



**Peer Global File Service
User Guide**

Copyright (c) 1993-2025 Peer Software, Inc. All Rights Reserved.
Updated Thursday, March 20, 2025

Table of Contents

Peer Global File Service Help	1
Terminology	1
Requirements and Prerequisites	8
Installing Peer Management Center	9
Installing Peer Management Center and Agents	9
Windows Installation Instructions	10
Securing Access to the Web Client	16
Installing Peer Agents	16
Peer Services Required for Web Client	20
Updating Peer Management Center and Peer Agents	21
Updating Peer Management Center	21
Updating Peer Agents	23
Uninstalling Peer Management Center and Agents	24
Running Peer Management Center	25
Launching the Rich Client	25
Accessing the Web Client	25
Peer Management Center User Interface	27
Main Window	29
Perspectives and Views	29
Views	31
Jobs Perspective	32
Agents View	35
Agents Toolbar	36
Alerts View	36
Brokers View	38
Dashboard	39
Job Alerts View	40
Jobs View	41
Jobs Toolbar	42
Runtime Views	43
Cloud Backup and Replication Job Runtime View	44
DFS-N Management Job Runtime View	45
File Collaboration Job Runtime View	46
Summary Tab	47
Session Tab	49
Event Log Tab	51
Quarantines Tab	52
Retries Tab	53
Alerts Tab	54
Participants Tab	55
Configuration Tab	57
File Replication Job Runtime View	58
File Synchronization Job Runtime View	59
Summary Views	60
Agent Summary View	61
Cloud Summary View	61
Collab, Sync, and Repl Summary View	63

Summary Tab	63
Edge Caching Tab.....	66
Reports Tab	66
Namespace Summary View.....	70
Topology Perspective.....	71
Menus and Toolbars	72
File Menu	73
Window Menu.....	73
Tools Menu	75
Help Menu	76
Toolbar	79
Tables	81
Basic Concepts	83
Email Alerts	83
File and Folder Filters	84
Creating and Applying File Filters.....	85
Predefined File Filters.....	85
Updating Predefined File Filters.....	86
Defining Filter Patterns.....	88
Using Wildcards in Filter Patterns.....	89
Automatically Excluded File Types.....	90
Excluding Temporary Files.....	90
Using Complex Regular Expressions in Filter Patterns.....	93
Filtering on Last Modified Date.....	93
Filtering on File Size.....	96
Filtering Folders.....	98
Folder Filter Examples.....	99
File Filter Usage Notes.....	100
List Filters	101
Creating Complex Filter Expressions.....	102
Saving and Managing List Filters.....	103
Removing List Filters.....	103
Logging and Alerts	103
Retrieving Log Files.....	105
SNMP Notifications	110
Tags	111
Creating Tags and Categories.....	111
Assigning Tags.....	111
Using Tags to Filter Resources.....	113
Web Client Users	114
Roles for Web Client Users.....	115
Base Roles	115
Base Role Permissions.....	116
Custom Roles	119
Custom Role Permissions.....	119
Advanced Topics	120
Analytics	121
File System Analytics.....	123
File System Analytics Scan Process.....	124
File Activity Analytics.....	124
PeerIQ	125
Proactive Monitoring.....	127
Conflicts, Retries, and Quarantines	128

DFS Namespaces	129
Using DFS Namespaces with Jobs	131
DFS Namespace Failover and Failback	131
Edge Caching	132
File Metadata Synchronization	133
Multi-Protocol	134
Network of Brokers	137
Managing Peer Agents	137
Peer Agent Connection Statuses	140
Re-enabling a Disabled Agent Within a Job	140
Editing an Agent Configuration	141
Broker	144
General	146
Logging	149
Performance	149
VM Options	151
Viewing Agent Properties	152
Editing Agent Properties	154
Updating a Peer Agent	156
PeerGFS API	156
Accessing the PeerGFS API	157
Testing the PeerGFS API	157
Integrating Your Own Tools and Scripts with the PeerGFS API	158
API Quick Reference	159
API Categories	159
Status Codes	159
Scheduled Replication	159
Smart Data Seeding	160
Storage Capacity	162
TLS Encryption and Certificate Management	163
Preferences	163
Configuring Preferences	165
Analytics	166
Proactive Monitoring Preferences	169
Step 1: General Configuration	170
Step 2: Telemetry Options	172
Step 3: Agent Locations	174
Step 4: Health Checker Setup	175
Set Up the Health Checker	178
Select Jobs	182
Schedule Task	183
Update Health Checker	184
Step 5: Confirmation	186
Disabling and Re-enabling Proactive Monitoring	188
Cloud Backup and Replication Job Preferences	191
Cloud Backup and Replication	192
Database Connections	194
Destination Credentials	197
Email Alerts	200
File Retries and Source Snapshots	203
File and Folder Filters	205
Performance	208
Proxy Configuration	210
Replication and Retention Policies	213

SNMP Notifications	215
Scan Manager.....	217
Collaboration, Replication, and Synchronization Job Preferences	219
Collab, Sync, and Replication	219
Application Enhancements.....	221
DFS-N Management.....	224
Edge Caching.....	226
Edge Caching	226
Email Alerts	227
Master Data Service.....	231
Pinning Filters	233
SNMP Notifications.....	237
Utilization Policies.....	240
Volume Policies	244
Email Alerts.....	254
File Retries.....	260
File and Folder Filters.....	261
Locking	265
Performance.....	266
Real-time Event Detection.....	267
SNMP Notifications.....	269
Scan Manager.....	273
Scheduled Replication Filters.....	275
DFS-N Management Job Preferences	280
Email Alerts.....	282
SNMP Notifications.....	286
Email Configuration	289
General Configuration	292
General Configuration.....	293
Agent Connectivity.....	295
Broker Configuration.....	297
Email Alerts.....	298
Proxy Configuration.....	302
Software Updates.....	304
Tags Configuration.....	306
Web and API Configuration.....	308
Licensing	309
MED Configuration	311
Storage Configuration	317
SMB Access Configuration.....	317
Storage Platforms.....	319
Amazon FSxN Configurations.....	319
Amazon FSxN Advanced Options.....	324
Dell Configurations.....	325
Dell PowerScale Configurations.....	329
Dell PowerScale Advanced Options.....	332
Dell Unity Configurations.....	337
Dell Unity Advanced Options.....	339
NetApp ONTAP Configurations.....	342
NetApp ONTAP Advanced Options.....	347
Nutanix Files Configurations.....	348
Nutanix Files Advanced Options.....	352
Real-time Event Detection	354
SNMP Configuration	356

User Management	358
Accessing User Management	358
Managing Web Client Users	360
Managing Internal Users	360
Creating an Internal User	361
Editing an Internal User	364
Deleting an Internal User	364
Managing Active Directory Users and Groups	365
Configuring Active Directory Authentication	365
Adding an Active Directory User or Group	368
Editing an Active Directory User or Group	371
Deleting an Active Directory User and Group	371
Managing Roles for Web Client Users	372
Creating a Custom Role	373
Editing a Custom Role	377
Deleting a Custom Role	378
Assigning Tags to a Role	378
Cloud Backup and Replication Jobs	378
Overview	379
Before You Create Your First Cloud Backup and Replication Job	379
Creating a Cloud Backup and Replication Job	379
Step 1: Job Type and Name	380
Step 2: Source Storage Platform	382
Step 3: Management Agent	383
Step 4: Proxy Configuration	384
Step 5: Storage Information	390
Amazon FSxN	391
Dell PowerScale	393
Dell Unity	397
NetApp ONTAP	399
Nutanix Files	401
Step 6: Source Paths	403
Step 7: File and Folder Filters	405
Step 8: Destination	406
Step 9: Destination Credentials	408
Azure Blob Storage Credentials	408
Amazon S3 Credentials	410
NetApp StorageGRID Credentials	412
Nutanix Objects Credentials	414
Wasabi Credentials	415
Step 10: Container or Bucket Details	416
Azure Blob Storage Container Details	417
Amazon S3 Bucket Details	419
NetApp StorageGRID Bucket Details	421
Nutanix Objects Bucket Details	423
Wasabi Bucket Details	425
Step 11: Replication and Retention Policy	427
Step 12: Replication Schedule	428
Scheduled Scans	429
Batched Real-Time	431
Continuous Data Protection	432
Step 13: Retention	433
Step 14: Source Snapshots	434
Step 15: Miscellaneous Options	435

Step 16: Email Alerts	438
Step 17: SNMP Notifications	440
Step 18: Confirmation	442
Running a Cloud Backup and Replication Job	444
Starting a Cloud Backup and Replication Job	444
Stopping a Cloud Backup and Replication Job	446
Monitoring Cloud Backup and Replication Jobs	447
Deleting a Cloud Backup and Replication Job	448
Recovering Data	449
Search Options	452
Search by Name	452
Search by Snapshot	455
Search by Point in Time	456
Search by Latest Replication	458
Recovery Options	458
DFS-N Management Jobs	462
Creating a DFS-N Management Job	463
Step 1: Job Type	464
Step 2: Management Agent	465
Step 3: Agent Verification	466
Step 4: Namespace Name	468
Step 5: Namespace Servers	469
Step 6: Namespace Settings	471
Step 7: Namespace Folders	473
Step 8: Email Alerts	478
Step 9: SNMP Notifications	481
Step 10: Review	483
Step 11: Results	484
Importing an Existing Namespace	486
Running a DFS-N Management Job	496
Starting a DFS-N Management Job	496
Stopping a DFS-N Management Job	498
Managing DFS Namespaces	499
Adding a Namespace Server	499
Adding a Namespace Folder	504
Adding a Folder Target	511
Linking a DFS Namespace to File Collaboration and File Synchronization Jobs	516
Creating a File Collaboration or File Synchronization Job from a Namespace Folder	516
Linking an Existing Namespace Folder to an Existing File Collaboration or File Synchronization Job	523
File Collaboration Jobs	527
Overview	528
Before You Create Your First File Collaboration Job	528
Creating a File Collaboration Job	529
Step 1: Job Type and Name	529
Step 2: Participants	531
Management Agent	533
Storage Platform	534
Storage Information	535
Amazon FSxN	536
Dell PowerScale	537
Dell Unity	541
NetApp ONTAP	543
Nutanix Files	545

Windows File Server.....	547
Windows File Server Advanced Options.....	548
Path	551
Edge Caching	552
Master Data Service.....	555
Volume Policy	558
Utilization Policy.....	561
Pinning Filter	563
Step 3: Master-Edge Assignment.....	567
Step 4: File Metadata.....	570
Step 5: Application Support.....	572
Step 6: Email Alerts.....	573
Step 7: DFS Namespace.....	576
Step 8: Save Job.....	580
Editing a File Collaboration Job	581
Participants.....	583
Adding and Deleting a Participant.....	584
Editing a Participant.....	590
Master-Edge Assignment.....	592
General	594
File and Folder Filters.....	596
Scheduled Replication Filters.....	597
Conflict Resolution.....	598
Delta Replication.....	600
File Metadata.....	602
File Locking.....	604
Application Support.....	609
Target Protection.....	609
Email Alerts.....	611
SNMP Notifications.....	613
Tags	614
DFS-N Link.....	615
Editing Multiple Jobs.....	616
Running and Managing a File Collaboration Job	618
Overview	618
Job Initialization Process.....	619
Initial Synchronization Process.....	619
Starting a File Collaboration Job.....	620
Stopping a File Collaboration Job.....	622
Auto-Restarting a File Collaboration Job.....	622
Host Connectivity Issues.....	624
Removing a File from Quarantine.....	625
Manual Retries.....	626
File Replication Jobs	626
Overview	627
Before You Create Your First File Replication Job	627
Creating a File Replication Job	627
Step 1: Job Type and Name.....	628
Step 2: Source Agent.....	629
Step 3: Storage Platform.....	631
Step 4: Storage Information.....	632
Amazon FSxN	633
Dell PowerScale	635
Dell Unity	638

Linux File Server	640
Linux File Server Advanced Options	641
NetApp ONTAP	642
Nutanix Files	644
Windows File Server.....	646
Windows File Server Advanced Options.....	647
Step 5: Source Path.....	650
Step 6: Destination Agent.....	654
Step 7: Destination Path.....	655
SMB Path	660
Step 8: File Metadata.....	660
SMB File Metadata.....	661
NFS File Metadata.....	662
Multi-Protocol Metadata.....	664
Step 9: Email Alerts.....	666
Step 10: Save Job.....	669
Editing a File Replication Job	670
Participants.....	672
Adding and Deleting a Participant.....	672
Editing a Participant.....	677
General	679
File and Folder Filters.....	681
Scheduled Replication Filters.....	682
Conflict Resolution.....	684
Delta Replication.....	686
File Metadata.....	687
SMB File Metadata.....	688
NSF File Metadata.....	690
Multi-Protocol Metadata.....	692
Job Relay	694
File Locking.....	696
Application Support.....	701
Target Protection.....	702
Email Alerts.....	703
SNMP Notifications.....	706
Tags	707
Editing Multiple Jobs.....	708
File Synchronization Jobs	710
Overview	710
Before You Create Your First File Synchronization Job	711
Creating a File Synchronization Job	711
Step 1: Job Type and Name.....	711
Step 2: Participants.....	713
Management Agent.....	714
Storage Platform	716
Storage Information.....	717
Amazon FSxN	718
Dell PowerScale.....	720
Dell Unity	724
Linux File Server.....	726
Linux File Server Advanced Options.....	728
NetApp ONTAP	729
Nutanix Files	731
Windows File Server.....	733

Windows File Server Advanced Options.....	734
Path	736
Edge Caching	742
Master Data Service.....	745
Volume Policy	748
Utilization Policy.....	751
Pinning Filter	753
Step 3: Master-Edge Assignment.....	757
Step 4: File Metadata.....	761
SMB File Metadata.....	761
NFS File Metadata.....	763
Mixed Permissions.....	765
Step 5: Email Alerts.....	767
Step 6: DFS Namespace.....	769
Step 7: Save Job.....	772
Editing a File Synchronization Job	772
Participants.....	775
Adding and Deleting a Participant.....	775
Editing a Participant.....	782
Master-Edge Assignment.....	784
General	786
File and Folder Filters.....	789
Scheduled Replication Filters.....	791
Conflict Resolution.....	792
Delta Replication.....	794
File Metadata.....	796
SMB File Metadata.....	797
NFS File Metadata.....	799
Multi-Protocol File Metadata.....	801
File Locking.....	803
Application Support.....	805
Target Protection.....	806
Email Alerts.....	808
SNMP Notifications.....	810
Tags	811
DFS-N Link.....	811
Editing Multiple Jobs.....	812
Running and Managing a File Synchronization Job	814
Overview	814
Job Initialization Process.....	815
Initial Synchronization Process.....	815
Starting a File Synchronization Job.....	816
Stopping a File Synchronization Job.....	818
Auto-Restarting a File Synchronization Job.....	818
Host Connectivity Issues.....	820
Removing a File from Quarantine.....	821
Manual Retries.....	821

Peer Global File Service Help

Using This Help File

This help file is designed to be used online. It is cross-linked so that you can find more relevant information on any subject from any location. If you prefer reading printed manuals, a PDF version of this help file is available from our website. The PDF version 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

Copyright (c) 1993-2025 Peer Software, Inc. All Rights Reserved. Although we try to provide quality information, Peer Software makes no claims, promises, or guarantees about the accuracy, completeness, or adequacy of the information contained in this document. Peer Software, Peer Management Center, Peer Global File Service, PeerGFS, PeerSync, and their respective logos are trademarks or registered trademarks of Peer Software, Inc. Microsoft, Azure, Windows, Windows Server, and their respective logos are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. NetApp, the NetApp logo, Data ONTAP, and FPolicy are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. "Amazon Web Services", "AWS", "Amazon S3", "Amazon Simple Storage Service", "Amazon SNS", "Amazon Simple Notification Service", and their respective graphics, logos, and service names are trademarks, registered trademarks or trade dress of Amazon Web Services LLC and/or its affiliates in the U.S. and/or other countries. Dell, EMC, Celerra, Isilon, VNX, Unity and other trademarks are trademarks of Dell Inc. or its subsidiaries. Nutanix is a trademark of Nutanix, Inc., registered in the United States and other countries. All other trademarks are the property of their respective companies. Peer Software vigorously protects and defends its trade name, trademarks, patents, designs, copyrights, and other intellectual property rights. Unless otherwise specified, no person has permission to copy, redistribute, reproduce, or republish in any form the information in this document.

Last updated: Thursday, March 20, 2025

PeerGFS Version 6.2

Terminology

Introduction

Before getting started, familiarize yourself with key concepts and terminology.

Terms

Term	Definition
Active-Active	Two or more file servers that host 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 <i>Peer Agent</i> .
Cloud Backup and Replication job	A type of job where a single participating host has a designated set of folders and files to be replicated to a cloud storage device.
DFS (Distributed File System)	A set of client and server services that allow an organization using Microsoft Windows servers to organize many distributed SMB file shares into a distributed file system.
DFS namespace (DFS-N)	A namespace that enables you to group shared folders located on different servers into one or more logically structured namespaces.
DFS Namespaces	A Windows Server feature that allows multiple SMB shares across different file servers (and even locations) to be combined into a single unified namespace. DFS Namespaces simplifies access to files, especially in large, distributed environments. When combined with Peer file synchronization technology, DFS Namespaces can provide redundancy to file shares across file servers and locations.

Term	Definition
DFS-N Management job	A type of job that enables the creation and management of DFS namespaces.
Event	A single operation performed by a user on a file server.
Failback	The process of redirecting previously displaced users from a secondary file server back to the primary after a failure state has been resolved.
Failover	The process of redirecting users from one file server to a secondary in the event of a failure.
File access event	An event that is triggered from the opening or closing of a file.
File change event	An event that causes a file to be changed in some way, for example, file modify, file delete, file rename, file attribute change.
File Collaboration job	A type of job that combines file synchronization with distributed file locking to prevent version conflicts across multiple active file servers.
File Collaboration	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.

Term	Definition
session	
File filter	A type of filter used to include or exclude specific files from replication and locking.
File lock conflict	A file collaboration condition that exists when two users open a file at the same time, and both hold exclusive locks on the file.
File Replication job	A type of job that involves real-time and/or scheduled copying of files and folders from one file server to another.
File Synchronization job	A type of job that involves multi-directional real-time replication so that two or more file servers are always up to date with each other.
Filter	Three types of filters: file, folder, and list.
Filter expression	See <i>list filter</i> .
Folder filter	A type of filter used to include or exclude specific folders (and the files they contain) from replication and locking.
Heart beat	A communication mechanism used between Peer Management Center and all connected Peer Agents to ensure that Peer Agents are alive and responsive. Heartbeats share information about the Peer Agent host server with Peer Management Center, aid in verifying when a Peer Agent is no longer available, and signal when a disconnected Peer Agent has reconnected. All heartbeat information is sent through the Peer Management Broker.

Term	Definition
Host	A server that a Peer Agent is installed upon.
Initialization process	The steps executed whenever a job is started in Peer Management Center. The steps for an initialization process are different for each job type.
Initial synchronization process	The background process that occurs at the start of a file collaboration session, where the directory watch set is recursively scanned on all participating hosts, file conflict resolution is performed, and any files that require updating are synchronized with the most current copy of the file.
List filter	A type of filter used to show or hide information from various views in Peer Management Center.
Management Agent	A server running the Peer Agent. Can manage storage devices or a DFS namespace.
Master host	In file synchronization and collaboration, the master host will always win in a split-brain scenario.
Malicious Event Detector (MED)	Leverages the same real-time event detection that powers all job types to detect and alert administrators to malicious user and application behavior. For more information, see Introduction to Peer MED in our knowledge base.

Term	Definition
Participant	A participant consists of an Agent and the volume/share/export folder to be replicated. Applies to File Collaboration, File Replication, and File Synchronization jobs.
Peer Agent (or Agent)	A lightweight piece of software that is installed on Windows Servers to perform the storage and file management functions used by the entire Peer Global File Service solution suite. Typically installed on or alongside the file servers that will be managed by Peer Management Center.
Peer Management Center (PMC)	The focal component of Peer Global File Service. Responsible for configuration, management, and monitoring of Peer Agents and the various solutions configured in Peer Global File Service. Peer Management Center runs as three parts: a Windows Service that is always running, along with a rich client application and a web server component, both used for configuration and monitoring.
Peer Management Broker	The central messaging system of Peer Global File Service. The Peer Management Broker serves to connect Peer Management Center and Peer Agents, forming a Peer Management Center "network" that can be cast over local or wide-area networks via TCP/IP. One or more Peer Management Brokers are deployed in a Peer Management Center environment.
Quarantined file	A file that has been removed from a File Collaboration or File Synchronization job as a result of a lock or replication conflict that cannot 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	The requirement for a minimum of two participants must be available and connected. If that number dips to one or less, the quorum will not be met. Applies to File Collaboration, File Replication, and File Synchronization jobs.
Real-time event detection	A key technology that backs all job types in Peer Management Center. Peer Management Center receives notifications as end users interact with the file servers that are being monitored. These notifications will usually result in replication or locking between file servers.

Term	Definition
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 preseeded 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, see Smart Data Seeding or contact Peer Support .
SMB / CIFS	Server Message Block or Common Internet File System, an application-layer protocol used for providing shared access to file data and other networked resources.
Source host	The file server hosting a file from which file access or change event originated.
Target host	One or more Management Agents of file servers where file access and change events will be propagated to.
TLS	Transport Layer Security, a successor to Secure Socket Layer (SSL) that secures network traffic between a client and server.
UNC Path	A UNC path can be used to access network resources and must be in the format specified by the Universal Naming Convention. A UNC path always starts with two backslash characters (\\).
View	Individual sections of Peer Management Center's user interface, each providing unique information and control.

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

Requirements and Prerequisites

Environmental Requirements

Before you begin using using Peer Global File Service, ensure that the [environmental requirements](#) have been satisfied.

Platform-Specific Prerequisites

In addition to the environmental requirements, specific platform prerequisites must be met for using Peer Global File Service:

- [Amazon FSx for NetApp ONTAP Prerequisites](#)
- [Dell PowerScale Prerequisites](#)
- [Dell Unity Prerequisites](#)
- [NetApp ONTAP Prerequisites](#)

- [Nutanix Files prerequisites](#)

Dell Platform Configuration

If you plan to run a Peer Agent with a Dell PowerScale or Unity storage device, refer to the following guides for detailed configuration instructions:

- [Dell PowerScale Syslog Configuration Guide](#) (for PowerScale OneFS v9.5 and above)
- [Dell PowerScale Configuration Guide](#) (for older PowerScale OneFS versions prior to v9.5; using Legacy CEE)
- [Dell Unity Configuration Guide](#)

Installing Peer Management Center

Topics in this section provide information about:

[Installing Peer Management Center and Peer Agents](#)

[Updating Peer Management Center and Peer Agents](#)

[Uninstalling Peer Management Center and Peer Agents](#)

For information about Peer Global File System licensing, see [Licensing](#).

Installing Peer Management Center and Agents

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

- **Peer Management Center installer:** This installer installs both Peer Management Center and [Peer Management Broker](#) on the same server. Peer Management Broker handles the communication between Peer Management Center and Peer Agents. For instructions on installing Peer Management Center (PMC), see [Windows Installation Instructions](#)
- **Peer Agent installer:** This installer contains the Peer Agent installation files. You must install an Agent on each server you plan to include in any of your jobs. See [Installing Peer Agents](#) for installation instructions.

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

Windows Installation Instructions

Installing Peer Management Center and Peer Management Broker

To install Peer Management Center and Peer Management Broker:

1. Download the Peer Management Center installer (**PMC_Installer_win64.exe**) to the server you want to host Peer Management Center.
2. Run the installer and follow the installation wizard instructions.
3. During the installation, you will be prompted to configure access to the **Peer Management Center Web Service** and the **Peer Management API Service**. The web service allows users to access Peer Management Center via a web browser; the API service allows access to Peer Management Center through REST API calls. If you do not enable web access during the initial installation or want to modify settings at a later date, you can modify them through [Web and API Configuration in Preferences](#).

Setup - Peer Management Center 5.0.0.20220601

Peer Management Center Web and REST API Configuration

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

Hostname or IP:

Enable HTTPS Web Access using port:

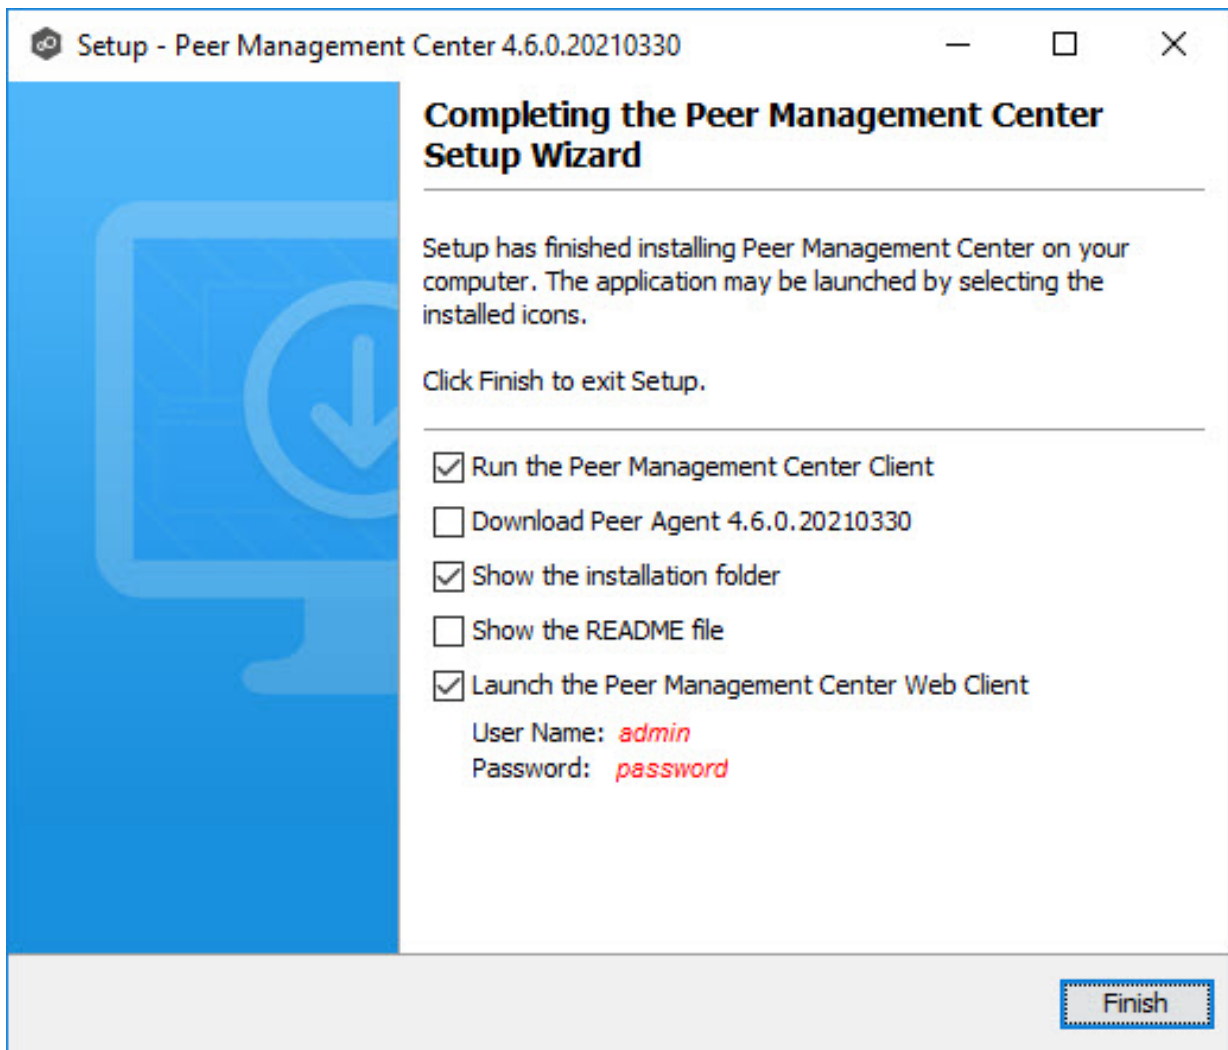
Enable HTTPS REST API using port:

< Back Next > Cancel

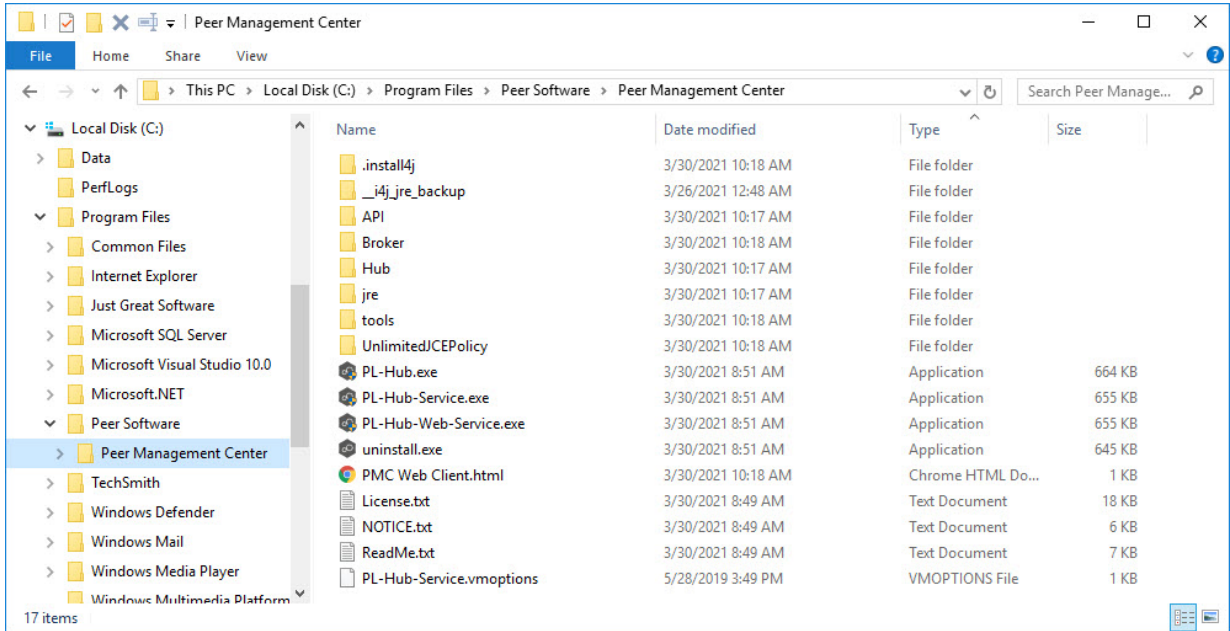
Field	Description
Hostname or IP	<p>Enter the hostname or IP address via which the services can be accessed:</p> <ul style="list-style-type: none"> • Enter localhost or 127.0.0.1 if you want the services to be accessible only to users of the local server via the loopback interface. • Enter 0.0.0.0 to make the services accessible via all network interfaces. • Enter a specific IP address to restrict access to a specific network interface.
Enable HTTPS Web Access	Select this to enable secure access to the web service, and then enter a port number.
Enable HTTPS REST API	Select this to enable secure access to the REST API service, and then enter a port number. The REST API port cannot be the same

Field	Description
	as the web service port.

- If you enabled access to the web client, see the [Securing Access to the Web Client](#) for additional information about securing access to the web client.
- On the final page of the installation wizard, you have several options; we recommend that, at minimum, you select the first option.
- If you enabled the Peer Management Center Web Service in Step 3 and selected the **Launch the Peer Management Center Web Client**, on the final page of the installation wizard, the default username and password for accessing the web client is displayed. After [logging in to the web client](#), you should [change the password](#) immediately.

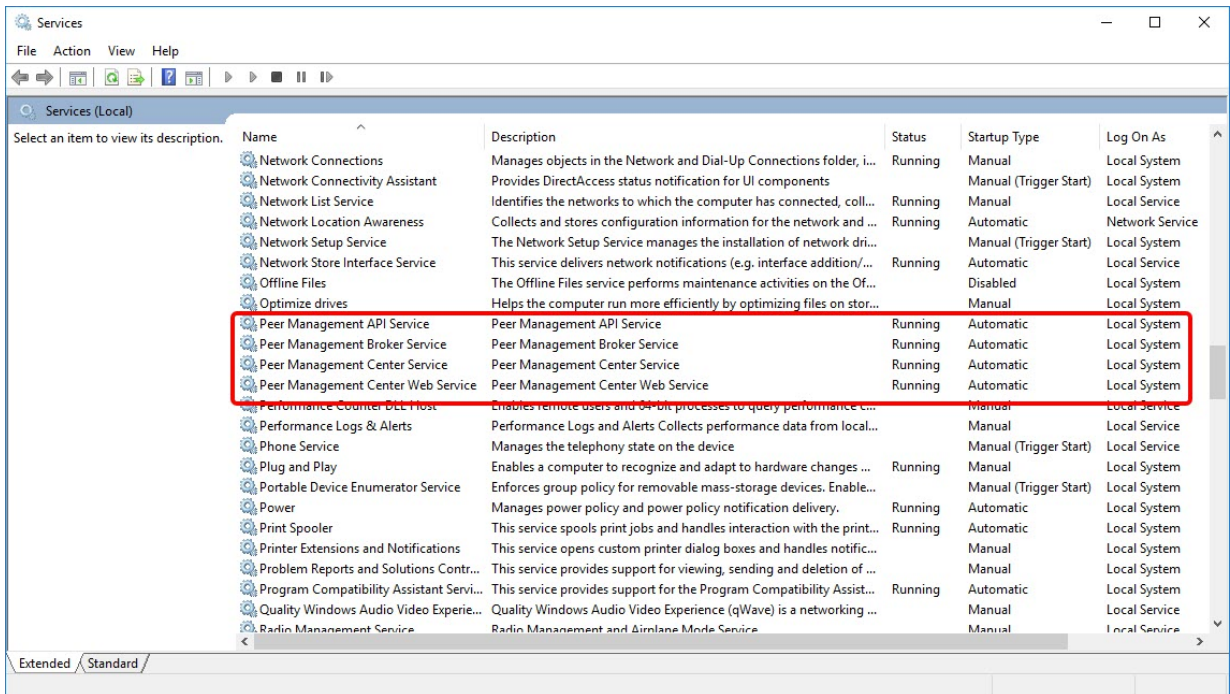


When the installation is complete, the Peer Management Center installation folder contains the following files and folders:



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

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

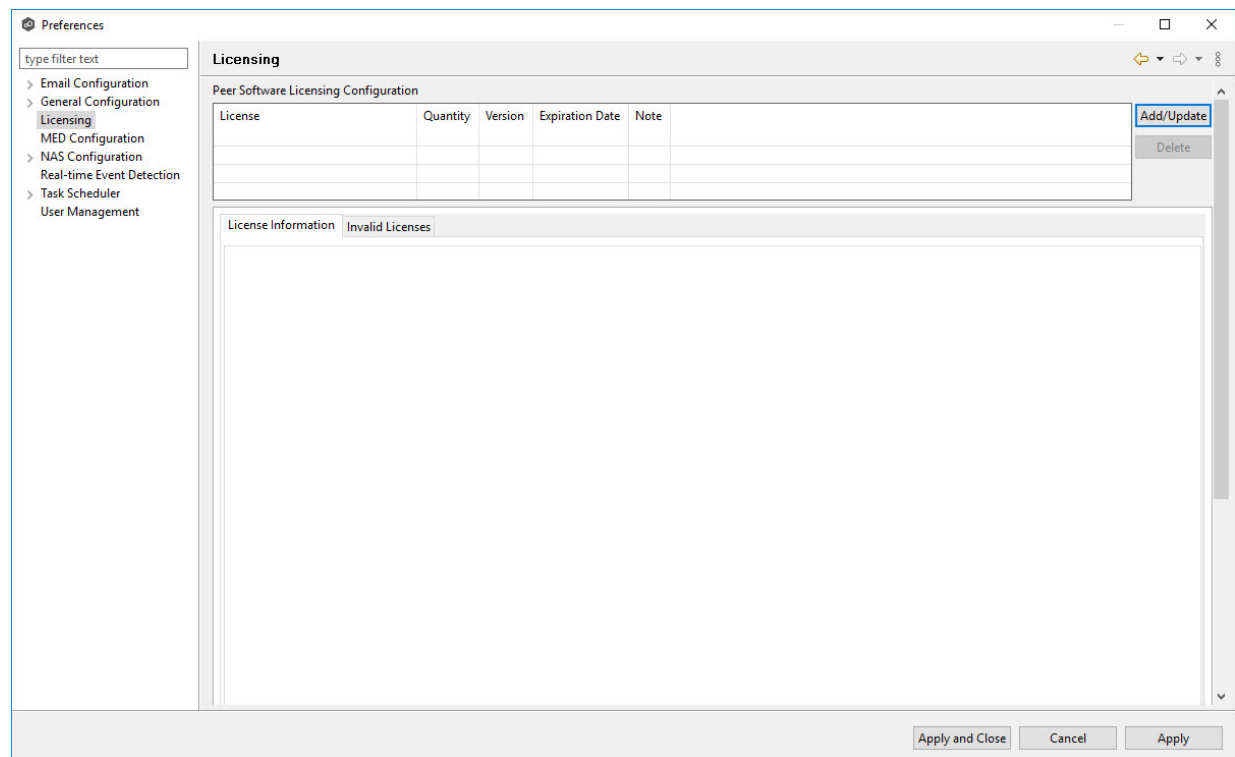


7. If you didn't select the option to launch Peer Management Center Client on the last page of the installation wizard, launch it using one of the following methods:

- Select **Peer Management Center** from the Windows **Start** menu.
- Double-click the **PL-Hub.exe** executable in the Peer Management Center installation directory.

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

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

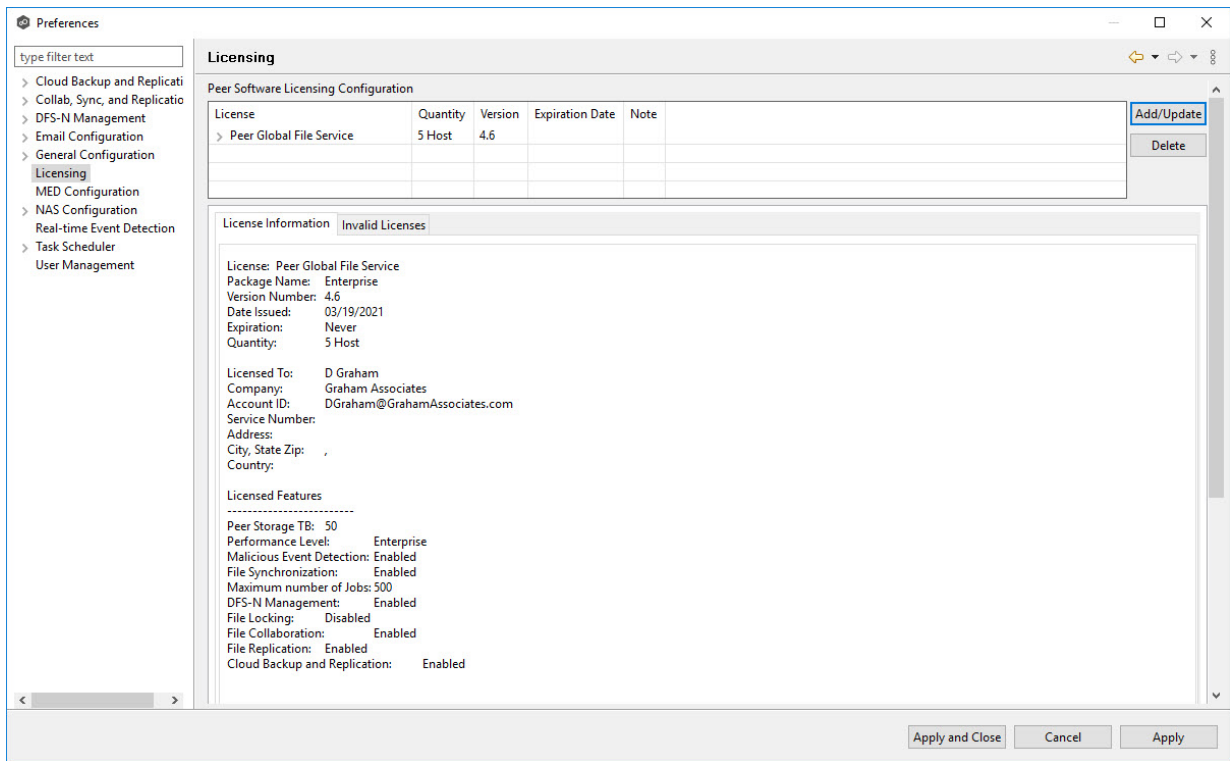


9. Click **Add/Update**.


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

11. Click **Open**.

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



12. Click **Apply and Close**.



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

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

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

Secure Encrypted TLS Connections

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

Installing Peer Agents

You'll need to install a Peer Agent on each server you intend to incorporate into your jobs. Following the installation of the Peer Agent software, you should verify that the **Peer Agent Service** is running and can successfully connect to a [Peer Management Broker](#).

Note: For customers using clustered file server roles with Windows Failover Cluster, please review the Peer Software knowledge base article [Using a Peer Agent in a Windows Failover Cluster](#).

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

1. Download the Peer Agent installer (**P-Agent_Installer_win64.exe**) to the server you want to host the Agent.
2. Run the installer and follow the wizard instructions.

During installation, you will be prompted for:

- The Peer Management Broker hostname (computer name, fully qualified domain name, or IP address) of the server where Peer Management Broker is running.
- The security protocol (either SSL or TLS).
- The TCP/IP port number of the server where Peer Management Broker is running. The default port for TLS communication is 61617.

Enter the same values that you entered when [installing Peer Management Center and Peer Management Broker](#).

Setup - Peer Agent 4.6.0.20210316

Peer Management Broker Configuration

Please provide the hostname, protocol and TCP/IP port of the Broker.

Hostname: ?

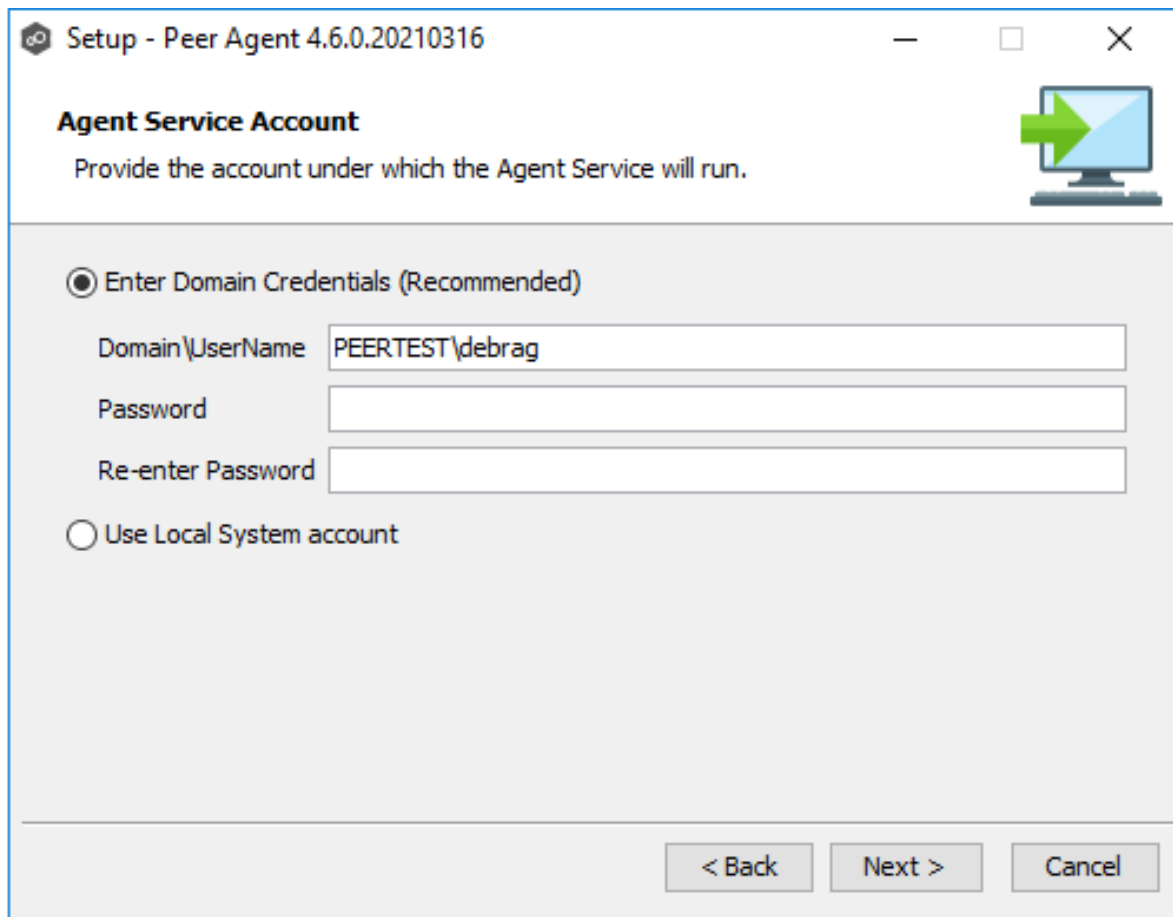
Protocol: ?

Port: ?

Please note that enabling SSL/TLS may reduce performance.


< Back Next > Cancel

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

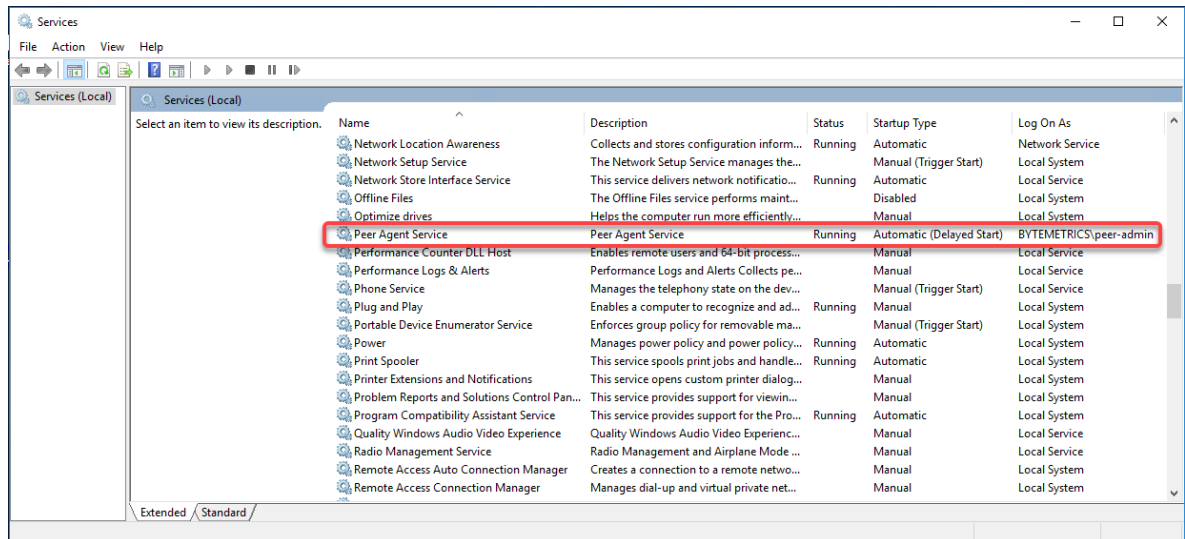


The screenshot shows a Windows installation wizard window titled "Setup - Peer Agent 4.6.0.20210316". The window has standard Windows window controls (minimize, maximize, close) in the top right corner. The main heading is "Agent Service Account" with a sub-instruction: "Provide the account under which the Agent Service will run." To the right of the text is a green arrow pointing right, overlaid on a computer monitor icon. Below this, there are two radio button options. The first option, "Enter Domain Credentials (Recommended)", is selected. It includes three text input fields: "Domain\UserName" containing "PEERTEST\debrag", "Password" (empty), and "Re-enter Password" (empty). The second option, "Use Local System account", is unselected. At the bottom of the window, there are three buttons: "< Back", "Next >", and "Cancel".

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

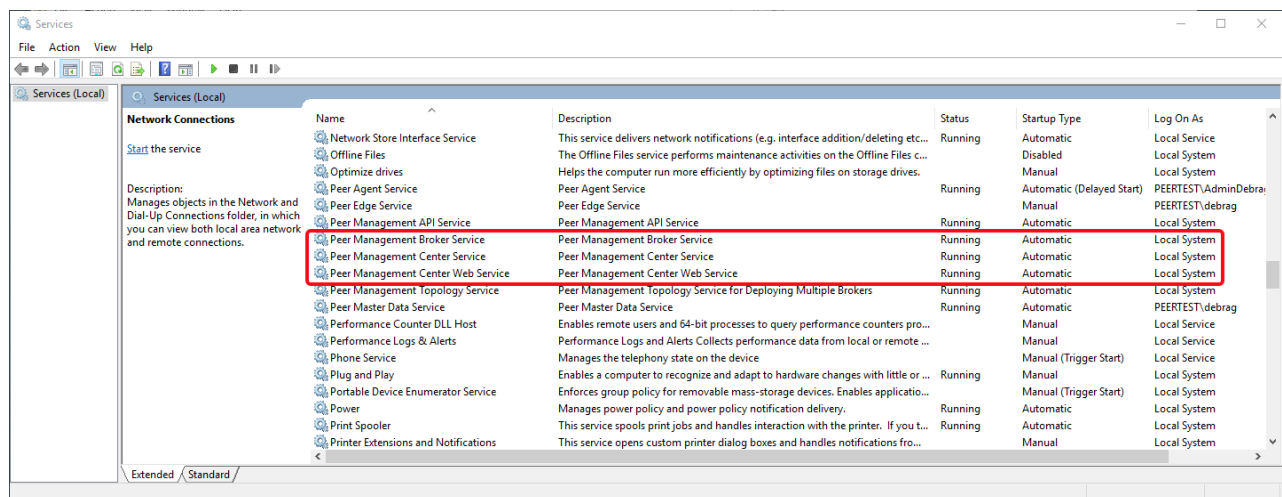


4. After the installation finishes, the Peer Agent is installed as a Windows service. You will need to verify that the **Peer Agent Service** is running, and that it was able to successfully connect to [Peer Management Broker](#). You can do this by opening the Windows Services Panel (services.msc) and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



Peer Services Required for Web Client

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



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

Updating Peer Management Center and Peer Agents

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

For details about updating Peer Management Center and Peer Agents, see:

- [Updating Peer Management Center](#)
- [Updating Peer Agents](#)

Updating Peer Management Center

Overview

There are two ways to check for software updates:

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


Note: When upgrading Peer Management Center (PMC) on a Windows Failover Cluster, the process closely resembles the steps involved in the initial installation.

Manually Checking for and Installing an Update

To manually check for an update:

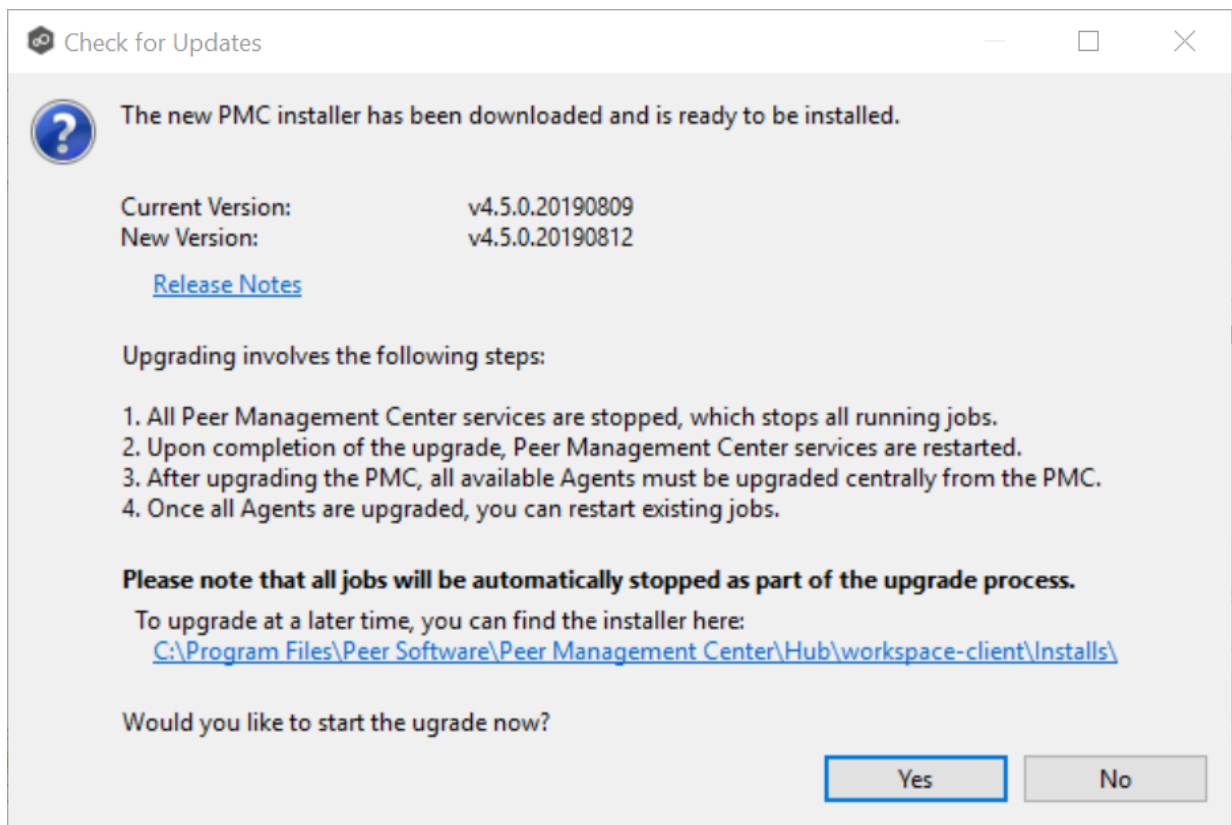
1. From the **Help** menu, select **Check for Updates**.

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



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

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



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

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

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

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

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

5. After the Peer Management Center upgrade is installed, update the Peer Agents. See [Updating Peer Agents](#) for details.


Updating Peer Agents

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

Note: For customers using clustered file server roles with Windows Failover Cluster, please review the Peer Software knowledge base article [Using a Peer Agent in a Windows Failover Cluster](#). The steps for upgrading Agents tied to clustered file server roles are the same as installing these Agents for the first time.

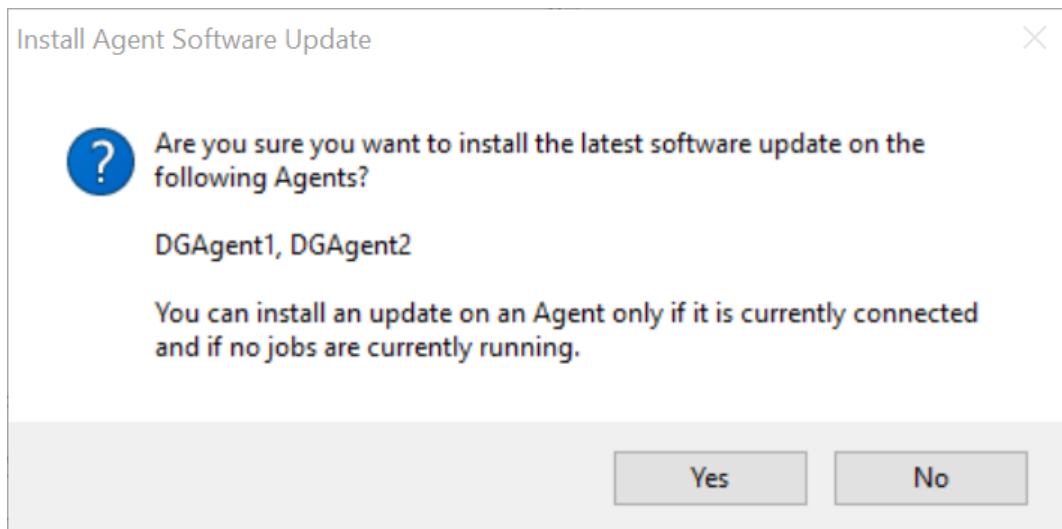
To update Peer Agents:

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




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

A confirmation dialog appears.



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

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



Uninstalling Peer Management Center and Agents

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 and Peer Agents.

Running Peer Management Center

Peer Management Center offers two graphical user interface options for accessing the application:

- **Rich client:** This option involves [installing](#) and running the rich client application directly on the server where Peer Management Center is deployed. For more details, see [Launching the Rich Client](#). Note that the rich client is available only on Windows and is not supported on Linux-based systems.
- **Web client:** Alternatively, you can directly log in via a web browser, either on the local Peer Management Center server or remotely from any system connected to the network that has access to the server. The web client provides the same robust functionality as the rich client and is highly responsive, closely mirroring its capabilities. This enables users to effectively manage and monitor jobs from remote locations without the need to directly access the Peer Management Center server. For more details, see [Accessing the Web Client](#).

Launching the Rich Client

After completing the Peer Management Center installation process, you have the option to access either the rich client or the web client. Note that the rich client is available only on Windows and is not supported on Linux-based systems.

To launch the Peer Management Center rich client, use one of the following methods:

- Utilize the shortcut generated during the installation process. This shortcut is commonly found on the desktop.
- Navigate to the Microsoft Windows Start menu, locate **Peer Software**, and select **Peer Management Center Client** from the menu.

Accessing the Web Client

Options for Accessing the Web Client

Upon completing the installation of Peer Management Center, configuring the Peer Management Center Web Client Service, and starting the [required Peer services](#), users can access the Peer Management Center web client through various methods:

- **Log in via a web browser:** Users can directly log in via a web browser on the local Peer Management Center server, as outlined in the subsequent section.
- **Remote network access:** Access is available from any system on the network with connectivity to the Peer Management Center server.
- **Launching the Peer Management Center Client:** Users can also access the Peer Management Center Web Client by launching the Peer Management Center Client (rich client application) and choosing **Open Peer Management Center Web Client** from the **Window** menu. Note that the rich client is available only on Windows.

Logging into the Web Client

To access the Peer Management Center web client:

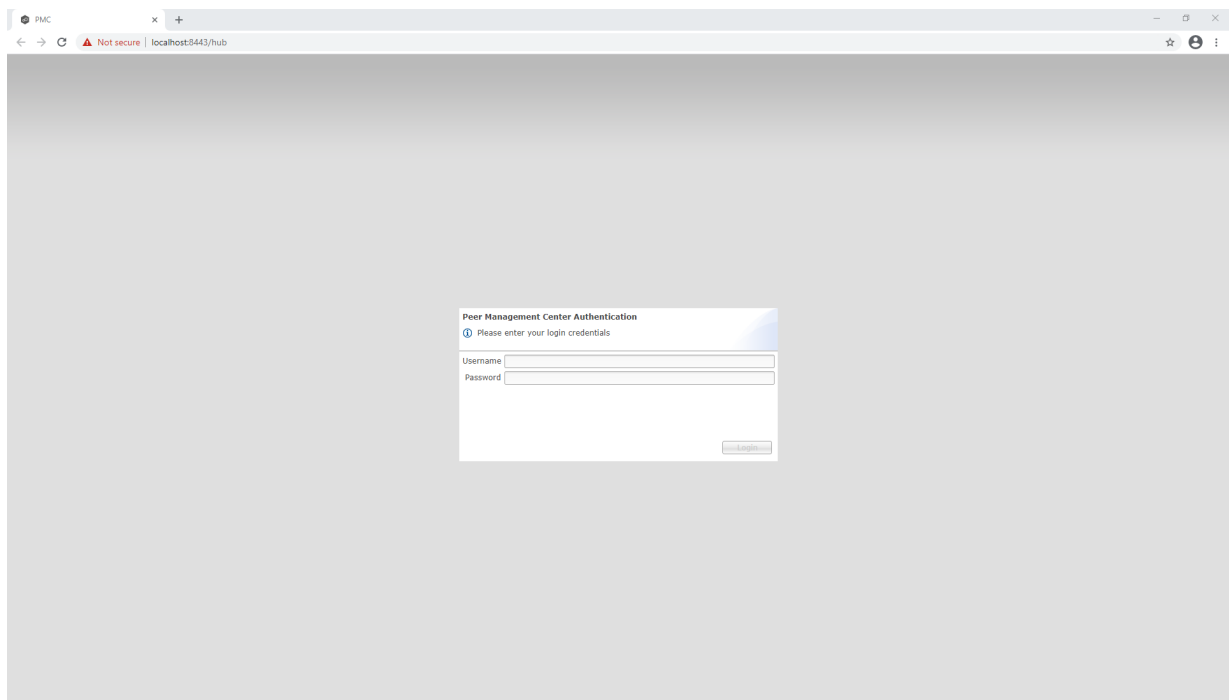
1. Open a web browser.
2. Enter one of the following URLs in the address field:

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

Notes:

- The URL required to access Peer Management Center depends on the configuration of the web service during installation. If you're unsure of the correct URL, you may need to reach out to the administrator who installed Peer Management Center for clarification.
- By default, the HTTPS port is set to 8443. However, if a different port is used in your environment, you should replace 8443 with the appropriate port number.
- **/hub** is required after the port number to reach the PMC client UI. Additionally, **/hub** needs to be appended after the port number to access the PMC client UI.
- To modify the web service configuration, navigate to the [General Configuration](#) page of [Preferences](#). For detailed instructions, refer to [Web and API Configuration](#) topic.

After the URL is entered, the login page is displayed.



3. Enter a user name and password.

- The default user name is **admin**; the default password for most installs is **password**. If you have deployed the PMC virtual appliance via the Microsoft Azure Marketplace, your password will be your Azure Subscription ID. If you have deployed the PMC virtual appliance via the Amazon AWS Marketplace, your default password will be the EC2 instance ID of the virtual appliance.
- For security reasons, we highly recommend that the user immediately changes the **admin** password. See [Editing an Internal User](#) for more information on changing account passwords.
- If logging in with an Active Directory account, enter the user name in this format: `username@mydomain.local`.

4. Click **Login**.

Peer Management Center User Interface

Peer Management Center is a comprehensive management console designed for configuring and deploying jobs, as well as providing summary and runtime information for these jobs. This section will guide you through the user interface, helping you familiarize yourself with the

various elements and functionalities available to you. By understanding the layout and features of the interface, you'll be better equipped to navigate and utilize Peer Management Center effectively.

The key components of the user interface include:

- [Main window](#) - This central hub serves as your primary workspace, where you'll access essential tools and functionalities to accomplish your tasks efficiently. Understanding the layout and components of the main window is crucial for navigating Peer Management Center.
- [Menus and toolbars](#) - Learn how to access different sections and functionalities through the main navigation menus and toolbar.
- [Perspectives and views](#) - Understand how perspectives and views are organized within the interface to display relevant information and tools.
- [Tables](#) - Tables are used to organize and present data in a structured format. Learn how to use tables to quickly find relevant information.

To explore ways to customize the interface to suit your preferences and workflow, see [Preferences](#).

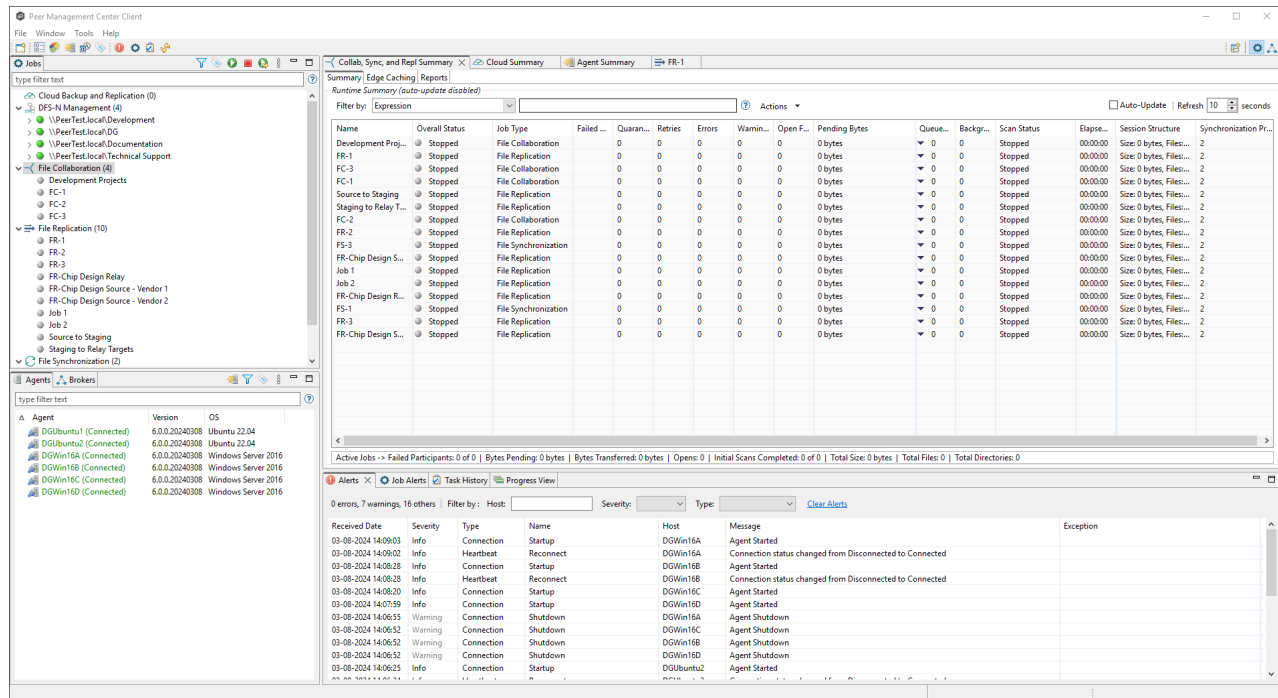
Accessing the interface

Peer Management Center offers two graphical user interface options for accessing the application:

- **Rich client:** The rich client application runs directly on the server where Peer Management Center is deployed. For more details, see [Launching the Rich Client](#). Note that the rich client is available only on Windows and is not supported on Linux-based systems.
- **Web client:** You can run the web client on the local Peer Management Center server or remotely from any system connected to the network that has access to the server. The web client provides the same robust functionality as the rich client and is highly responsive, closely mirroring its capabilities. This enables users to effectively manage and monitor jobs from remote locations without the need to directly access the Peer Management Center server. If you are using the web client, access to functionality is role-based. For more information about roles, see [Managing Roles for Web Client Users](#).

Main Window

The main window is your workspace, where you'll interact with content, create, edit, or view documents, and perform tasks specific to Peer Management Center. This area contains different perspective and views, depending on the context of your work.



Perspectives and Views

Overview




A **perspective** is a collection of [views](#) designed for a specific set of tasks or workflow. Switching between perspectives allows you to focus on different aspects of your work without cluttering the workspace with unnecessary tools or windows. Peer Management Center has two perspectives:


- The [Jobs](#) perspective gives you the ability to configure jobs and Agents, start and stop jobs, view summary and runtime information for these jobs, and view various alerts. This is the default perspective that appears when you open Peer Management Center.
- The [Topology](#) perspective gives you the ability to view and configure sites and brokers from a single view in the PMC, allowing you to create paths for Agents to communicate more directly with one another.

You can customize a perspective by displaying, hiding, detaching, and rearranging views to suit your preferences and workflow requirements. To revert the perspective to its default layout, use the **Reset Perspective** command on the **Window** menu.

Opening a Perspective




Here are two methods to open a perspective:

- Choose **Open Perspective** from the **Window** menu.
- Click the **Open Perspective** button located on the right side of the toolbar in the main window    , and then select a perspective:



Switching Between Perspectives

While multiple perspectives can be open simultaneously, only one perspective remains active at any given time.

To switch between perspectives, click the other perspective button:   

Resetting a Perspective

Perspectives initially arrive with a predefined layout, yet they're adaptable. You can modify the arrangement of views, add or remove views, and so forth, and these adjustments persist with those perspectives. To restore the original layout, simply choose **Reset Perspective** from the **Window** menu, reverting the layout back to its default configuration.

Views

Overview

A [view](#) is a graphical component within a perspective that displays specific information or provides functionality relevant to the current task or workflow. A view can contain one or more tabs. For example, the **Collab, Sync and Repl Summary** view contains three tabs:

The screenshot displays the Peer Management Center Client interface. The main window is titled 'Collab, Sync, and Repl Summary' and contains a table of job status. The table has columns for Name, Overall Status, Job Type, Failed, Quarantined, Retries, Errors, Warnings, Open Files, Pending Bytes, Queue, Background, Scan Status, Elapsed, Session Structure, and Synchronization. The table lists various jobs such as 'Development Proj...', 'FR-1', 'FC-3', 'FC-1', 'Source to Staging', 'Staging to Relay T...', 'FR-2', 'FR-3', 'FR-Chip Design S...', 'Job 1', 'Job 2', 'FR-Chip Design R...', 'FS-1', 'FR-3', and 'FR-Chip Design S...'. All jobs are listed as 'Stopped'.

