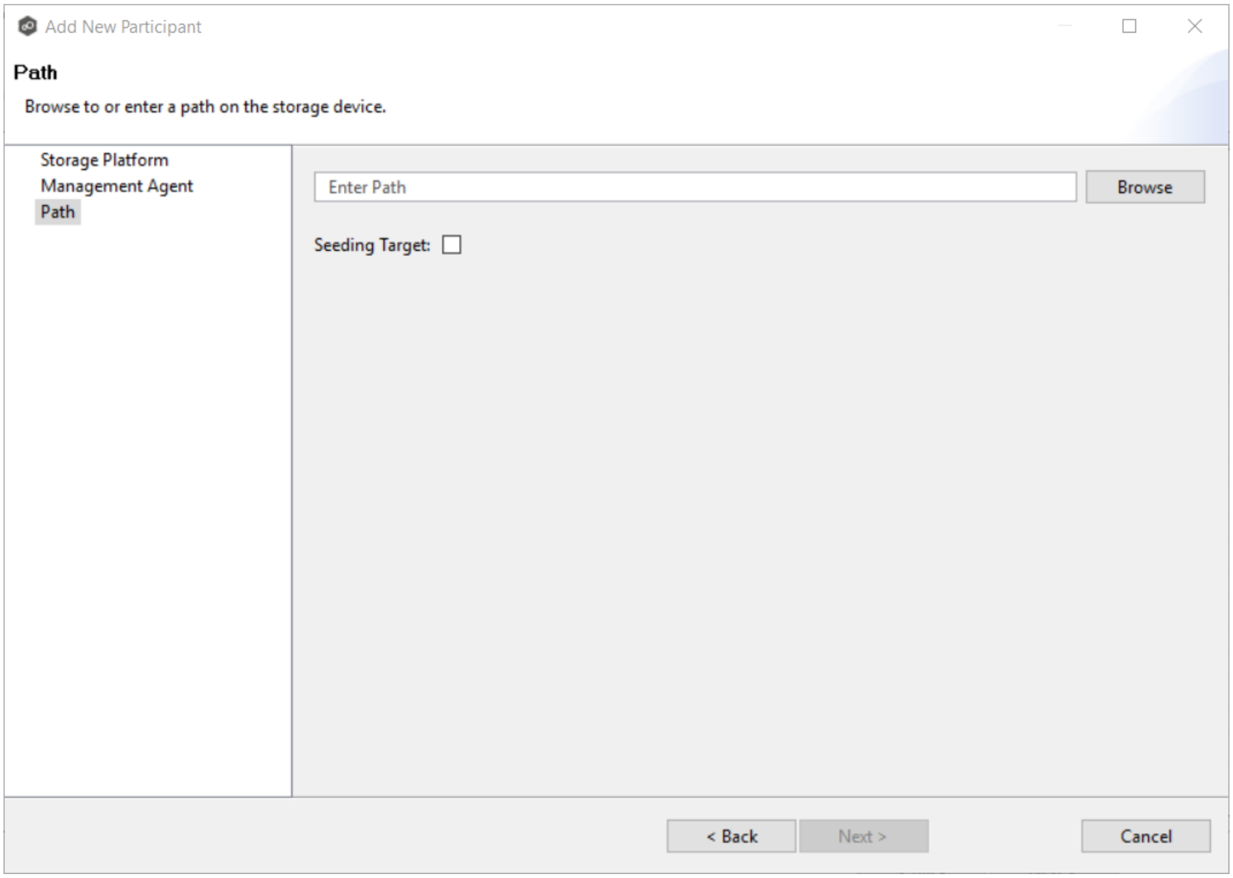

Add New Participant

Path

Browse to or enter a path on the storage device.

If you selected **Browse**, the **Folder Browser** dialog appears:

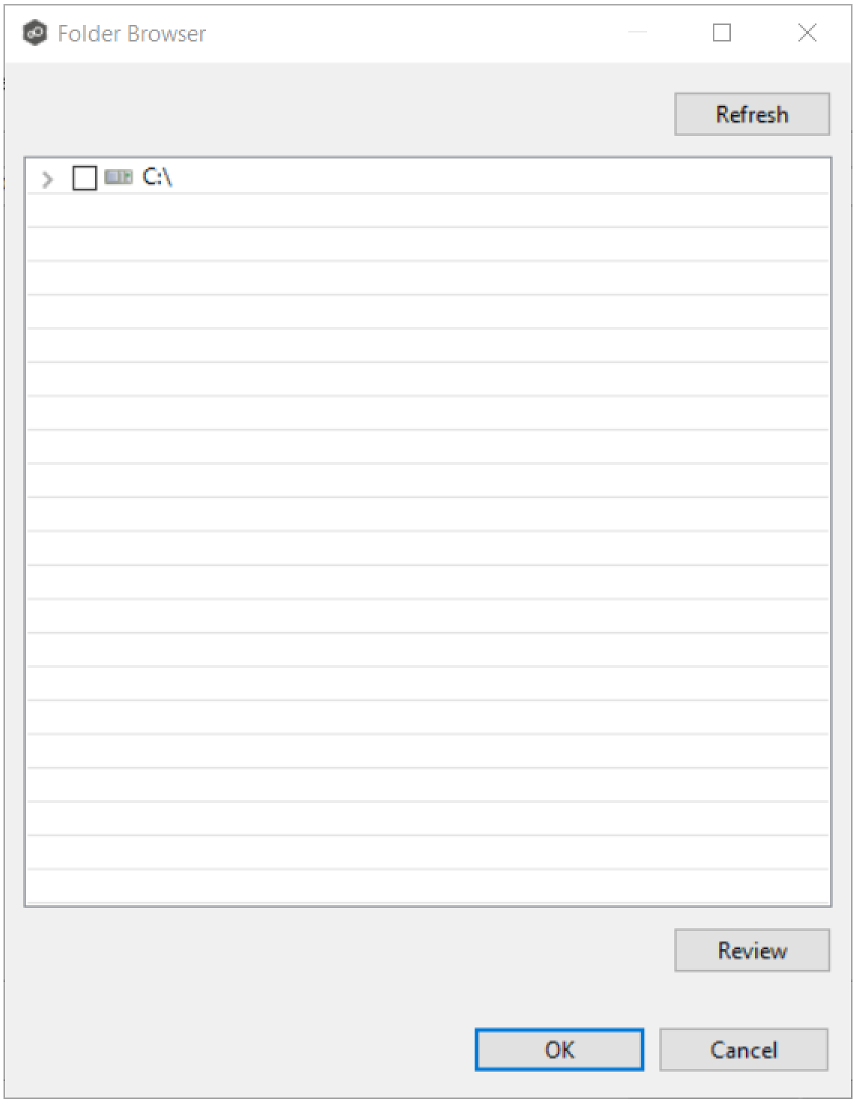

a. Expand the folder tree.

b. Select the appropriate volume/share/folder/

c. (Optional) Click the **Review** button to see your selection.

d. Click **OK**.

2. Click **Finish** to complete the wizard for this participant.

3. 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 of the participants, continue with Step 3: File Metadata.

Step 3: 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 of the 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 (DACL, SACL, Owner.) 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 (DACL, SACL, Owner) will be synchronized during the initial scan.

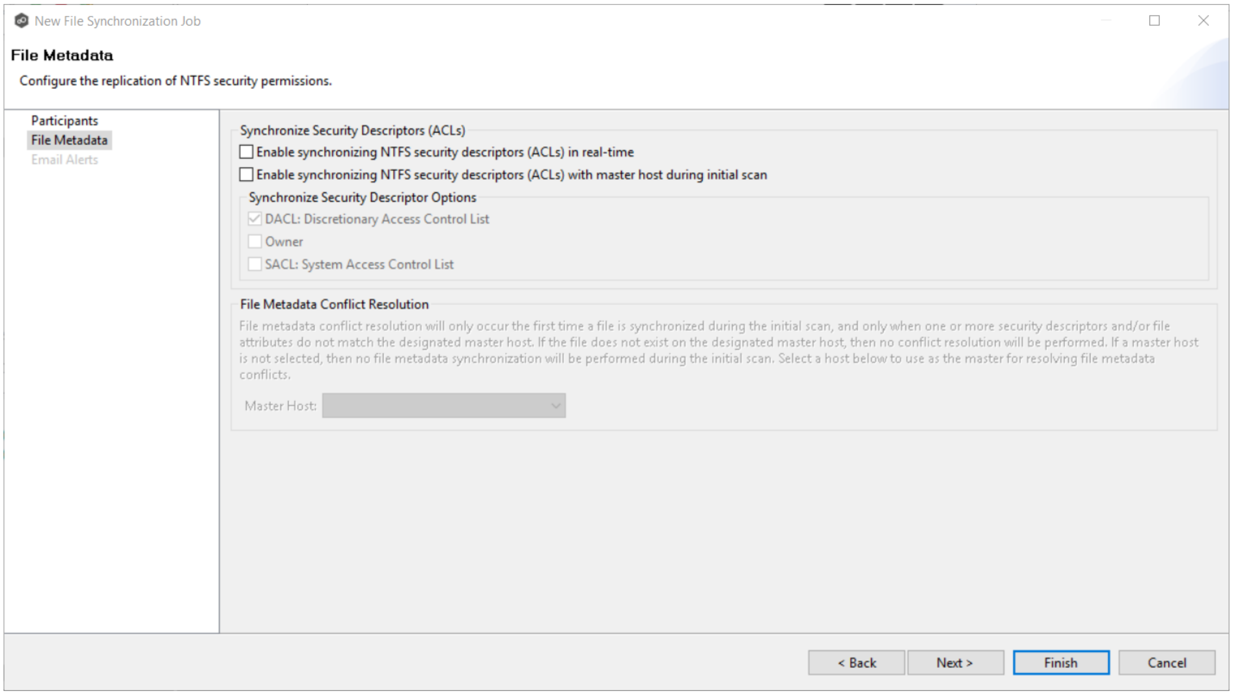

New File Synchronization Job

File Metadata

Configure the replication of NTFS security permissions.

Synchronize Security Descriptors (ACLs)

□ Enable synchronizing NTFS security descriptors (ACLs) in real-time

□ Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan

Synchronize Security Descriptor Options

☑ DACL: Discretionary Access Control List

□ Owner

□ SACL: System Access Control List

File Metadata Conflict Resolution

File metadata conflict resolution will only occur the first time a file is synchronized during the initial scan, and only when one or more security descriptors and/or file attributes do not match the designated master host. If the file does not exist on the designated master host, then no conflict resolution will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan. Select a host below to use as the master for resolving file metadata conflicts.

Master Host:

2. Click **OK** in the informational dialog that appears after selecting a metadata option.

3. Select which security descriptor components (DACL, SACL and Owner) are synchronized.

   In general, you will usually need to synchronize only DACLs. If you need to synchronize SACLs or Owner, then the user that a Peer Agent service is run under on each participating host must have permission to read and write SACLs and Owner.

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.

   Conflict resolution for file metadata 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 will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan.

5. Click **Next**.

   The Email Alerts page is displayed.

Step 4: 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 thenbe 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.

Apply an Existing Email Alert

To apply an existing email alert to the job.

1. Click the **Select** button.

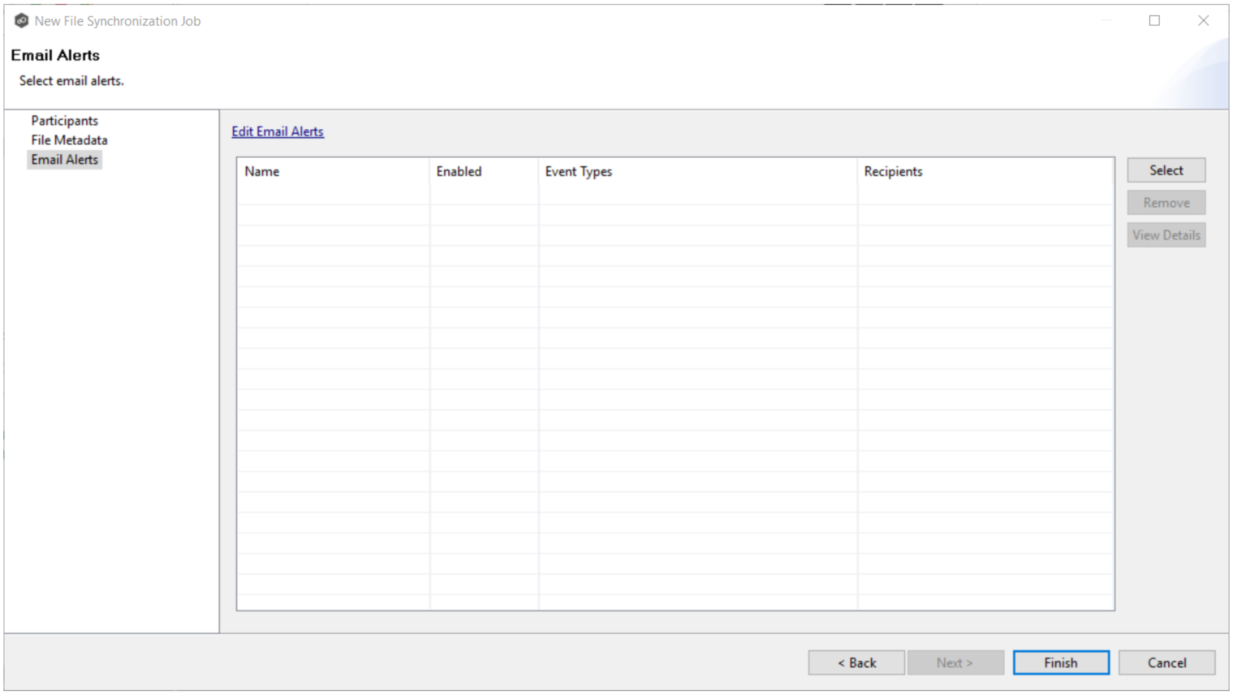

New File Synchronization Job

Email Alerts

Select email alerts.

Name	Enabled	Event Types	Recipients

The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.

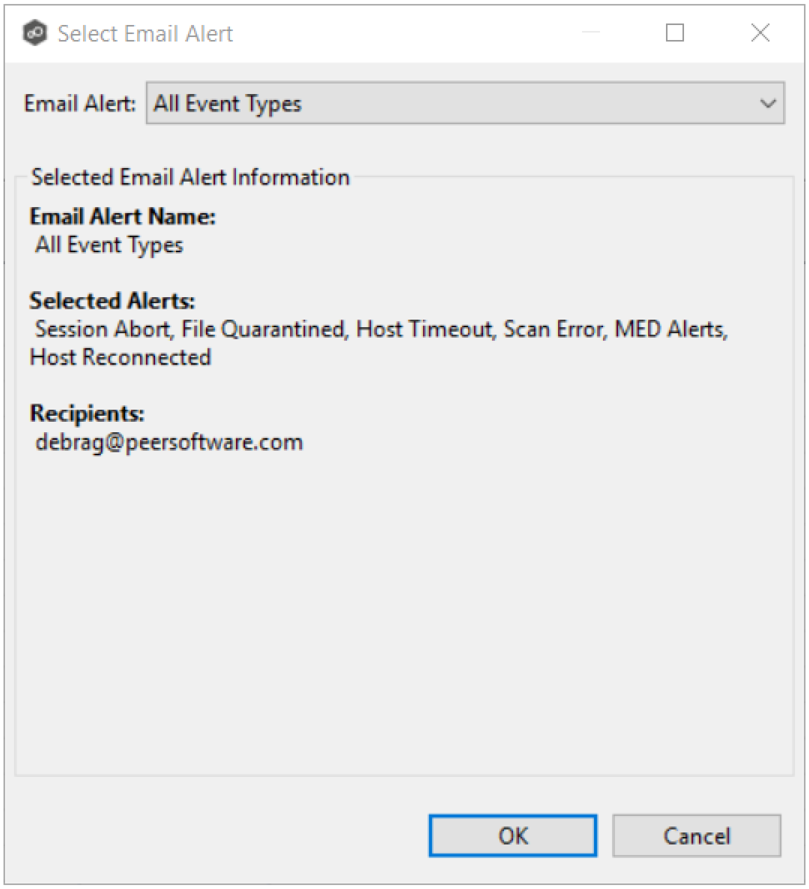

Email Alert: All Event Types

Selected Email Alert Information

Email Alert Name:
All Event Types

Selected Alerts:
Session Abort, File Quarantined, Host Timeout, Scan Error, MED Alerts, Host Reconnected

Recipients:
debrag@peersoftware.com

3. Click **OK**.

   The alert is listed on the **Email Alerts** page.

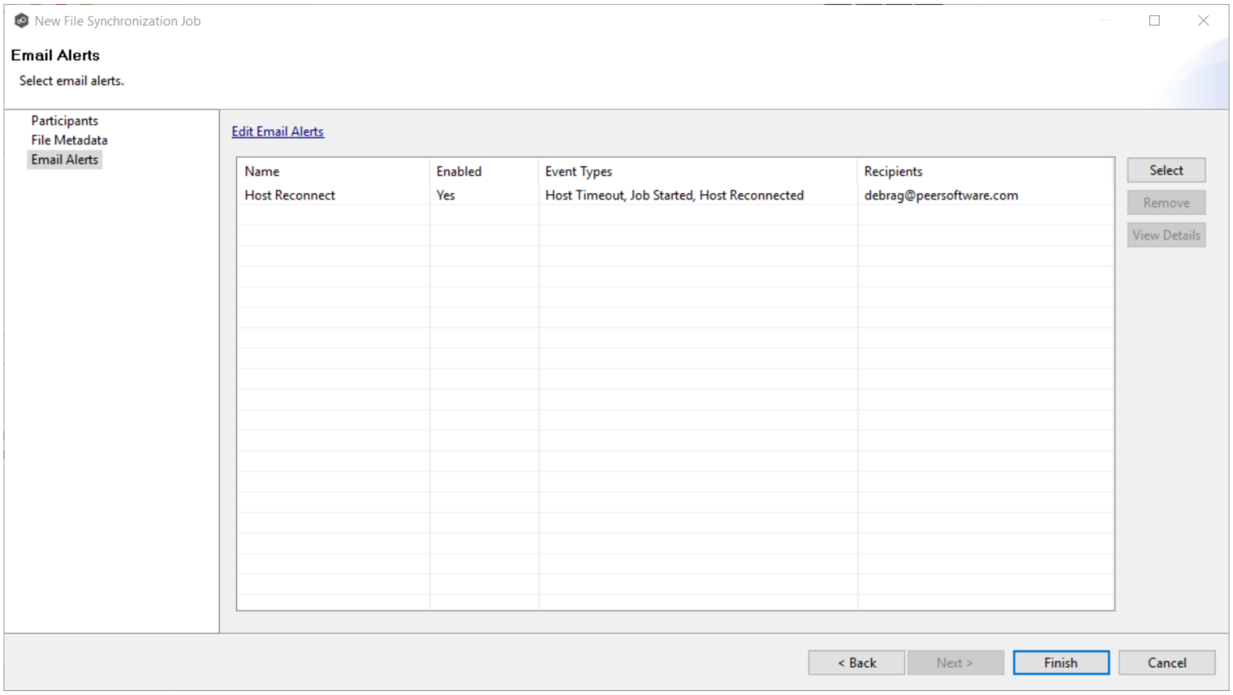

Name	Enabled	Event Types	Recipients
Host Reconnect	Yes	Host Timeout, Job Started, Host Reconnected	debrag@peersoftware.com

4. (Optional) Repeat steps 1-3 to apply additional alerts.

5. After applying email alerts, continue to Step 5: Save Job.

Create an Email Alert

To create an alert or edit an existing alert from the **Email Alerts** page of the **Create Job** wizard:

1. Click the **Edit Email Alerts** link.

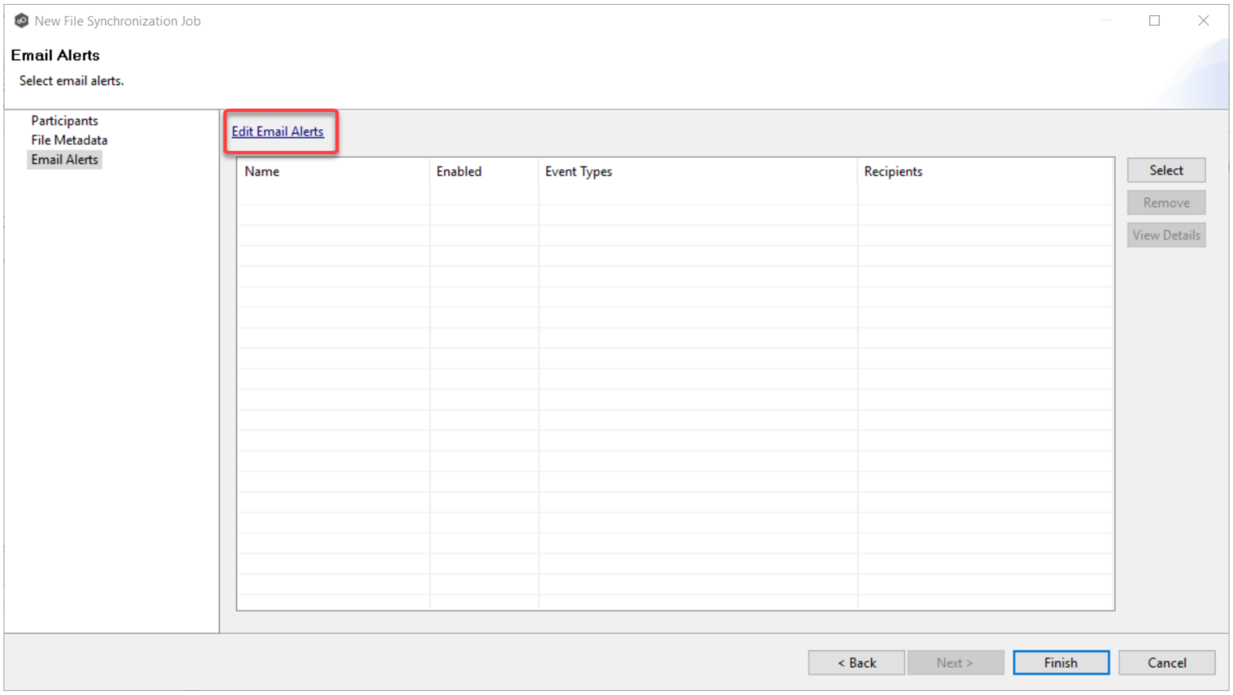

The **Collaboration Email Alerts** dialog appears.

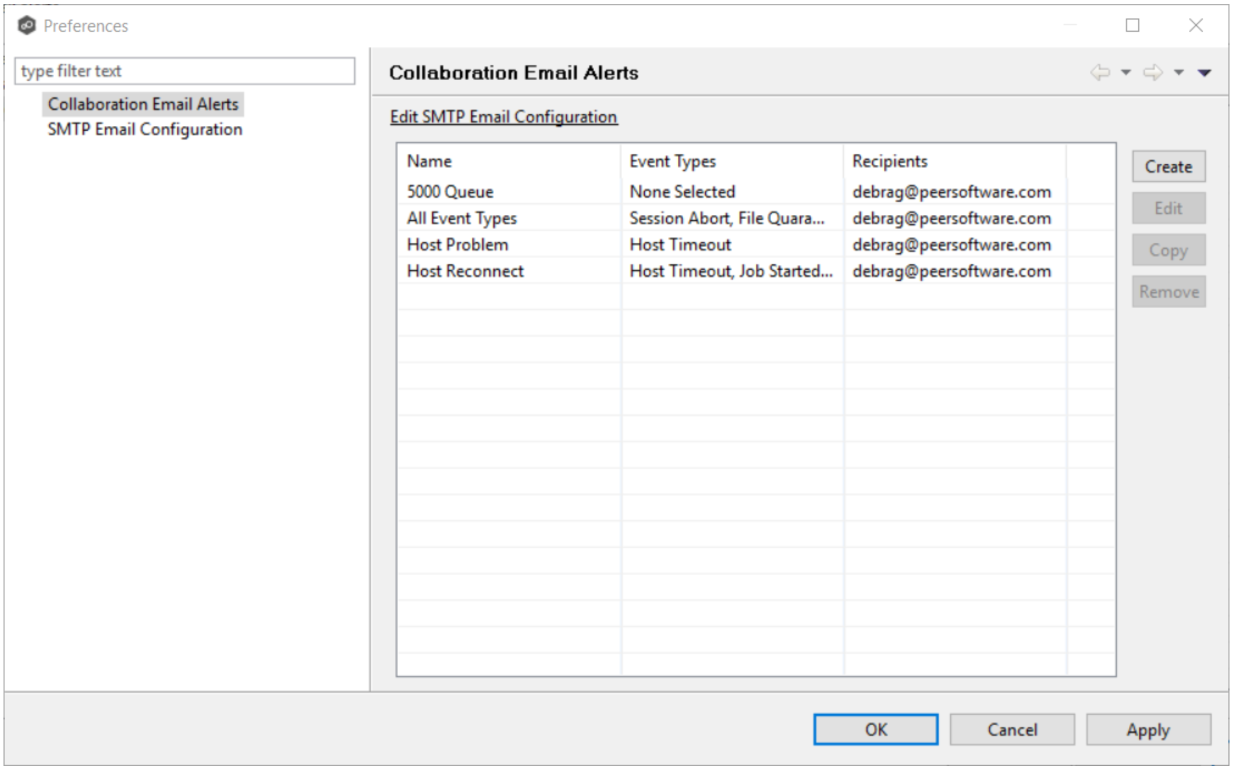

Name	Event Types	Recipients
5000 Queue	None Selected	debrag@peersoftware.com
All Event Types	Session Abort, File Quara...	debrag@peersoftware.com
Host Problem	Host Timeout	debrag@peersoftware.com
Host Reconnect	Host Timeout, Job Started...	debrag@peersoftware.com

2. Click **Create**.

The **Create Email Alert** dialog appears.

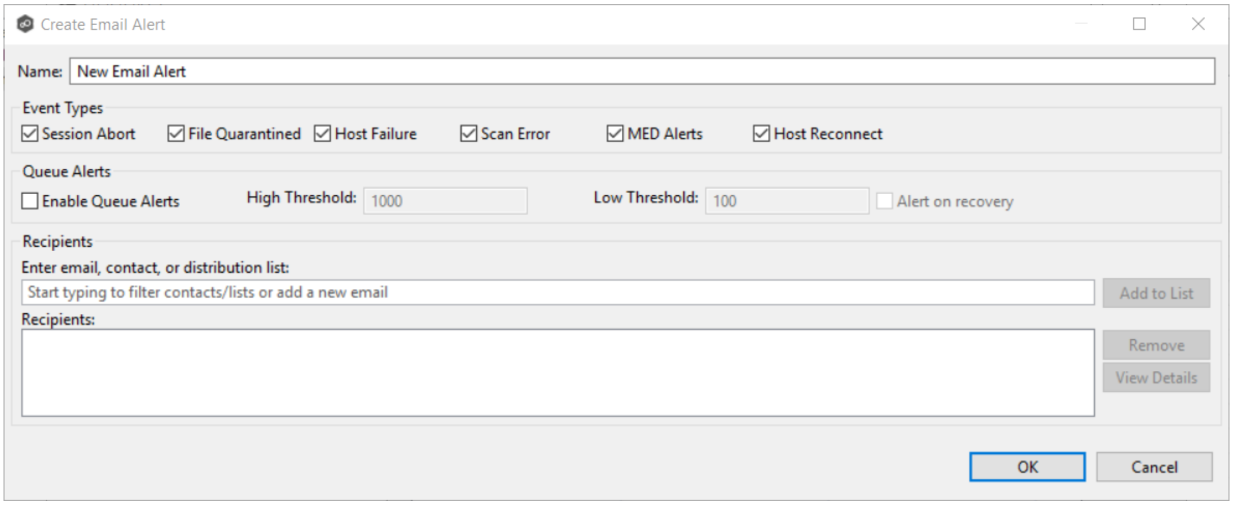

Create Email Alert

Name: New Email Alert

Event Types

☑ Session Abort ☑ File Quarantined ☑ Host Failure ☑ Scan Error ☑ MED Alerts ☑ Host Reconnect

Queue Alerts

□ Enable Queue Alerts High Threshold: 1000 Low Threshold: 100 □ Alert on recovery

Recipients

Enter email, contact, or distribution list:

Start typing to filter contacts/lists or add a new email

Recipients:

3. Continue with Step 4 of the instructions in [Email Alerts]() in the **Preferences** section.

Step 5: Save Job

Now that you have completed the first four 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 Synchronization job. It is now listed in the **Jobs** view under **File Synchronization** and a job run-time view appears in the **Summaries** area. You can start the job from either place.

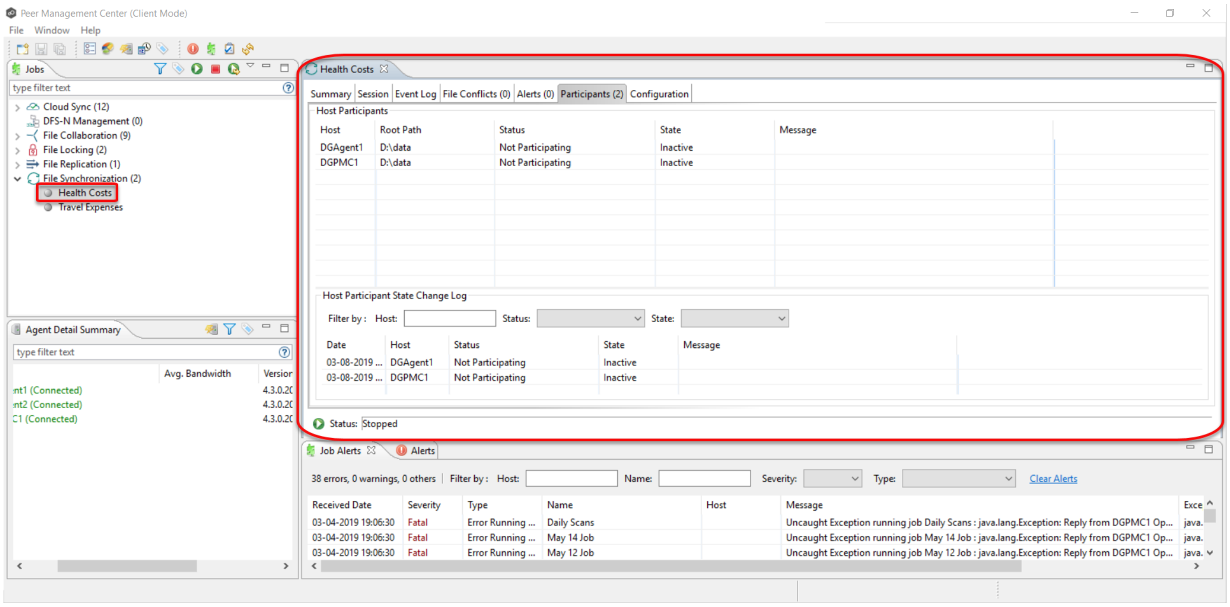

	Avg. Bandwidth	Versior
:nt1 (Connected)		4.3.0.2(
:nt2 (Connected)		4.3.0.2(
C1 (Connected)		4.3.0.2(

Host Participants

Host	Root Path	Status	State	Message
DGAgent1	D:\data	Not Participating	Inactive	
DGPMC1	D:\data	Not Participating	Inactive	

Host Participant State Change Log

Filter by : Host: Status: State:

Date	Host	Status	State	Message
03-08-2019 ...	DGAgent1	Not Participating	Inactive	
03-08-2019 ...	DGPMC1	Not Participating	Inactive	

Status: Stopped

38 errors, 0 warnings, 0 others | Filter by : Host: Name: Severity: Type:

Received Date	Severity	Type	Name	Host	Message	Exce
03-04-2019 19:06:30	Fatal	Error Running ...	Daily Scans		Uncaught Exception running job Daily Scans : java.lang.Exception: Reply from DGPMC1 Op...	java.
03-04-2019 19:06:30	Fatal	Error Running ...	May 14 Job		Uncaught Exception running job May 14 Job : java.lang.Exception: Reply from DGPMC1 Op...	java.
03-04-2019 19:06:30	Fatal	Error Running ...	May 12 Job		Uncaught Exception running job May 12 Job : java.lang.Exception: Reply from DGPMC1 Op...	java.