Below the table, there is an 'Alerts' panel showing 28 errors, 0 warnings, and 4 others. The alerts are filtered by 'Host' and 'Severity'. The alerts list includes columns for Received Date, Severity, Type, Name, Host, and Message. The messages describe various errors such as 'Uncaught Exception running job DGWin16A_Technical Support', 'Quorum lost for DFS-N job DGWin16A_Development', and 'Operations-StartSession Command: Error starting session on agent DGWin16A_Invalid Ag...'. The alerts are dated from 03-08-2024 13:35:38 to 03-08-2024 13:35:36.

You can display, hide, detach, and rearrange views to suit your preferences and workflow requirements.

Displaying Views

You can open views by:

- Selecting a view from the **Window** menu.
- Clicking the **View** button in a toolbar and selecting an option from the **View** menu.

Resizing Views

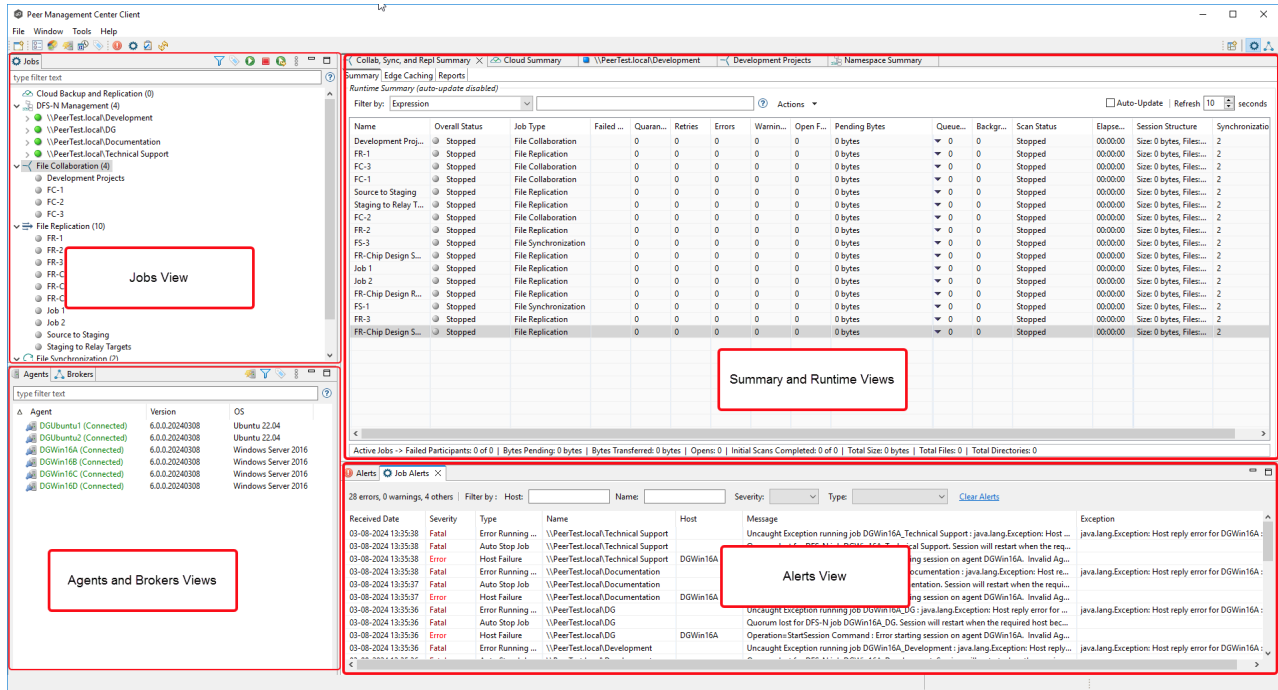
You can resize views in a variety of ways:

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

Jobs Perspective

The **Jobs perspective** can be divided into four quadrants; each quadrant displays information in panels called [views](#). There are various types of views. For example, some views display a combination of real-time file I/O activity, history, and configuration information for a specific job; others display a summary of information about all jobs of a specific type.

The **Jobs** perspective can be divided into four quadrants. See [Views in the Jobs Perspective](#) for a description of the views in the Jobs perspective.



Views in the Jobs Perspective

The views in the Jobs perspective are described in the following table.

Quadrant	Description
Upper left	Contains one view, the Jobs view , which displays a list of all jobs, grouped by job type. The toolbar in this view allows you to start and stop jobs.
Bottom left	<p>Contains two views:</p> <ul style="list-style-type: none"> • The Agents view displays a list of known Peer Agents and connection status for each. Individual Peer Agents can be updated and restarted from this view as well by right-clicking on one or more items and selecting the appropriate item from the pop-up menu. • The Brokers view displays a list of all known brokers installed in your environment, alongside the Agents connected to each respective broker. You can conveniently right-click on an Agent within this view to execute the same actions available in the Agents view.
Upper right	<p>Several types of views are displayed in this area, including:</p> <ul style="list-style-type: none"> • A dashboard that provides metrics and key performance indicators. • Summaries of jobs by job type. • An Agent Summary view, which displays a list of all known Peer Agents deployed and detailed status information that can be used to assess the health of the environment. • Runtime statistics for individual jobs.
Lower right	<p>Contains a variety of views, including:</p> <ul style="list-style-type: none"> • The Alerts view, which displays a list of Peer Management Center alerts that have occurred with detailed information about each alert. Alerts relating to Peer Agent connection status changes are reported here. • The Jobs Alerts view, which displays a list of job-specific alerts that have occurred. Alerts relating to the automatic stopping and restarting of jobs are displayed here.

The **Agents** view is displayed in the lower left quadrant of the Peer Management Center interface. It lists all known Peer Agents installed in your environment and displays the current [connection status](#) for each: Connected, Disconnected, Pending, or Unknown. The color of an Agent serves as a visual aid, making it easy to identify its version status—red indicates that the Agent is disconnected or needs updating, while green signifies that the Agent is connected and current.

When you right-click on an Agent, a menu of Agent-related commands is displayed. See [Managing Peer Agents](#) for more information.

This view is automatically displayed when Peer Management Center is started.

The screenshot displays the Peer Management Center Client interface. The 'Agents' view is highlighted with a red box, showing a list of agents with columns for Agent, Version, and OS. The main window shows a table of jobs and their statuses.

Name	Overall Status	Job Type	Failed	Quaran...	Retries	Errors	Warnin...	Open F...	Pending Bytes	Queue...	Backgr...	Scan Status	Elapse...	Session Structure
Development Projects	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-4	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FC-1	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FC-2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Source to Staging	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Staging to Relay Targets	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-3	Stopped	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-4	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Source - Vendor 1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FS-2	Stopped	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-5	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-6	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Source - Vendor 1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Source - Vendor 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Relay	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FS-1	Stopped	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-3	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-4	Stopped	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Source - Vendor 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...

Filtering a List of Agents


If you have a lengthy list of agents, you may want to filter it to view a smaller set. To filter a large list of agents, utilize the **Filter** field located below the Agents view toolbar. To filter a large list of Agents, use the **Filter** field located below the [Agents view toolbar](#). For further guidance on filtering agents, refer to the List Filters section.

Updating Peer Agent Software

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

Agents Toolbar

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



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

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

The screenshot displays the Peer Management Center interface. The main window shows a 'Summary Reports' section with a table of jobs. The table columns include Name, Overall Status, Job Type, Failed Hosts, Quarant., Retries, Errors, Warnings, Open Files, Pending Bytes, Pending Events, Queued Items, Background..., Scan Status, Elapsed TL..., and Session Structure. The jobs are categorized into File Collaboration, File Replication, and File Synchronization. Some jobs are marked as 'Halted (Quorum Lost)' or 'Stopped'.

Below the jobs table, there is an 'Alerts' view. The alerts are filtered by 'Host' and 'Severity'. The alerts table has columns for Received Date, Severity, Type, Name, Host, Message, and Exception. The alerts include various events such as 'Agent Started', 'User attempting to restart disconnected Agent', 'Started Blob Tiering Maintenance on: DGAGENT1', 'Started Nightly Delete Shadow Copies Job on: DGAGENT1', 'Started Blob Tiering Database Maintenance on: DGAGENT1', 'Started Process Retention Policy on: DGAGENT1', 'Started Nightly Purge Destination Job on: DGAGENT1', 'Connection status changed from Pending to Disconnected', 'Peer Agent on DGAgent2 has disconnected', 'Peer Agent on DGAgent2 connect status has changed from Connected to Disconnected. PL...', 'Connection status changed from Connected to Pending', 'Agent Started', and 'Connection status changed from Disconnected to Connected'.

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

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

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

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

Displaying the Alerts View

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

The **Brokers** view is displayed in the lower left quadrant, adjacent to the **Agents** view. It presents all known brokers installed in your environment in a hierarchical tree structure, with each broker represented as a node in the tree. It also displays the current status of each broker:

- Red - Disconnected
- Update - Update Required
- Yellow - Warning

If a broker needs to be updated, you can right-click and select **Install Software Update** without having to navigate to the Topology perspective.

By clicking the arrow icon next to a broker, you can expand that node to see the Agents connected to that particular broker. You can conveniently right-click on an Agent within this view to execute many of the same actions available in the Agents view.

This view is automatically displayed when Peer Management Center is started.

The screenshot displays the Peer Management Center Client interface. The main window shows a hierarchical tree structure of brokers under the 'Jobs' tab. The 'File Replication (13)' folder is expanded, showing various replication jobs and their statuses. A red box highlights the 'Agents' view in the bottom-left corner, which shows a list of agents connected to the selected broker. The 'Alerts' view in the bottom-right corner shows a list of alerts with columns for Received Date, Severity, Type, Name, Host, Message, and Exception.

Name	Overall Status	Job Type	Failed ...	Quaran...	Retries	Errors	Warnin...	Open F...	Pending Bytes	Queue...	Backgr...	Scan Status	Elapse...	Session Structure
Development Projects	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-4	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FC-1	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Source to Staging	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FC-2	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Staging to Relay Targets	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-3	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-4	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-5	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-6	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Source - Vendor 1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FS-2	Stopped	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-5	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-6	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Relay	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Source - Vendor 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
Job 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Relay	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FS-1	Stopped	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-3	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-4	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...
FR-Chip Design Source - Vendor 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...

The Dashboard is divided into two sections:

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

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

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

The screenshot shows the Peer Management Center Client interface. The main window is titled 'Dashboard' and has a refresh interval set to 10 seconds. The dashboard is divided into two main sections: 'Collab, Sync, and Repl' and 'Agents'.

The 'Collab, Sync, and Repl' section includes a 'Summary View' table with the following data:

1 Running with Quarantines!		Active Statistics		Active Watch Set	
Running with Quarantines	1	Failed Participants	0 of 2	Total Size:	713.23 MB
Running with Disconnected Agents	0	Bytes Pending:	0 bytes	Total Files:	11920
Lost Quorum	0	Bytes Transferred:	0 bytes	Total Directories:	116
Not Running - Stopped	12	Open:	0		
Running in Good State	1	Initial Scans Completed:	2 of 2		

The 'Agents' section includes an 'Agent Summary' table with the following data:

All connected		Agent Memory Load		Top Connectivity Offenders	
Connected	●	0-49%	●		
Needing Upgrade	●	50-74%	●		
Disconnected	●	75-100%	●		

Below the dashboard is an 'Alerts' section with a table of recent alerts:

Received Date	Severity	Type	Name	Host	Message	Exception
04-21-2021 01:31:43	Info	Connection	Restart Agent Service	DGAgent2	Agent Started	
04-21-2021 01:00:00	Info	Agent	Scheduled Task Started	DGAgent1	Started Blob Tiering Maintenance on: DGAGENT1	
04-20-2021 23:59:00	Info	Agent	Scheduled Task Started	DGAgent1	Started Nightly Delete Shadow Copies Job on: DGAGENT1	
04-20-2021 23:29:59	Info	Agent	Scheduled Task Started	DGAgent1	Started Blob Tiering Database Maintenance on: DGAGENT1	
04-20-2021 22:37:00	Info	Agent	Scheduled Task Started	DGAgent1	Started Process Retention Policy on: DGAGENT1	
04-20-2021 20:09:59	Info	Agent	Scheduled Task Started	DGAgent1	Started Nightly Purge Destination Job on: DGAGENT1	
04-20-2021 16:02:20	Error	Heartbeat	Disconnect	DGAgent2	Connection status changed from Pending to Disconnected	
04-20-2021 16:02:20	Warning	Heartbeat	Peer Agent on DGAgent2 has disc...	DGAgent2	Peer Agent on DGAgent2 connect status has changed from Connected to Disconnected. PL...	
04-20-2021 16:02:00	Warning	Heartbeat	Missed Heartbeat	DGAgent2	Connection status changed from Connected to Pending	
04-20-2021 14:25:00	Info	Connection	Startup	DGAgent2	Agent Started	
04-20-2021 14:25:00	Info	Heartbeat	Reconnect	DGAgent2	Connection status changed from Disconnected to Connected	
04-20-2021 14:24:59	Info	Connection	Startup	DGAgent2	Agent Started	

Displaying the Dashboard

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

- Select **Show Dashboard** from the **Window** menu.
- Click the **Show Dashboard** icon in the main [Peer Management Center toolbar](#).

- Set the Dashboard to launch automatically at start. See [General Configuration](#).

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

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

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

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

The screenshot shows the Peer Management Center Client interface. The main window displays a table of jobs with columns: Name, Overall Status, Job Type, Failed Hosts, Quorum, Retries, Errors, Warnings, Open Files, Pending Bytes, Pending Events, Queued Items, Background..., Scan Status, Elapsed Tl., and Session Structure. The Job Alerts window is open at the bottom right, showing a list of alerts with columns: Received Date, Severity, Type, Name, Host, Message, and Exception.

Received Date	Severity	Type	Name	Host	Message	Exception
04-21-2021 02:49:01	Error	Configuration	FS-5	DGAgent2.DG...	Disabled auto restart because maximum # (13) of restarts attempts was exceeded for hosts...	
04-21-2021 02:47:59	Error	Start Job	FS-5	DGAgent1	Host Reconnect Startup Error. Operations\Install Watch Directories : ERROR Registering job...	
04-21-2021 02:47:01	Info	Auto Start Job	FS-5	DGAgent1	Auto starting job. DGAgent1 host is now available.	
04-21-2021 02:46:00	Error	Start Job	FS-5	DGAgent1	Host Reconnect Startup Error. Operations\Install Watch Directories : ERROR Registering job...	
04-21-2021 02:45:01	Info	Auto Start Job	FS-5	DGAgent1	Auto starting job. DGAgent1 host is now available.	
04-21-2021 02:44:01	Error	Configuration	FS-1	DGAgent2.DG...	Disabled auto restart because maximum # (13) of restarts attempts was exceeded for hosts...	
04-21-2021 02:44:00	Error	Start Job	FS-5	DGAgent1	Host Reconnect Startup Error. Operations\Install Watch Directories : ERROR Registering job...	
04-21-2021 02:43:01	Info	Auto Start Job	FS-5	DGAgent1	Auto starting job. DGAgent1 host is now available.	
04-21-2021 02:42:56	Error	Start Job	FS-1	DGAgent1	Host Reconnect Startup Error. Operations\Install Watch Directories : ERROR Registering job...	
04-21-2021 02:42:01	Info	Auto Start Job	FS-1	DGAgent1	Auto starting job. DGAgent1 host is now available.	
04-21-2021 02:41:58	Error	Start Job	FS-5	DGAgent1	Host Reconnect Startup Error. Operations\Install Watch Directories : ERROR Registering job...	
04-21-2021 02:41:01	Info	Auto Start Job	FS-5	DGAgent1	Auto starting job. DGAgent1 host is now available.	

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

Displaying Job Alerts

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

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

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

The screenshot displays the Peer Management Center Client interface. The main window is titled "Peer Management Center Client" and shows a dashboard with various tabs like "Dashboard", "Agent Summary", "Cloud Summary", etc. The "Jobs" view is active, showing a list of jobs grouped by type. The "Alerts" view is also visible at the bottom, showing a list of alerts with columns for Received Date, Severity, Type, Name, Host, Message, and Exception.

Name	Overall Status	Job Type	Failed Hosts	Quaran...	Retries	Errors	Warnings	Open Files	Pending Bytes	Pending Events	Queued Items	Background...	Scan Status	Elapsed Ti...	Session Structure
FC-4	Running	File Collaboration	8	0	0	1	0	12 KB	0	0	0	0	Scanning Directories -	00:44:05	Size: 367.32 MB, ...
FR-4	Stopped	File Replication	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:00	Size: 0 bytes, File...
FC-2	Running	File Collaboration	0	0	0	3	0	0 bytes	0	0	0	0	Completed - 00:00:41	00:44:23	Size: 345.92 MB, ...
FC-5	Stopped	File Collaboration	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:00	Size: 0 bytes, File...
FS-1	Stopped	File Synchronization	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:04	Size: 0 bytes, File...
FS-3	Stopped	File Synchronization	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:00	Size: 0 bytes, File...
FR-1	Stopped	File Collaboration	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:00	Size: 0 bytes, File...
FR-2	Contacting participants ...	File Replication	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:07	Size: 0 bytes, File...
FR-4	Contacting participants ...	File Synchronization	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:01	Size: 0 bytes, File...
FS-5	Installing Directory Watchers ...	File Synchronization	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:20	Size: 0 bytes, File...
FR-1	Stopped	File Replication	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:00	Size: 0 bytes, File...
FS-6	Stopped	File Synchronization	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:00	Size: 0 bytes, File...
FR-2	Stopped	File Synchronization	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:00	Size: 0 bytes, File...
FR-3	Stopped	File Replication	0	0	0	0	0	0 bytes	0	0	0	0	Stopped	00:00:00	Size: 0 bytes, File...

Active Jobs -> Failed Participants: 0 of 2 | Bytes Pending: 12 KB | Bytes Transferred: 325.01 MB | Opens: 0 | Initial Scans Completed: 1 of 6 | Total Size: 713.23 MB | Total Files: 11917 | Total Directories: 116

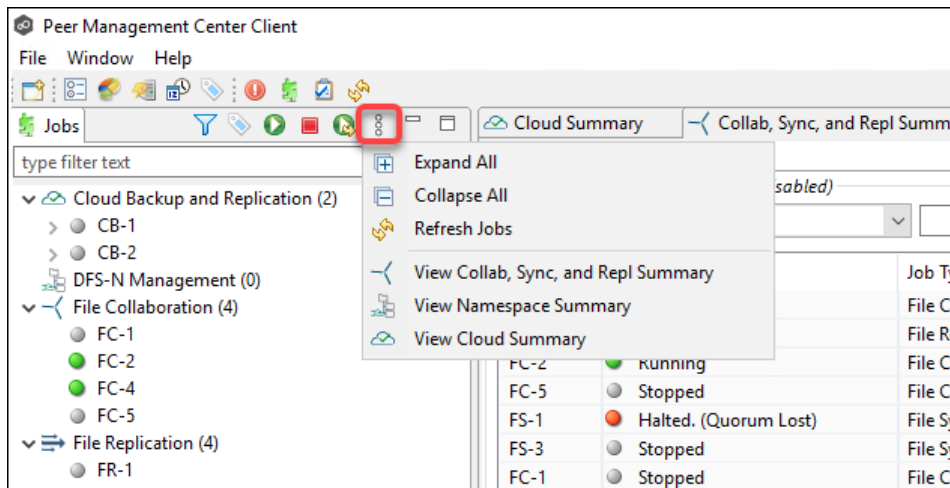
Received Date	Severity	Type	Name	Host	Message	Exception
04-20-2021 14:25:00	Info	Connection	Startup	DGAgent2	Agent Started	
04-20-2021 14:25:00	Info	Heartbeat	Reconnect	DGAgent2	Connection status changed from Disconnected to Connected	
04-20-2021 14:24:37	Warning	Connection	Shutdown	DGAgent2	Agent Shutdown	
04-20-2021 13:47:45	Warning	Agent	Connection	DGAgent2	Connection to Broker DGPMAC1 was lost for 1 minute 13 seconds	
04-20-2021 13:47:45	Warning	Agent	Connection	DGAgent1	Connection to Broker DGPMAC1 was lost for 1 minute 13 seconds	

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

- Double-clicking any job name in the list will display a [runtime view](#) of that job.
- Double-clicking any job type name in the list will display a [summary view](#) of that job type.


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

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



Jobs Toolbar

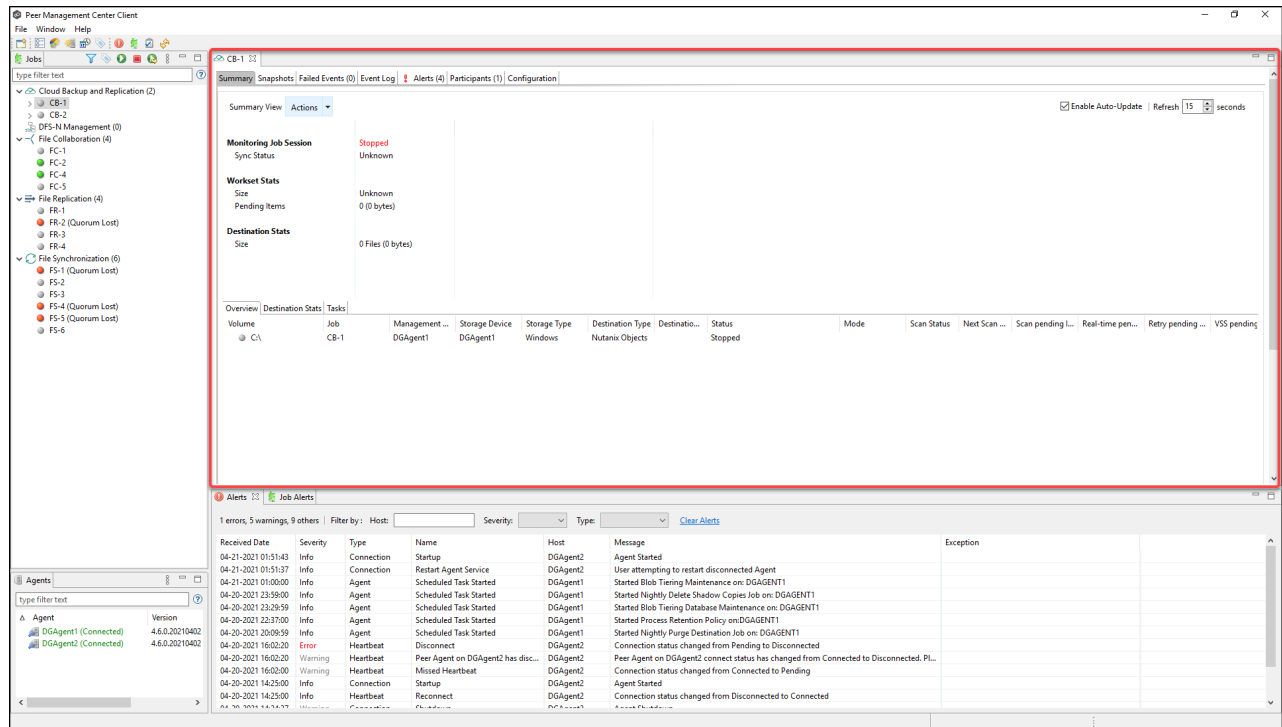
The following buttons are available on the toolbar within the **Jobs** view:



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

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

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



The runtime views include:

- [Cloud Backup and Replication job](#)
- [DFS-N Management job](#)
- [File Collaboration job](#)
- [File Replication job](#)
- [File Synchronization job](#)

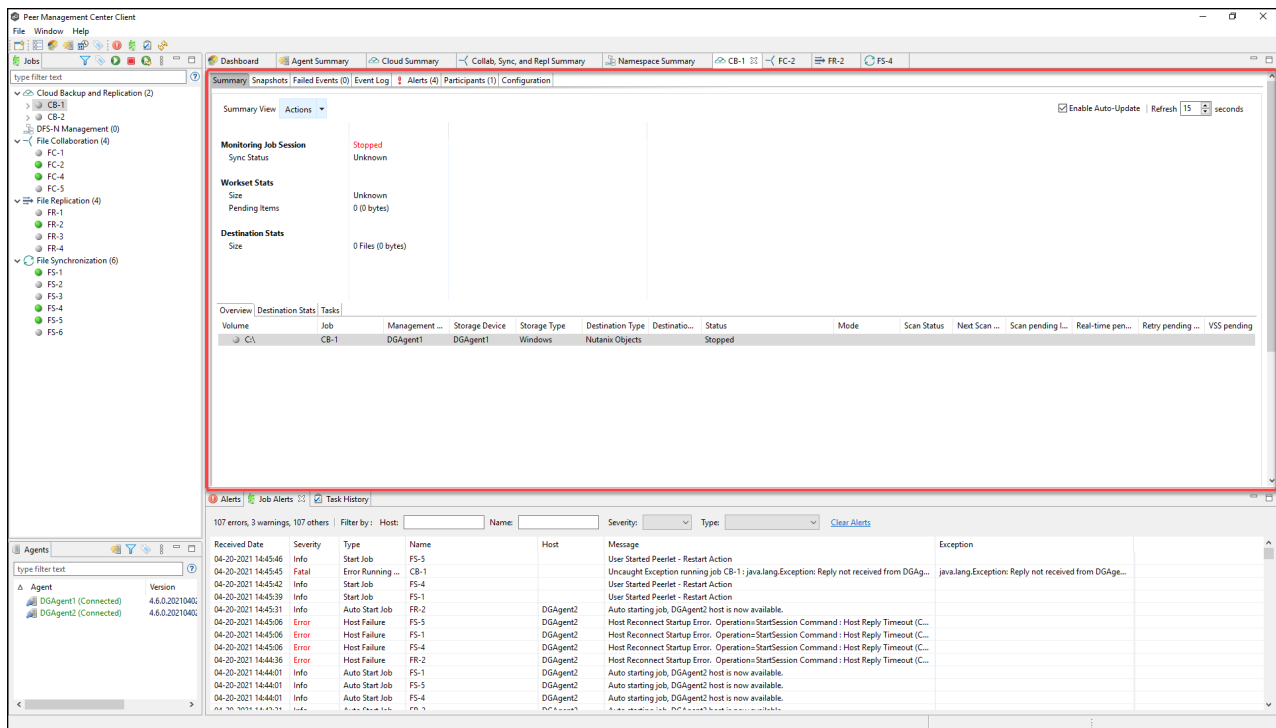
Cloud Backup and Replication Job Runtime View

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

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

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

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

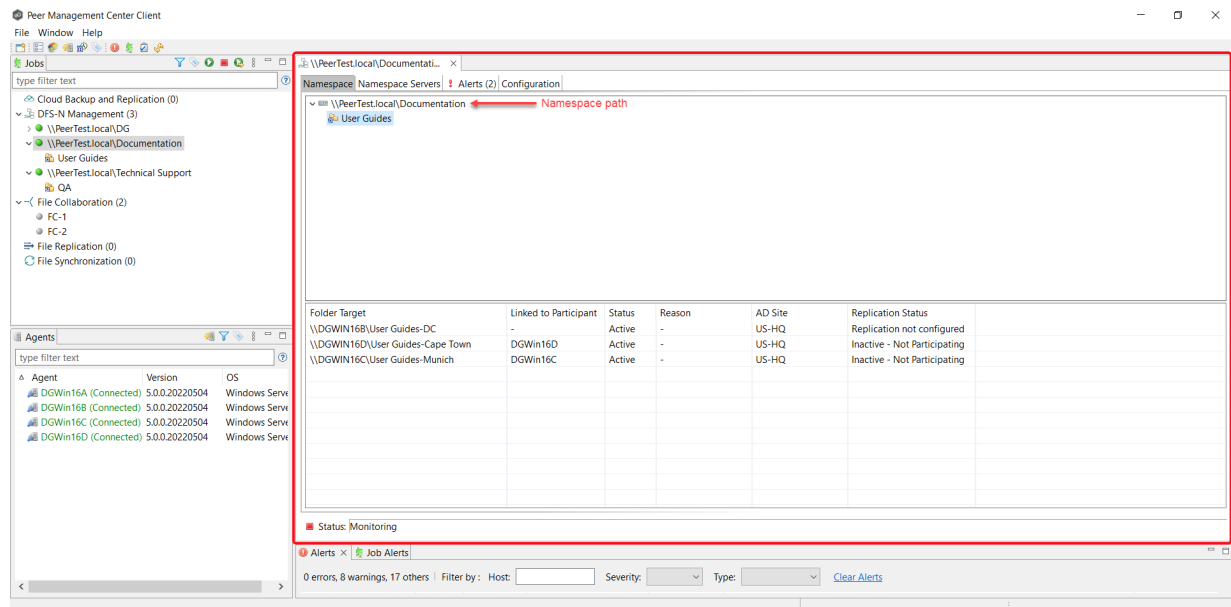


DFS-N Management Job Runtime View

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

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

- **Namespace** – The top panel of the tab displays the namespace folders in a tree structure. The namespace path is shown at the top of the tree. The bottom panel displays the folder targets linked to the selected namespace folder.
- **Namespace Servers** – Displays a list of the namespace servers and folder targets for the namespace selected in the top panel.
- **Alerts** – Displays a log of alerts that were issued for the job.
- **Configuration** – Displays a summary of the job configuration.



File Collaboration Job Runtime View

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

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

The view contains the following eight tabs:


- [Summary](#) - Displays overall statistics for the selected job.
- [Session](#) - Displays active open files and files that are currently in transit between [participating hosts](#).
- [Event Log](#) - Displays a list of all runtime activity that has occurred within the selected job.

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

The screenshot displays the Peer Management Center Client interface. The main window shows the 'Summary' tab for a job, with a red box highlighting the job status and activity metrics. Below this, a timeline chart shows the job's progress from 03:00 PM to 02:00 PM. At the bottom, an 'Alerts' table lists various error messages.

Received Date	Severity	Type	Name	Host	Message	Exception
04-20-2021 14:28:11	Error	Start Job	FS-1	DGAgent1	Host Reconnect Startup Error. Operations=Install Watch Directories: ERROR Registering job...	
04-20-2021 14:28:31	Info	Auto Start Job	FS-5	DGAgent1	Auto starting job, DGAgent1 host is now available.	
04-20-2021 14:28:20	Error	Start Job	FS-5	DGAgent1	Host Reconnect Startup Error. Operations=Install Watch Directories: ERROR Registering job...	
04-20-2021 14:27:31	Info	Auto Start Job	FR-2	DGAgent1	Auto starting job, DGAgent1 host is now available.	
04-20-2021 14:27:18	Error	Start Job	FR-2	DGAgent1	Host Reconnect Startup Error. Operations=Install Watch Directories: ERROR Registering job...	
04-20-2021 14:27:01	Info	Auto Start Job	FS-1	DGAgent1	Auto starting job, host is now available.	
04-20-2021 14:26:59	Info	Start Job	CB-1		User Started Peerlet	
04-20-2021 14:26:24	Error	Start Job	FS-1	DGAgent1	Host Reconnect Startup Error. Operations=Install Watch Directories: ERROR Registering job...	
04-20-2021 14:25:38	Info	Start Job	FS-5		User Started Peerlet - Restart Action	
04-20-2021 14:25:35	Info	Start Job	FS-4		User Started Peerlet - Restart Action	
04-20-2021 14:25:32	Info	Start Job	FR-2		User Started Peerlet - Restart Action	
04-20-2021 14:25:28	Info	Start Job	FS-1		User Started Peerlet	

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



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

Activity

This section presents statistics on pending activity:

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

Replication Status

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

- **Bytes Transferred** – Total number of bytes transferred for all real-time Add, Bulk Add, Modify, and Scan synchronization. This does not include bulk delete, security or renames.
- **Added** – Total number of files and folders added in real-time. Move your cursor over the field to see the breakdown for folders and files.
- **Updated** – Total number of files synchronized by initial scan or real-time.
- **Deleted** – Total number of files and folders deleted.
- **Renamed** – Total number of files and folders renamed. Move your cursor over the field to see the breakdown for folders and files.
- **File Metadata Updates** – Total number of real-time and scan metadata updates for folders and files.

Background Scan

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

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

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

Development Projects

Summary | Session | Event Log | Quarantines (0) | Retries (0) | Alerts (0) | Participants (2) | Configuration

Open Files (0) (Auto-Update Disabled)

Session Status: Running | Filter by Host: | Filter by: | Actions

Auto-Update | Refresh 10 seconds

File Path	Host	Is Source	User Name	Sync. Status	Sync. Position	Transfer Rate	File Size	Last Modified	Date Opened	Message	Attribut...	Process Name	Is Directory	Access Level	Share Mode	Is Open	Success	Client IP
-----------	------	-----------	-----------	--------------	----------------	---------------	-----------	---------------	-------------	---------	-------------	--------------	--------------	--------------	------------	---------	---------	-----------

Status: Running

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

Component	Description
Open Files table	<p>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 header will sort by that column in ascending or descending order.</p> <p>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.</p>
Session Status field	<p>Field indicating the current status of the session. Valid values are:</p> <ul style="list-style-type: none"> • Stopped: Session is stopped. • Starting: Session is starting up. • Collaborating: Real-time event detection is enabled.
Filter by Host list	<p>A drop-down list of participating hosts to filter on. Selecting a specific host will filter the open files to show files on that host only.</p>
Filter by list	<p>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.</p>
Actions menu	<p>Refresh View: Refreshes the entire Open Files table to show the latest list of file transfers and locks.</p>

The **Event Log** tab allows you to view recent file event history for the currently running File Collaboration, File Synchronization, or File Replication job. You can specify the maximum number of events to show 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 of Peer Management

Center. 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.

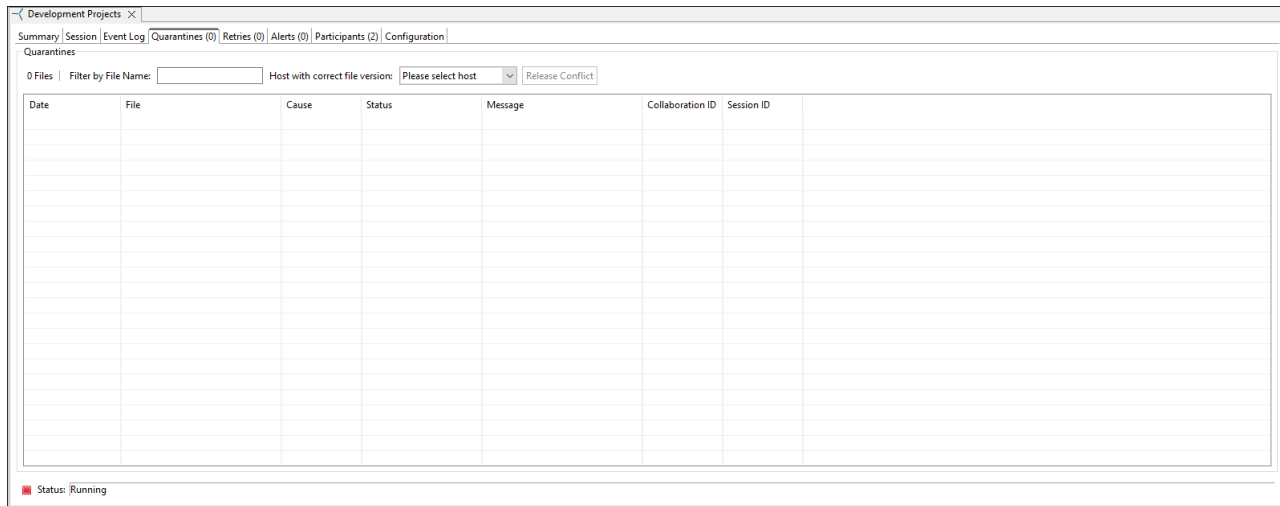
You can click any column header to sort by the column. For example, clicking the **File** column will sort by file name and you will be able to view all file events for that file in chronological order. Warnings are displayed in light gray, errors are displayed in red, and fatal errors are displayed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

Date	Severity	Type	Host	Is Source	Collaboration ID	File	Comments	Message	Username	File Size	Delta Size	Modified Time	Start Date	End Date	Root Path	Event ID	Client IP
04-02-2024 00:20:52	INFO	Scan Complete		true				Scan completed for pat...					04-02-2024 00:20:52				
04-02-2024 00:20:50	INFO	Scan Start		true				Scan Type: Full Directory					04-02-2024 00:20:50				
04-02-2024 00:20:37	INFO	Watch Directory	DGWin16C	true									04-02-2024 00:20:36	04-02-2024 00:20:37	C:\Data\Projects		
04-02-2024 00:20:37	INFO	Watch Directory	DGWin16B	true									04-02-2024 00:20:36	04-02-2024 00:20:37	C:\Data\Projects		
04-02-2024 00:20:36	INFO	Install File Dis...	DGWin16C	true									04-02-2024 00:20:36	04-02-2024 00:20:36	C:\Data\Projects		
04-02-2024 00:20:36	INFO	Install File Dis...	DGWin16B	true									04-02-2024 00:20:36	04-02-2024 00:20:36	C:\Data\Projects		
04-02-2024 00:20:31	INFO	Job Started	DGWin16C	true									04-02-2024 00:20:28	04-02-2024 00:20:31	C:\Data\Projects		
04-02-2024 00:20:31	INFO	Job Started	DGWin16B	true									04-02-2024 00:20:28	04-02-2024 00:20:31	C:\Data\Projects		

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

Option	Description
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.

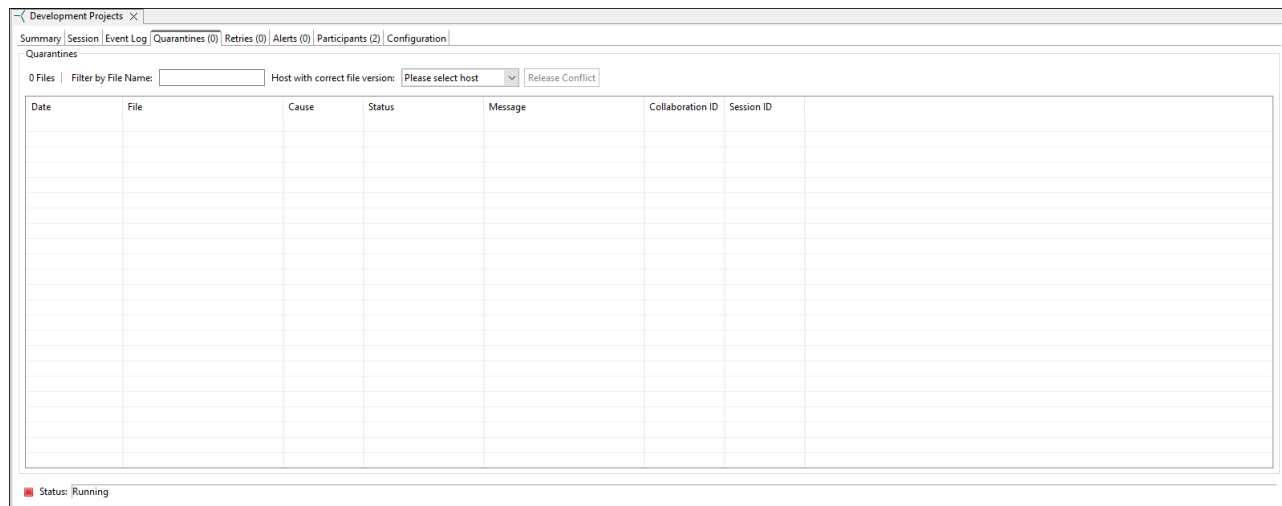
The **Quarantines** tab displays a list of files (a) for which file conflicts cannot be automatically resolved or (b) retries have failed after the maximum number of attempts. Files in this list will no longer be synchronized or protected with file locks until a winning file is picked through the PeerGFS user interface.



The context menu for the table contains the following actions:

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

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



The context menu for the table contains the following actions:

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

The **Alerts** tab allows you to view any alerts relevant to the running File Collaboration, File Synchronization, or File Replication job. You can specify the maximum number of alerts to view in the table by adjusting the Display Alerts spinner located in the top right corner of the panel. The alerts are also written to a tab delimited file named **fc_alert.log** within the subdirectory 'Hub/logs' within the installation directory of Peer Management Center.

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 job. 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.

Received Date	Severity	Type	Host	Message
04-20-2021 14:26:20	WARNING	Application	DGAgent1	Recursive reparse point folder detected during scan, excluding folder path=C:\Users\Public\Documents\...
04-20-2021 14:26:20	WARNING	Application	DGAgent1	Recursive reparse point folder detected during scan, excluding folder path=C:\Users\Public\Documents\...
04-20-2021 14:26:20	WARNING	Application	DGAgent1	Recursive reparse point folder detected during scan, excluding folder path=C:\Users\Public\Documents\...
04-20-2021 14:24:41	ERROR	Application	DGAgent2	Agent service on host DGAgent2 was shutdown while job was running.
04-20-2021 13:58:33	WARNING	Application	DGAgent1	Recursive reparse point folder detected during scan, excluding folder path=C:\Users\Public\Documents\...
04-20-2021 13:58:33	WARNING	Application	DGAgent1	Recursive reparse point folder detected during scan, excluding folder path=C:\Users\Public\Documents\...
04-20-2021 13:58:33	WARNING	Application	DGAgent1	Recursive reparse point folder detected during scan, excluding folder path=C:\Users\Public\Documents\...
04-20-2021 13:56:17	WARNING	Application	DGAgent1	Unsupported Host Configuration: 8.3 short file name is enabled for Host DGAgent1
04-20-2021 13:52:37	ERROR	Application	DGAgent2	Host Reconnect Startup Error. Operation=StartSession Command : Host Reply Timeout (Connected)
04-20-2021 13:49:30	ERROR	Application	DGAgent2	Host Reconnect Startup Error. Operation=StartSession Command : Host Reply Timeout (Connected)

The context menu for the table contains the following actions:

Action	Description
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** tab is divided into two sections:

- [Host Participants](#)
- [Host Participant State Change Log](#)

Host	Root Path	Status	State	Message	Local Root Path	Status Date	Message Time	Edge Caching Role	Total Disk Space	Free Disk Sp...	Trash Bin Size	Storage Version	Disk Space Last Up...	Trash Bin Size Last ...	Last Started	Last Stopped
DGWin16B	C:\Data\Projects	Participating	Active			04-02-2024 00:24:38		None	59.45 GB	37.95 GB	0 bytes	Windows Server 2016 Dat...	Tue Apr 02 01:01:5...	Tue Apr 02 00:21:5...	Tue Apr 02 00:21...	N/A
DGWin16C	C:\Data\Projects	Participating	Active			04-02-2024 00:24:38		None	59.45 GB	42.88 GB	0 bytes	Windows Server 2016 Dat...	Tue Apr 02 01:02:4...	Tue Apr 02 00:21:4...	Tue Apr 02 00:21...	N/A

Date	Severity	Host	Status	State	Message	Exception	Ref ID
04-02-2024 00:24:38	Unknown	DGWin16B	Participating	Active			
04-02-2024 00:24:38	Unknown	DGWin16C	Participating	Active			
04-02-2024 00:20:27	Unknown	DGWin16C	Participating	Active			
04-02-2024 00:20:27	Unknown	DGWin16B	Participating	Active			
04-02-2024 00:17:04	Unknown	DGWin16B	Not Participating	Inactive			
04-02-2024 00:17:04	Unknown	DGWin16C	Not Participating	Inactive			

Host Participants

The **Host Participants** section contains a table that displays all the current [host participants](#) for the selected File Collaboration, File Synchronization, or File Replication 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:

Action	Description
Disable Host Participant	Temporarily disables the selected participant from taking part in the 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 job that is currently being viewed. The cancellation of the auto-restart functionality for the host will only be in effect until the next time you start the 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.

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 a job, or when a host came back online.

The **Host Participant State Change Log** is a log of all host participant status changes (e.g., Participating, Not Participating) 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:

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

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

The screenshot shows the 'Configuration' tab of the Peer Global File Service configuration wizard. The interface is divided into several sections:

- Currently Running Configuration Summary:** Shows the job name 'Development Projects' and its status as 'Running'.
- Selected Participants and Configurations:** Lists the storage platforms for 'DGWin168' and 'DGWin16C'.
- Edge Caching Configuration:** Includes 'General Settings' (Job Name, ID, Type, Block Size, Checksums, etc.) and 'Tags' (currently None).
- Selected File Filters:** Shows 'Windows Default (General) - 0' and 'File Collaboration Sync Only (Synchronization Only) - 4' with their respective included/excluded patterns and filters.
- Selected Scheduled Replication Filters:** Currently shows 'No Scheduled Replication Filters Selected'.
- Locking:** Configures options like 'Allow Write Access During Synchronization' (true), 'Exclusive Target Lock' (false), and 'Include MS Office User Lock Information' (true).
- Conflict Resolution:** Sets 'Latest Modified Time (Truncate Milliseconds=true)', 'Quarantine Multi-edit Conflicts' (true), and 'Offline Folder Rename Detection' (false).
- Delta Replication Settings:** Configures 'Enable Block/Byte Synchronization' (true), 'Checksum Transfer Size' (256 KB), and 'Delta Block Transfer Size' (1024 KB).
- File Metadata Settings:** Controls synchronization of attributes and ACLs.
- Logging Alerts:** Enables logging and sets severity to 'INFO'.
- Target Protection Settings:** Enables target protection and sets backup retention (3 days).
- Email Alerts:** Currently disabled.
- SNMP Notifications:** Currently disabled.

File Replication Job Runtime View

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

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

- **Summary** - Displays overall statistics for the selected job.
- **Session** - Displays active open files and files that are currently in transit between [participating hosts](#).
- **Event Log** - Displays a list of all runtime activity that has occurred within the selected job.
- **Quarantines** - Displays a list of all files that are quarantined for the session or are in conflict between two or more participating hosts.
- **Retries** - Displays a list of files that are currently in the Retries list.
- **Alerts** - Displays a list of all job alerts specifically tied to the selected job.
- **Participants** - Displays a list of all hosts participating in the selected job.
- **Configuration** - Displays a summary of all configurable options for the selected job.

The screenshot displays the Peer Management Center Client interface. The main window shows the 'Summary' tab for a File Replication job. The job status is 'Stopped'. The 'Watch Set' section shows 0 total files and 0 total folders. The 'Activity' section shows 0 active opens, 0 files pending, 0 bytes pending, and 0 metadata pending. The 'Replication Status' section shows 0 bytes transferred, 0 added, 0 updated, and 0 metadata updates. The 'Alerts' tab at the bottom shows 111 errors, 3 warnings, and 113 others. The alerts table includes the following data:

Received Date	Severity	Type	Name	Host	Message	Exception
04-20-2021 14:47:25	Info	Start Job	FS-1		User Started Peerlet - Restart Action	
04-20-2021 14:47:07	Info	Stop Job	FS-1		User Stopped Peerlet - Restart Action	
04-20-2021 14:47:02	Info	Auto Start Job	FS-5	DGAgent2	Auto starting job, DGAgent2 host is now available.	
04-20-2021 14:47:01	Info	Start Job	FS-1		User Started Peerlet - Restart Action	
04-20-2021 14:46:58	Info	Start Job	FS-4		User Started Peerlet - Restart Action	
04-20-2021 14:46:54	Info	Start Job	FR-2		User Started Peerlet - Restart Action	
04-20-2021 14:46:50	Error	Host Failure	FS-5	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command : Host Reply Timeout (C...	
04-20-2021 14:46:47	Error	Host Failure	FS-4	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command : Host Reply Timeout (C...	
04-20-2021 14:46:44	Error	Host Failure	FS-1	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command : Host Reply Timeout (C...	
04-20-2021 14:46:36	Error	Host Failure	FR-2	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command : Host Reply Timeout (C...	
04-20-2021 14:45:46	Info	Start Job	FS-5		User Started Peerlet - Restart Action	
04-20-2021 14:45:45	Fatal	Error Running ...	CB-1		Uncaught Exception running job CB-1: java.lang.Exception: Reply not received from DGAge... java.lang.Exception: Reply not received from DGAge...	

File Synchronization Job Runtime View

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

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

- **Summary** - Displays overall statistics for the selected job.
- **Session** - Displays active open files and files that are currently in transit between [participating hosts](#).
- **Event Log** - Displays a list of all runtime activity that has occurred within the selected job.
- **Quarantines** - Displays a list of all files that are quarantined for the session or are in conflict between two or more participating hosts.
- **Retries** - Displays a list of files that are currently in the Retries list.
- **Alerts** - Displays a list of all job alerts specifically tied to the selected job.
- **Participants** - Displays a list of all hosts participating in the selected job.
- **Configuration** - Displays a summary of all configurable options for the selected job.

The screenshot displays the Peer Management Center Client interface. The main window shows the runtime view for a File Synchronization job named 'FS-4'. The job status is 'Stopped'. The summary view includes the following statistics:

Category	Value	Sub-category	Value
Job Status	Stopped	Elapsed Time	
Watch Set	0	Total Folders	0
Activity			
Active Opens	0	File Quarantines	0
Files Pending	0	Renames Pending	0
Bytes Pending	0 bytes	Deletes Pending	0
Metadata Pending	0	File Retries	0
Replication Status			
Bytes Transferred	0 bytes	Delta-level Savings	0
Added	0	Renamed	0
Updated	0	Deleted	0
Metadata Updates	0		

Below the summary view is a timeline graph showing activity from 01:25 PM to 02:50 PM. The legend indicates: Priority Queue (blue), Replication Queue (orange), and Active Opens (green).

The Alerts tab shows 118 errors, 3 warnings, and 120 others. The alerts table is as follows:

Received Date	Severity	Type	Name	Host	Message	Exception
04-20-2021 14:50:11	Info	Start Job	FS-5		User Started Peerlet - Restart Action	
04-20-2021 14:50:06	Error	Host Failure	FS-5	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command: Host Reply Timeout (C...	
04-20-2021 14:50:05	Info	Start Job	FS-1		User Started Peerlet - Restart Action	
04-20-2021 14:50:02	Info	Start Job	FS-4		User Started Peerlet - Restart Action	
04-20-2021 14:49:58	Info	Start Job	FR-2		User Started Peerlet - Restart Action	
04-20-2021 14:49:36	Error	Host Failure	FS-4	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command: Host Reply Timeout (C...	
04-20-2021 14:49:06	Error	Host Failure	FR-2	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command: Host Reply Timeout (C...	
04-20-2021 14:49:02	Info	Auto Start Job	FS-5	DGAgent2	Auto starting job. DGAgent2 host is now available.	
04-20-2021 14:48:32	Info	Auto Start Job	FS-4	DGAgent2	Auto starting job. DGAgent2 host is now available.	
04-20-2021 14:48:29	Error	Host Failure	FS-1	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command: Host Reply Timeout (C...	
04-20-2021 14:48:06	Error	Host Failure	FS-5	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command: Host Reply Timeout (C...	
04-20-2021 14:48:02	Error	Host Failure	FS-4	DGAgent2	Host Reconnect Startup Error. Operations=StartSession Command: Host Reply Timeout (C...	

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

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

The screenshot shows the Peer Management Center Client interface. The main pane displays the 'Summary Edge Caching Runtime Summary (auto-update enabled)' view. The table below represents the data shown in the screenshot:

Name	Overall Status	Job Type	Failed Hosts	Quaran...	Retries	Errors	Warnin...	Open F...	Pendin...	Queue...	Backgr...	Scan Status	Elapse...	Session Stru
FR-2	⊖ Halted.	File Replication	0	0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
User Guides	⊕ Running	File Collaboration	0	0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00...	00:22:25	Size: 0 bytes,
Projects	⊕ Running	File Collaboration	0	0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00...	00:19:39	Size: 0 bytes,
FC-2	⊕ Running	File Collaboration	0	0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00...	00:32:28	Size: 0 bytes,
FS-1	⊕ Running	File Synchronization	0	0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00...	00:28:41	Size: 0 bytes,
FC-3	⊖ Halted.	File Collaboration	0	0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
FS-2	⊖ Halted.	File Synchronization	0	0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
FC-4	⊖ Halted.	File Collaboration	0	0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
FS-3	⊖ Halted.	File Synchronization	0	0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
FC-1	⊕ Running	File Collaboration	0	0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:02...	00:30:09	Size: 630.22 MB,
FR-1	⊕ Running	File Replication	0	0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00...	00:28:49	Size: 0 bytes,

Summary statistics at the bottom of the table: Active Jobs -> Failed Participants: 0 of 4 | Bytes Pending: 0 bytes | Bytes Transferred: 0 bytes | Opens: 0 | Initial Scans Completed: 6 of 11 | Total Size: 630.22 MB | Total Files: 100 | Total Directories: 8

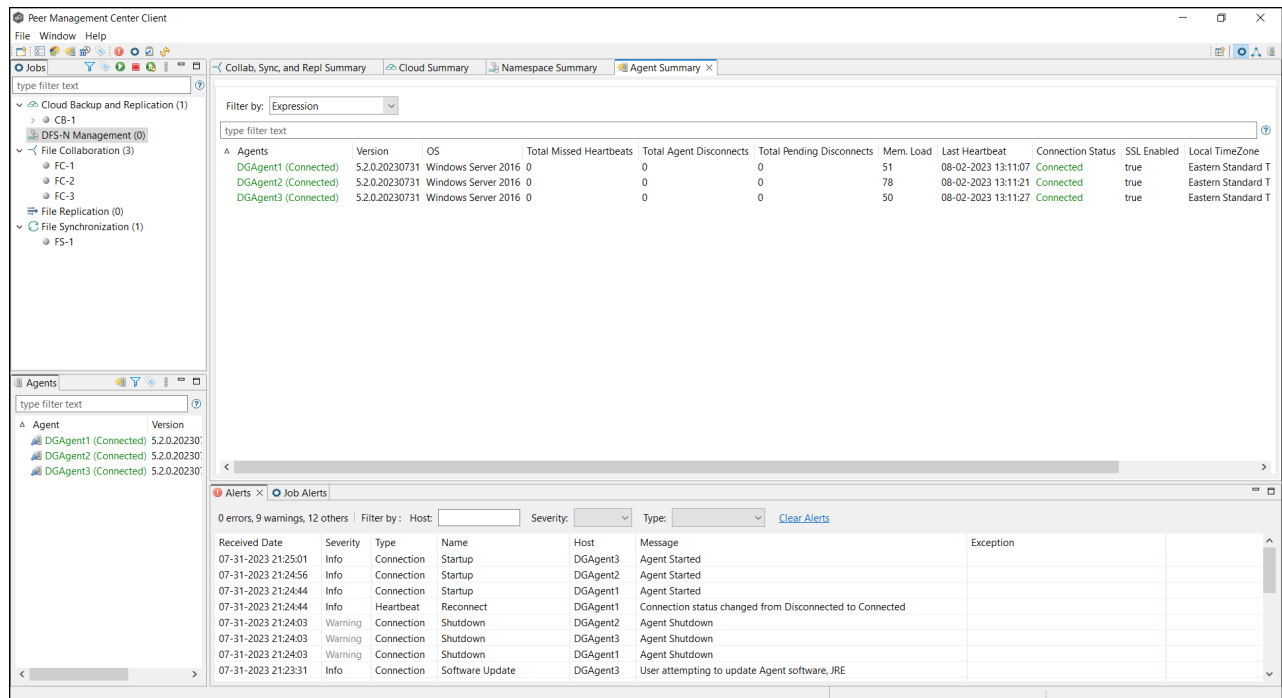
The summary views include:

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

Agent Summary View

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

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



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

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

Cloud Summary View

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

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

This view has four tabs:

- **Volume Summary** – Displays the volumes associated with jobs. The color of the icon next to a volume name quickly indicates the status of the job associated with that volume—a green icon indicates an active job; a gray icon indicates an inactive job, and a red icon indicates a problem with a job.
- **Job Summary** – Displays the status of all Cloud Backup and Replication jobs.
- **Destination Statistics** – Displays the total number of files that have been replicated since the first run of the jobs and other statistics.
- **Tasks** – Displays a high-level view of activities such as snapshots, recovery processes, and background events for all Cloud Backup and Replication jobs.

The screenshot shows the Peer Management Center Client interface. The main window is titled "Volume Summary" and displays a table of job details. The table has the following columns: Volume, Job, Management, Storage Device, Storage Type, Destination Type, Destination, Status, Mode, Scan Status, Next Scan, Scan pending, Real-time pen..., Retry pending, and VSS pending. Two rows are visible in the table:

Volume	Job	Management	Storage Device	Storage Type	Destination Type	Destination	Status	Mode	Scan Status	Next Scan	Scan pending	Real-time pen...	Retry pending	VSS pending
CA	CB-1	DGAgent1	DGAgent1	Windows	Nutanix Objects		Stopped							
SVM01_cifs1	CB-2	DGAgent1	SVM01	NetApp - cDOT	Amazon S3									

Below the table is an Alerts section showing a list of 106 errors, 3 warnings, and 103 others. The Alerts table has the following columns: Received Date, Severity, Type, Name, Host, Message, and Exception. The bottom left shows a list of Agents including DGAgent1 and DGAgent2.

Collab, Sync, and Repl Summary View

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

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

This view has three tabs:

- [Summary](#)
- [Edge Caching](#)
- [Reports](#)

The screenshot shows the Peer Management Center Client interface. The main window displays the 'Collab, Sync, and Repl Summary' view. The interface includes a sidebar with 'Jobs' and 'Agents' sections, and a main table displaying job details.

Name	Overall Status	Job Type	Failed Hosts	Quaran...	Retries	Errors	Warnin...	Open F...	Pendin...	Queue...	Backgr...	Scan Status	Elapse...	Session Stru
FR-2	⊖ Halted	File Replication		0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
User Guides	⊕ Running	File Collaboration		0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00:...	00:25:05	Size: 0 bytes,
Projects	⊕ Running	File Collaboration		0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00:...	00:22:19	Size: 0 bytes,
FC-2	⊕ Running	File Collaboration		0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00:...	00:35:08	Size: 0 bytes,
FS-1	⊕ Running	File Synchronization		0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00:...	00:31:21	Size: 0 bytes,
FC-3	⊖ Halted	File Collaboration		0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
FS-2	⊖ Halted	File Synchronization		0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
FC-4	⊖ Halted	File Collaboration		0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
FR-1	⊖ Halted	File Synchronization		0	0	0	0	0	0 bytes	▼ 0	0	Stopped	00:00:00	Size: 0 bytes,
FC-1	⊕ Running	File Collaboration		0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:02:...	00:32:49	Size: 630.22 f
FR-1	⊕ Running	File Replication		0	0	0	0	0	0 bytes	▼ 0	0	Completed - 00:00:...	00:31:29	Size: 0 bytes,

At the bottom of the table, there is a summary bar: Active Jobs -> Failed Participants: 0 of 4 | Bytes Pending: 0 bytes | Bytes Transferred: 0 bytes | Opens: 0 | Initial Scans Completed: 6 of 11 | Total Size: 630.22 MB | Total Files: 100 | Total Directories: 8

Below the table, there is an 'Alerts' section with a filter by: Host: [] Name: [] Severity: [] Type: [] Clear Alerts

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

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

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

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

Column Descriptions

Key columns in this view are:

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

- **Queued Items** – Presents the number of items in just the Fast and Slow queue (does not include bulk adds).
- **Background Sync.** – Presents the number of initial and full scan items in queue.

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

Actions Menu

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

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

The **Edge Caching** tab presents information about jobs using Edge Caching in a single table.

The screenshot shows the Peer Management Center Client interface. The main window is titled "Peer Management Center Client" and has a menu bar with "File", "Window", and "Help". Below the menu bar are several tabs: "Collab, Sync, and Repl Summary", "Cloud Summary", and "Namespace Summary". The "Collab, Sync, and Repl Summary" tab is active, and within it, the "Edge Caching" sub-tab is selected. The interface is divided into three main sections:

- Left Pane:** A tree view showing the hierarchy of the Peer Management Center. It includes categories like "Cloud Backup and Replication (0)", "DFS-N Management (4)", "File Collaboration (6)", "File Replication (7)", and "File Synchronization (3)".
- Main Table:** A table displaying the status of Edge Caching jobs. The columns are: Agent/Vol..., Last Update, Temporary Storage, Edge Service, Volume Utilization (Total/Local), Cache Utilization (Total/Local), Local Files, Pinned Local Fil., Stubbed Files, Pinned Stubbe..., and Scan State. The table shows several jobs for agents DGWin16A, DGWin16D, DGWin16B, and DGWin16C.
- Job Details:** A section below the main table with columns for Job, Status, Items In Cache, Stubbed Files, and Pinned Local Files.
- Alerts Section:** Located at the bottom, it shows a list of alerts with columns for Received Date, Severity, Type, Name, Host, Message, and Exception. The alerts include information about agent startups, shutdowns, and connection status changes.

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

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

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

The screenshot displays the Peer Management Center Client interface. The main window shows a table of pending activity with columns: Name, File Sync Queue, Real-Time Queue, Queued Bytes, Mods, Adds, Metadata, Scan Queue, Deletes, Renames, and Event Queue. The table lists various jobs such as Development Projects, FC-1, FC-2, FC-3, FR-Engineering Relay 1, FR-2, FS-3, FR-Chip Design Relay 1, FS-1, FR-3, FC-4, FR-Chip Design Relay 2, FC-5, and FS-2. All values in the table are 0.

Below the table is an Alerts section showing 0 errors, 7 warnings, and 15 others. The alerts table has columns: Received Date, Severity, Type, Name, Host, Message, and Exception. The alerts include connection and shutdown events for agents DGWin16A, DGWin16B, DGWin16C, and DGWin16D.

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

Column Descriptions

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

Column	Description
Mods	The number of file update events waiting to be processed for each job.
Adds	The number of file add events waiting to be processed for each job.
Metadata	The number of metadata updates waiting to be processed for each job.
Scan Queue	The number of files waiting to be synchronized as discovered by either the initial or real-time scans.
Deletes	The number of files deleted on a source host that are waiting to be processed.
Renames	The number of files renamed on a source host that are waiting to be processed.
Event Queue	The number of events that are queued up to run for each job.
Slow Expedited Queue	The number of events that are queued in the Slow Expedited Queue for each job.
Fast Expedited Queue	The number of events that are queued in the Fast Expedited Queue for each job.
Scheduled Replication Pending	The number of events that are queued awaiting replication at a scheduled time or interval.
Scheduled Replication Processing	The number of events that are queued awaiting a validation scan to make sure that the source version is correct before being released for replication.
Scheduled Replication	The number of events that are queued awaiting an available replication slot.

Column	Description
n Transfers	

Actions Menu

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

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

Option	Description
Move Totals Row To Bottom	Moves the Totals row to the bottom of the table.

Namespace Summary View

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

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

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

The screenshot shows the Peer Management Center Client interface. The main window displays the **Namespace Summary** view, which is a table with the following columns: **Namespace Path/Folders**, **Management Status**, **State**, **Errors**, **Servers**, and **Total Folders/Targets**. The table is filtered by **Namespace**. The left sidebar shows a tree view of DFS-N Management jobs, and the bottom pane shows a list of alerts.

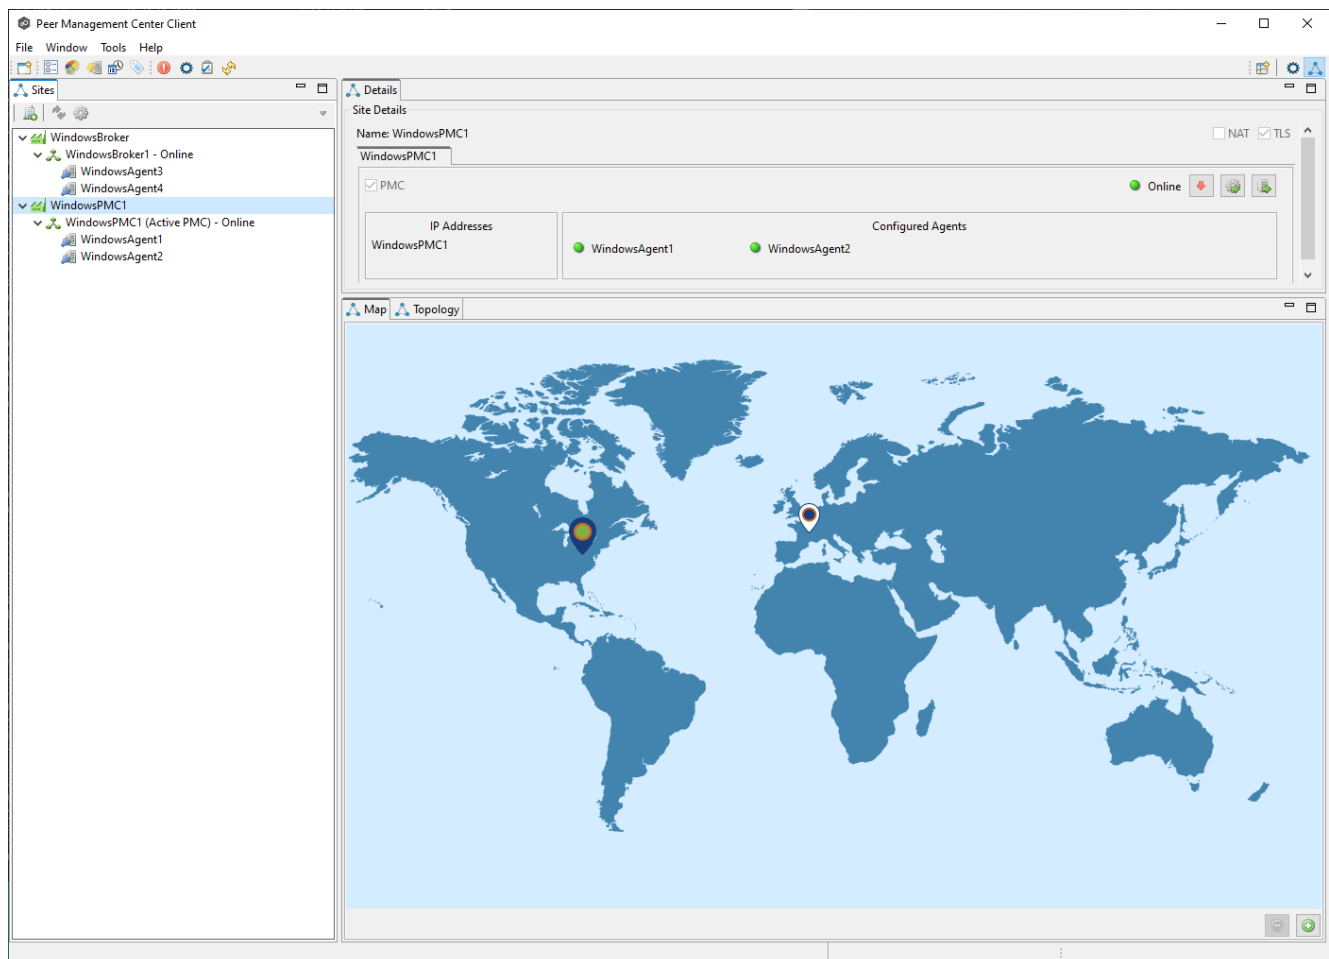
Namespace Path/Folders	Management Status	State	Errors	Servers	Total Folders/Targets
\\PeerTest.local\Development	Monitoring	Online	0	PEERTESTDC1	3 Folders 3 Targets 3 Targets
\\PeerTest.local\DG	Monitoring	Online	0	peertestdc2, peertestdc1	1 Folder 2 Targets
\\PeerTest.local\Documentation	Monitoring	Online	0	peertestdc1, peertestdc2	1 Folder 3 Targets
\\PeerTest.local\Technical Support	Monitoring	Online	0	peertestdc1, peertestdc2	2 Folders 2 Targets 2 Targets

Topology Perspective

The **Topology** perspective in the PMC gives customers the ability to view and configure sites and brokers from a single view in the PMC, allowing customers to create paths for Agents to communicate more directly with one another.

The Topology perspective contains four views:

- Sites
- Details
- Map
- Topology



Views in the Topology Perspective

The views in the Topology perspective are described in the following table.

Views	Description
Sites	Provides an overview of the PeerGFS deployment and status of the various components. It comprises a tree display containing the configured sites, brokers, and connected Agents.
Details	Provides information about brokers that are configured on a selected site from the Sites view. This display shows broker configurations settings, detailed status information, and Agents that are configured to connect to it.
Map	Provides a visual display of sites that have been configured and their geographical location. Sites on the map can be dragged into the desired location or selected to display the site details within the Details view.
Topology	Shows a logical schematic diagram of the PeerGFS deployment, displaying brokers, Agents, and the connections between them.

Menus and Toolbars

The main window of Peer Management Center features four menus and a toolbar:

- [File](#)
- [Window](#)
- [Tools](#)
- [Help](#)
- [Toolbar](#)

File Menu

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

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

Window Menu

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

Command	Description
Refresh	Refreshes all open views and tabs.
Open Perspective	Displays a submenu with the two perspectives : Jobs and Topology.
Reset Perspective	Resets all views currently displayed in the perspective to their default size and layout.

Command	Description
Show Dashboard	Displays the Dashboard in the Jobs perspective. The Dashboard displays metrics and key performance indicators from all running File Collaboration, File Replication, and File Synchronization jobs, and Peer Agents.
Show Agent Summary	Displays the Agent Summary view, which displays a list of all known Peer Agents deployed and their detailed status information, which can be used to assess the health of the environment.
Show Job Summary	<p>Displays a submenu with the following options:</p> <ul style="list-style-type: none"> • Cloud Summary - Displays the summary view for Cloud Backup and Replication jobs. • Collab, Sync, and Repl Summary - Displays the summary view for File Collaboration, File Synchronization, and File Replication jobs. • Namespace Summary - Displays the summary view for DFS-N Management jobs.
Show View	<p>Displays a submenu with the following options:</p> <ul style="list-style-type: none"> • Alerts - Displays the Alerts view, which displays Peer Management Center alerts such as Peer Agent connection status changes. • Job Alerts - Displays the Job Alerts view, which displays alerts such as job restarts. • Task History - Displays the Task History view, which displays the status of tasks such as Daily Cleanup. • Progress - Displays the Progress view, which displays information pertaining to any running background tasks within Peer Management Center.
Open Peer Management Center Web Client	Opens the Peer Management Center Web Client in a web browser.
Open Peer Management	Opens the Peer Global File Service API in a web browser.

Command	Description
Center API	

Tools Menu

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

Command	Description
Open Preferences	Select Open Preferences to access the page where you can configure global settings for Peer Management Center, as well as settings for individual job types.
Assign Tags	Select Assign Tags to view, tag, and assign resources to categories. This feature proves especially useful when managing a large number of resources.
Event Detection Analytics	Select the Analyze PMC Event Logs submenu to immediately trigger event detection analytics. While PeerGFS typically performs event detection analysis nightly, using this option ensures you receive the most up-to-date analytics. The output comprises a series of CSV and XLSX files, with the XLSX files being particularly significant. Among these files, one provides comprehensive statistics regarding PeerGFS's performance, while the other offers insights into the most active files, folders, and extensions for

Command	Description
	content replicated by PeerGFS. The output files are saved in the following location: <PMC Install Dir>\Hub\workspace\analytics\eda.
Compress DB on Restart	Select this option to compress the database when restarting Peer Management Center Service. This action is recommended when the database consumes a significant amount of disk space.

Help Menu

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

Command	Description
User Guide	Opens the PeerGFS User Guide.
Knowledge Base	Opens a web browser that displays the home page of the Peer Knowledge Base.
Support Portal	Opens the Support Portal on the Peer Software website.
Support Tools	Displays a submenu with the following options: <ul style="list-style-type: none"> • Retrieve PMC/Agent Logs • Retrieve Broker Statistics • Generate Thread Dump • Generate Memory Dump See Support Tools for details.
Download Peer Agent	Offers several options for downloading Peer Agent: <ul style="list-style-type: none"> • A Windows installer that is compatible with this version of the PMC. • A Linux installer that is compatible with this version of the PMC.

Command	Description
	<ul style="list-style-type: none">• An Ubuntu-based virtual appliance image for VMware ESX that includes a Linux-based Peer Agent.• An Ubuntu-based virtual appliance image for Nutanix AHV that includes a Linux-based Peer Agent.• An Ubuntu-based virtual appliance image for Microsoft Hyper-V that includes a Linux-based Peer Agent.
Download Broker	<p>Offers several options for downloading Peer Broker:</p> <ul style="list-style-type: none">• A Windows installer that is compatible with this version of the PMC.• A Linux installer that is compatible with this version of the PMC.• An Ubuntu-based virtual appliance image for VMware ESX that includes a Linux-based Peer Broker.• An Ubuntu-based virtual appliance image for Nutanix AHV that includes a Linux-based Peer Broker.• An Ubuntu-based virtual appliance image for Microsoft Hyper-V that includes a Linux-based Peer Broker. <p>See Getting Started with the Broker Virtual Appliance for more information on the Peer Broker Virtual Appliance.</p>
Download PeerIQ	<p>Offers several options for downloading PeerIQ:</p> <ul style="list-style-type: none">• An Ubuntu-based virtual appliance image for VMware ESX that includes a Linux-based PeerIQ.• An Ubuntu-based virtual appliance image for Nutanix AHV that includes a Linux-based PeerIQ.• An Ubuntu-based virtual appliance image for Microsoft Hyper-V that includes a Linux-based PeerIQ. <p>See Getting Started with the PeerIQ Virtual Appliance for more information on the Peer Agent Virtual Appliance.</p>

Command	Description
Check for Updates	(Rich client only) Checks for updates to Peer Management Center. Minor releases can be automatically downloaded and installed. Major releases require a new license key and must be requested from Peer Support .
Licenses	Displays the Licensing page in Preferences .
About Peer Management Center	Displays version information and installation details.


Support Tools

The following options are available from the Support Tools submenu:

Command	Description
Retrieve PMC/Agent Logs	Collects and retrieves all useful log files for specified Peer Agents, Peer Management Center, and all jobs. This information is assembled into a single encrypted zip file that can optionally be uploaded to Peer Support . The collection and retrieval of the log and support files is performed in the background, which might take a while, depending on content size and network speed. Upon completion, you are notified and can view the zip file.
Retrieve Broker Statistics	Displays detailed statistical information about all messaging that has transpired for all connections (Peer Agents and Peer Management Center) to Peer Management Broker . Peer Support can use these statistics to aid in diagnosing problems.
Generate Thread Dump	Displays options to generate a thread dump of the running PMC, Broker, or Topology Services. These can be used by Peer Support to debug certain issues.
Generate Memory Dump	Displays options to generate a memory dump of the running PMC or Topology Services. These can be used by Peer Support to debug certain issues.

Toolbar

Beneath the main window menu bar, you'll find a toolbar containing frequently used tools and actions represented by icons. These shortcuts allow for quick access to commonly performed tasks without navigating through menus.



Use the toolbar in the Peer Management Center main window to quickly launch these actions:

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

Tables

Tables are used throughout the Peer Management Center interface to present information effectively. You can efficiently organize and locate relevant data within tables by sorting and filtering. To sort a table, simply click on a column header. For instance, within the **Summary** tab of the **Collab, Sync, and Repl** view, you can sort columns to arrange data according to your preference.

Name	Overall Status	Job Type	Failed ...	Quaran...	Retries	Errors	Warnin...	Open F...	Pending Bytes	Queue...	Backgr...	Scan Status	Elapse...	Session Structure	Synchronization Pr...
Development Proj...	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FR-1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FC-3	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FC-1	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
Source to Staging	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
Staging to Relay T...	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FC-2	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FR-2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FS-3	Stopped	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FR-Chip Design S...	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
Job 1	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
Job 2	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FR-Chip Design R...	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FS-1	Stopped	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FR-3	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2
FR-Chip Design S...	Stopped	File Replication	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files...	2

Active Jobs -> Failed Participants: 0 of 0 | Bytes Pending: 0 bytes | Bytes Transferred: 0 bytes | Opens: 0 | Initial Scans Completed: 0 of 0 | Total Size: 0 bytes | Total Files: 0 | Total Directories: 0

When you right-click within a table, a **context menu** appears, offering additional operations for managing the table. This menu enables you to customize the display of columns, such as choosing which ones to hide or show. A particularly handy feature found in many context menus is the ability to copy detailed information for one or more rows simultaneously. This copied information can then be easily pasted into any text editor for further use.

The screenshot displays the Peer Global File Service interface. At the top, there are tabs for 'Collab, Sync, and Repl Summary', 'Cloud Summary', 'Agent Summary', and 'Development Projects'. Below these, there's a 'Runtime Summary (auto-update disabled)' section with a filter by 'Expression' and an 'Actions' dropdown. The main table lists various jobs with columns for Name, Overall Status, Job Type, Failed, Quarantines, Retries, Errors, Warnings, Open Files, Pending Bytes, Queue, Backgr..., Scan Status, Elapse..., Session Structure, and Synchronization Pr... A context menu is open over the 'Development Proj...' row, showing options like 'File Collaboration', 'Refresh View', 'Auto-Update', 'Show Details', 'Copy Details', 'Open', 'Open in New Tab', 'Start', 'Stop', 'Edit Job(s)', 'Purge All Quarantines', 'Clear Alerts', 'Clean Up Trash Bin', and 'Run Event Detection Analytics...'. Below the main table, there's an 'Alerts' section with a filter by 'Type' and a 'Clear Alerts' button. The 'Alerts' table shows a list of alerts with columns for Received Date, Severity, Type, and Exception.

Name	Overall Status	Job Type	Failed ...	Quaran...	Retries	Errors	Warnin...	Open F...	Pending Bytes	Queue...	Backgr...	Scan Status	Elapse...	Session Structure	Synchronization Pr...
Development Proj...	Stopped	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FR-1	Stopped	Refresh View	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FC-3	Stopped	Auto-Update	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FC-1	Stopped	Auto-Update	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
Source to Staging	Stopped	Show Details	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
Staging to Relay T...	Stopped	Copy Details	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FR-2	Stopped	Open	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FR-2	Stopped	Open in New Tab	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FS-3	Stopped	Start	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FR-Chip Design S...	Stopped	Stop	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
Job 1	Stopped	Edit Job(s)	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
Job 2	Stopped	Edit Job(s)	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FR-Chip Design R...	Stopped	Purge All Quarantines	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FS-1	Stopped	Clear Alerts	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FR-3	Stopped	Clean Up Trash Bin	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2
FR-Chip Design S...	Stopped	Run Event Detection Analytics...	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes, Files:...	2

Received Date	Severity	Type	Exception
03-08-2024 14:09:03	Info	Conn	
03-08-2024 14:09:02	Info	Heart	
03-08-2024 14:08:28	Info	Conn	
03-08-2024 14:08:28	Info	Heart	
03-08-2024 14:08:20	Info	Conn	
03-08-2024 14:07:59	Info	Conn	
03-08-2024 14:06:55	Warning	Conn	
03-08-2024 14:06:52	Warning	Conn	
03-08-2024 14:06:52	Warning	Conn	
03-08-2024 14:06:52	Warning	Conn	
03-08-2024 14:06:25	Info	Connection	Shutdown
03-08-2024 14:06:25	Info	Connection	Startup

Double-clicking on any row in most tables triggers a dialog or opens another view that presents detailed information specific to that row. For instance, double-clicking on a row within the Job Alerts table unveils comprehensive details regarding the corresponding alert.

The screenshot displays the Peer Global File Service interface. At the top, there are tabs for 'Alerts', 'Job Alerts', 'Task History', and 'Progress View'. Below these, there's an 'Alerts' section with a filter by 'Host', 'Severity', and 'Type', and a 'Clear Alerts' button. The 'Alerts' table shows a list of alerts with columns for Received Date, Severity, Type, and Exception. A detailed view of an alert is shown, with a 'PMCA Alert Details' dialog box open. The dialog box shows the alert's received date, severity, type, and message.

Received Date	Severity	Type	Exception
03-08-2024 14:09:03	Info	Co	
03-08-2024 14:09:02	Info	He	
03-08-2024 14:08:28	Info	Co	
03-08-2024 14:08:28	Info	He	
03-08-2024 14:08:20	Info	Co	
03-08-2024 14:07:59	Info	Co	
03-08-2024 14:06:55	Warning	Co	
03-08-2024 14:06:52	Warning	Co	
03-08-2024 14:06:52	Warning	Co	
03-08-2024 14:06:52	Warning	Co	
03-08-2024 14:06:25	Info	Co	
03-08-2024 14:06:24	Info	Heartbeat	Reconnect
03-08-2024 14:06:17	Info	Connection	Startup
03-08-2024 14:06:16	Info	Heartbeat	Reconnect
03-08-2024 14:05:49	Warning	Connection	Shutdown
03-08-2024 14:05:46	Warning	Connection	Shutdown
03-08-2024 14:05:40	Info	Connection	Software Update
03-08-2024 14:05:40	Info	Connection	Software Update
03-08-2024 14:05:40	Info	Connection	Software Update
03-08-2024 14:05:40	Info	Connection	Software Update
03-08-2024 14:05:40	Info	Connection	Software Update
03-08-2024 14:05:25	Info	Connection	Software Update
03-08-2024 14:05:25	Info	Connection	Software Update
03-08-2024 13:34:31	Warning	Agent	Connection

PMCA Alert Details

Received at: 03-08-2024 13:34:31

Severity: Warning

Category: Agent

Host Name: DGWin16A

Locally Generated at: 03-08-2024 13:34:31

Name: Connection

Message: Connection to Broker DGWin16A was lost for 11 minutes 32 seconds

Click outside of popup to close

Basic Concepts

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

- [Email Alerts](#)
- [File and Folder Filters](#)
- [List Filters](#)
- [Logging and Alerts](#)
- [SNMP Notifications](#)
- [Tags](#)
- [Web Client Users](#)

Email Alerts

Overview

An email alert notifies recipients when a certain type of event occurs, for example, file quarantined, session aborted, host failure, system alert. When an email alert is applied to a job, an alert is sent to all listed recipients whenever a selected event type is triggered by the job.

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

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

See [Email Configuration](#) for configuring an SMTP email connection. This must be configured before email alerts can be sent.

Managing Email Alerts

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

To manage email alerts:

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

The **Preferences** dialog appears.

2. Select the job type from the navigation tree and expand it.
3. Select **Email Alerts** from the navigation tree.

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

File and Folder Filters

Overview

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

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

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

Types of File Filters

There are three types of file filters:

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

For more information, see [Creating and Applying File Filters](#).

Creating and Applying File Filters

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

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

For more information about creating a file filter, see:

- [Creating File Filters for a Cloud Backup and Replication Job](#)
- [Creating File Filters for File Collaboration, File Locking File Replication, and File Synchronization Jobs](#)

Predefined File Filters

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

File and Folder Filters					
Name	Type	Exclusions	Inclusions	Date Filter	Size Filter
Default	General	~*.*; *.BAK; *.BCK; *.WBK; *.ASD...	None Selected	Include all dates	None
File Collaboration Sync Only	Synchronization Only	None Selected	*.LOG; *.EXE; *.DLL; *.OTF,...	Include all dates	None
Invalid Characters	General	<<.[]\$>>	None Selected	Include all dates	None
Locking Only	Locking Only	None Selected	*	Include all dates	None
MacOS Exclusions	General	*\ _MACOSX; *\ TemporaryItem...	None Selected	Include all dates	None
PEER_Autodesk AutoCAD [General]	General	*.BAK; *.DWL	None Selected	Include all dates	None
Synchronizing Only	Synchronization Only	None Selected	*	Include all dates	None
User Profile Exclusions	General	*\AppData\Roaming\Microsoft\...	None Selected	Include all dates	None

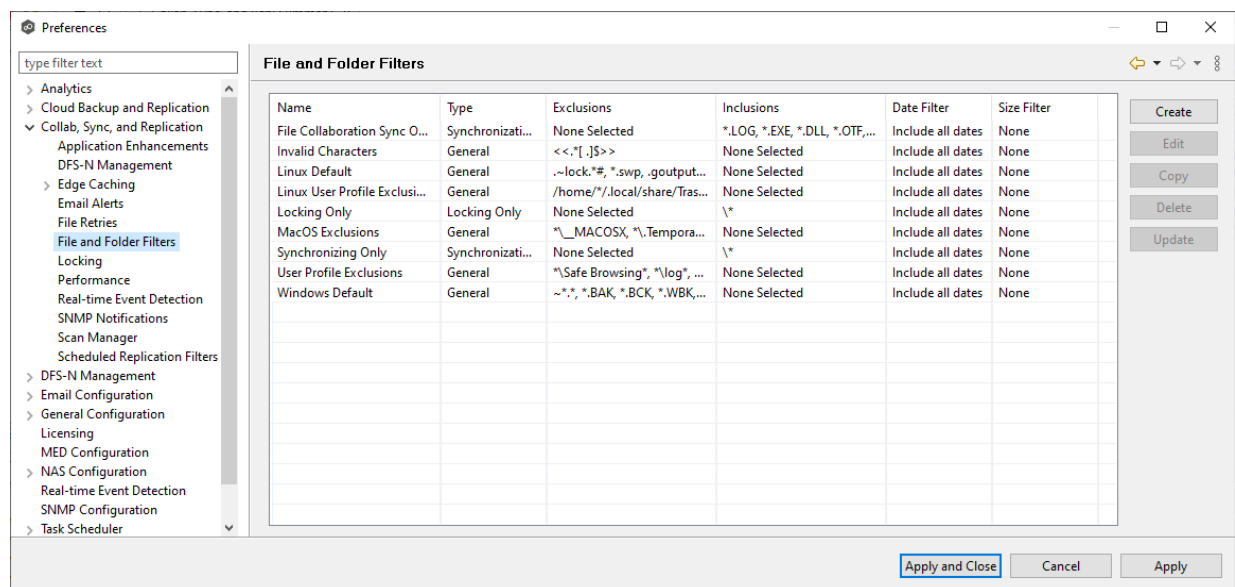
Two of the predefined filters, **Default** and **Invalid Characters**, are applied to all jobs by default. However, you can deselect a predefined filter for a specific job. Only the **Default** filter can be modified; none of the predefined file filters can be deleted.

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

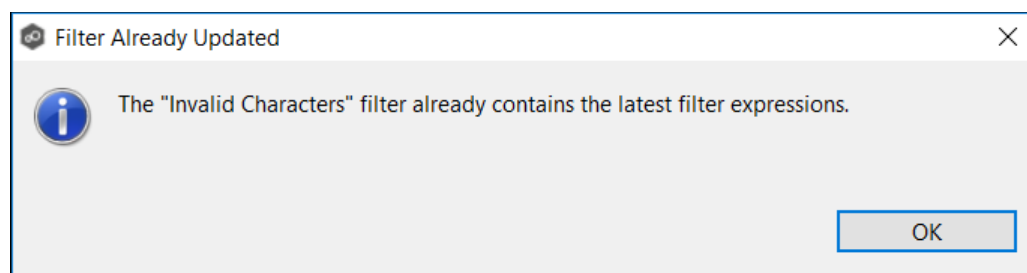
To upgrade a predefined filter:

1. Select **Preferences** from the **Window** menu.
2. Expand **Cloud Backup and Replication** or **Collab, Sync, and Repl Summary** in the navigation tree, and then select **File and Folder Filters**.

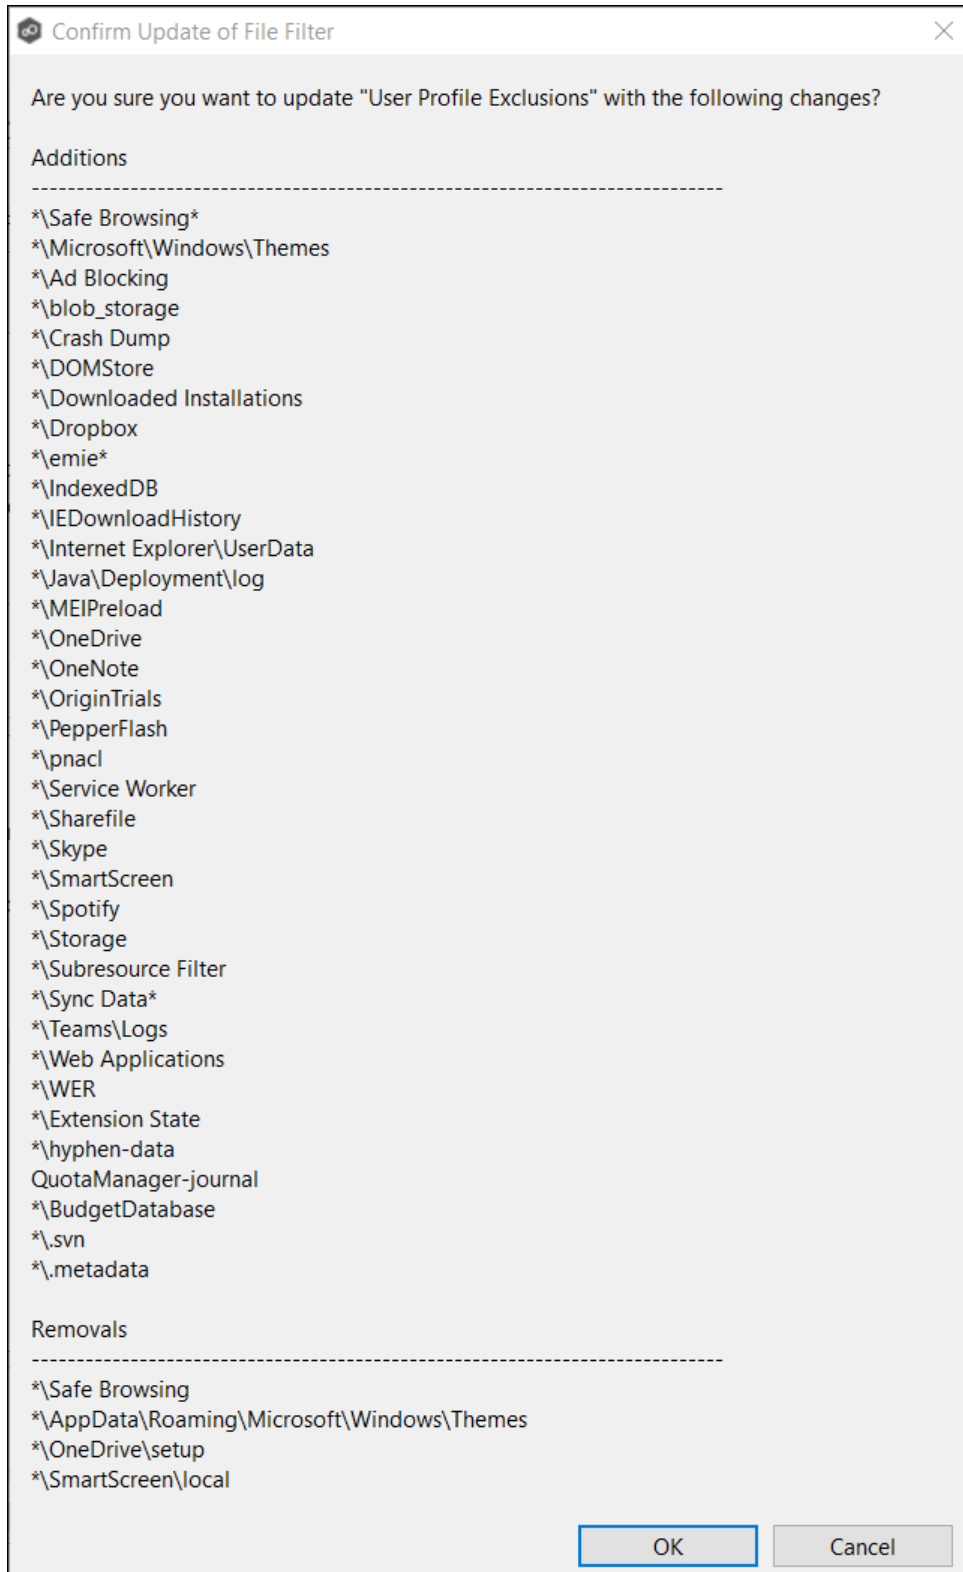
Existing file filters are listed in the **File and Folder Filters** table.



3. Select the filter to upgrade, and then click **Upgrade**.
4. If no changes are available, click **OK** to close the message that appears.



If an updated filter definition is available, a confirmation message lists the changes to the filter definition; click **OK** to install the updated definition.



Defining Filter Patterns

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

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

Create File Filter

Name:

Filter Type: **General** ▾

Auto Excluded

[View file types that are automatically excluded](#)

Excluded Patterns

Add Edit Delete

Included Patterns

Add Edit Delete

Included Last Modified Dates

Include all dates ▾

days

Excluded File Sizes

None ▾

bytes

OK Cancel

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

*	Matches zero or more characters of any value
?	Matches one character of any value

The following examples show the use of a wildcard:

- *.ext** Filter files that end with the **.ext** extension
- ext*** Filter files that begin with the string **ext**
- ext** Filter files that contain the string **ext**

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

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

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

Typically, AutoCAD files have the following extensions:

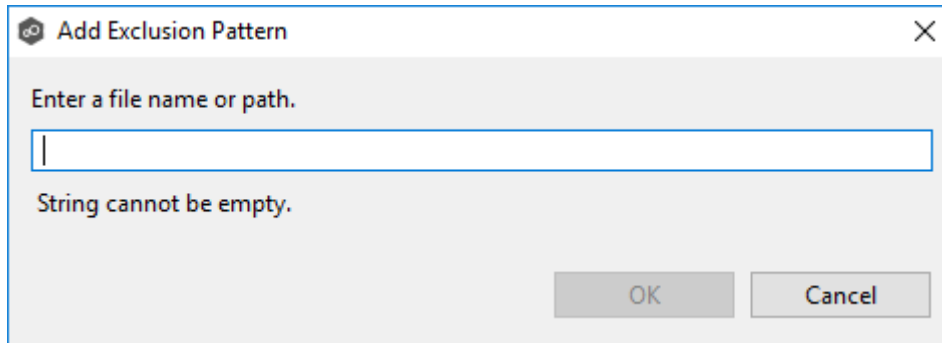
- .AC\$
- .SV\$
- .DWL

.BAK

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

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

The **Add Exclusion Pattern** dialog appears.



2. Enter ***.AC\$**, and then click **OK**.
3. Repeat Step 2 to add ***.BAK**, ***.DWL*** and ***.SV\$**.

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

Create File Filter

Name:

Filter Type:

Auto Excluded
[View file types that are automatically excluded](#)

Excluded Patterns

- *.AC\$
- *.BAK
- *.DWL*
- *.SV\$

Included Patterns

Included Last Modified Dates

days

Excluded File Sizes

bytes

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

Using Complex Regular Expressions in Filter Patterns

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

```
<<regEx>>
```

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

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

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

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

The following files and folders **MATCH** the above expression:

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

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

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

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

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

- Peer Management Center supports filtering on a file's last modified date but does not support filtering on a folder's last modified date.
- If you have a folder hierarchy that contains files that are all being filtered based on the last modified date, then all folders will still be created during the initial scan process on all hosts.

- If a file is excluded from collaboration based on its last modified date, then the initial scanning process will not synchronize the file even if the file's last modified time and size do not match, or the file does not exist on all hosts. However, the file will be synchronized, if and when the file is modified in the future, and if a user deletes or renames the file on any host, the file will be deleted or renamed from all other hosts where the file exists.
- A file filter cannot combine filtering on last modified date with inclusion or exclusion patterns or [file size](#). The last modified date is the sole criteria used to identify matching files.

Options for Included Last Modified Date Filter

The screenshot shows a 'Create File Filter' dialog box with the following sections:

- Name:** An empty text input field.
- Filter Type:** A dropdown menu set to 'General'.
- Auto Excluded:** A section with a link: [View file types that are automatically excluded](#).
- Excluded Patterns:** A large empty text area with 'Add', 'Edit', and 'Delete' buttons below it.
- Included Patterns:** A large empty text area with 'Add', 'Edit', and 'Delete' buttons below it.
- Included Last Modified Dates:** A dropdown menu set to 'Include older than' with a value of '0' days. This section is highlighted with a red box.
- Excluded File Sizes:** A dropdown menu set to 'None' with a value of '0' bytes.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom right.

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

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

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

Options for Excluded File Sizes

The image shows a 'Create File Filter' dialog box with the following sections:

- Name:** An empty text input field.
- Filter Type:** A dropdown menu set to 'General'.
- Auto Excluded:** A section with a link: [View file types that are automatically excluded](#).
- Excluded Patterns:** An empty list box with 'Add', 'Edit', and 'Delete' buttons below it.
- Included Patterns:** An empty list box with 'Add', 'Edit', and 'Delete' buttons below it.
- Included Last Modified Dates:** A dropdown menu set to 'Include all dates' and a text input field containing '0' followed by 'days'.
- Excluded File Sizes:** A dropdown menu set to 'None' and a text input field containing '0' followed by 'bytes'. This section is highlighted with a red border.

At the bottom of the dialog are 'OK' and 'Cancel' buttons.

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

Filtering Folders

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

Reduce the Number of Jobs Using Folder Filtering

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

Folder Filter Syntax

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

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

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

\Folder*\Folder

Example of a Simple Folder Filter

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

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

After consolidation:

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

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

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

Please note the following regarding regular expressions:

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

Additional Examples of Folder Filters

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

File Filter Usage Notes

Conflicting Patterns

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

Rename Operations

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

Folder Deletions

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

List Filters

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

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

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

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

Basic Filter Expressions

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

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

For information about creating more complex filter expressions using operators and labels, see [Creating Complex Filter Expressions](#).

Predefined List Filters

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

Creating Complex Filter Expressions

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

Using Operators

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

Using Labels

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

<label>:<search string>.

List of possible labels include:

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

Examples

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

Sales

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

cat:Location AND tag:"North America"

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

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

Saving and Managing List Filters

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

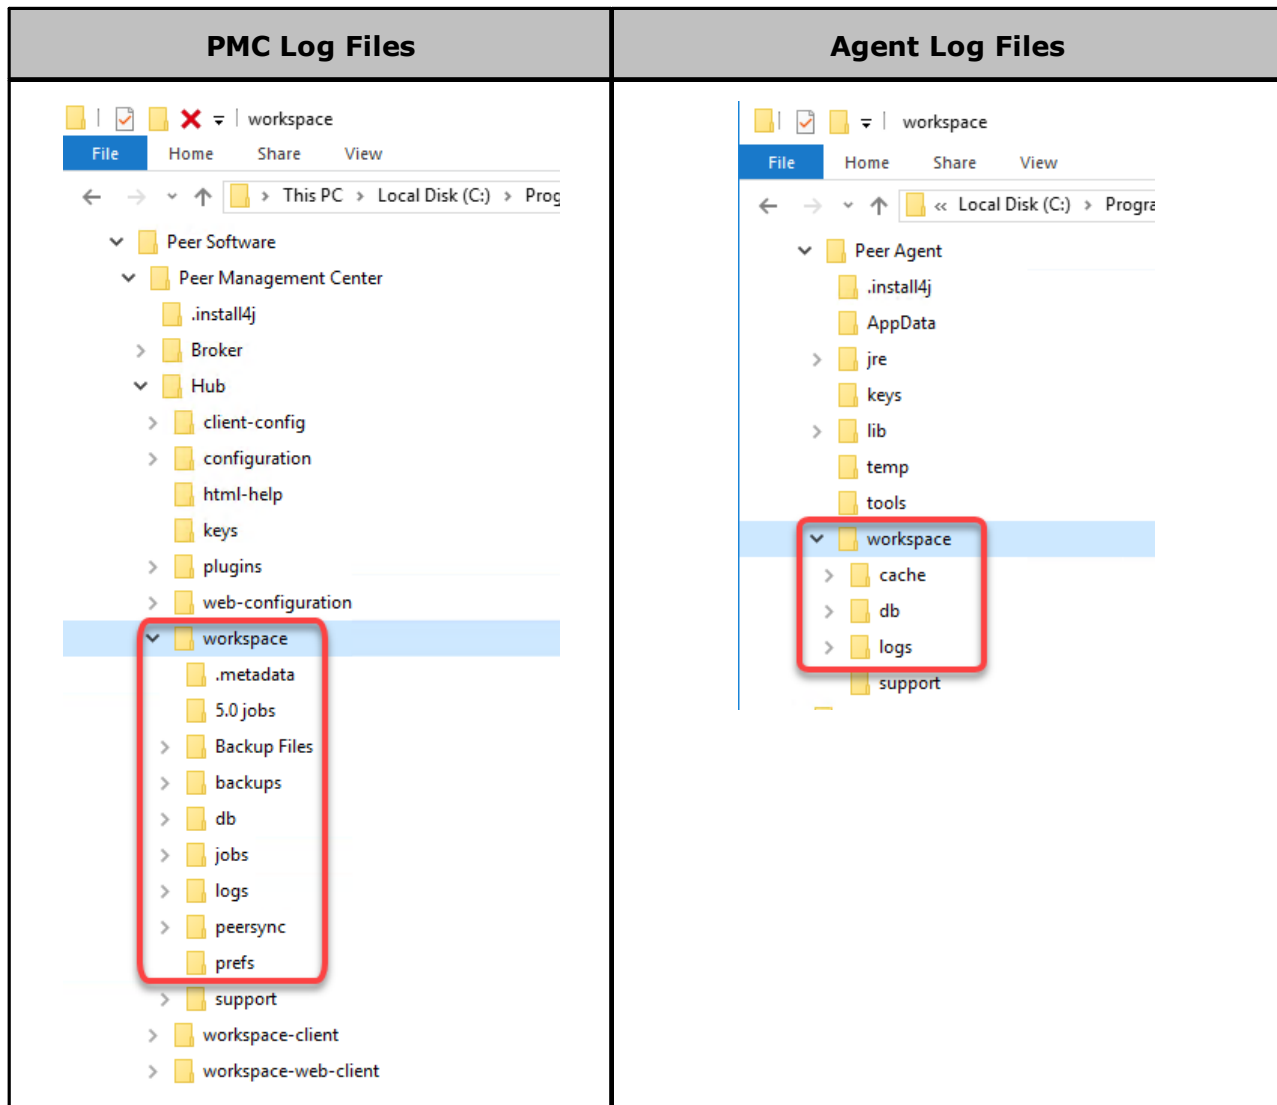
To remove a list filter and show all items in the list, click the pencil icon to the right of the filter expression.

Logging and Alerts

Logging

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

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



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

Job Alerts

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

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

Retrieving Log Files

To retrieve log files:

1. Open Peer Management Center.
2. From the **Help** menu, select **Support Tools**, and then select **Retrieve PMC/Agent Logs**.

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

Retrieve PMC/Agent Log Files

Log Collection Options

Include logs newer than days

Run Event Detection Analytics before log file collection

Include detailed log files

Collect system event logs

Include topology statistics

Agent Log Options

Exclude all Agent log files

Include all connected Agent log files

Include log files from the following connected Agents:

Agent
<input type="checkbox"/> DGWin16A
<input type="checkbox"/> DGWin16B
<input type="checkbox"/> DGWin16C
<input type="checkbox"/> DGWin16D

[Select All](#) [Clear Selected](#)

Encryption and Support Options

Encrypt log files

Upload log files and telemetry to Peer Software Support

Log retrieval can take a while based on network speed and log file sizes.
You will be notified when this operation completes.

Are you sure you want to proceed with this operation?

3. Select log collection options:

Option	Description
Include logs newer than X days	Use this option to restrict the logs retrieved to a certain time period.
Run Event Detection Analytics before log file collection	Select this option to run event detections analytics immediately before the log files are collected. PeerGFS can perform event detection analysis every night; however, this option ensures that the log bundle contains the most up-to-date analytics.
Include detailed log files	This option is selected by default. If selected, log collection includes all Peer-generated log files (for example, event log files, activity log files, and Agent output logs if Agents are selected in the Agent Log Options section). Detailed log collection enhances Peer Support's ability to troubleshoot using log files. However, if you want to reduce the size of log uploads, deselect this option. Only basic logs and statistics will be collected.
Collect system event logs	Select this option to retrieve Windows or Linux operating system event logs from the PMC and Agents.
Include topology statistics	Select this option to include topology statistics.

4. Select Agent log options:

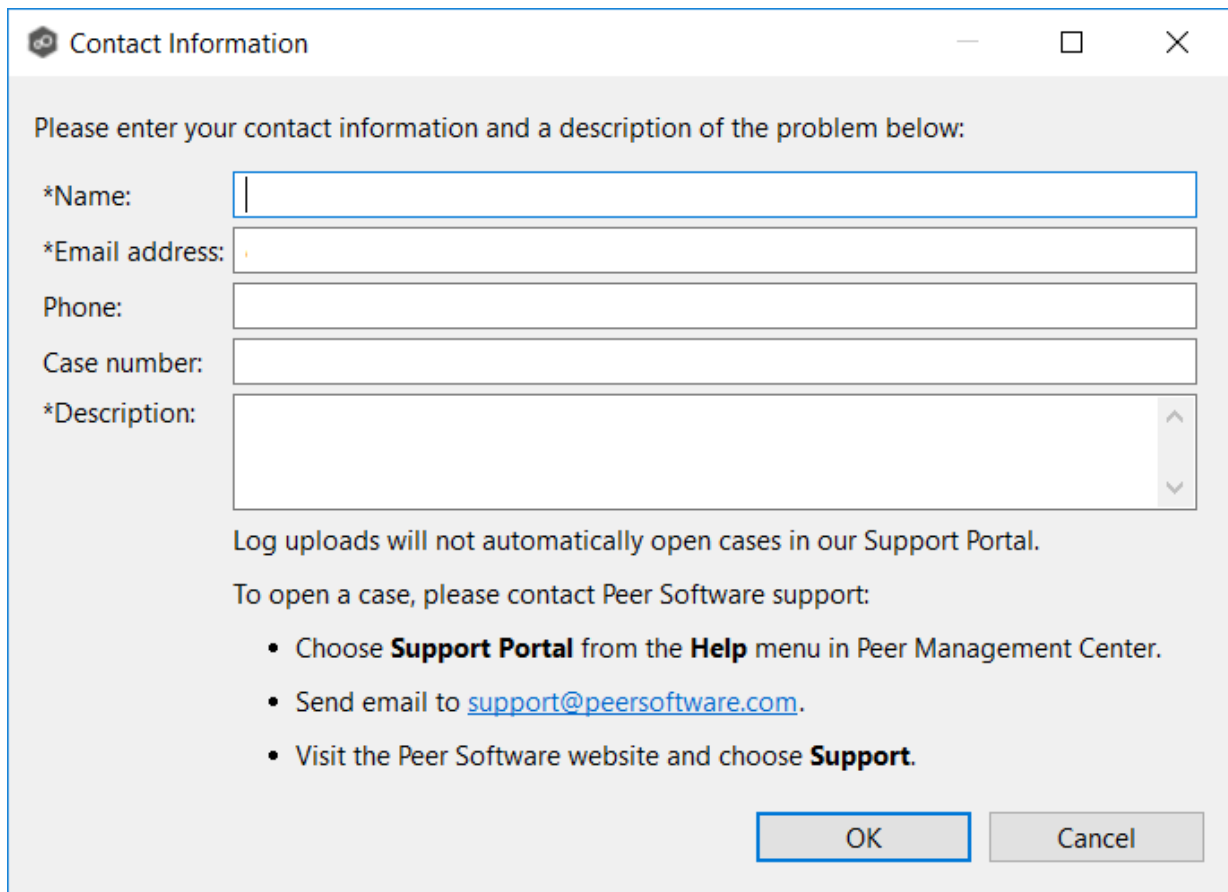
Option	Description
Exclude all Agent log files	Select this option if you do not want to retrieve log files for any Agent.
Include all connected Agent log files	Select this option if you want to retrieve log files for all connected Agents.
Include log files from the following connected Agents:	Select this option if you want to retrieve log files for selected connected Agents.

5. Select encryption and support options:

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

6. Enter your contact information and a description of the problem.

All fields are required. This information will be sent to [Peer Support](#).



Contact Information

Please enter your contact information and a description of the problem below:

*Name:

*Email address:

Phone:

Case number:

*Description:

Log uploads will not automatically open cases in our Support Portal.

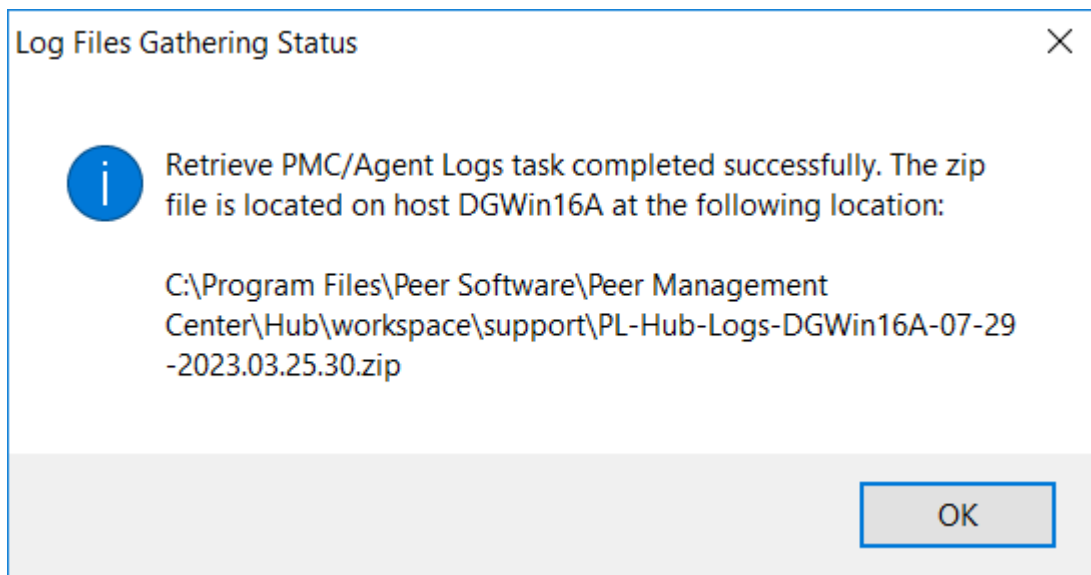
To open a case, please contact Peer Software support:

- Choose **Support Portal** from the **Help** menu in Peer Management Center.
- Send email to support@peersoftware.com.
- Visit the Peer Software website and choose **Support**.


OK Cancel

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

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



Log Files Gathering Status

 Retrieve PMC/Agent Logs task completed successfully. The zip file is located on host DGWin16A at the following location:

C:\Program Files\Peer Software\Peer Management Center\Hub\workspace\support\PL-Hub-Logs-DGWin16A-07-29-2023.03.25.30.zip

OK

8. Click **OK**.

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

SNMP Notifications

Overview

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

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

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

Managing SMNP Notifications

To manage SMNP notifications:

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

The **Preferences** dialog appears.

2. Select the job type from the navigation tree and expand it.
3. Select **SMNP Notifications** from the navigation tree.

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

Tags

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

You can assign tags to:

- Jobs
- Resources
- Roles
- Agents

You can also assign resources to tags. See [Using Tags to Filter Resources](#).

Creating Tags and Categories

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

Assigning Tags

You can:

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

Assigning Tags to Jobs

- During job creation - You can assign tags during the creation of a job from the [Tags](#) page of the job creation wizard.

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

Assigning Tags to Resources

To assign tags to one or more resources:

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

The **Assign/Unassign resources** dialog appears.

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

Tip: To select multiple resources, press the Shift key on the keyboard when selecting resources.

6. Click **Save**.
7. Repeat the preceding steps for all the tags that need to be assigned to one or more resources.

Assigning Tags to Roles

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

To assign tags to user roles:

1. Create tags and categories as outlined in Step 1 above.
2. Assign tags to one or more jobs as outlined in Step 2 above.
3. Go to [User Management](#) in the [Preferences](#) page.
4. Select the desired role to which you wish to assign specific job tags.
5. Click the **Edit** button.

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

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

Assigning Resources to One or More Tags

To assign resources to one or more tags:

1. Click the **Assign Tags** button from a summary view, [Jobs view](#), or [Agent Summary view](#) toolbar.
2. In the **Assign Tags** dialog, click the **Resources** radio button.
3. On the left-hand side, click inside the **Resource Name Filter** or **Type Filter** fields and press the CTRL + Space keys on the keyboard to list all possible filters and predefined labels, which can be selected to refine your search quickly.
4. Select the resource that needs to be assigned to one or more tags.
5. Click the **Edit** button to the right.
6. From the **Unassigned Tags** table on the left side, select the tags that need to be assigned the selected Resource, and then click the right-arrow button (Add one) to move it to the table on the right side (hold down the Shift key on the keyboard when selecting tags to select more than one).
7. Click the **Save** button to commit your changes, and then close the dialog.
8. Repeat the preceding steps for all the resources that need to be assigned to one or more tags.

Using Tags to Filter Resources

You can use tags to filter resources:

- Filter jobs
- Filter agents

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

Filter Jobs

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

Example:

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

tag:"North America" AND cat:Location

Filter Agents

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

Web Client Users

Peer Management Center offers two interfaces:

- **A rich client interface:** Rich client users have unrestricted access to all Peer Management Center functionality. This interface is accessible only on the server where Peer Management Center is installed.
- **A web client interface:** Web client users' access to Peer Management Center functionality is governed by their assigned [role](#). The level of access and permissions granted to web client users is determined by their role. For more information about roles, see [Managing Roles for Web Client Users](#).

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

- Internal users - Users whose access to the web client is authenticated through the internal PMC database.
- Active Directory (AD) users and groups - Users whose access to the web client is authenticated through Active Directory.

For information about managing web client user, see [User Management](#).

Roles for Web Client Users

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

Web client users can have a [predefined \(base\) role](#) or a [custom role](#).

In contrast, a user who accesses Peer Management Center through the rich client does not have a role. All Peer Management Center functionality is accessible to a rich client user.

For more information about web roles, see [Managing Roles for Web Client Users](#).

A **base role** is a predefined role set by the PMC. There are three base roles, each with a predefined set of permissions:

- **Administrator** - This role has complete access to all the functionality of Peer Management Center.
- **Power User** - This role has view-only access to jobs and the **Agent Summary** view. This role can start and stop jobs but cannot create, edit, or delete jobs, access settings in Preferences, or assign tags.
- **Help Desk** - This role has view-only access to jobs. Specifically, Help Desk users are limited to view-only access to the following:
 - The [Jobs view](#)
 - The [runtime views](#)
 - The **Summary** and **Session** tabs of an individual job view.

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

Base roles cannot be modified or deleted, with one exception: tags can be assigned to base roles. For a list of the permissions associated with base roles, see [Base Role Permissions](#).

Base Role Permissions

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

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

Functionality	Administr ator	Power User	Help Desk
Job View	Edit	Edit	View-only
Log Dump Action	Edit	Edit	
Memory Dump Action	Edit	Edit	
New Job Action	Edit		
Participant View	Edit	Edit	
Permission Mode	Edit	Edit	
PMC Alert View	Edit	Edit	
PMC Download Agent	Edit	Edit	
PMC Refresh Perspective	Edit	Edit	
PMC View Progress	Edit	Edit	
Preferences	Edit		
Runtime Summary Interface	Edit	Edit	View-only
Session View	Edit	Edit	View-only
Status Agent Tree View	Edit	View-only	
Tag Resources Dialog	Edit		
Thread Dump Action	Edit	Edit	

PeerSync Management Job Permissions

The following table outlines the permissions for PeerSync Management jobs.

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

A **custom role** allows you to customize and fine-tune the access that a user has to Peer Management Center resources. This is useful if you have multiple types or levels of users that need different types of access. For example, if you have multiple tiers of help desk staff, creating custom roles based on the standard Help Desk role allows you to provide them with varying levels of access to Peer Management Center.

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

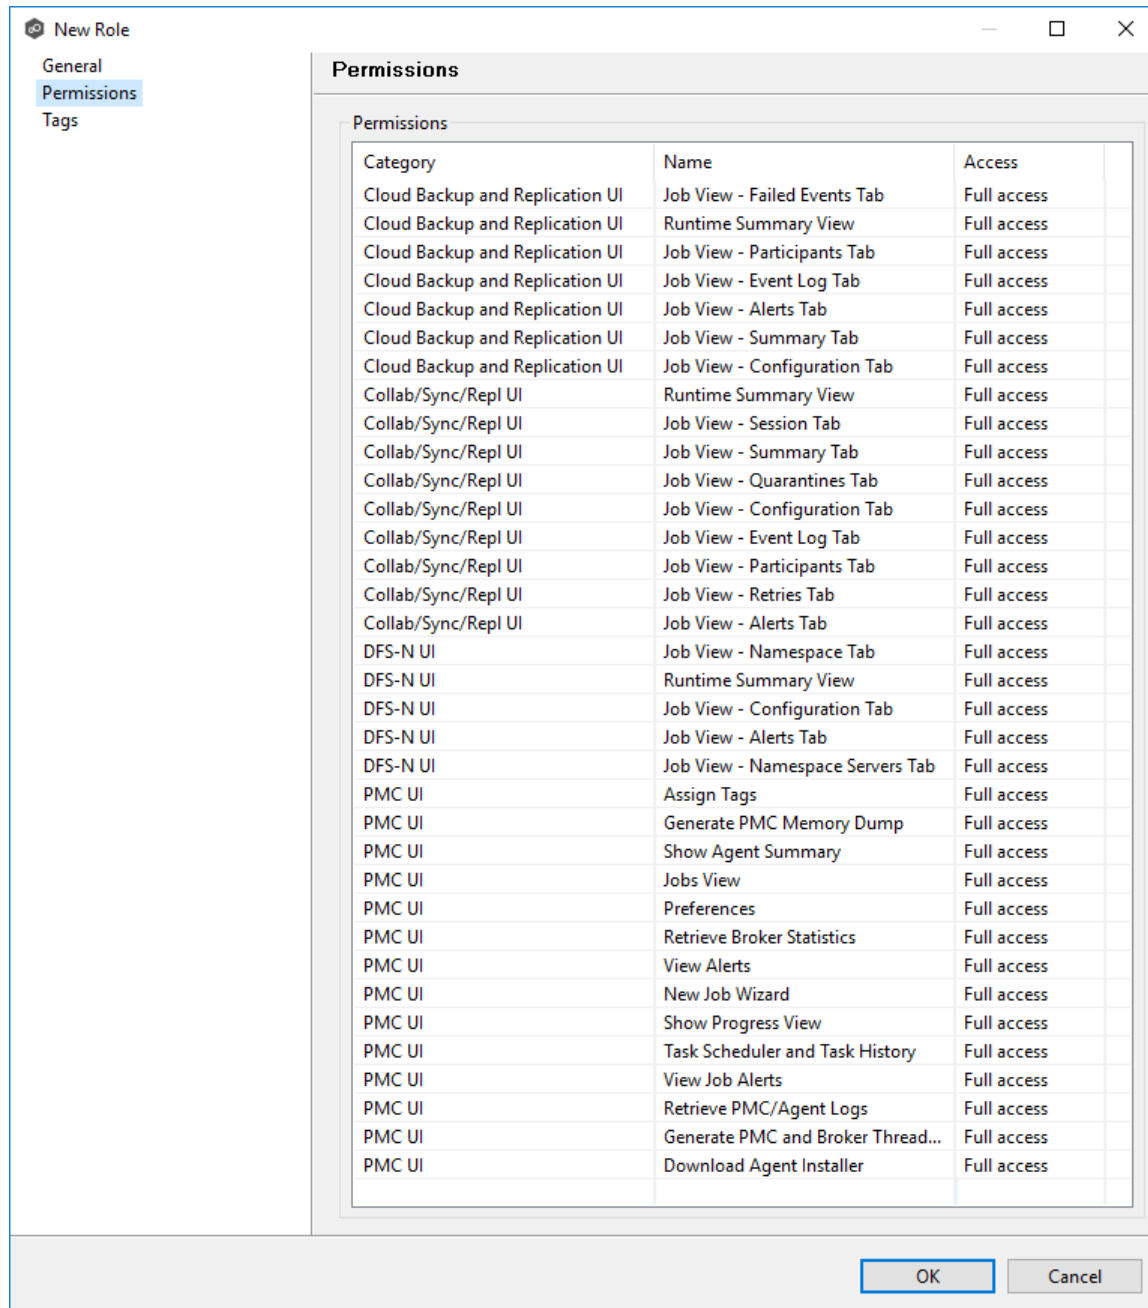
For more information, see [Creating a Custom Role](#).

Custom Role Permissions

You can [create custom roles](#) in User Management and specify the permissions you want associated with the role.

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

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



Advanced Topics

This section discusses the following topics:

- [Analytics](#)

- [Proactive Monitoring](#)
- [Conflicts, Retries, and Quarantines](#)
- [DFS Namespaces](#)
- [Edge Caching](#)
- [File Metadata Synchronization](#)
- [Managing Peer Agents](#)
- [PeerGFS API](#)
- [Scheduled Replication Filters](#)
- [Smart Data Seeding](#)
- [Storage Capacity](#)
- [TLS Certificates](#)

Analytics

PeerGFS offers four distinct analytics capabilities tailored to varying needs and levels of access:

1. Anonymous Diagnostic Data:

Purpose: This feature allows PeerGFS to collect [anonymized diagnostic data](#) to aid in improving the product's performance, reliability, and features.

Accessibility: This information is sent to Peer Software and used exclusively for internal purposes.

Opt-out: Users have the option to disable this feature via the Analytics preferences page if they prefer not to share diagnostic information.

2. PeerIQ:

Purpose: [PeerIQ](#) is a virtual appliance-based analytics system. It serves as a comprehensive platform for monitoring the health and performance of PeerGFS and the storage environment and is the foundation for Peer Software's overall analytics capabilities.

Features: PeerIQ provides extensive analytics capabilities, including real-time monitoring, performance metrics, and insights into PeerGFS and the storage environment. Includes File System Analytics and File Activity Analytics

Accessibility: Access to PeerIQ is determined by the user's subscription license purchased from Peer Software.

3. File System Analytics:

Purpose: [File System Analytics](#) is part of PeerIQ and leverages scan-based data from Peer Agents to provide insights and trends about the unstructured data residing on the storage infrastructure.

Features: This capability offers detailed analysis of file systems, including file types, sizes, data age, data access patterns, and other relevant metrics.

Accessibility: Access to File System Analytics is determined by the user's subscription license purchased from Peer Software.

4. File Activity Analytics:

Purpose: [File Activity Analytics](#) is part of PeerIQ and leverages real-time-based data from Peer Agents to provide insights and trends about how your unstructured data and storage infrastructure is being utilized.

Features: This capability offers analysis of user activity over time, including a list of the top users by activity.

Accessibility: Access to File Activity Analytics is determined by the user's subscription license purchased from Peer Software.

5. Proactive Monitoring:

Purpose: [Proactive Monitoring](#) is an option available exclusively to PeerGFS customers with dedicated Technical Account Managers (TAMs). It offers comprehensive insights to the Peer Software TAM team into the health and performance of the PeerGFS environment, including Agent information, job information, and overall PMC information.

Accessibility: This option is available to customers with dedicated TAMs and is used to provide regular reviews of the status of their PeerGFS environment. Data collected is securely uploaded to a Peer Software-owned storage account in Microsoft Azure and is accessible only to authorized Peer Software employees.

These analytics capabilities offer a range of options for monitoring, analyzing, and optimizing the PeerGFS environment, catering to different user requirements and preferences. See [Analytics Preferences](#) for more information about enabling and configuring these analytics capabilities.

File System Analytics

The primary purpose of File System Analytics (FSA) is to gather file system data and send it to PeerIQ for analysis and visualization. The data is then displayed on the File System Analytics pages within the [PeerIQ](#) interface.

FSA offers advanced capabilities for analyzing unstructured data within your storage infrastructure. Leveraging the same scan engine utilized by the standalone File System Analyzer tool, FSA enhances PeerIQ with the following features:

- **Expanded Protocol Support:**

File System Analytics supports both SMB (Server Message Block) and NFS (Network File System) protocols, ensuring comprehensive coverage across diverse storage environments.

- **Centralized Reporting:**

PeerIQ provides a centralized and user-friendly set of reports, consolidating insights from all file servers within your environment. This simplifies the process of accessing and analyzing data across multiple storage systems.

- **Detailed File Type and Extension Reports:**

File System Analytics offers detailed reports on file types and extensions. This allows users to gain insights into the composition of their data and identify trends related to specific file formats.

- **Data Aging:**

File System Analytics provides information on file timestamps and aging, facilitating informed decision-making regarding storage management and optimization.

- **Historical Analysis:**

Administrators can access historical data, enabling them to track file type and data age changes over time.

- **Scheduled Scans:**

Scans are scheduled to run every Saturday at 9 PM by default, covering all volumes managed in real-time by PeerGFS. The [File System Analytics scan process](#) ensures regular and consistent data collection for analysis. Future releases of PeerIQ will include options for adjusting scan schedules and coverage areas to better align with specific user requirements.

See [Analytics Preferences](#) for information about enabling and configuring File System Analytics.

The scan process used by File System Analytics (FSA) typically follows these steps:

1. **FSA Scan Timing:** By default, the FSA scan runs every Saturday at 9 PM. The duration of the scan can vary, potentially taking minutes, hours, or even days to complete depending on the size and complexity of the file system. The FSA traverses through the file system hierarchy, analyzing each directory, subdirectory, and file encountered during the scan. During the traversal, the FSA collects data about the file system structure, file properties, and other relevant information. It collects various metadata and attributes associated with files and directories, such as size, permissions, timestamps, and file types. This data may include details about file types, extensions, sizes, access patterns, and ownership.
2. **Zip File Creation by Agents:** Once the Agents complete the scan, the Agents generate a zip file containing the scan results. However, the zip file is not automatically sent to PeerIQ.
3. **PMC Zip File Check:** Peer Management Center (PMC) checks daily to see if there is a zip file ready to be sent. If a zip file is found, the PMC initiates the process to send it to PeerIQ.
4. **Alerts for Zip File Transfer:** When the zip file is sent, an alert is generated in the Alerts view in the PMC. This alert indicates whether the transfer was successful or not, providing visibility into the status of the data transfer process.

By following this process, File System Analytics ensures that the scan results are effectively transferred to PeerIQ for further analysis and visualization, with alerts providing notification of the transfer status.

File Activity Analytics

The primary purpose of File Activity Analytics (FAA) is to gather statistics about how users and clients are accessing file server data and send them to PeerIQ for analysis and visualization. The data is then displayed on the File Activity Analytics pages within the [PeerIQ](#) interface.

FAA offers advanced capabilities for analyzing unstructured data within your storage infrastructure. Leveraging the same real-time engine as the standalone File Activity Analyzer tool, FAA enhances PeerIQ with the following features:

- **Expanded Protocol Support:**

File Activity Analytics supports both SMB (Server Message Block) and NFS (Network File System) protocols, ensuring comprehensive coverage across diverse storage environments.

- **Centralized Reporting:**

PeerIQ provides a centralized and user-friendly set of reports, consolidating insights from all file servers within your environment. This simplifies the process of accessing and analyzing data across multiple storage systems.

- **User Activity:**

File Activity Analytics offers a report showing user behavior over time across all file servers as well as a list of the top ten users by overall activity. This provides insights into which users are most active and their patterns of activity.

- **Historical Analysis:**

Administrators can access historical data, enabling them to track user activity changes over time.

- **Real-time Statistics Gathering:**

Unlike with the scan-based engine of File System Analytics, File Activity Analytics relies on statistics that are collected and sent to PeerIQ on an ongoing basis.

Upcoming releases of PeerIQ will add additional user-based and client-IP-based insights as well as details on most commonly accessed extensions and folder paths.

See [Analytics Preferences](#) for information about enabling and configuring File Activity Analytics.

PeerIQ

Peer Software's PeerIQ is a powerful virtual appliance provided to subscription customers, offering a self-hosted dashboard and analytics environment for monitoring the health and performance of PeerGFS and the replication environment. Here's a detailed overview of PeerIQ's capabilities and features:

1. **Deployment Options:**

PeerIQ is available as a virtual appliance in multiple formats, including VMware ova, Hyper-V vhd, and Nutanix AHV qcow2. This allows flexibility in deployment across various virtualization platforms.

2. **Ease of Setup:**

Setting up PeerIQ is quick and straightforward, requiring a simple configuration on the PeerIQ side and enabling a single option on the [Analytics](#) preference page. This streamlined setup process ensures rapid deployment and integration with the existing infrastructure.

3. **Web-Based Interface:**

PeerIQ's dashboards are accessible through a web browser, providing a visual and interactive interface for monitoring and analyzing telemetry data. The interface offers real-time updates, with data refreshed automatically every few minutes by default.

4. Comprehensive Monitoring:

PeerIQ provides intelligent insights into the performance and health of PeerGFS and the replication environment. Key information includes agent details (disk space and memory utilization), job information (watch set growth and overall performance), overall PMC information (disk space, memory utilization, queue backlogs, and quarantine counts), and details about volumes monitored by PeerGFS (free and consumed space).

5. License Utilization:

PeerIQ offers visibility into PeerGFS license utilization, providing projections about when the licensed capacity might be exhausted. This helps in capacity planning and optimizing resource allocation.

6. File Type and Age Insights:

Using data from File System Analytics, PeerIQ offers insights into the types and ages of files stored on the storage infrastructure, facilitating better understanding of data composition, and aiding in data management decisions.

7. User Activity Insights:

Using data from File Activity Analytics, PeerIQ offers insights file and folder access patterns, facilitating better understanding of end user activity, and aiding in data management decisions.

8. Enhanced Historical Data:

PeerIQ consolidates information available in the Peer Management Center interface, providing historical data not found in the PMC interface. This enables users to track trends over time and gain deeper insights into system performance and behavior.

9. Future Roadmap:

PeerIQ's future development roadmap includes plans to provide proactive information about the PeerGFS environment, additional trends and insights about data stored on file servers, and additional analysis of user and application data usage across the storage infrastructure. These enhancements will further enrich PeerIQ's capabilities and provide valuable insights for system administrators.

For more information about PeerIQ, see articles about [PeerIQ](#) in our knowledge base.

Proactive Monitoring

Proactive Monitoring for PeerGFS, available to customers with dedicated Technical Account Managers (TAMs), offers comprehensive insights into the health and performance of the PeerGFS environment. Here's a breakdown of the data collected:

1. Agent Information:

- **Disk Space Utilization:** Monitoring the disk space usage of all agents ensures that sufficient storage is available for operations and prevents potential disk space-related issues.
- **Memory Utilization:** Tracking memory usage helps identify any memory-intensive processes or potential memory leaks on the Agents, ensuring optimal performance.

2. Job Information:

- **Watch Set Growth:** Monitoring the growth of watch sets helps track changes in the data being monitored by PeerGFS, enabling proactive management of storage requirements and system resources.
- **Overall Performance:** Analyzing job performance metrics provides insights into the efficiency and effectiveness of data replication and synchronization processes, helping optimize system performance.

3. Overall Peer Management Center Information:

- **Disk Space and Memory Utilization:** Monitoring the PMC's disk space and memory usage ensures the smooth operation of the management console and helps prevent performance degradation due to resource constraints.
- **Queue Backlogs:** Tracking queue backlogs helps identify any bottlenecks or delays in job processing, enabling timely intervention to maintain system efficiency.
- **Quarantine Counts:** Monitoring the number of files in quarantine provides insights into potential issues with data integrity or synchronization errors, facilitating prompt resolution to maintain data consistency.

This data are securely uploaded to a Peer Software-owned storage account in Microsoft Azure and is accessible only to authorized Peer Software employees. Regular reviews of this data by PeerGFS experts enable proactive monitoring and management of the PeerGFS environment, ensuring optimal performance, reliability, and data integrity for customers.

See [Proactive Monitoring Preferences](#) for information about enabling and configuring Proactive Monitoring. For information about data sent to Proactive Monitoring, see [What types of data are uploaded to Proactive Monitoring?](#) in our knowledge base.

Conflicts, Retries, and Quarantines

Making unstructured data active at multiple locations increases the chance of users making conflicting changes to different copies of the same file. The real-time synchronization and locking engines built into Peer Global File Service are designed to prevent these conflicts by ensuring that only one user can modify a file at a time while also making sure that all locations always have the most up to date version of a file. There are scenarios, however, where the synchronization and locking engines may not be able to prevent version conflicts. Such scenarios include network outages and file system issues.

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

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

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

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

File retries can occur for any of the following reasons:

- The transfer of a file between locations is interrupted for any reason.

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

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

Two users have the same file open at two different locations prior to a Peer Global File Service job being enabled. When starting the job, PeerGFS will track this file as a potential conflict. If only one or no users make a change to the file, this conflict will automatically be resolved. If both users make a change, the conflict will become a quarantine.

DFS Namespaces

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

Overview

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

The elements that make up a DFS namespace are:

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

content are stored. For example, if a user navigates to `\Domain.local\MyNamespace\MyFolder`, the user is transparently redirected to `\NYC-FS.Domain.local\MyFolder` or `\\LA-FS.Domain.local\MyFolder`, depending on which site the user is currently accessing. Adding multiple folder targets increases the availability of the folder in the namespace.

For more information about DFS namespaces, see [DFS Namespaces overview](#) on Microsoft's website.

Managing DFS Namespaces through PeerGFS

PeerGFS enables you to create a namespace and manage various activities related to it, such as creating namespace folders, adding folder targets, and linking the namespace to a File Collaboration or File Synchronization job. You could manage DFS namespace using Microsoft tools; however, you can manage [DFS namespaces](#) through a dedicated job type in Peer Management Center, the [DFS-N Management job](#).

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

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

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

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

- [Using DFS Namespaces with Jobs](#)
- [DFS Namespace Failover and Failback](#)
- [DFS-N Management Jobs](#)
 - [Creating a DFS-N Management Job](#)
 - [Managing DFS Namespaces](#)

- [Linking a DFS Namespace to File Collaboration or File Synchronization Job](#)

Using DFS Namespaces with Jobs

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

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

- You can create a new namespace or import an existing one by first creating a DFS-N Management job and then later [linking the namespace to a File Collaboration or File Synchronization job](#).
- You can create a new namespace or import an existing one when creating a File Collaboration or File Synchronization job, thus automatically linking namespace folder targets to the watch sets of the collaboration or synchronization participants. A DFS-N Management job is automatically created during this process.
- You can also import an existing namespace by right-clicking on the Namespace Summary view which will guide you through the import of a namespace into PeerGFS. [Importing an existing namespace](#) will automatically create a DFS-N Management job which can then be linked to a File Collaboration or File Synchronization job.

Before creating jobs that use namespaces, you may want to [configure DFS preferences](#).

See [Managing DFS Namespaces](#) for information about adding namespace servers, namespace folders, or folder targets to a DFS namespace.

DFS Namespace Failover and Failback

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

Failover

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

automatic failover to another folder target, see [DFS-N Management](#) in [Collaboration, Replication, and Synchronization Job Preferences](#).

Failback

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

1. Kicks off a rescan to ensure the disconnected site or file server is brought back in sync with the others.
2. Re-enables the associated folder target once the re-scan is complete. Once this is done, DFS Namespaces begins to direct end users back to this file server.

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

Edge Caching

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

Overview

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

Edge Caching offers flexible edge storage management with:

- The ability to assign an amount or percentage of available storage to be used on the edge storage device.

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

Fundamental Concepts

A **master participant** has a complete set of **hydrated** files and no stub files. An **edge participant** contains a subset of the complete, hydrated files on a master participant, while the rest of the files will be stub files that don't take up any space. Users can retrieve stubbed files directly from a master participant as needed. The goal of Edge Caching is to keep as much as possible cached locally on edge participants for rapid access.

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

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

File Metadata Synchronization

Overview

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

The Security Descriptor Information elements that can be synchronized are:

- **Owner:** NFTS Creator-Owner. By default, the owner is whoever created the object. The owner can modify permissions and give other users the right to take ownership.

- **DACL:** Discretionary Access Control List. It identifies the users and groups that are assigned or denied access permissions to a file or folder.
- **SACL:** System Access Control List. It enables administrators to log attempts to access a secured file or folder and is used for auditing.

File Metadata Conflict Resolution

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

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

ACL Requirements

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

Multi-Protocol

Overview

PeerSoftware's multi-protocol functionality allows both NFS and SMB clients to interact with the same set of data within PeerGFS, eliminating the need for separate jobs for different protocols. This enables seamless data sharing and access between Linux and Windows systems.

What Do We Mean by Multi-Protocol?

Multi-protocol in PeerGFS refers to the capability of a job to detect real-time events from both NFS (NFSv3 and NFSv4) clients and CIFS (SMB) clients simultaneously. This allows the same

data to be accessed via different protocols (NFS for Linux clients and SMB for Windows clients) without the need to create separate jobs for each protocol.

How Does It Work?

To enable multi-protocol functionality, both a share and an export must point to the same directory to be monitored by PeerGFS. These configurations allow both NFS and SMB clients to access the same data set. The primary security style configured on the NAS system determines the metadata settings for the PeerGFS job.

- **Share:** A network share that provides access to the data for SMB clients.
- **Export:** A directory export that provides access to the data for NFS clients.

These elements work together to provide unified access to the data set, whether the client is Linux-based (NFS) or Windows-based (SMB).

What Type of Jobs Does Multi-Protocol Apply To?

Multi-protocol functionality can be applied to the following job types in PeerGFS:

- **File Synchronization**
- **File Replication**

What Are the Requirements to Use Multi-Protocol?

- **Supported Client:** PeerGFS Multi-Protocol support is available for Linux clients only.
- **Supported NAS Platforms:**
 - NetApp ONTAP (cDOT)
 - Amazon FSXN
 - Nutanix Files
 - Dell PowerScale (Syslog only)

Ensure your system meets these prerequisites before configuring multi-protocol support.

What Limitations or Restrictions Exist?

While multi-protocol functionality offers flexibility, there are some limitations to keep in mind:

- **Transfer Method:** NFS will perform file transfers, maintaining name capitalization during the process.
- **File Collaboration Jobs:** File Collaboration job types are not supported with multi-protocol functionality.
- **Job Overlapping:** Job relay for File Replication jobs (overlapping jobs) is not supported.
- **File Permissions:** The user determines which file permissions (metadata) will be transferred. This ensures predictable scans.
- **Security Style Limitation:** A multi-protocol job can only support one security style at a time—either Linux security (NFSv4, NFSv3, Posix) or Windows ACLs, but not both simultaneously.
- **Mixed File Permissions:** Mixed file permissions (SMB & NFS) will be supported on verified platforms, but with the limitation of not supporting both security styles simultaneously.