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
- Delta Replication
- DFS-N
- File Filters
- File Locking
- General
- Logging and Alerts

- [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.

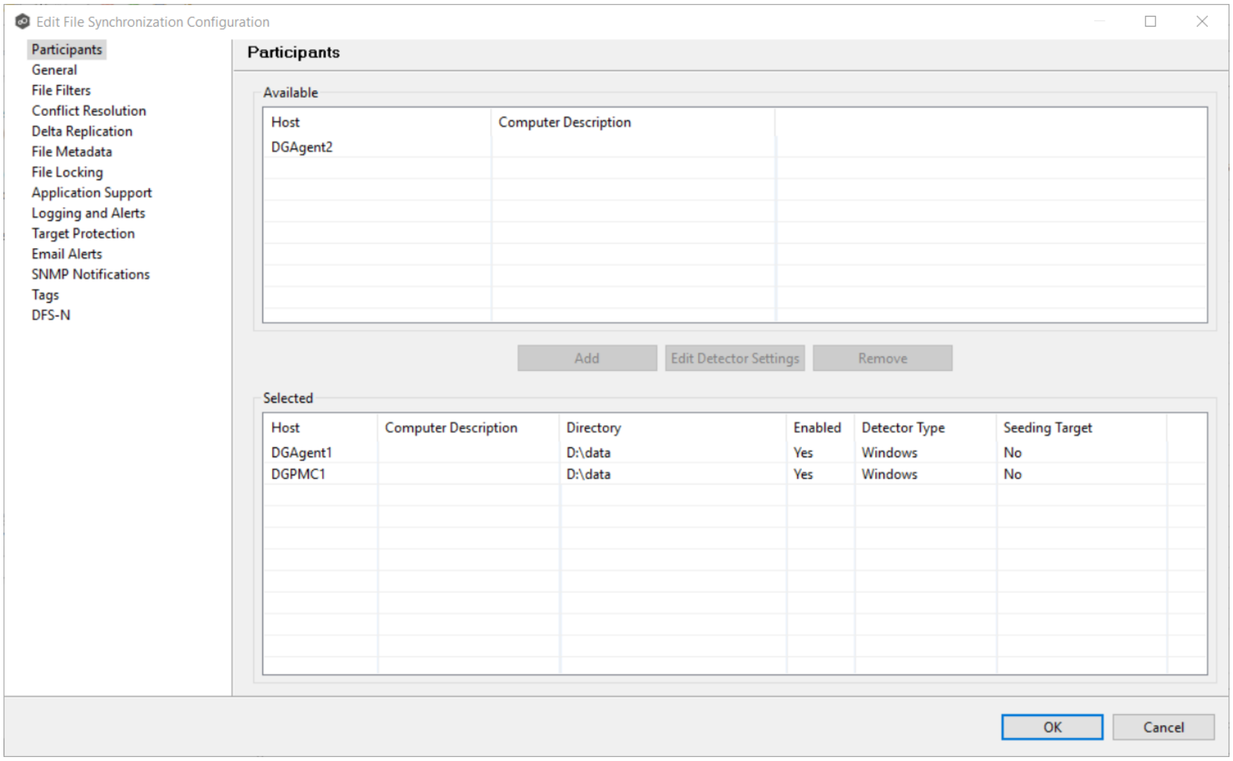

Participants

Available

Host	Computer Description
DGAgent2	

Selected

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent1		D:\data	Yes	Windows	No
DGPMC1		D:\data	Yes	Windows	No

3. Select a configuration item in the navigation tree and make the desired changes:

   - Participants
   - General

- File Filters
- 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 Configuration** wizard allows you to:

- Add and remove participants from a job.
- Modify a participant's attributes.
- Modify a participant's detector settings.

The **Participants** page in the **Edit File Synchronization Configuration** wizard has two tables: the **Available** table and the **Selected** table. The **Available** table lists the available hosts and the **Selected** table lists hosts that have already been added to the job. The **Computer Description** field displays the name of the server that the Peer Agent is running on.

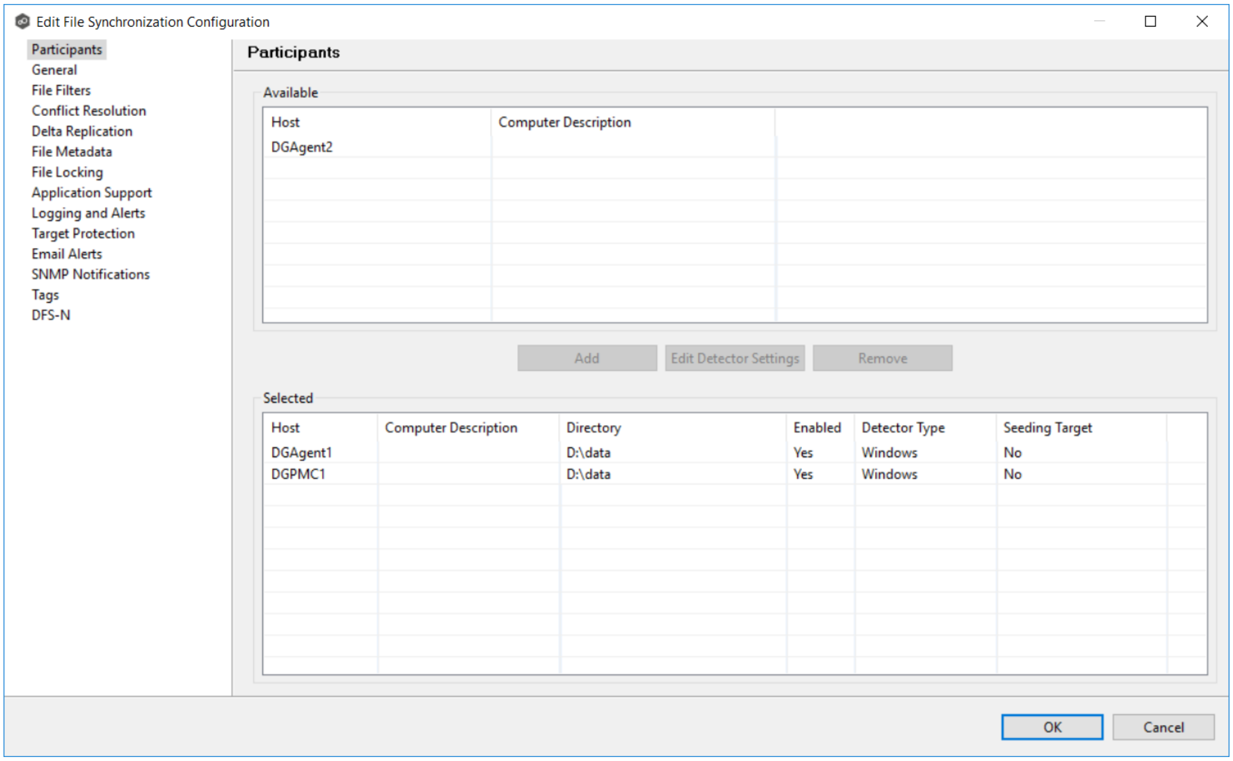

Host	Computer Description
DGAgent2	

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent1		D:\data	Yes	Windows	No
DGPMC1		D:\data	Yes	Windows	No

This topic describes adding and removing participants in a File Synchronization job.

Adding a Participant

To add a participant:

1. Click the participant in the **Available** table.

   To be available, a host must have Peer Agent installed and successfully connect to the Peer Management Broker. If a particular host is not displayed in the list, try restarting the Peer Agent Windows Service on that host, and if it successfully connects to the Peer Management Center Broker, then the list will be updated with the computer name of that host.

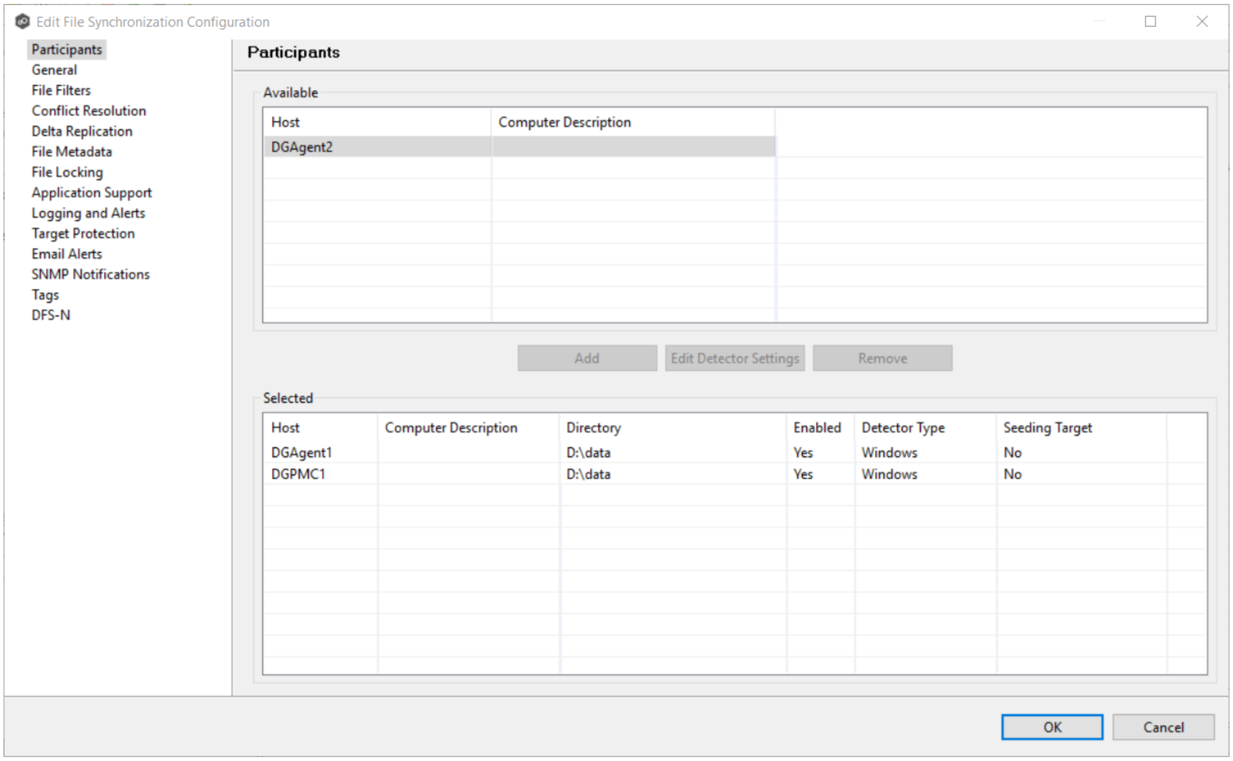

Host	Computer Description
DGAgent2	

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent1		D:\data	Yes	Windows	No
DGPMC1		D:\data	Yes	Windows	No

2. Click the **Add** button.

   The participant is moved to the **Selected** table.

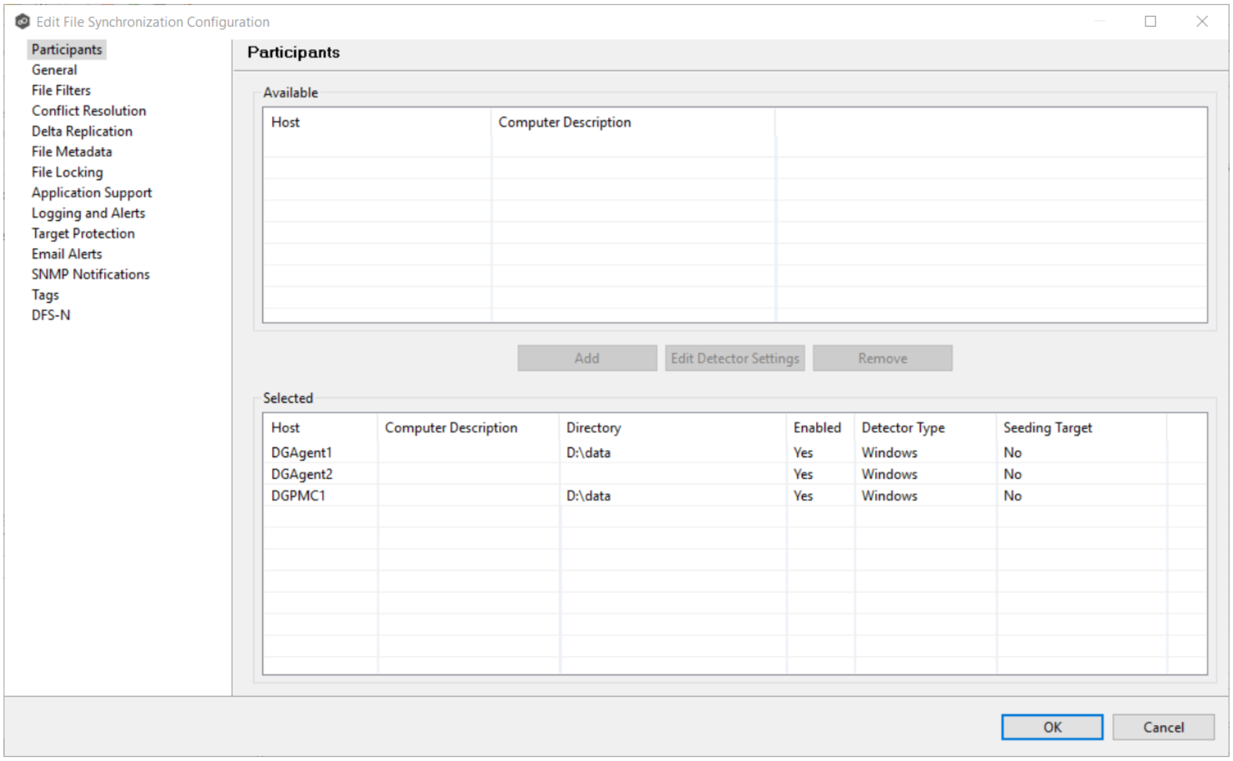

Host	Computer Description

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent1		D:\data	Yes	Windows	No
DGAgent2			Yes	Windows	No
DGPMC1		D:\data	Yes	Windows	No

3. (Optional) Enter the computer's name in the **Computer Description** column.

4. Enter the path to the folder to be watched in the **Directory** column.

5. (Optional) Modify whether the participant is a seeding target.

6. (Optional) Modify the participant's detector settings.

7. Click **OK** or select another item to modify.

Removing a Participant

To remove a participant:

1. Click the participant in the **Selected** table.

2. Click the **Remove** button.

   The participant is moved to the **Available** table.

   Note: A File Synchronization job must have at least two participants, so if after removing a participant, there is only a single participant, you must add another participant to the job.

3. Click **OK** or select another item to modify.

For more information on smart data seeding, see Smart Data Seeding in Advanced Topics or contact support@peersoftware.com.

To set a host as a smart data seeding target:

1. Select the host in the **Selected** table.

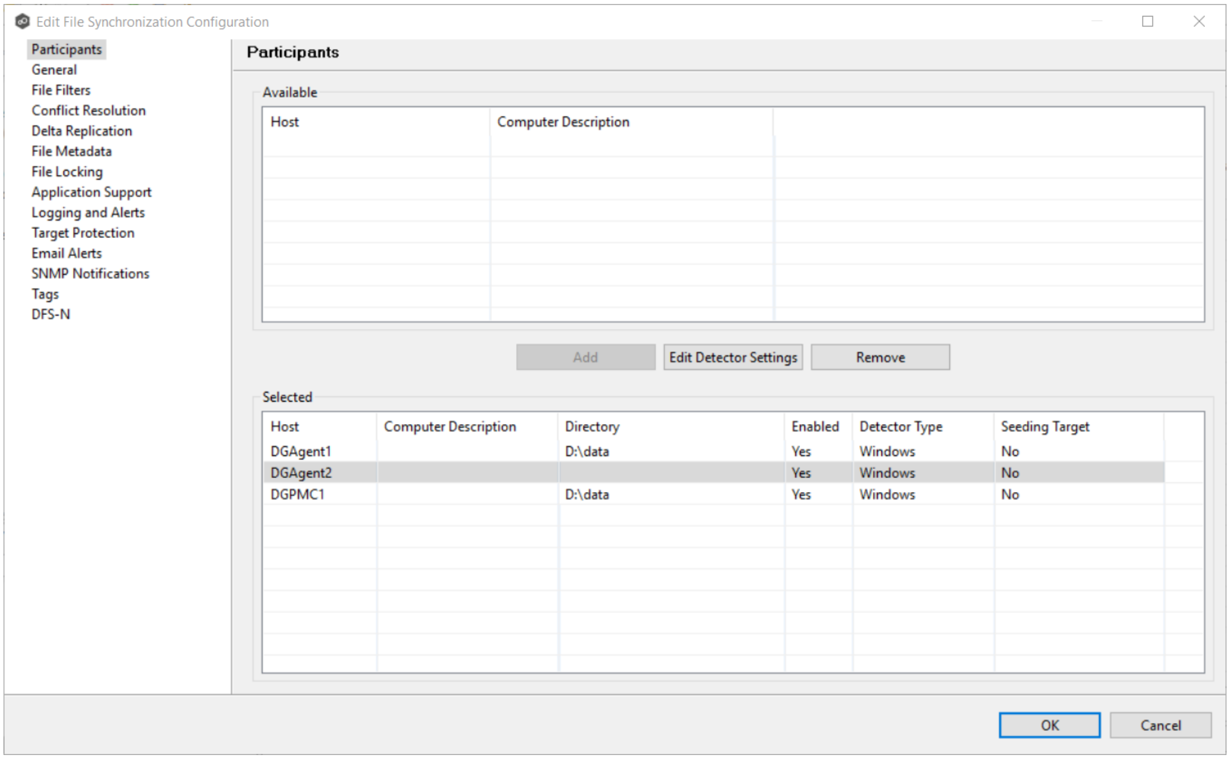

Participants

Available

Host	Computer Description

Selected

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent1		D:\data	Yes	Windows	No
DGAgent2			Yes	Windows	No
DGPMC1		D:\data	Yes	Windows	No

2. Select **Yes** in the **Seeding Target** column.

3. Review the information in the message dialog that appears:

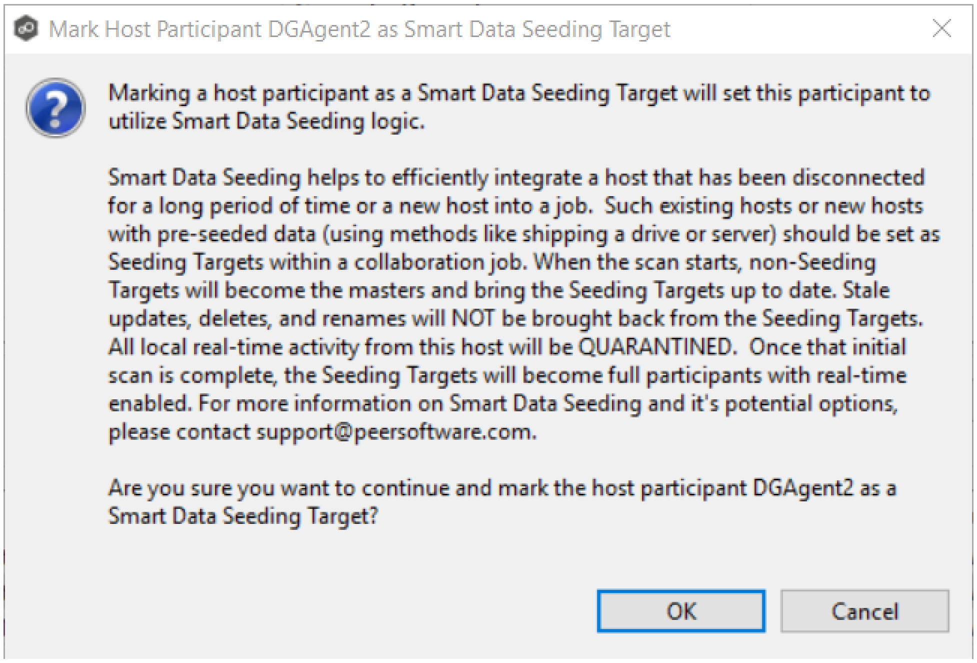

Mark Host Participant DGAgent2 as Smart Data Seeding Target

Marking a host participant as a Smart Data Seeding Target will set this participant to utilize Smart Data Seeding logic.

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 job. Such existing hosts or new hosts with pre-seeded data (using 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 from this host 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 it's potential options, please contact support@peersoftware.com.

Are you sure you want to continue and mark the host participant DGAgent2 as a Smart Data Seeding Target?

4. Click **OK**.

The value in the **Seeding Target** column is updated.

In addition to global real-time detection options that apply to all jobs, you can set additional detection-related options for a specific File Synchronization job. For example, you can exclude real-time events by certain users. This is helpful if you are trying to prevent events generated from backup and/or archival tools from triggering activity.

To modify the detector settings for a host:

1. Select the host in the **Selected** table.

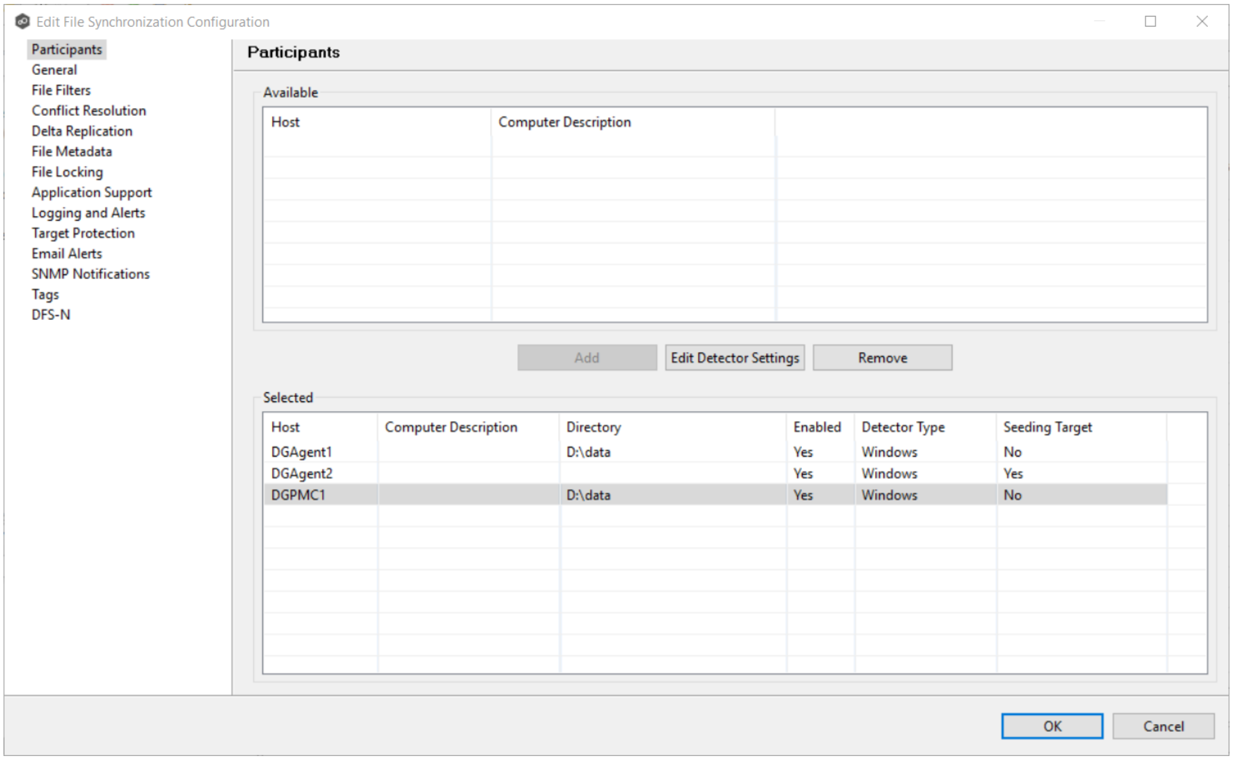

Host	Computer Description

Host	Computer Description	Directory	Enabled	Detector Type	Seeding Target
DGAgent1		D:\data	Yes	Windows	No
DGAgent2			Yes	Windows	Yes
DGPMC1		D:\data	Yes	Windows	No

2. Click **Edit Detector Settings**.

   The information you are prompted to enter varies, depending on the type of storage platform.

3. Modify the values as needed.

4. Click **OK**.

General

The **General** page in the **Edit File Synchronization Configuration** wizard presents miscellaneous settings pertaining to a File Synchronization job. You may want to consult with Peer Software's support team before modifying these values.

To modify these settings:

1. Enter the values recommended by Peer Software Support.

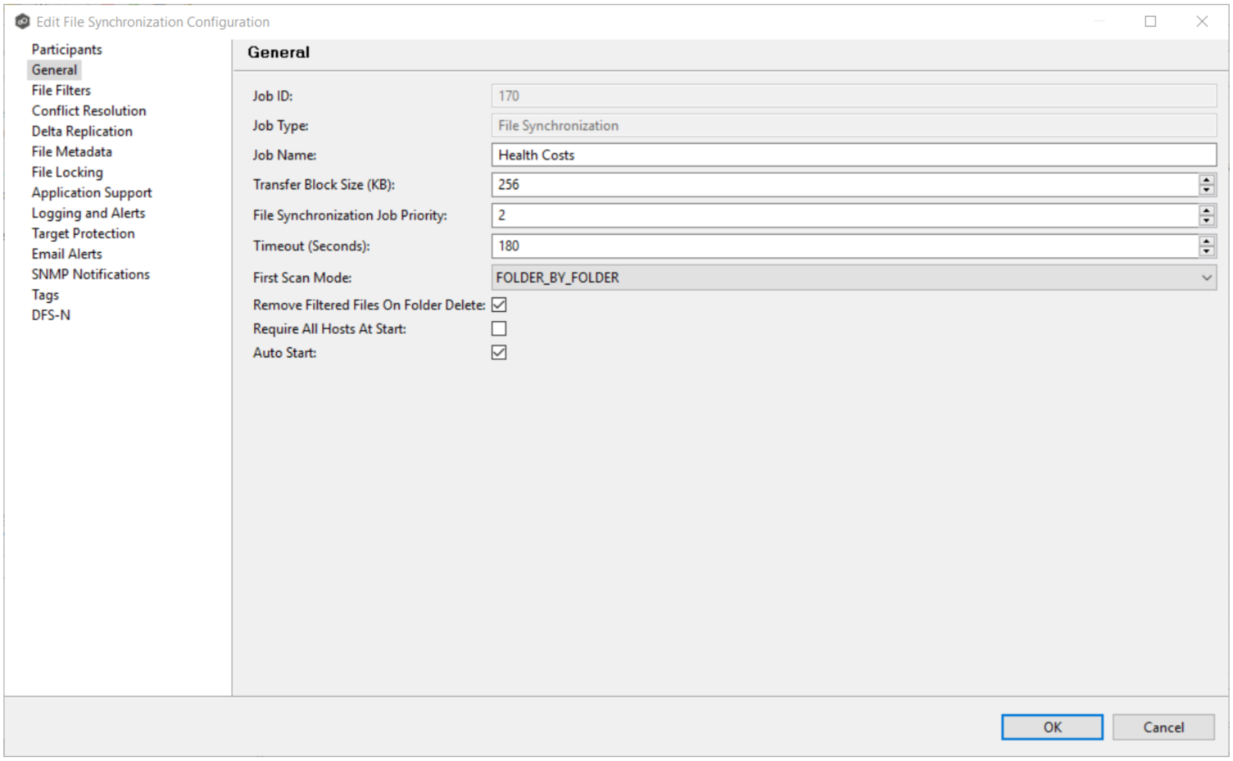

General

Job ID: 170
Job Type: File Synchronization
Job Name: Health Costs
Transfer Block Size (KB): 256
File Synchronization Job Priority: 2
Timeout (Seconds): 180
First Scan Mode: FOLDER_BY_FOLDER
Remove Filtered Files On Folder Delete: ☑
Require All Hosts At Start: □
Auto Start: ☑

2. Click **OK**.

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.

Job ID	Unique, system-generated job identifier that cannot be edited.
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.
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 results 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.

File Filters

This topic describes how to apply a file filter to a File Synchronization job. File filters are defined at a global level in Preferences and can then be applied to individual jobs.

- For a description of how to create a file filter, see File Filters in the **Preferences** section for File Synchronization jobs.

- For a discussion of how file filters work, see File and Folder Filters in the **Basic Concepts** section.

The **File Filters** page in the **Edit File Synchronization Configuration** wizard allows you to select which file filters to apply to a File Synchronization job. When the job is run, each selected filter is combined into one large filter (by combining all exclusions and inclusions together). In general, you should have at least one default global file filter that is applied to all jobs and possibly other file filters that apply to specific jobs. However, for most environments, only a single default global file filter is necessary.