Who Will Want to Use Multi-Protocol?

PeerGFS's Multi-Protocol (MKP) functionality is designed for organizations that require **seamless cross-platform data access** and **collaboration** between **Linux (NFS)** and **Windows (SMB)** systems. Typical users who would benefit from this feature include:

- **Enterprises with Mixed Environments:** Organizations that have a combination of Linux and Windows workloads needing concurrent access to shared data.
- **Data Centers with Diverse Workloads:** Enterprises running applications like electronic design automation (EDA), media and entertainment workflows, or other data-intensive tasks that involve both Linux-based compute instances and Windows-based applications.
- **IT Administrators:** IT teams managing mixed environments that need to ensure both Linux and Windows systems can efficiently access and work with the same data set, without the complexity of managing separate data sets or jobs for each protocol.
- **Collaborative Teams:** Teams that require file synchronization and collaboration across platforms (NFS for Linux and SMB for Windows), ensuring that all users—regardless of OS—can access the same data in real-time.
- **Data Replication and Backup:** Organizations with multi-protocol environments that need to replicate or back up data to multiple platforms simultaneously, such as cloud backup services supporting both Linux and Windows-based systems.

This functionality simplifies data access management and enhances workflow efficiency in heterogeneous IT environments, especially for organizations with both Linux and Windows clients that need to share the same data.

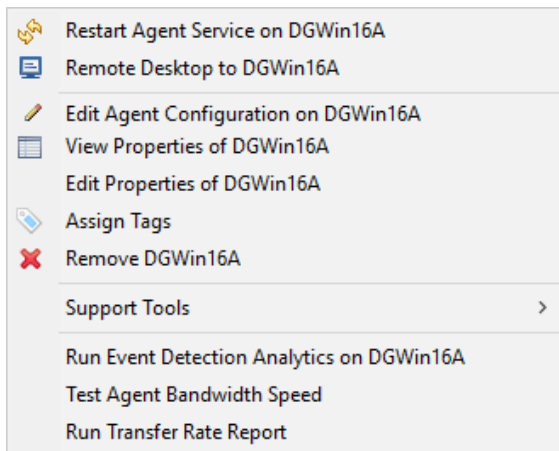
Network of Brokers

The Peer Management Broker is a key technology that is used by PeerGFS to facilitate communication between Agents and the PMC. With the Network of Brokers capability, customers can deploy multiple instances of the Peer Management Broker across their infrastructure to better optimize and control the flow of replication traffic.

The Network of Brokers capability significantly enhances an active-passive configuration by enabling automatic failover from a primary PMC (server) to a backup PMC (server). While it is possible to establish an active-passive configuration without Network of Brokers technology, utilizing Network of Brokers facilitates automatic failover instead of requiring manual intervention for the failover process.

For information using the Network of Brokers capability, see [Getting Started with Network of Brokers](#). For more information about setting up failover to a backup PMC, see the article [Achieving high availability for the PMC through active-passive configuration](#) in our knowledge base.

Managing Peer Agents



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

Options

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

Option	Description
Support Tools	Displays a list of tools that can be used to assist Peer Software Support.
Run Event Detection Analytics on Agent	Runs the Event Detection Analytics tool for the selected job, which looks at real-time activity that has been occurring on that specific Agent.
Test Agent Bandwidth Speed	Runs a bandwidth speed test to be performed in the background if the selected Peer Agent is connected. You are notified at completion with the results of the test.
Run Transfer Rate Report (not available on Web Client)	(Rich client only) Displays a time series performance chart of average transfer rate for the selected Peer Agent over the last 24 hours.

Support Tools

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

Command	Description
Retrieve Log Files from Agent	Retrieves log files for the selected Agent. The log files contain information that Peer Support uses in debugging issues. The log files are encrypted and are located in the support folder of the Peer Management Center installation directory. They can optionally be uploaded to the Peer Support team.
Generate Thread Dump on Agent	Generates a thread dump file for the selected Agent, which can be used by Peer Support to debug certain issues. The debug file is located in the Peer Agent installation directory.
Generate Memory Dump on Agent	Generates a memory dump file for the selected Agent, which can be used by Peer Support to debug certain issues. The debug file is located in the Peer Agent installation directory.

Peer Agent Connection Statuses

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

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

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

Agent can have the following statuses:

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

Re-enabling a Disabled Agent Within a Job

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

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

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

The screenshot shows the 'Participants' tab in the Peer Management Center. The 'Host Participants' table is as follows:

Host	Root Path	Status	State	Message
DellIT110a	\\svm9x-1\cifs1\Departments\Sales	Disabled	Disabled	Malicious Event Detection (MED) - Bait File Alert (A...
DellIT3610b	C:\Departments\Sales			

The 'Participant Details' popup for DellIT110a shows the following information:

- Host Name: DellIT110a
- 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 appld=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=DellIT110a|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

At the bottom of the interface, the status is shown as 'Halted. (Quorum Lost)'.

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


Editing an Agent Configuration

The ability to remotely manage the configuration of connected [Peer Agents](#) is available from within Peer Management Center.

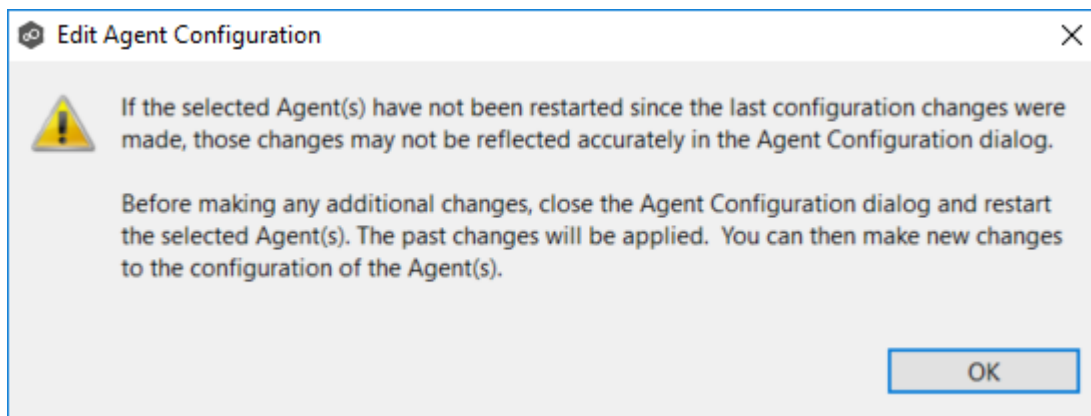
Note: For customers using clustered file server roles with Windows Failover Cluster, review this knowledge base article: [Using a Peer Agent in a Windows Failover Cluster](#). Custom Agent settings must be applied to each potential node in the cluster that may host the Peer Agent. Contact [Peer Support](#)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

To edit an Agent's configuration:

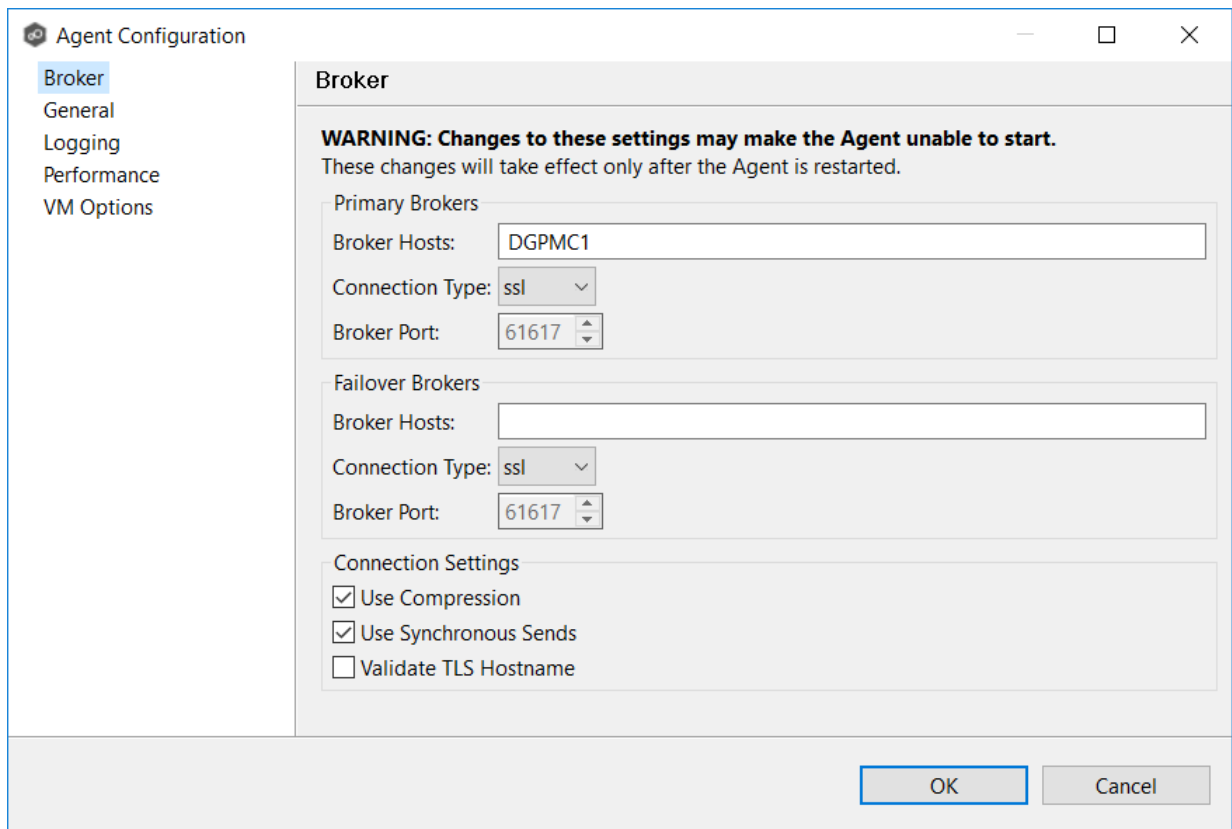
1. Right-click the connected Peer Agent in the **Agents** view:



2. Select **Edit Agent Configuration**.
3. If the following dialog appears, Click **OK**:



The **Agent Configuration** dialog appears.



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

- [Broker](#)
- [General](#)
- [Logging](#)
- [Performance](#)
- [VM Options](#)

5. Click **OK**.

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

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

The settings in the **Broker** page apply to communication between the selected Peer Agent and broker(s) only; they do not apply to communication between Peer Management Center and a broker.

Agent Configuration

Broker

WARNING: Changes to these settings may make the Agent unable to start.
These changes will take effect only after the Agent is restarted.

Primary Brokers

Broker Hosts: 192.168.169.154

Connection Type: ssl

Broker Port: 61617

Failover Brokers

Broker Hosts:

Connection Type: ssl

Broker Port: 61617

Connection Settings

Use Compression

Use Synchronous Sends

Validate TLS Hostname

OK Cancel

Options

Primary Brokers

Options	Description
Broker Hosts	Enter the IP address or fully qualified host name of the server running the primary broker.

Options	Description
	This option will also accept a comma-separated list of IPs or FQDNs. Agent will connect to any of the primary brokers in the order that they in listed. Agent will try to failover to a primary broker first before trying the failover brokers.
Connecti on Type	Select the type of connection to use when communicating with the primary broker(s). Types include SSL (encrypted using TLS v1.3 by default) and TCP (not encrypted).
Broker Port	The port on which to communicate with the primary broker(s).

Failover Brokers

Options	Description
Broker Hosts	Enter the IP address or fully qualified host name of the server running the secondary broker. This option will also accept a comma-separated list of IPs or FQDNs. Agent will connect to any of the failover brokers in the order that they in listed but only after failing to connect to all primary brokers.
Connecti on Type	Select the type of connection to use when communicating with the failover broker (s). Types include SSL (encrypted using TLS v1.3 by default) and TCP (not encrypted).
Broker Port	The port on which to communicate with the failover broker(s).

Connection Settings

Option	Description
Use Compres sion	Enable to compress all communication between the selected Agent and broker(s).
Use Synchro	Enable to always send messages from an Agent to broker(s) in synchronous mode. If not enabled, then messages between Agent

Option	Description
nous Sends	and broker(s) will always be sent asynchronously. Note: Enabling this will affect the performance of communication between the broker(s) and the Agent, especially over connections with high latency.
Validate TLS Hostname	Enable if you are using your own certificates and would like certificate hostnames to be validated between an Agent and broker(s).

The **General** page has three sets of options:

- [Workspace](#)
- [Location Information](#)
- [General](#)

The screenshot shows the 'Agent Configuration' dialog box with the 'General' tab selected. The left sidebar contains the following options: Broker, General (highlighted), Logging, Performance, and VM Options. The main area is titled 'General' and contains three sections: 'Workspace' with a text field for 'Agent Workspace Directory' containing the value 'workspace'; 'Location Information' with three text fields for 'Location:', 'Longitude:', and 'Latitude:'. The 'Statistics' section has two checked checkboxes: 'Enable Server Statistics' and 'Enable Disk Statistics'. At the bottom right, there are 'OK' and 'Cancel' buttons.

Options

Workspace

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

Location Information

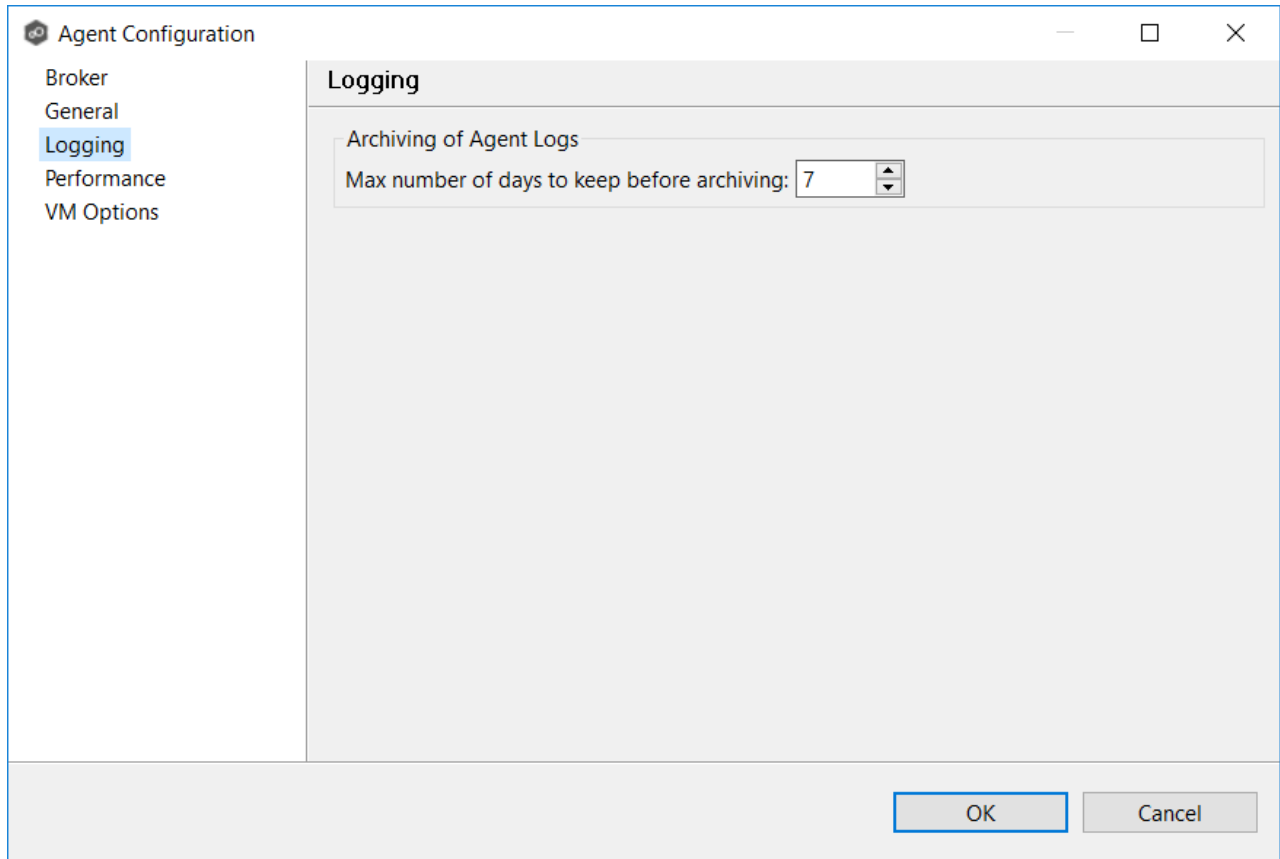
Agent location information is used by [Proactive Monitoring](#). If you change any location values, you must restart the Agent Service.

Option	Description
Location	Enter the city, state, and country where this Agent is located.
Longitude	Enter the longitude coordinates of this Agent.
Latitude	Enter the latitude coordinates of this Agent.

Statistics

These statistics are useful to identify performance bottlenecks. Statistics are collected every 60 seconds and stored in a database that can be sent to [Peer Support](#).

Option	Description
Enable Server Statistics	Select this to collect statistics about network latency, CPU usage, and memory usage.
Enable Disk Statistics	Select this to collect statistics about disk latency. Two key components will be monitored: the workspace folder located in the Agent's installation directory and the watch sets of all jobs tied to that Agent.

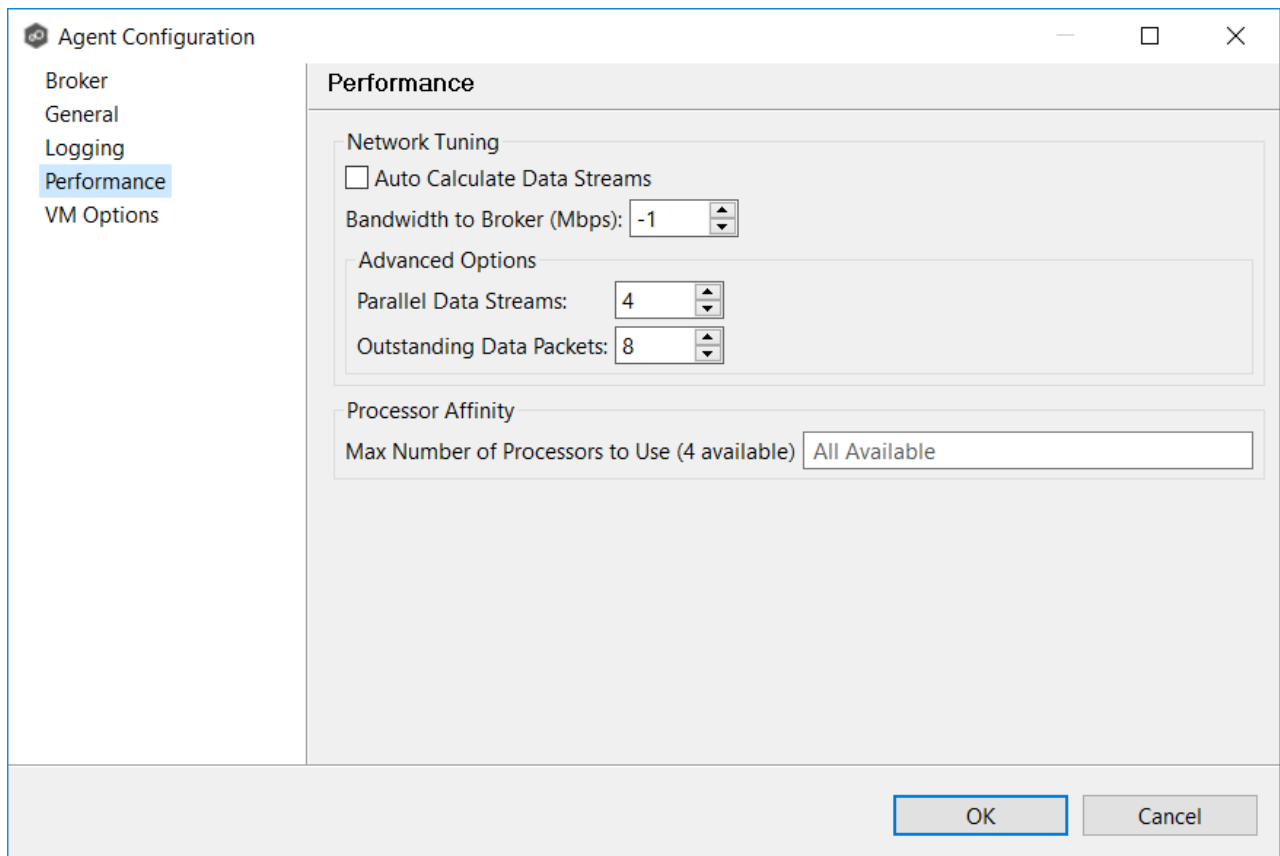


Option

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

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

- [Network Tuning](#) - These settings control the number of parallel streams of data that can be sent between the Agent and the Broker. In active, latent environments, adjusting these settings can improve performance or limit the data throughput between the Agent and the Broker.
- [Processor Affinity](#) - Allows you to specify the number of processors that the Agent should use.



Options

Network Tuning


Field	Description
Auto Calculate Data Streams	Select this checkbox if you want the number of blob command threads to be calculated rather than using the value in the Parallel Data Streams field. The optimum number of Agent parallel data streams is calculated based on network performance, using the value of Bandwidth to Broker (Mbps) in the calculation along with latency between the Broker and Agent.

Field	Description
Bandwidth to Broker (Mbps)	Enter the bandwidth in megabits per second that you want to use for the connection between the Agent and Peer Management Broker. The default value is -1, which means use all available bandwidth.
Parallel Data Streams	Enter the maximum number of threads to handle data transfer between each Agent and the Peer Management Broker. Increasing this typically improves replication performance but also increases memory consumption. The default value is 4. The minimum is 1; the maximum is 100.
Outstanding Data Packets	Modify this setting only at the instruction of the Peer Support team as it can lead to increased memory consumption. Enter the maximum number of blocks of data to be buffered to be sent to the Agent. The default number is 8; the maximum size is 100.

Processor Affinity

Field	Description
Max Number of Processors to Use (x available)	Enter the number of processors that the Agent process can use on the server where it is installed. This number should be less than or equal to the number of processors available on the server. The default value is -1, which means use all available processors.

The option on the page allows you to tune the maximum amount of system memory that the Peer Agent service will use on the server where it is installed. If you change the value, you must restart the Agent Service.



Options

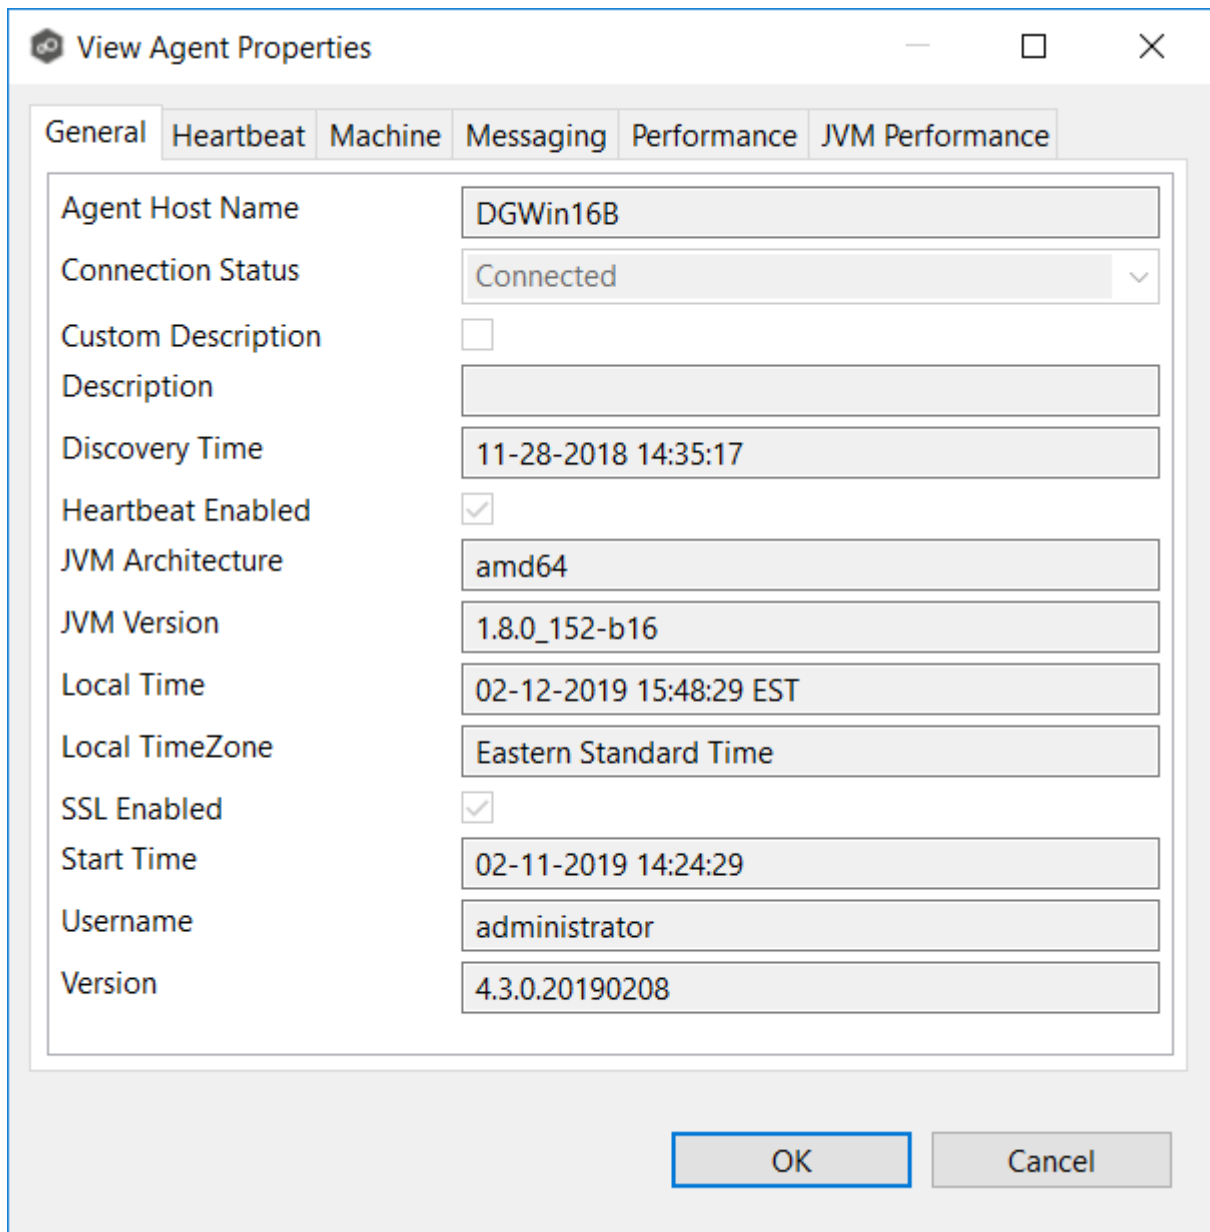
Field	Description
Maximum Agent Memory (in MB)	Enter the maximum amount of memory in megabytes that the JVM portion of the Agent service can use. We recommend a minimum value of 2048 MB on 64-bit Agent servers with a recommended maximum of 16384 MB. We strongly recommend that this value be set to no lower than 2 GB.

Viewing Agent Properties

To view the properties of an Agent:

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

The **View Agent Properties** dialog opens.



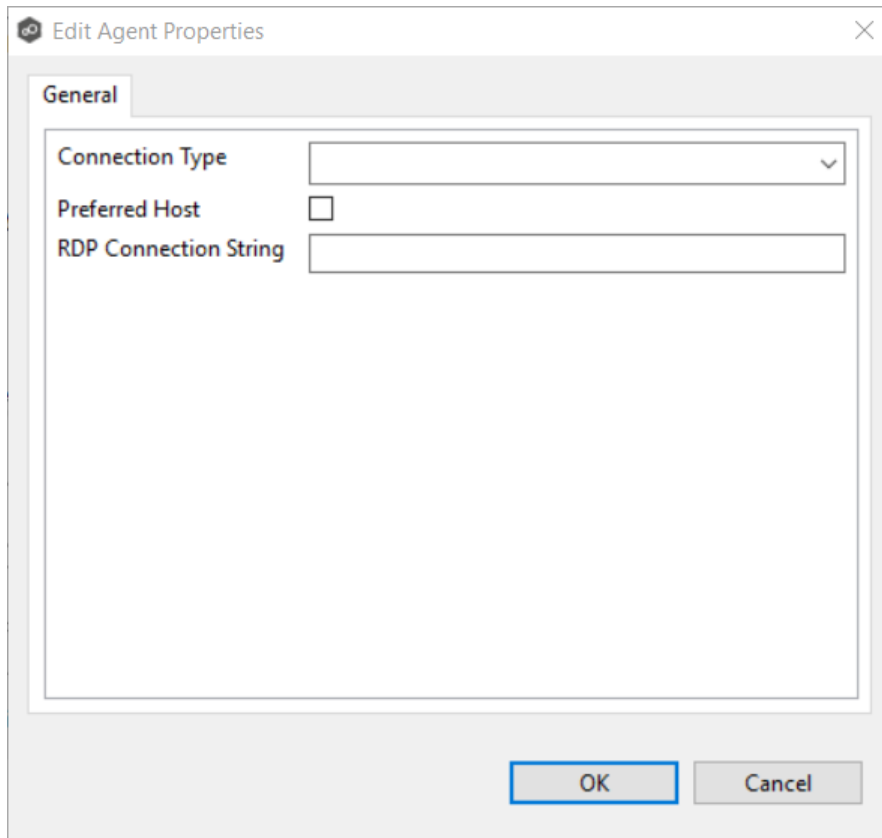
This dialog displays Peer Agent and host machine information across the following tabs:

Tab	Description
General	Displays general Peer Agent run-time information such as discovery time, local time, TLS use, Peer Agent start up time, Peer Agent version, and the user name Peer Agent service is running as.

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

Editing Agent Properties

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



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

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

Updating a Peer Agent

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

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

PeerGFS API

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

Currently, the API allows users to:

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

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

Accessing the PeerGFS API

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

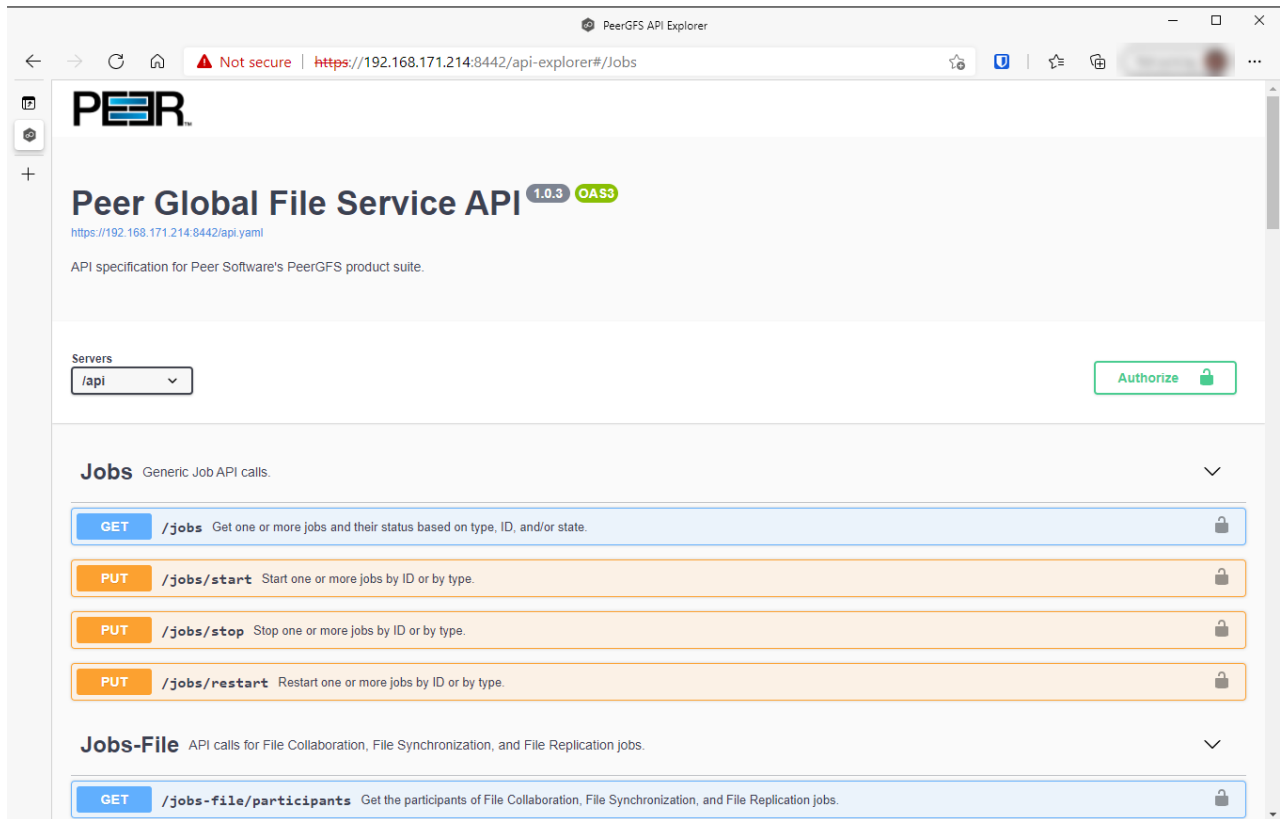
- A web URL hosted by the API service. This URL is defined as a combination of a PMC server name or IP and a port, as specified by the [Web and API Configuration](#) settings in [Preferences](#).
- Local (aka basic) authentication with a user name and password that is passed into a script or the API web interface. This user name and password is used to authenticate the user with the PeerGFS API service.

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

Testing the PeerGFS API

To test the PeerGFS API, you can use the API web interface:

1. Open a browser and navigate to the API endpoint (e.g., `https://<PMC IP or name>:8442`).
2. From there, you can execute and test various API calls.



Integrating Your Own Tools and Scripts with the PeerGFS API

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

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

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

The manifest should appear on the left with the front-end mockup on the right.

5. Use the Swagger Editor to generate client code.

API Quick Reference

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

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

The API calls are divided into four categories:

- **Jobs** - Generic job-related calls
- **Jobs-File** - Job-specific calls
- **Agents** - Agent-related calls
- **PMC** - Calls related to alerts, tasks, and logs
- **Storage Devices** - Calls reporting available storage devices

The PeerGFS API has three status codes:

200 - Success

401 - Unauthorized

404 - Job(s) not found

Scheduled Replication

Scheduled replication is a feature that allows you to delay replication of certain files and folders, allowing you to manage bandwidth and prioritize the replication of critical data in real-time. By using scheduled replication, you can reduce the impact of replication on network

bandwidth and ensure that critical data are replicated in real-time while less critical data are replicated on a schedule that makes sense for your organization. This can help to ensure that your data replication processes are efficient and effective, and that your network resources are being used as efficiently as possible.

To use scheduled replication, you create a **scheduled replication filter** that identifies the files and folders you want to replicate at a later time. The filter is based on file type. Once the filter is applied to a job, any files or folders that meet the criteria will be queued for replication at a scheduled time or interval that you specify.

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

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

Smart Data Seeding

Overview

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

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

Smart Data Seeding provides the ability to set one or more participants in a Smart Data Seeding mode. Smart Data Seeding hosts are considered the hosts from where files will be copied to all the other participants in the session. When a host is in Smart Data Seeding mode, it follows the rules of the job's Smart Data Seeding Mode configuration (see below). Initial scans run in a one-way mode to avoid bringing back deleted files. It is not recommended to have active ([Active-Active](#)) users on the target hosts. Once the initial scan is completed, the Smart Data Seeding host(s) are set back to their default full collaboration mode with no user interaction or final scan.

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

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

Smart Data Seeding Options

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

Mode	Description
PASSIVE (Default)	Initial scan will be one-way only with any host in Smart Data Seeding mode: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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): <ul style="list-style-type: none"> • Real-time activity on Smart Data Seeding host is enabled with latest modified file wins, regardless of whether the latest file is on the Smart Data Seeding host.

Mode	Description
REACTIV ATION	<p>Initial Scan will be one-way only with any host in Smart Data Seeding mode:</p> <ul style="list-style-type: none"> • Real-time activity on Smart Data Seeding host is enabled with Quarantine (Added and Updated Files will be quarantined during the scan). • Unsynchronized file updates will be quarantined during real-time. • Unsynchronized file renames will be restored. • Unsynchronized deletes will be restored.

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

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

Storage Capacity

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

Total capacity consumed is defined by the total number of unique TBs under management across all participants rather than the total capacity used by all participants. In this unique TB model, a 1 TB file that is synchronized across 10 participants only counts as 1 TB and not 10 TBs. For example, if your licensed storage capacity is 100 TB and you have a job with 5 participants totaling 20 unique TBs, you have used a total of 20% of your storage capacity, not 100%.

TLS Encryption and Certificate Management

PeerGFS secures data transmission between the PMC, Agents, and brokers using TLS encryption. TLS is also used to secure access to the PMC web user interface and API service. It relies on keys and certificates for host authentication and data encryption.

By default, PeerGFS includes preconfigured certificates. However, you can replace some or all of these with new certificates, either self-signed or issued by a Certificate Authority (CA).

PeerGFS includes utility scripts to simplify TLS certificate generation and replacement. For details on using these scripts, see the knowledge base article [Replace TLS Certificates](#).

Preferences

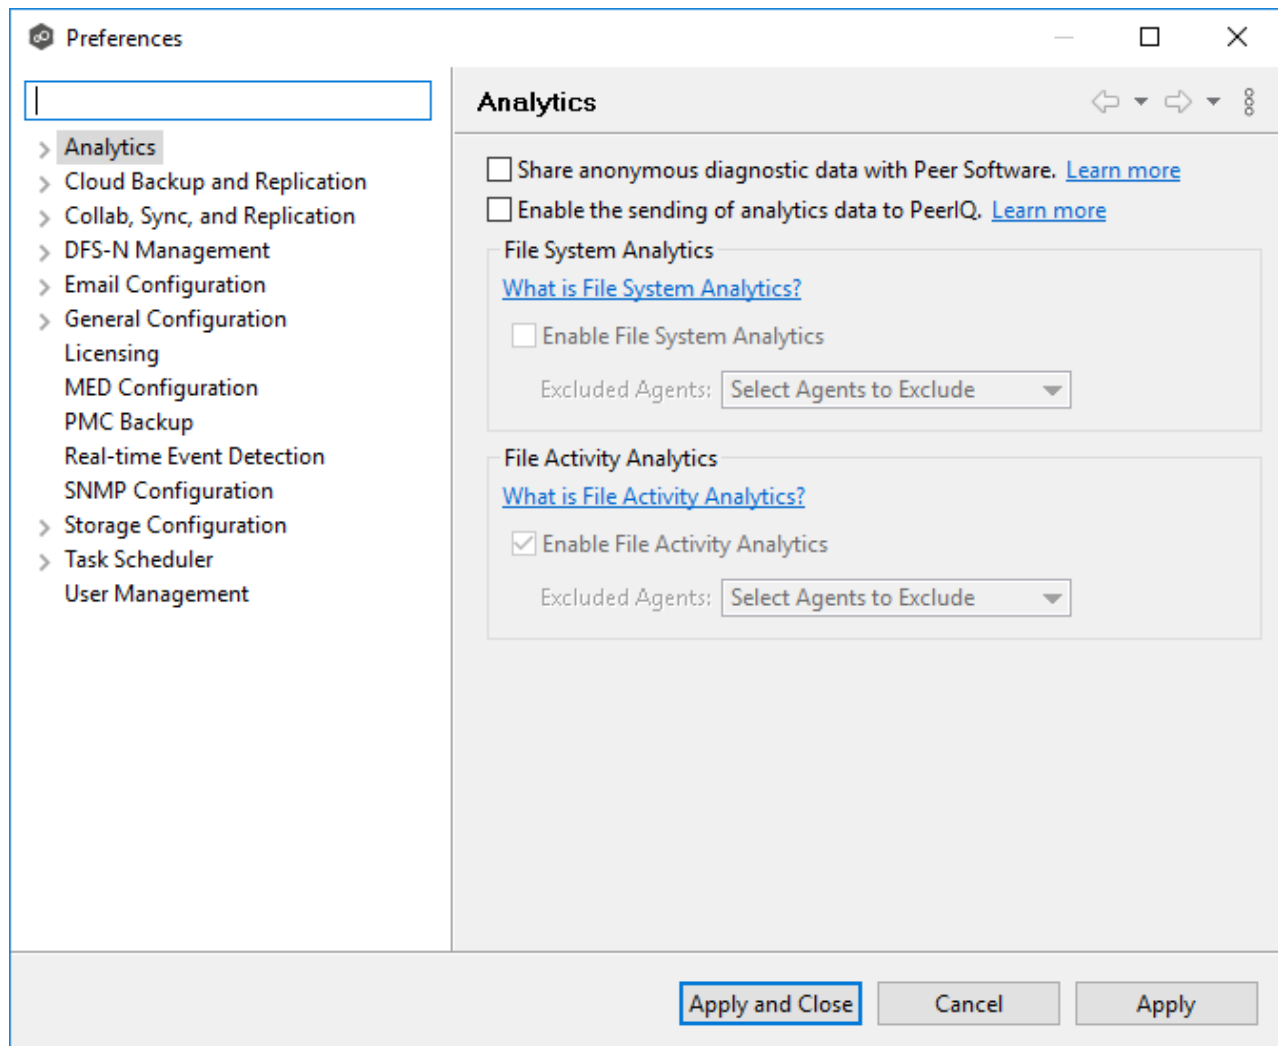
The **Preferences** dialog serves as a central hub for configuring settings in Peer Software's environment. It allows users to set up global settings that apply program-wide, as well as settings specific to particular job types. Peer Software recommends configuring these settings before creating any jobs or fine-tuning individual aspects of a job. Here's a breakdown of how the Preferences dialog is organized:

- Global settings - These settings apply program-wide and affect all aspects of the Peer Software environment. They include configurations such as system defaults, user preferences, and general behavior settings that are not specific to any particular job type.
- Job type-specific settings - In addition to global settings, the Preferences dialog also allows users to configure settings specific to each job type supported by Peer Software's environment. These settings tailor the behavior and parameters of individual job types, ensuring that each job type operates optimally according to its unique requirements and objectives.

By organizing settings in this manner, Peer Software provides users with flexibility and granularity in configuring their environment. Users can define global defaults while also fine-tuning settings for specific job types as needed, enabling efficient and customized management of tasks and workflows within the Peer Software ecosystem.

To access Preferences:

1. From the **Tools** menu, select **Open Preferences**.



Configuring Global Settings

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

- [Email Configuration](#)
- Contacts and Distribution Lists
- System Alerts

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

Configuring Job Type Specific Settings

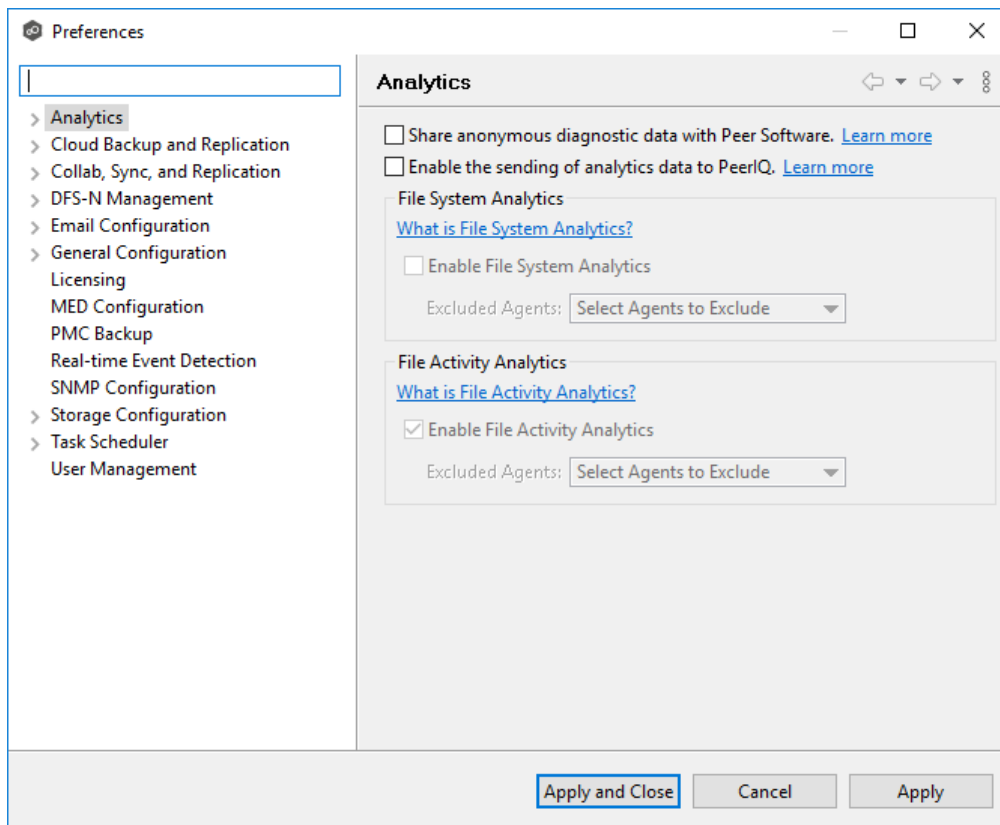
Job Type	Setting
Cloud Backup and Replication	<ul style="list-style-type: none">• Email Alerts• File and Folder Filters• Proxy Configuration.
File Collaboration, File Replication, and File Synchronization	<ul style="list-style-type: none">• Email Alerts• File and Folder Filters
DFS-N Management	<ul style="list-style-type: none">• Email Alerts• File and Folder Filters

Configuring Preferences

To modify preferences:

1. Use the navigation tree on the left to select a category.

The corresponding options will appear on the right side of the dialog. For example, select **General Configuration** in the navigation tree to view and configure general program-wide settings.



2. After making changes, click:

- **Apply and Close** to save the new settings and exit.
- **Cancel** to discard changes and exit.
- **Apply** to save your changes while keeping the **Preferences** dialog open.

Analytics

The **Analytics** page provides users with the ability to enable or disable various analytics settings within the Peer Software environment:

- **Anonymous Diagnostic Data:** This feature allows Peer Software to collect anonymized diagnostic data to aid in improving the performance, reliability, and features of PeerGFS. Information is sent to Peer Software and used exclusively for internal purposes.
- **PeerIQ:** PeerIQ provides extensive analytics capabilities, including real-time monitoring, performance metrics, and insights into the PeerGFS environment. PeerIQ also includes the following capabilities:

- **File System Analytics:** The primary purpose of FSA is to gather file system data from Peer Agents and send it to PeerIQ for analysis and visualization.
- **File Activity Analytics:** The primary purpose of FAA is to gather file activity data from Peer Agents and send it to PeerIQ for analysis and visualization.
- **Proactive Monitoring:** This feature offers comprehensive insights into the health and performance of the PeerGFS environment, including agent information, job information, and overall PMC information.

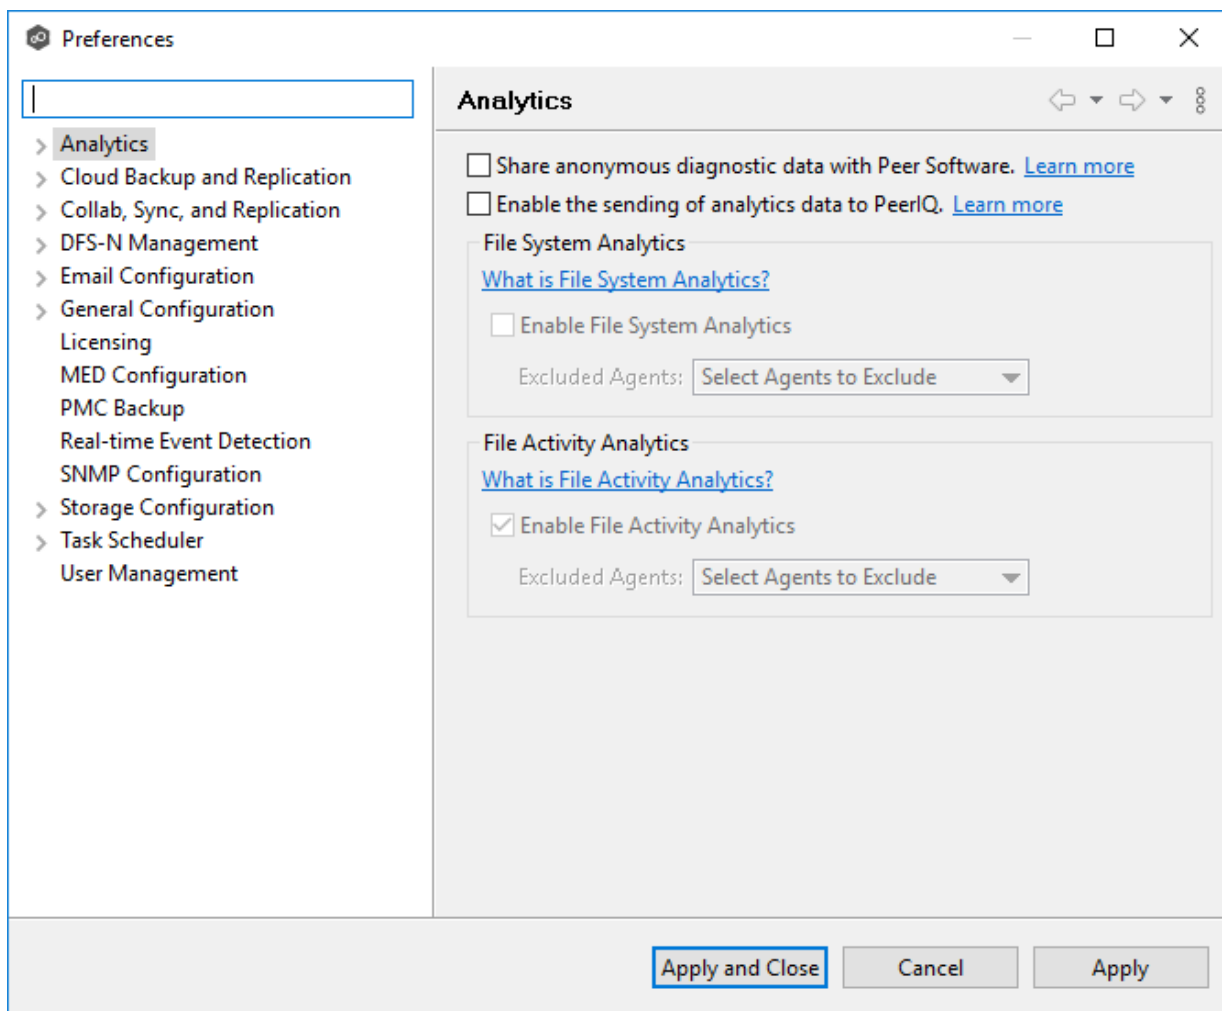
Users can toggle the Analytics settings on or off, enabling or disabling the collection and processing of analytics data according to their preferences. This allows users to control the extent to which analytics data are gathered within the environment.

The [Anonymous Diagnostic Data](#) and [PeerIQ](#) (with [File System Analytics](#) and [File Activity Analytics](#)) are activated on this page. To enable and configure [Proactive Monitoring](#), navigate to the Proactive Monitoring subpage.

For information about Analytics, see [Analytics](#) in [Advanced Topics](#).

To enable the first two Analytics features:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **Analytics** in the navigation tree.



3. Select options as needed.

Option	Description
Share anonymous diagnostic data with Peer Software.	Select this option to share anonymous diagnostic information with Peer Software. This information will help us improve PeerGFS. No customer-identifiable information is sent. More details can be found in our knowledge base.
Enable the sending of analytics data to Peer IQ.	Select this option to enable the flow of PeerGFS telemetry to PeerIQ. PeerIQ, a virtual appliance-based analytics engine, offers a set of dashboards for monitoring the health and performance of PeerGFS and the replication environment.
Enable File System Analytics	<p>Check Enable File System Analytics to set up scheduled weekly file server scans and to send the results of those scans to PeerIQ.</p> <p>To exclude specific agents from participating in file system scans and data transmission, select them from the Excluded Agents menu.</p>
Enable File Activity Analytics	<p>Check Enable File Activity Analytics to turn on the collection of real-time-based summary statistics and to send these statistics to PeerIQ.</p> <p>To exclude specific agents from participating in summary statistic collection and data transmission, select them from the Excluded Agents menu.</p>

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

Proactive Monitoring Preferences

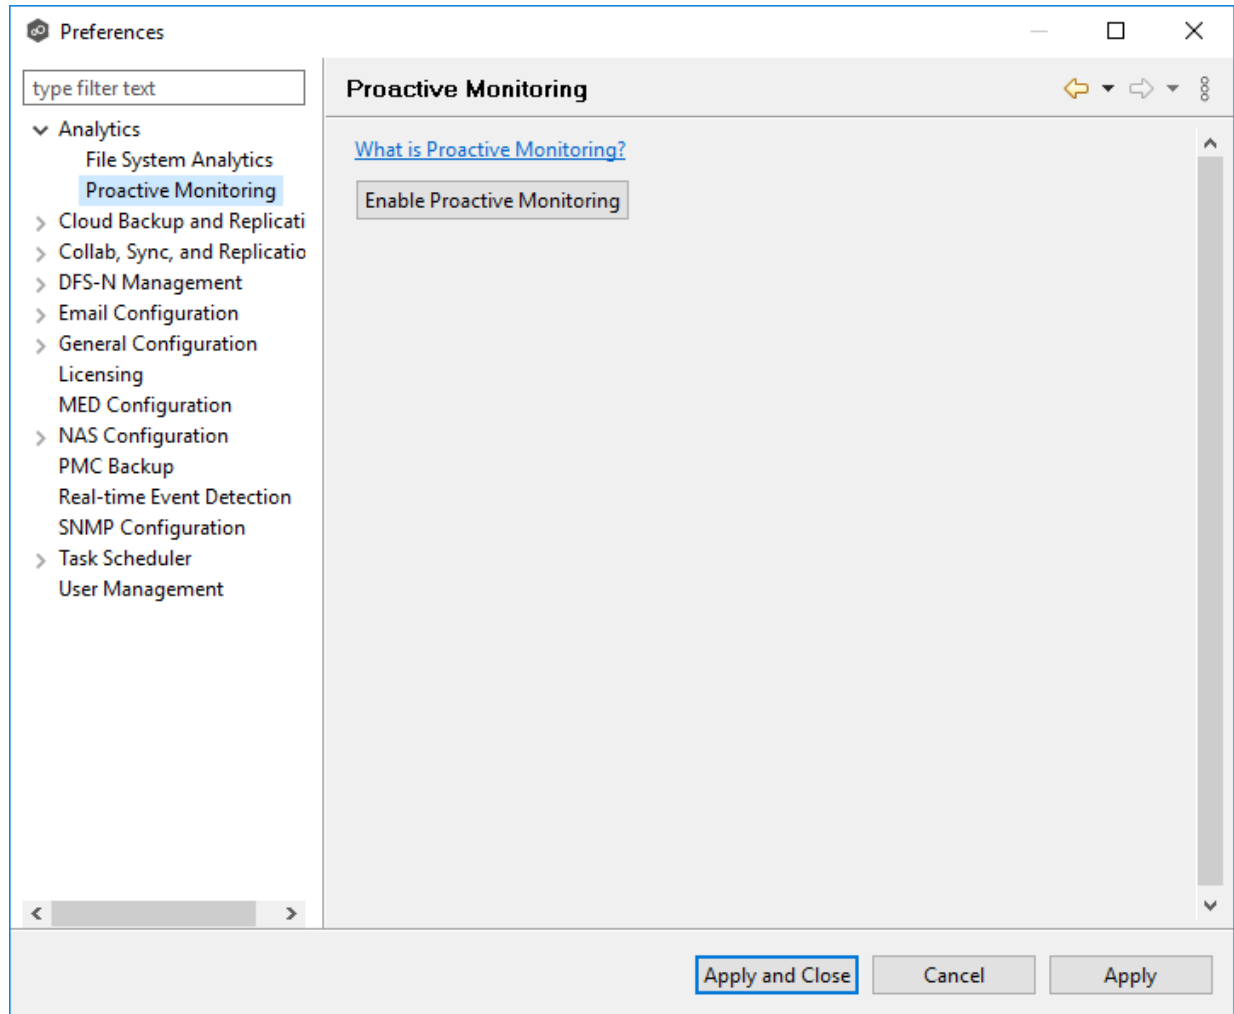
Proactive Monitoring, available to customers with dedicated Technical Account Managers (TAMs), provides comprehensive insights into the health and performance of the PeerGFS environment. For more information about Proactive Monitoring, see [Analytics](#) in [Advanced Topics](#).

To set up Proactive Monitoring:

1. Select **Preferences** from the **Tools** menu, .

2. Select **Analytics** in the navigation tree, and then select **Proactive Monitoring**.

The **Proactive Monitoring** page is displayed.



3. Click the **Enable Proactive Monitoring** button.
4. Continue configuring Proactive Monitoring on the [General Configuration](#) page.

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

1. Enter your Subscription ID if the field is not auto-filled. Contact Peer Software if you do not know your Subscription ID. Once entered, it cannot be changed.

The screenshot shows a window titled "Set Up Proactive Monitoring" with a "General Configuration" tab selected. The window contains a sidebar with navigation options: "General Configuration", "Telemetry Options", "Agent Locations", "Health Checker Setup", and "Confirmation". The main area is titled "General Configuration" and contains the instruction "Enter information required to enable Proactive Monitoring." Below this, there are three input fields: "Subscription ID" (with a blurred value and a link "What is my Subscription ID?"), "Environment Name" (with the value "DGPMC1"), and "Upload Interval" (with a spinner set to "30" and the unit "minutes"). At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. If the **Environment Name** field's auto-filled value doesn't match the name of the server or environment where the PMC server is installed, enter the correct name.

Note: Changing the name here will also change the name in the **Environment Name** field in the [General Configuration](#) preferences page.

3. In the **Upload Interval** field, specify the number of minutes between uploads of data to Proactive Monitoring.

The default is 30 minutes; the minimum is 15 minutes, and the maximum is 180 minutes.

4. Click **Next** to proceed to the [Telemetry Options](#) page.

The **Telemetry Options** page allows you to select detailed telemetry data to upload to be used by the Peer Software Technical Account Management team if you signed up for Proactive Monitoring.

The detailed telemetry data are divided into three categories:

- PMC Details
- Agent Details
- Job Details

1. Select the data to be uploaded:

Set Up Proactive Monitoring

Telemetry Options

Review and select the types of information to upload for Proactive Monitoring.

General Configuration
Telemetry Options
 Agent Locations
 Health Checker Setup
 Confirmation

[What types of data are uploaded to Proactive Monitoring?](#)

PMC Details:

- Include IP Information
- Include Statistical Information

Agent Details:

- Include IP Information
- Include Agent Names
- Include Agent Locations
- Include Storage Information

Job Details:

- Include Job Names
- Include MED Alerts

< Back Next > Finish Cancel

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

Category	Standard Data Upload	Include Options
PMC Details	Includes details about this PMC deployment. Details include service memory consumption, replication backlog, quarantines, license consumption, and watch set size.	<ul style="list-style-type: none"> • IP Information - The IP address of the server that this PMC is installed on. • Statistical information - In-depth statistics about the queues and performance of PeerGFS's replication engine.
Agent Details	Includes the details about the Agents that are connected to this PMC. Details include service and server memory consumption, replication throughput, uptime, operating system, and disconnect counts.	<ul style="list-style-type: none"> • IP information - The IP addresses of the servers that the Agents are installed on. • Agent Names - The names assigned to the Agents (typically the name of each Agent's Windows Server). If this option is not checked, random strings will be used in the Proactive Monitoring system to represent each Agent. • Agent Locations - The locations (the latitude, longitude, city, state, and country) of the Agents. You can enter the locations while running this wizard or in the Agent Configuration dialog later. If you choose to include Agent location data, you will be prompted to enter Agent location information on the next page of this wizard (the Agent Locations page of this wizard). No location information is automatically determined—it must be manually entered. • Storage Information - Information specific to the storage platforms that each Agent is managing, including available and used disk space.
Job Details	Includes the details about the file collaboration, synchronization and/or replication jobs configured within this PMC. Per-job details include	<ul style="list-style-type: none"> • Job Names - The names of the file collaboration, synchronization, and replication jobs configured in this PMC. If this option is not checked, random strings will be used in the Proactive Monitoring system to represent each job.
		<ul style="list-style-type: none"> • MED Alerts - If MED alerts are enabled in this PMC, this option will include any alerts for use in the Proactive Monitoring system.

2. Click **Next**.

If you selected the **Include Agent Locations** option, the [Agent Locations](#) page appears; otherwise the [Set up Health Checkup](#) page appears.

This step is optional.

The **Agent Locations** page allows you to set location details for each Agent. The **Agent Locations** page appears only if you selected the **Include Agent Locations** option on the previous wizard page. If an Agent's location has already been set through its [Agent Configuration](#), those values will automatically appear on this page.

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

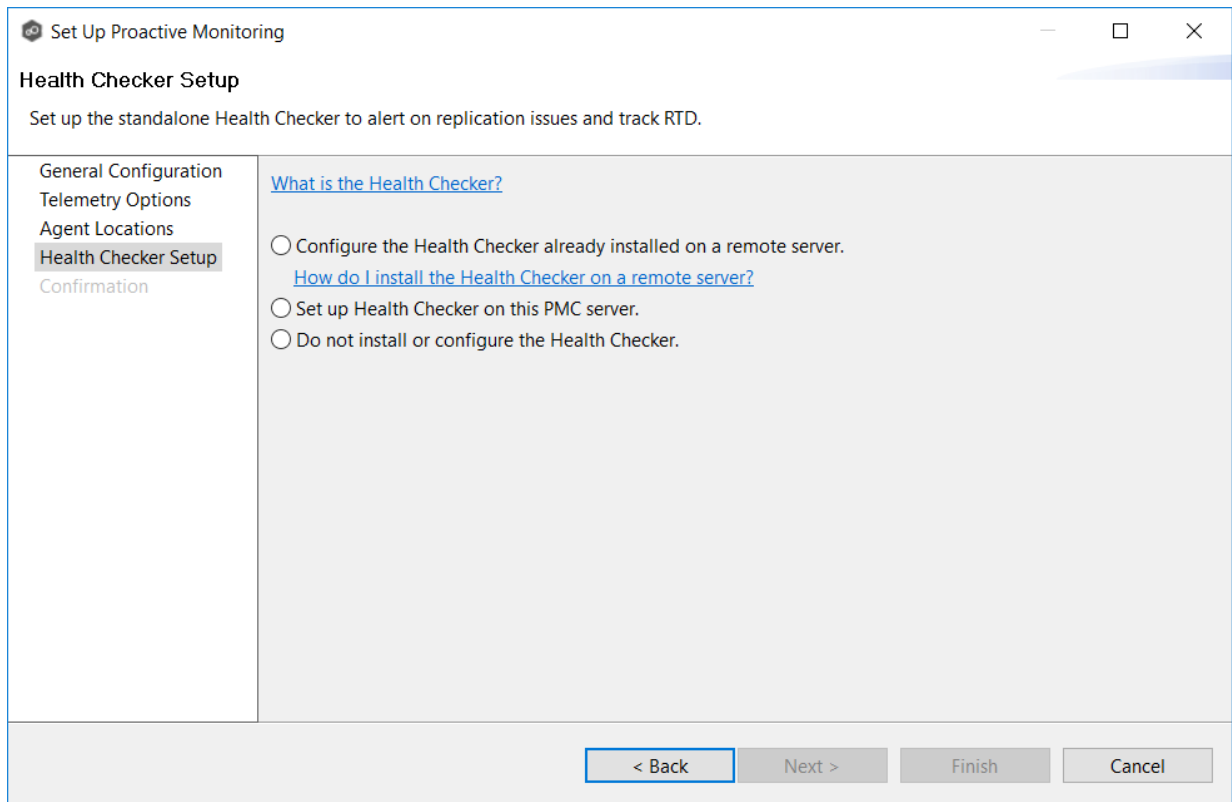
1. For each Agent that you want to set location details, enter the following information in the appropriate fields:
 - **Location** - Enter the city, state, and country of where the Agent server is installed.
 - **Longitude** - Enter the longitude of where the Agent server is installed.
 - **Latitude** - Enter the longitude of where the Agent server is installed.

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

are running on). For full functionality, we recommend installing it externally. If installed locally, you won't receive failure alerts if the PMC server goes down.

To set up Health Checker:

1. Select a setup option.



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

2. Click **Next**.
3. Your next step depends on the option you selected in Step 1:

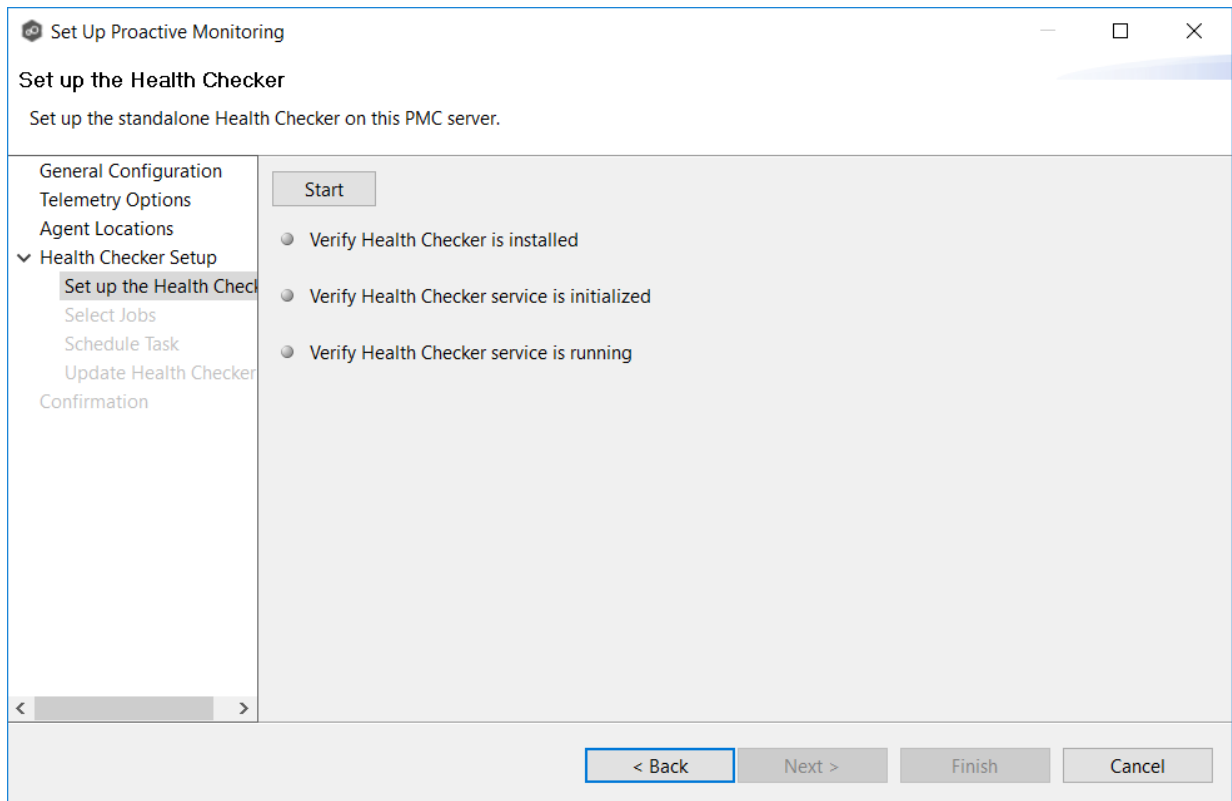
If you selected this option	Continue on this wizard page
Configure the Health Checker already installed on a remote server.	Select Jobs
Set up Health Checker on this PMC server.	Set up the Health Checker
Do not install or configure the Health Checker.	Confirmation

Set Up the Health Checker

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

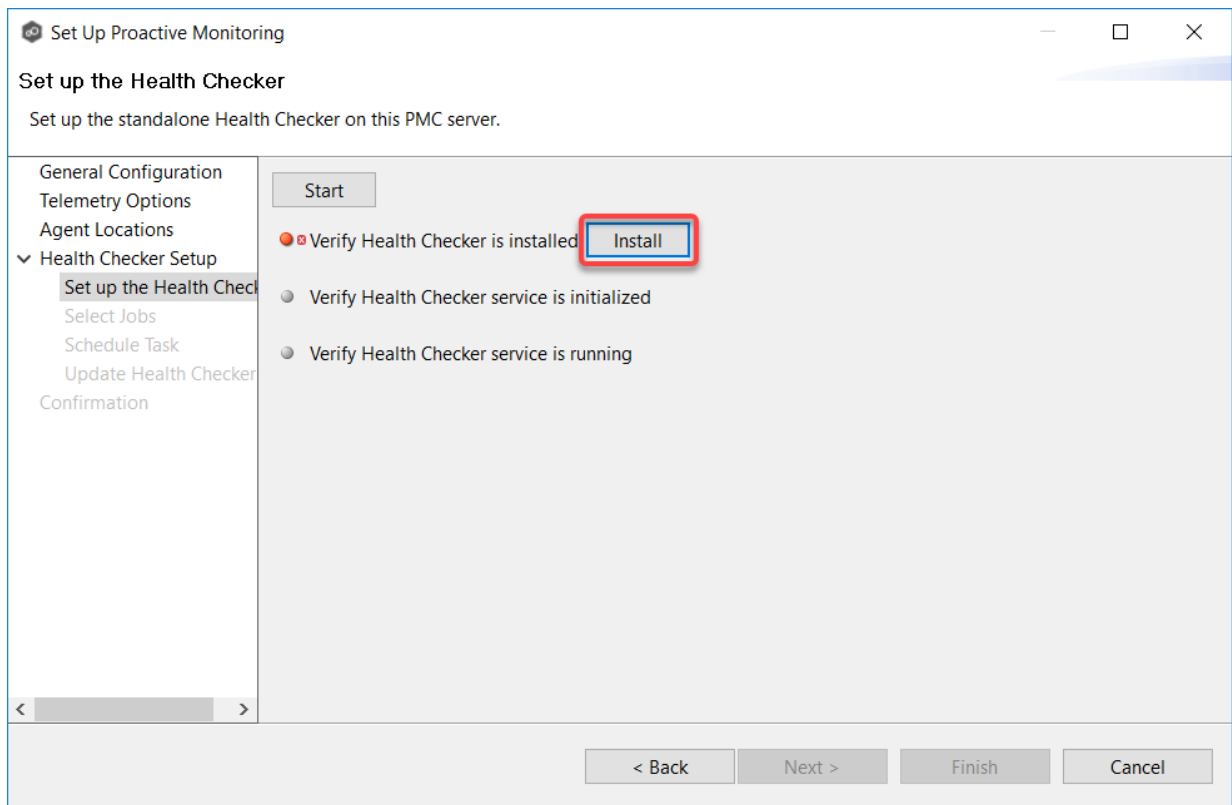
From this page, you install Health Checker. If you have previously installed Health Checker on the PMC server, it will verify that Health Checker was successfully installed and initialized and is currently running.