To modify which file filters are applied to a File Synchronization job:

1. Select the checkbox next to the filters you want applied to the job.

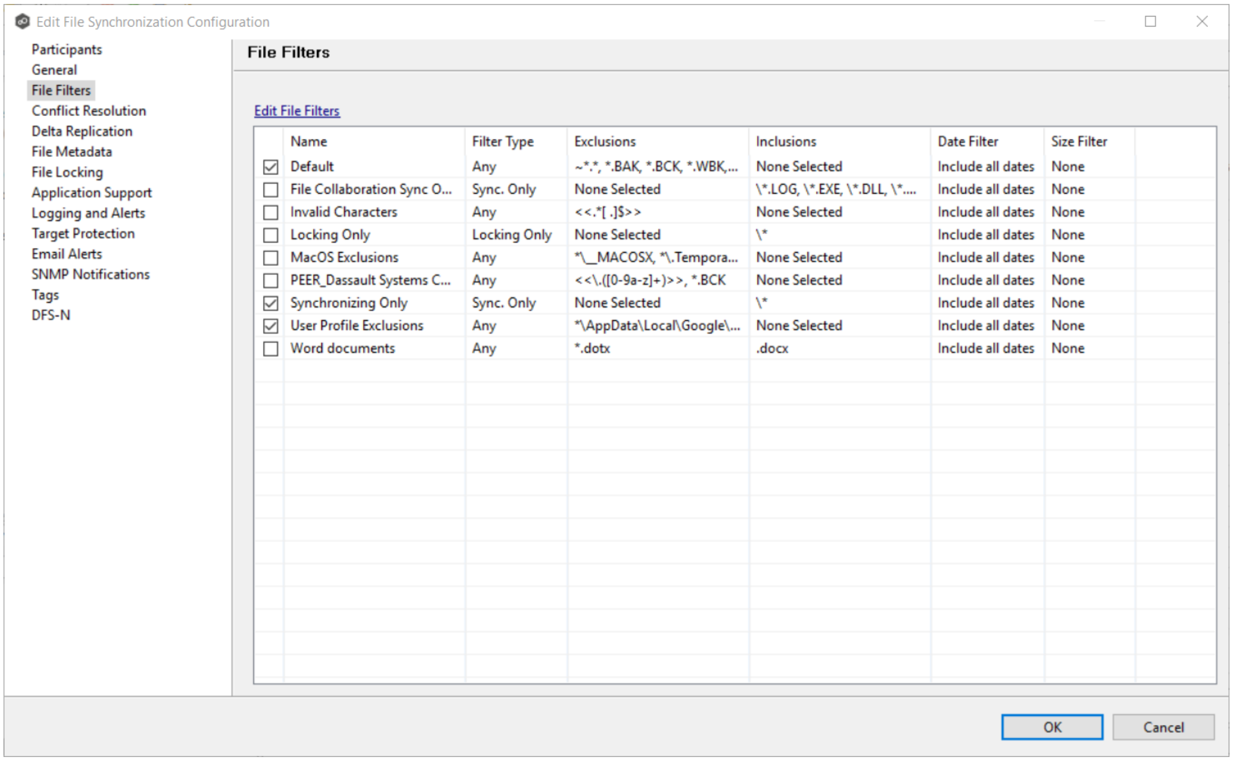

	Name	Filter Type	Exclusions	Inclusions	Date Filter	Size Filter
☑	Default	Any	~*.*, *.BAK, *.BCK, *.WBK,...	None Selected	Include all dates	None
□	File Collaboration Sync O...	Sync. Only	None Selected	*.LOG, *.EXE, *.DLL, *....	Include all dates	None
□	Invalid Characters	Any	<<.*[.]$>>	None Selected	Include all dates	None
□	Locking Only	Locking Only	None Selected	*	Include all dates	None
□	MacOS Exclusions	Any	*__MACOSX, *\\.Tempora...	None Selected	Include all dates	None
□	PEER_Dassault Systems C...	Any	<<\\.([0-9a-z]+)>>, *.BCK	None Selected	Include all dates	None
☑	Synchronizing Only	Sync. Only	None Selected	*	Include all dates	None
☑	User Profile Exclusions	Any	*\AppData\Local\Google\...	None Selected	Include all dates	None
□	Word documents	Any	*.dotx	.docx	Include all dates	None

2. Click **OK**.

Conflict Resolution

The **Conflict Resolution** page in the **Edit File Synchronization Configuration** wizard allows you to specify the file conflict resolution options to use during the initial scan when a file conflict exists for a file between two or more hosts.

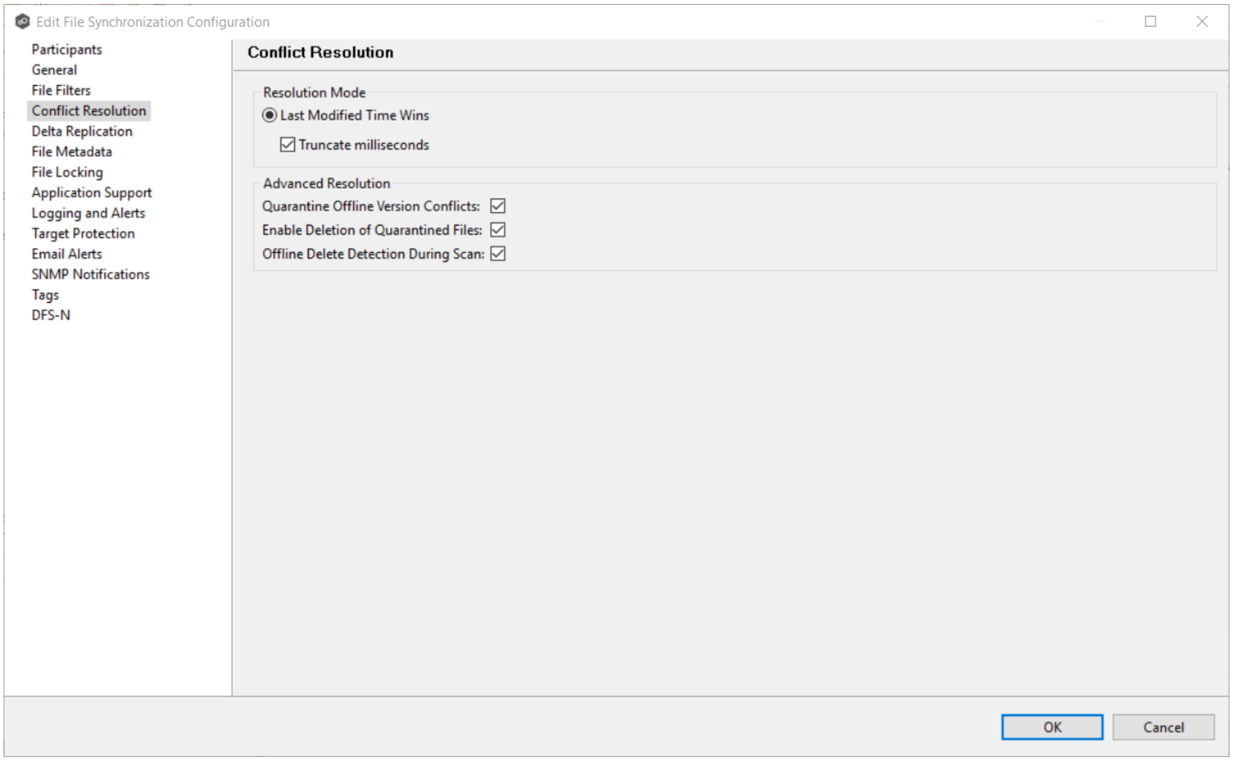

Edit File Synchronization Configuration

Conflict Resolution

Resolution Mode

● Last Modified Time Wins

☑ Truncate milliseconds

Advanced Resolution

Quarantine Offline Version Conflicts: ☑

Enable Deletion of Quarantined Files: ☑

Offline Delete Detection During Scan: ☑

To modify conflict resolution settings for the File Synchronization job:

1. Select the resolution mode:

Last Modified Time Wins	A file's modification time will be used to designate an instance as a resolution candidate. The later the modification time, the greater the likelihood for a file's selection. Option: **Truncate milliseconds:** When comparing the time stamps of a file on two or more hosts, truncate the millisecond value from each time stamp.
None (Manual Resolution)	This is an advanced option. Contact Peer Software to enable. When selected, any file conflicts that are encountered during the initial synchronization process will result in quarantines that are added to the File Conflict List. These file conflicts must be resolved manually by selecting the host with the correct version of the file from the conflict list.

All the types listed above have the potential for producing multiple resolution candidates. A collaboration session can be configured with any one of the available conflict options. If a option produces more than one candidate for a conflicted file, a winner will be selected arbitrarily.

2. Select the **Advanced Resolution** options you want applied.

Quarantine Offline Version Conflicts	Enable 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.
Enable Deletion of Quarantined Files	If a file that is quarantined is deleted, Peer Management Center will process the delete event and remove the quarantine when this option is enabled.
Offline Delete Detection During Scan	If this option is enabled and target protection is enabled, and it can be determined that a file or folder has been deleted since the session was stopped, then the file or folder will be deleted from all hosts. If this option is not enabled then the deleted file or folder will be brought back to any host where it was removed.

3. Click **OK**.

Delta Replication

The **Delta Replication** page in the **Edit File Synchronization Configuration** wizard 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 the use of 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.

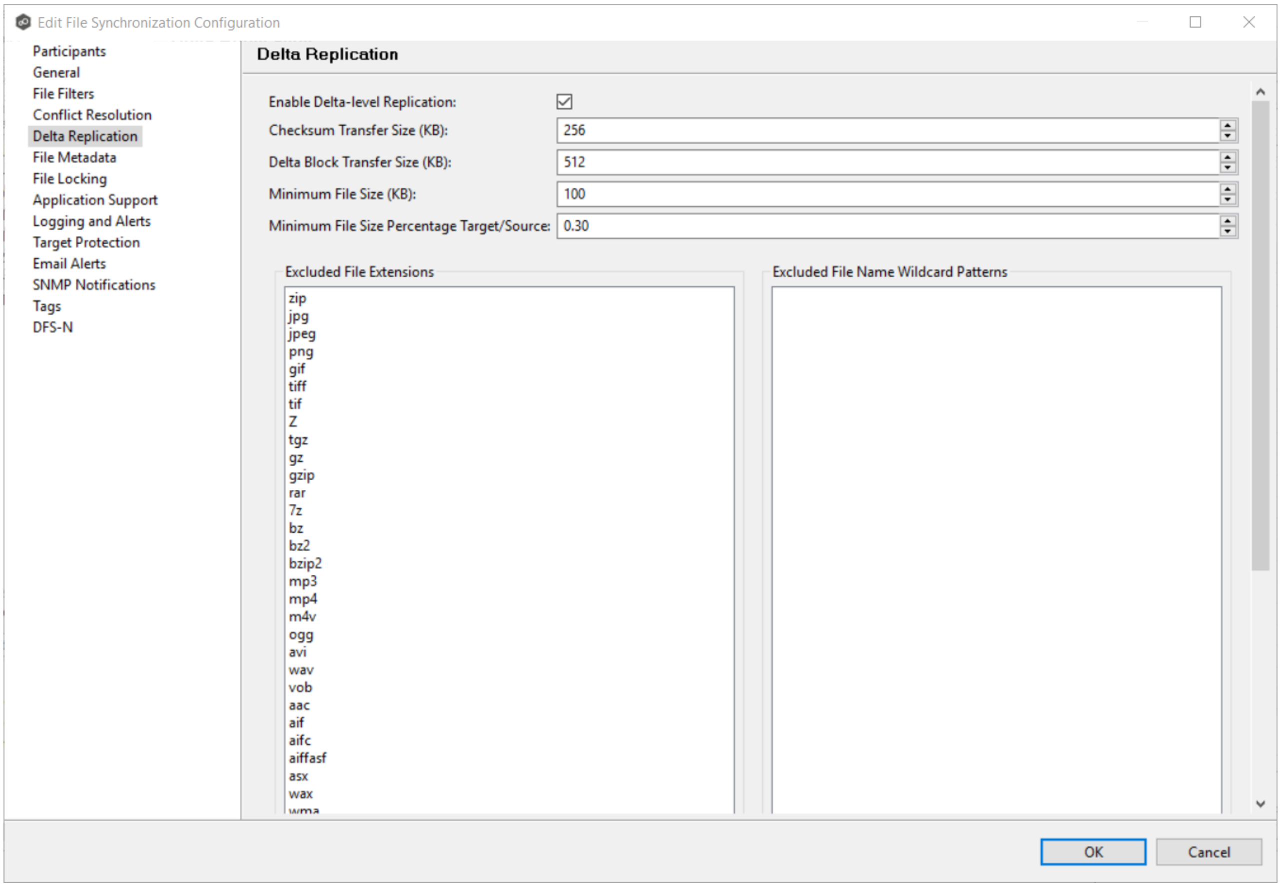

Delta Replication

Enable Delta-level Replication: ☑

Checksum Transfer Size (KB): 256

Delta Block Transfer Size (KB): 512

Minimum File Size (KB): 100

Minimum File Size Percentage Target/Source: 0.30

Excluded File Extensions

zip
jpg
jpeg
png
gif
tiff
tif
Z
tgz
gz
gzip
rar
7z
bz
bz2
bzip2
mp3
mp4
m4v
ogg
avi
wav
vob
aac
aif
aifc
aiffasf
asx
wax
wma

Excluded File Name Wildcard Patterns

To modify delta-level replication options:

1. Modify the following the fields as necessary.

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 wildcard 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 Wildcard Patterns	Enter a list of file name wildcard patterns to be excluded from delta encoding. If a file name matches any wildcard 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**.

File Metadata

The **File Metadata** page in the **Edit File Synchronization Configuration** wizard allows you to modify your file metadata synchronization settings and presents additional options for metadata replication. See File Metadata Replication in Advanced Topics for more information about file metadata replication.

To enable file metadata replication:

1. Select when you want the metadata synchronized (you can select one or both of the 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 (DACL, SACL, Owner.) 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 (DACL, SACL, Owner) will be synchronized during the initial scan.

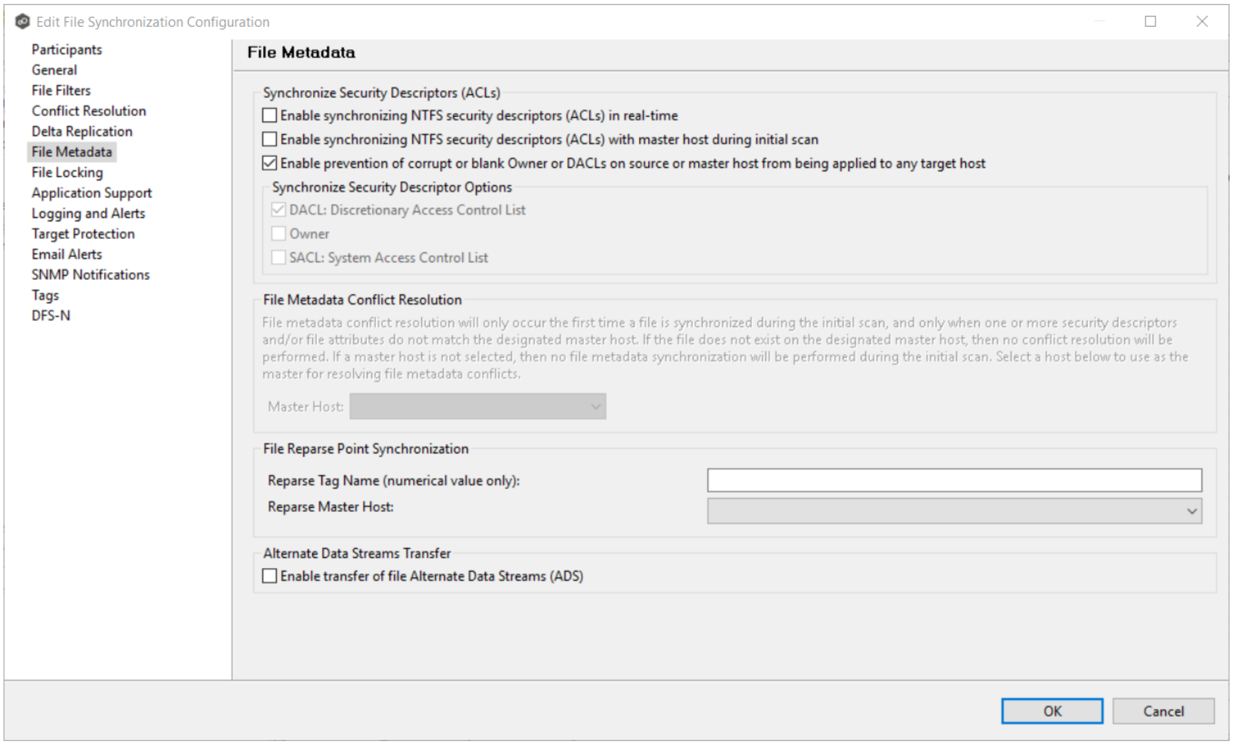

File Metadata

Synchronize Security Descriptors (ACLs)

□ Enable synchronizing NTFS security descriptors (ACLs) in real-time
□ Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan
☑ Enable prevention of corrupt or blank Owner or DACLs on source or master host from being applied to any target host

Synchronize Security Descriptor Options

☑ DACL: Discretionary Access Control List
□ Owner
□ SACL: System Access Control List

File Metadata Conflict Resolution

File metadata conflict resolution will only occur the first time a file is synchronized during the initial scan, and only when one or more security descriptors and/or file attributes do not match the designated master host. If the file does not exist on the designated master host, then no conflict resolution will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan. Select a host below to use as the master for resolving file metadata conflicts.

Master Host:

File Reparse Point Synchronization

Reparse Tag Name (numerical value only):
Reparse Master Host:

Alternate Data Streams Transfer

□ Enable transfer of file Alternate Data Streams (ADS)

2. Click **OK** in the informational dialog that appears after selecting a metadata option.

3. Select which security descriptor components (DACL, SACL and Owner) are synchronized.

   In general, you will usually need to synchronize only DACLs. If you need to synchronize SACLs or Owner, then the user that a Peer Agent service is run under on each participating host must have permission to read and write SACLs and Owner.

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.

   Conflict resolution for file metadata 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 will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan.

5. (Optional) Enter values for one or both of the 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/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 verse 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.

6. (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.

7. Click **OK**.

File Locking

The **File Locking** page in the **Edit File Synchronization Configuration** wizard presents options pertaining to how source and target files are locked by Peer Management Center.

To modify file locking options:

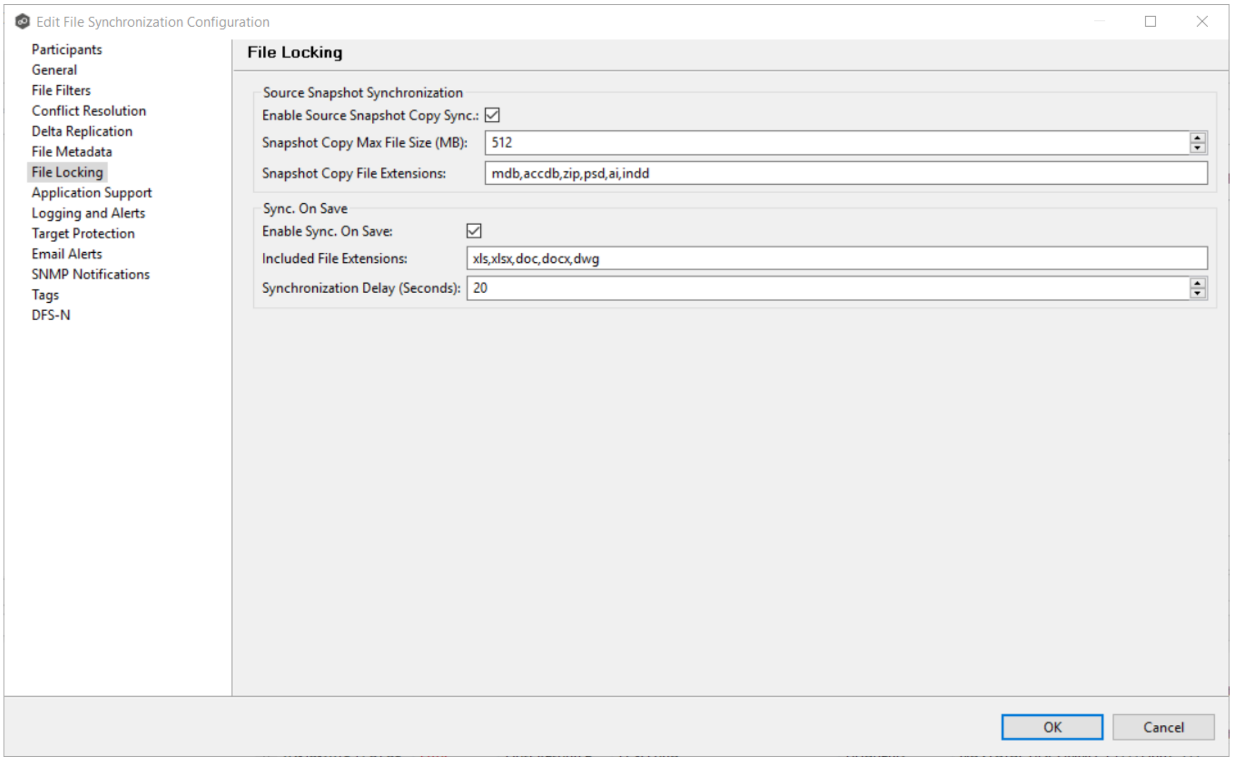

File Locking

Source Snapshot Synchronization

Enable Source Snapshot Copy Sync.: ☑

Snapshot Copy Max File Size (MB): 512

Snapshot Copy File Extensions: mdb,accdb,zip,psd,ai,indd

Sync. On Save

Enable Sync. On Save: ☑

Included File Extensions: xls,xlsx,doc,docx,dwg

Synchronization Delay (Seconds): 20

Modify these fields as needed:

Enable Source Snapshot 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 Synchronization job, you have the option of selecting applications that are automatically optimized. You can modify your selections when editing the job.

To modify which applications are optimized:

1. Select the applications to be optimized.

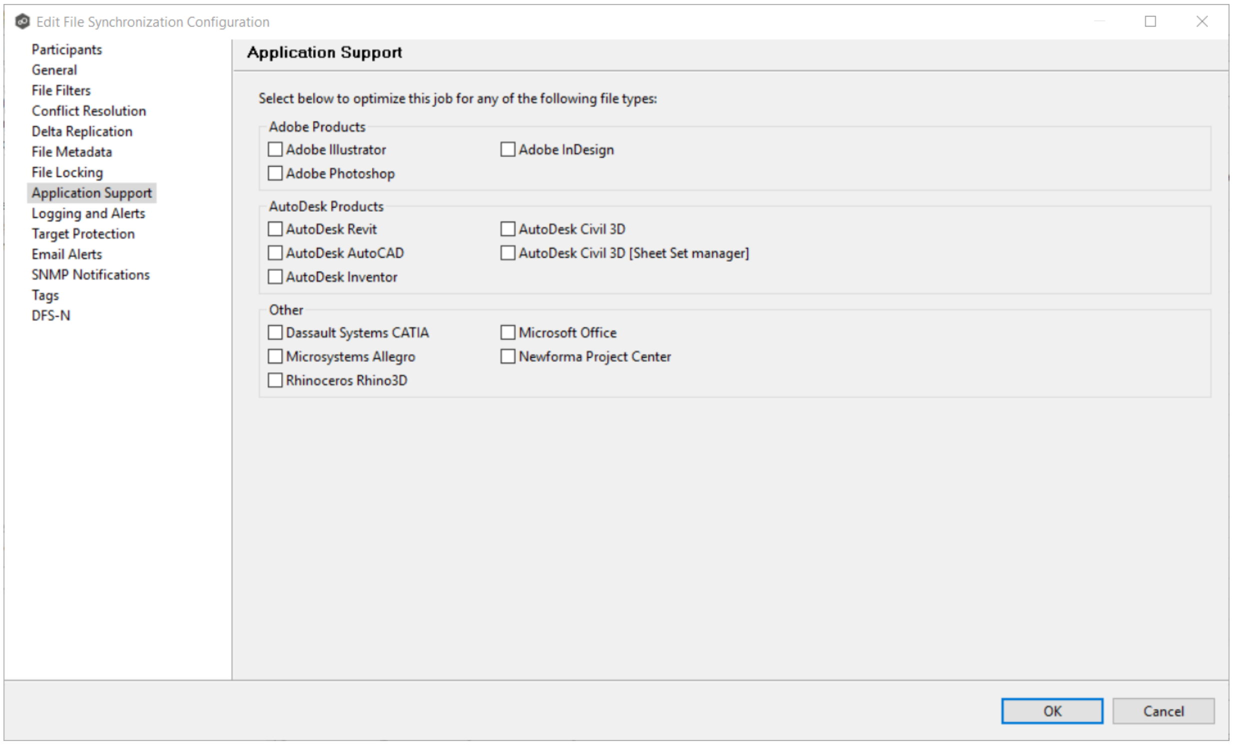

Application Support

Select below to optimize this job for any of the following file types:

Adobe Products

- ☐ Adobe Illustrator
- ☐ Adobe InDesign
- ☐ Adobe Photoshop

AutoDesk Products

- ☐ AutoDesk Revit
- ☐ AutoDesk Civil 3D
- ☐ AutoDesk AutoCAD
- ☐ AutoDesk Civil 3D [Sheet Set manager]
- ☐ AutoDesk Inventor

Other

- ☐ Dassault Systems CATIA
- ☐ Microsoft Office
- ☐ Microsystems Allegro
- ☐ Newforma Project Center
- ☐ Rhinoceros Rhino3D

2. Click **OK**.

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

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 synchronization activity is logged for all severity levels. You can enable or disable file event logging as well as select the level of granularity.

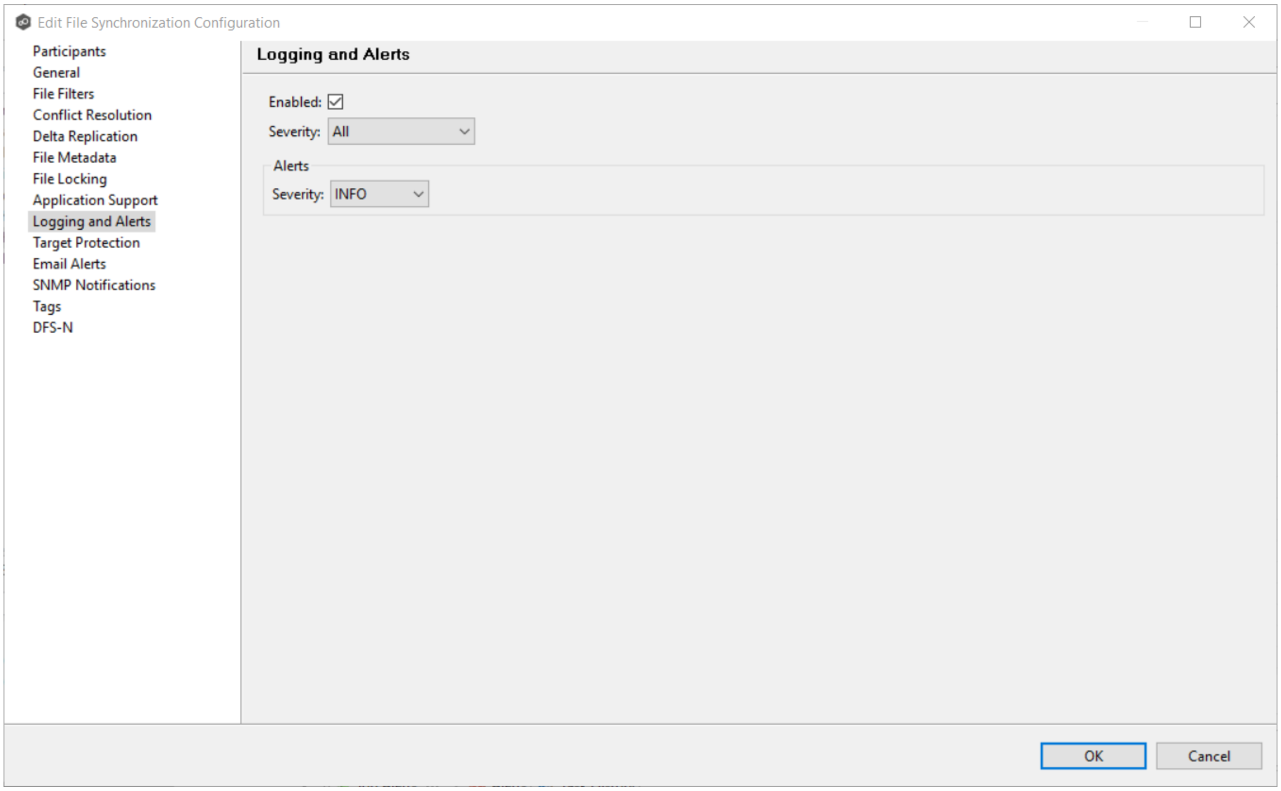

Logging and Alerts

Enabled: ☑

Severity: All

Alerts

Severity: INFO

Below is a list of logging fields and their descriptions:

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.
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

Configured in the screen shown above, various types of alerts will 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 File Synchronization Configuration dialog.

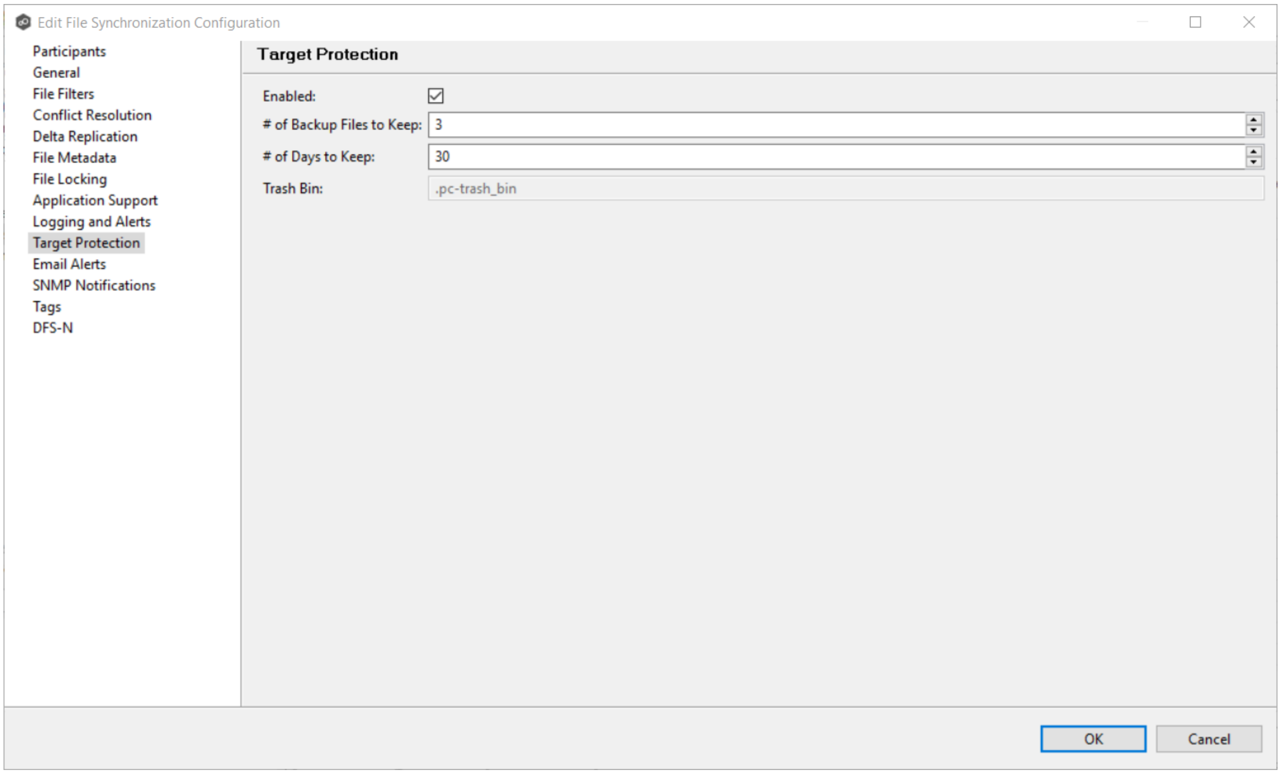

Enabled: ☑
# of Backup Files to Keep: 3
# of Days to Keep: 30
Trash Bin: .pc-trash_bin

Modify the fields as needed:

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 Configuration** wizard 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.

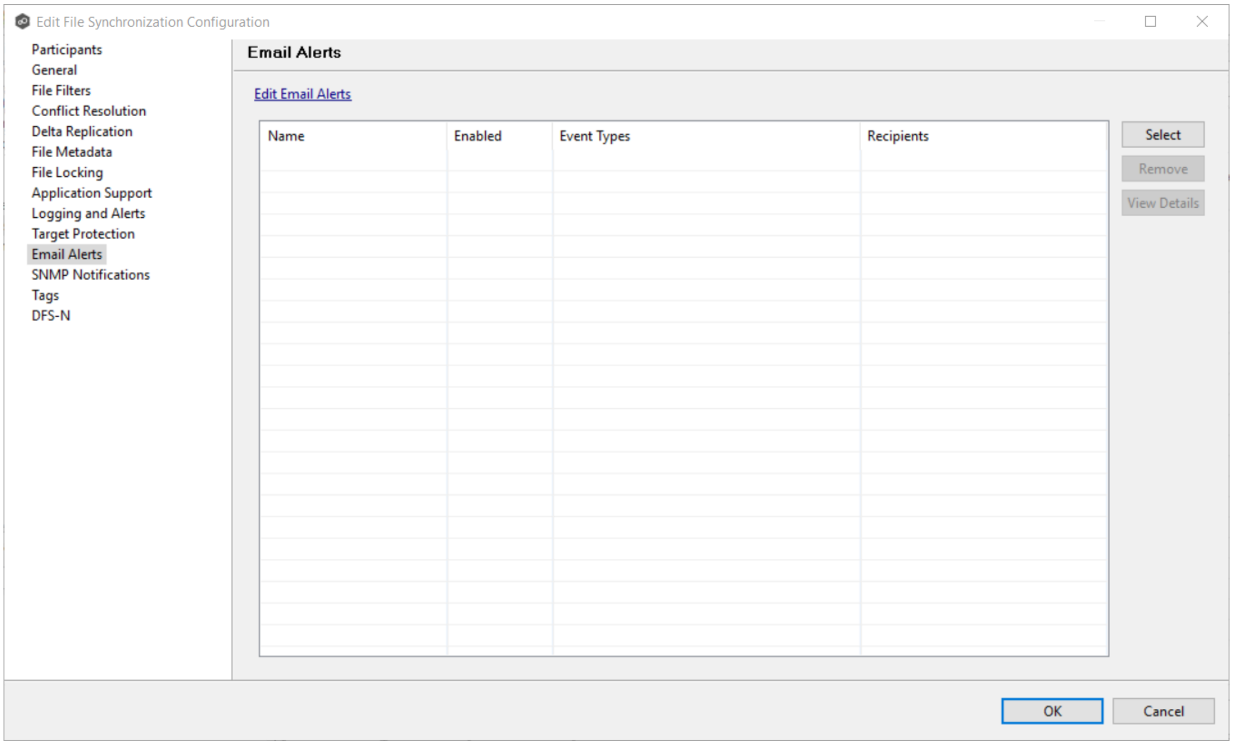

The **Select Email Alert** dialog opens.

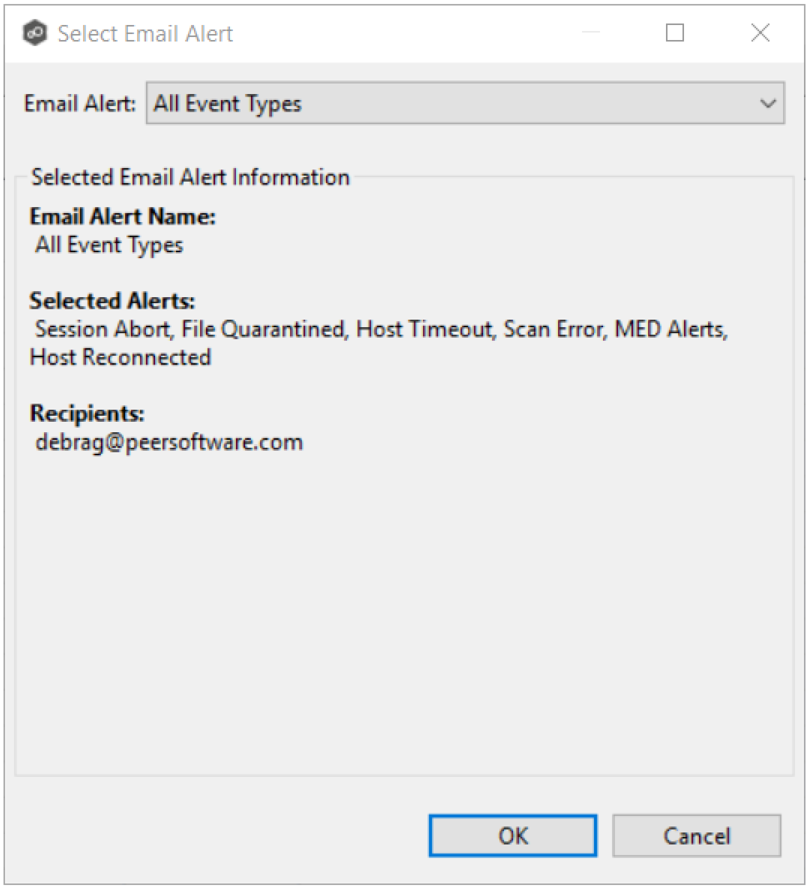

Email Alert: All Event Types

Selected Email Alert Information

Email Alert Name:
All Event Types

Selected Alerts:
Session Abort, File Quarantined, Host Timeout, Scan Error, MED Alerts, Host Reconnected

Recipients:
debrag@peersoftware.com

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**.

SNMP Notifications

The **SNMP Notifications** page in the **Edit File Synchronization Configuration** wizard allows you to select which SNMP notifications 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**.

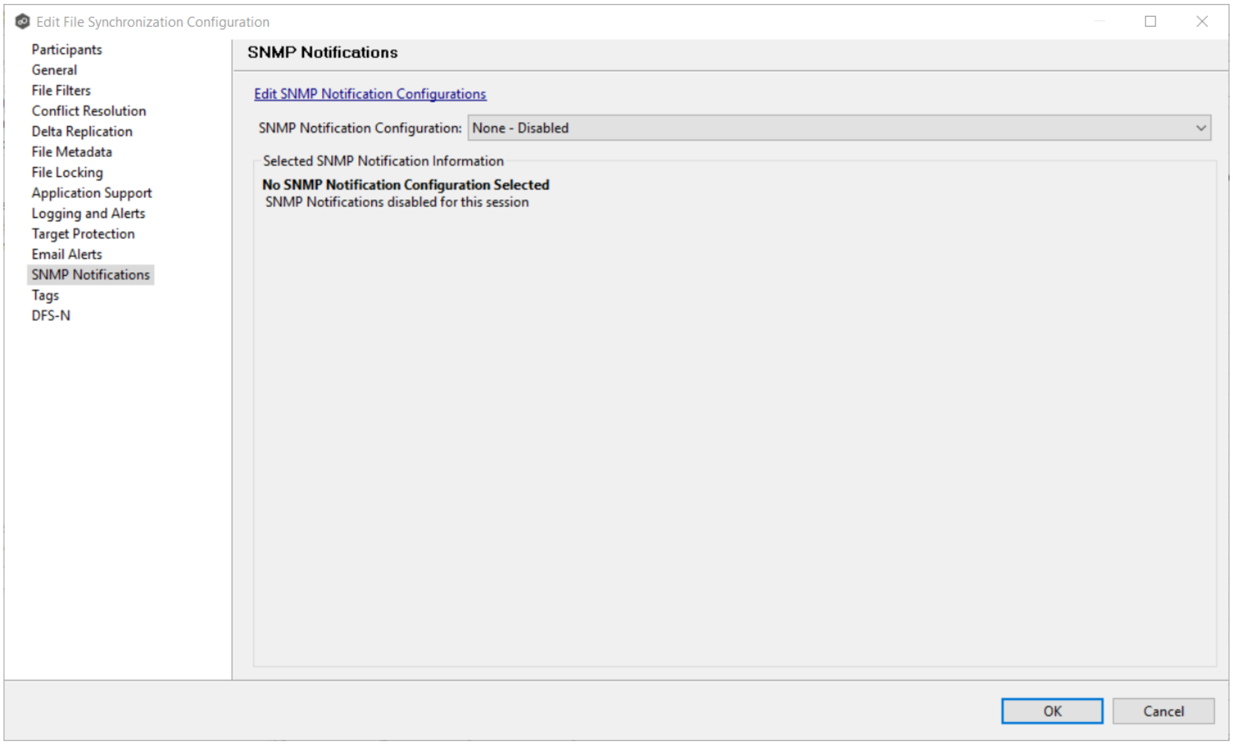

SNMP Notification Configuration: None - Disabled

Selected SNMP Notification Information

No SNMP Notification Configuration Selected
SNMP Notifications disabled for this session

2. Click **OK**.

Tags

The **Tags** page in the **Edit File Synchronization Configuration** wizard allows you to 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.

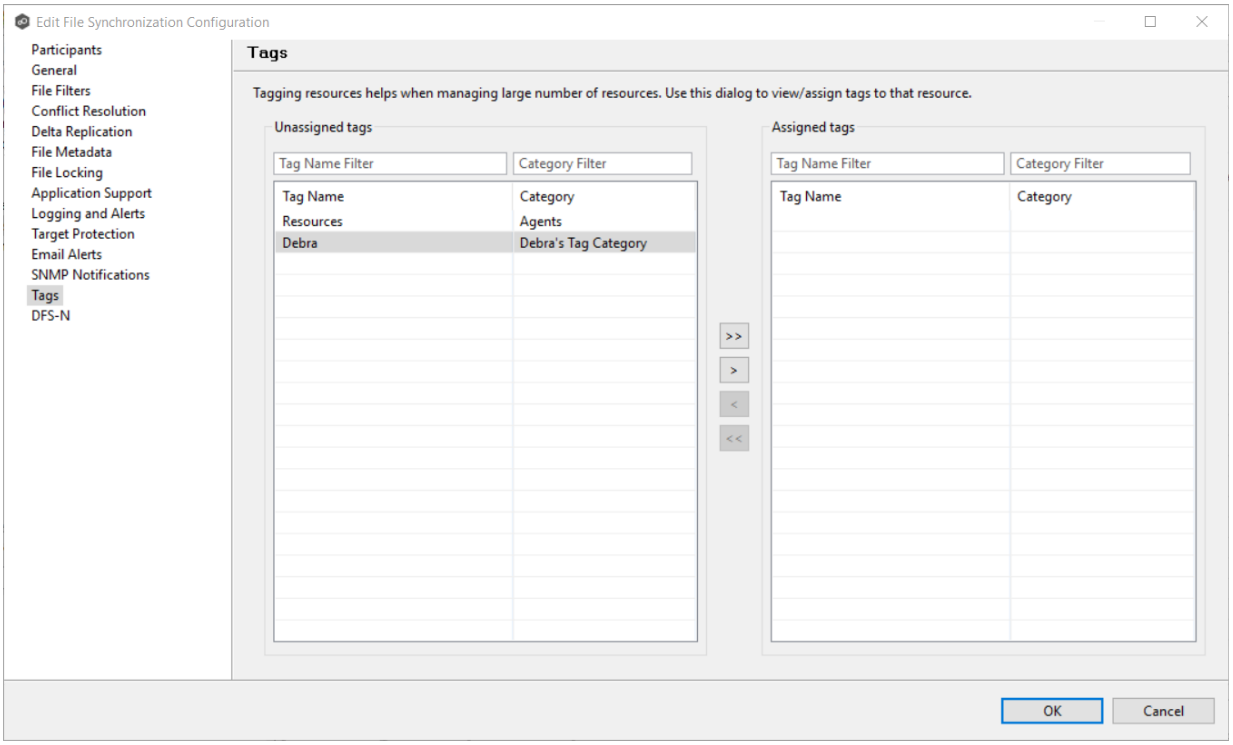
Edit File Synchronization Configuration

Tags

Tagging resources helps when managing large number of resources. Use this dialog to view/assign tags to that resource.

Unassigned tags

Tag Name	Category
Resources	Agents
Debra	Debra's Tag Category

Assigned tags

Tag Name	Category

DFS-N

The **DFS-N** page in the **Edit File Synchronization Configuration** wizard presents options for linking a DFS namespace folder to this job. See Link a Namespace Folder with an Existing File Collaboration or Synchronization Job for more information.

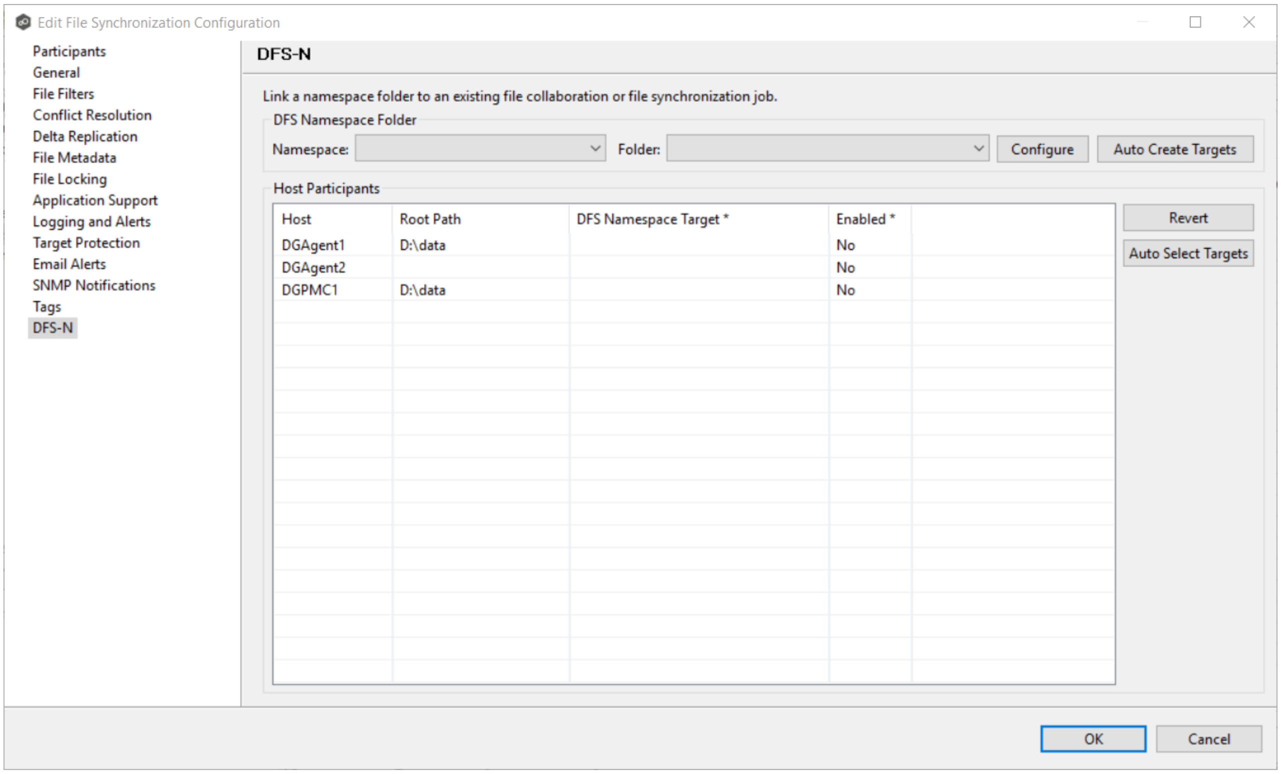

Link a namespace folder to an existing file collaboration or file synchronization job.

Namespace:

Folder:

Host	Root Path	DFS Namespace Target *	Enabled *
DGAgent1	D:\data		No
DGAgent2			No
DGPMC1	D:\data		No

Editing Multiple Jobs

The 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 the 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:

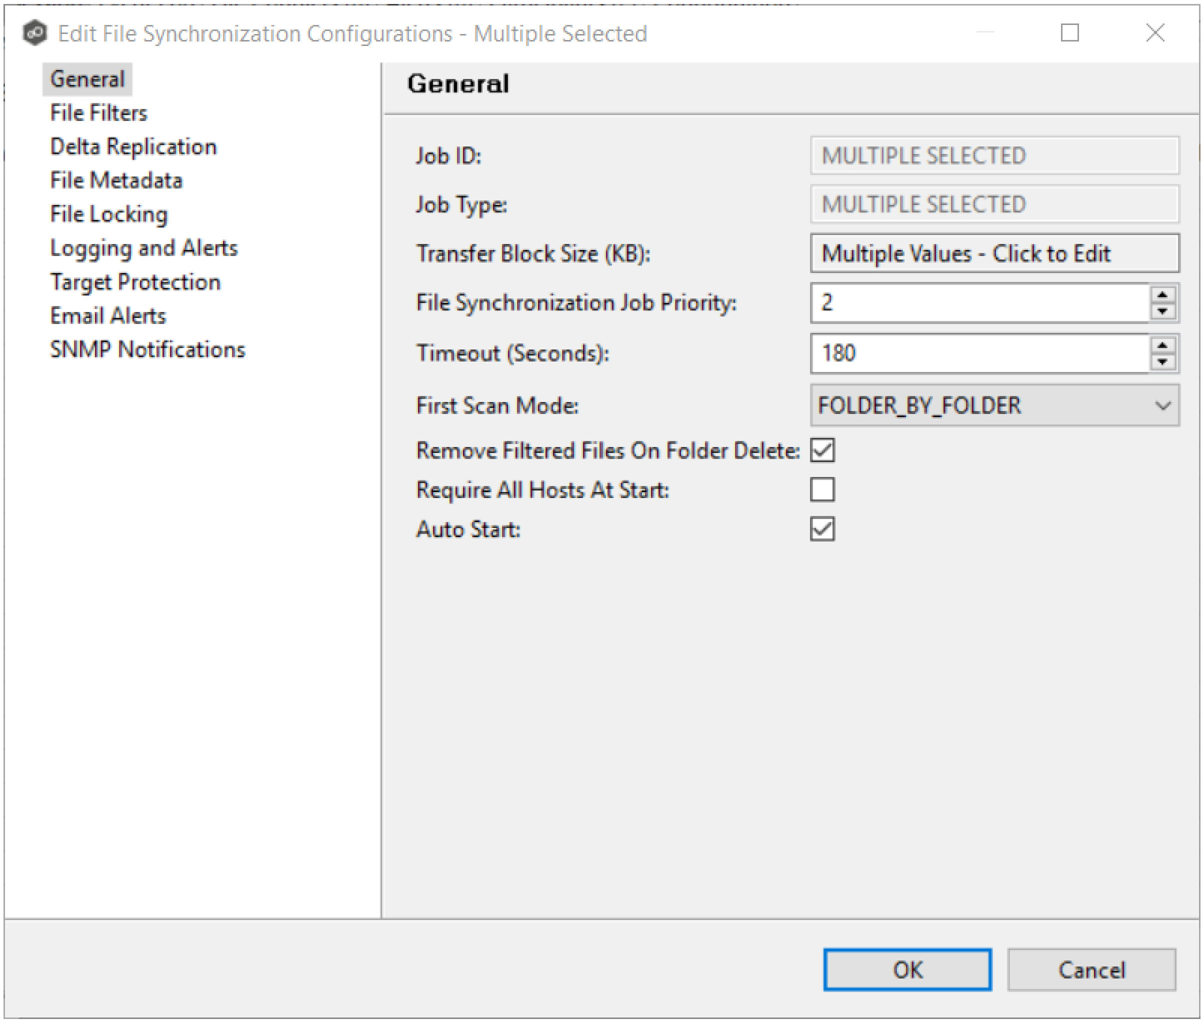

Edit File Synchronization Configurations - Multiple Selected

General

Job ID: MULTIPLE SELECTED

Job Type: MULTIPLE SELECTED

Transfer Block Size (KB): Multiple Values - Click to Edit

File Synchronization Job Priority: 2

Timeout (Seconds): 180

First Scan Mode: FOLDER_BY_FOLDER

Remove Filtered Files On Folder Delete: ☑

Require All Hosts At Start: □

Auto Start: ☑

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:

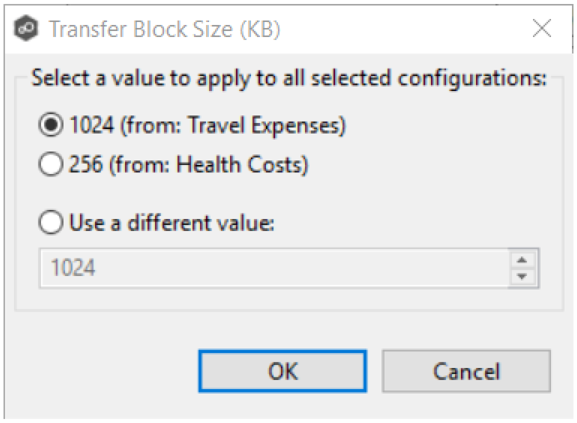

Transfer Block Size (KB)

Select a value to apply to all selected configurations:

● 1024 (from: Travel Expenses)

○ 256 (from: Health Costs)

○ Use a different value:

1024

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 check box 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
- Managing File Conflicts

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 File 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 **File Synchronization Job Summary** view, and then choose **Start** from the pop-up menu.

- Open a job and then click the **Start/Stop** button in the bottom left corner of the job's **Summary** tab (shown below).

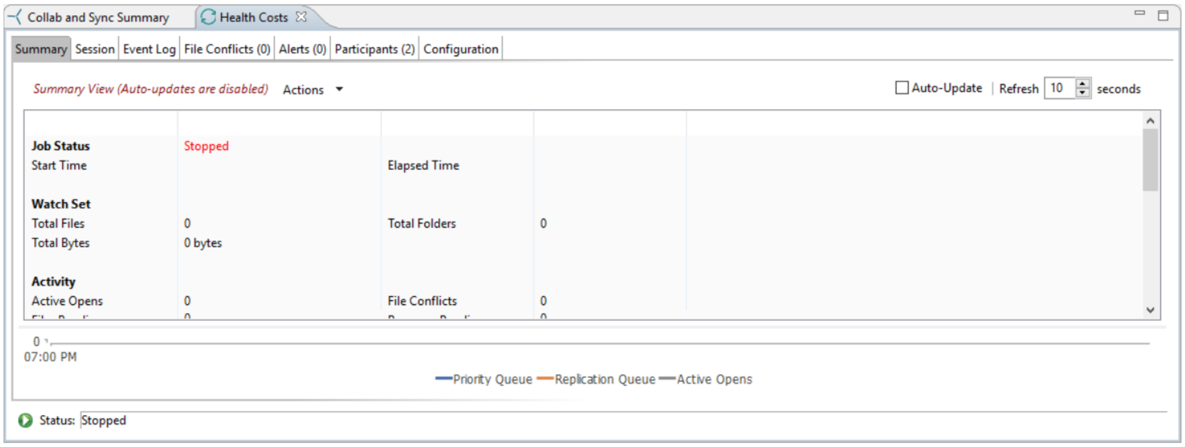

Summary View (Auto-updates are disabled)

☐ Auto-Update | Refresh 10 seconds

Job Status	Stopped		
Start Time		Elapsed Time	
Watch Set			
Total Files	0	Total Folders	0
Total Bytes	0 bytes		
Activity			
Active Opens	0	File Conflicts	0

Status: Stopped

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.

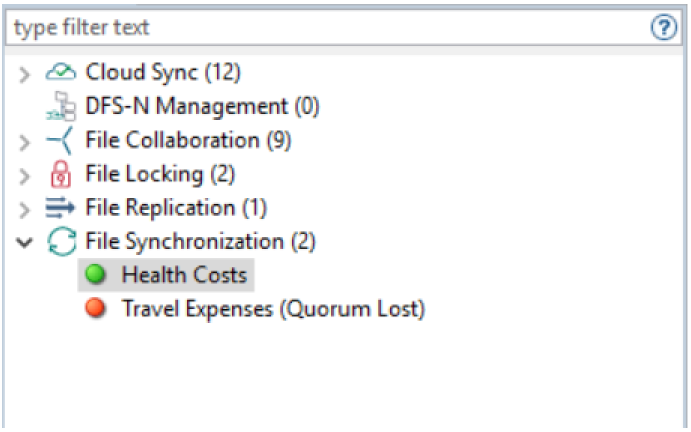

Stopping a File Synchronization Job

You can stop a File Synchronization job at any time by clicking the **Stop** button. 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 **File Collab, Sync, and Locking** in the navigation tree.

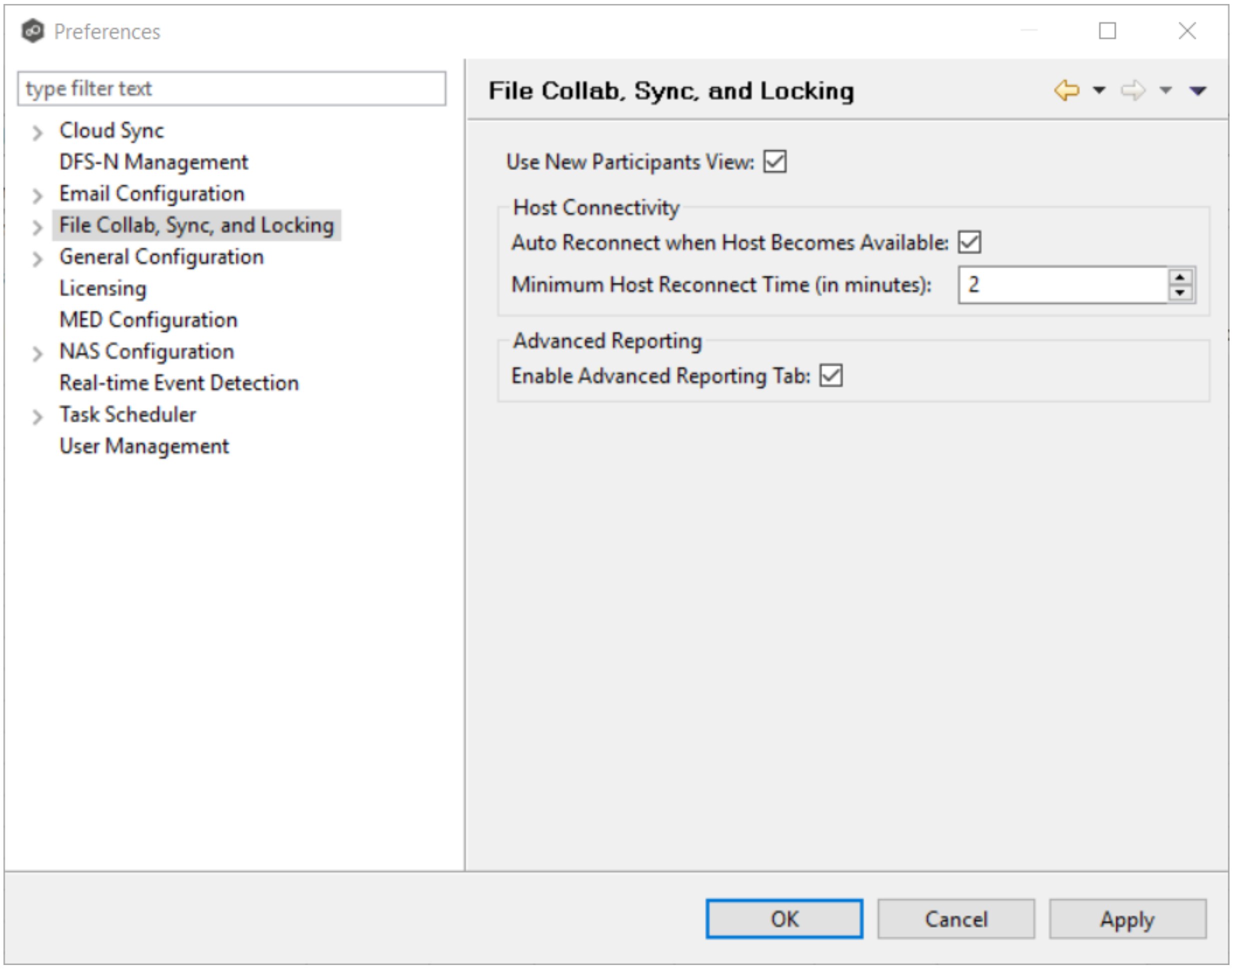

File Collab, Sync, and Locking

Use New Participants View: ☑

Host Connectivity

Auto Reconnect when Host Becomes Available: ☑

Minimum Host Reconnect Time (in minutes): 2

Advanced Reporting

Enable Advanced Reporting Tab: ☑

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.

Unavailable Hosts

If a host becomes unavailable while a File Synchronization job is running, and is unreachable within the configured timeout period (specified within 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, then the host will be pinged, and if still no response, the host will be taken out of the running session, a FATAL event will be logged, and the **Participants** view for the job will be 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) will be sent.

If auto-restart not enabled, you will need to stop and start the File Synchronization job in order 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 jobs.

Quorum

In order 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 will automatically be terminated. If email alerts and/or SNMP notifications are configured and enabled for **Session Aborts**, then the appropriate message(s) will be sent.

Managing File Conflicts

Files conflicts can occur for the following reasons:

- A file was modified on two or more hosts.
- A general I/O failure occurs on the source host after the file has been modified, but before the file is synchronized to all target hosts. In this case, the file is automatically quarantined.

When a file conflict is detected, the file is placed in the File Conflicts list and assigned a conflict status, which determines how the conflict is resolved. The File Conflicts list is displayed in the File Conflicts view.

For an elaboration on conflict scenarios, see file conflict scenarios.

- A job is started and Initial Scan Logic is performed on a file:

  If the file has never been synchronized by Peer Management Center and if file sizes and last modified times do not match on all participating hosts, or if the file does not exist on one or more hosts, then the file will be synchronized based on the file conflict resolution strategy, which is typically the most recent last modified time.

  Files that have previously been synchronized by Peer Management Center where just a single file's last modified timestamp is newer than the last recorded timestamp, then that file will be synchronized to all other hosts. However, if two or more files have a more recent last modified timestamp than was last recorded timestamp, then the file will be quarantined (this is the default behavior and can be disabled by deselecting the file conflict resolution strategy **Quarantine Offline Version Conflicts** option).

- A single user has a file opened before starting a collaboration job:

  A file conflict will be created with a status of **Pending Initial Synchronization**. After the user closes the file, if all file sizes and timestamps match, then the file conflict is removed and no synchronization is performed. However, if any file last modified times or file sizes do not match, the file will be synchronized or quarantined based on the file conflict resolution strategy and according to the initial scan logic detailed above. Once the file is synchronized, the file conflict will be removed.

- Two or more users have a file open before starting a collaboration job

  A file conflict will be created with a status of **Pending Conflict Resolution**. After the users close all files, the conflict will be removed if the last modified timestamp matches on all files. Otherwise, if the file has never been synchronized by Peer Management Center, then the file conflict status will be updated to Quarantined. However, if the file has previously been synchronized by Peer Management Center, then the file will be synchronized or quarantined based on the file conflict resolution strategy and according to the initial scan logic detailed above.

- Two or more users open a file at the same time

  In the rare situation when two users open a file at the same time, or in-and-around the same time and Peer Management Center is unable to obtain corresponding locks on target hosts before this happens (this is dependent on WAN latency and other factors), then a file conflict will be created with a status of **Pending Conflict Resolution**. After all users close the files, file lock conflict resolution will be performed as follows:

  - If all files last modified timestamps and file sizes match, then the file conflict will be removed.

  - o If only a single file has been modified, then the file that changed is synchronized or quarantined based on the configured file conflict resolution strategy and according to the initial scan logic detailed above.
  - o If two or more files have been modified since it was opened, then the file conflict status will be updated to quarantined.

When a file conflict occurs, the status is set to one of the the following statuses:

- **Pending Conflict Resolution** if the file has already been verified or synchronized by the initial synchronization process.
- **Pending Initial Synchronization** if the file hasn't been verified or synchronized.
- **Quarantined** if the conflict is a result of a fatal I/O error on the source.

Note: If a File Synchronization job is stopped before a file conflict with a status of **Pending Conflict Resolution** is resolved, then that file is automatically quarantined the next time the File Synchronization job is started.