1. Click the **Start** button.

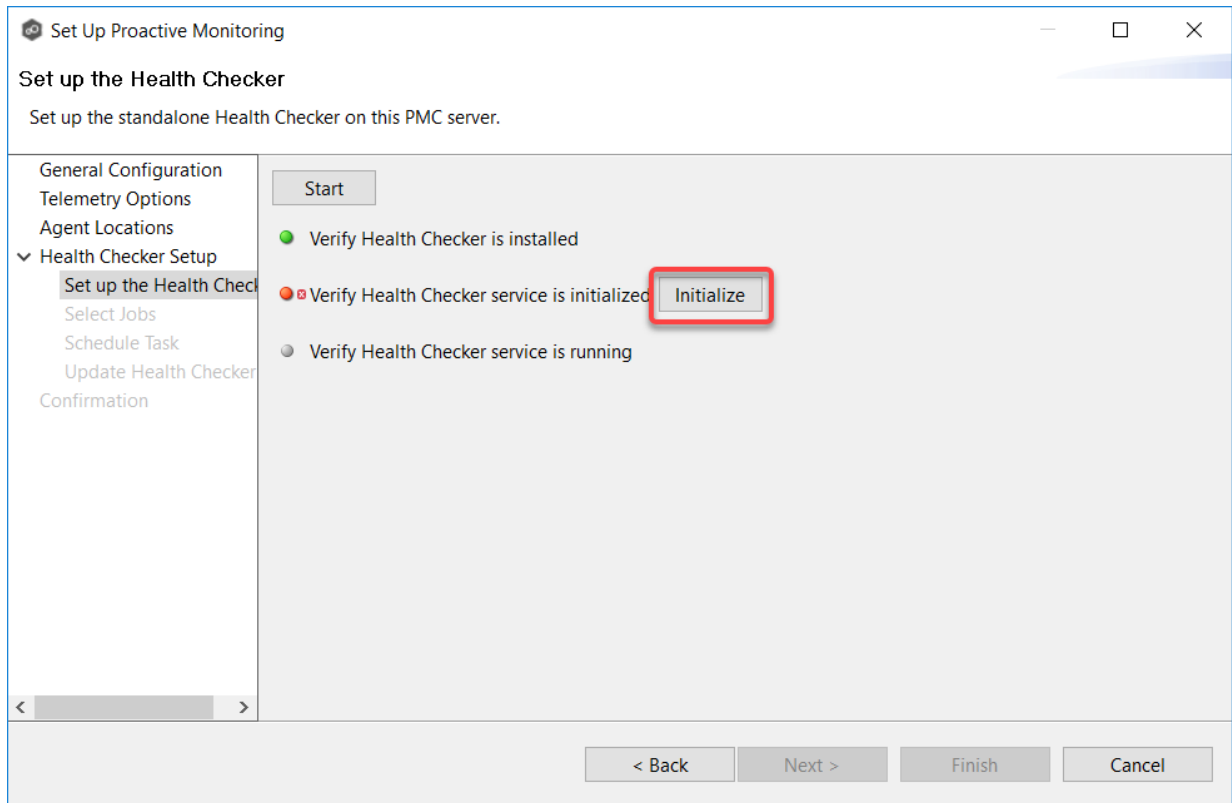


As the wizard performs the setup process, it communicates results via colored dots. A green dot indicates that successful completion of that setup stage. A red dot indicates that action is needed.

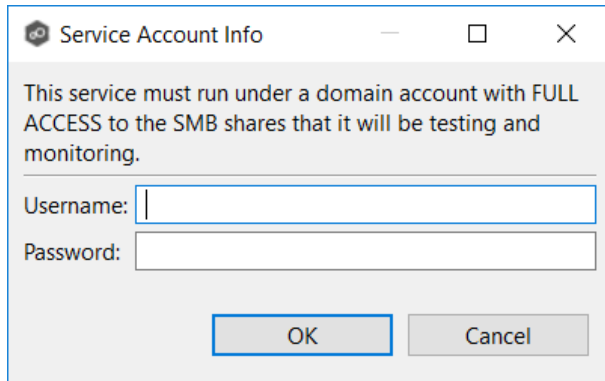
2. The wizard first checks to see if the Health Checker has already been installed. If the dot next to **Verify Health Checker is installed** turns red, Health Checker has not yet been installed. Click the **Install** button that appears. It runs a silent packaged installer for the Health Checker. Click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



- The wizard next checks to see if the Health Checker service has been initialized. If the dot next to **Verify Health Checker service is initialized** turns red, the Health Checker service has not yet been initialized. Click the **Initialize** button that appears.

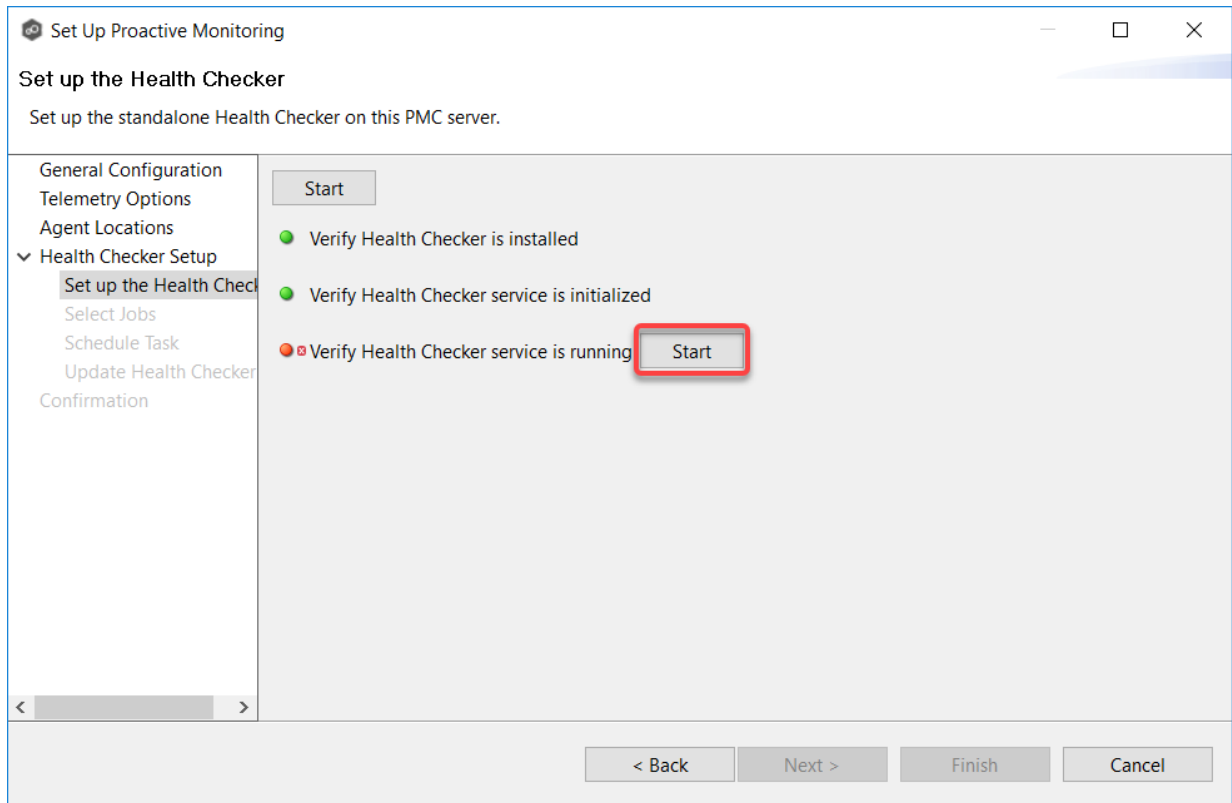


In the **Service Account Info dialog**, enter the user name and password for the service, and then click **OK**.

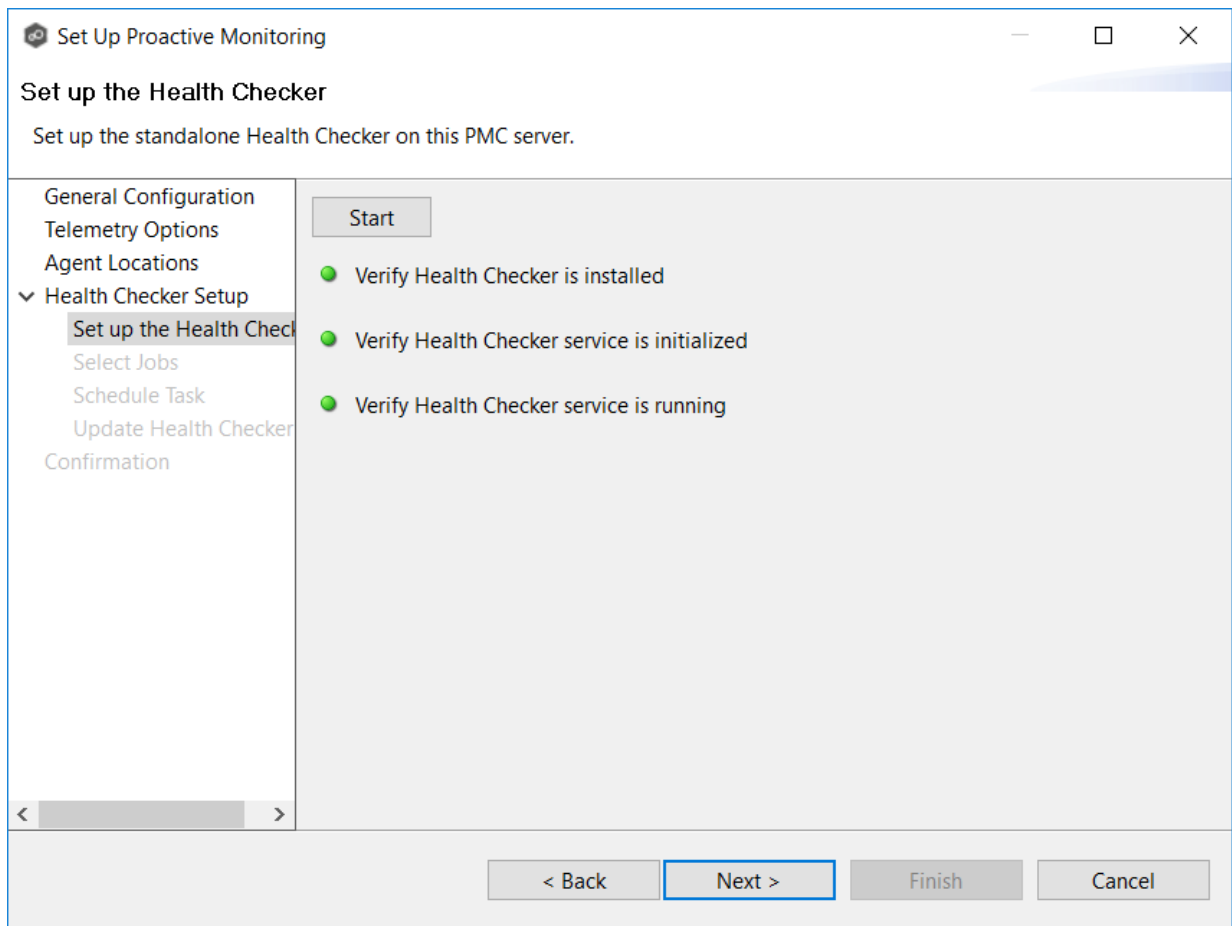


Click **OK** in the **Success** message that appears.

4. The wizard next checks to see if the Health Checker Service is running. If the dot next to **Verify Health Checker service is running** turns red, click the **Start** button that appears. Click **OK** in the **Success** message that appears.



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



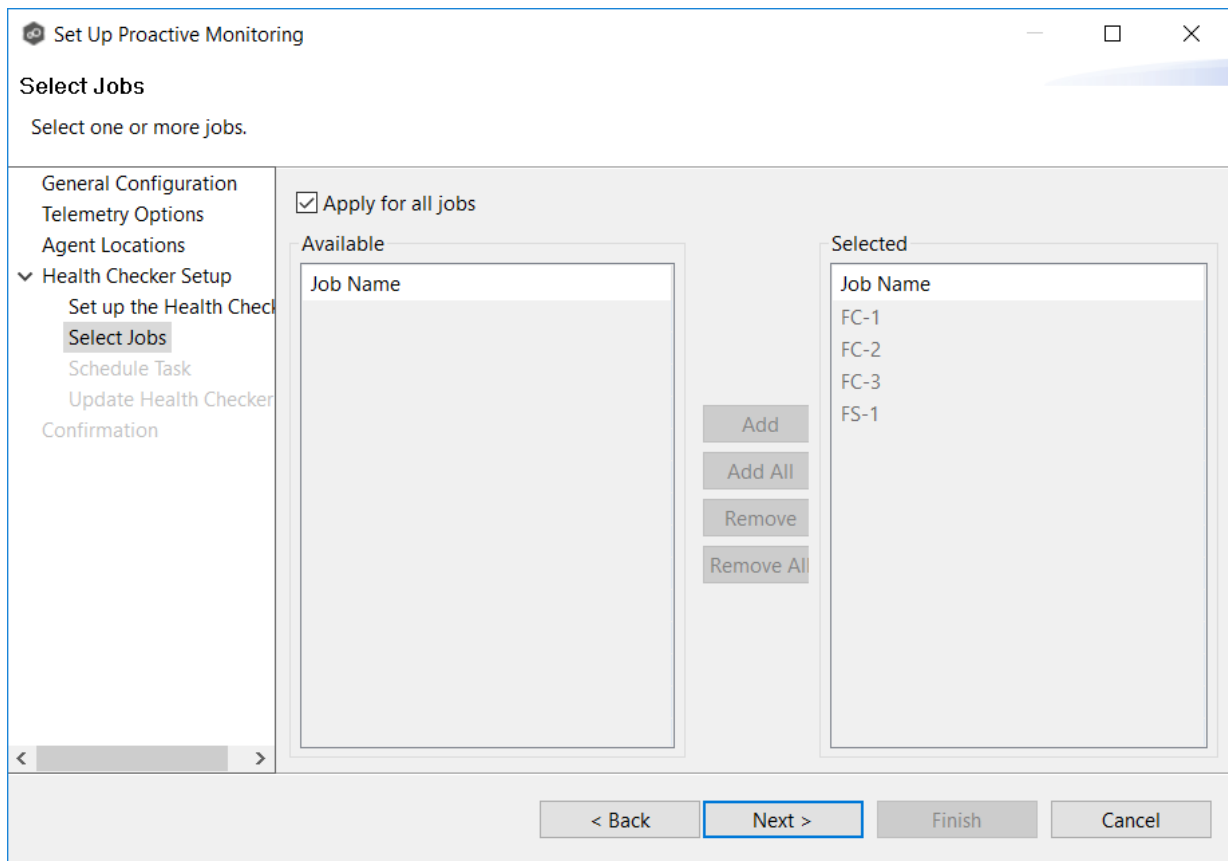
The [Select Jobs](#) page appears.

Select Jobs

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

To select jobs:

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



2. Click **Next**.

The [Update Health Checker](#) page appears.

Schedule Task

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

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

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

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

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

4. Click **Next**.

The [Configure External Health Checker](#) page appears.

Update Health Checker

The **Update Health Checker** page allows you to specify the criteria for the jobs and shares that Health Checker should monitor.

To configure the Health Checker:

1. If you installed Health Checker on an external server, enter the path in UNC format to Health Checker **workspace** folder in the **Path to Health Checker Configuration** field.

The Health Checker **workspace** is a subfolder of the Health Checker installation folder. If Health Checker is installed on the PMC server, the path is automatically detected and filled in.

2. In the **SMB Username** field, enter the user name if the server hosting the Health Checker requires account credentials. In most cases, a locally installed Health Checker will not need a user name.
3. In the **SMB Password** field, enter the password if the server hosting the Health Checker requires account credentials. In most cases, a locally installed Health Checker will not need a password.
4. In the **Extensions to Test** field, enter the extensions for the file types that you want Health Checker to monitor. Separate the extensions with a semicolon.
5. Select the **Include Participants that are inactive** checkbox if you want the Health Checker to monitor inactive participants.
6. In the **Monitor jobs that are** section, select which job states should be included in Health Checker monitoring.

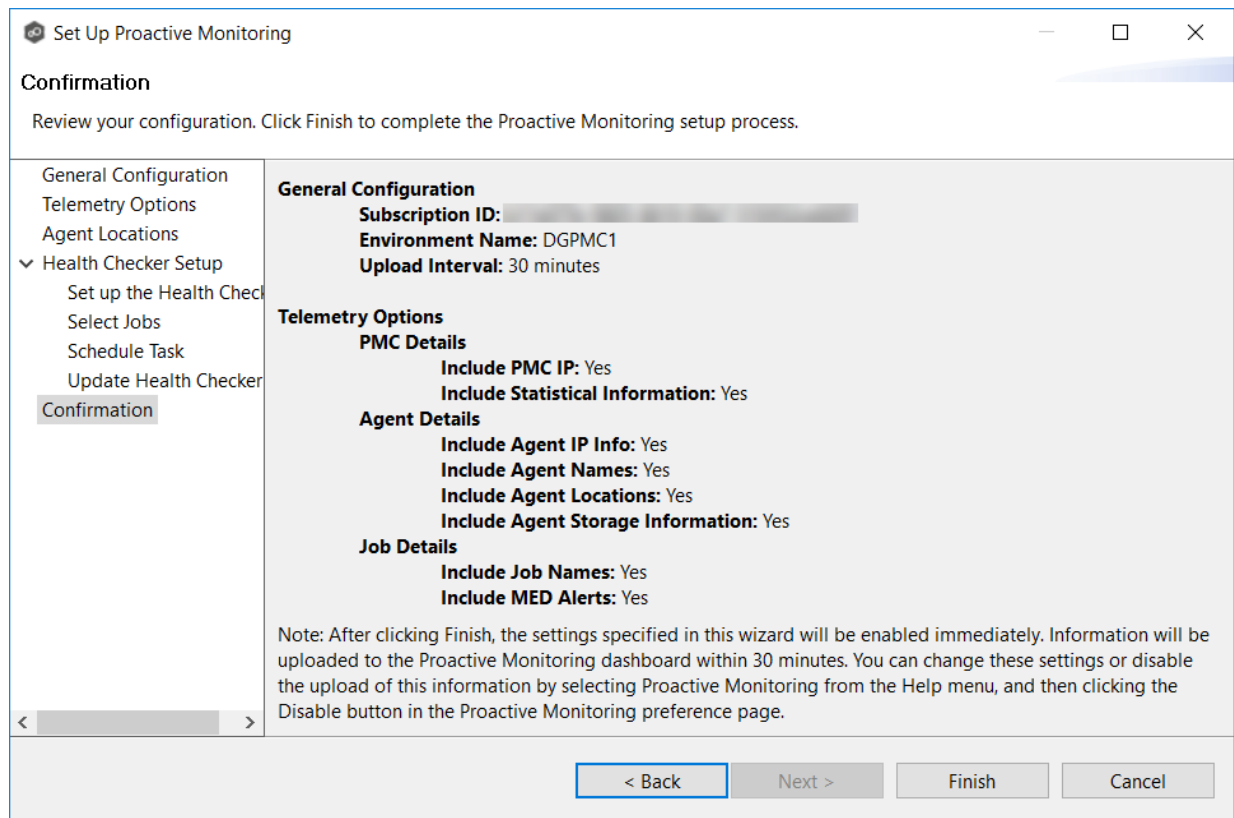
We recommend checking all but the **Stopped** category. The **Stopped** category includes jobs that are stopped for any reason other than a quorum lost. For example, if you have manually stopped a job, you may not want Health Checker monitoring it.

7. Click **Next**.

The [Confirmation](#) page appears.

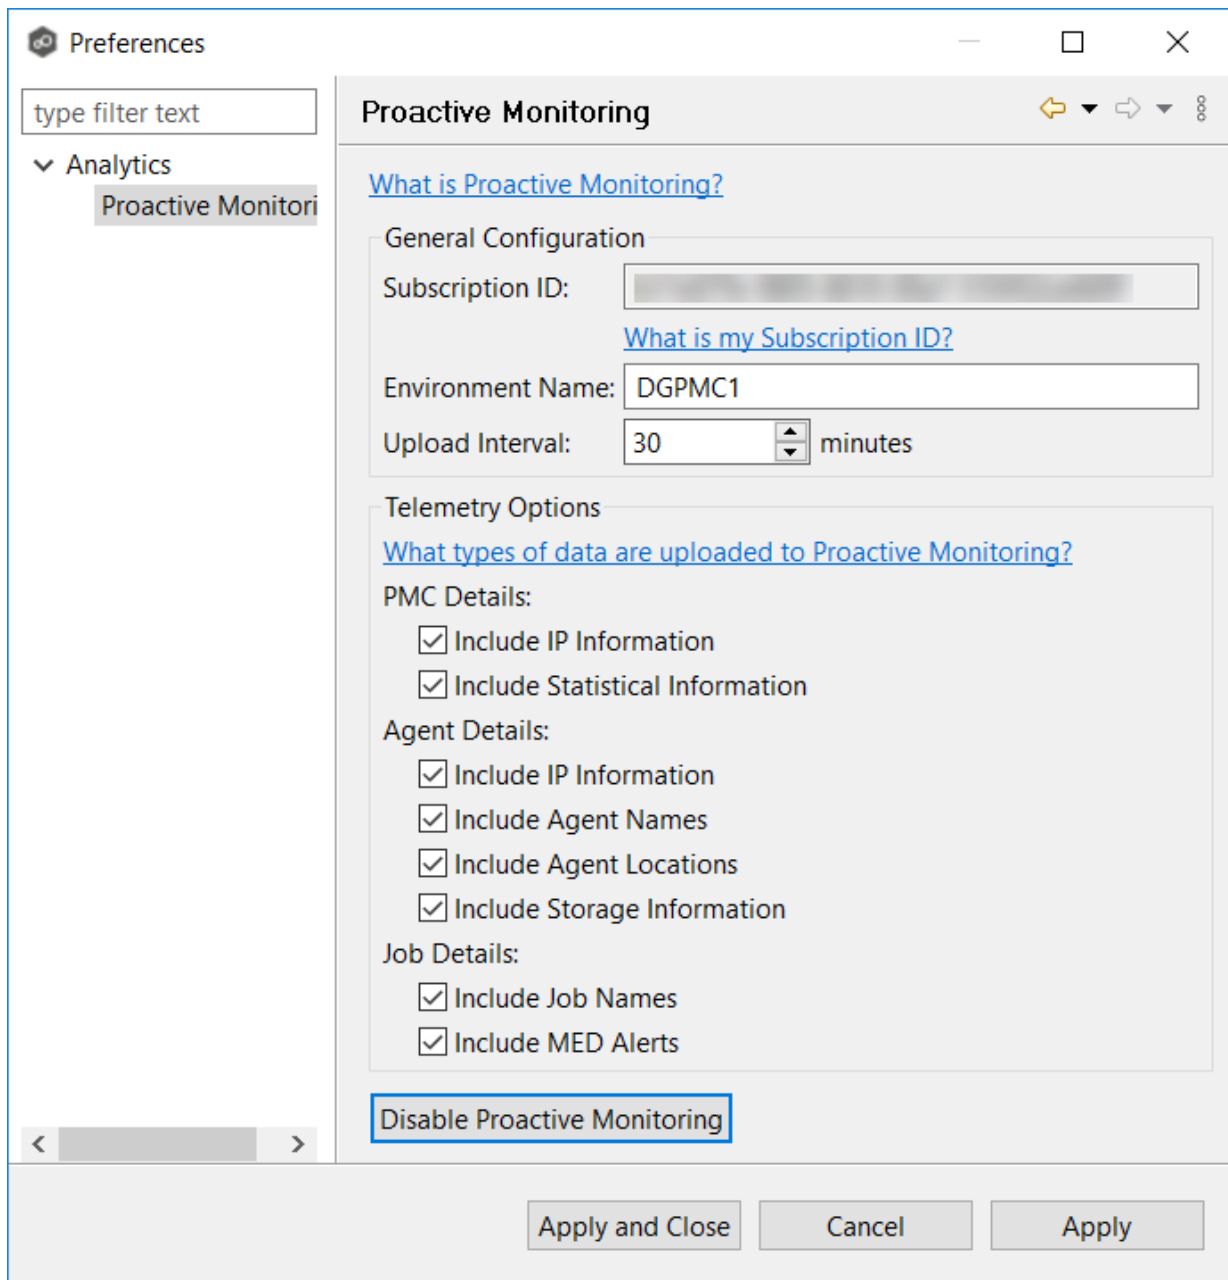
The **Confirmation** page displays a summary of the settings you have selected.

1. Review the configuration.



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

The settings are displayed in the Proactive Monitoring preferences page.



Note: If you modified the location of an Agent, you will be prompted to restart the Agent. Click **Restart Later** if you do not want to restart the Agent services because you have other jobs running; otherwise, click **Restart Now**.

3. Click **Apply and Close**.
4. Notify your Peer Software Technical Account Manager that the setup of Proactive Monitoring is complete.

Effects of disabling Proactive Monitoring

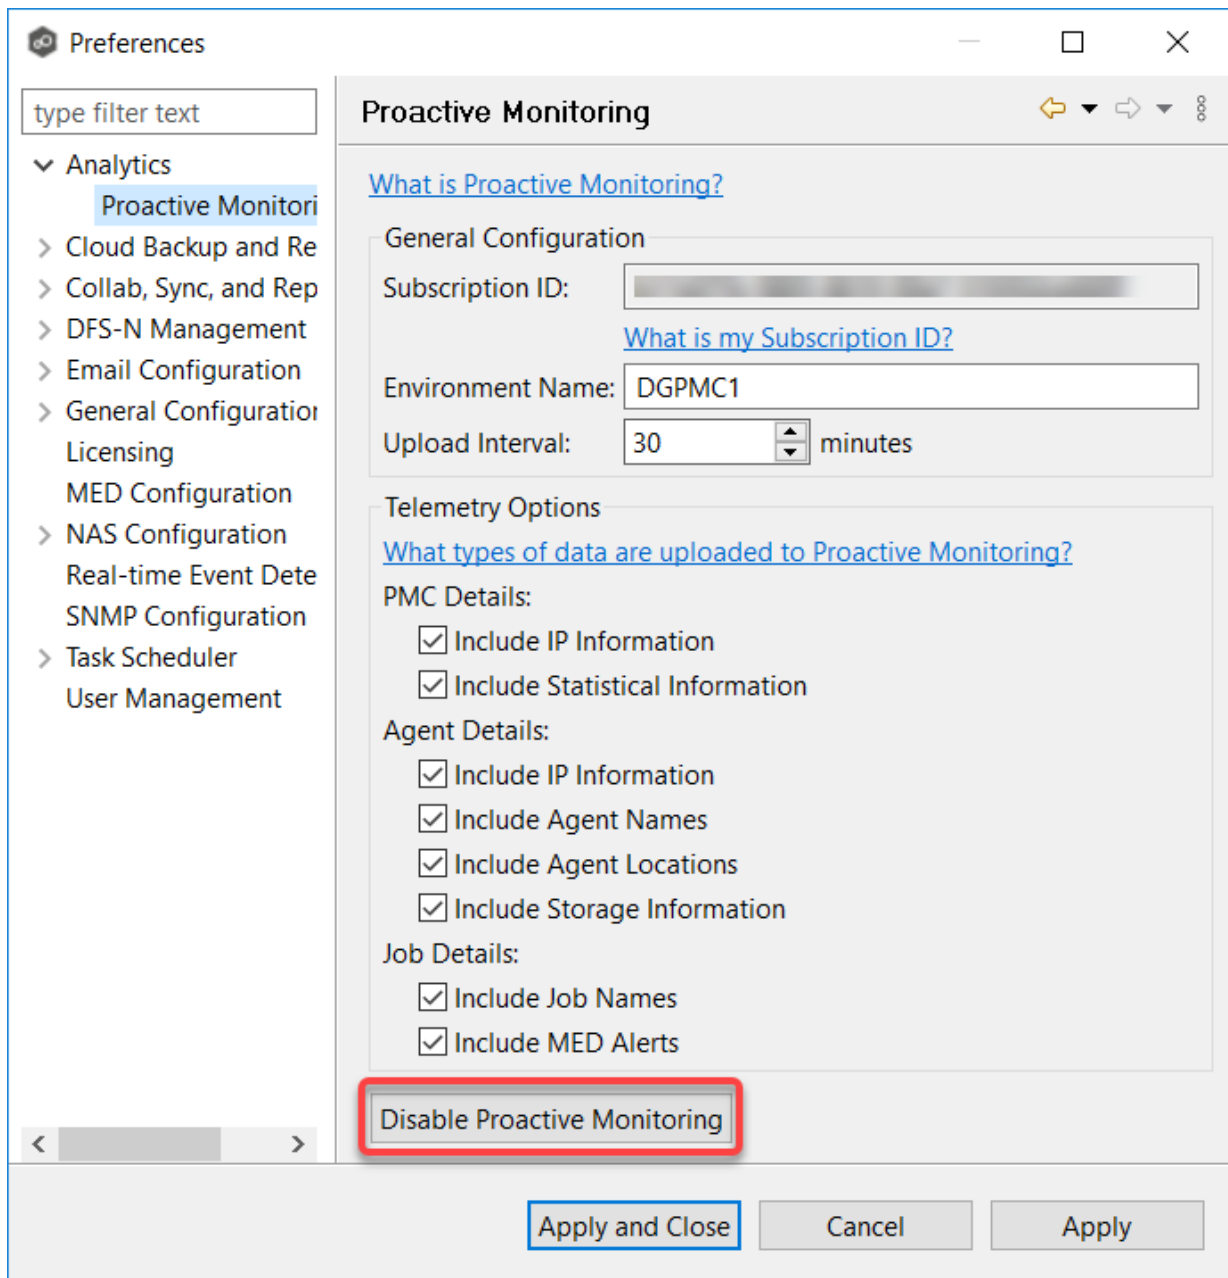
- Your Peer Software Technical Account Manager will no longer be able to check on the status of your PeerGFS environment.
- Data will no longer be uploaded to Peer Software. If disabled, you can request Peer Software to delete your data.

To disable Proactive Monitoring:

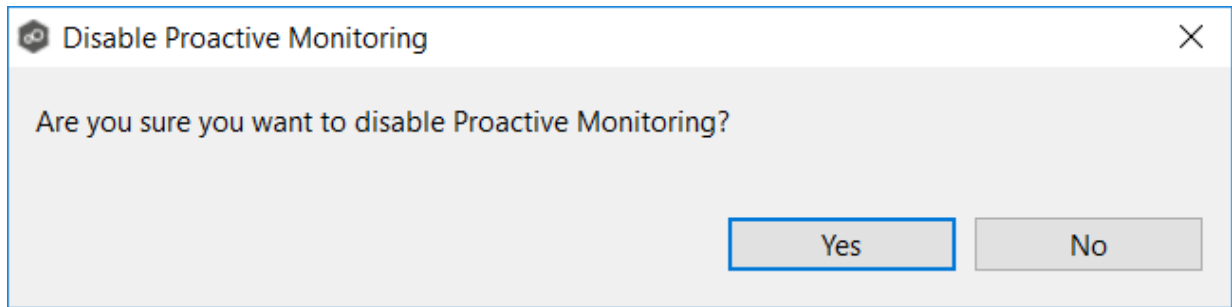
1. Select **Preferences** from the **Window** menu.

The **Preferences** dialog appears.

2. In the navigation tree, select **Analytics**, and then select **Proactive Monitoring**.

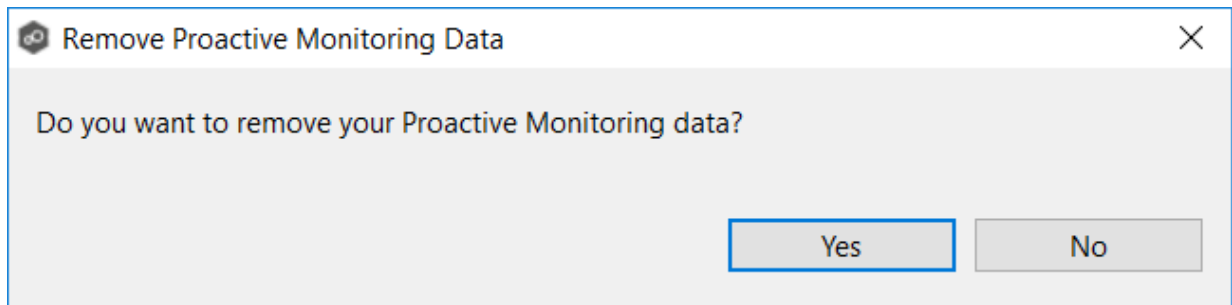


3. Click the **Disable Proactive Monitoring** button.
4. In the confirmation dialog, click **Yes** to confirm.

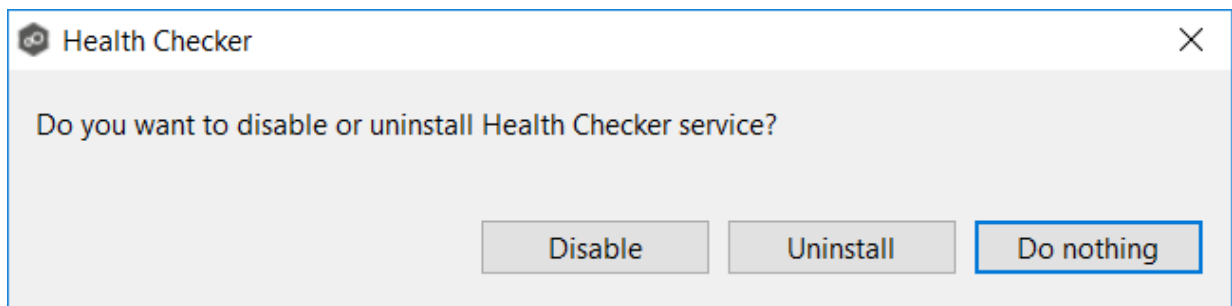


5. In the **Remove Proactive Monitoring Data** dialog, click **Yes** if you want your data deleted.

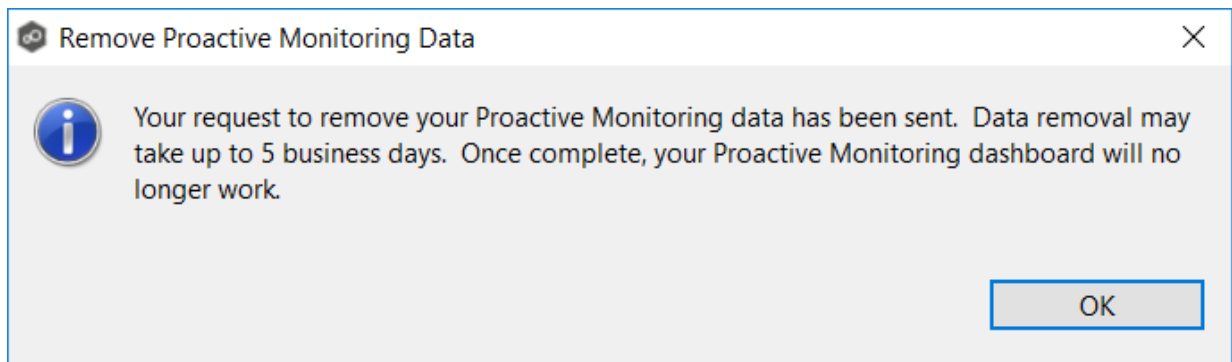
Peer Software will delete your data from Microsoft Azure. This action removes both current and historical data, which will be irrecoverable



6. If Health Checker was installed locally, choose an option in the **Health Checker** dialog:



- **Disable:** Stops the Health Checker Service and prevents it from starting automatically..
 - **Uninstall:** Removes the Health Checker service and all files.
 - **Do Nothing:** Keeps Health Checker installed and running.
7. Click **OK** in the dialog that appears if you chose to delete Proactive Monitoring data.



The **Preferences** page reappears.

8. Click **Apply and Close** or **Apply**.

Re-enabling Proactive Monitoring

To re-enable Proactive Monitoring, rerun the **Set up Proactive Monitoring** wizard. Click the **Enable Proactive Monitoring** button in the [Proactive Monitoring preferences](#) page.

Cloud Backup and Replication Job Preferences

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

You can modify the following Cloud Backup and Replication settings:

- [Cloud Backup and Replication](#)
- [Database Connections](#)
- [Destination Credentials](#)
- [Email Alerts](#)
- [File Retries and Source Snapshots](#)
- [File and Folder Filters](#)

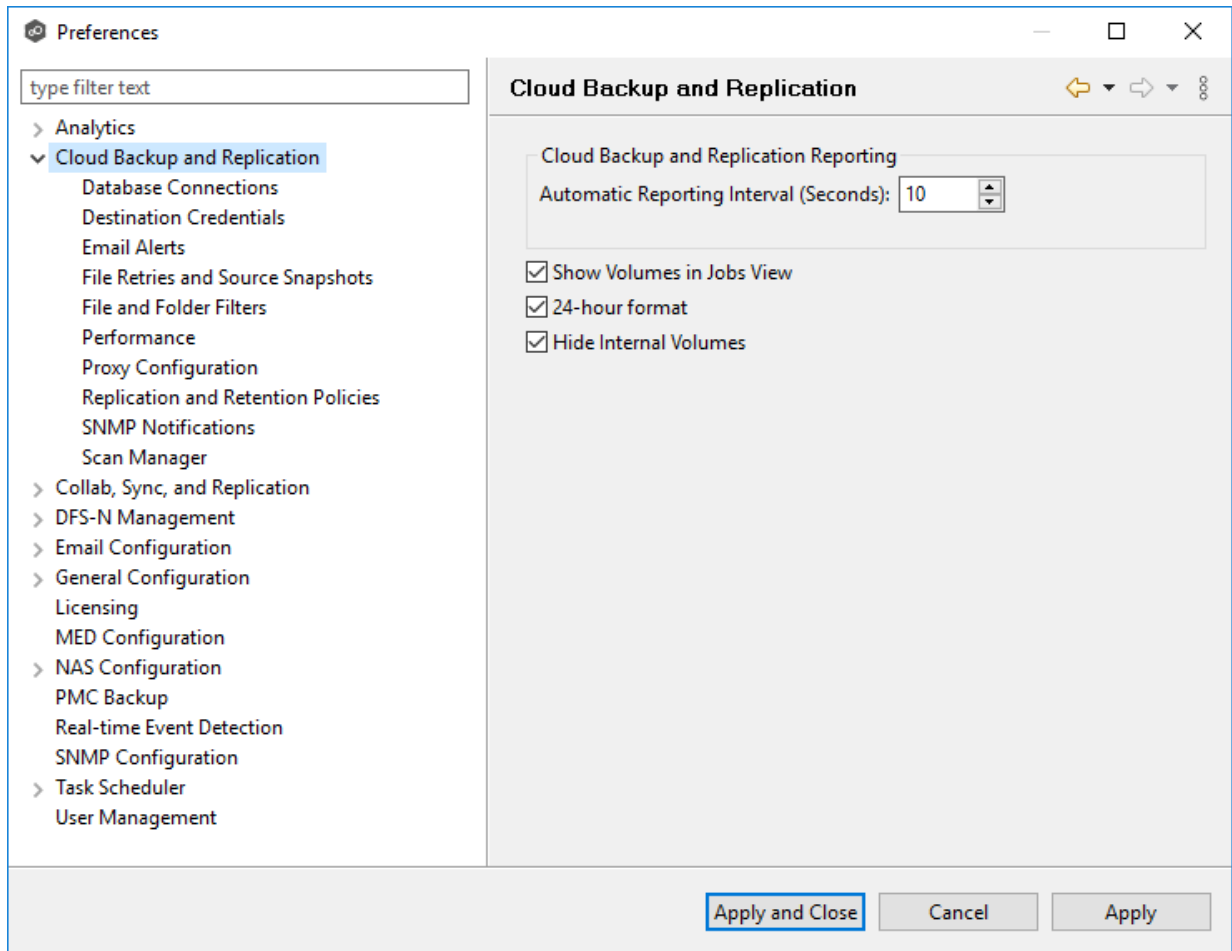
- [Performance](#)
- [Proxy Configuration](#)
- [Replication and Retention Policies](#)
- [SNMP Notifications](#)
- [Scan Manager](#)

Cloud Backup and Replication

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

To modify these settings:

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



3. Modify the settings as needed.

Automatic Reporting Interval (Seconds)	Each Peer Agent automatically reports its statistics to Peer Management Center at regular intervals. Select the number of seconds between these intervals. The default is 10 seconds.
Show Volumes in Jobs View	Select this checkbox if you want volumes to be displayed in the Jobs view.
24-hour format	Select this checkbox if you want times to be displayed in a 24-hour format rather than a 12-hour format.
Hide Internal Volumes	Select this checkbox if you don't want internal volumes displayed when choosing which volumes to replicate.

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

Create Database Connection

Configure connection information for MS SQL Server.

*Database Connection Name:

*Management Agent:

*DB Host Name:

Port:

Instance Name:

*Database Name:

Authentication: Integrated Credentials

Username:

Password:

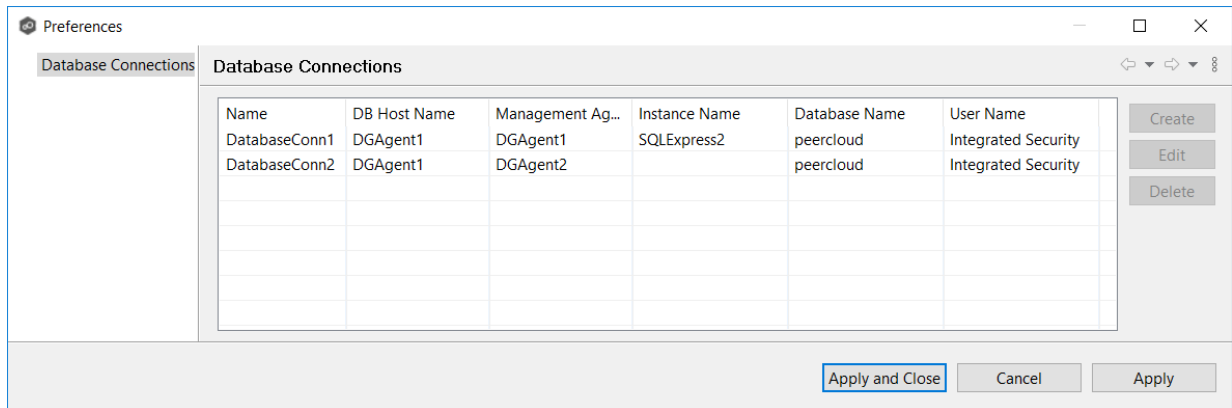
4. Enter the required values.

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

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

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

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



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

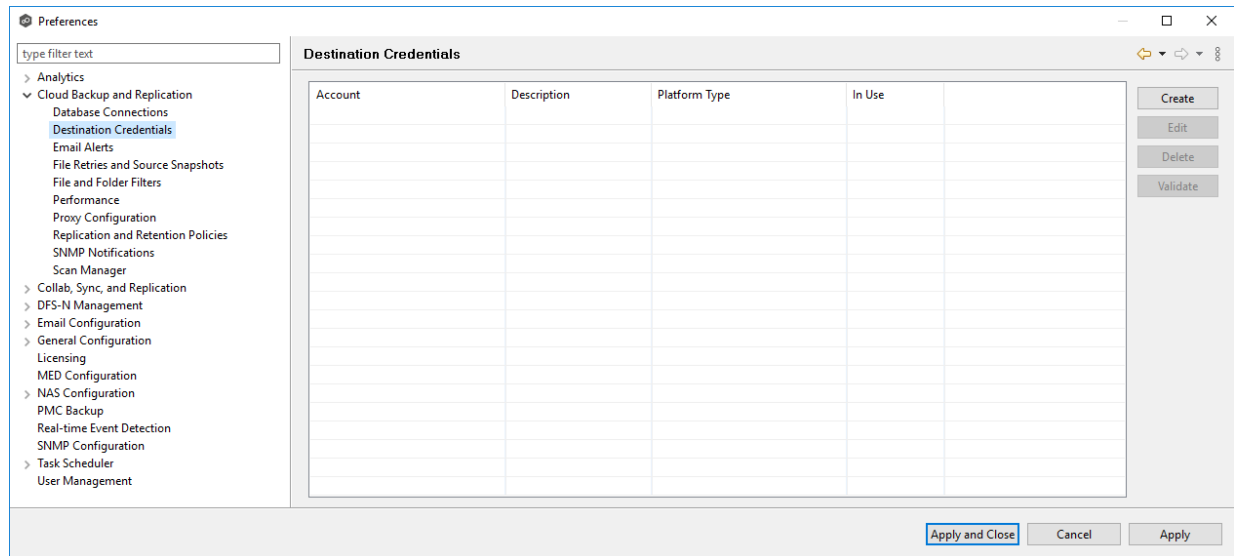
Destination Credentials

When you create a Cloud Backup and Replication job, you can select existing destination storage account credentials to apply to the job or you can create new credentials and apply them to the job. This [Preferences](#) page lists the existing credentials. From this page, you can view, create, edit, and delete credentials. However, you cannot edit or delete credentials while they are applied to a job.

To create new destination storage account credentials:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Destination Credentials**.

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



3. Click the **Create** button.

The **Storage Account** dialog appears.

Create Destination Credentials

Platform:

- Microsoft Azure Blob Storage
- Amazon S3
- NetApp StorageGRID
- Nutanix Objects
- S3 Compatible
- File System

*Description:

*Account:

*Shared Key: Show Key

*Endpoint Type: ▾

Use SSL

4. Enter the required values. For information about the required values, see [Step 8: Destination Credentials](#) in [Creating a Cloud Backup and Replication Job](#).
5. Click **OK**.

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

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

Create Email Alert

Name:

Temporarily Disable

Event Types

Job Start Job Stop Job Failure Participant Failure

System Event Malicious Event

Report Types

Scan Destination Snapshot

Recipients

Enter email, contact, or distribution list:

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.

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

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

- Select the report types to be sent.

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

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

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

- Click **OK**.

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

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

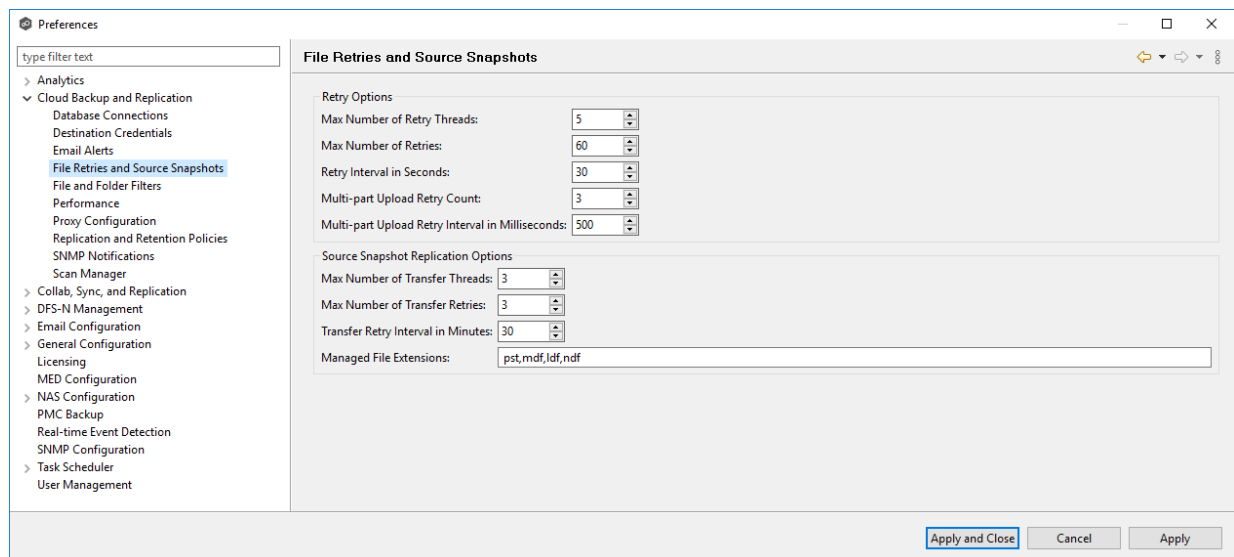
File Retries and Source Snapshots

This page allows you to specify two sets of options:

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

To modify these options:

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



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

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

4. Modify the **Source Snapshot Replication Options** as needed:

Option	Description
Max Number of Transfer Threads	Enter the maximum number of threads available for replicating files from a source snapshot.
Max Number	Enter the maximum number of retries to perform on a file or folder that has failed to be replicated from a source snapshot. If the

Option	Description
of Transfer Retries	number of retries is exceeded, the file or folder will be added to the Failed Events view and will need to be manually processed.
Transfer Retry Interval in Minutes	Enter the number of minutes to wait between retries of the failed replication of a file or folder from a source snapshot.
Managed File Extensions	Enter the extensions for managed files that should be read from a source snapshot.

5. Click **Apply and Close** or **Apply**.

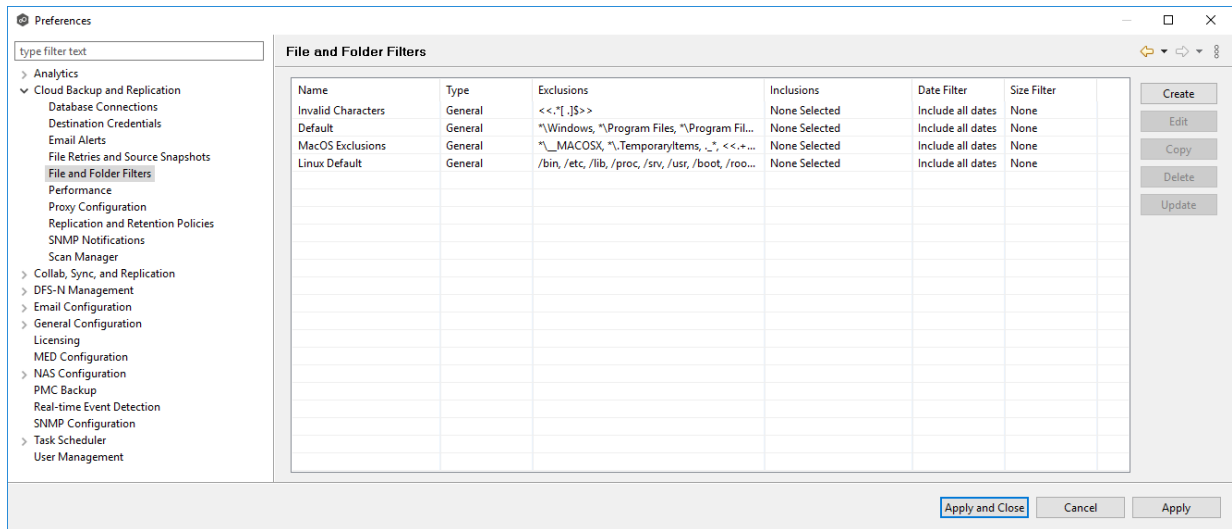
File and Folder Filters

When you create a Cloud Backup and Replication job, you can select existing file filters to apply to the job or you can create new file filters and apply them to the job. This [Preferences](#) page lists the existing file filters. From this page, you can view, create, edit, update, and delete file filters. However, you cannot edit or delete a file filter while it is applied to a job. See [File and Folder Filters](#) in the [Basic Concepts](#) section for more information about file and folder filters.

To create a file filter:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **File and Folder Filters**.

Any existing Cloud Backup and Replication file filters are listed in the **File Filters** table.



3. Click the **Create** button.

The **Create File Filter** dialog appears.

Create File Filter

Name:

Filter Type: **General** ▾

Auto Excluded
[View file types that are automatically excluded](#)

Excluded Patterns

Add **Edit** **Delete**

Included Patterns

Add **Edit** **Delete**

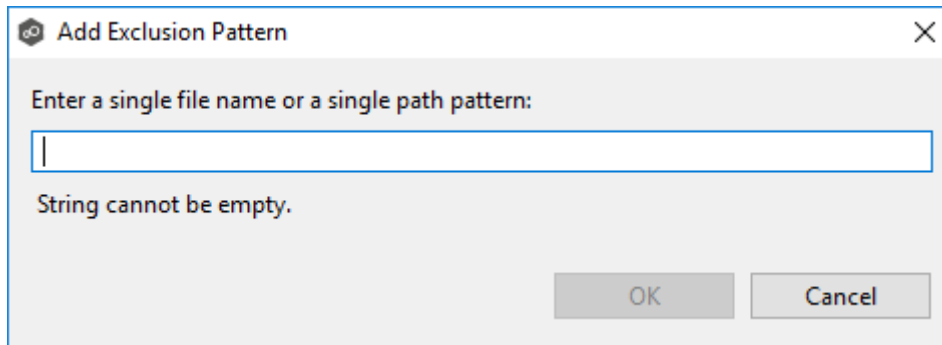
Included Last Modified Dates: **Include all dates** ▾
 days

Excluded File Sizes: **None** ▾
 bytes

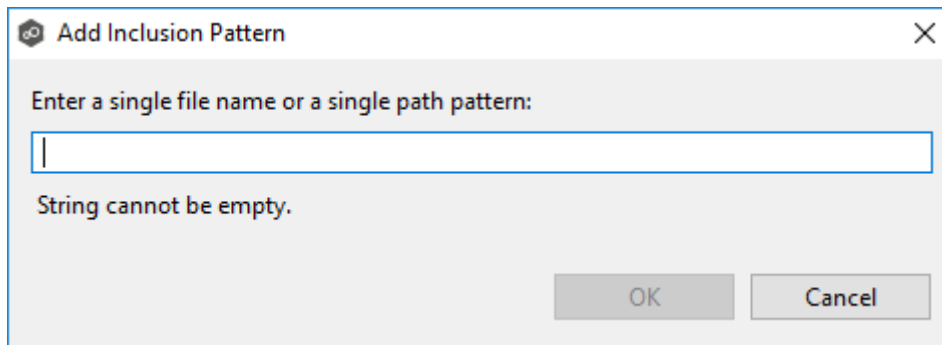
OK **Cancel**

4. Enter a unique name for the filter.
5. Select the [filter type](#).
6. (Optional) In the **Excluded Patterns** section, click the **Add** button to enter a filter pattern for files that you want excluded from the job. Repeat to add more filter patterns.

See [Defining Filter Patterns](#) for information about filters patterns.



7. (Optional) In the **Included Patterns** section, click the **Add** button to enter a filter pattern for files that you want included in the job. Repeat to add more filter patterns.



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

Note: A filter cannot use **Included Last Modified Dates** in conjunction with any other excluded or included patterns.

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

Note: A filter cannot use **Excluded File Sizes** in conjunction with any other excluded or included patterns.

10. Click **OK**.

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

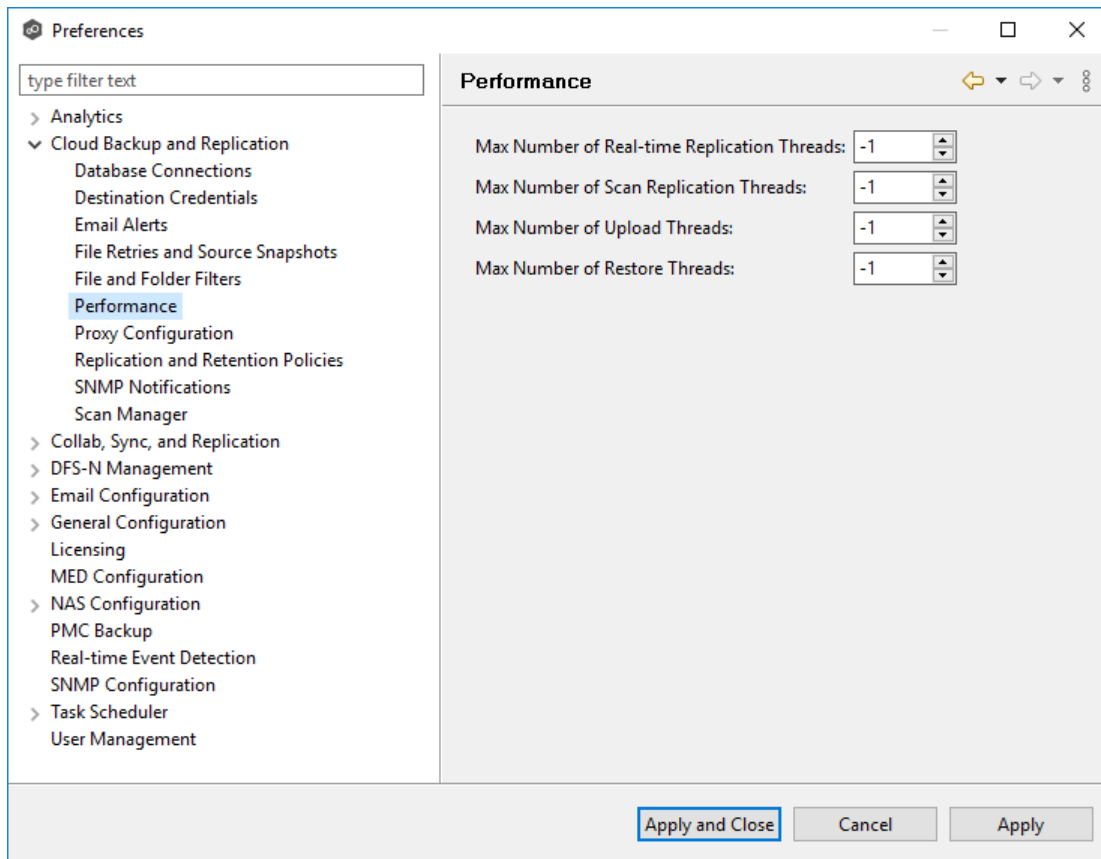
11. Click **Apply and Close** or **Apply**.

Performance

Performance settings allow you to adjust the performance of Cloud Backup and Replication jobs.

To modify the Cloud Backup and Replication performance settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Performance**.



3. Modify the settings as needed:

Setting	Description
Max Number of Real-time Replication Threads	Enter the maximum number of threads available for replicating files as they are updated in real-time on the source storage device.
Max Number of Scan Replication Threads	Enter the maximum number of threads available for replicating files during scheduled and on-demand scans of the source storage device.

Setting	Description
Max Number of Upload Threads	Enter the maximum number of threads available for uploading files to the destination storage device.
Max Number of Restore Threads	Enter the maximum number of threads available for restoring from the destination storage device.

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

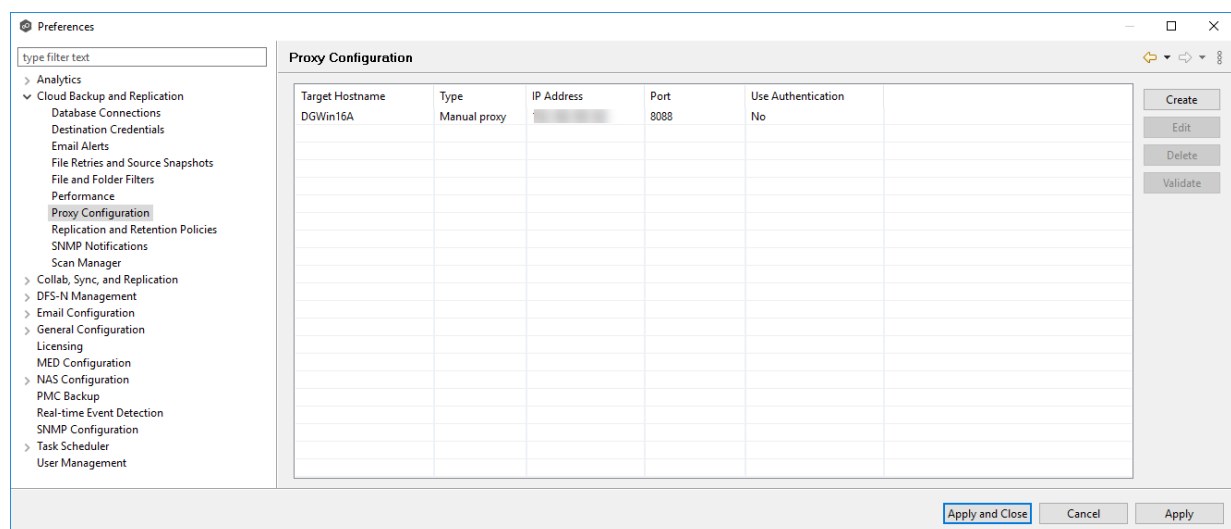
Proxy Configuration

The **Proxy Configuration** page offers the capability to establish a proxy for use with Microsoft Azure Blob Storage, Amazon S3, and S3 Compatible storage accounts.

To set up a proxy configuration:

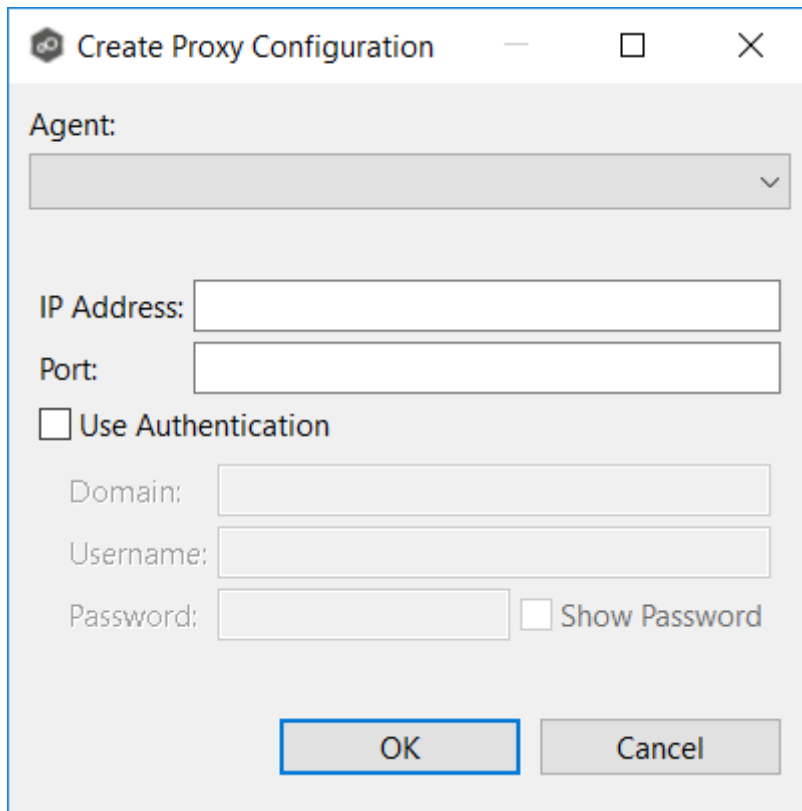
1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Proxy Configuration**.

Any existing proxies are listed in the **Proxy Configuration** table.



3. Click the **Create** button.

The **Create Proxy Configuration** dialog appears.



The screenshot shows a dialog box titled "Create Proxy Configuration". It features a dropdown menu for "Agent", text input fields for "IP Address" and "Port", a checkbox for "Use Authentication", and text input fields for "Domain", "Username", and "Password". A "Show Password" checkbox is located next to the password field. The dialog has "OK" and "Cancel" buttons at the bottom.

4. Select the Agent that manages your storage device.
5. Enter values for the following fields:

Field	Description
IP Address	Enter the IP address or fully qualified domain name of the proxy server.
Port	Enter the port number.
Use Authentication	Select if your proxy server requires authentication. This option does not apply for proxy servers connecting to an Azure storage device

Replication and Retention Policy Wizard

Replication and Retention Policy

✖ You must enter a name for the policy.

Replication and Retention
Replication Schedule
Retention
Source Snapshots

*Name:

Enable Backup with Destination Snapshots

< >

< Back Next > Finish Cancel

4. Enter a name, and then click **Next**.
5. Complete the wizard.

See [Step 11: Replication and Retention Policy](#) in [Creating a Cloud Backup and Replication Job](#) for assistance in completing the wizard.

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

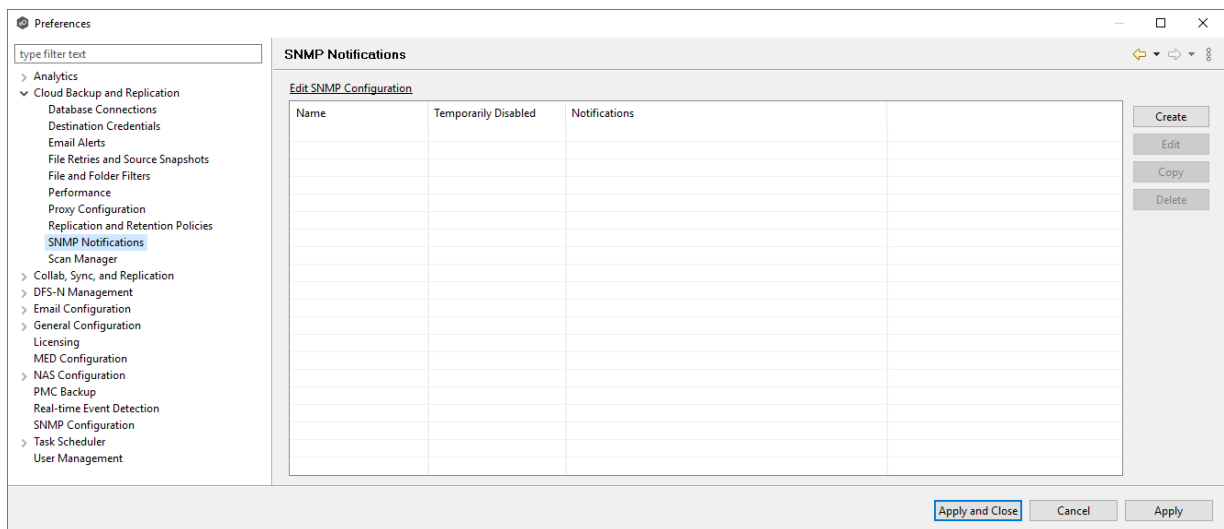
SNMP Notifications

When you create a job, you can select an existing SNMP notification to apply to the job or you can create a new notification and apply it to the job. You cannot edit or delete an SNMP notification while it is applied to a job. See [SNMP Notifications](#) in the [Basic Concepts](#) section for more information about SNMP notifications.

To create an SNMP notification:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **SNMP Notifications**.

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



3. Click the **Create** button.

The **Add SNMP Notification** dialog appears.

4. Select the types of events that will trigger the generation of an SNMP trap:

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

Event Type	Description
Event	

5. Click **OK**.

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

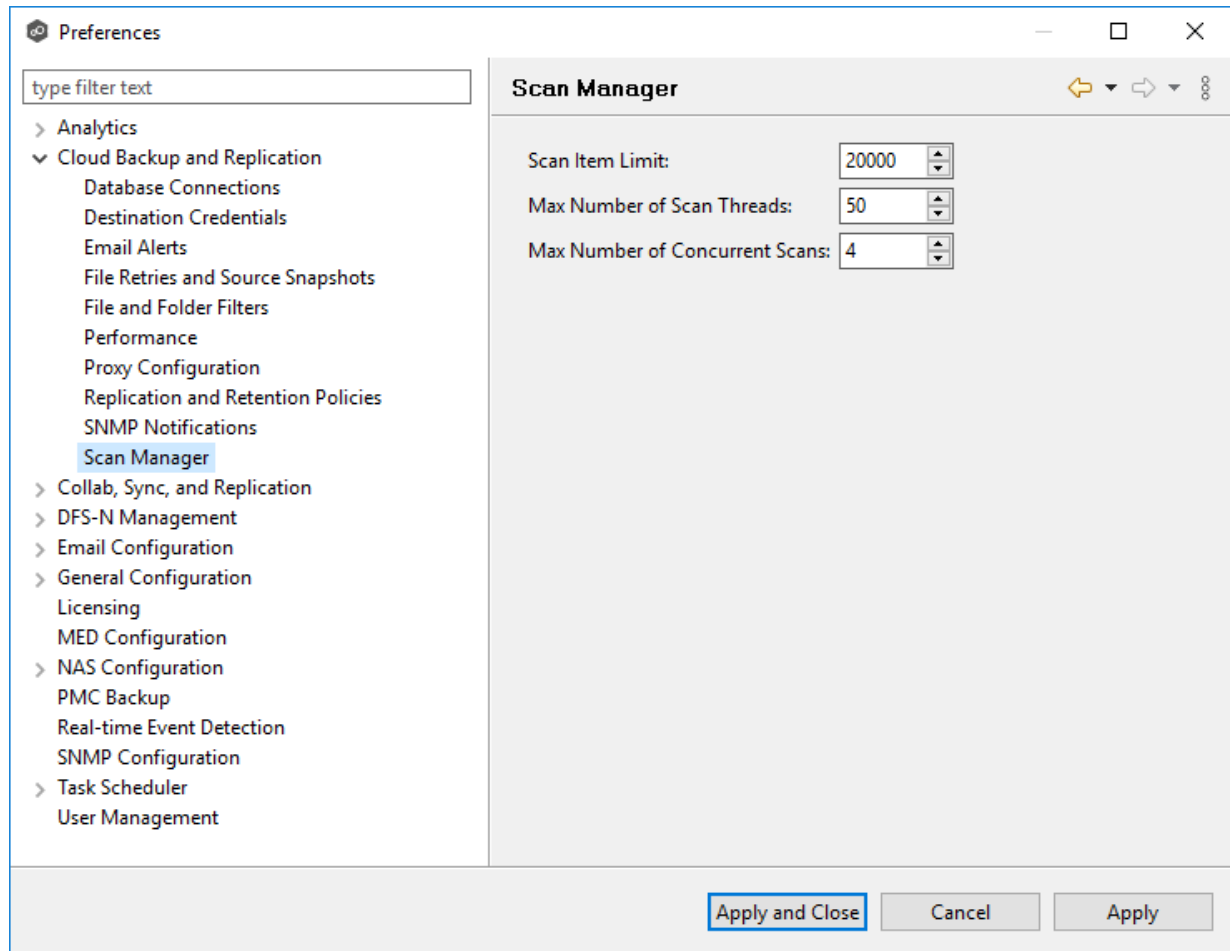
6. Click **Apply and Close** or **Apply**.

Scan Manager

The Cloud Backup and Replication Scan Manager is responsible for handling all scheduled and on-demand scans of the source storage device.

To modify the Scan Manager settings for Cloud Backup and Replication jobs:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Cloud Backup and Replication** in the navigation tree, and then select **Scan Manager**.



3. Modify the settings as needed.

Setting	Description
Scan Item Limit	Enter the maximum number of files and folders to obtain from a folder structure at a time during a scan.
Max Number of Scan Threads	Enter the maximum number of threads available for scanning files and folders. Set the number to at least the maximum number of jobs running on any single Management Agent.
Max Number of Concurrent Scans	Enter the maximum number of scans that can run in parallel. If the number of active scan threads is greater than this number, scan threads will process on a rotating basis. Increasing this number can increase scan performance but will also increase system memory and CPU utilization.

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

Collaboration, Replication, and Synchronization Job Preferences

You can modify the following settings for File Collaboration, File Synchronization, and File Replication jobs:

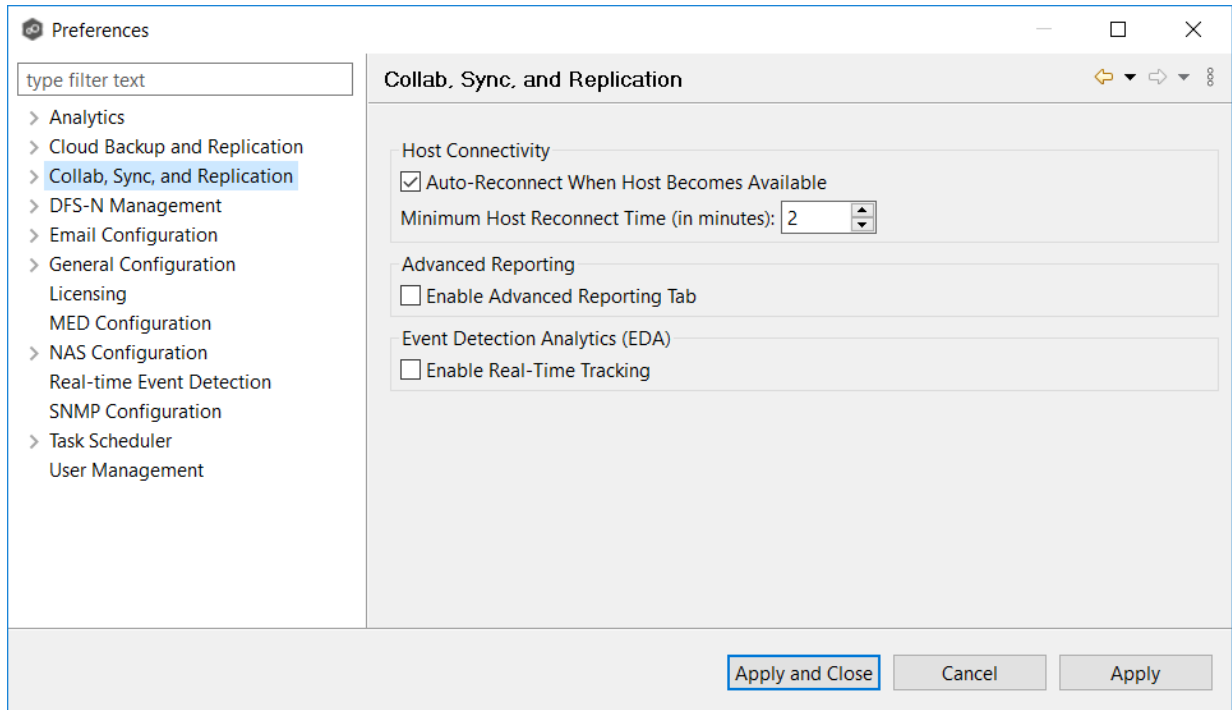
- [Collab Sync, and Replication](#)
- [DFS-N Management](#)
- [Edge Caching](#)
- [Email Alerts](#)
- [File Retries](#)
- [File and Folder Filters](#)
- [Locking](#)
- [Performance](#)
- [Real-time Event Detection](#)
- [Revit Enhancements](#)
- [SMNP Notifications](#)
- [Scan Manager](#)
- [Scheduled Replication Filters](#)

Collab, Sync, and Replication

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

To modify these settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **Collab, Sync, and Replication** in the navigation tree.



3. Modify the settings as needed.

Option	Description
Auto Reconnect When Host Becomes Available	When an Agent reconnects to Peer Management Center after a failure, the Agent is automatically re-enabled in any associated jobs. Highly recommended.
Minimum Host Reconnect Time (in minutes)	Enter the minimum number of minutes to wait after an Agent reconnects before re-enabling it in any associated jobs.
Enable Advanced Reporting Tab	Enables the Reporting tab in the global Collab, Sync, and Repl Summary view.
Enable Real-Time Tracking	<p>Enables Event Detection Analytics to track and report common activity processed by Peer Global File Service.</p> <p>If enabled, every 24 hours, an Excel-based report will be written to disk that shows top folders, files, extensions, and users by total processed activity over the previous 24-hour window. These reports are stored under the installation folder of Peer Management Center and can be reviewed by Peer Software Technical Support when uploading log files.</p>

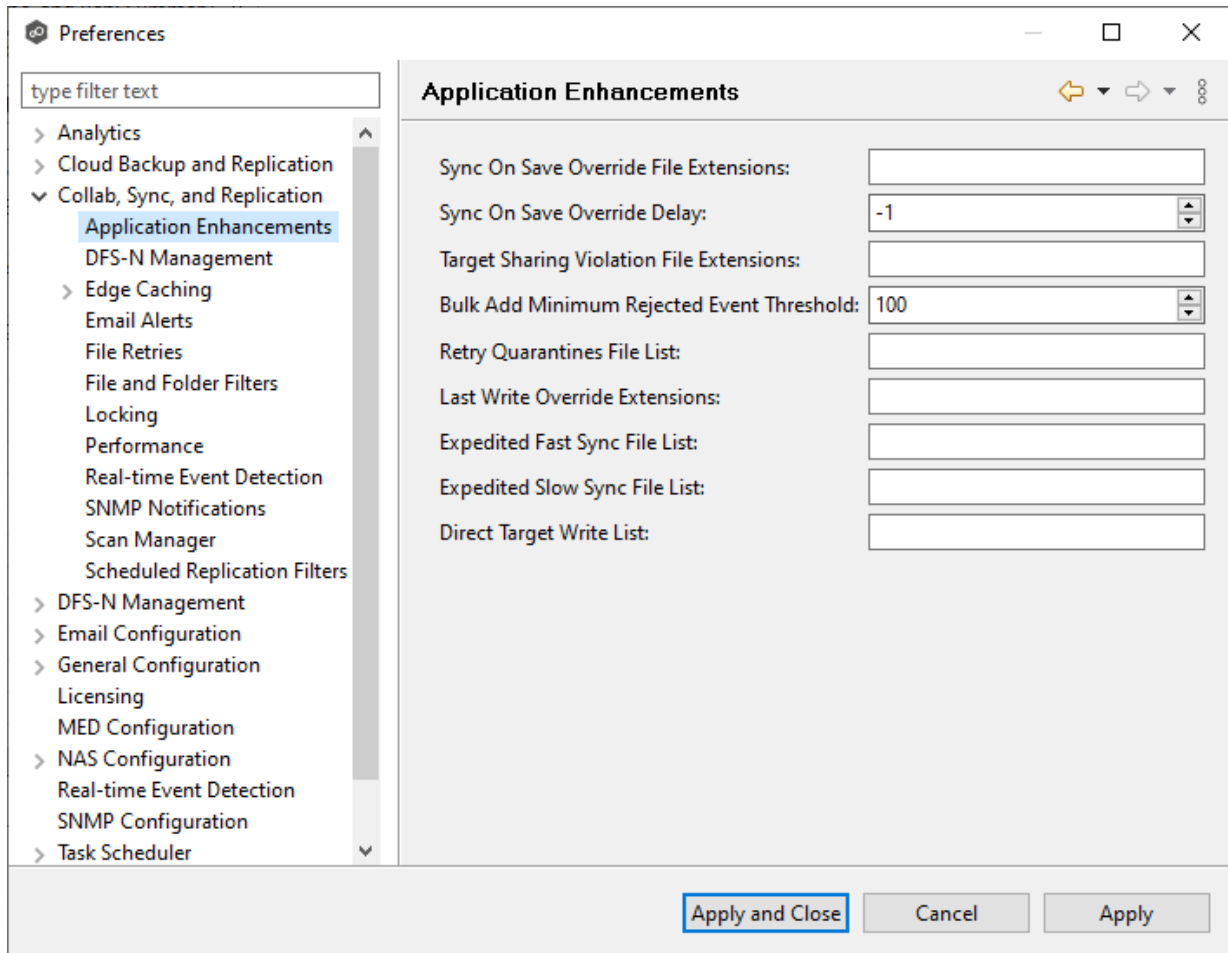
4. Click **Apply and Close** or **Apply**.

Application Enhancements

The settings on this page fine-tune how the specified file types are replicated. Most of the settings on this page are automatically configured based on selections in the **Application Support** page for the job. Consult with the [Peer Support](#) team before changing any settings as modifying values may cause unexpected results.

To configure advanced settings for Application Enhancements:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Application Enhancements**.



Y

3. Modify the options as needed.

Option	Description
Sync On Save Override File Extensions	Extensions configured here will overwrite the Sync. On Save values configured in the interface for the job. In addition, these extensions use the delay value in Sync On Save Override Delay setting instead of the delay value configured in the interface. If no delay value is set, it will default to using a one second delay. Extensions configured in this list will still be processed via Sync. On Save even if they also exist in the user defined non-collaborative extension list (under the Window > Preferences menu option). Extensions in the normal Sync. On Save list that also exist in this list will not be processed.

Option	Description
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.
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 be processed 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.

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

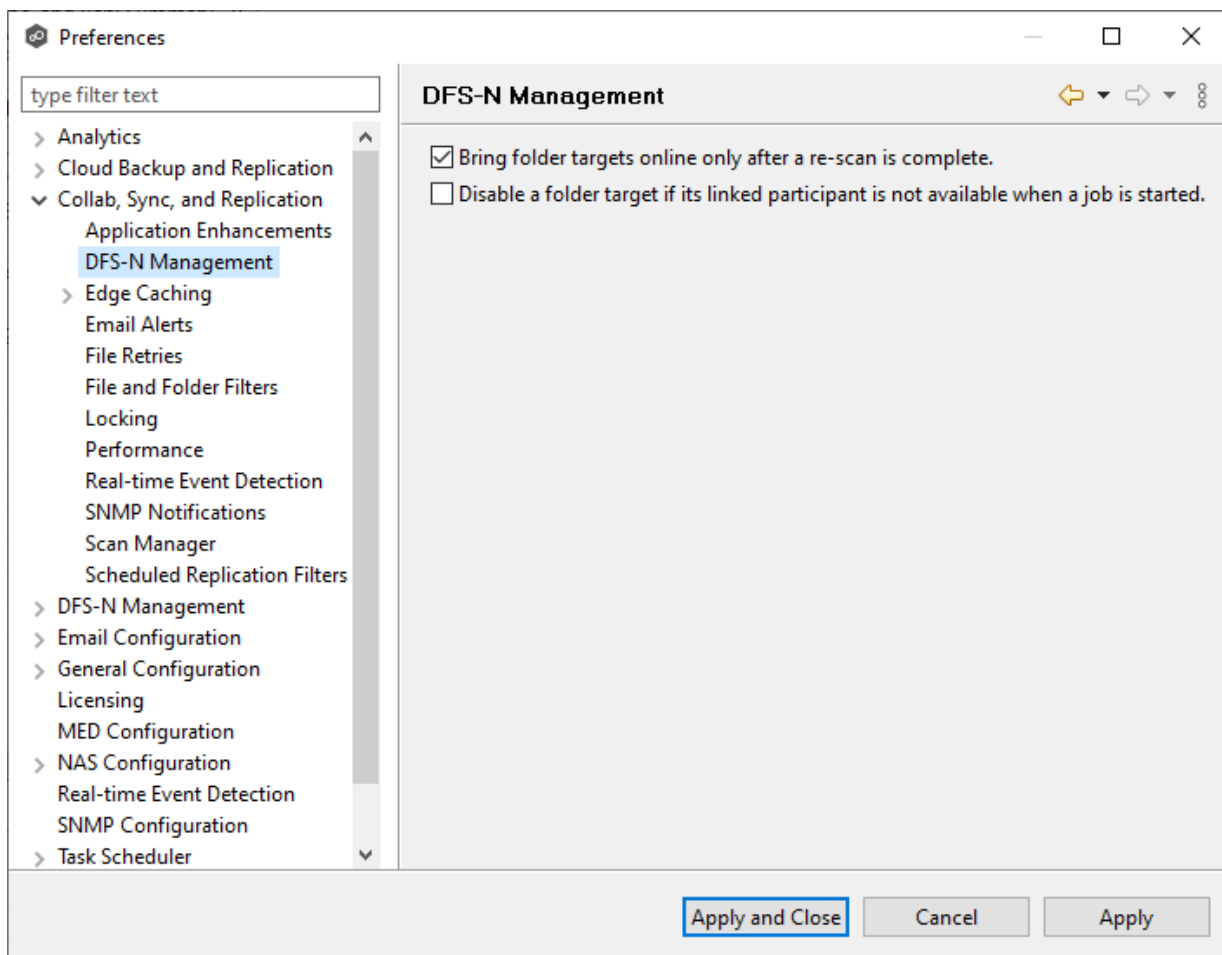
DFS-N Management

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

These settings control the [failover and failback](#) capabilities for [PeerGFS-managed namespaces](#) that are [linked to File Collaboration and File Synchronization jobs](#).

To modify these settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **DFS-N Management**.



3. Select options as needed.

Option	Description
Bring folder targets online only after a re-scan is complete	Re-enable a disabled folder target in a PeerGFS-managed DFS namespace only when it has been rescanned and is back in sync after an outage. Highly recommended.
Disable a folder target if its linked participant is not available when a job is started	If a File Collaboration or File Synchronization job is started and a participant is not available, automatically disable its associated folder target in a managed DFS namespace.

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

Edge Caching

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

These settings control the following aspect of jobs that use [Edge Caching](#):

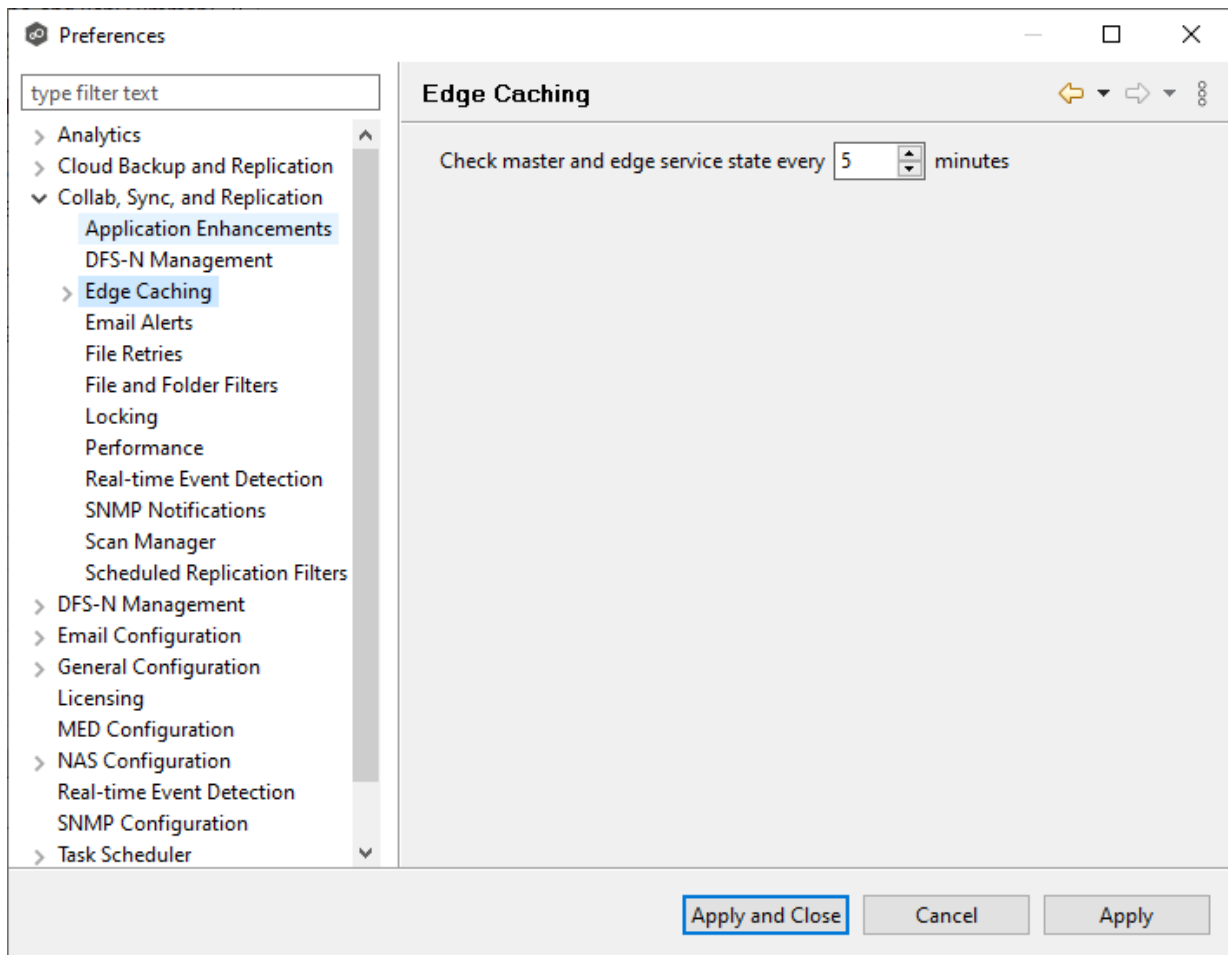
- [Edge Caching](#)
- [Email Alerts](#)
- [Master Data Service](#)
- [Pinning Filters](#)
- [Utilization Policies](#)
- [Volume Policies](#)

The [Peer Master Data Service](#) and Peer Edge Service are used by Edge Caching. The Peer Master Data service is a web service that handles requests from edge participants for files on a master participant. It runs on Master participants and is [configurable](#). The primary job of the Peer Edge service is to service reparse requests for Peer stub files and read and/or rehydrate stub files.

Use this page to set the frequency that these services are checked to see if they are operational.

To change the frequency:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Edge Caching**.
3. Change the frequency.



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

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

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

To create an Edge Caching email alert:

1. Select **Open Preferences** from the **Tools** menu.

Create Edge Caching Email Alert

Name:

Temporarily Disable

Caching Scan Alerts

Scan Start/End Scan Fatal Errors Scan Errors

Volume Alerts

Cache Size Exceeded Low Disk Space

Cache Safe Percentage Exceeded

Service Alerts

Master/Edge Services Health Monitoring

Recipients

Enter email, contact, or distribution list:

Recipients:

5. Enter a name for the alert.
6. Select the caching scan event types to be alerted.

Event Type	Description
Scan Start/End	Sends a notification when a caching scan is started or stopped.
Scan Fatal Errors	Sends a notification when a fatal error occurs during a caching scan.
Scan Errors	Sends a notification when errors occur during a caching scan.
Cache Size Exceeded	Sends a notification when the amount of volume disk space used by Edge Caching exceeds the size specified by the Cache Size option in

Event Type	Description
	the volume policy.
Low Disk Space	Sends a notification the volume disk space falls below the size specified by the Disk space is less than X option in the volume policy.
Cache Safe Percentage Exceeded	Sends a notification when the percentage specified by the Cache usage exceeds X% of cache size option in the volume policy.
Master/Edge Services Health Monitoring	Sends a notification if either the Peer Master Data Service or the Peer Edge Service goes down.

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

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


- Click **OK**.

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

The screenshot shows the 'Email Alerts' configuration window. The table below represents the data shown in the 'Edit Email Configuration' table:


Name	Temporarily Disa...	Event Types	Recipients
Fatal Errors with Edge Cac...	No	Scan Fatal Errors	AriF@peersoftware.com
Low Disk Space	No	Low Disk Space	AriF@peersoftware.com

If you selected a master participant currently being used in a job, the **In Use Configuration** dialog appears. Click **OK** to close the dialog.



Otherwise, the **Edit Master Data Service** dialog appears. The first two fields on this page are automatically populated:

- **Protocol:** This field lists the protocol that will be used to transfer file content between master participants and edge participants. HTTPS is currently the only available option as it requires only a single open firewall port on each master participant and does a better job of handling latency over WAN-based connections.
- **Agent Name:** This field lists the name of the Agent.



5. (Optional) Enter a value for **Agent Alias**; the value can be a hostname, FDQN, or IP address.

A value for **Agent Alias** is required only if the name of the Agent cannot be converted to an IP address via DNS. If an alias name is entered, it will be used by the Edge Service on edge participants to connect with the Master Data Service. If no alias name is entered, the name of the Agent will be used.

6. (Optional) Modify the default value for **Port** if you are use a different port.

If you modify the port number, the Master Data Service will be restarted, and the new port number will take effect immediately.

7. Click **OK**.

The revised Master Data Service is listed in the Master Data Service table.

8. Click **Apply and Close** or **Apply**.

The new settings will be applied to all Edge Caching-enabled jobs.


When you create an Edge Caching-enabled job, you can select existing pinning filters to apply to the job or you can create new pinning filters and apply them to the job. This [Preferences](#) page lists the existing pinning filters. From this page, you can create, edit, and delete pinning filters. However, you cannot edit or delete a pinning filter while it is applied to a job.

A pinning filter specifies whether specific files or files in a particular directory are always stubbed or always local on the edge participant. A pinning filter similar is to a utilization policy—it can be applied to multiple jobs. If there is a conflict between a pinning filter and utilization policy (where, for example, you might have something set to be always stubbed), the pinning filter will take precedence.

To create a pinning filter:

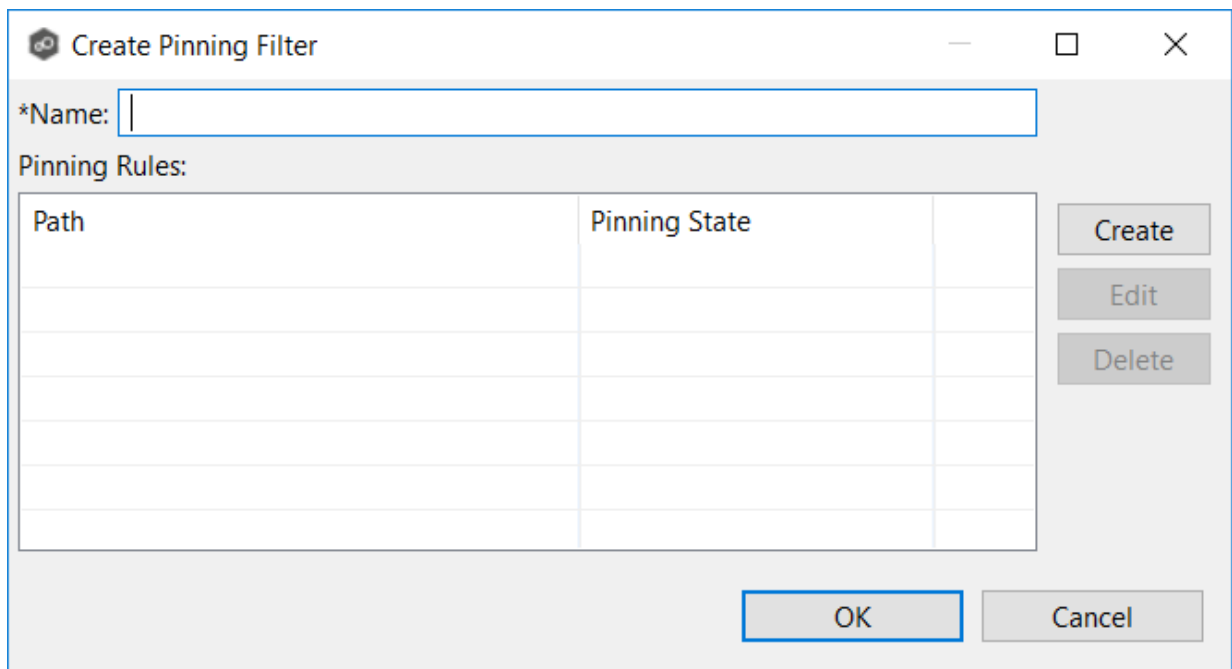
1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Edge Caching**.
3. Select **Pinning Filters**.

Any existing pinning filters are listed in the **Pinning Filters** table.

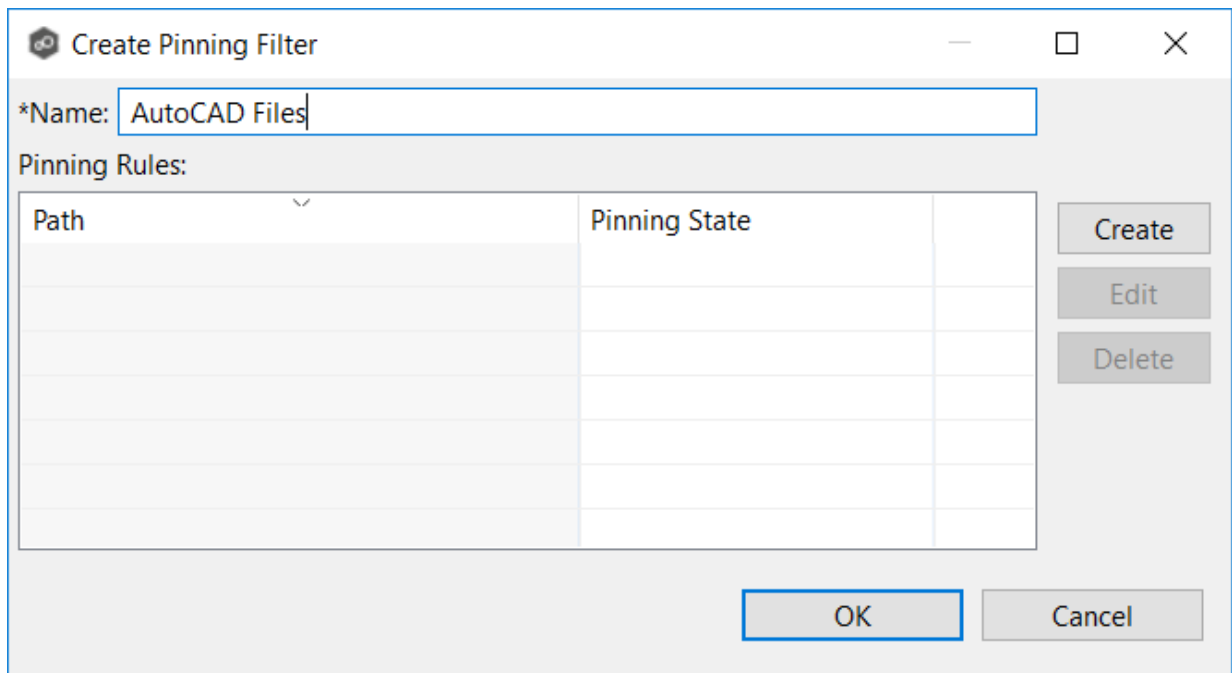


4. Click **Create**.

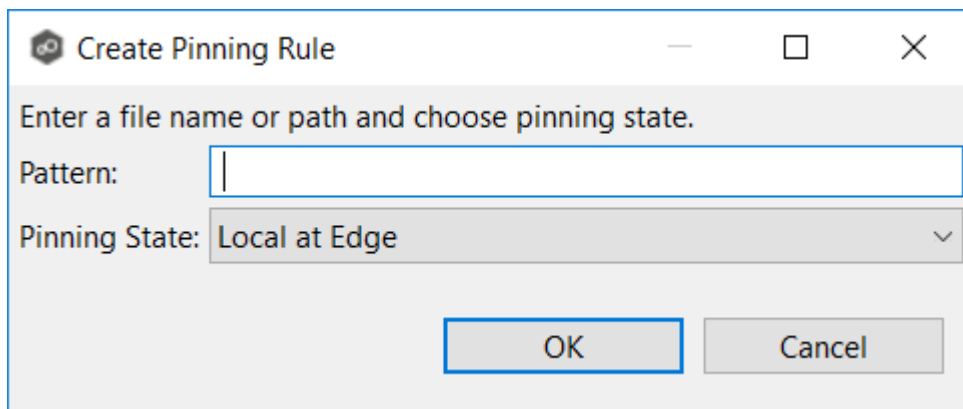
The **Create Pinning Filter** dialog appears.



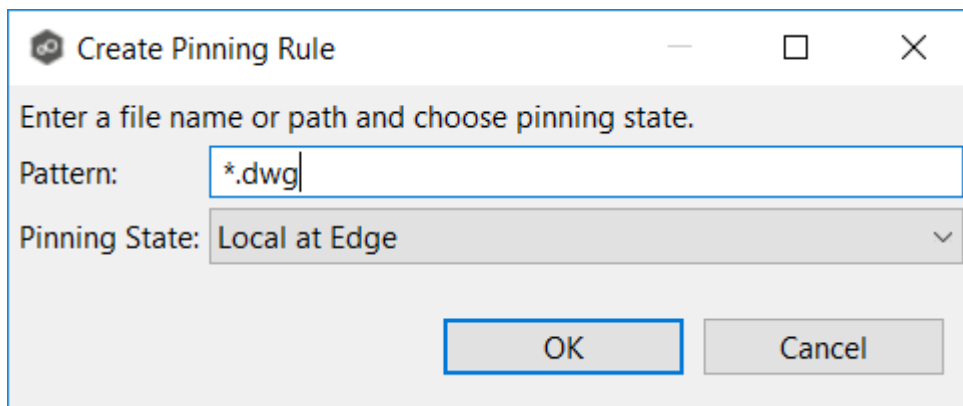
5. Enter a name for the pinning filter.



6. Click **Create** to add a pinning rule to the filter.



7. Enter a file name or enter a path.



8. Choose a pinning state:

- Select **Local at Edge** if you want the specified files or path to always be local and never stubbed.
- Select **Stubbed at Edge** if you want the specified files or path to always be stubbed at edge.

9. Click **OK**.

The rule appears in the **Pinning Rules** table.

Create Pinning Filter

*Name: AutoCAD Files

Pinning Rules:

Path	Pinning State
*.dwg	Local at Edge

Create

Edit

Delete

OK

Cancel

10. If you want to add additional rules to the pinning filter, repeat Steps 6-9.

11. Click **OK** to close the **Create Pinning Filter** dialog.

The new filter is listed in the **Pinning Filters** table and can now be applied to Edge Caching-enabled jobs.

Event Type	Description
Scan Start/End	Sends a notification when a caching scan is started or stopped.
Scan Fatal Errors	Sends a notification when a fatal error occurs during a caching scan.
Scan Errors	Sends a notification when errors occur during a caching scan.
Cache Size Exceeded	Sends a notification when the amount of volume disk space used by Edge Caching exceeds the size specified by the Cache Size option in the volume policy.
Low Disk Space	Sends a notification the volume disk space falls below the size specified by the Disk space is less than X option in the volume policy.
Cache Safe Percentage Exceeded	Sends a notification when the percentage specified by the Cache usage exceeds X % of cache size option in the volume policy.
Master/Edge Services Health Monitoring	Sends a notification if either the Peer Master Data Service or the Peer Edge Service goes down.

6. Click **OK**.

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

Create Utilization Policy

*Name:

File Size

Keep files local if less than

Stub files if greater than

Time Period

Keep recently used files local based on a dynamic set of rules

Keep recently used files local based on the following rules:

Stub files if not modified within the past

Stub files if not accessed within the past

Stubbing Override

Select to override time period rules and Stub at Edge pinning rule:

Stub files if not accessed within the past

Advanced Options

Do not hydrate files during caching scan

5. Enter a name for the policy.
6. (Optional) In the **File Size** section, select one or both options:

Field	Description
Keep files local if less than X size	Select this option if you want files under a specified size to remain local.
Stub files if greater than X size	Select this option if you want files over a specified size to be stubbed.

7. (Optional) In the **Time Period** section, select one of the options:

Field	Description
Keep recently used files local based on a dynamic set of rules	Select this option if you want Edge Caching to control when to stub files based on last accessed and last modified times. Edge Caching dynamically adjusts the accessed and modified time periods it uses based on the total amount of space that Edge Caching is actively using on a volume.
Keep recently used files local based on the following rules	Select this option if you want to specify the dates used when stubbing files based on the last accessed and last modified.


8. If you selected the **Keep recently used files local based on the following rules** option in **Time Period**, two additional options are enabled; select the options to be used and specify the time periods to be used:

Field	Description
Stub files if not modified within the past X time period	Select this option and specify a time range to use the last modified date of a file to determine if it should be kept local or stubbed.
Stub files is not accessed within the past X time period	Select this option and specify a time range to use the last accessed date of a file to determine if it should be kept local or stubbed.


9. (Optional) in the **Stubbing Override** section, select the checkbox and a time period if you want to use the last accessed date of all recently hydrated files to determine if they should be kept local or re-stubbed.

10. Click **OK**.

The new policy is listed in the **Utilization Policies** table and can now be applied to jobs.



6. If you selected a policy currently being used in a job, the **In Use Configuration** message appears.



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

The **Edit Volume Policy** dialog appears.

Edit Volume Policy (DGAgent2 - C:\)

Volume Policy
Define volume policy and set temporary storage path.

Volume Policy
Email Alerts
SNMP Notifications

Select Different Utilization Policy

Cache Size

Use up to 75 % of this volume

Use up to 10 GB of this volume

Cache Threshold Alerts

Send alerts when:

Disk space is less than 500 MB

Cache usage exceeds 80 % of the cache size

Caching Scan Schedule

Scan every 1 day(s) at 10:00:00 PM

Define schedule
Runs every Monday, Friday every week [Edit](#)

*Temporary Storage Path: C:\ [Browse](#)
C:\PeerTempPath

< Back Next > Finish Cancel

8. If you want to associate a different utilization policy with this volume policy, click **Select Different Utilization Policy** link, select the policy, and then click **OK**.

Select Different Utilization Policy (DGAgent2 - C:\)

Selecting a different utilization policy will affect all jobs with the same Agent and volume.

Utilization Policy:

Current Configuration

Agent: DGAgent2

Volume: C:\

Current Utilization Policy: Default Utilization Policy

OK Cancel

9. In the **Cache Size** section, choose an option for setting the cache size.

- Use up to X % of this volume
- Use up to X size of this volume

Edit Volume Policy (DGAgent2 - C:\)

Volume Policy
Define volume policy and set temporary storage path.

Volume Policy
Email Alerts
SNMP Notifications

Select Different Utilization Policy

Cache Size

- Use up to % of this volume
- Use up to of this volume

Cache Threshold Alerts
Send alerts when:

Disk space is less than

Cache usage exceeds % of the cache size

Caching Scan Schedule

- Scan every day(s) at
- Define schedule

*Temporary Storage Path:

C:\PeerTempPath

< Back Next > **Finish** Cancel

10. In the **Cache Threshold Alerts** section, enter values for automatic alerts about free disk space and cache usage.

Peer Management Center will automatically display alerts in the **Alerts** tab and send alerts via email (if configured) when:

- The amount of free disk space on the volume falls below the specified value. For example, if a 1TB volume has 500MB of free space and the threshold is set to 512MB, an alert will be sent.
- Cache usage on the volume exceeds the specified percentage of the cache size. For example, if the cache size is set to 80%, equating to 750 GB, Edge Caching will start sending alerts when it has used 600 GB.

Edit Volume Policy (DGAgent2 - C:\)

Volume Policy
Define volume policy and set temporary storage path.

Volume Policy
Email Alerts
SNMP Notifications

Select Different Utilization Policy

Cache Size

Use up to 75 % of this volume

Use up to 10 GB of this volume

Cache Threshold Alerts

Send alerts when:

Disk space is less than 500 MB

Cache usage exceeds 80 % of the cache size

Caching Scan Schedule

Scan every 1 day(s) at 10:00:00 PM

Define schedule

*Temporary Storage Path: C:\ Browse

C:\PeerTempPath

< Back Next > **Finish** Cancel

11. In the **Caching Scan Schedule**, set the frequency that the volume should be scanned to optimize the balance of fully hydrated vs stubbed files.

This scan can be run daily at a specified time, or you can define a more customized schedule.

Volume Policy

Define volume policy and set temporary storage path.

Volume Policy

Email Alerts

SNMP Notifications

Select Different Utilization Policy

Cache Size

Use up to 75 % of this volume

Use up to 10 GB of this volume

Cache Threshold Alerts

Send alerts when:

Disk space is less than 500 MB

Cache usage exceeds 80 % of the cache size

Caching Scan Schedule

Scan every 1 day(s) at 10:00:00 PM

Define schedule

*Temporary Storage Path: C:\ Browse

C:\PeerTempPath

< Back Next > Finish Cancel

12. In the **Temporary Storage Path** field, enter a path or browse to the location you want to be used as temporary storage space.

The temporary storage space will be used to store the content of stub files as they are being rehydrated. The content of files undergoing rehydration are referred to as **file blocks**. File blocks are fixed-length chunks of data that are read into memory when requested by an application. Edge Caching will create a subfolder named **.PeerTempPath** under the location that you specify and temporarily store the file blocks in that subfolder.

For optimal performance, we recommend that the temporary storage space be on the same volume as the watch set. If that is not possible, it should be on a high-performance disk.

Edit Volume Policy (DGAgent2 - C:\)

Volume Policy
Define volume policy and set temporary storage path.

Volume Policy
Email Alerts
SNMP Notifications

Select Different Utilization Policy

Cache Size

Use up to 75 % of this volume

Use up to 10 GB of this volume

Cache Threshold Alerts

Send alerts when:

Disk space is less than 500 MB

Cache usage exceeds 80 % of the cache size

Caching Scan Schedule

Scan every 1 day(s) at 10:00:00 PM

Define schedule


*Temporary Storage Path: C:\ Browse

C:\PeerTempPath

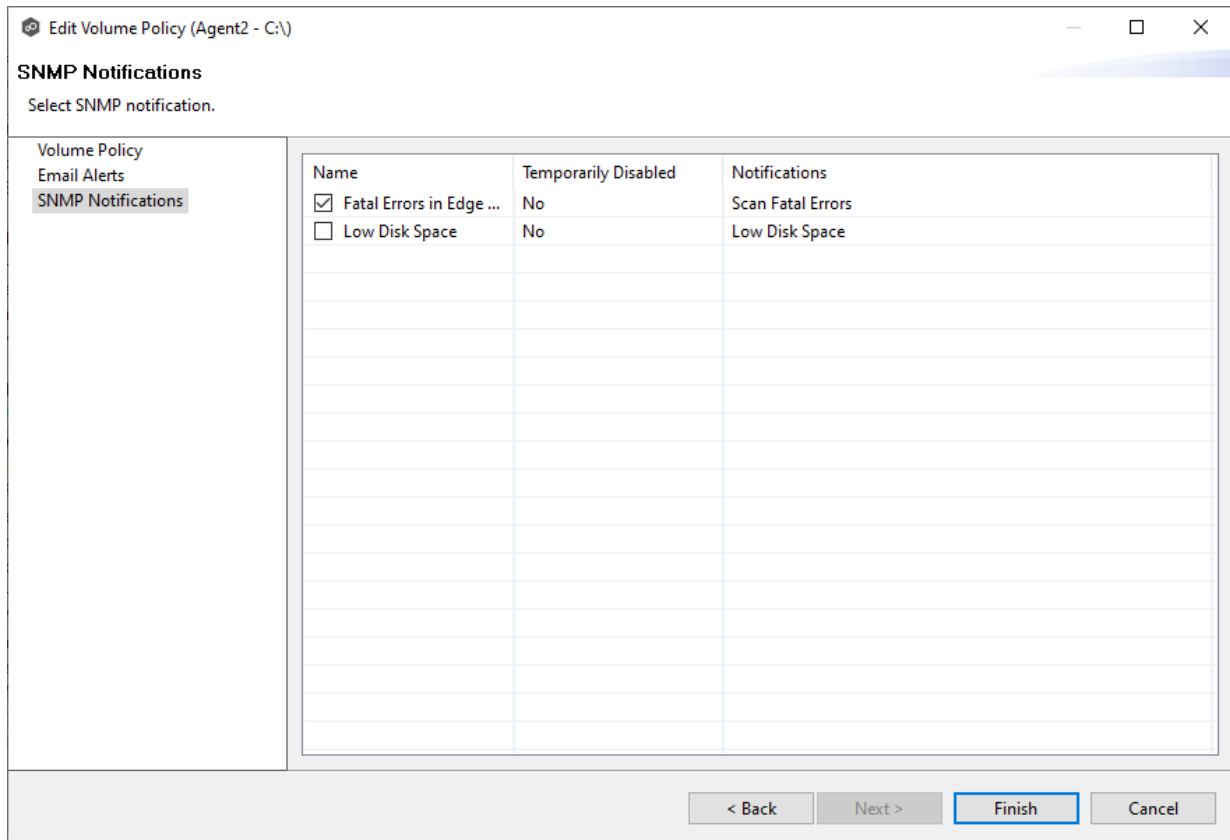
< Back Next > **Finish** Cancel

13. Click **Next**.

14. (Optional) Select one or more email alerts to be associated with this volume policy, and then click **Next**.

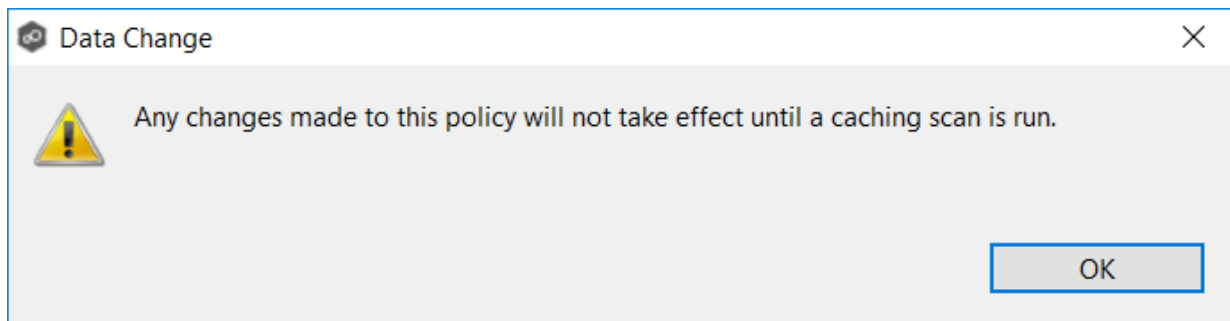


15. (Optional) Select one or more SNMP notifications to be associated with this volume policy.



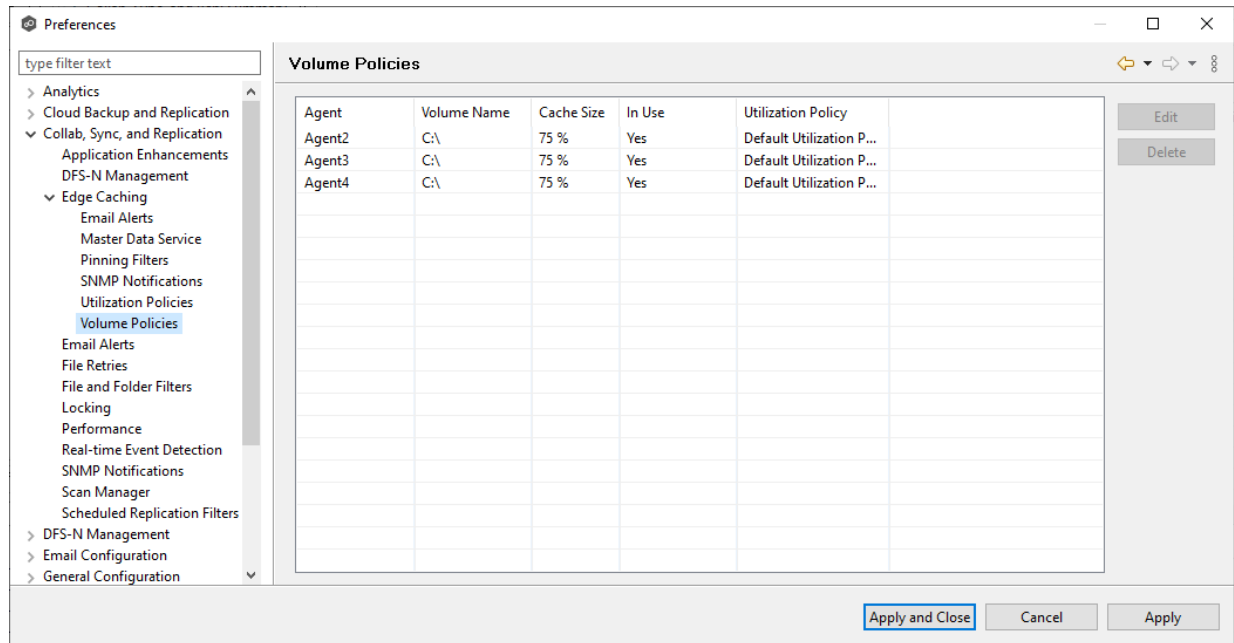
16. Click **Finish**.

The following message is displayed:



17. Click **OK**.

The revised volume policy is listed in the **Volume Policies** table. It will be used after jobs using the policy are restarted.



18. Click **Apply and Close** or **Apply**.

Email Alerts

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

Tip: If you are performing maintenance, you can temporarily disable an email alert by clicking in the **Temporarily Disabled** column in the **Email Alerts** table and selecting **Yes**. No email alerts will be sent for that type of alert until you reenables the alert.

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

To create an email alert:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Email Alerts**.

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

Create Email Alert

Name:

Temporarily Disable

Event Types

Job Start Job Stop Job Failure Participant Failure Participant Reconnect File Quarantine
 Scan Error Malicious Event

Queued Items

<input type="checkbox"/> Number of Queued Items	<input type="checkbox"/> Size of Queued Items:
Exceeds: <input type="text" value="5000"/> Items	Exceeds: <input type="text" value="10240"/> MB
Recovers Below: <input type="text" value="1000"/> Items	Recovers Below: <input type="text" value="1024"/> MB
<input type="checkbox"/> Alert on Recovery	<input type="checkbox"/> Alert on Recovery

Reports

Scan

Batch Email Alerts

Quarantined Files

Recipients

Enter email, contact, or distribution list:

Recipients:

4. Enter a name for the alert.
5. Select the types of events that will trigger an email alert to be sent.

Event Type	Description
Job Start	Sends an alert when the job starts.
Job Stop	Sends an alert when the job stops.
Job Failure	Sends an alert when the job is aborted because of lack of quorum due to one or more failed participants.
Participant Failure	Sends an alert when a participant timeout occurs, and the participant is taken out of the running job.
Participant Reconnect	Sends an alert when a participant reconnects to the job and the job resumes with the reconnected participant.
File Quarantine	Sends an alert when a file is marked as quarantined because a file conflict was not able to be resolved.
Scan Error	Sends an alert when an error occurs during the initial synchronization process .
Malicious Event	Sends an alert when Peer Malicious Event Detection (MED) detects potentially malicious activity. For more information, see MED Configuration .

6. Select options in the **Queued Items** section.


Option	Description
Number of Queued Items section	<p>Select this checkbox if you want alerts to be sent regarding the number of items in the queue. This is useful to notify you about when there is a queue backlog potentially due to latency issues.</p> <p>If you select this checkbox, you need to enter two values that work in tandem:</p> <ul style="list-style-type: none"> • Exceeds X Items - Enter the highest number of queued items before an email alert is sent. The default value is 5000. • Recovers Below X Items - Enter a value. The default value is 1000. <p>An alert is sent the first time that the Queued Items counter has items greater than the value set in Exceeds X Items. The counter's value is displayed in the Queued Items column in the in the Collab, Sync, and Repl Summary view. The counter's value is a combination of the Real-time and File Sync queues.</p> <p>Another alert will not be sent until the Queued Items counter has dropped below the Recovers Below x Items value and then exceeds the Exceeds X Items value again. This prevents multiple or redundant alerts from being sent.</p>
Alert on Recovery	<p>Select this option if you want an alert to be sent when the number of queued items has fallen below the Recovers Below value.</p>
Size of Queued Items section	<p>Select this if you want alerts to be sent based on the total data size of queued items for a job. This is useful to notify you about when there is a queue backlog potentially due to bandwidth issues.</p> <p>If you select this checkbox, you need to enter two values that work in tandem:</p> <ul style="list-style-type: none"> • Exceeds X MB - Enter the highest number of queued items before an email alert is sent. The default value is 10240 MB. • Recovers Below X MB - Enter a value. The default value is 1024 MB. <p>An alert is sent the first time that the Pending Bytes for a job has items greater than the value set in Exceeds X MB. The counter's value is displayed in the Pending Bytes column in the Collab, Sync, and Repl Summary view.</p>
	<p>Another alert will not be sent until the Pending Bytes counter has dropped below the Recovers Below x MB value and then exceeds the Exceeds X MB value again. This prevents multiple or redundant alerts from being sent.</p>

File Retries

File retries settings enable you to configure the frequency of attempts and the maximum number of attempts. These settings apply to all File Collaboration, File Replication, and File Synchronization jobs. For more information about file retries, see [Conflicts, Retries, and Quarantines](#).

To modify the file retries settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **File Retries**.



3. Modify the settings as needed.

Setting	Description
Enable Retries	Select this checkbox to enable the retry of failed file transfers. If this option is not enabled, files that would have been candidates for retries will be automatically quarantined.
Maximum Number of File Retries	Enter the maximum number of attempts to retry a failed file transfer before it is quarantined.
Delay Between Retries in Seconds	Enter the number of seconds to wait between retries of a failed file transfer.

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

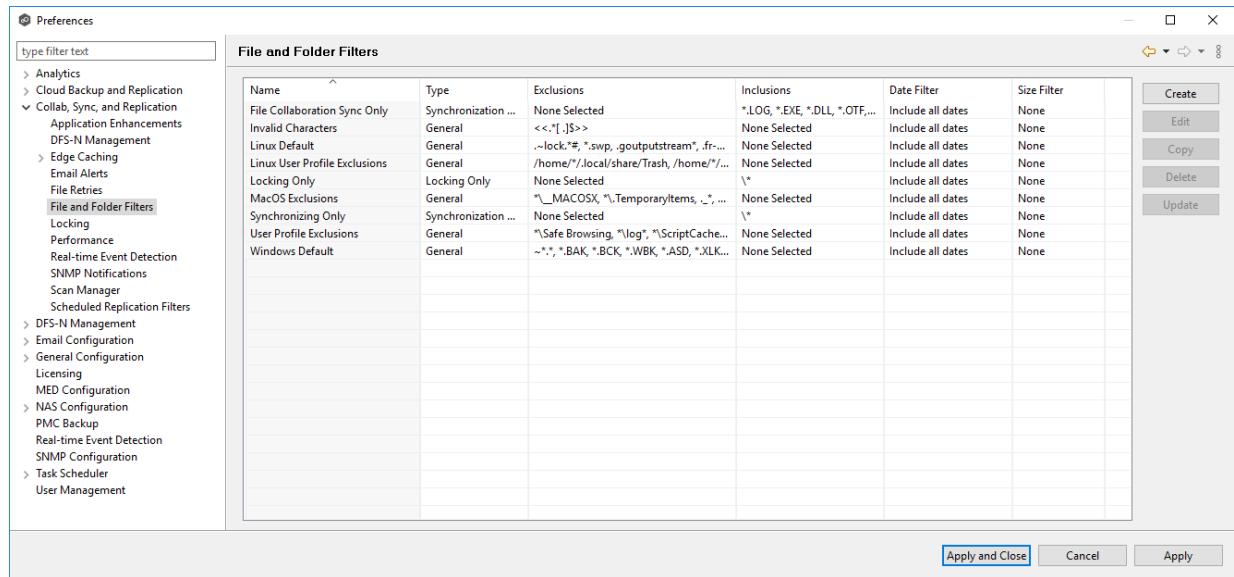
File and Folder Filters

When you create a File Collaboration, File Synchronization, or File Replication job, you can select existing file and folder filters to apply to the job or you can create new file filters and apply them to the job. This [Preferences](#) page lists the existing file and folder filters. From this page, you can view, create, edit, update, and delete file filters. However, you cannot edit or delete a file filter while it is applied to a job. See [File and Folder Filters](#) in the [Basic Concepts](#) section for more information about file and folder filters.

To create a file and folder filter:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **File and Folder Filters**.

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



3. Click **Create**.

Create File Filter

Name:

Filter Type: **General** ▾

Auto Excluded
[View file types that are automatically excluded](#)

Excluded Patterns

Add Edit Delete

Included Patterns

Add Edit Delete

Included Last Modified Dates
Include all dates ▾
 days

Excluded File Sizes
None ▾
 bytes

OK Cancel

4. Enter a unique name for the filter.
5. Select the [filter type](#).

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

See [Defining Filter Patterns](#) for information about filter patterns.

7. (Optional) Click **Add** to enter a filter pattern for files that you want included in the job. Repeat to add more filter patterns.
8. (Optional) Select a value for [Included Last Modified Dates](#).

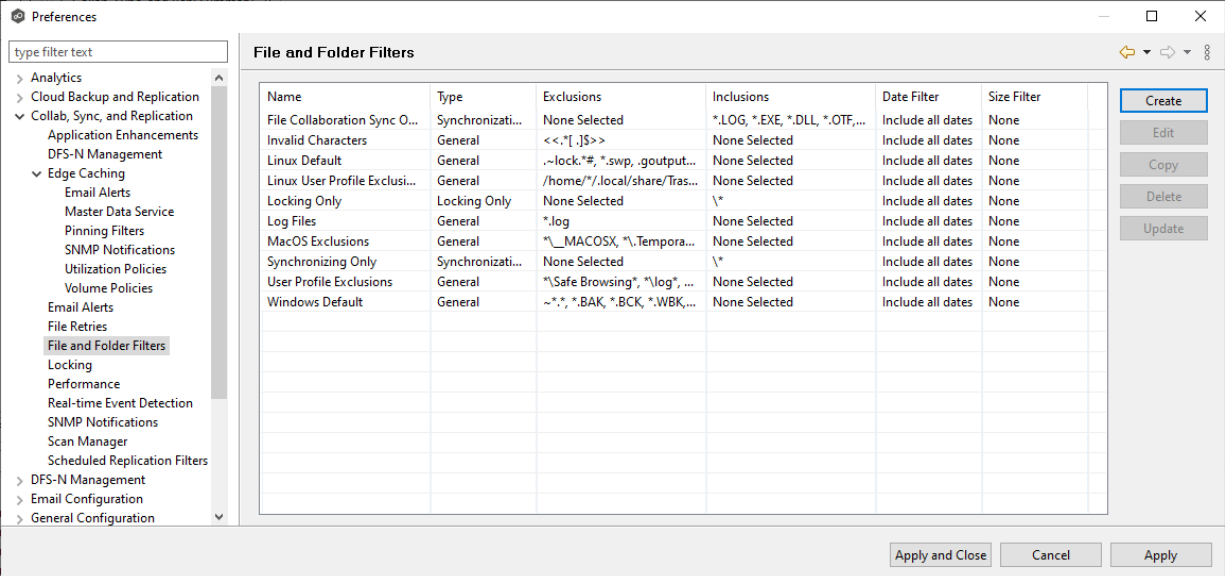
Note: A filter cannot use **Included Last Modified Dates** in conjunction with any other excluded or included patterns.

9. (Optional) Select a value for [Excluded File Sizes](#). Note: This cannot be combined with any other filter criteria

Note: A filter cannot use **Excluded File Sizes** in conjunction with any other excluded or included patterns.

10. Click **Apply and Close** or **Apply**.

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



The screenshot shows the 'File and Folder Filters' table with the following data:

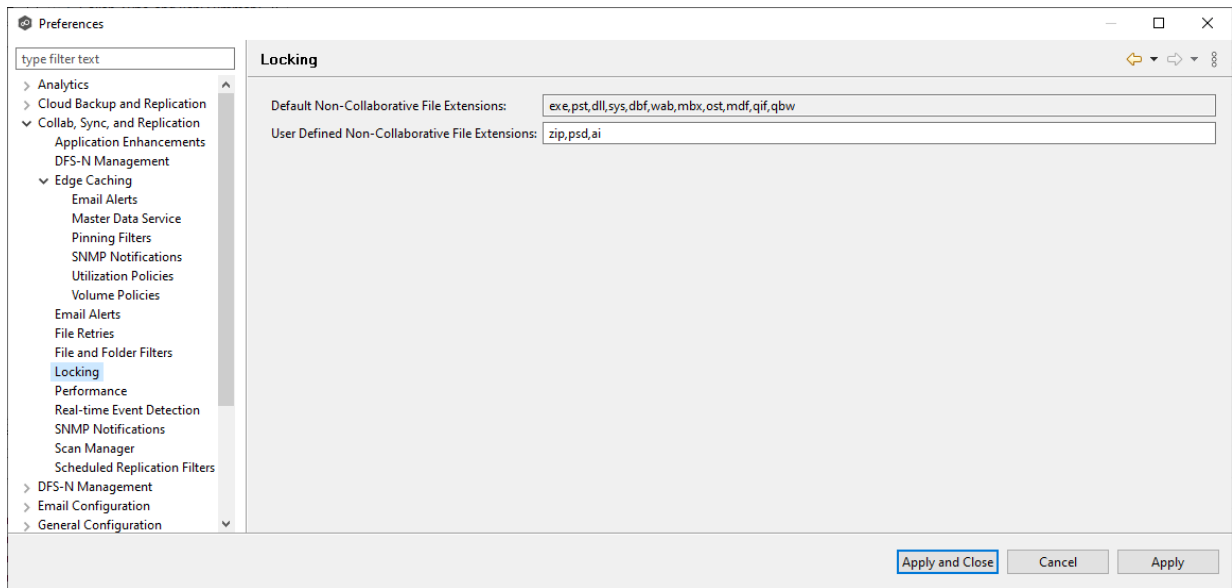
Name	Type	Exclusions	Inclusions	Date Filter	Size Filter
File Collaboration Sync O...	Synchronizati...	None Selected	*.LOG, *.EXE, *.DLL, *.OTF, ...	Include all dates	None
Invalid Characters	General	<<.[_!\$>	None Selected	Include all dates	None
Linux Default	General	./lock.*#, *.swp, .goutput...	None Selected	Include all dates	None
Linux User Profile Exclusi...	General	/home/~/local/share/Tras...	None Selected	Include all dates	None
Locking Only	Locking Only	None Selected	*	Include all dates	None
Log Files	General	*.log	None Selected	Include all dates	None
MacOS Exclusions	General	*\ _MACOSX, *\ Tempora...	None Selected	Include all dates	None
Synchronizing Only	Synchronizati...	None Selected	*	Include all dates	None
User Profile Exclusions	General	*\Safe Browsing*, *.log*, ...	None Selected	Include all dates	None
Windows Default	General	~*.*, *.BAK, *.BCK, *.WBK, ...	None Selected	Include all dates	None

Locking

An option is available to mark certain file types as non-collaborative, changing the way locks on the specified file types are handled. These settings apply to all File Collaboration, File Synchronization, and File Replication jobs. These settings are critical for certain file types so that the job can correctly read these files, ensuring that managed file types are synchronized in a consistent and usable state.

To modify the locking settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Locking**.



3. Modify the options as needed.


Option	Description
Default Non-Collaborative File Extensions	Non-editable. Displays the default, comma-separated list of file extensions of non-collaborative file types (e.g., database files). Write access to source files of these types is denied while the files are being synchronized.
User Defined Non-Collaborative File Extensions	Displays an editable, comma-separated list of file extensions of non-collaborative file types (e.g., database files). Write access to the source files of these types is denied while the files are being synchronized.

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

Performance

To customize the performance settings of File Collaboration, File Synchronization, and File Replication jobs:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Performance**.



3. Modify the options as needed.


Option	Description
Real-Time Expedited Threads	Enter the maximum number of threads for controlling file locking and renames.
Real-Time Background Threads	Enter the maximum number of threads for controlling the replication of file content.
Background Scan Synchronization Threads	Enter the maximum number of threads for processing the differences found by background scans.
Multipart Transfer Threads	Enter the maximum number of threads to be used for processing chunks of large files in parallel.
Debug Mode	Select to enable debug mode for the various types of threads.

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

Real-time Event Detection

To modify the File Collaboration real-time detection settings:

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



3. Modify the options as needed.

Option	Description
Change Dispatch Quiet Period (Seconds)	The number of seconds to wait before acting on a file modification, rename, or delete.
Bulk Add Delay (Seconds)	Controls when the bulk add logic is triggered. This is used to help deprioritize mass copying or adding of files to a directory.
Bulk Add Minimum Rejected Threshold	The minimum number of file adds that must occur within the Bulk Add Delay for bulk add logic to be triggered.
Bulk Add Max Events	The maximum number of file adds to lump together in one batch.

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

Add SNMP Notification

Name:

Temporarily Disable

Event Types

Job Start Job Stop Job Failure Participant Failure

File Quarantine Scan Error Malicious Event

Queued Items

Number of Queued Items Size of Queued Items:

Exceeds: Items Exceeds: MB

Recovers Below: Items Recovers Below: MB

Alert on Recovery Alert on Recovery

4. Select the types of events that will trigger the generation of an SNMP trap:

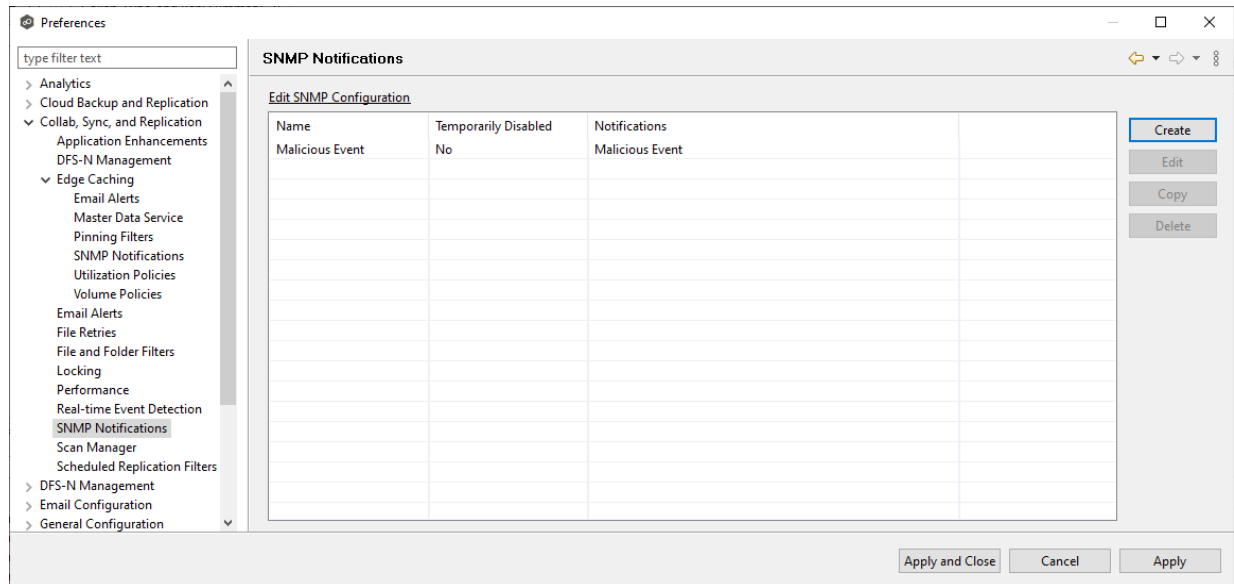
Event T	Description
Job Start	Sends a notification when the job starts.
Job Stop	Sends a notification when the job stops.
Job Failure	Sends a notification when the job is aborted because of lack of quorum due to one or more failed participants.
Partici pant Failur e	Sends a notification when a participant timeout occurs, and the participant is taken out of session.
Partici pant Recon nect	Sends a notification when a participant reconnects to the job and the job resumes with the reconnected participant.
File Quara ntine	Sends a notification when a file is marked as quarantined because a file conflict was not able to be resolved.
Scan Error	Sends a notification when an error occurs during the initial synchronization process .
Malici ous Event	Sends a notification when Peer MED detects potentially malicious activity. For more information, see MED Configuration .