The resolution strategies for the three conflict statuses are as follows:

Conflict Status	Resolution Strategy
Pending Conflict Resolution This status is assigned to files that have already been verified or synchronized by the session via the initial synchronization process.	When all files in use are closed by users on the source hosts, the files will be analyzed to determine if a file conflict has occurred as follows: • If more than one file has been modified, then the file will be quarantined by updating the file conflict status to Quarantined. • If only one file as been modified, then that file will be used as the source, synchronized with all other participating hosts, and removed from the File Conflicts list. • If no files have been modified, then no action will be taken and the file will be removed from the File Conflict list.
Pending Initial Synchronizatio	When all files in use are closed by users on the source hosts, then standard file conflict resolution will be performed based

n This status is assigned to files that have not been verified or synchronized by session via the initial synchronization process.	on the configured File Conflict Resolvers. However, if the **Quarantine Offline Multi-Edits** option is enabled, then if a file is modified on two or more hosts while the collaboration session is not running, and the last modified timestamps are all newer then the last timestamp recorded by the collaboration session, then the file will be quarantined.
Quarantined	A file will be quarantined when a file conflict with **Pending Conflict Resolution** status cannot be resolved or a fatal I/O error occurs. Quarantined files need to be explicitly removed from the File Conflicts list.

.

Once a file is marked as **Quarantined**, the file no longer participates in collaboration, and thus changes to any version of the file will not be propagated to other hosts. However, subsequent file activity on a quarantined file will be logged in the Event Log as a warning so that you can determine who modified the file while it was quarantined.

Quarantined files are saved to disk and will survive session restarts. The File Conflicts list displays the time and date of the quarantine along with an error message indicating the reason for the quarantine.

A Quarantined File event is also logged in the Event Log and you can obtain a more detailed reason for the quarantine by analyzing the Event Log file(s). In addition, if email alerts and/or SNMP notifications are configured and enabled for **File Quarantines**, then the appropriate message(s) will be sent.

Removing a File from Quarantine

You must explicitly remove a file from quarantine in order to have it participate in the collaboration session once again.

You may also chose to perform no action, in which case, the file is removed from the File Conflict 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 removing the file conflict, then the Status field in the File Conflict table is updated to reflect the error.

To remove a file (or multiple files) from quarantine:

1. Select the file in the File Conflicts list.
2. Select the host with the correct version.
3. Click the **Release Conflict** button.

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 the 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 the Peer Management Center or select the **New** menu item 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 the 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 mange 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.

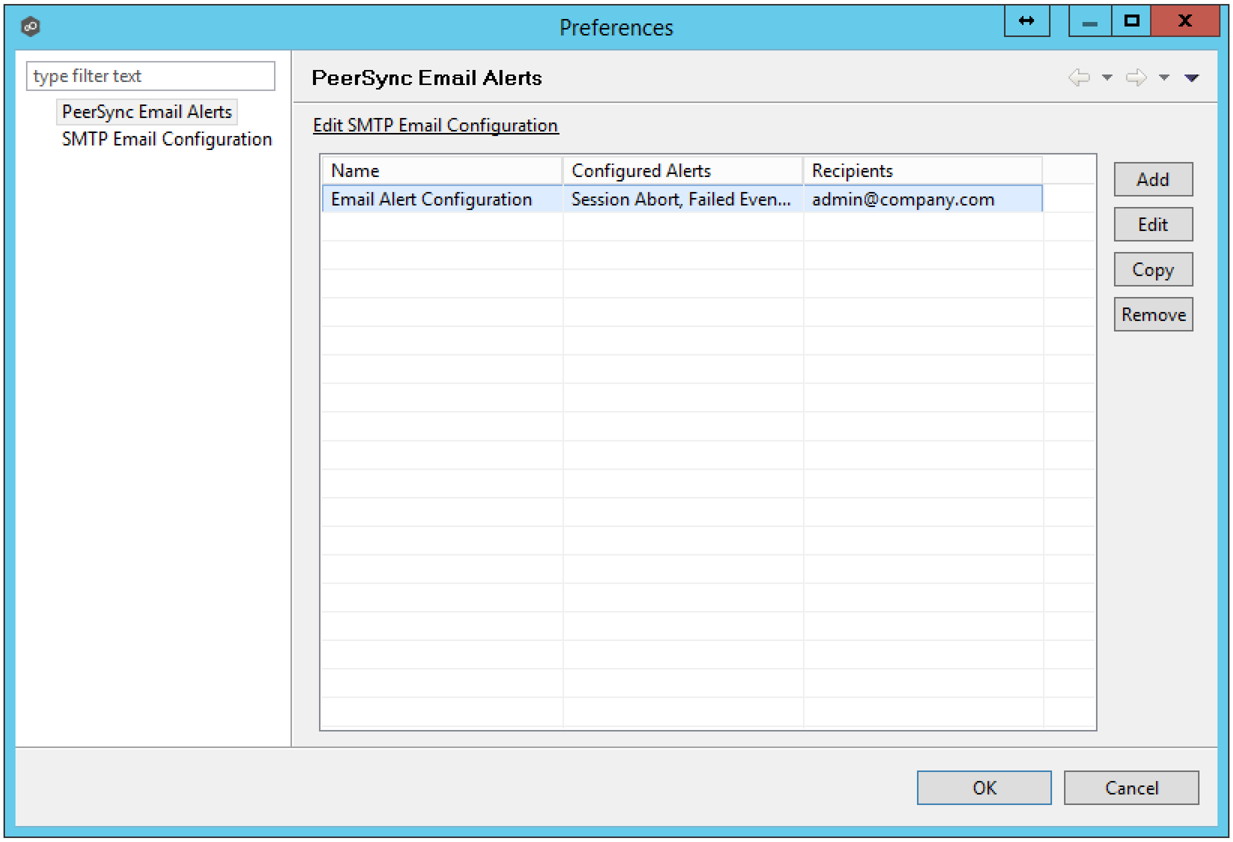

Name	Configured Alerts	Recipients
Email Alert Configuration	Session Abort, Failed Even...	admin@company.com

Upon adding or editing an email alert, the following dialog is displayed:

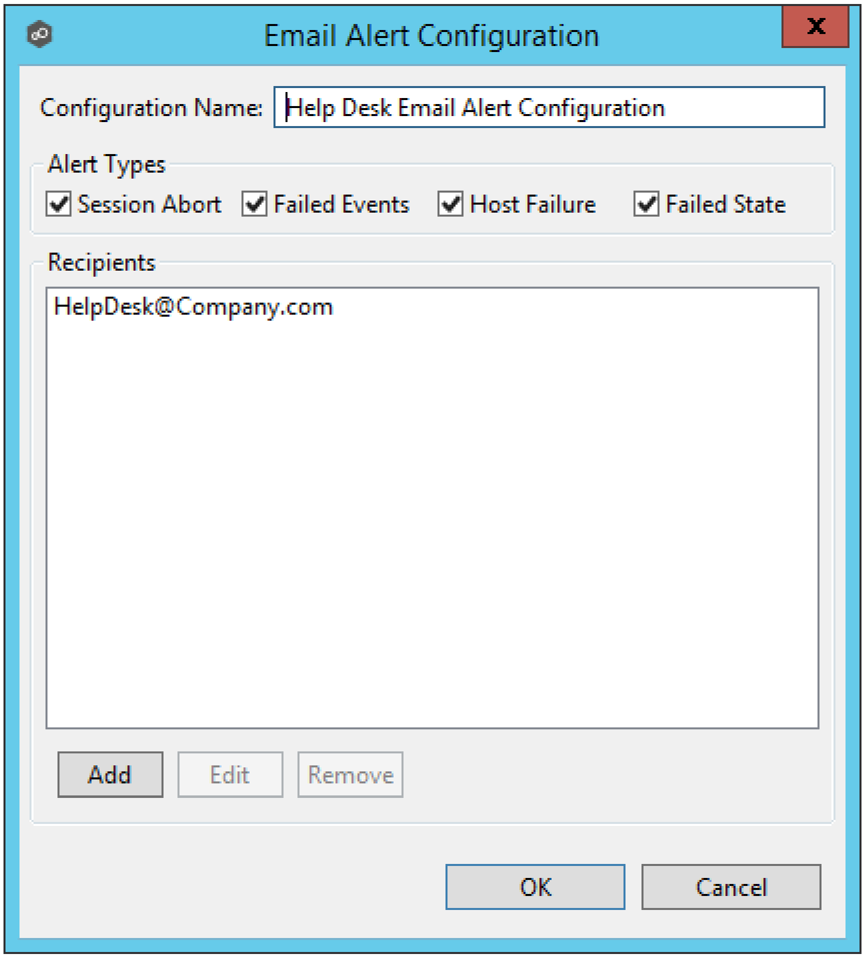

Email Alert Configuration

Configuration Name: Help Desk Email Alert Configuration

Alert Types

☑ Session Abort ☑ Failed Events ☑ Host Failure ☑ Failed State

Recipients

HelpDesk@Company.com

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 has to be installed as a Service and running version 9.3.0 or newer.
- Peer Agent has to 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 has to 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 the PMC.

- 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.

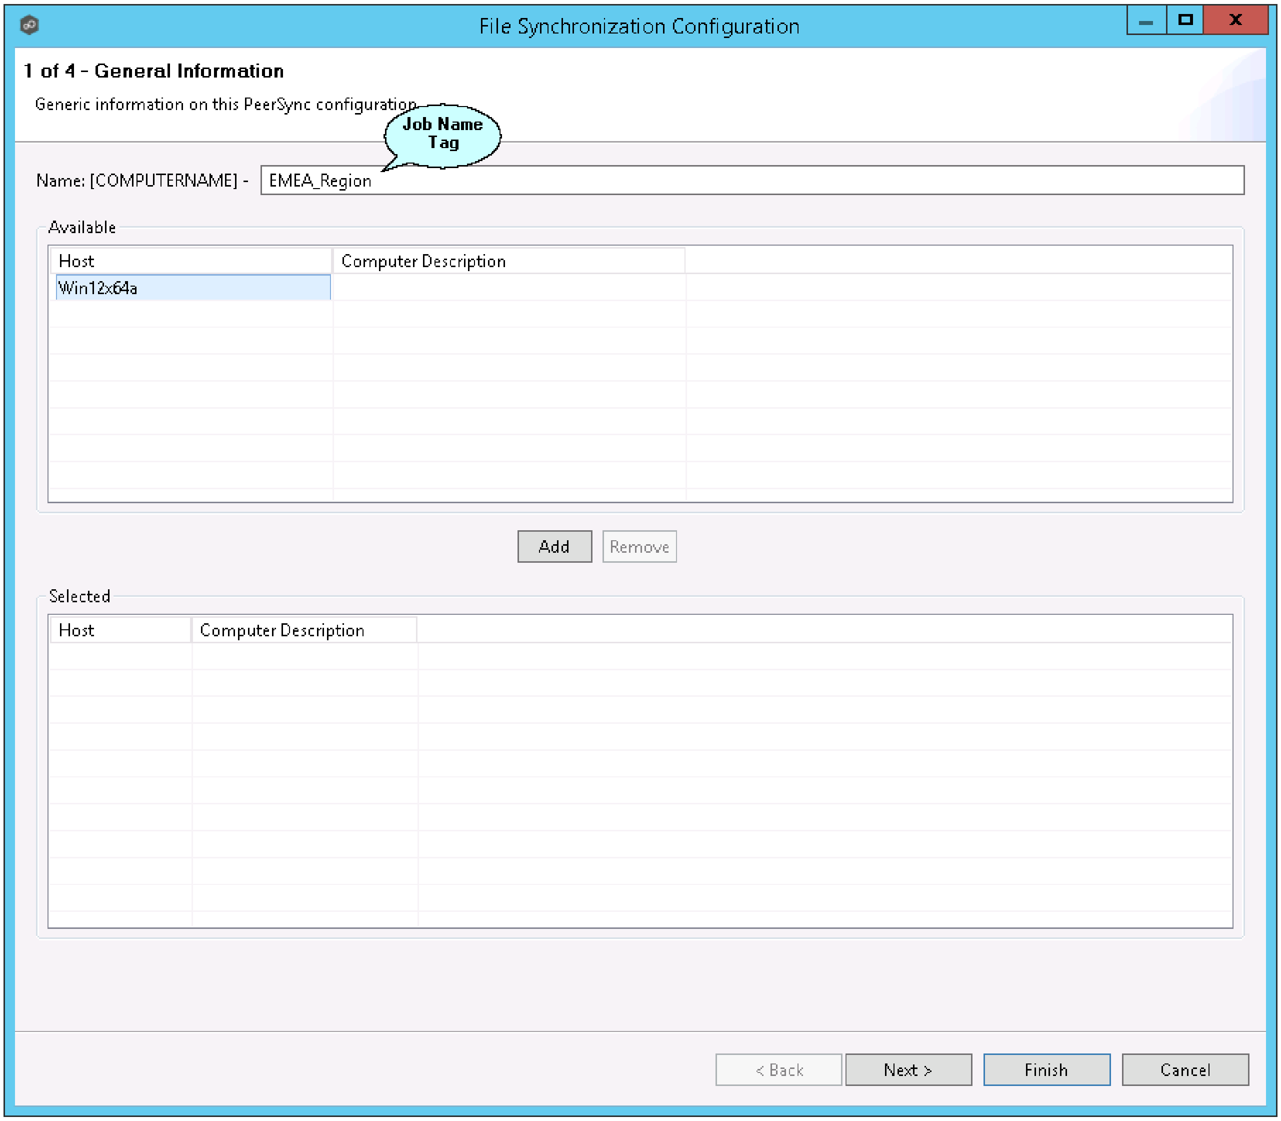

File Synchronization Configuration

1 of 4 - General Information

Generic information on this PeerSync configuration

Name: [COMPUTERNAME] - EMEA_Region

Available

Host	Computer Description
Win12x64a	

Selected

Host	Computer Description

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 the 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**.

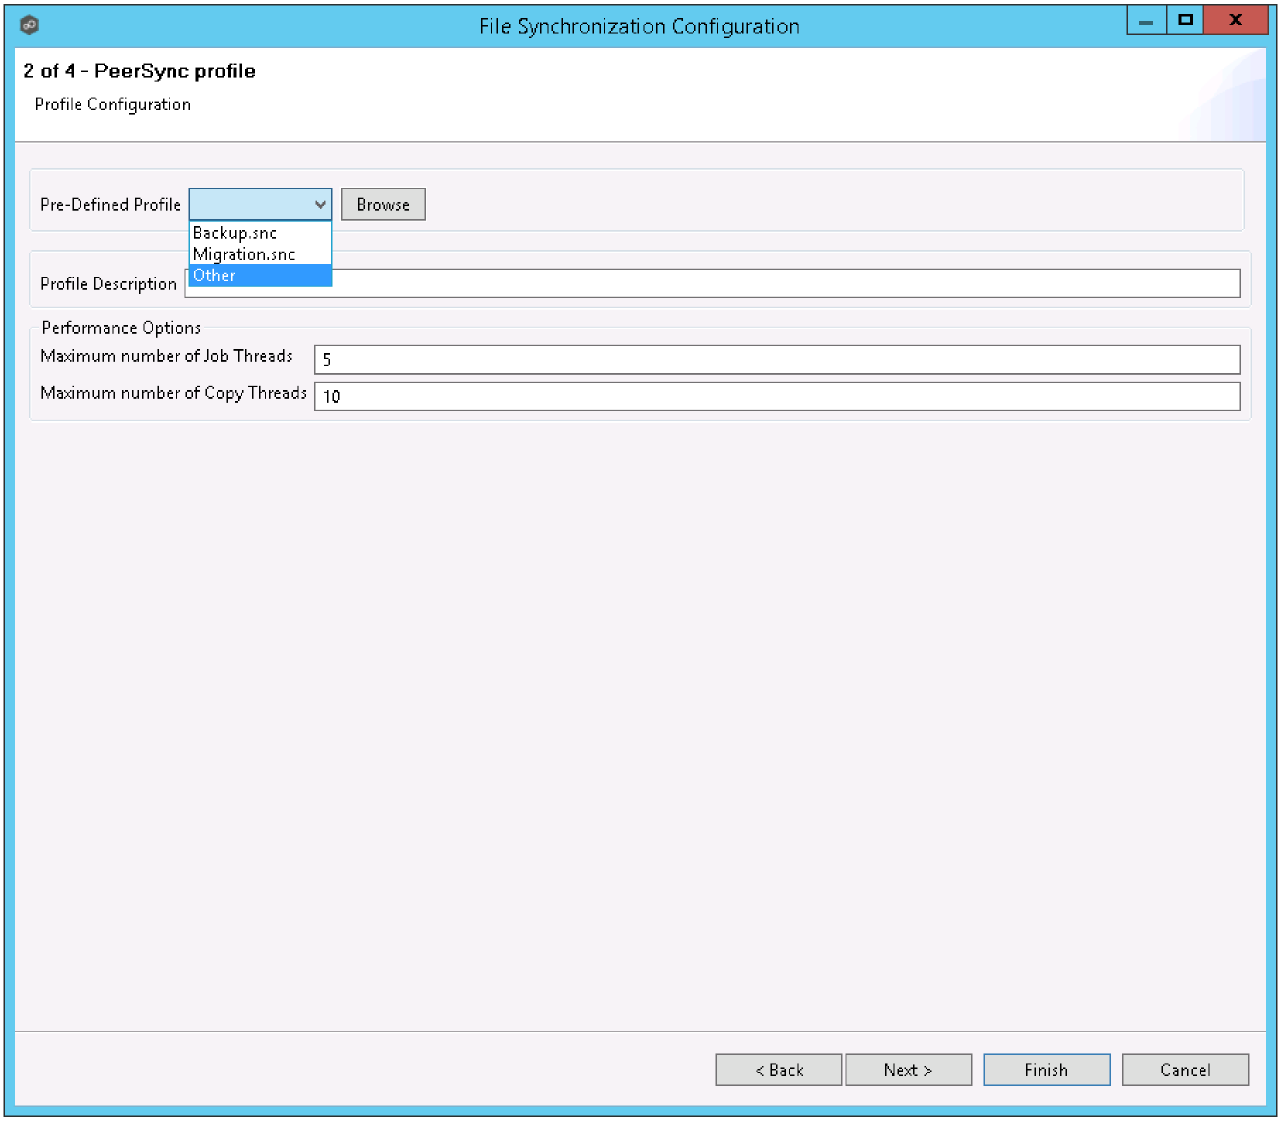

2 of 4 - PeerSync profile

Profile Configuration

Pre-Defined Profile: Backup.snc / Migration.snc / Other

Profile Description

Performance Options

Maximum number of Job Threads: 5

Maximum number of Copy Threads: 10

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	Maximum number of events that can be processed parallel to one another.

Profile Description	A textual description of the current profile.
number of Copy Threads	

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.

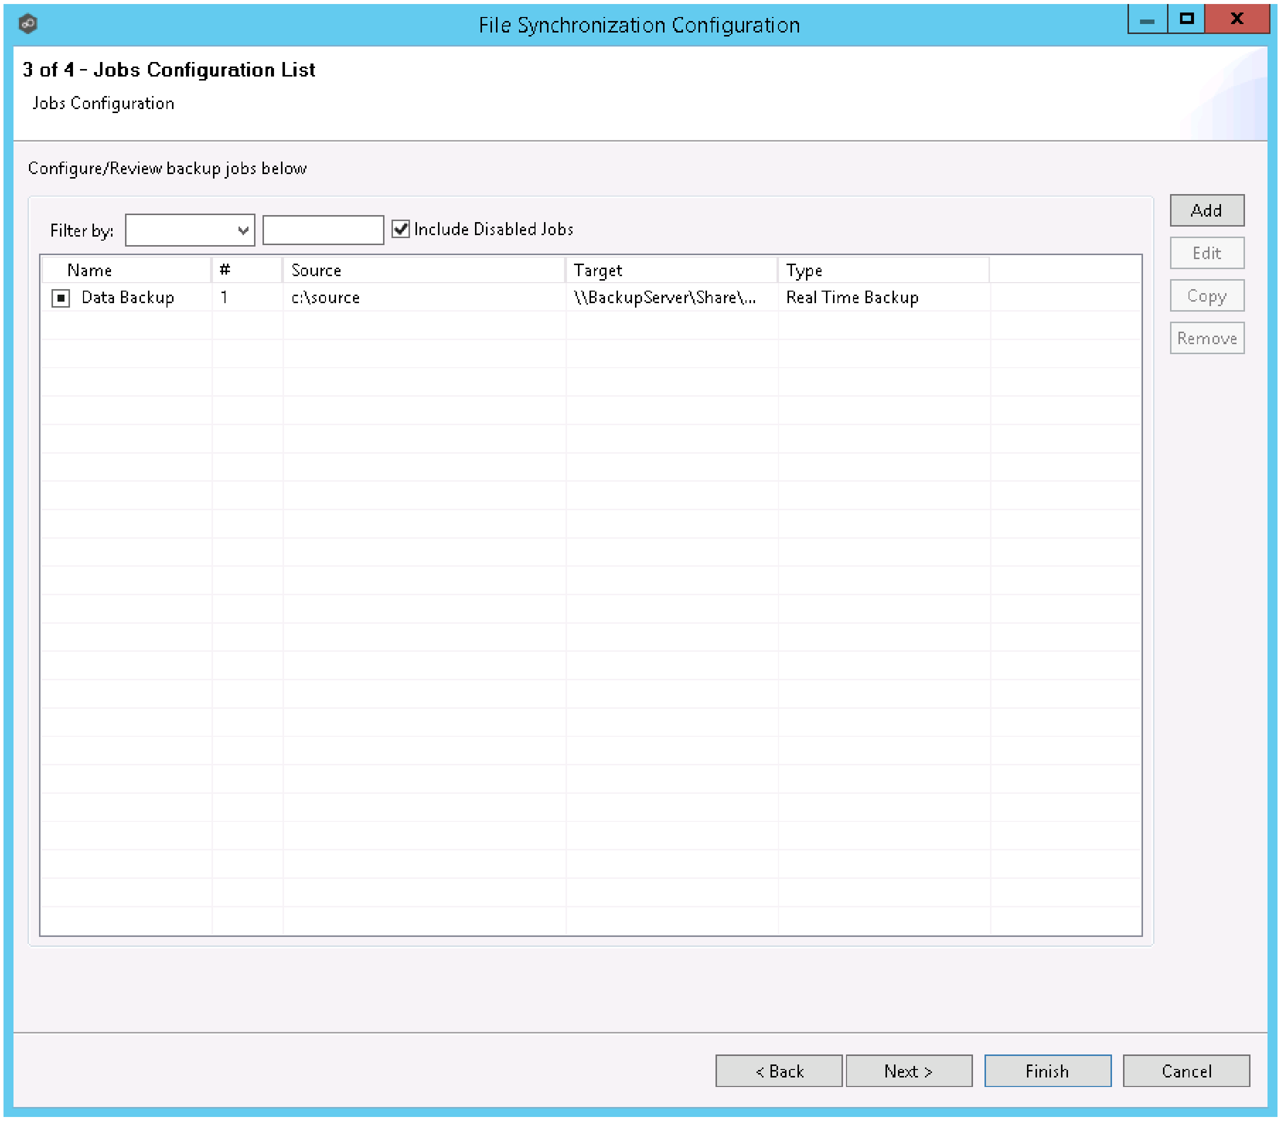

3 of 4 - Jobs Configuration List

Jobs Configuration

Configure/Review backup jobs below

Filter by: ☑ Include Disabled Jobs

Name	#	Source	Target	Type
■ Data Backup	1	c:\source	\\BackupServer\Share\...	Real Time Backup

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.

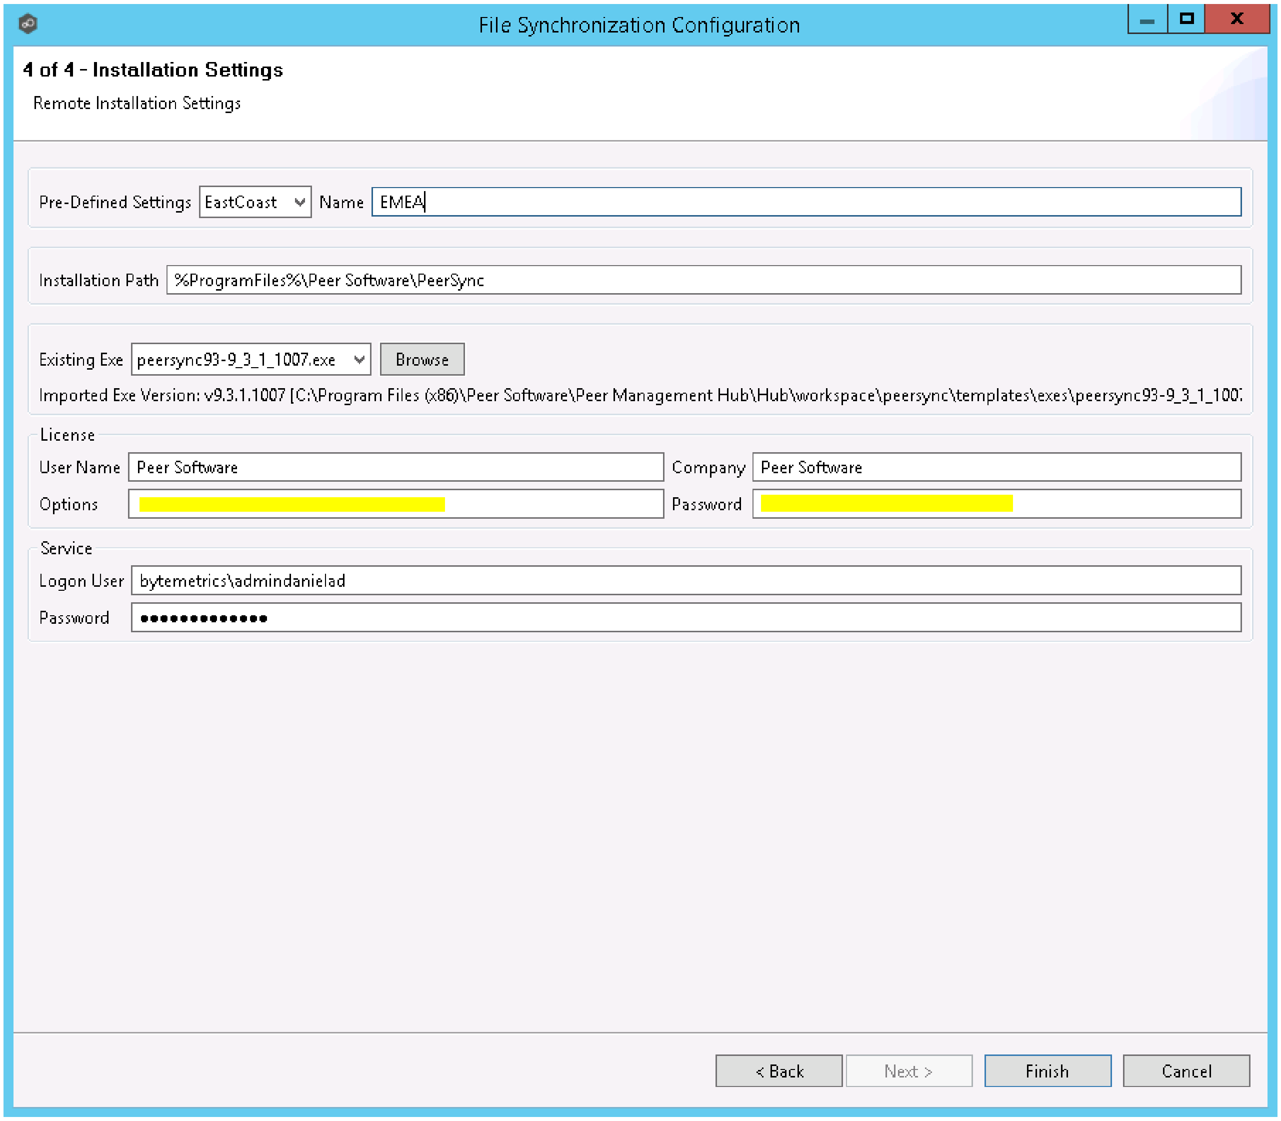

4 of 4 - Installation Settings

Remote Installation Settings

Pre-Defined Settings: EastCoast Name: EMEA

Installation Path: %ProgramFiles%\Peer Software\PeerSync

Existing Exe: peersync93-9_3_1_1007.exe

Imported Exe Version: v9.3.1.1007 [C:\Program Files (x86)\Peer Software\Peer Management Hub\Hub\workspace\peersync\templates\exes\peersync93-9_3_1_1007

License

User Name: Peer Software Company: Peer Software

Options: [redacted] Password: [redacted]

Service

Logon User: bytemetrics\admindanielad

Password: ●●●●●●●●●●●●●●

Pre-Defined 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 if 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.

Pre-Defined 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 if a new name has been given to the Installation Settings.
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 has to 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.

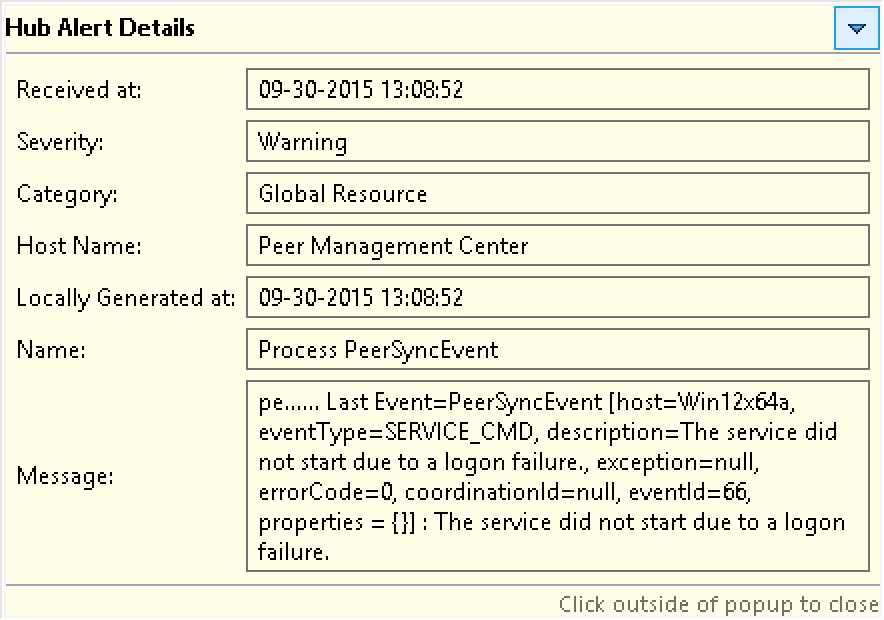

Hub Alert Details

Received at:	09-30-2015 13:08:52
Severity:	Warning
Category:	Global Resource
Host Name:	Peer Management Center
Locally Generated at:	09-30-2015 13:08:52
Name:	Process PeerSyncEvent
Message:	pe...... Last Event=PeerSyncEvent [host=Win12x64a, eventType=SERVICE_CMD, description=The service did not start due to a logon failure., exception=null, errorCode=0, coordinationId=null, eventId=66, properties = {}] : The service did not start due to a logon failure.

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.