5. Select options in the **Queued Items** section.

Option	Description
Number of Queued Items section	<p>Select this checkbox if you want alerts to be sent regarding the number of items in the queue. This is useful to notify you about when there is a queue backlog potentially due to latency issues.</p> <p>If you select this checkbox, you need to enter two values that work in tandem:</p> <ul style="list-style-type: none"> • Exceeds X Items - Enter the highest number of queued items before an email alert is sent. The default value is 5000. • Recovers Below X Items - Enter a value. The default value is 1000. <p>A notification is sent the first time that the Queued Items counter has items greater than the value set in Exceeds X Items. The counter's value is displayed in the Queued Items column in the in the Collab, Sync, and Repl Summary view. The counter's value is a combination of the Real-time and File Sync queues.</p> <p>Another notification will not be sent until the Queued Items counter has dropped below the Recovers Below x Items value and then exceeds the Exceeds X Items value again. This prevents multiple or redundant alerts from being sent.</p>
Alert on Recovery	<p>Select this option if you want a notification to be sent when the number of queued items has fallen below the Recovers Below value.</p>
Size of Queued Items section	<p>Select this if you want notifications to be sent based on the total data size of queued items for a job. This is useful to notify you about when there is a queue backlog potentially due to bandwidth issues.</p> <p>If you select this checkbox, you need to enter two values that work in tandem.</p> <ul style="list-style-type: none"> • Exceeds X MB - Enter the highest number of queued items before an email alert is sent. The default value is 10240 MB. • Recovers Below X MB - Enter a value. The default value is 1024 MB. <p>An alert is sent the first time that the Pending Bytes for a job has items greater than the value set in Exceeds X MB. The counter's value is displayed in the Pending Bytes column in the Collab, Sync, and Repl Summary view.</p>
	<p>Another alert will not be sent until the Pending Bytes counter has dropped below the Recovers Below x MB value and then exceeds the Exceeds X MB value again. This prevents multiple or redundant alerts from being sent.</p>

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

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

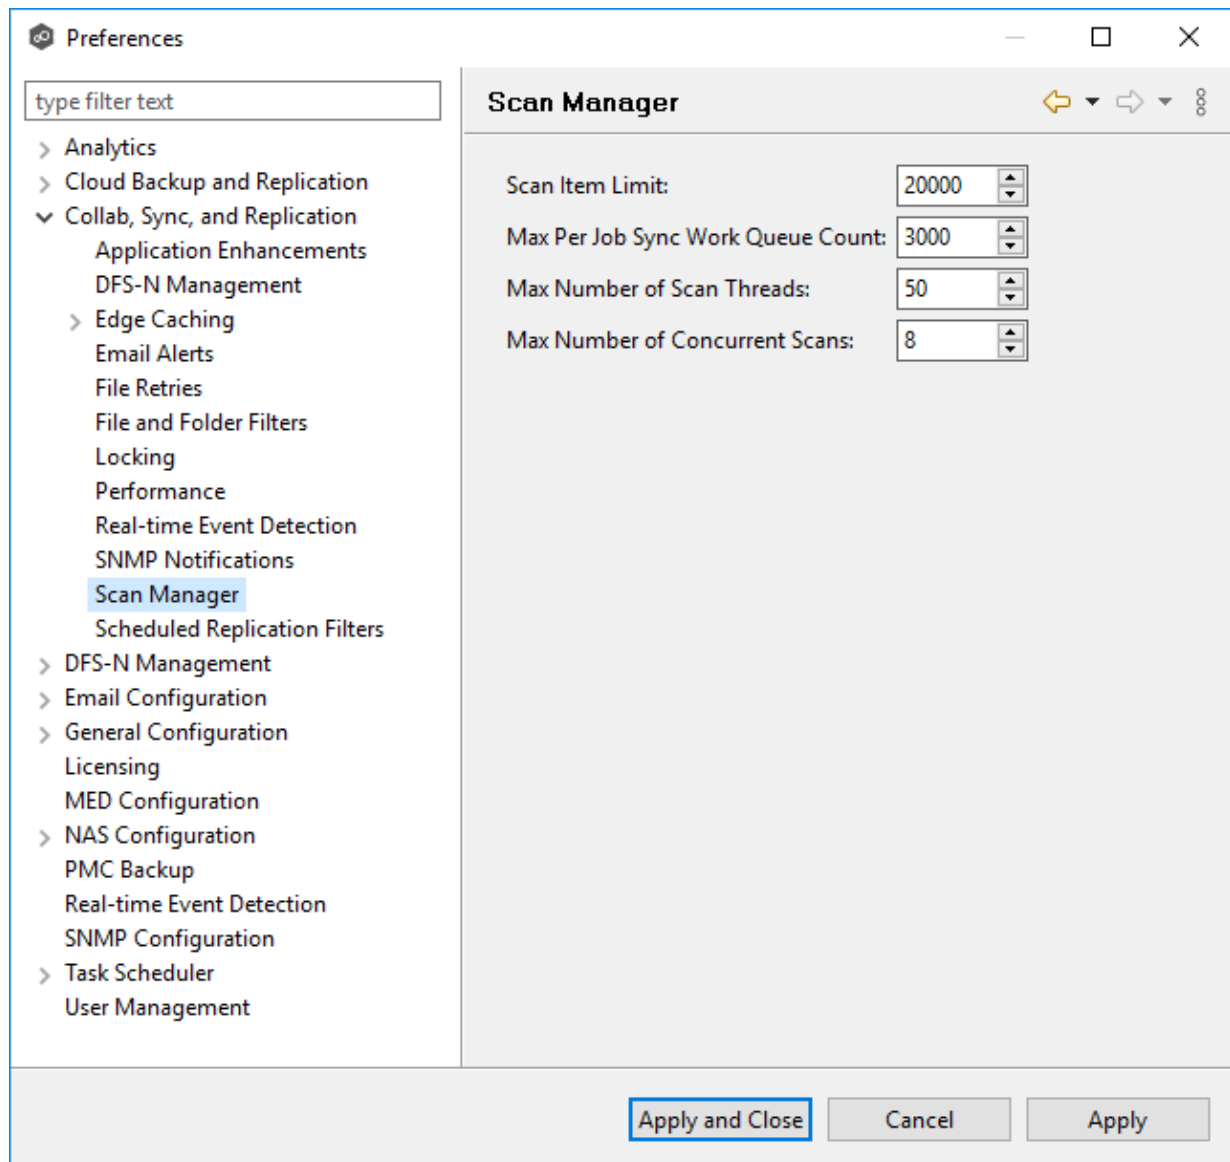


Scan Manager

Several options are available to tune the way scans are performed for File Collaboration, File Synchronization, and File Replication jobs.

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

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Scan Manager**.



3. Modify the options as needed.

Option	Description
Scan Item Limit	The maximum number of file and folder scan results that are returned in one scan iteration during a job's initial scan. This value is used to constrain the amount of memory used when performing initial scans with a large number of jobs.
Max Per Job Sync Work Queue Count	The per job maximum number of pending file synchronization tasks that are queued in memory before pausing the current scan. This value only has an effect on jobs with large numbers of files that must be synchronized during initial synchronization .
Max Number of Scan Threads	The maximum number of threads that can be created to scan folders and files. This number should be set to at least the number of jobs that you are running.
Max Number of Concurrent Scans	The maximum number of scan threads that can be actively working at the same time. This differs from the Max Number of Scan Threads in that not all created scan threads can be simultaneously doing work. For example, if 20 scan threads are configured but only 10 can run concurrently, 10 of the 20 threads will be paused at any one time, waiting for a time slot to continue working. Each of the 20 scan threads will get a chance to work in a round-robin fashion.

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

Scheduled Replication Filters

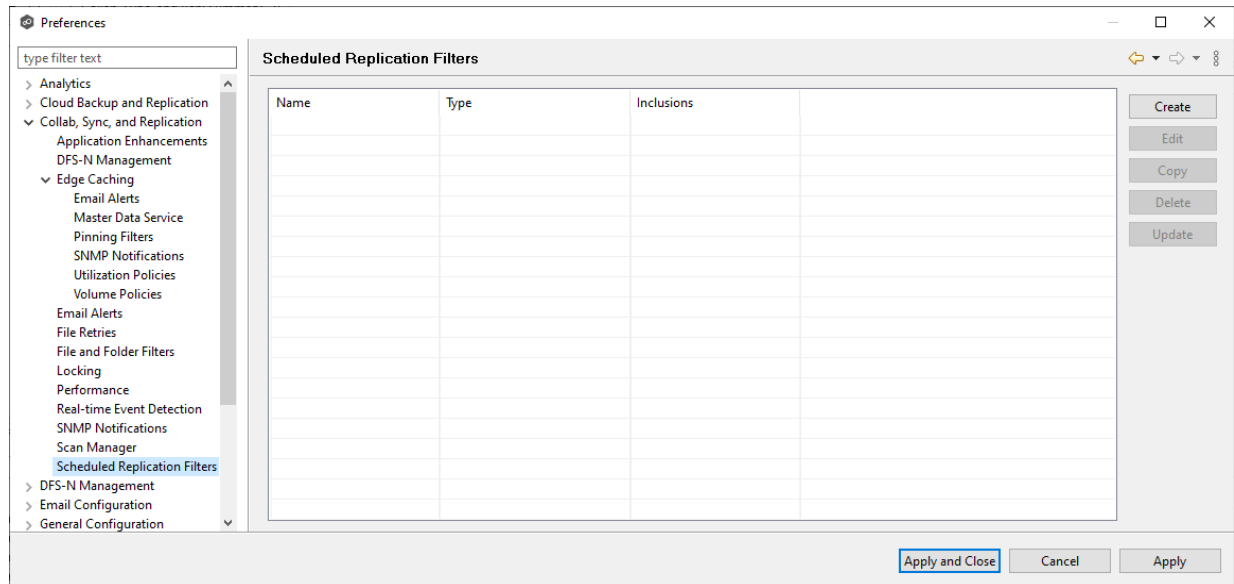
When you create a job, you can select existing scheduled replication filters to apply to the job or you can create new scheduled replication filters and apply them to the job. This [Preferences](#) page lists the existing scheduled replication filters. From this page, you can view, create, edit, and delete scheduled replication patterns. However, you cannot edit or delete a scheduled replication filter while it is applied to a job. See [Scheduled Replication](#) in the [Advanced Topics](#) section for more information about scheduled replication.

To create a scheduled replication filter:

1. Select **Open Preferences** from the **Tools** menu.

- Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Scheduled Replication Filters**.

Any existing scheduled replication filters are listed in the **Scheduled Replication** table.



- Click **Create**.

The screenshot shows a dialog box titled "Create Scheduled Replication Filter". It has a standard Windows window title bar with minimize, maximize, and close buttons. The dialog is divided into several sections:

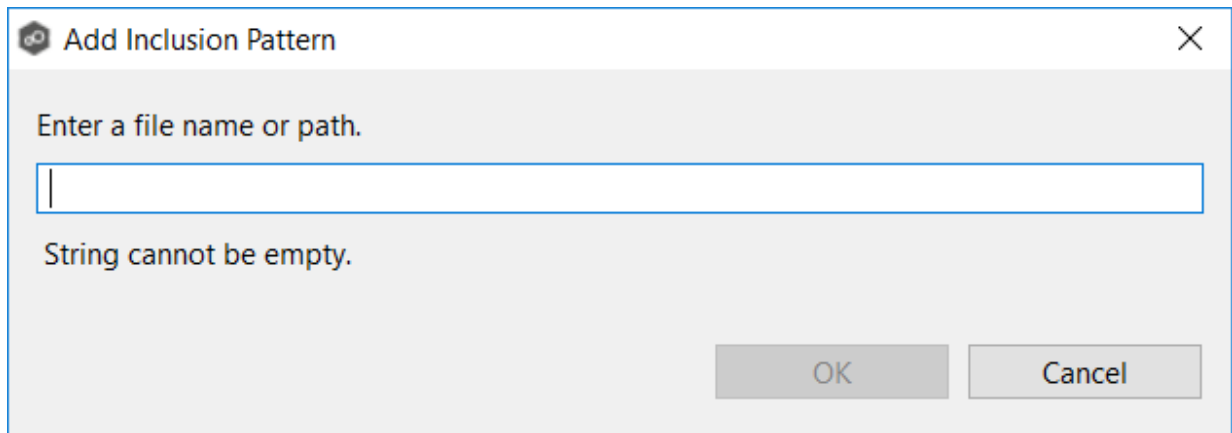
- Name:** A text input field that is currently empty.
- Filter Type:** A dropdown menu set to "Scheduled Replication".
- Included Patterns:** A large empty text area for listing patterns. Below it are three buttons: "Add", "Edit", and "Delete".
- Scheduling Options:**
 - A radio button labeled "Process every:" is selected, followed by a dropdown menu showing "1 minute".
 - An unselected radio button labeled "Process on a schedule".
 - Below "Process on a schedule", there are two radio buttons: "Daily" (unselected) and "Weekly" (selected).
 - A section labeled "Day(s):" contains a list of days with checkboxes: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday. All checkboxes are currently unchecked.
 - A section labeled "Time:" has a dropdown menu set to "None".

At the bottom right of the dialog are two buttons: "OK" and "Cancel".

4. Enter a unique name for the filter.

5. Click **Add** under **Included Patterns** to enter a filter pattern for files that you want to delay replication. Repeat to add more filter patterns.

A **filter pattern** is a character string that defines a logical expression that is evaluated to determine which files and folders match the pattern. A filter pattern can contain [complex regular expressions](#) and [wildcards](#).



The screenshot shows a dialog box titled "Add Inclusion Pattern" with a close button (X) in the top right corner. The dialog contains the text "Enter a file name or path." above a text input field. Below the input field, the text "String cannot be empty." is displayed. At the bottom right of the dialog are two buttons: "OK" and "Cancel".

6. Click **OK**.

The pattern appears in the **Included Patterns** field.

Create Scheduled Replication Filter

Name:

Filter Type:

Included Patterns

Scheduling Options

Process every:

Process on a schedule

Daily Weekly

Day(s):

Sunday

Monday

Tuesday

Wednesday

Thursday

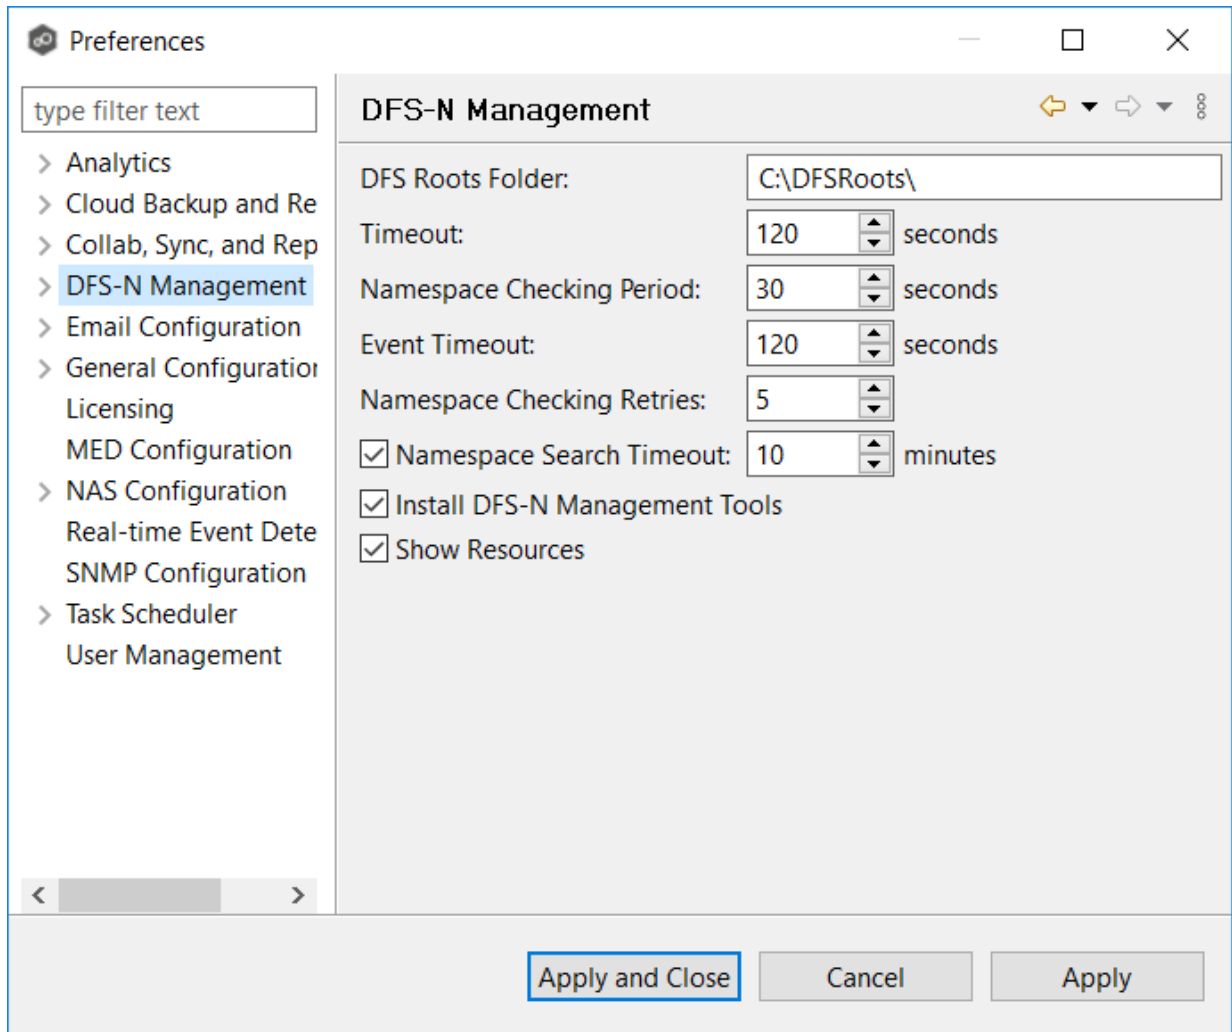
Friday

Saturday

Time:

7. Select a scheduling option:

- **Interval** - Process at a specified interval.



3. Modify settings as needed.

Setting	Description
DFS Roots Folder	The default folder for namespace roots is C:\DFSRoots\. If you want a different folder to be used, enter the local path to the folder.
Timeout	Enter the number of seconds to wait for a response from any Agent. The default value is 120 seconds.
Namespace Checking Period	Enter the number of seconds to delay between checking namespace information calls. This check catches any changes made to a namespace using the Microsoft DFS Management tool.

Setting	Description
	Selecting a low value will negatively affect performance but will reflect changes to the user interface more quickly. The default value is 120 seconds.
Event Timeout	Enter the number of seconds to wait before marking an event containing DFS namespace information from the Agent as timed out. The default value is 120 seconds.
Namespace Checking Retries	Enter the maximum number of times for checking namespace information if the namespace is not found. Once the maximum number is exceeded, the job is stopped. The default value is 5 retries.
Namespace Search Timeout	When a user tries to import a namespace, PeerGFS searches for the namespace. This may take some time, depending on the environment. Enter the number of minutes to wait before timing out. The default value is 10 minutes.
Install DFS-N Management Tools	Select this option if you want Microsoft's DFS-N Management tools installed when creating or importing a namespace.
Show Resources	Select this option if you want to display individual namespace folders under each namespace in the Jobs view.

- Click **Apply and Close** or **Apply**.

Email Alerts

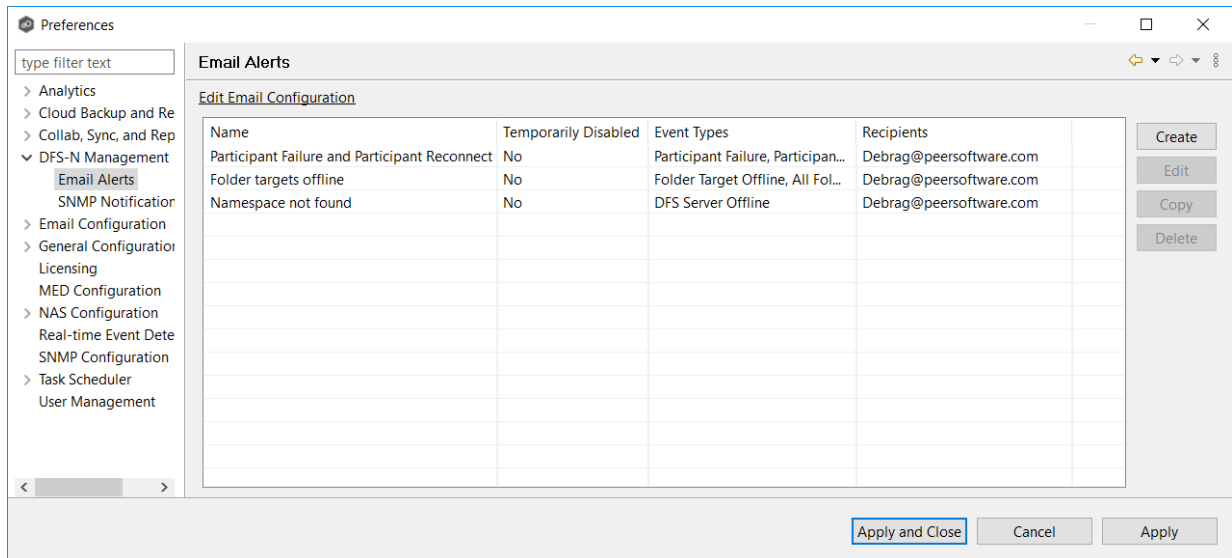
When you create a job, you can select existing email alerts to apply to the job or you can create new email alerts and apply them to the job. This [Preferences](#) page lists the existing email alerts. From this page, you can view, create, edit, and delete email alerts. However, you cannot edit or delete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

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

To create an email alert:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **DFS-N Management** in the navigation tree, and then select **Email Alerts**.

Any existing DFS-N Management email alerts are listed in the **Email Alerts** table.



3. Click the **Create** button.

The **Create Email Alert** dialog appears.

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

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

Event Type	Description
Job Start	Sends an alert when the job starts.
Job Stop	Sends an alert when the job stops.
Job Failure	Sends an alert when the job stops unexpectedly.
Participant Failure	Sends an alert when the Management Agent job disconnects or stops responding.

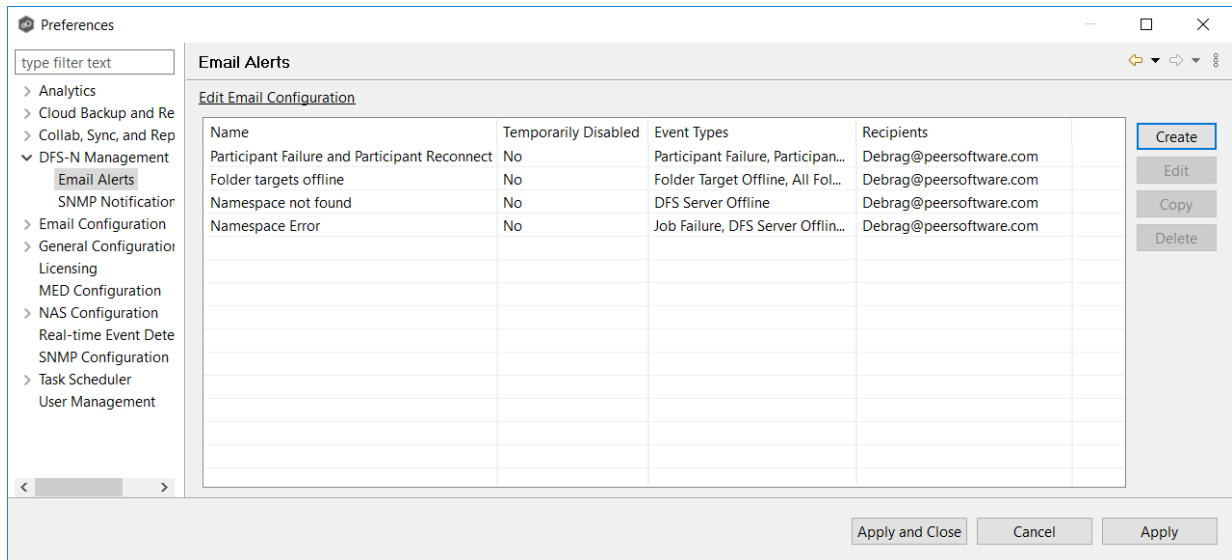
Event Type	Description
Participant Reconnect	Sends an alert when the Management Agent reconnects.

6. Select the DFS-N event types.

Event Type	Description
Namespace Offline	Sends an alert when a namespace goes offline.
Namespace Not Found	Sends an alert when a namespace is not found.
Folder Target Offline	Sends an alert when a folder target goes offline.
All Folder Targets Offline	Sends an alert when all folder targets go offline
DFS Server Offline	Sends an alert when a DFS server goes offline.

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

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



8. Click **OK**.

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

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

SNMP Notifications


When you create a job, you can select an existing SNMP notification to apply to the job or you can create a new notification and apply it to the job. You cannot edit or delete an SNMP notification while it is applied to a job. See [SNMP Notifications](#) in the [Basic Concepts](#) section for more information about SNMP notifications.

Note that the [SNMP source address, destination, and prefix must be configured](#) before notifications can be set.

To create an SNMP notification:

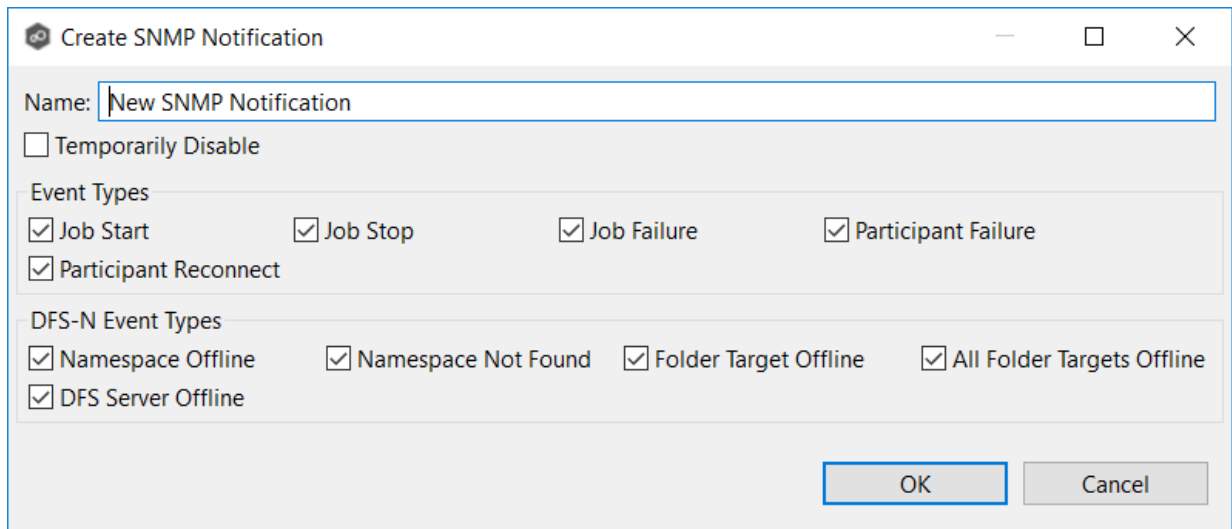
1. Select **Open Preferences** from the **Tools** menu.
2. Expand **DFS-N Management** in the navigation tree, and then select **SNMP Notifications**.

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



3. Click the **Create** button.

The **Create SNMP Notification** dialog appears.



4. Select the types of events that will trigger the generation of an SNMP trap.

Event Type	Description
Job Start	Sends a notification when the DFS-N Management job starts.
Job Stop	Sends a notification when the DFS-N Management job stops.

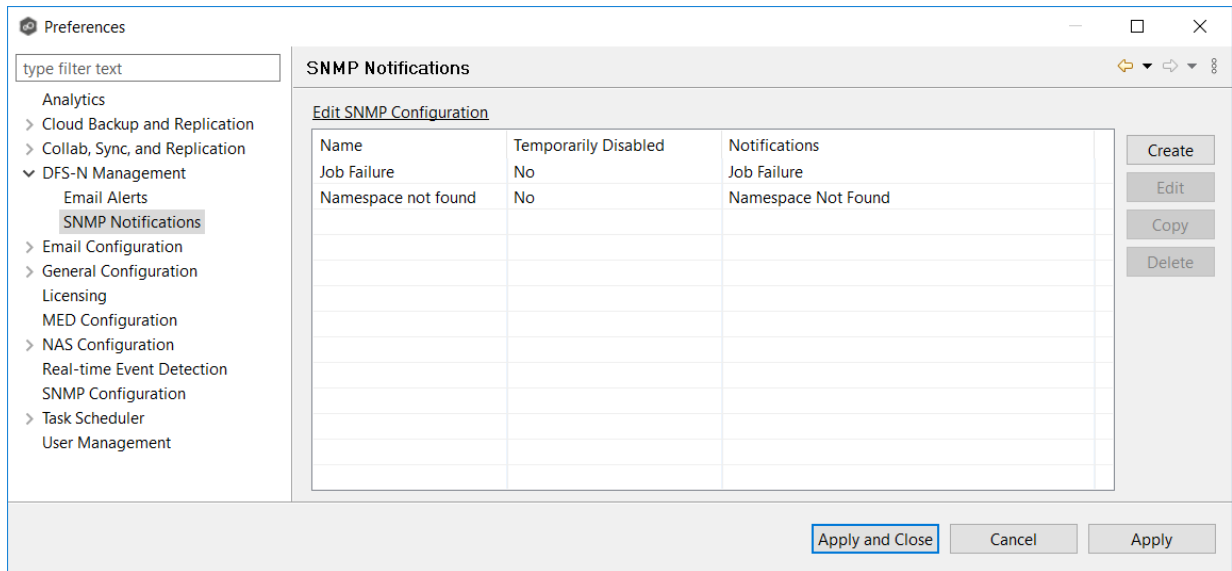
Event Type	Description
Job Failure	Sends a notification when the DFS-N Management job stops unexpectedly.
Participant Failure	Sends a notification when the Management Agent of the DFS-N Management job disconnects or stops responding.
Participant Reconnect	Sends a notification when the Management Agent of the DFS-N Management job reconnects.

5. Select the DFS-N event types that will trigger the generation of an SNMP trap.

Event Type	Description
Namespace Offline	Sends a notification when a namespace goes offline.
Namespace Not Found	Sends a notification when a namespace is not found.
Folder Target Offline	Sends a notification when a folder target goes offline.
All Folder Target Offline	Sends a notification when all folder targets go offline
DFS Server Offline	Sends a notification when a DFS server goes offline.

6. Click **OK**.

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



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

Email Configuration

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

To configure email settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **Email Configuration** in the navigation tree.

The screenshot shows the 'Preferences' dialog box with the 'Email Configuration' section selected. The left sidebar lists various configuration categories, with 'Email Configuration' highlighted. The main area contains the following settings:

- SMTP Email Configuration**
 - *SMTP Host: smtp.office365.com
 - *SMTP Port: 587
 - Encryption:
 - *Encryption Type: TLS
 - *Username: debrag@peersoftware.com
 - *Password: [masked]
 - *Sender Email: debrag@peersoftware.com
 - Use Recommended Office 365 Settings:
 - Test Email Settings (button)
- Batch Email Alerts for Quarantined Files**
 - Batch Quiet Period (in seconds): 60
 - Maximum Number of Alerts: 1000

At the bottom of the dialog, there are three buttons: 'Apply and Close', 'Cancel', and 'Apply'.

3. Enter values for the following fields:

Field	Description
SMTP Host	Enter the host name or IP address of the SMTP mail server through which Peer Management Center will send emails.
SMTP Port	Enter the TCP/IP connection port (default is 25 and 465 for encryption) on which the mail server is hosting the SMTP service. We recommend that you leave the default setting unless your email provider specifies otherwise.
Encryption	Select this checkbox if the SMTP mail server requires an encrypted connection.
Encryption Type	If encryption is enabled, an encryption method must be selected. TLS and SSL are the available options. If you do not know which one your mail server requires, try one, and then the other.
Username	Enter the user name to authenticate as on the SMTP mail server.
Password	Enter the password for the user specified above.
Sender Email	Enter the email address to appear in the From field of any sent emails. This email address sometimes needs to have a valid account on the SMTP mail server.
Use Recommended Office 365 Settings	Select this checkbox if you are connecting to an Office 365 SMTP server. Follow Microsoft's Direct Send recommendations to set up email configuration with an Office 365 SMTP server.

- (Recommended) Click **Test Email Settings**, enter an email address, and then click **OK**.

It is highly recommended that you test your SMTP settings before saving them. You will be prompted for an email address to send the test message to. Upon submission, Peer Management Center will attempt to send a test message using the specified settings.

- Enter values for the fields in the **Batch Email Alerts for Quarantined Files** section:

Field	Description
Batch Quiet Period (in seconds)	Enter the number of seconds to wait before releasing a batch of alerts.
Maximum Number of Alerts	Enter the maximum number of alerts that should be sent in a single email.

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

General Configuration

The **General Configuration** settings affect the overall operation of Peer Management Center, Peer Agents, the Peer Broker, and other general operations. They are not specific to jobs or job types.

You can modify the following settings:

[General Configuration](#)

[Agent Connectivity](#)

[Broker Configuration](#)

[Email Alerts](#)

[Proxy Configuration](#)

[Software Updates](#)

[Tags Configuration](#)


[Web and API Configuration](#)

General Configuration

To modify General Configuration settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **General Configuration** in the navigation tree.

The first page of the General Configuration options is displayed.



3. Modify the first four settings as needed:

Setting	Description
Launch Dashboard at Start	Select this option if you want the Dashboard to be automatically displayed when Peer Management Center is started.

Setting	Description
Launch Summary Views at Start	Select this option if you want the Summary views to be automatically displayed when Peer Management Center is started. Summary views will be displayed for all job types, even for job types without currently running jobs.
Always Run Tasks in Background	Select this option to run tasks like log gathering and Agent updates in the background, preventing these tasks from blocking the use of the Peer Management Center client while they run.
Auto Expand Job Resources	Select this option if you want all jobs with associated resources to start expanded in the Jobs view. Currently only available for Cloud Backup and Replication jobs and DFS-N Management jobs.
Display Job Types When No Jobs Are Configured	Select this option if you want to display a job type in the Jobs view, even when no jobs of that type have been configured.

4. Select options for alerts regarding the operation of Peer Management Center in the [Alerts view](#):

Option	Description
Severity	Select one of these options: <ul style="list-style-type: none"> • INFO • DEBUG • TRACE
Auto Display Alerts View	Select this option if you want the alerts to be automatically displayed in the Alerts view .

5. Select options for managing browsing files and folders on remote file systems in the **Browsing Files/Folders** section:

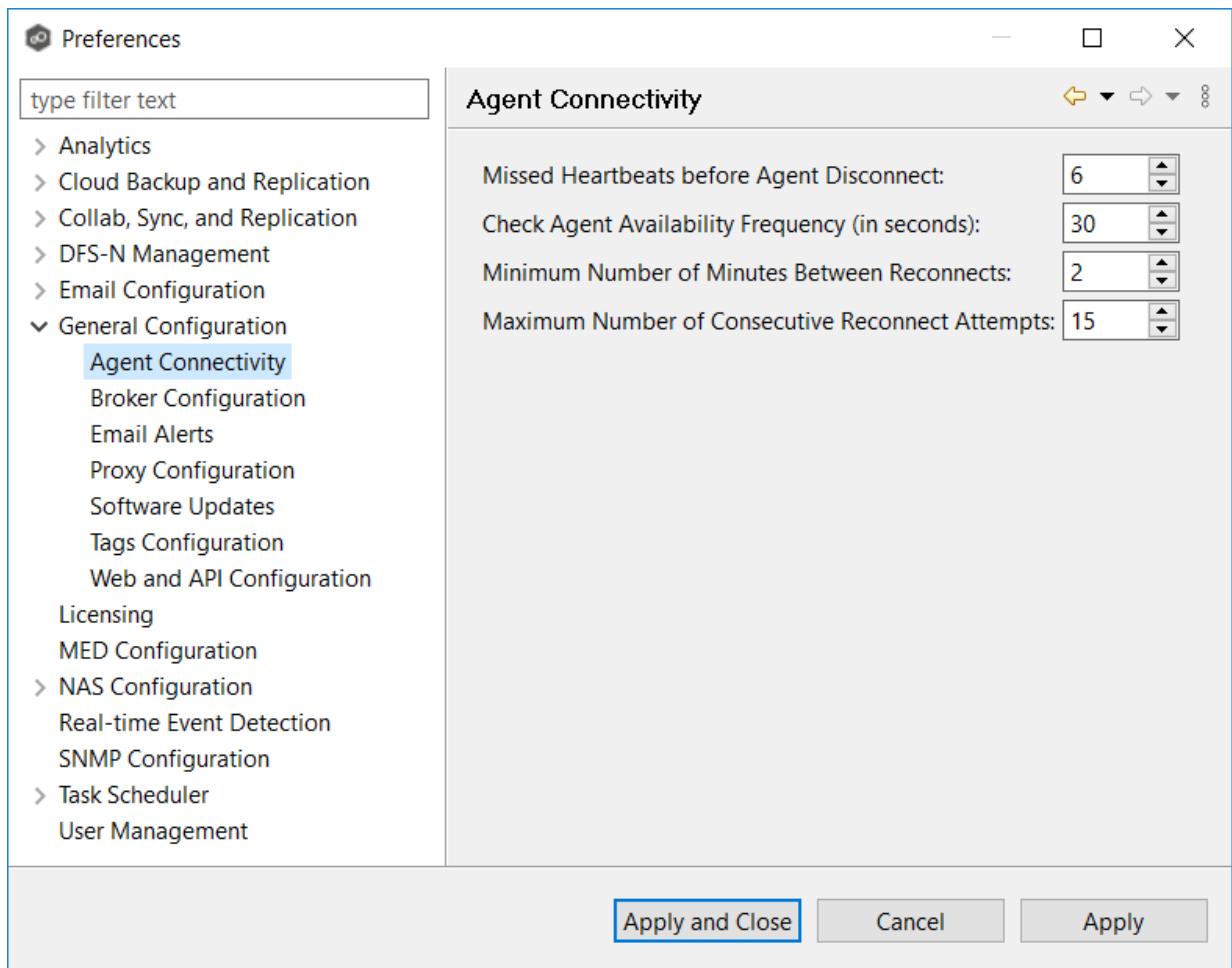
Option	Description
Remote Browser Page Size	Enter the maximum page size for the remote file system browser. This browser is used for selecting paths during the creation of most new jobs.
Show System Folders	Select this checkbox to show system folders in the remote file system browser.
Show Hidden Folders	Select this checkbox to show hidden folders in the remote file system browser.

6. (Optional) Enter the name of your PMC server or environment in the **Environment Name** field; if left blank, reports and dashboards will use the name of the PMC server.
7. Click **Apply and Close** or **Apply**.

Agent Connectivity

To modify Agent Connectivity settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **General Configuration** in the navigation tree, and then select **Agent Connectivity**.



3. Modify the settings as needed:

Option	Description
Missed Heartbeats before Agent Disconnect	Enter the maximum number of heartbeats that can be missed on a host before Peer Management Center labels the Agent as disconnected. If a running job hits a timeout when communicating with a specific Agent, Peer Management Center will check this status to decide if the Agent should be dropped from the job.
Check Agent Availability Frequency (in seconds)	Enter the frequency (in seconds) that Peer Management Center should check whether an Agent is back online.
Minimum Number of Minutes Between Reconnects	Enter the minimum number of minutes that must elapse before Peer Management Center attempts to retry reconnecting to the Agent.
Maximum Number of Consecutive Reconnect Attempts	Enter the maximum number of attempts that Peer Management Center tries to reintegrate a previously connected agent into one or more jobs. Once the maximum number of attempts has been reached, you must manually reintegrate the Agent into affected jobs, typically by restarting the affected jobs.

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

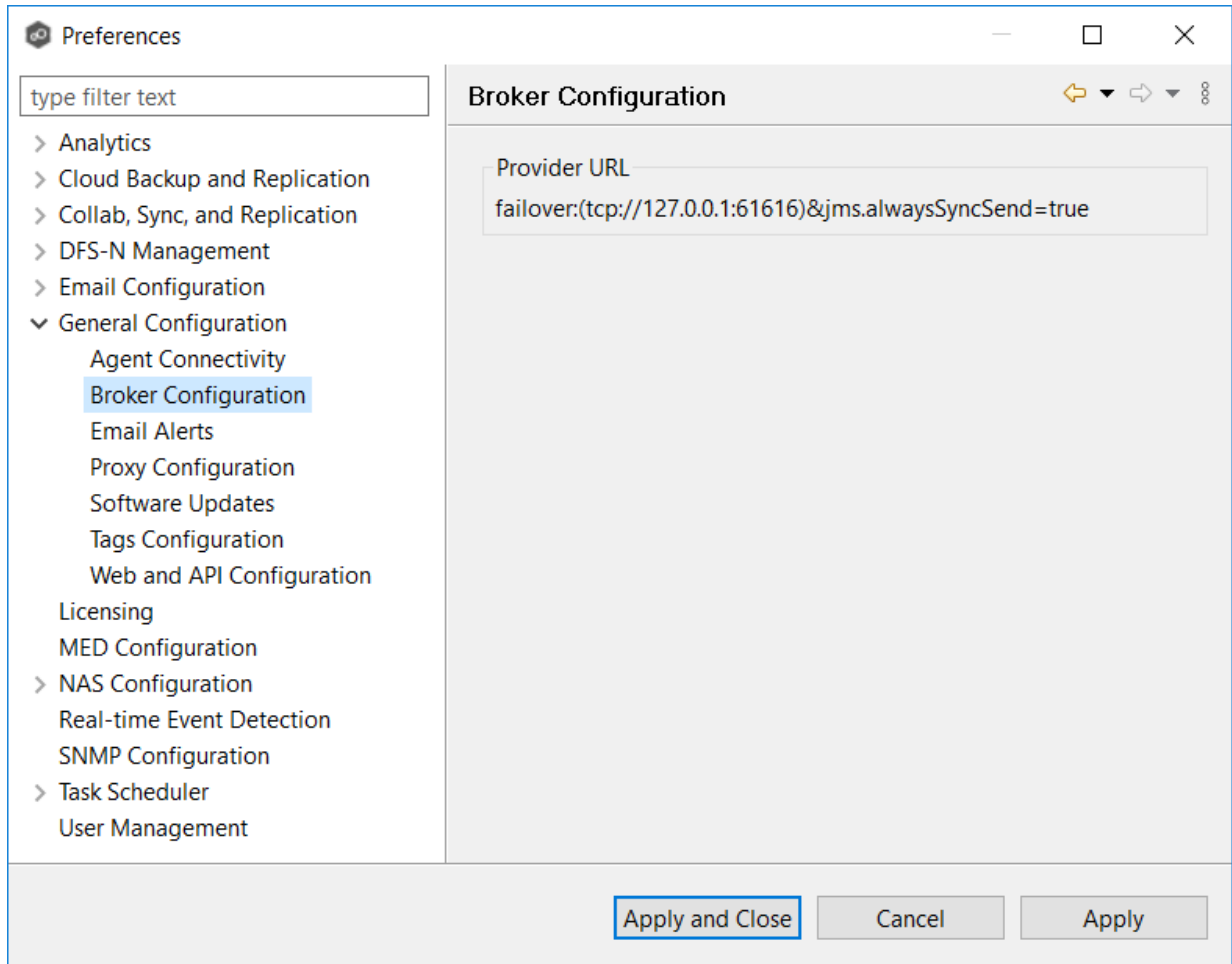
Broker Configuration

The **Broker Configuration** page displays a non-editable field that shows the URL used by the Peer Management Center service to connect to the Broker service.

To view the Broker Configuration URL:

1. Select **Open Preferences** from the **Tools** menu.

2. Expand **General Configuration** in the navigation tree, and then select **Broker Configuration**.



3. Click **Apply and Close** or **Apply**.

Email Alerts

System email alerts notify recipients when certain types of system events occur, for example, low memory, low disk space, disconnected agents. This [Preferences](#) page lists the existing system email alerts. From this page, you can create, edit, and delete system email alerts. You can also disable and enable alerts. See [Email Alerts](#) in the [Basic Concepts](#) section for more information about email alerts.

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

Create Email Alert

Name:

Temporarily Disable

Event Types

Low Memory Low PMC Disk Space Agent Install Disk Space Agent Disconnects

Licensing Warnings

Recipients

Enter email, contact, or distribution list:

Recipients:

4. Enter a name for the alert.
5. Select the **Enable** checkbox if you want to enable the alert.
If you choose not to enable the alert, you can enable it later.
6. Select the type of events for which you want alerts sent:

Event Type	Description
Low Memory	Sends an alert when Peer Management Center or connected Agent services are low on memory.
Low PMC Disk Space	Sends an alert when the space on the disk where Peer Management Center software is installed running low.
Agent Install Disk Space	Sends an alert when the space on the disk where the Peer Agent software is installed is running low.
Agent Disconnects	Sends an alert whenever an Agent is disconnected.
License Warnings	Sends an alert when a license has expired or when a license violation is about to occur (for example, when storage usage reaches 95% of maximum storage and when storage usage exceed license limits).

7. Enter alert recipients, and then click **Add to List**.
8. Click **Apply and Close** or **Apply**.

The new alert is listed in the **Email Alerts** table.

The screenshot shows the 'Preferences' dialog box with the 'Proxy Configuration' section selected in the left-hand menu. The 'Proxy Configuration' section contains the following fields and options:

- IP Address: [Text Input Field]
- Port: [Text Input Field]
- Use Authentication
- Domain: [Text Input Field]
- Username: [Text Input Field]
- Password: [Text Input Field] Show Password

At the bottom of the dialog box, there are three buttons: 'Apply and Close' (highlighted with a blue border), 'Cancel', and 'Apply'.

3. Enter values for the following fields:

Field	Description
IP Address	Enter the IP address or fully qualified domain name of the proxy server.
Port	Enter the port number.
User Authentication	Select if your proxy server requires authentication.

4. If your proxy server requires authentication, select the **Use Authentication** checkbox, and then supply the necessary values:

Field	Description
Domain	Enter the domain name on the proxy server.
Username	Enter the user name for the proxy server.
Password	Enter the password for the proxy server.

5. Click **Apply and Close** or **Apply**.

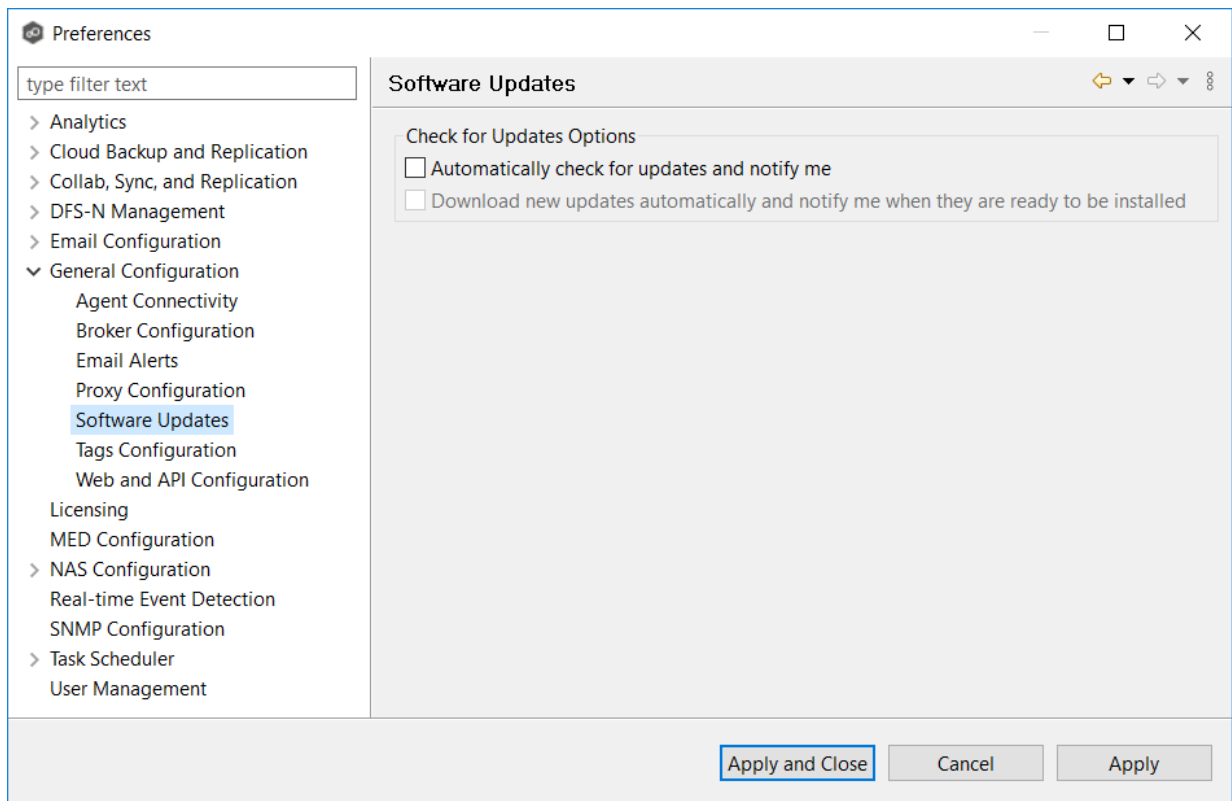
Software Updates

You can configure Peer Management Center to automatically check for updates and download the updates. Peer Management Center checks for updates every evening at 11 p.m. local time. Only minor updates are automatically downloaded; if a major update is available, a notification appears. Major releases require a new license key and must be requested from Peer Software Support.

You can also manually check for updates. See [Updating Peer Management Center](#) for information about manually checking for updates.

To configure Peer Management Center to automatically check for updates:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **General Configuration** in the navigation tree, and then select **Software Updates**.

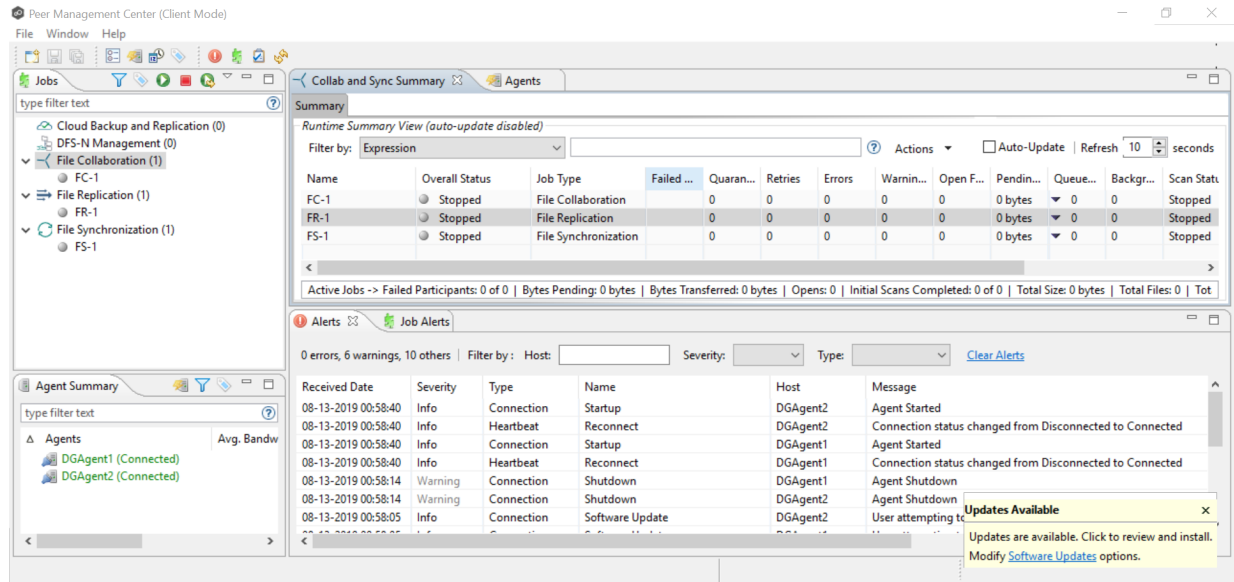


3. Select update options:

- **Automatically check for updates and notify me** - Select this option if you want to automatically check for updates.
- **Download new updates automatically and notify me when they are ready to be installed** - Select this option if you want to automatically check for and download available updates.

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

Whenever updates are available, a notification appears in the lower right corner of Peer Management Center.



5. Click the notification to review and proceed with the update. See [Updating Peer Management Center](#) for details.


Tags Configuration

The **Tags Configuration** page in Preferences is the starting place for creating [tags](#) and categories that can later be assigned to resources. See [Assigning Tags](#) for more information about assigning to resources.

To create a tag:


1. Select **Open Preferences** from the **Tools** menu.
2. Expand **General Configuration** in the navigation tree, and then select **Tags Configuration**.

Any existing tags are listed in the **Tags** table.



3. Click the **Create** button.

The **New Tag** dialog appears.



4. Enter a name for a tag.
5. Select a category or create a new category.
6. Click **OK**.

The tag appears in the **Tags** table.

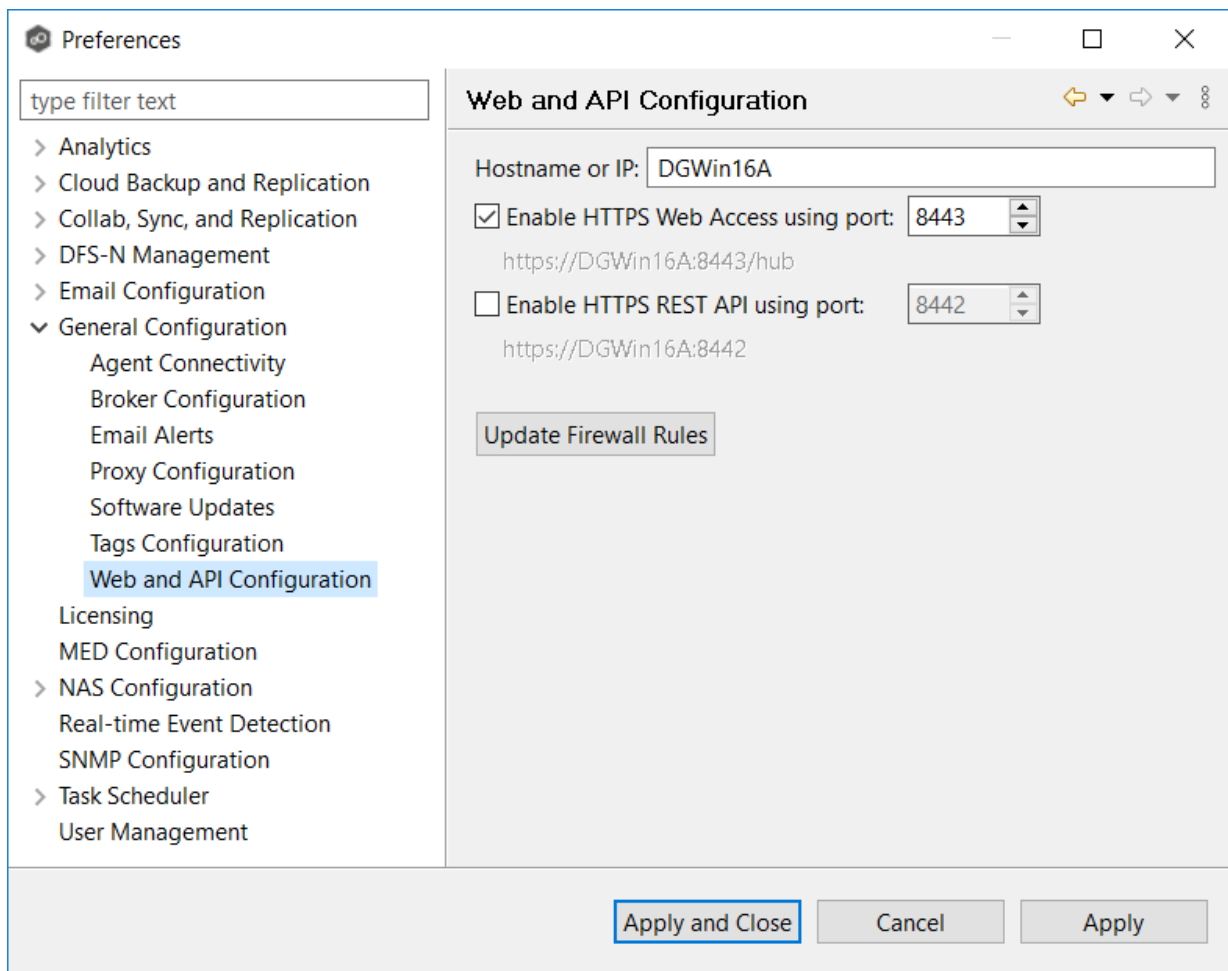
7. Click **Apply and Close** or **Apply**.

Web and API Configuration

As part of the Peer Management Center installation process, you are prompted to configure access to the web and API services. If you do not enable access during the initial installation or want to modify settings at a later date, you can modify them in Preferences.

To modify web and API settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **General Configuration** in the navigation tree, and then select **Web and API Configuration**.



3. Modify the configuration options.

Option	Description
Hostname or IP	Enter the hostname or IP address via which the services can be accessed: <ul style="list-style-type: none">• Enter localhost or 127.0.0.1 if you want the services to be accessible only to users of the local server via the loopback interface.• Enter 0.0.0.0 to make the services accessible via all network interfaces.• Enter a specific IP address to restrict access to a specific network interface.
Enable HTTPS Web Access	Select this checkbox to enable HTTP access to the web service using the specified port.
Enable HTTPS REST API	Select this checkbox to enable HTTPS access to the REST API service using the specified port.

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

Licensing

Peer Global File System is licensed by the number of unique [participants](#) and by the number of terabytes in the [watch sets](#).

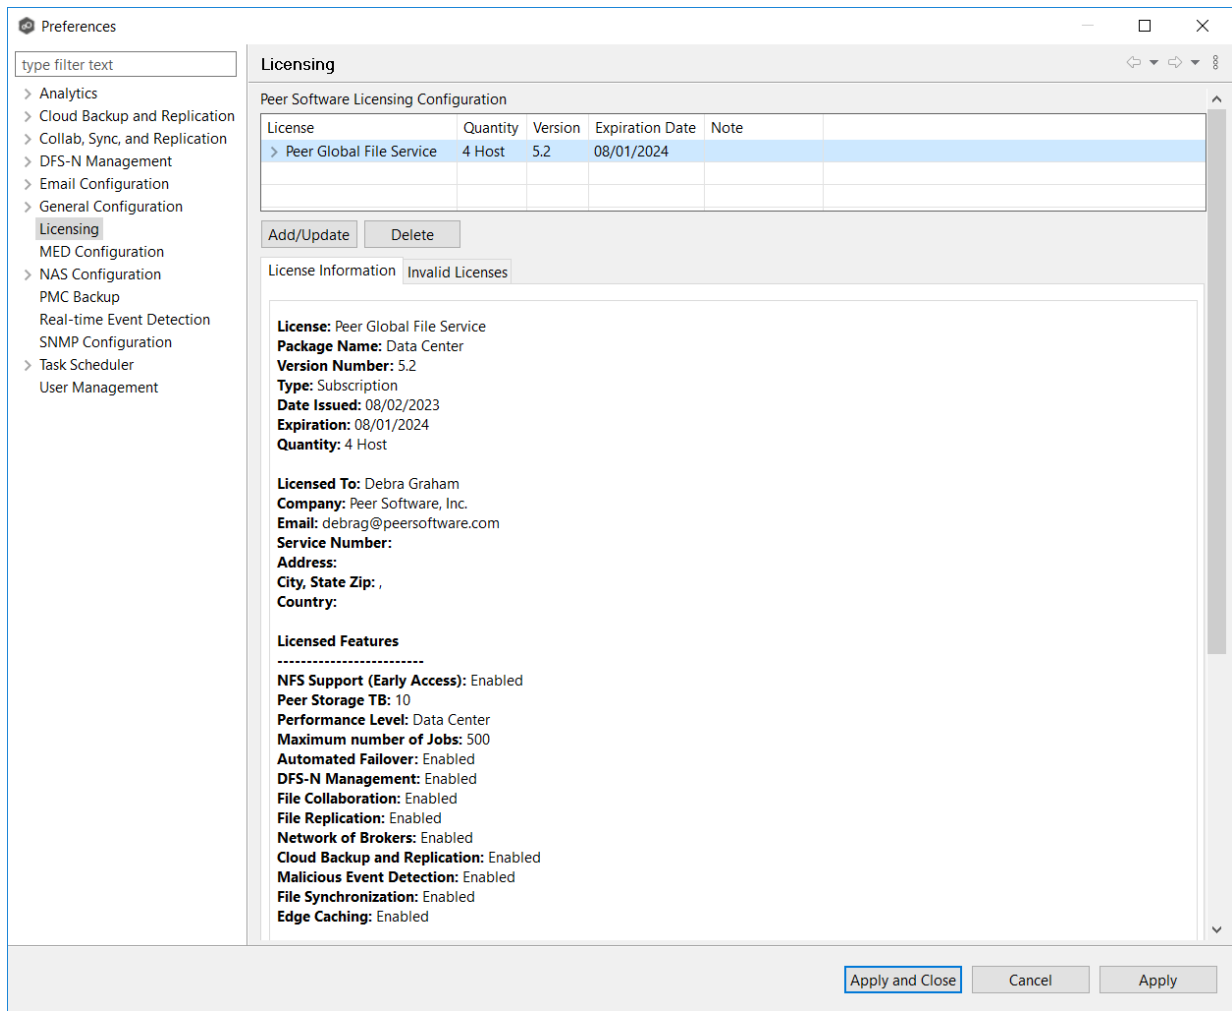
Installing or Upgrading a License File

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. Select **Open Preferences** from the **Tools** menu.
2. Select **Licensing** in the navigation tree.

Existing valid licenses are listed in the **Peer Software Licensing Configuration** table.



The screenshot shows the 'Preferences' window with the 'Licensing' section selected. The 'Peer Software Licensing Configuration' table is displayed, showing one license entry:

License	Quantity	Version	Expiration Date	Note
> Peer Global File Service	4 Host	5.2	08/01/2024	

Below the table, the 'Add/Update' button is highlighted. The 'License Information' tab is active, showing the following details:

License: Peer Global File Service
Package Name: Data Center
Version Number: 5.2
Type: Subscription
Date Issued: 08/02/2023
Expiration: 08/01/2024
Quantity: 4 Host

Licensed To: Debra Graham
Company: Peer Software, Inc.
Email: debrag@peersoftware.com
Service Number:
Address:
City, State Zip: ,
Country:

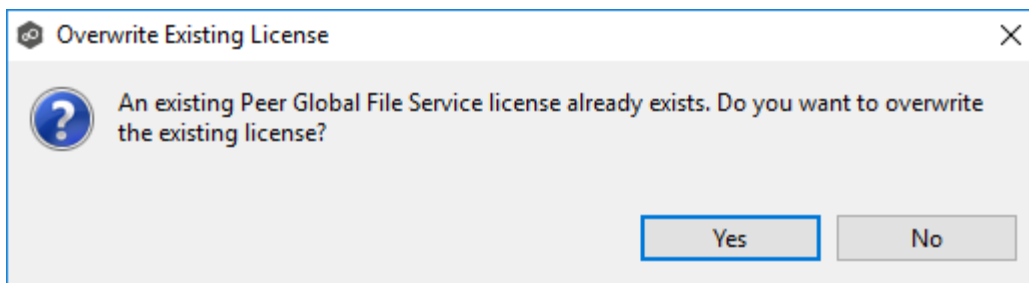
Licensed Features

NFS Support (Early Access): Enabled
Peer Storage TB: 10
Performance Level: Data Center
Maximum number of Jobs: 500
Automated Failover: Enabled
DFS-N Management: Enabled
File Collaboration: Enabled
File Replication: Enabled
Network of Brokers: Enabled
Cloud Backup and Replication: Enabled
Malicious Event Detection: Enabled
File Synchronization: Enabled
Edge Caching: Enabled

At the bottom of the window, the 'Apply and Close' button is highlighted.

3. Click the **Add/Update** button to browse for a license file.
4. Select the license file, and then click **Open**.

If you are prompted with a message that an existing license already exists, click **Yes** to overwrite the existing license.



After successful installation of the license, it is listed in the table, along with the license quantity, version, and an expiration date (if applicable). You can now create, configure, and run jobs using the new license.

Note: You will need to restart existing jobs if any of the following applies:

- The software version is different (typically when upgrading to a new version).
 - The software package level is different.
 - The new license is insufficient for the number of existing hosts.
5. Click the license in the table to view details about the license.
 6. Click **Apply and Close** or **Apply**.

Deleting a License File

To delete a license.

1. From the **Windows** menu, select **Preferences**.
2. Select **Licensing** in the navigation tree.
3. Select the license you want to delete.
4. Click the **Delete** button.

Any job types enabled by that license will be hidden from Peer Management Center.

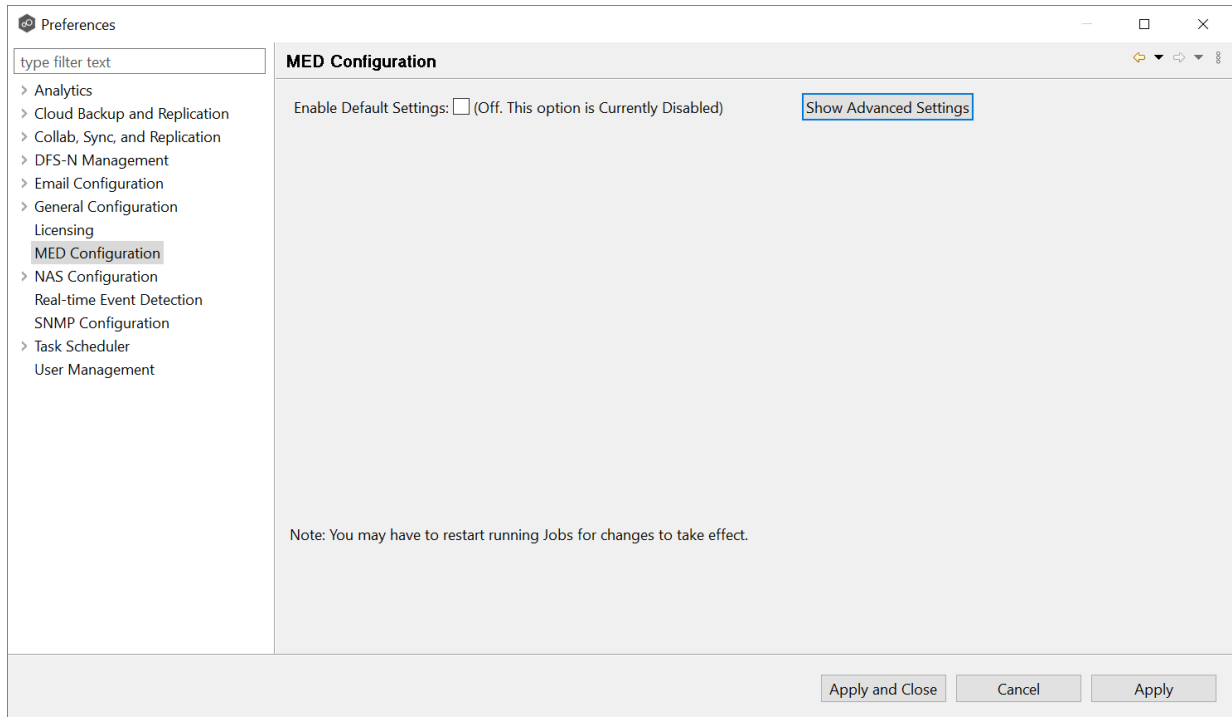
MED Configuration

Peer's Malicious Event Detection (MED) real-time engine can spot unwanted activity being executed on storage platforms by ransomware, viruses, malware, hackers, or rogue users. MED technology provides alerting capabilities, as well as the ability to minimize the amount of encrypted or deleted content from being replicated to remote locations. Once MED is enabled and jobs are restarted, these capabilities apply to all jobs. For more information, see our knowledge base article [Introduction to Peer MED](#).

Peer MED deploys three different mechanisms for spotting malicious activity, each of which can be enabled and tuned independently. These settings are configured on a global level.

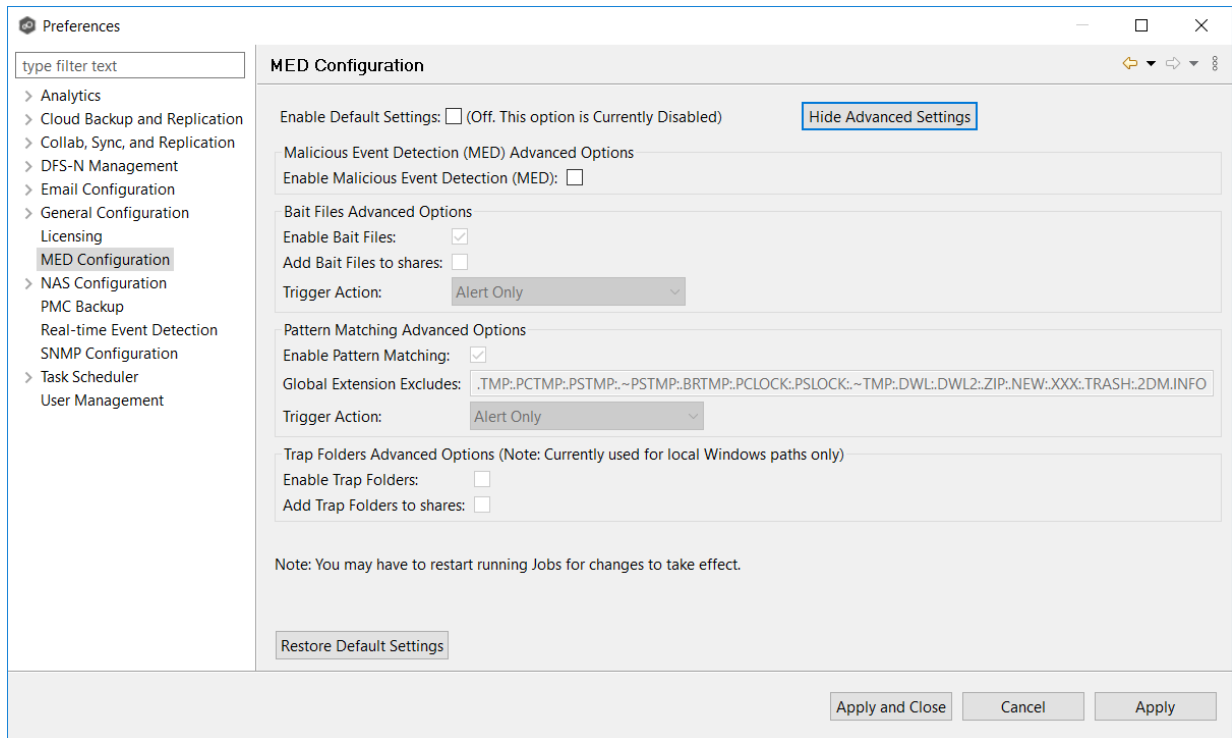
To modify MED settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **MED Configuration** in the navigation tree.