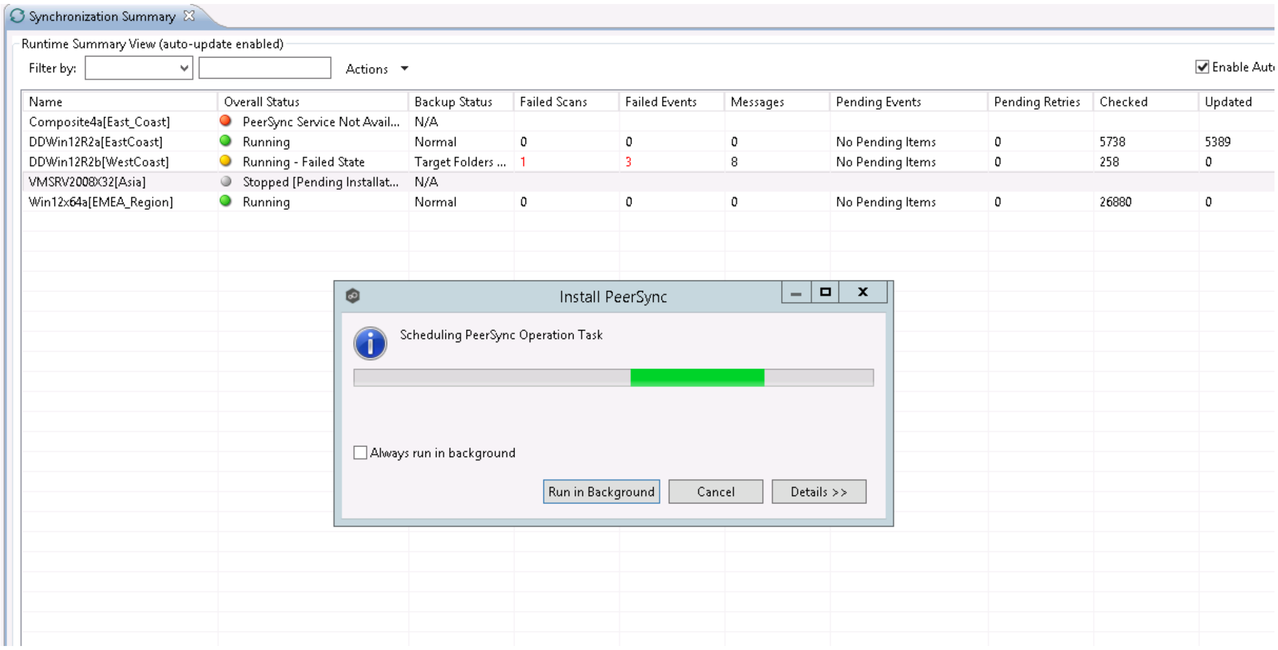

Name	Overall Status	Backup Status	Failed Scans	Failed Events	Messages	Pending Events	Pending Retries	Checked	Updated
Composite4a[East_Coast]	🔴 PeerSync Service Not Avail...	N/A							
DDWin12R2a[EastCoast]	🟢 Running	Normal	0	0	0	No Pending Items	0	5738	5389
DDWin12R2b[WestCoast]	🟡 Running - Failed State	Target Folders ...	1	3	8	No Pending Items	0	258	0
VMSRV2008X32[Asia]	● Stopped [Pending Installat...	N/A							
Win12x64a[EMEA_Region]	🟢 Running	Normal	0	0	0	No Pending Items	0	26880	0

Install PeerSync

Scheduling PeerSync Operation Task

□ Always run in background

Logging and Alerts

Use the following dialog to enable or disable logging and alerts, including specifying event types to log.

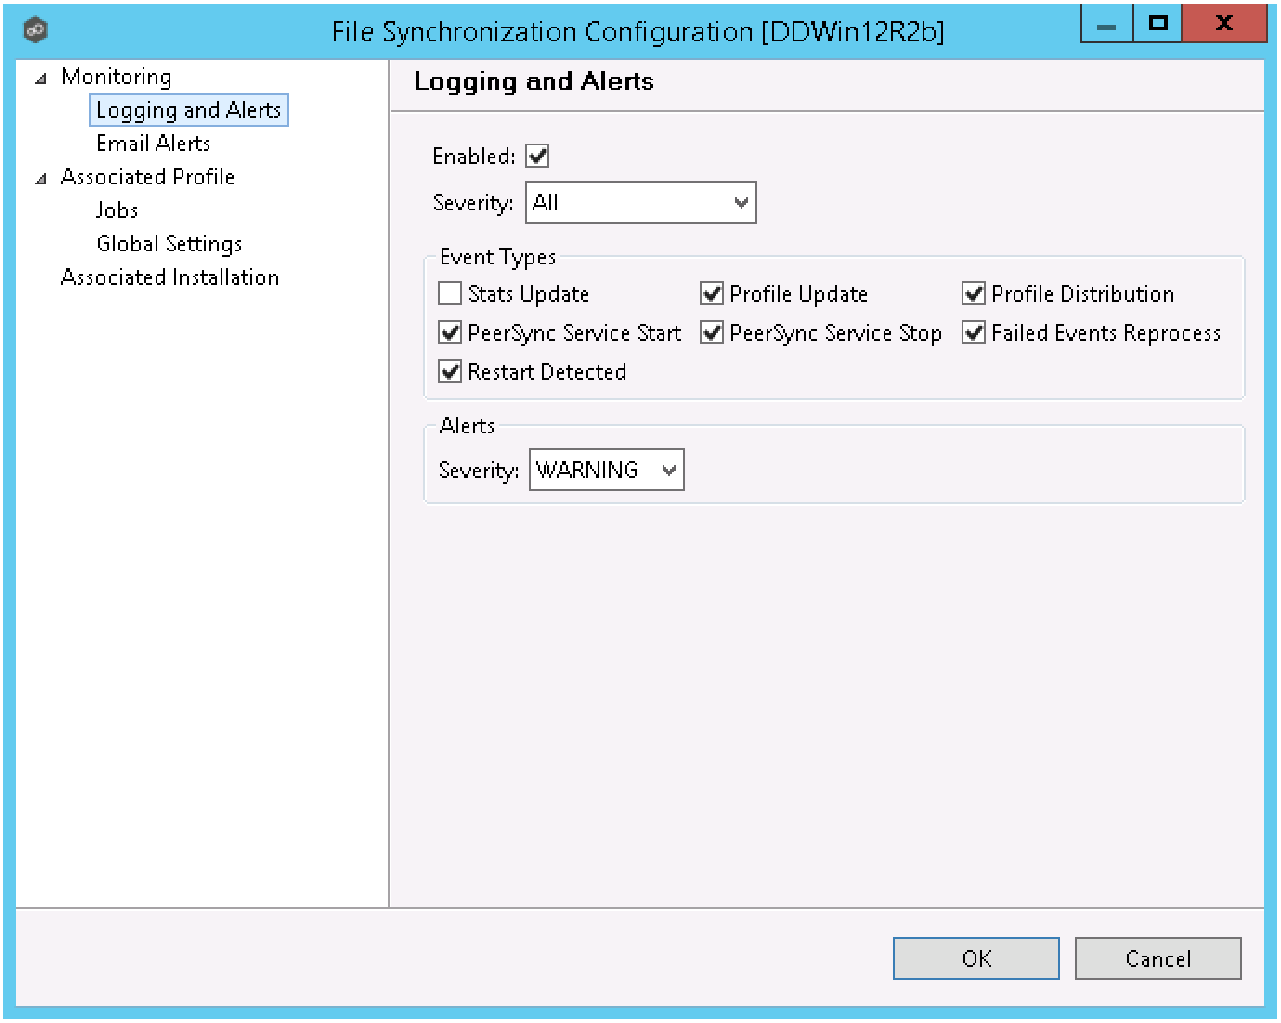

Logging and Alerts

Enabled: ☑

Severity: All

Event Types

☐ Stats Update ☑ Profile Update ☑ Profile Distribution
☑ PeerSync Service Start ☑ PeerSync Service Stop ☑ Failed Events Reprocess
☑ Restart Detected

Alerts

Severity: WARNING

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.

Stats Update	Log when PeerSync Stats are received (This could generate large amount of Log Entries).
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.

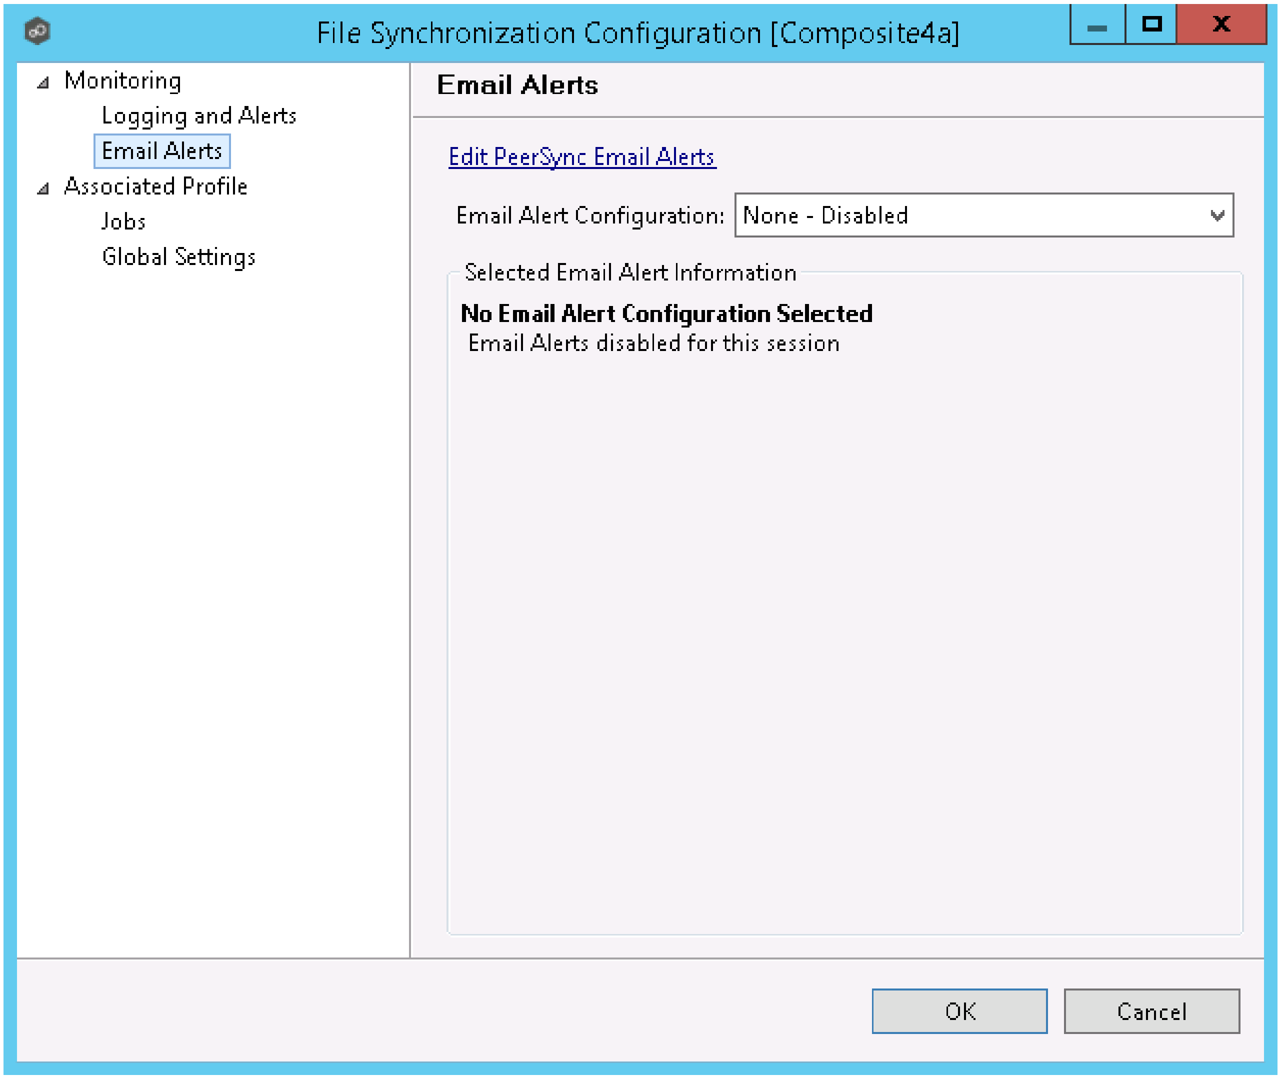

Email Alert Configuration: None - Disabled

Selected Email Alert Information

No Email Alert Configuration Selected
Email Alerts disabled for this session

To enable email alerts for this particular 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.

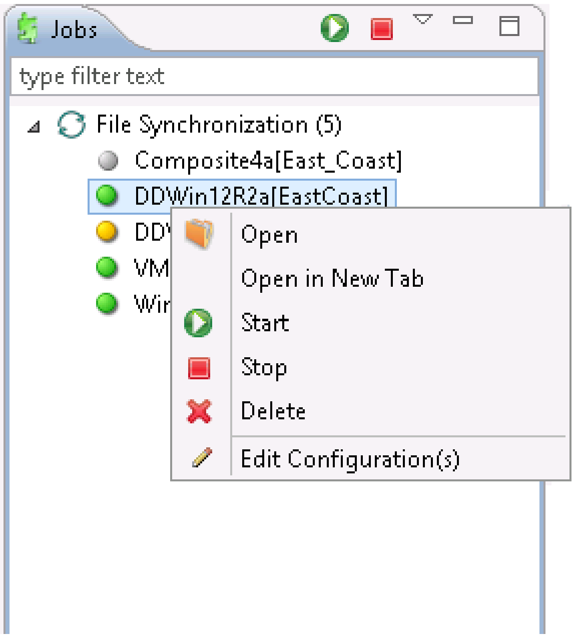

PeerSync Management Summary

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.

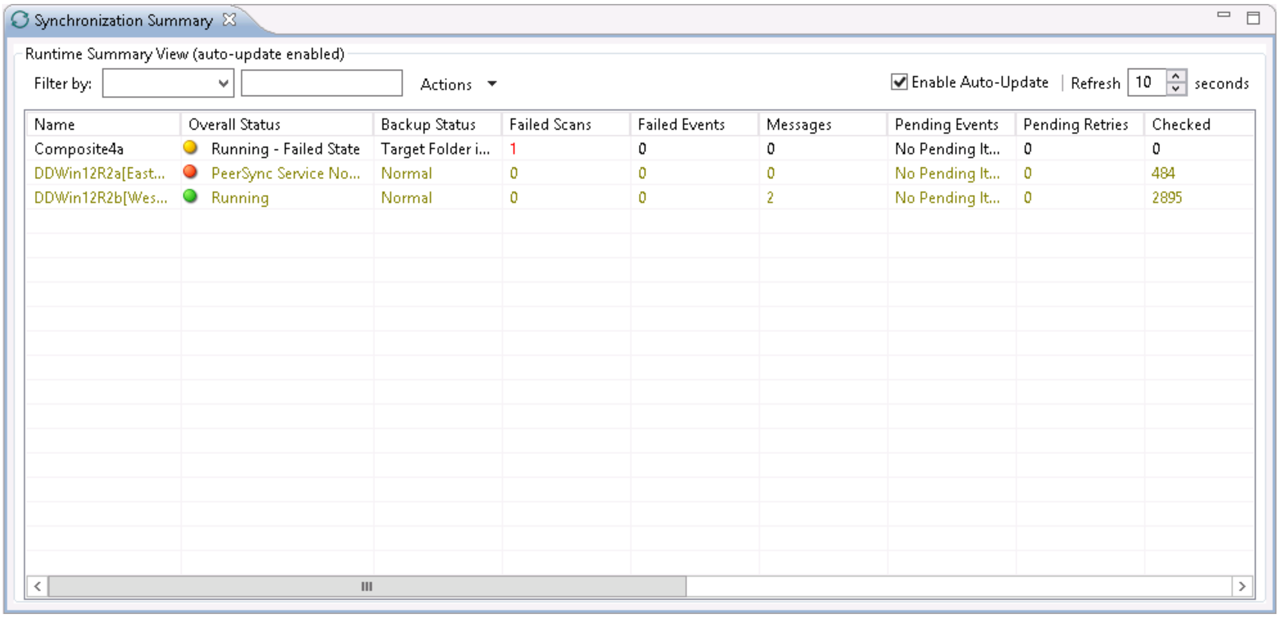

Runtime Summary View (auto-update enabled)

☑ Enable Auto-Update | Refresh 10 seconds

Name	Overall Status	Backup Status	Failed Scans	Failed Events	Messages	Pending Events	Pending Retries	Checked
Composite4a	● Running - Failed State	Target Folder i...	1	0	0	No Pending It...	0	0
DDWin12R2a[East...	● PeerSync Service No...	Normal	0	0	0	No Pending It...	0	484
DDWin12R2b[Wes...	● Running	Normal	0	0	2	No Pending It...	0	2895

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 check box. 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.

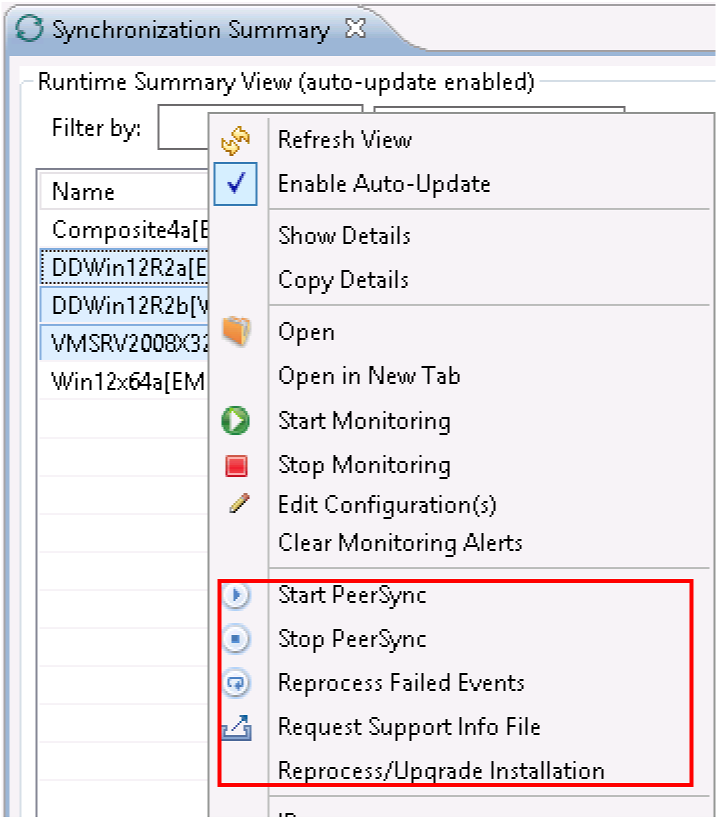

Start PeerSync	Send a Start command to the PeerSync service instance running on the selected jobs' participant. This can be performed while the associated File Synchronization Job is running.
Stop PeerSync	Send a Stop command to the PeerSync service instance running on the selected jobs' participant. 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 document editor.
Export Entire Table to File	Dump the entire contents of the table to a text file that can be viewed in any document 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.

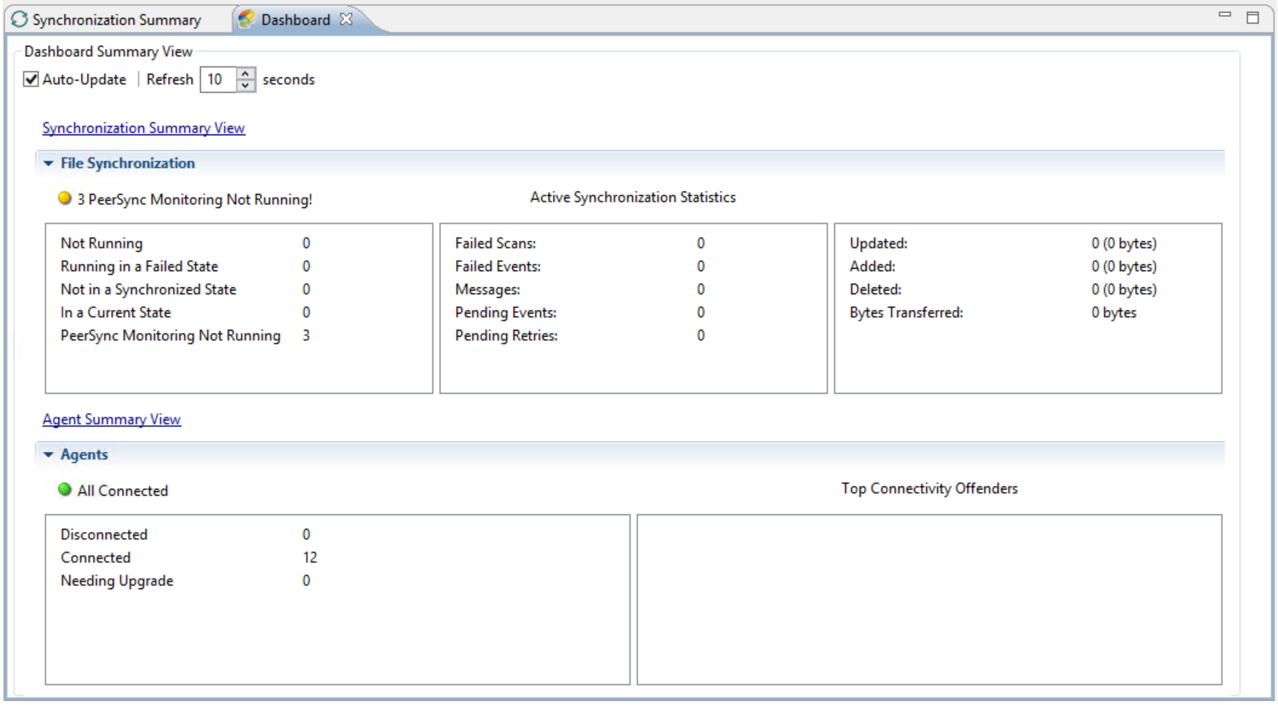

Dashboard Summary View

☑ Auto-Update | Refresh 10 seconds

Synchronization Summary View

File Synchronization

● 3 PeerSync Monitoring Not Running!

Active Synchronization Statistics

Not Running	0	Failed Scans:	0	Updated:	0 (0 bytes)
Running in a Failed State	0	Failed Events:	0	Added:	0 (0 bytes)
Not in a Synchronized State	0	Messages:	0	Deleted:	0 (0 bytes)
In a Current State	0	Pending Events:	0	Bytes Transferred:	0 bytes
PeerSync Monitoring Not Running	3	Pending Retries:	0		

Agent Summary View

Agents

● All Connected

Top Connectivity Offenders

Disconnected	0
Connected	12
Needing Upgrade	0

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 check box.

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 the 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 in the PMC client through the **Import** button.

- After having made your entire configuration changes either through the PMC client 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 distribute to the PeerSync host through the Distribute button.

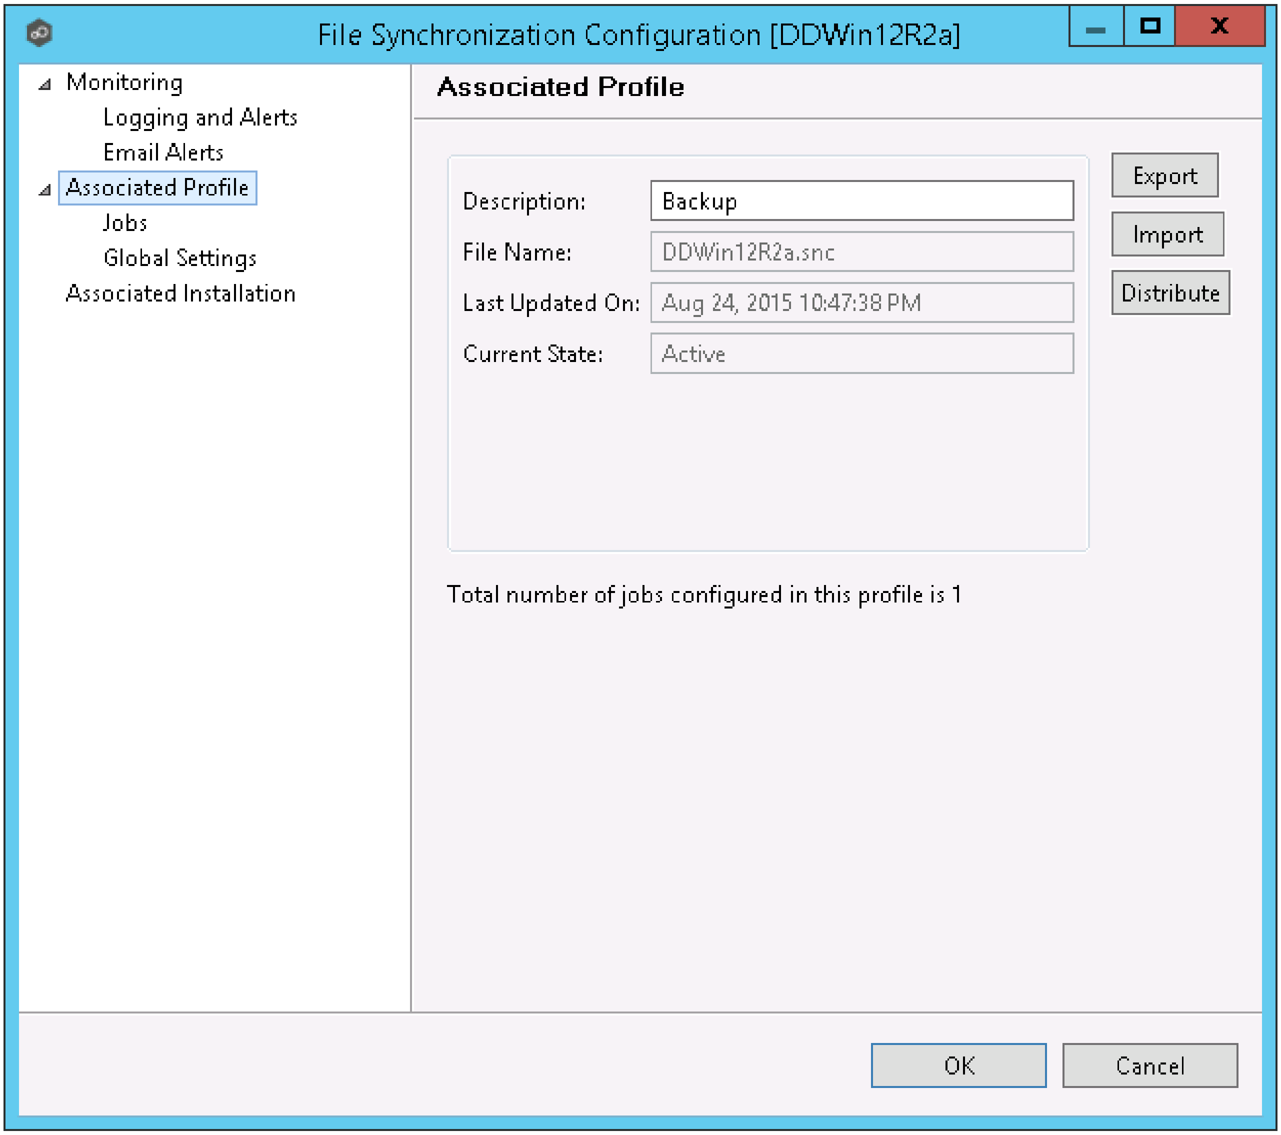

File Synchronization Configuration [DDWin12R2a]

Associated Profile

Description: Backup

File Name: DDWin12R2a.snc

Last Updated On: Aug 24, 2015 10:47:38 PM

Current State: Active

Total number of jobs configured in this profile is 1

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

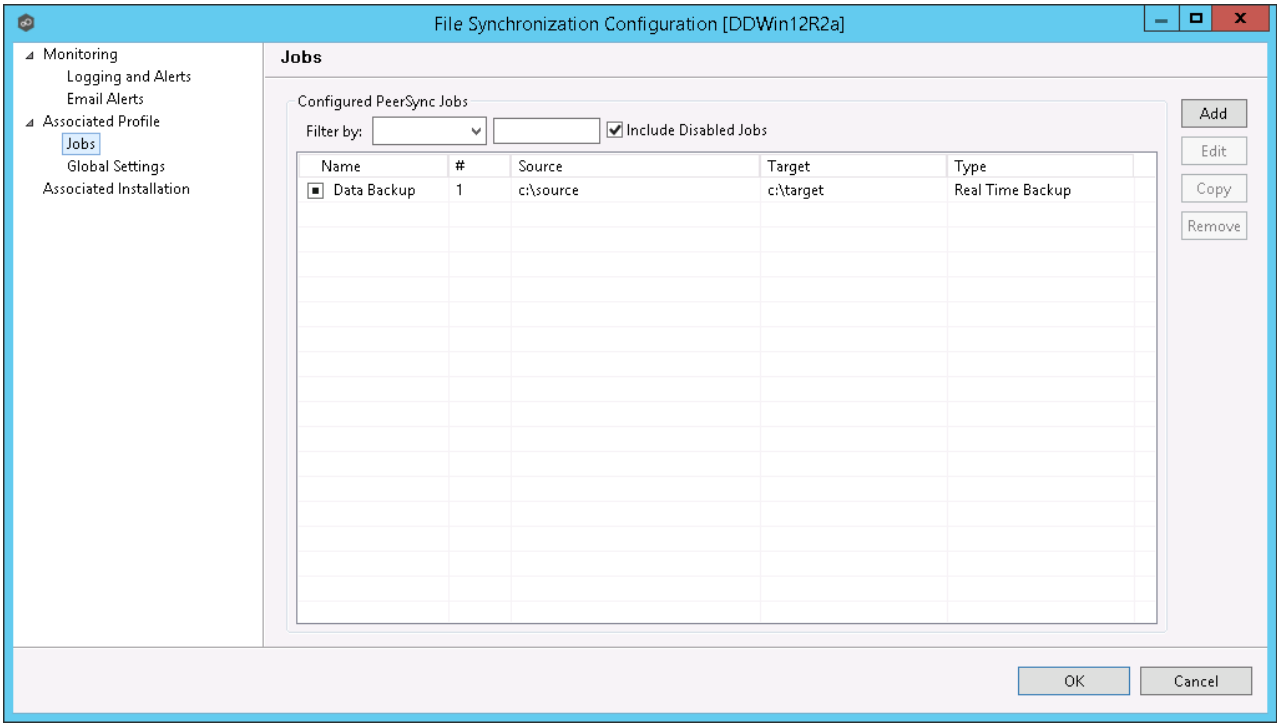

Filter by:

☑ Include Disabled Jobs

Name	#	Source	Target	Type
▣ Data Backup	1	c:\source	c:\target	Real Time Backup

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.

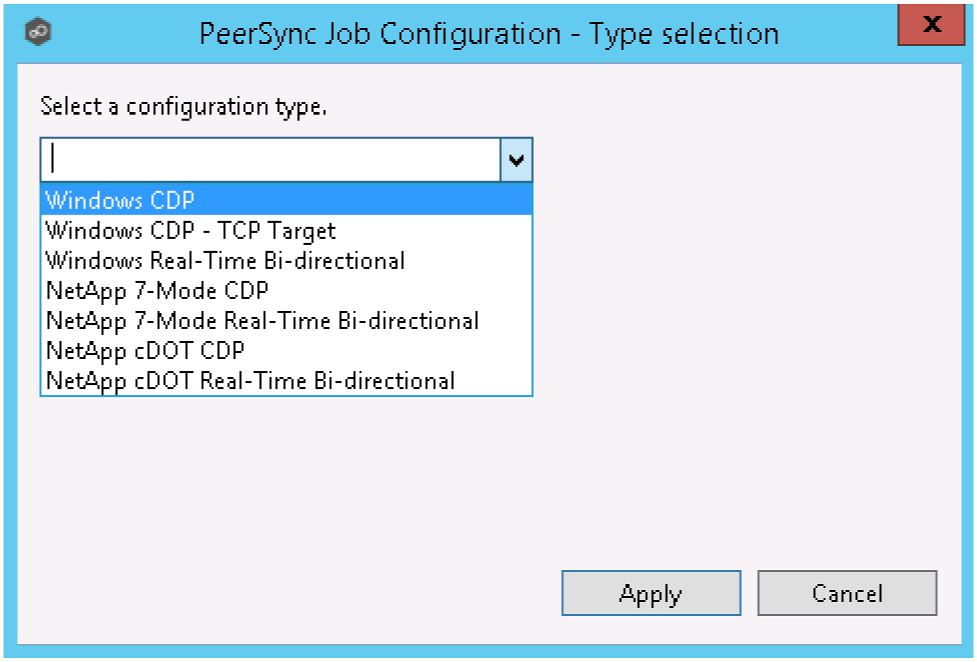

PeerSync Job Configuration - Type selection
Select a configuration type.
Windows CDP
Windows CDP - TCP Target
Windows Real-Time Bi-directional
NetApp 7-Mode CDP
NetApp 7-Mode Real-Time Bi-directional
NetApp cDOT CDP
NetApp cDOT Real-Time Bi-directional
Apply
Cancel

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 a left menu tree.

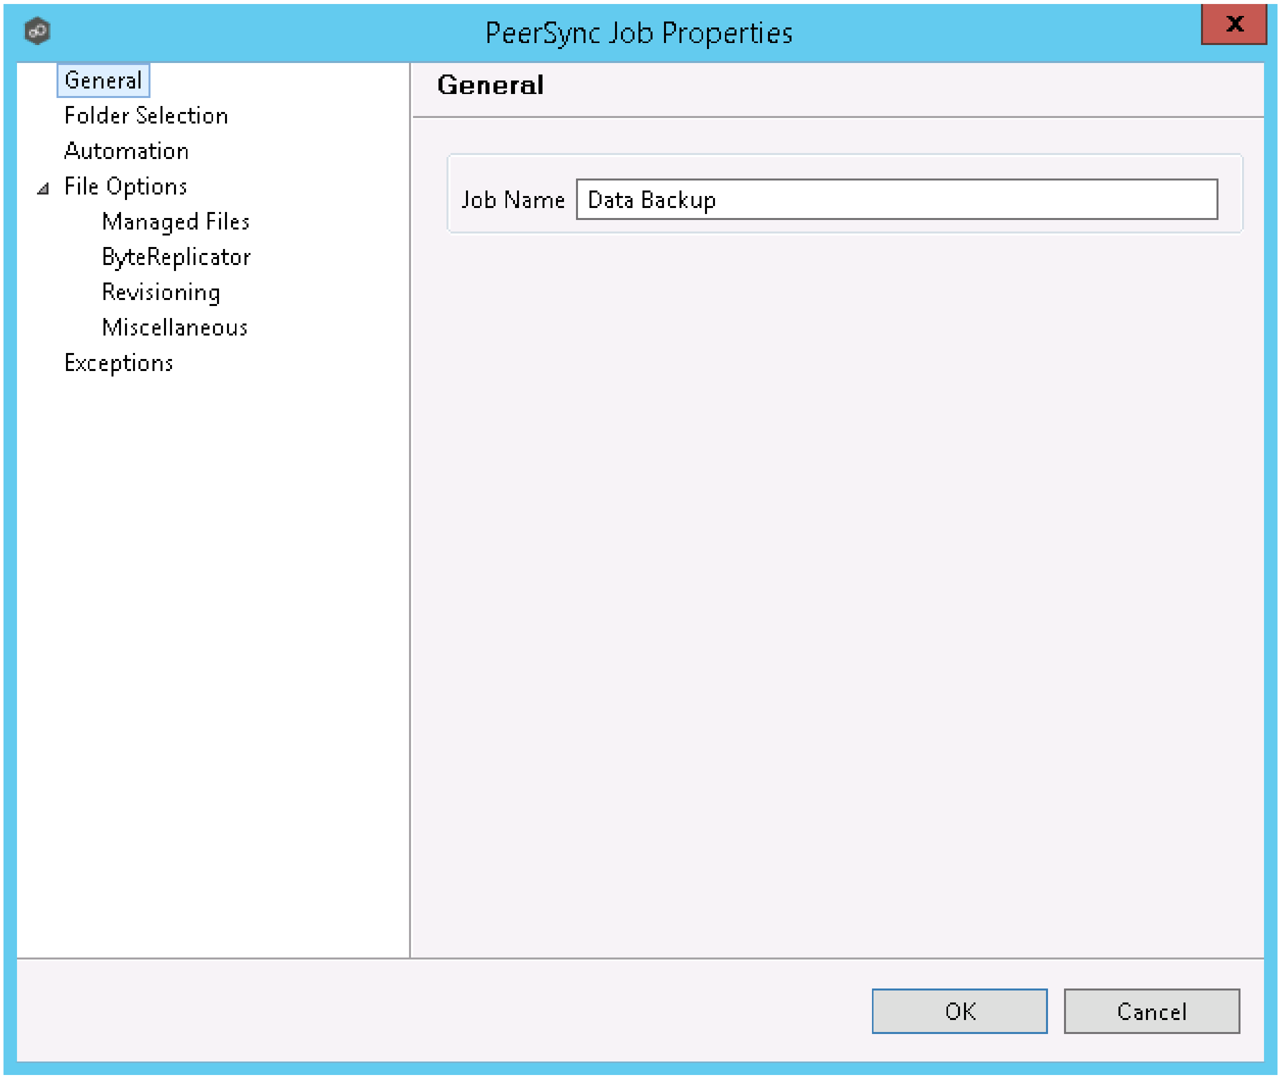
PeerSync Job Properties

General

Job Name: Data Backup

Enable/Disable Job

To enable or disable a job, click the check box 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 have to 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.

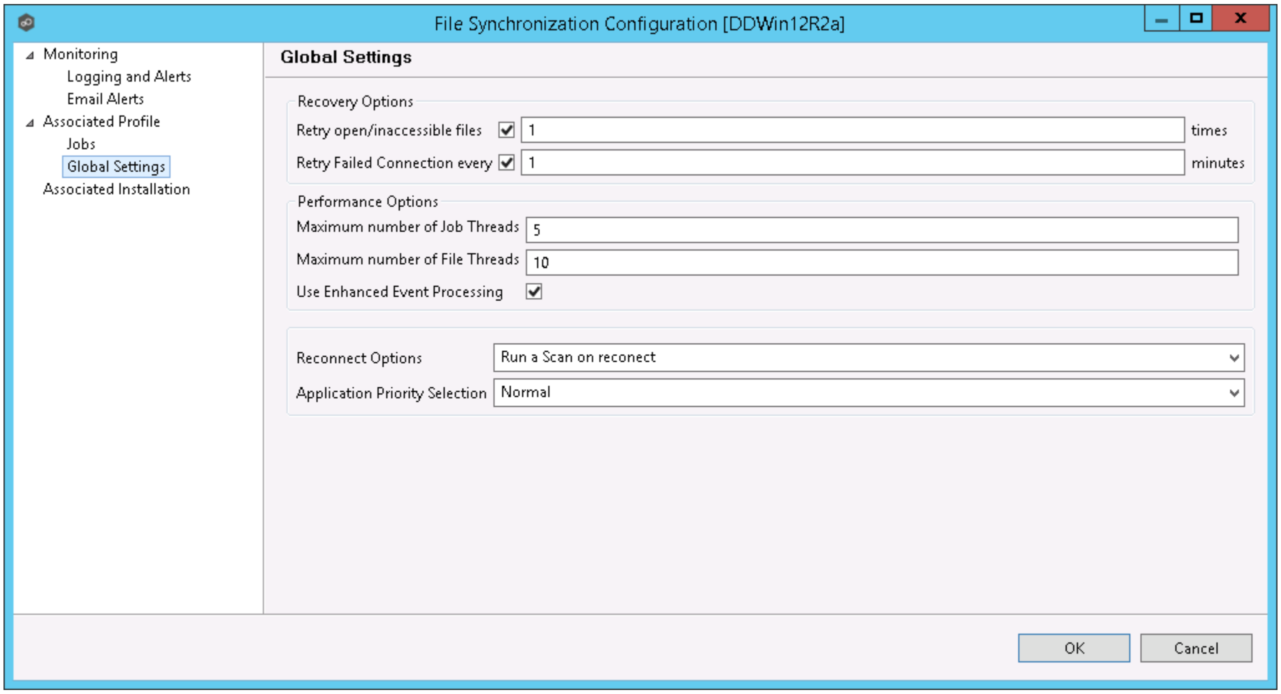

Global Settings

Recovery Options

Retry open/inaccessible files ☑ 1 times

Retry Failed Connection every ☑ 1 minutes

Performance Options

Maximum number of Job Threads 5

Maximum number of File Threads 10

Use Enhanced Event Processing ☑

Reconnect Options: Run a Scan on reconect

Application Priority Selection: Normal

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.

Recovery Options	These changes will update how we retry failed or inaccessible files as well as the interval in which we retry Failed Connections.
Applicati on 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 along side 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.

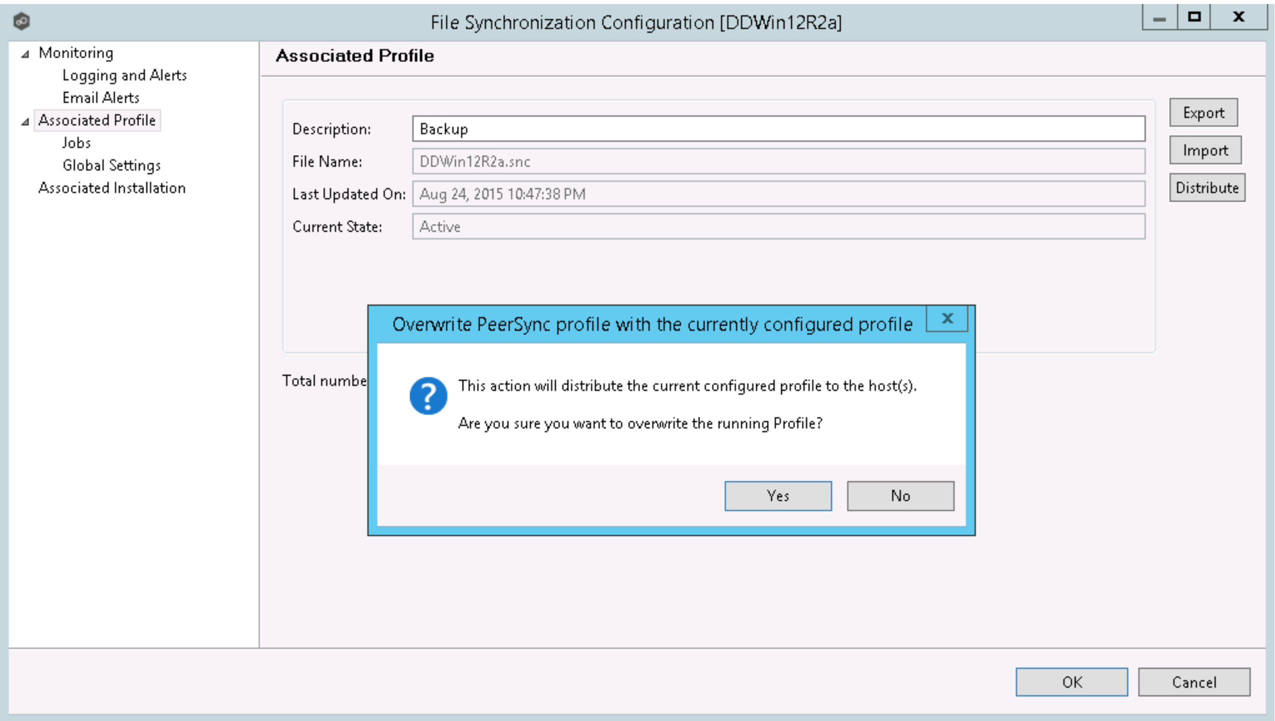

Description: Backup
File Name: DDWin12R2a.snc
Last Updated On: Aug 24, 2015 10:47:38 PM
Current State: Active

Overwrite PeerSync profile with the currently configured profile

This action will distribute the current configured profile to the host(s).

Are you sure you want to overwrite the running Profile?

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 has to be running in order to successfully perform this action.

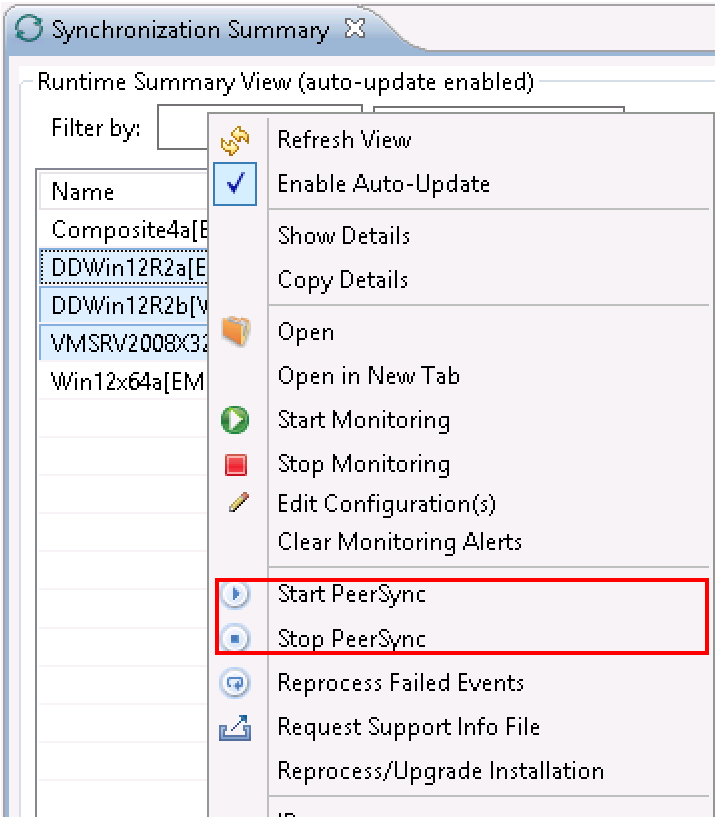

For information on the additional PeerSync multi-job global actions, see Synchronization Summary View.

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

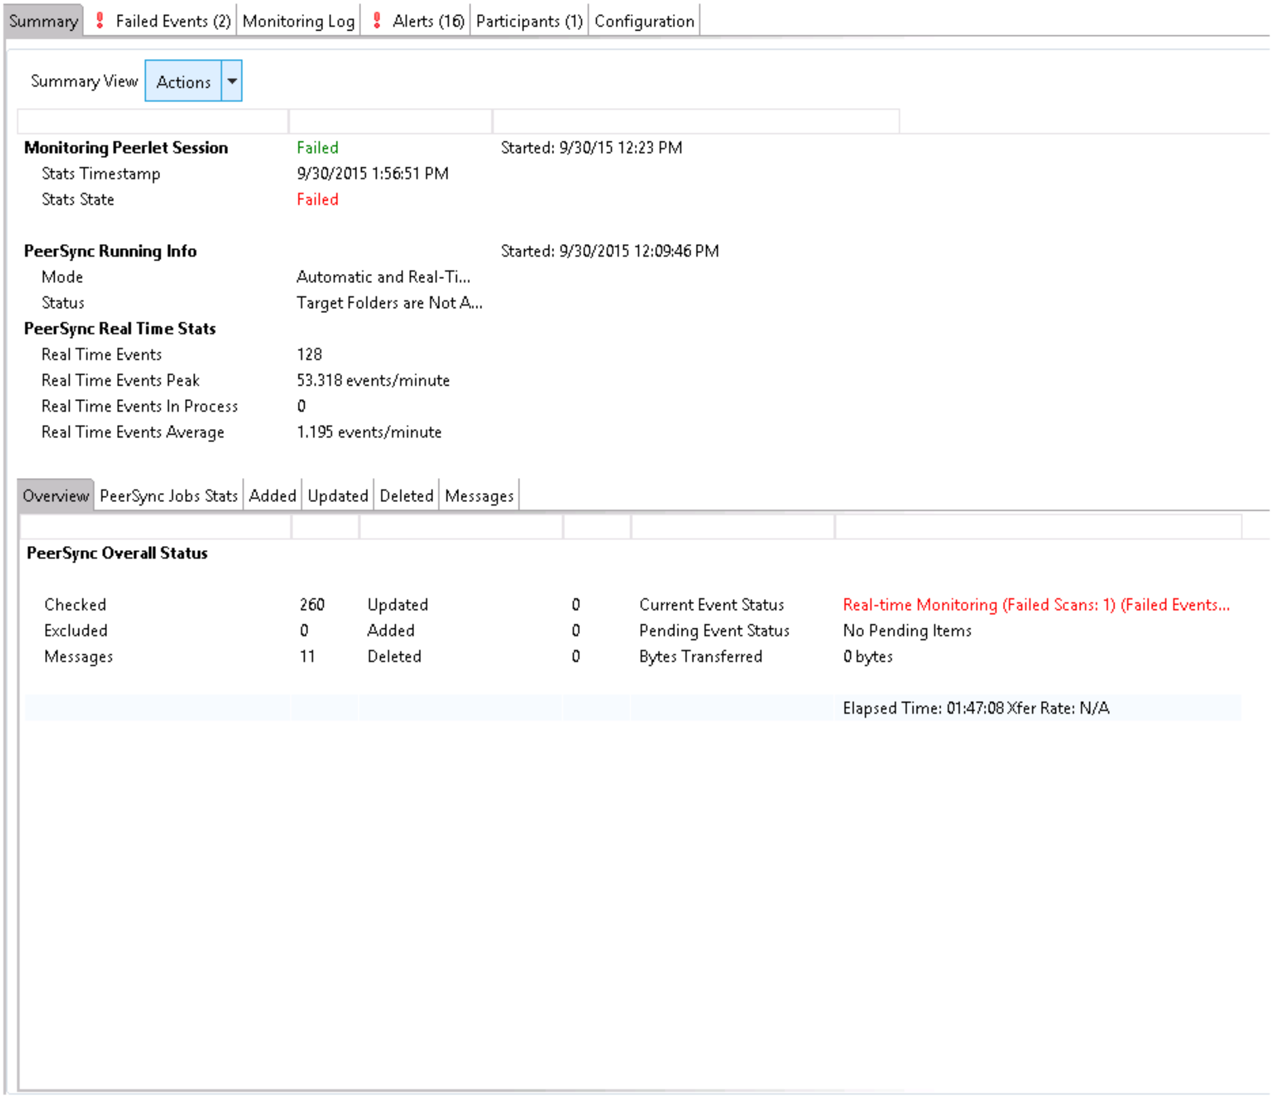

Monitoring Peerlet Session	Failed	Started: 9/30/15 12:23 PM
Stats Timestamp	9/30/2015 1:56:51 PM	
Stats State	Failed	
PeerSync Running Info		Started: 9/30/2015 12:09:46 PM
Mode	Automatic and Real-Ti...	
Status	Target Folders are Not A...	
PeerSync Real Time Stats		
Real Time Events	128	
Real Time Events Peak	53.318 events/minute	
Real Time Events In Process	0	
Real Time Events Average	1.195 events/minute	

PeerSync Overall Status

Checked	260	Updated	0	Current Event Status	Real-time Monitoring (Failed Scans: 1) (Failed Events...
Excluded	0	Added	0	Pending Event Status	No Pending Items
Messages	11	Deleted	0	Bytes Transferred	0 bytes
					Elapsed Time: 01:47:08 Xfer Rate: N/A

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.

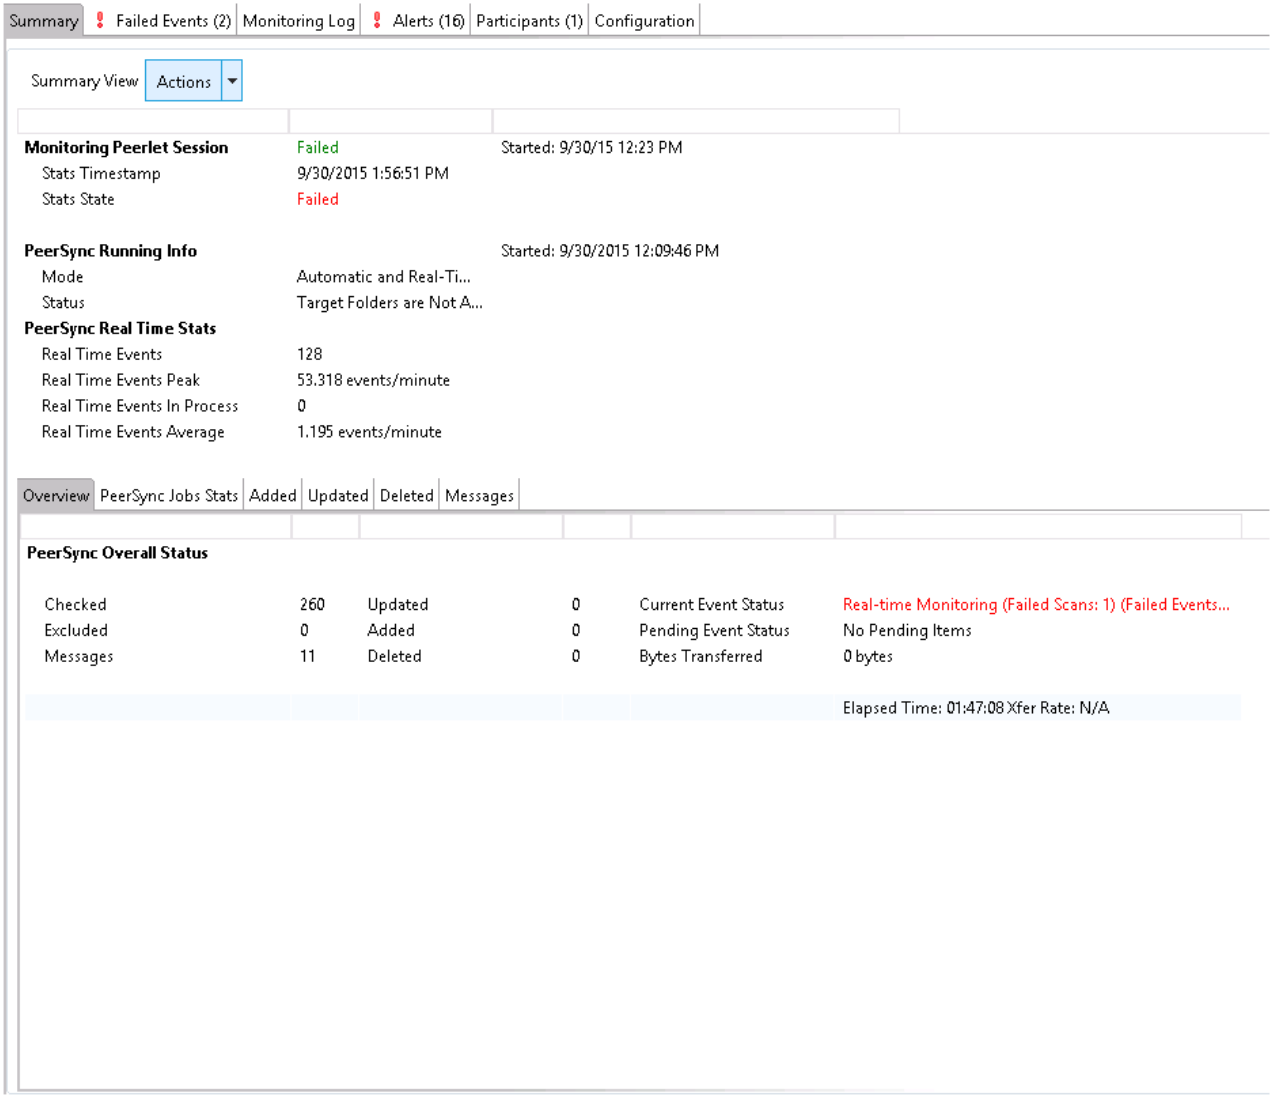

Monitoring Peerlet Session	Failed	Started: 9/30/15 12:23 PM
Stats Timestamp	9/30/2015 1:56:51 PM	
Stats State	Failed	
PeerSync Running Info		Started: 9/30/2015 12:09:46 PM
Mode	Automatic and Real-Ti...	
Status	Target Folders are Not A...	
PeerSync Real Time Stats		
Real Time Events	128	
Real Time Events Peak	53.318 events/minute	
Real Time Events In Process	0	
Real Time Events Average	1.195 events/minute	

PeerSync Overall Status

Checked	260	Updated	0	Current Event Status	Real-time Monitoring (Failed Scans: 1) (Failed Events...
Excluded	0	Added	0	Pending Event Status	No Pending Items
Messages	11	Deleted	0	Bytes Transferred	0 bytes
					Elapsed Time: 01:47:08 Xfer Rate: N/A

On this page, you can right-click to display the PeerSync **Actions** menu:

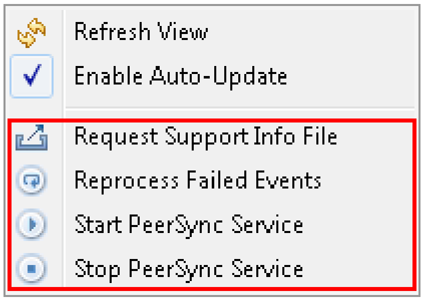

On the bottom half of the page, you will find a set of tabs showing more granular information regarding this PeerSync session.

- 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.

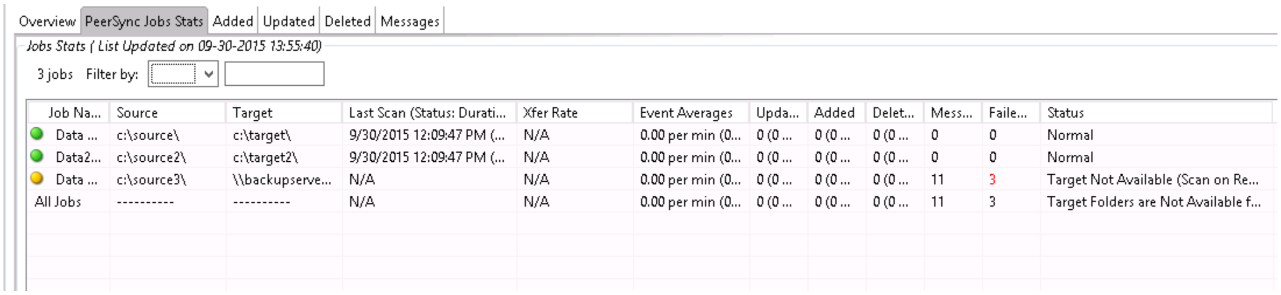

Jobs Stats (List Updated on 09-30-2015 13:55:40)

3 jobs Filter by:

Job Na...	Source	Target	Last Scan (Status: Durati...	Xfer Rate	Event Averages	Upda...	Added	Delet...	Mess...	Faile...	Status
● Data ...	c:\source\	c:\target\	9/30/2015 12:09:47 PM (...	N/A	0.00 per min (0...	0 (0 ...	0 (0 ...	0 (0 ...	0	0	Normal
● Data2...	c:\source2\	c:\target2\	9/30/2015 12:09:47 PM (...	N/A	0.00 per min (0...	0 (0 ...	0 (0 ...	0 (0 ...	0	0	Normal
● Data ...	c:\source3\	\\backupserve...	N/A	N/A	0.00 per min (0...	0 (0 ...	0 (0 ...	0 (0 ...	11	3	Target Not Available (Scan on Re...
All Jobs	----------	----------	N/A	N/A	0.00 per min (0...	0 (0 ...	0 (0 ...	0 (0 ...	11	3	Target Folders are Not Available f...

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.

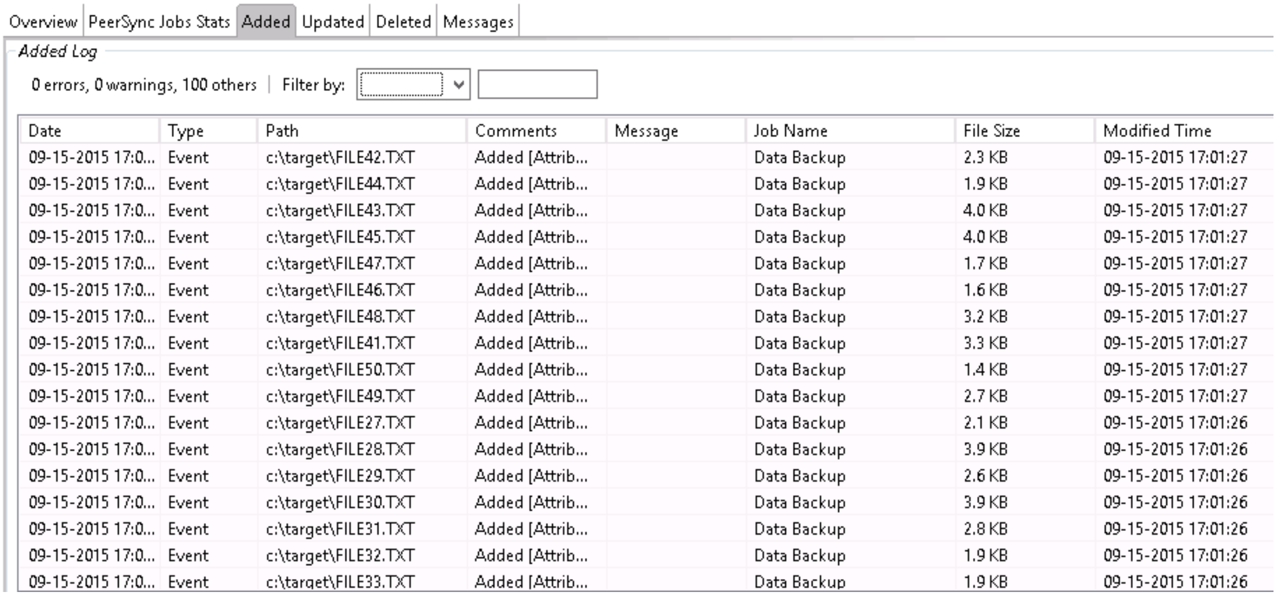

Added Log

0 errors, 0 warnings, 100 others | Filter by:

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.

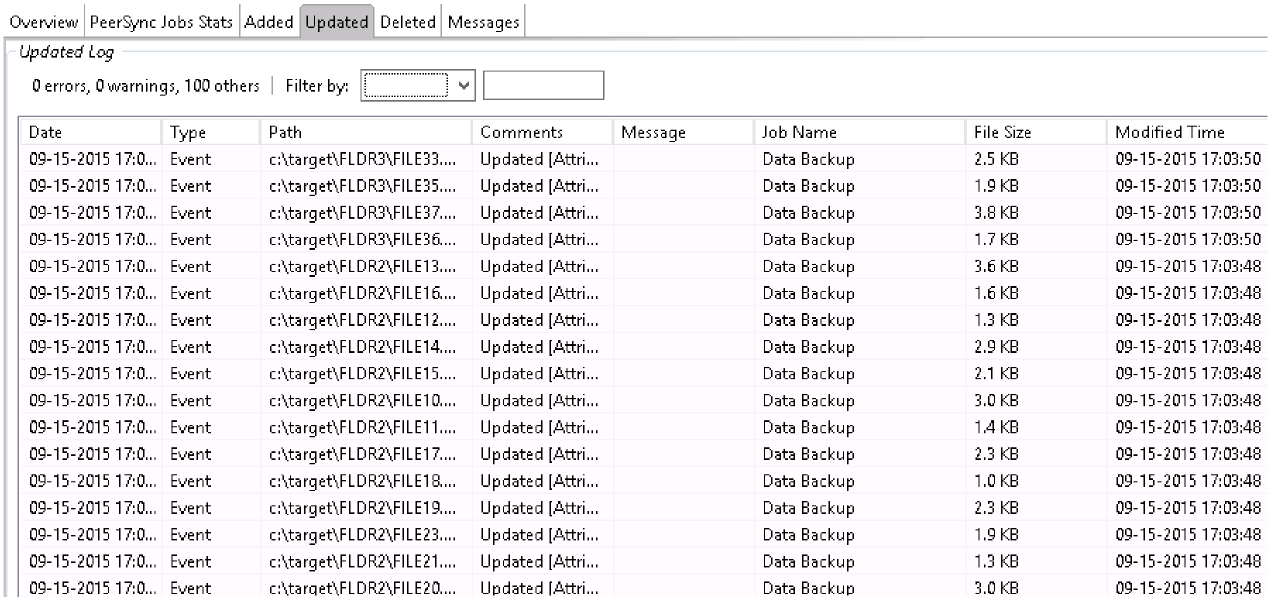

Updated Log

0 errors, 0 warnings, 100 others | Filter by:

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.

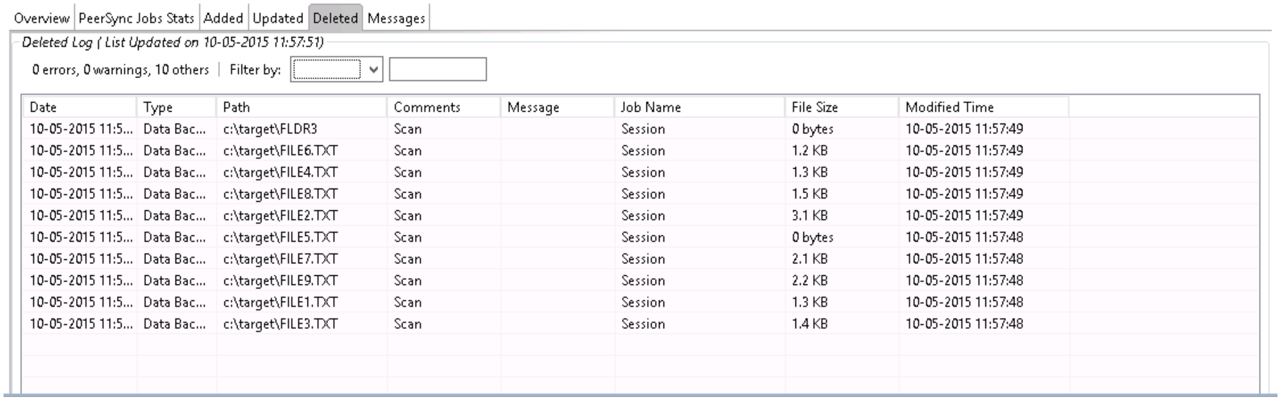

Deleted Log (List Updated on 10-05-2015 11:57:51)

0 errors, 0 warnings, 10 others | Filter by:

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.

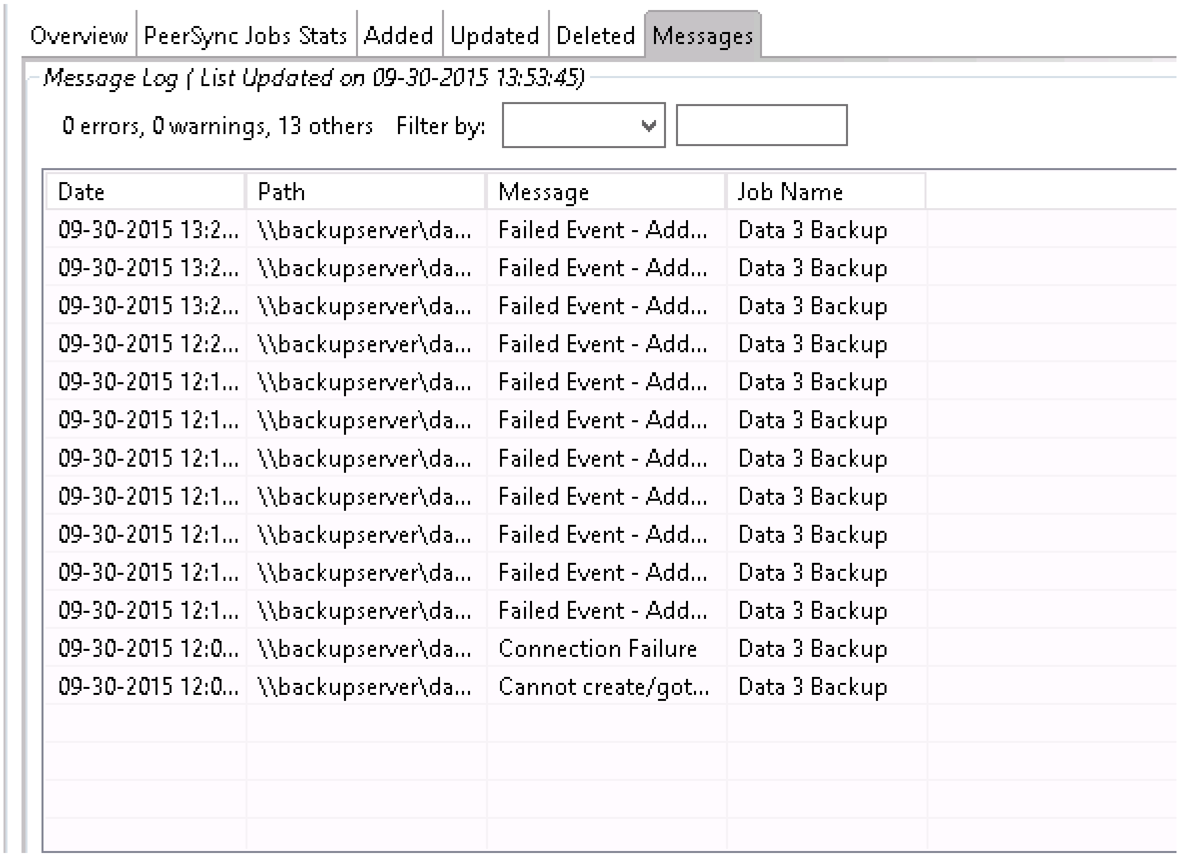

Message Log (List Updated on 09-30-2015 13:53:45)

0 errors, 0 warnings, 13 others Filter by:

Date	Path	Message	Job Name
09-30-2015 13:2...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 13:2...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 13:2...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 12:2...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup
09-30-2015 12:0...	\\backupserver\da...	Connection Failure	Data 3 Backup
09-30-2015 12:0...	\\backupserver\da...	Cannot create/got...	Data 3 Backup

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.

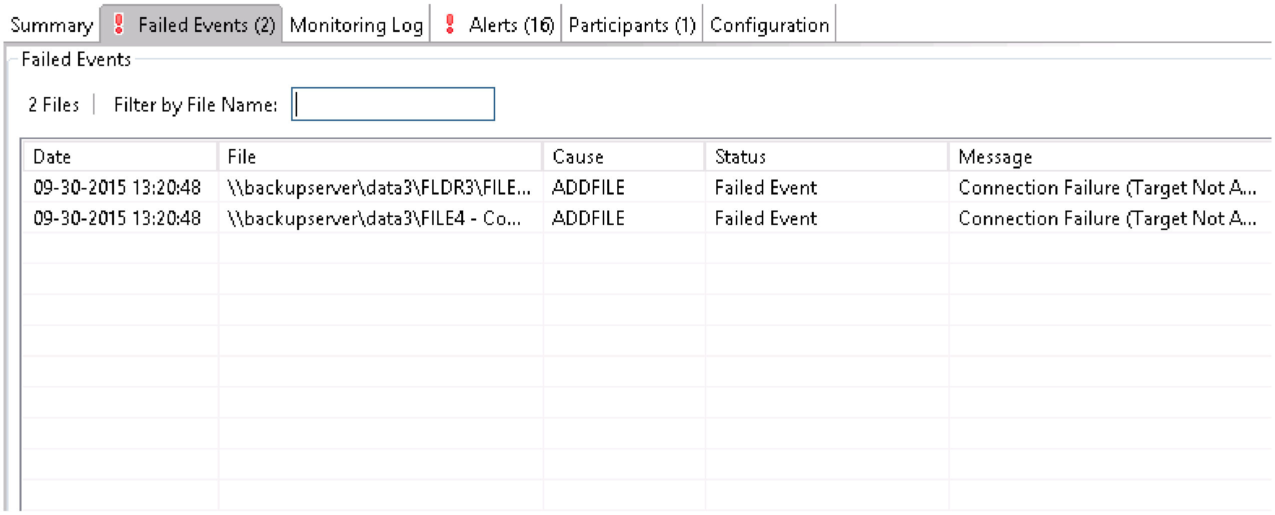

2 Files | Filter by File Name:

Date	File	Cause	Status	Message
09-30-2015 13:20:48	\\backupserver\data3\FLDR3\FILE...	ADDFILE	Failed Event	Connection Failure (Target Not A...
09-30-2015 13:20:48	\\backupserver\data3\FILE4 - Co...	ADDFILE	Failed Event	Connection Failure (Target Not A...

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.

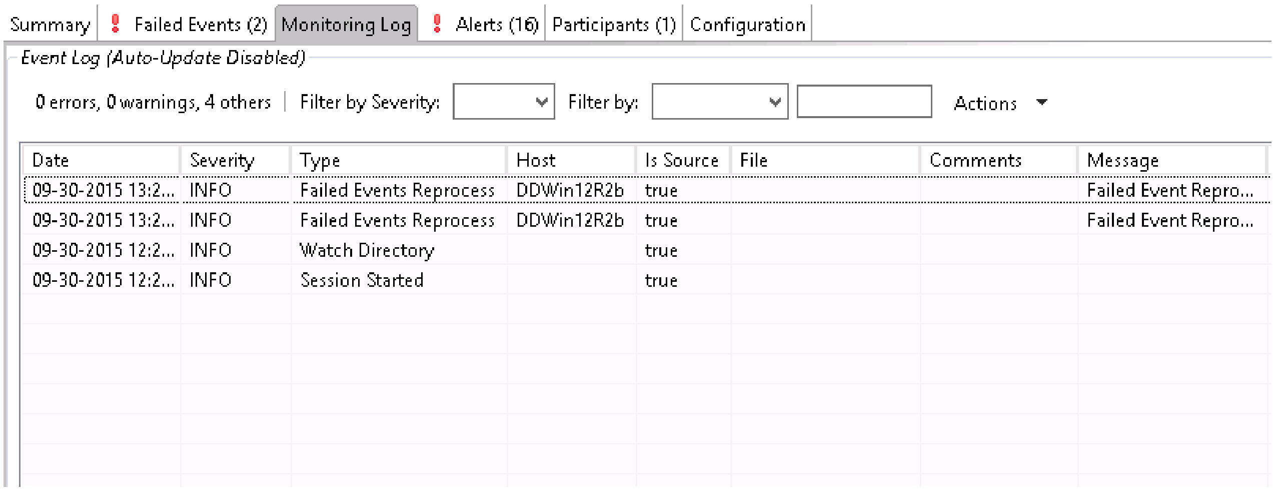

Date	Severity	Type	Host	Is Source	File	Comments	Message
09-30-2015 13:2...	INFO	Failed Events Reprocess	DDWin12R2b	true			Failed Event Repro...
09-30-2015 13:2...	INFO	Failed Events Reprocess	DDWin12R2b	true			Failed Event Repro...
09-30-2015 12:2...	INFO	Watch Directory		true			
09-30-2015 12:2...	INFO	Session Started		true			

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.

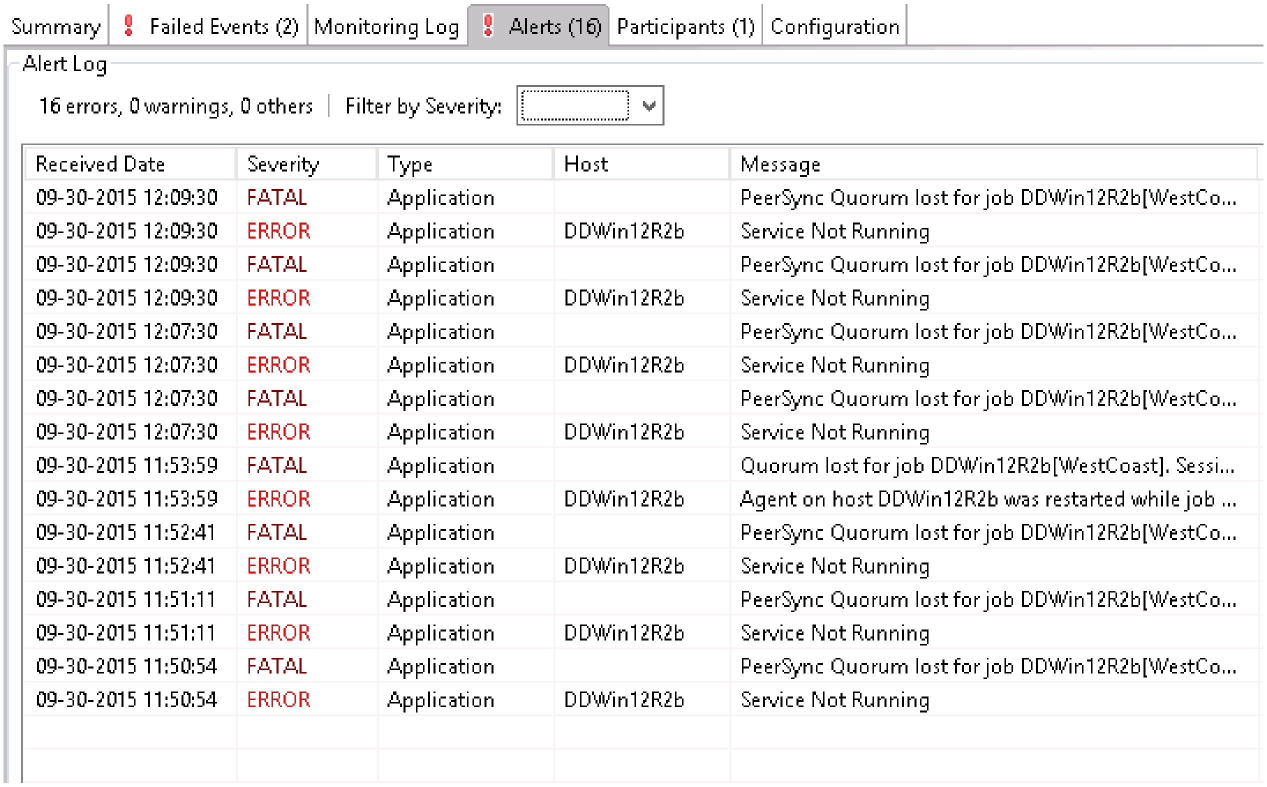

16 errors, 0 warnings, 0 others | Filter by Severity:

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 particular 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.

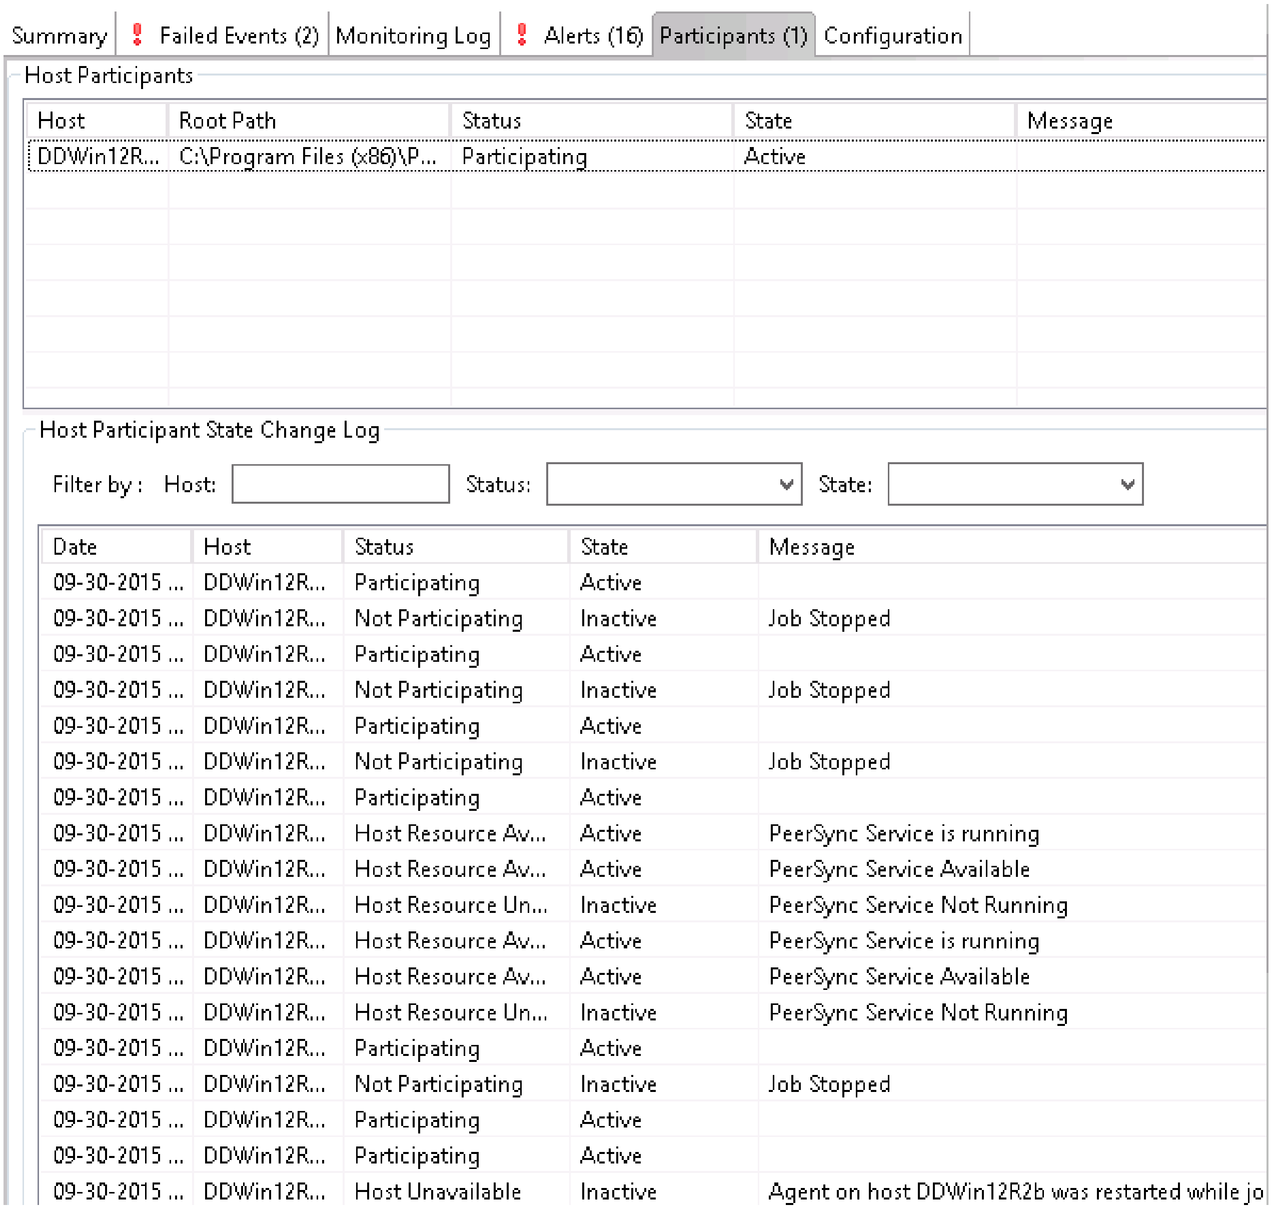

Host	Root Path	Status	State	Message
DDWin12R...	C:\Program Files (x86)\P...	Participating	Active	

Host Participant State Change Log
Filter by : Host: Status: State:

Date	Host	Status	State	Message
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Not Participating	Inactive	Job Stopped
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Not Participating	Inactive	Job Stopped
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Not Participating	Inactive	Job Stopped
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Host Resource Av...	Active	PeerSync Service is running
09-30-2015 ...	DDWin12R...	Host Resource Av...	Active	PeerSync Service Available
09-30-2015 ...	DDWin12R...	Host Resource Un...	Inactive	PeerSync Service Not Running
09-30-2015 ...	DDWin12R...	Host Resource Av...	Active	PeerSync Service is running
09-30-2015 ...	DDWin12R...	Host Resource Av...	Active	PeerSync Service Available
09-30-2015 ...	DDWin12R...	Host Resource Un...	Inactive	PeerSync Service Not Running
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Not Participating	Inactive	Job Stopped
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Host Unavailable	Inactive	Agent on host DDWin12R2b was restarted while jo

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.

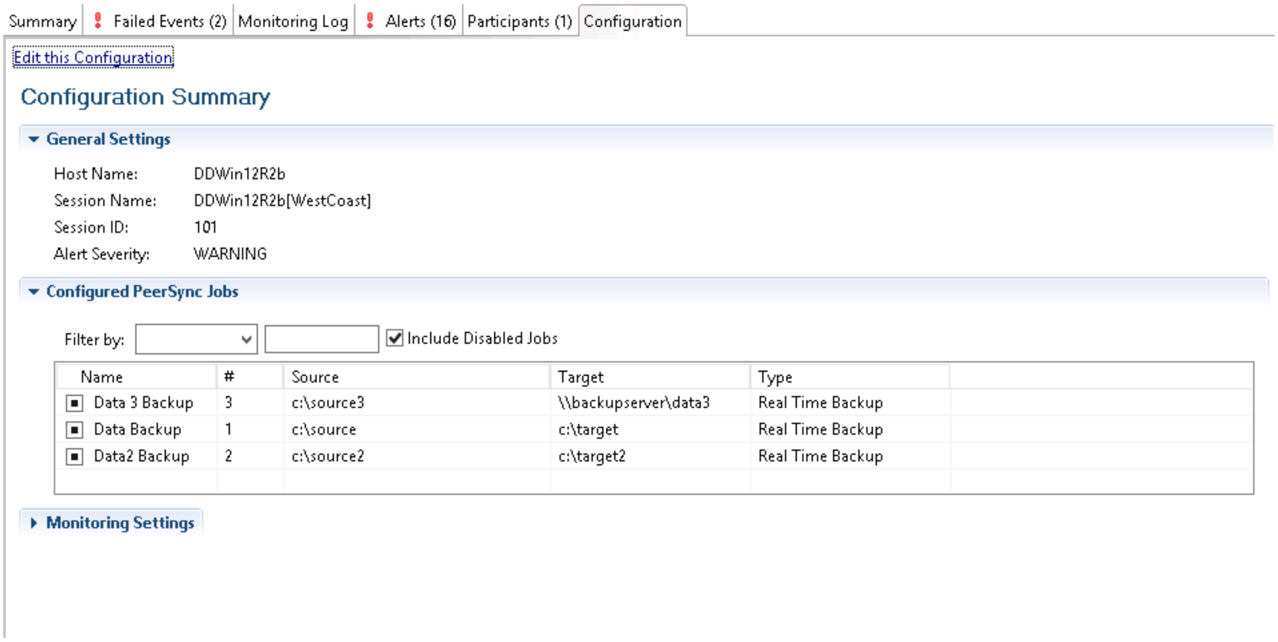

Configuration Summary

General Settings

Host Name: DDWin12R2b
Session Name: DDWin12R2b[WestCoast]
Session ID: 101
Alert Severity: WARNING

Configured PeerSync Jobs

Filter by: ☑ Include Disabled Jobs

Name	#	Source	Target	Type
■ Data 3 Backup	3	c:\source3	\\backupserver\data3	Real Time Backup
■ Data Backup	1	c:\source	c:\target	Real Time Backup
■ Data2 Backup	2	c:\source2	c:\target2	Real Time Backup

Monitoring Settings

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.

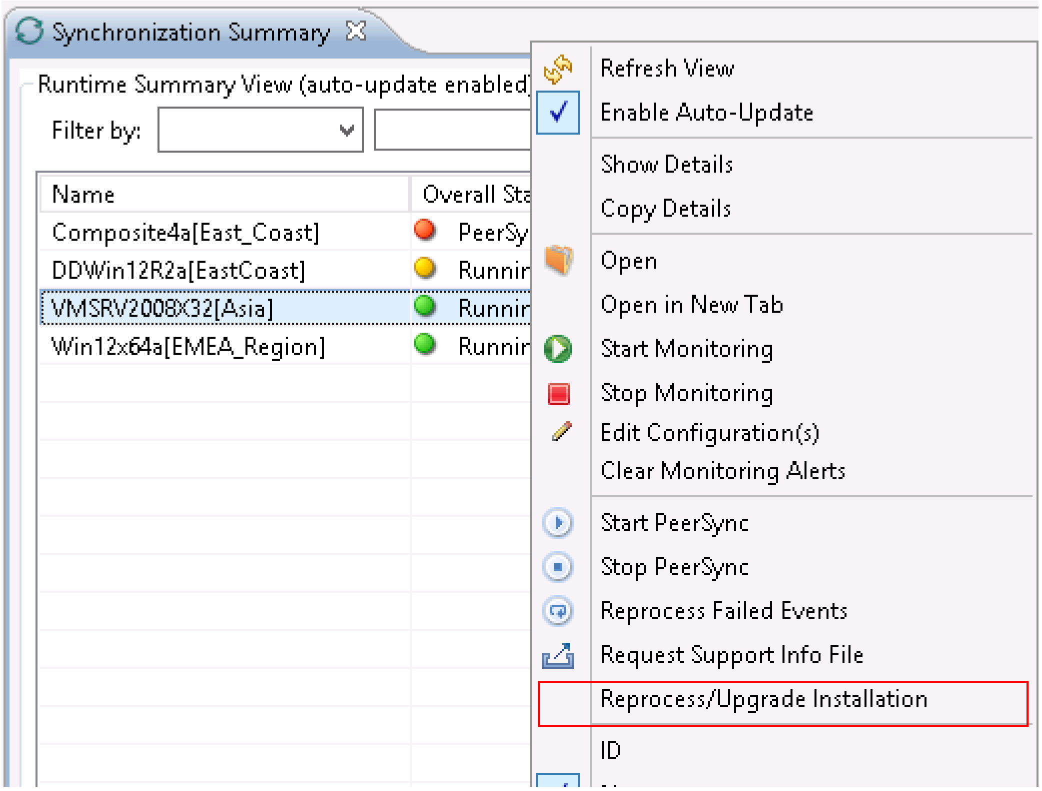

Name	Overall Sta
Composite4a[East_Coast]	🔴 PeerSy
DDWin12R2a[EastCoast]	🟡 Runnir
VMSRV2008X32[Asia]	🟢 Runnir
Win12x64a[EMEA_Region]	🟢 Runnir

The installation settings should be common for ALL the PeerSync Management PeerSync instances in order to successfully install PeerSync.

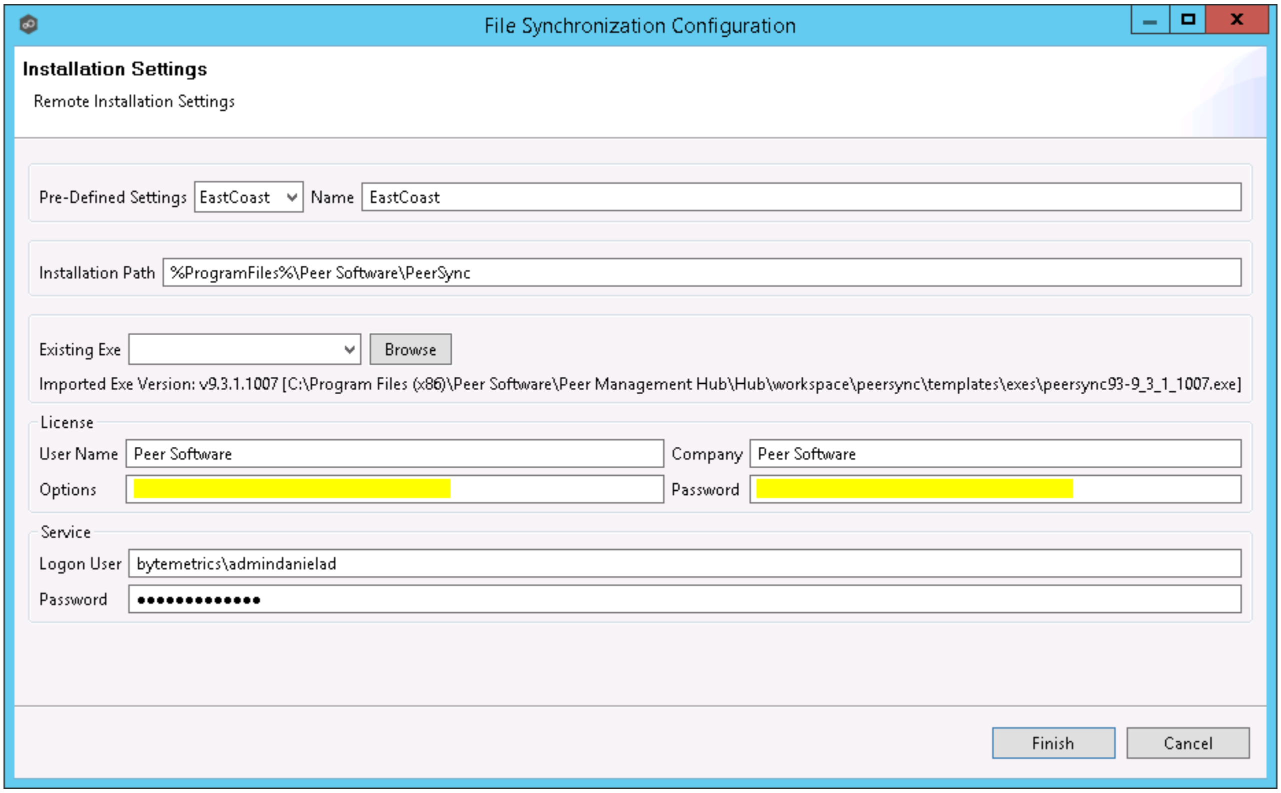

File Synchronization Configuration

Installation Settings

Remote Installation Settings

Pre-Defined Settings: EastCoast
Name: EastCoast

Installation Path: %ProgramFiles%\Peer Software\PeerSync

Existing Exe:
Browse

Imported Exe Version: v9.3.1.1007 [C:\Program Files (x86)\Peer Software\Peer Management Hub\Hub\workspace\peersync\templates\exes\peersync93-9_3_1_1007.exe]

License
User Name: Peer Software
Company: Peer Software
Options: [redacted]
Password: [redacted]

Service
Logon User: bytemetrics\admindanielad
Password: ●●●●●●●●●●●●●●

See Installation Settings for information on the settings on this page.

Index