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

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



4. Modify the options as needed:

- [Primary MED Options](#)
- [Bait File Advanced Options](#)
- [Trap Folders Advanced Options](#)

5. Click **Apply and Close** or **Apply**.

Primary MED Options

The main options are as follows:

Option	Description
Enable Default Settings	Enables/disables Peer MED using default settings. By default, all three MED mechanisms are enabled.
Show/Hide Advanced Settings	Shows/hides options for each of the three MED mechanisms.

Option	Description
Enable Malicious Event Detection (MED)	The master on/off switch for MED. If unchecked, all MED mechanisms will be disabled.
Restore Default Settings	Restores all defaults across the three MED mechanisms.

Bait File Advanced Options

Bait files are files of common types, inserted into the file system in a way that hides them from users. Though hidden, these bait files are likely to be accessed by automated processes (like ransomware) or by mass deletions of entire folder structures. As soon as these files are touched, an action is triggered.

The options for bait files are:

Option	Description
Enable Bait Files	Enables/disables bait file creation and monitoring.
Add Bait Files to shares	At the start of each job, creates bait files under the root of each participant's watch directory. To see the watch directory for a job, review Host Participants and Directories .
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.

Pattern Matching Advanced Options

These options

The options for pattern matching are:

Option	Description
Enable Pattern Matching	Enables/disables pattern matching creation and monitoring.
Global Extensions Exclude	A list of file extensions to exclude from all patterns. This list can also be updated by right-clicking on specific MED alerts within the PMC's Job Alerts view.
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 when malicious activity is detected. These actions are:

Action	Description
Alert Only	Triggers an alert in Peer Management Center. If email alerts are configured for MED Alerts and enabled for a job, an email will also be sent. See Email Alerts in the Basic Concepts section for more information about email alerts. If SNMP traps are configured for MED Alerts and enabled for a job, an SNMP trap will also be sent. See SNMP Notifications in the Basic Concepts section for more information about SNMP notifications.
Alert and Disable Host	Triggers an alert while also removing the afflicted Agent from the job in which the malicious activity was detected. Once disabled, Agents will need to be manually re-enabled for collaboration to resume. See Re-enabling a Disabled Agent Within a Job for details.
Alert and Stop Job	Triggers an alert while also stopping the job where the malicious activity was detected. Jobs will need to be restarted in order for collaboration to resume.
Alert, Disable Host and Stop	Triggers an alert, removes the afflicted Agent from the job where the malicious activity was detected, and stops the job.

Action	Description
Job	This option is the most aggressive and will require administrators to re-enable Agents as well as restart jobs. See Re-enabling a Disabled Agent Within a Job for details.

An example of an alert as displayed in Peer Management Center is as follows:

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 appld=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

Click outside of popup to close

Trap Folders Advanced Options

On Windows file servers, Peer MED can be configured to create hidden, recursive folders that attempt to trap or slowdown ransomware as it enumerates a folder structure. As with the bait files, these folders cannot be seen by users but will be accessible by automated processes. If bait files (above) are enabled, a bait file will be placed within each trap folder, and an action will be triggered as soon as these files are touched.

Options for trap folders are:

Option	Description
Enable Trap Folders	Enables/disables the creation and monitoring of trap folders.
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.

Storage Configuration

Enter topic text here.

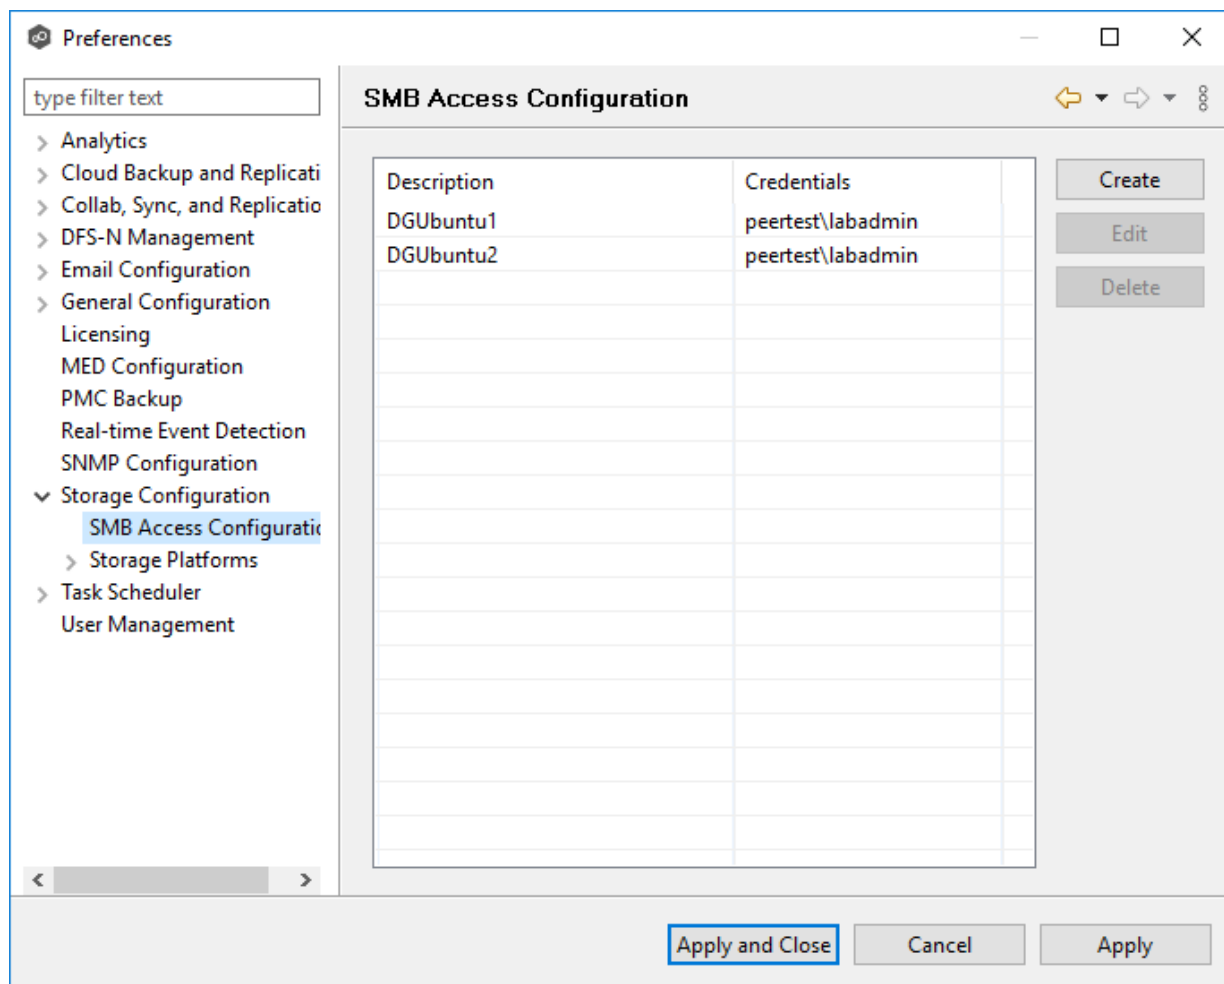
SMB Access Configuration

SMB access configuration enables a Linux Agent to access data on the file server using the SMB file sharing protocol. It is required in the following situations when creating a multi-protocol job:

- Synchronizing Windows-style permissions
- Synchronizing both Windows-style permissions and Linux-styled permission (mixed permissions)
- Using Nutanix Files.

To configure SMB access settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **SMB Access Configuration** in the navigation tree.



3. Click **Create**.
4. In the **New SMB Access Configuration** dialog, supply the requested information.

The 'New SMB Access Configuration' dialog box contains three input fields: 'Description:' (empty), 'Domain Username:' (containing 'Domain\Username'), and 'Domain Password:' (empty with a visibility toggle icon). 'OK' and 'Cancel' buttons are at the bottom.

Field	Description
Description	Enter a description of the intended use of the configuration.
Domain Username	Enter the domain username in the format: <code>Domain\Username</code>
Domain Password	Enter the password for the domain account.

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

a.

Storage Platforms

This section contains information about configuring your NAS for use with Peer Global File System:

- [Amazon FSxN Configurations](#)
- [Dell Configurations](#)
- [NetApp ONTAP Configurations](#)
- [Nutanix Configurations](#)

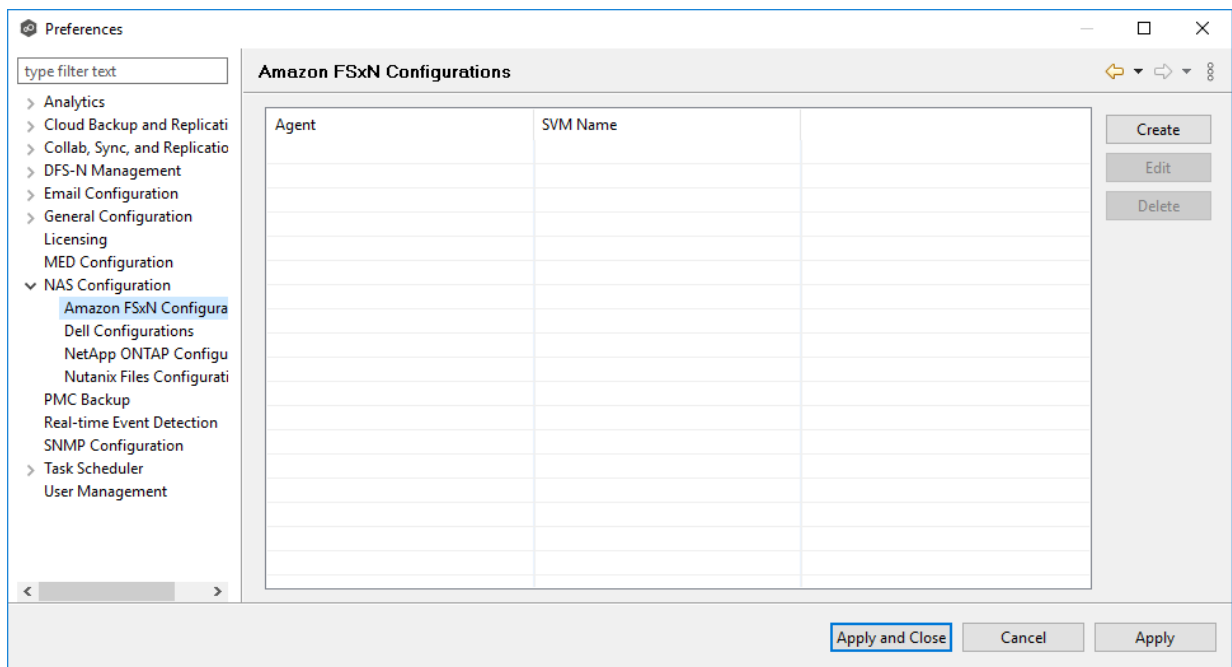
Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more Amazon FSxN file system within most available job types. In order to work with

Amazon FSxN, Peer Management Center utilizes the FPolicy API integrated into the NetApp operating system that powers FSxN. For detailed information about Amazon FSxN prerequisites and configuration, see [Amazon FSxN Prerequisites](#).

To create a new Amazon FSxN configuration:

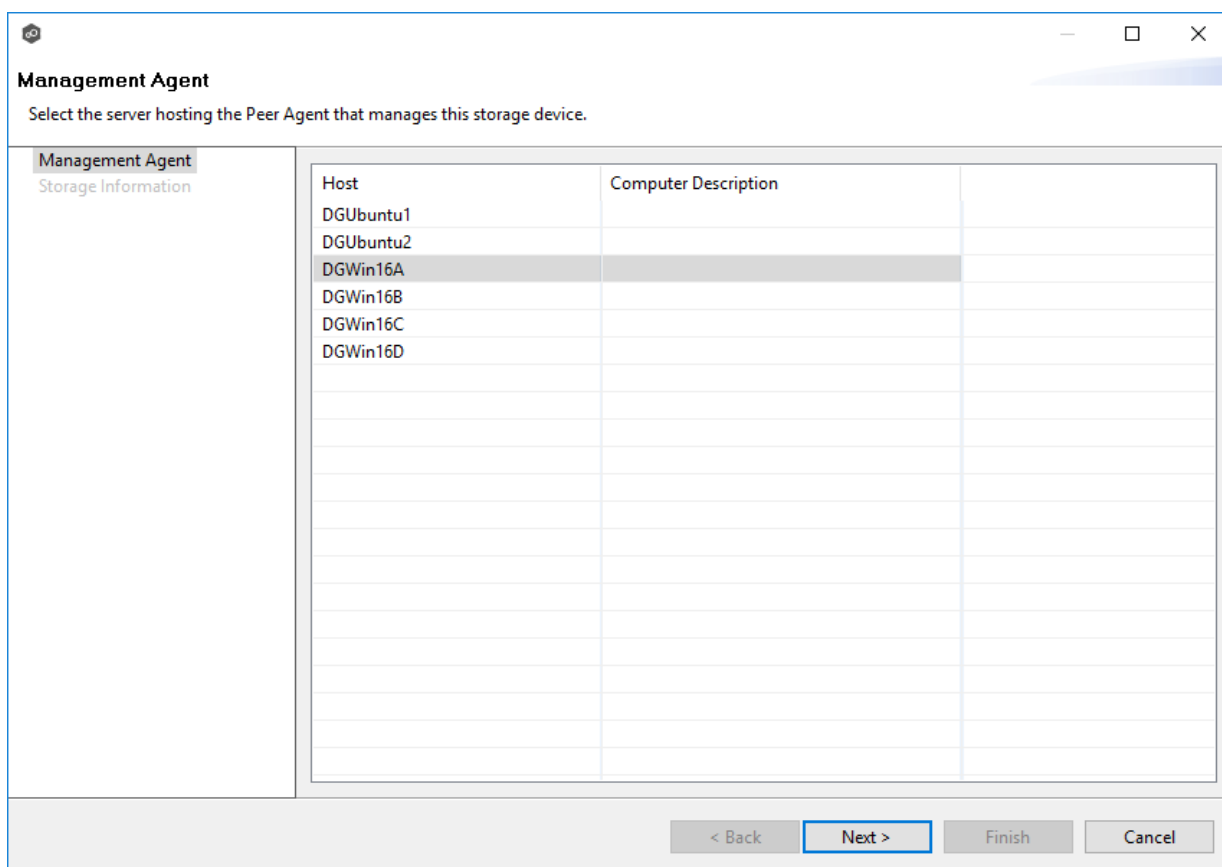
1. Select **Open Preferences** from the **Tools** menu.
2. Select **NAS Configuration** in the navigation tree.
3. Select **Amazon FSxN Configurations**.

The **Amazon FSxN Configurations** page is displayed. It lists any existing configurations.



4. Click the **Create** button.

The **Management Agent** page appears.



5. Select a Management Agent, and then click **Next**.

The **Storage Information** page appears.

The screenshot shows a window titled "Storage Information" with a subtitle "Enter the required information to connect to the storage". On the left, there is a sidebar with "Management Agent" and "Storage Information" (the latter is selected). The main area is divided into a "Credentials" section and a "Validation" section. The "Credentials" section contains five input fields: "*SVM Name:", "*SVM User Name:", "*SVM Password:", "SVM Management IP:", and "*Peer Agent IP:". The "*Peer Agent IP:" field is a dropdown menu. An "Advanced" button is located at the bottom right of the "Credentials" section. The "Validation" section contains a "Validate" button and a message: "You must enter the name." Below this is a note: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Amazon FSxN environments." At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

6. Enter the required values in **Credentials**.

Field	Description
SVM Name	Enter the name, FQDN, or IP address of the Storage Virtual Machine (SVM) hosting the data to be replicated.
SVM User Name	Enter the user name for the account managing the Storage Virtual Machine. The account must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. The account must not be a cluster management account.
SVM Management IP	Optional. Enter the IP address used to access the management API of the Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already permit management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually enter the address.

7. (Optional) Click [Advanced](#) and enter the required values.
8. Click **Validate**.
9. Click **Next**.
10. Click **OK**.

Amazon FSxN Advanced Options

The following configuration options are available for Amazon FSxN devices:

Amazon FSx for NetApp ONTAP Options

Advanced FPolicy Settings for host: DGWin16A and SVM: SVM1

*SVM Username:

*SVM Password:

SVM Management IP:

*Agent IP for SVM Conn.:

Filtered Extensions:

Admin Share Override:

Additional Properties:

NOTE: Any changes made to these Advanced FPolicy settings will be used with every other session in which this Agent is connecting to a NetApp storage device

Validate OK Cancel

Option	Description
SVM Username	Enter the account name of the VSAdmin or similar account on the Storage Virtual Machine (SVM) that has the appropriate access to ONTAPI.
SVM Password	Enter 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. Enter the IP address used to access the management API of the Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already permit management access, this field is not required.
Agent IP for SVM Conn.	Enter the IP address through which this Peer Agent will connect to the configured SVM. This address MUST be an IP address.
Filtered Extensions	Optional. Enter a comma-separated list of file extensions to exclude (without a leading asterisk (*)).
Admin Share Override	Optional. Enter the administrative-type share that you created on the cDOT SVM. To take advantage of performance improvements when using this option, the share must be created at the root of the SVM's namespace (/). Ideally, it should be named something similar to PMCShare\$ to prevent users from viewing it.
Additional Properties	Optional. Advanced settings that should only be used when directed by the Peer Software Support team.

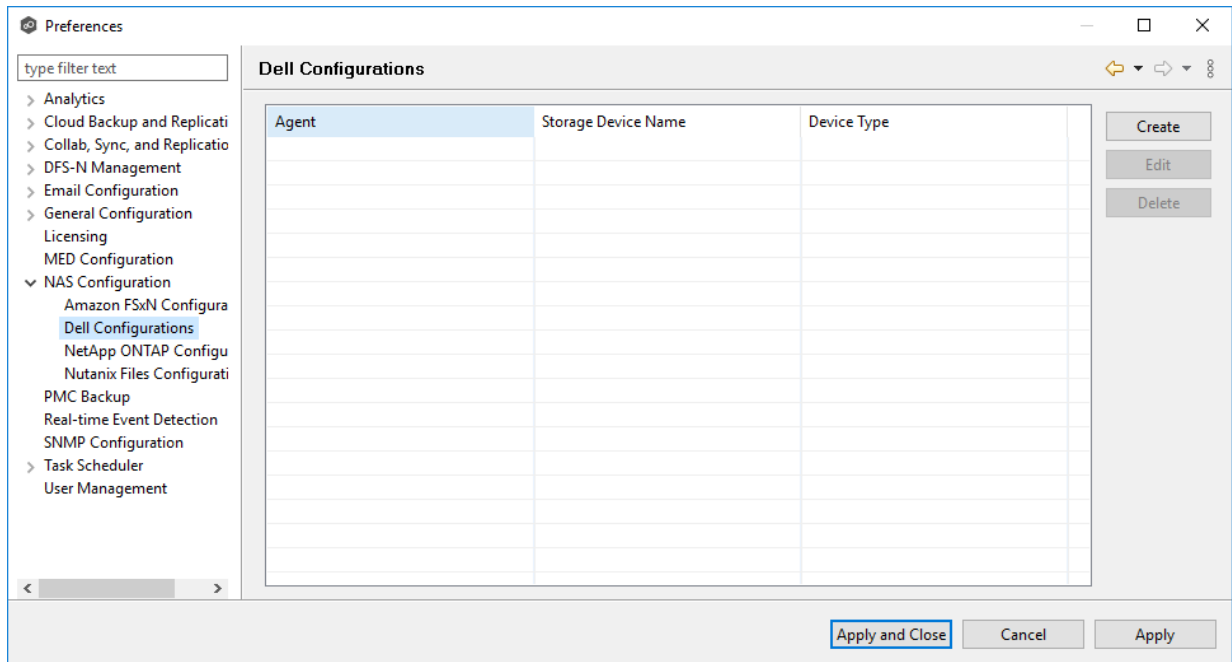
Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more Dell storage devices within most available job types. These Dell devices can be running PowerScale or Unity. For detailed information about Dell prerequisites, see [Dell Prerequisites](#).

To create a new Dell PowerScale configuration:

1. Select **Open Preferences** from the **Tools** menu.

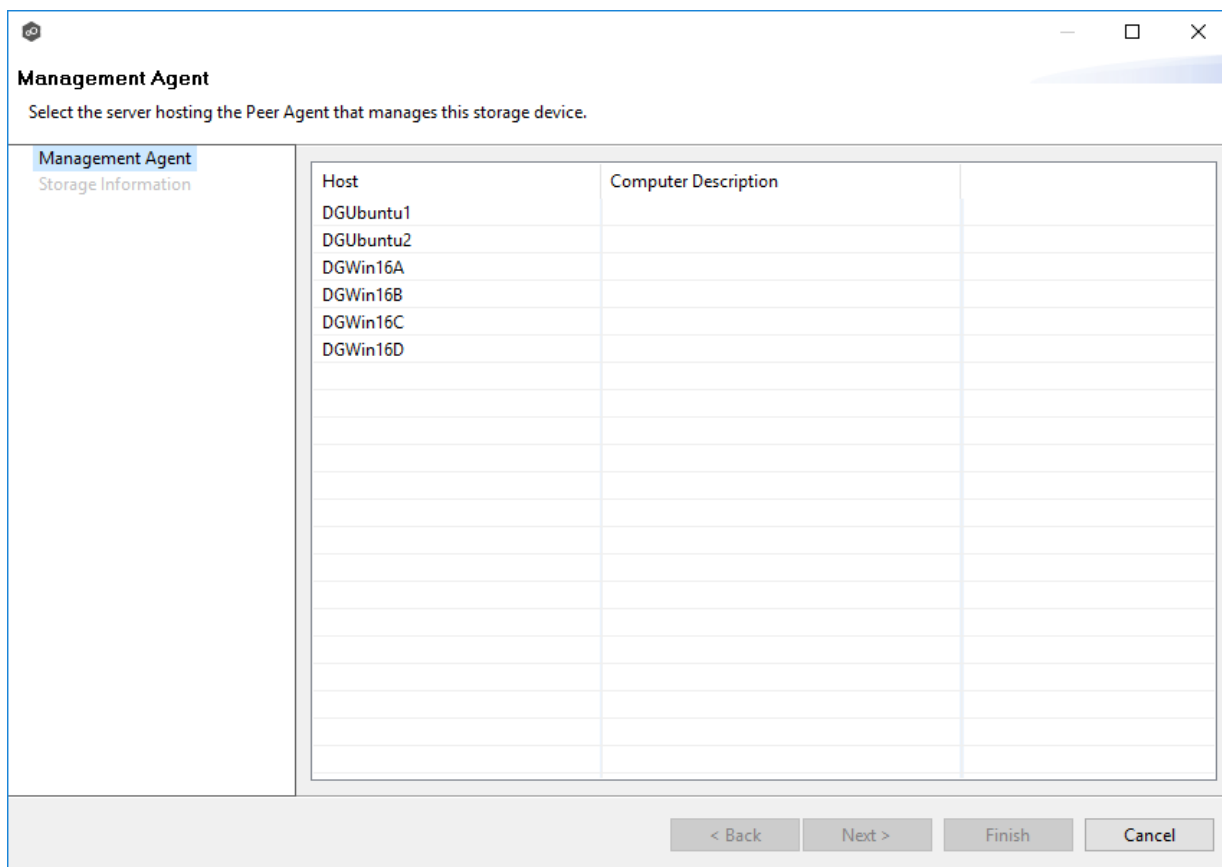
2. Select **NAS Configuration** in the navigation tree.
3. Select **Dell Configurations**.

The **Dell Configurations** page is displayed. It lists any existing configurations.



4. Click the **Create** button.

The **Management Agent** page appears.



5. Select a Management Agent, and then click **Next**.

The **Storage Information** page appears. The fields in the **Credentials** section vary, depending on the selected platform type; **PowerScale** is selected by default.

6. Select the device type from the **Device Type** drop-down.

Storage Information
Enter the required information to connect to the storage

Management Agent
Storage Information

Credentials

Device Type: PowerScale

*Cluster Name:

*Cluster Management IP:

*Cluster Username:

*Cluster Password:

Cluster Access Zone:

Connection Type: Syslog RabbitMQ

Advanced

Validate You must enter a Cluster Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments.

< Back Next > Finish Cancel

7. In the credentials page that appears, enter the required values:

[Dell PowerScale Configuration](#)

[Dell Unity Configuration](#)

8. (Optional) Click the **Advanced** button if you want to specify advanced options, and then enter the required values:

[Dell PowerScale Advanced Options](#)

[Dell Unity Advanced Options](#)

9. Click **Validate**.

10. Click **Next**.

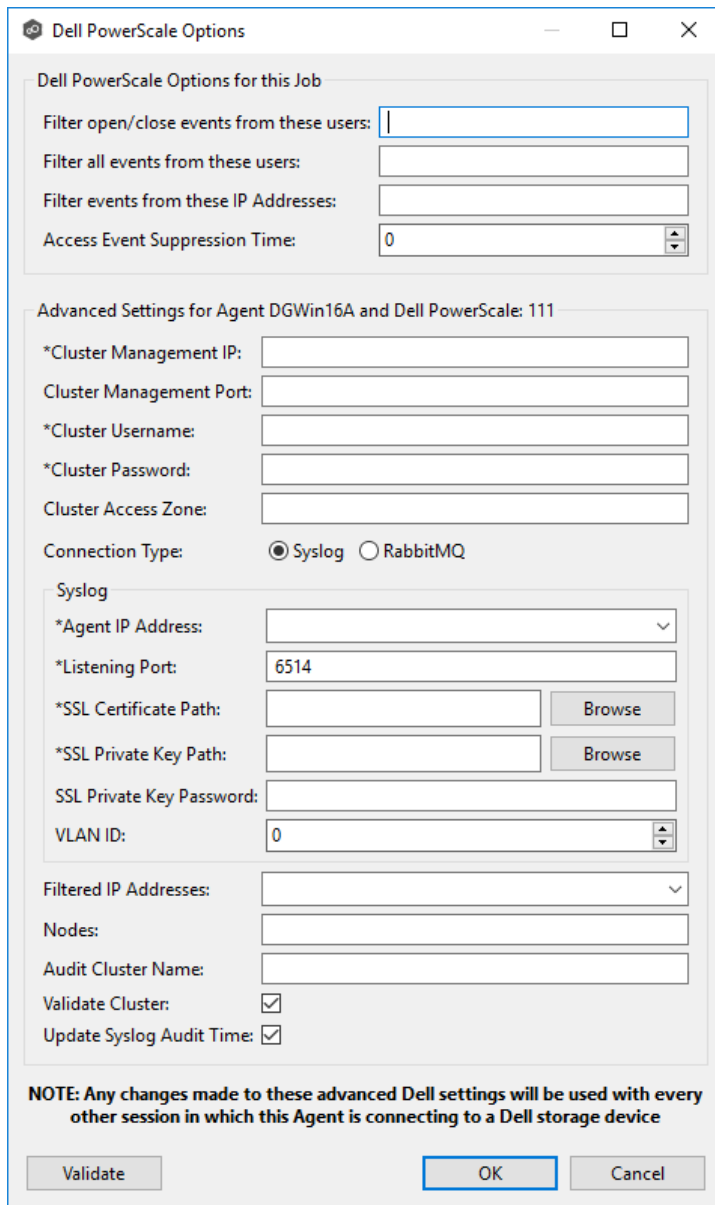
11. Click **OK**.

Dell PowerScale Configurations

1. Enter the required values.

The information required will vary, depending on whether you select **Syslog** or **RabbitMQ** as the connection type, due to the distinct protocols and mechanisms they employ for communication.

Syslog:



The screenshot shows the 'Dell PowerScale Options' dialog box. It is divided into two main sections: 'Dell PowerScale Options for this Job' and 'Advanced Settings for Agent DGWin16A and Dell PowerScale: 111'. The 'Advanced Settings' section is further divided into 'Syslog' and 'RabbitMQ' options. The 'Syslog' section is currently selected, and the 'Connection Type' is set to 'Syslog'.

Dell PowerScale 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:

Advanced Settings for Agent DGWin16A and Dell PowerScale: 111

- *Cluster Management IP:
- Cluster Management Port:
- *Cluster Username:
- *Cluster Password:
- Cluster Access Zone:
- Connection Type: Syslog RabbitMQ

Syslog

- *Agent IP Address:
- *Listening Port:
- *SSL Certificate Path:
- *SSL Private Key Path:
- SSL Private Key Password:
- VLAN ID:

Filtered IP Addresses:

- Nodes:
- Audit Cluster Name:
- Validate Cluster:
- Update Syslog Audit Time:

NOTE: Any changes made to these advanced Dell settings will be used with every other session in which this Agent is connecting to a Dell storage device

RabbitMQ:

Storage Information
— □ ×

Enter the required information to connect to the storage

Management Agent

Storage Information

Credentials

Device Type: PowerScale

*Cluster Name: DFDF

*Cluster Management IP:

*Cluster Username:

*Cluster Password:

Cluster Access Zone:

Connection Type: Syslog RabbitMQ

Advanced

Validate

 You must enter a valid Management IP Address.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments.

< Back
Next >
Finish
Cancel

Field	Description
Cluster Name	Enter the name of the PowerScale cluster hosting the data to be replicated.
Cluster Management IP	Enter the IP address to use to access the REST-based API integrated into the PowerScale cluster. Required only if multiple Access Zones are in use on the cluster.
Cluster Username	Enter the user name for the account managing the PowerScale cluster.
Cluster Password	Enter the password for account managing the PowerScale cluster.
Cluster Access Zone	Optional. Enter the name of the access zone that is being monitored.

Field	Description
Connection Type	<p>Select the appropriate method for sending real-time event notifications to the Agent:</p> <ul style="list-style-type: none"> • Opt for Syslog if the storage device directly transmits notifications to the Agent. • Opt for RabbitMQ if you're utilizing the CEE framework to dispatch notifications to the Agent.

If you selected **Syslog**, you will need to provide values for the following fields:

Field	Description
Agent IP Address	Select the IP address of the server hosting the Agent that manages the PowerScale cluster. The cluster must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.
Listening Port	Enter the port over which the Agent will receive TLS-based syslog events from the PowerScale cluster.
SSL Certificate Path	Enter the path to the certificate to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private Key Path	Enter the path to the private key to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private Key Password	[Optional] If your private key is protected with a password, enter it here. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .

2. (Optional) Click [Advanced](#) and enter the required values.
3. Click **Validate**.

4. Click **Finish**.

The options are divided into two groups:

- [Dell PowerScale Options for this job](#)
- [Dell PowerScale Advanced Settings](#)

The information required will vary, depending on whether you select **Syslog** or **RabbitMQ** as the connection type, due to the distinct protocols and mechanisms they employ for communication.

Syslog

Dell PowerScale Options

Dell PowerScale 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:

Advanced Settings for Agent DGWin16A and Dell PowerScale: 111

*Cluster Management IP:

Cluster Management Port:

*Cluster Username:

*Cluster Password:

Cluster Access Zone:

Connection Type: Syslog RabbitMQ

Syslog

*Agent IP Address:

*Listening Port:

*SSL Certificate Path:

*SSL Private Key Path:

SSL Private Key Password:

VLAN ID:

Filtered IP Addresses:

Nodes:

Audit Cluster Name:

Validate Cluster:

Update Syslog Audit Time:

NOTE: Any changes made to these advanced Dell settings will be used with every other session in which this Agent is connecting to a Dell storage device

RabbitMQ

Dell PowerScale Options

Dell PowerScale 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:

Advanced Settings for Agent DGWin16A and Dell PowerScale: 111

*Cluster Management IP:

Cluster Management Port:

*Cluster Username:

*Cluster Password:

Cluster Access Zone:

Connection Type: Syslog RabbitMQ

Filtered IP Addresses:

Nodes:

Audit Cluster Name:

Validate Cluster:

Update CEE Log Time:

NOTE: Any changes made to these advanced Dell settings will be used with every other session in which this Agent is connecting to a Dell storage device

Validate Cancel

Dell PowerScale Options for this Job

The following configuration options are available for Dell PowerScale devices:

Option	Description
Filter open/close events from these users	Enter a comma-separated list of user account names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Filter all events from these users	Enter a comma-separated list of user account names to exclude from all event detection. For example, if "USER"1 is excluded, any activity generated by USER1 will be ignored, e.g., file is open and modified.
Filter events from these IP Addresses	Enter a comma-separated list of IP addresses to exclude from all event detection. For example, if "192.168.0.100" is excluded, any activity generated by 192.168.0.100 will be ignored, e.g., file is open and modified.
Access Event Suppression Time	Enter a value that represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will be adjusted based on the selected NAS platform. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

Dell PowerScale Advanced Settings

The following advanced settings are available for Dell PowerScale devices:

Option	Description
Cluster Management IP	Enter the IP address to use to access the REST-based API integrated into the PowerScale cluster. Required only if multiple Access Zones are in use on the cluster.
Cluster Management Port	Optional. Enter the port number to use to access the REST-based API integrated into the PowerScale cluster. Default value is 8080.
Cluster Username	Enter the user name used to sign into the PowerScale cluster.
Cluster Password	Enter the password used to sign into the PowerScale cluster.
Cluster Access Zone	Optional. Enter the name of the access zone that is being monitored.
Connection Type	<p>Select the appropriate method for sending real-time event notifications to the Agent:</p> <ul style="list-style-type: none"> • Opt for Syslog if the storage device directly transmits notifications to the Agent. • Opt for RabbitMQ if you're utilizing the CEE framework to dispatch notifications to the Agent.
Filtered IP Addresses	Optional. Enter the IP addresses you wish to filter events from. We recommend that you include the IP address of the CEE Server.
Nodes	Optional. Enter a comma-delimited listed of additional node IP address to query for open files. These addresses must be accessible from the CEE Server where the Agent is running.
Audit Cluster Name	Optional. Enter the hostname that is set in the PowerScale audit system configuration.
Cluster Access Zone	Optional. Enter the name of the access zone that is being monitored.
Validate Cluster	Select this option if you want the cluster to undergo validation both during registration and periodically by a maintenance thread.
Update CEE Log Time	Select this option if you want the audit log time on the PowerScale cluster to be set to the start time of the first job to communicate with this cluster

If you selected **Syslog**, you will need to provide values for the following fields:

Field	Description
Agent IP Address	Select the IP address of the server hosting the Agent that manages the PowerScale cluster. The cluster must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.
Listening Port	Enter the port over which the Agent will receive TLS-based syslog events from the PowerScale cluster.
SSL Certificate Path	Enter the path to the certificate to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private Key Path	Enter the path to the private key to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private Key Password	[Optional] If your private key is protected with a password, enter it here. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
VLAN ID	Enter the tag of the VLAN to use when the Agent connects to the PowerScale cluster. This should either be set to 0 when using untagged interfaces, or to a value between 1-4094.

Dell Unity Configurations

1. Enter the required values.

Storage Information
— □ ×

Enter the required information to connect to the storage

Management Agent

Storage Information

Credentials

Device Type:

*CIFS Server Name:

*Unisphere Management IP:

*Unisphere Username:

*Unisphere Password:

You must enter a CIFS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments.

Field Description	Description
CIFS Server Name	Enter the name of the NAS server hosting the data to be replicated.
Unisphere Management IP	Enter the IP address of the Unisphere system used to manage the Unity storage device. The IP address should not point to the NAS server.
Unisphere Username	Enter the user name for the Unisphere account managing the Unity storage device.

Field Description	Description
Unisphere Password	Enter the password for the Unisphere account managing the Unity storage device.

2. (Optional) Click [Advanced](#) and enter the required values.
3. Click **Validate**.
4. Click **Finish**.

The options are divided into two groups:

- [Dell Unity Options for this Job](#)
- [Dell Unity Advanced Settings](#)

Dell Unity Options

Dell Unity 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:

Advanced Settings for Agent DGWin16A and Dell Unity: CVCV

*Unisphere Management IP:

Unisphere Management Port:

*Unisphere Username:

*Unisphere Password:

Filtered IP Addresses:

Validate Unisphere:

NOTE: Any changes made to these advanced Dell settings will be used with every other session in which this Agent is connecting to a Dell storage device

Validate OK Cancel

Dell Unity Options for this Job

The following configuration options are available for Dell Unity devices:

Option	Description
Filter open/close events from these users	Enter a comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Filter all events from these users	Enter a comma-separated list of user names to exclude from all event detection. For example, if "USER1" is excluded, any activity generated by USER1 will be ignored, e.g., file is open and modified.
Filter events from these IP Addresses	Enter a comma-separated list of IP addresses to exclude from all event detection. For example, if "192.168.0.100" is excluded, any activity generated by 192.168.0.100 will be ignored, e.g., file is open and modified.
Access Event Suppression Time	Enter a value that represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will be adjusted based on the selected NAS platform. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

Dell Unity Advanced Settings

The following advanced settings are available for Dell EMC Unity devices:

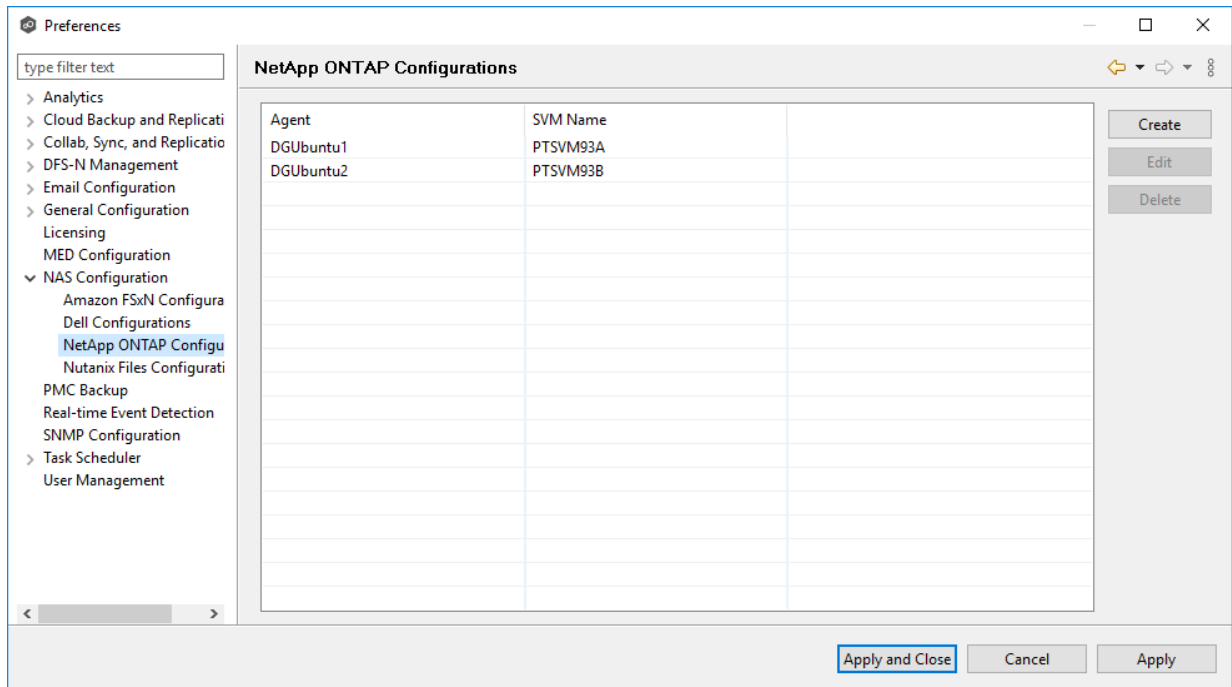
Option	Description
Unisphere Management IP	Enter the IP address of the Unisphere system used to manage the Unity storage device. The IP address should not point to the NAS server.
Unisphere Management Port	Optional. Enter the Unisphere Management port number of the Unity system. Default value is 443.
Unisphere Username	Enter the user name for the Unisphere account managing the Unity storage device.
Unisphere Password	Enter the IP address of the Unisphere system managing the Unity storage device. The IP address should not point to the NAS server.
Filtered IP Addresses	Optional. Enter the IP addresses you wish to filter events from. We recommend that you include the IP address of the CEE Server.
Validate Unisphere	Select this option if you want the cluster to undergo validation both during registration and periodically by a maintenance thread.

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more NetApp storage devices within most available job types. These NetApp devices can be running clustered Data ONTAP. In order to work with NetApp devices, Peer Management Center leverages the FPolicy API built into the NetApp device. For detailed information about NetApp prerequisites and configuration, see [NetApp Prerequisites](#).

To create a new NetApp ONTAP configuration:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **NAS Configuration** in the navigation tree.
3. Select **NetApp ONTAP Configurations**.

The **NetApp ONTAP Configurations** page is displayed. It lists any existing configurations.



4. Click the **Create** button.

The **Management Agent** page appears.

The screenshot shows a window titled "Storage Information" with a subtitle "Enter the required information to connect to the storage". On the left, there is a sidebar with two options: "Management Agent" and "Storage Information", with "Storage Information" selected. The main area is divided into two sections. The top section is titled "Credentials" and contains five input fields: "*SVM Name:", "*SVM User Name:", "*SVM Password:", "SVM Management IP:", and "*Peer Agent IP:". The "*Peer Agent IP:" field is a dropdown menu. An "Advanced" button is located at the bottom right of this section. The bottom section contains a "Validate" button and a message: "You must enter the name." Below this is a link: "Having trouble connecting? Please verify that all [prerequisites](#) are met for NetApp ONTAP environments." At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

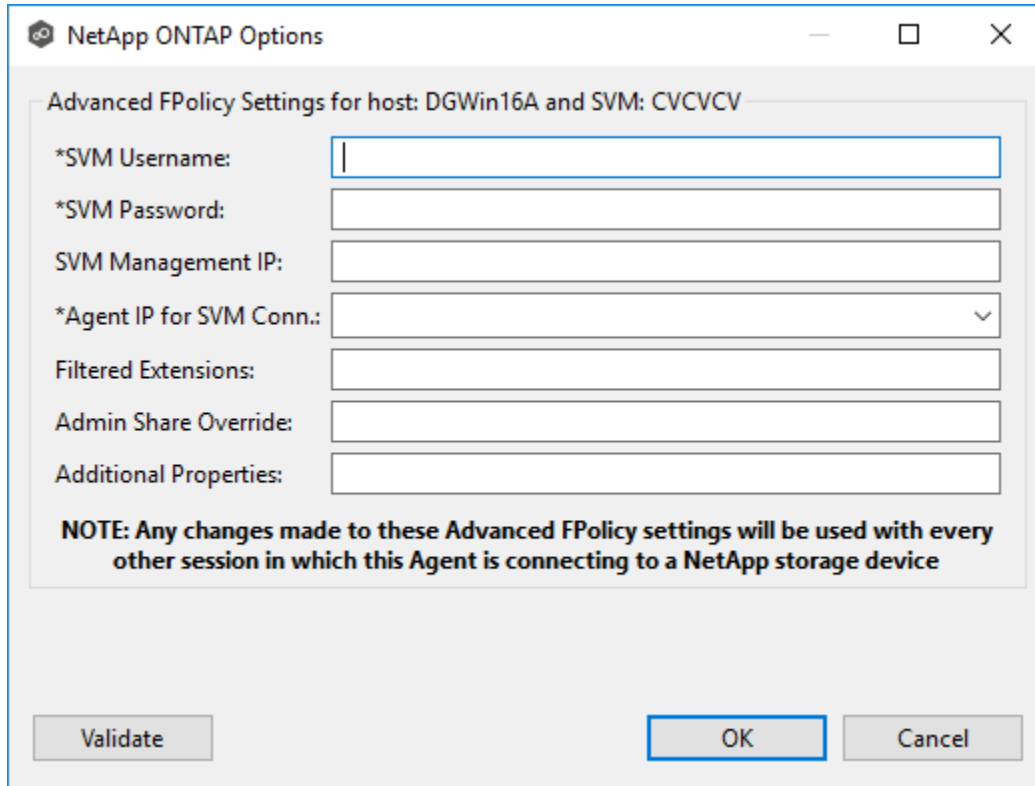
6. Enter the required values in **Credentials**.

Field	Description
SVM Name	Enter the name, FQDN, or IP address of the Storage Virtual Machine (SVM) hosting the data to be replicated.
SVM Username	Enter the user name for the account managing the Storage Virtual Machine. The account must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. The account must not be a cluster management account.
SVM Management IP	Optional. Enter the IP address used to access the management API of the Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already permit management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.

7. (Optional) Click [Advanced](#) and enter the required values.
8. Click **Validate**.
9. Click **Next**.
10. Click **OK**.

NetApp ONTAP Advanced Options

The following configuration options are available for NetApp cDOT devices:



The image shows a dialog box titled "NetApp ONTAP Options" with a standard Windows window border. The main content area is titled "Advanced FPolicy Settings for host: DGWin16A and SVM: CVCVCV". It contains several input fields: "*SVM Username:" (text box), "*SVM Password:" (text box), "SVM Management IP:" (text box), "*Agent IP for SVM Conn.:" (dropdown menu), "Filtered Extensions:" (text box), "Admin Share Override:" (text box), and "Additional Properties:" (text box). Below these fields is a bolded note: "NOTE: Any changes made to these Advanced FPolicy settings will be used with every other session in which this Agent is connecting to a NetApp storage device". At the bottom of the dialog are three buttons: "Validate", "OK", and "Cancel".

NetApp ONTAP Options

Advanced FPolicy Settings for host: DGWin16A and SVM: CVCVCV

*SVM Username:

*SVM Password:

SVM Management IP:

*Agent IP for SVM Conn.:

Filtered Extensions:

Admin Share Override:

Additional Properties:

NOTE: Any changes made to these Advanced FPolicy settings will be used with every other session in which this Agent is connecting to a NetApp storage device

Validate OK Cancel

Option	Description
SVM Username	The account name of the VSAdmin or similar account on the SVM that has the appropriate access to ONTAPI.
SVM Password	The password of the VSAdmin or similar account on the SVM that has the appropriate access to ONTAPI. This value will be encrypted.
SVM Management IP	Optional. Enter the IP address used to access the management API of the Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already permit management access, this field is not required.
Agent IP for SVM Conn.	The IP address over which this Peer Agent will connect to the configured SVM. This MUST be an IP address.
Filtered Extensions	Optional. A comma-separated list of file extensions to exclude (without a leading asterisk (*)).
Admin Share Override	Optional. Enter the administrative-type share that you created on the cDOT SVM. To take advantage of performance improvements when using this option, the share must be created at the root of the SVM's namespace (/). Ideally it should be named to something similar to PMCShare\$ to prevent users from being able to see it.
Additional Properties	Optional. Advanced settings that should only be used when directed by the Peer Software Support team.

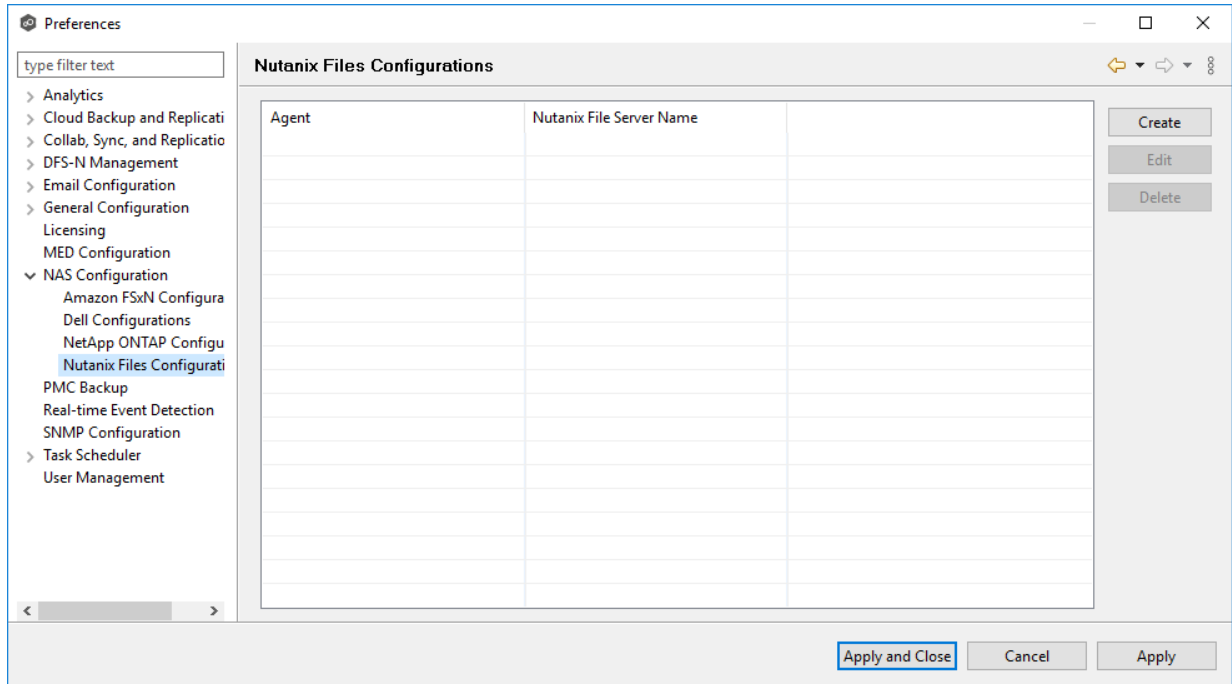
Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more Nutanix Files (formerly Acropolis File Services or AFS) clusters within most available job types. For detailed information about Nutanix prerequisites, see [Nutanix Prerequisites](#).

To create a new Nutanix Files configuration:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **NAS Configuration** in the navigation tree.

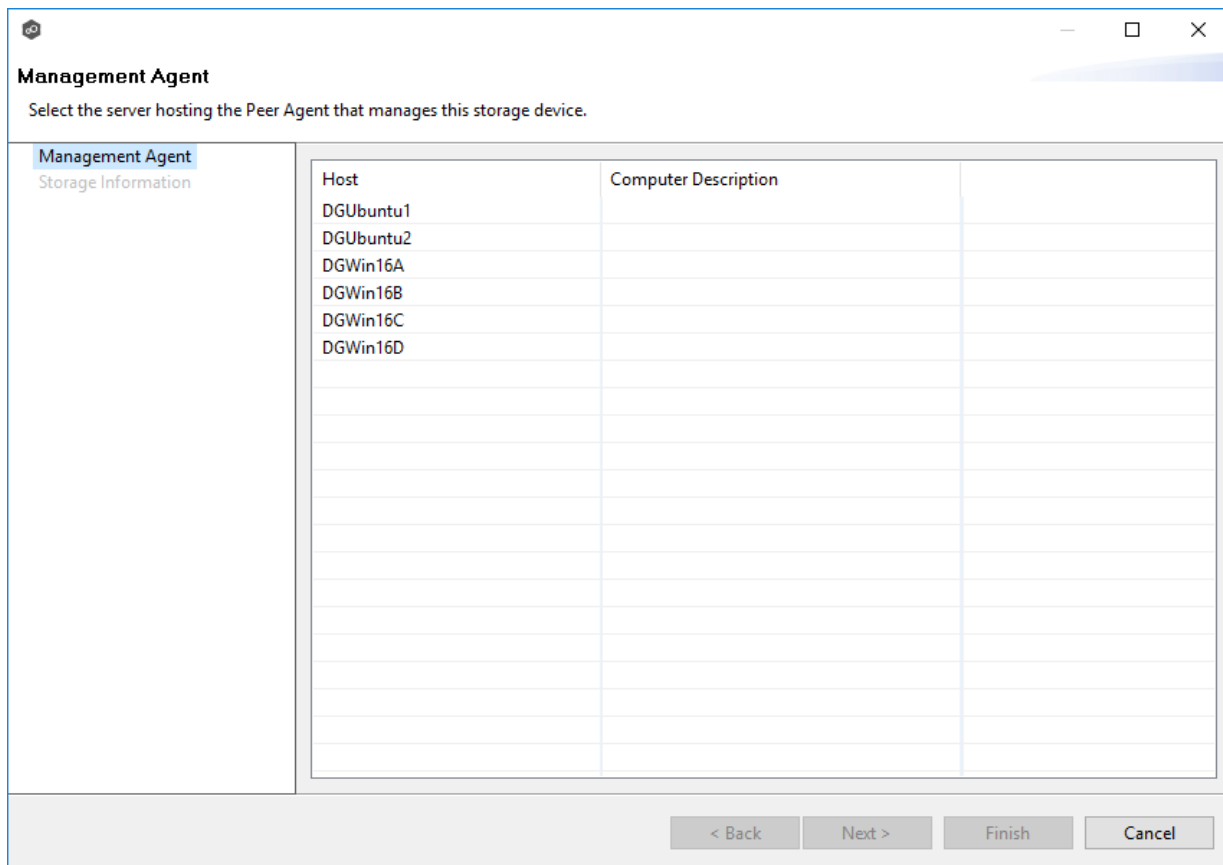
3. Select **Nutanix Configurations**.

The **Nutanix Files Configurations** page is displayed. It lists any existing configurations.



4. Click the **Create** button.

The **Management Agent** page appears.



5. Select a Management Agent, and then click **Next**.

The **Storage Information** page appears.

Storage Information
Enter the required information to connect to the storage

Management Agent
Storage Information

Credentials

*Nutanix File Server Name:

*Username:

*Password:

*Peer Agent IP:

Advanced

Validate You must enter the name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Nutanix Files environments.

< Back Next > Finish Cancel

6. Enter the required values in **Credentials**.

Field	Description
Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Enter the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually enter the address. The IP address should not point to the Files cluster itself.

7. (Optional) Click [Advanced](#) button and then enter the required values.
8. Click **Validate**.
9. Click **Next**.
10. Click **OK**.

Nutanix Files Advanced Options

The options are divided into two groups:

- [Nutanix Files Options for this Job](#)
- [Advanced Settings](#)

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:

Advanced Settings for Agent DGWin16A and Nutanix File Server Name: FS1

Peer Agent IP:

Username:

Password:

NOTE: Any changes made to these advanced Nutanix Files settings will be used with every other session in which this Agent is connecting to a Nutanix storage device

Nutanix Files Options for this Job

The following configuration options are available for Nutanix Files devices:

Option	Description
Filter open/close events from these users	Enter a comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Filter all events from these users	Enter a comma-separated list of user names to exclude from all event detection. For example, if "USER1" is excluded, any activity generated by USER1 will be ignored, e.g., file is open and modified.
Filter events from these IP Addresses	Enter a comma-separated list of IP addresses to exclude from all event detection. For example, if "192.168.0.100" is excluded, any activity generated by 192.168.0.100 will be ignored, e.g., file is open and modified.
Access Event Suppression Time	Enter a value that represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will be adjusted based on the selected NAS platform. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

Advanced Settings

The following advanced settings are available for Nutanix Files devices:

Option	Description
Peer Agent IP	Enter the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually enter the address. The IP address should not point to the Files cluster itself.
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.

Real-time Event Detection

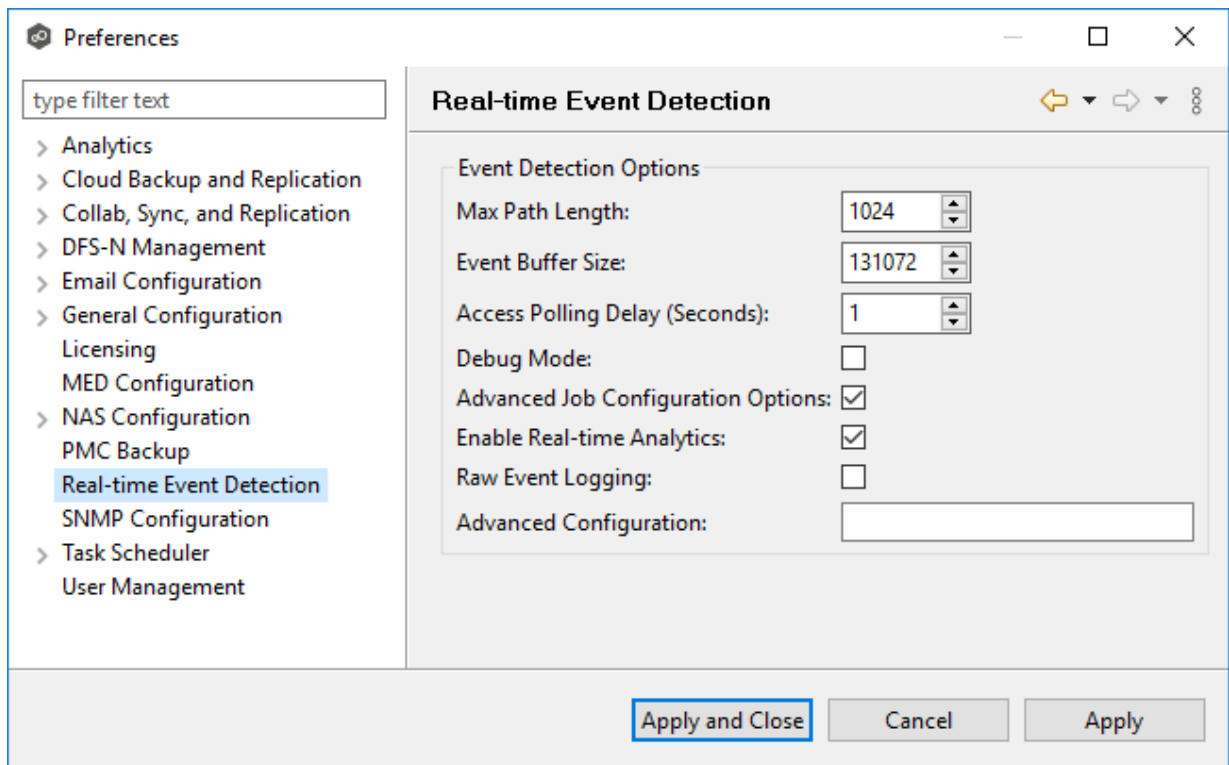
Several options are available to tune the way real-time event detection occurs. These options apply to all job types, except for DFS-N Management.

Note: There are also real-time event detection settings applicable to most job types in Peer Management Center. See [Real-time Event Detection](#) in the [File Collab, Sync, and Repl, and Locking Preferences](#) topic for more information.

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

1. Select **Open Preferences** from the **Tools** menu.
2. Select **Real-time Detection** in the navigation tree.

The following page is displayed.



3. Modify values as needed:

Option	Description
Max Path Length	Specify 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	Specify the buffer size to used by the Peer Agents to communicate with various Windows and enterprise NAS platform APIs.
Access Polling Delay (Seconds)	Specify the time interval that Peer Agent will poll a file server for its open files list.
Debug Mode	Select to enable debug logging for real-time detection. This logs additional information that is often useful in troubleshooting issues but can increase overhead.

Option	Description
Advanced Job Configuration Options	Select to enable advanced job-level options tied to real-time event detection.
Enable Real-time Analytics	Select to enable the collection of real-time data for trend analysis. Enabled by default; disable only when instructed by Peer Software support.
Raw Event Logging	Select to enable 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	Specify 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.

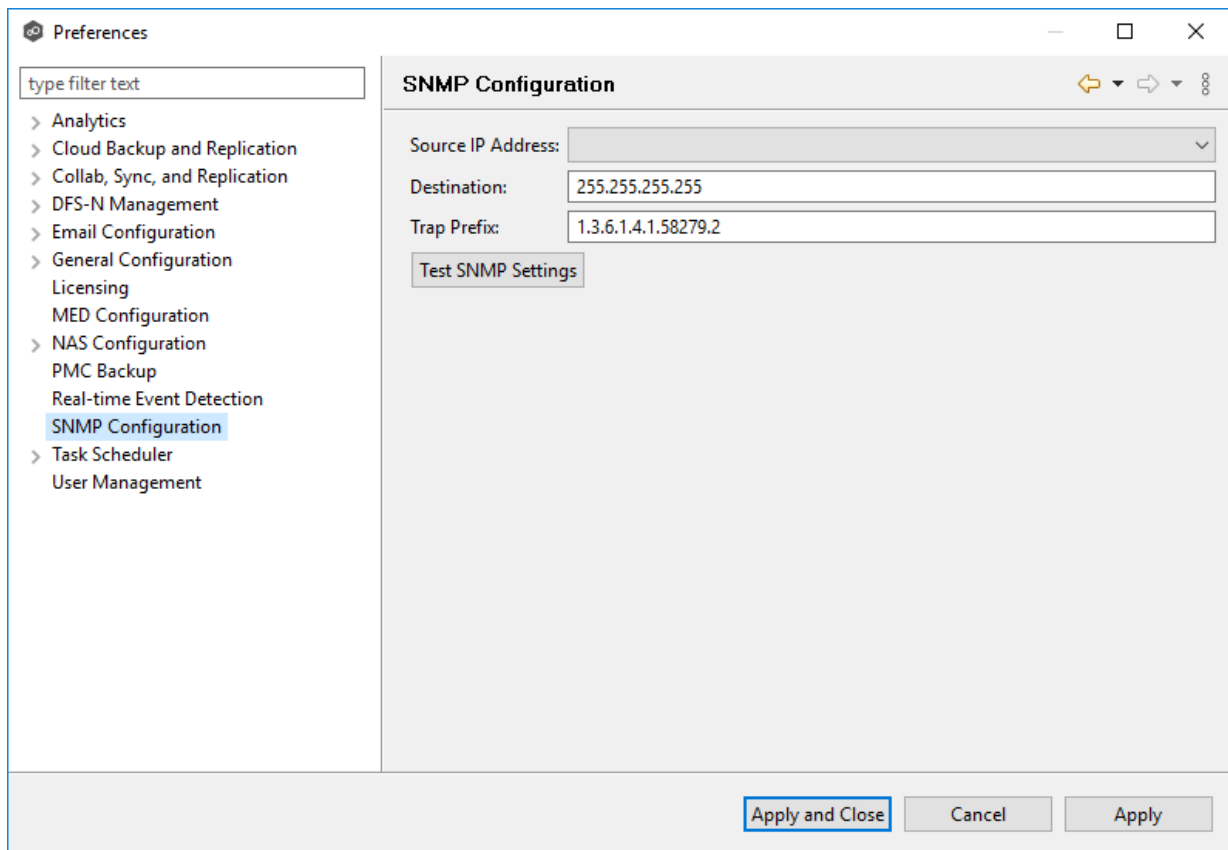
4. Click **Apply and Close** or **Apply**.

SNMP Configuration

Before Peer Management Center can send SNMP notifications on behalf of any job, a few key SNMP settings must be configured.

To configure SNMP settings:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **SNMP Configuration** in the navigation tree.



3. In the **Source IP Address** field, select or manually enter the IP address over which the trap will be sent.
4. In the **Destination** field, enter the destination host name, IP address, or broadcast address.
5. For **Trap Prefix**, enter a prefix that will help to identify whether the message is coming from different instances of Peer Management Center or from different jobs. In the default prefix, "1.3.5.1.4.1" represents IANA-registered private enterprises, "58279" is reserved for Peer Software, and the trailing ".2" represents Peer Management Center.
6. Click **Test SNMP Settings**, and then click **OK** in the **Test Connection** dialog.
You can verify the result by checking in your SNMP management tool.
7. Click **Apply and Close** or **Apply**.

User Management

The **User Management** page allows you to manage users of the web client interface. From this page, you can [manage web client users](#), [manage roles](#), and [configure Active Directory authentication](#).

Note: Any rich client user can access the User Management page, but only web client users with an Administrator role can access it.

Accessing User Management

To access User Management:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **User Management** from the navigation tree.

The **User Management** page is displayed:

Preferences

type filter text

- > Analytics
- > Cloud Backup and Re
- > Collab, Sync, and Rep
- > DFS-N Management
- > Email Configuration
- > General Configuratio
- Licensing
- MED Configuration
- > NAS Configuration
- Real-time Event Dete
- SNMP Configuration
- > Task Scheduler
- User Management**

User Management

Roles

Power User	Create
Administrator	Edit
Help Desk	Delete

Users

Internal Users

admin	Create
	Edit
	Delete

Active Directory Users and Groups

Active Directory Users

	Add
	Edit
	Delete

Active Directory Groups

	Add
	Edit
	Delete

Active Directory Authentication

Authentication will not work until the URL and credentials are entered.

LDAP Server URLs:

LDAP Search Domain (Optional):

LDAP Credentials:

Managing Web Client Users

Web client users are users that access Peer Management Center through the web client. All web client users have a role.

You add, modify, and delete web users through the [User Management](#) page in **Preferences**.

Who Can Manage Web Users?

Management of web users can be performed through the rich client or through the web client by a user with an **Administrator** role.

For more information, see:

- [Accessing User Management](#)
- [Managing Internal Users](#)
- [Managing Active Directory Users and Groups](#)
- [Configuring Active Directory Authentication](#)

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

Managing internal users involves:

- [Creating an internal user](#)
- [Editing an internal user](#)
- [Deleting an internal user](#)

Default Internal User

Peer Management Center includes a **default internal user** with access to all web client functionalities: the **admin** user. This user does not need to be created. The properties of this internal user are as follows:

Username: admin

Password*: password (should be changed immediately upon first log-in for security reasons)

Role: Administrator

The admin user cannot be deleted or have its role changed. For security reasons, it is essential to [change the admin user's password immediately upon first log-in](#).

*While default password for most installs is **password**, that is not the case with marketplace virtual appliances. If you have deployed the PMC virtual appliance via the Microsoft Azure Marketplace, your password will be your Azure Subscription ID. If you have deployed the PMC virtual appliance via the Amazon AWS Marketplace, your default password will be the EC2 instance ID of the virtual appliance.

Creating an Internal User

To add an internal user, follow these steps:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **User Management** from the navigation tree.

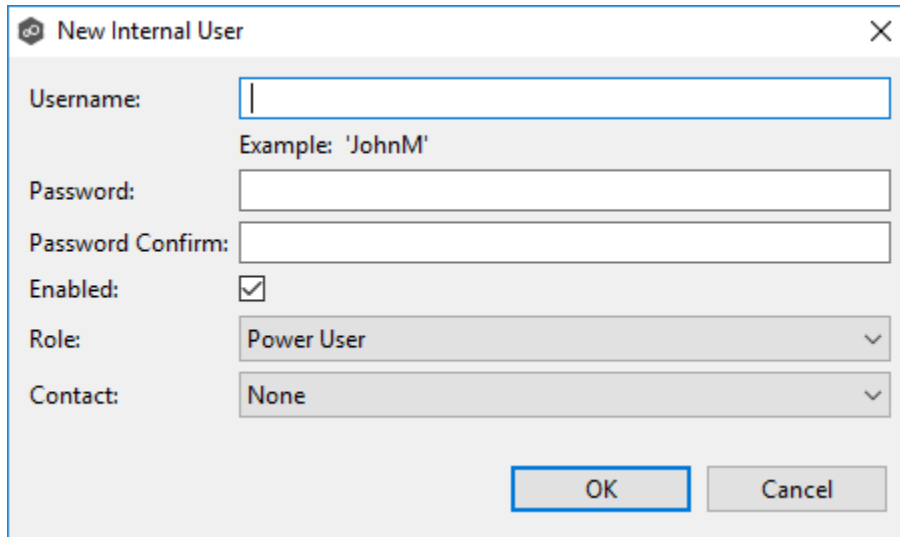
The screenshot shows the 'Preferences' window with the 'User Management' section selected in the left-hand navigation pane. The 'User Management' section is divided into several sub-sections:

- Roles:** A list box containing 'Power User', 'Administrator', and 'Help Desk'. To the right are buttons for 'Create', 'Edit', and 'Delete'.
- Users:** A sub-section titled 'Internal Users' with a list box containing 'admin'. To the right are buttons for 'Create', 'Edit', and 'Delete'.
- Active Directory Users and Groups:** A sub-section containing two sub-sections:
 - Active Directory Users:** An empty list box with 'Add', 'Edit', and 'Delete' buttons to its right.
 - Active Directory Groups:** An empty list box with 'Add', 'Edit', and 'Delete' buttons to its right.
- Active Directory Authentication:** A section with a warning: **Authentication will not work until the URL and credentials are entered.** It includes input fields for 'LDAP Server URLs', 'LDAP Search Domain (Optional):' (with 'Optional' entered), and 'LDAP Credentials'. Below these are buttons for 'Add/Update LDAP Credentials' and 'Test'.

At the bottom of the window, there are three buttons: 'Apply and Close' (highlighted with a blue border), 'Cancel', and 'Apply'. A scrollbar is visible on the left side of the main content area.

3. Select the **Create** button for Internal Users.

The **New Internal User** dialog appears.



The screenshot shows a dialog box titled "New Internal User" with a close button (X) in the top right corner. The dialog contains the following fields and controls:

- Username:** A text input field with a cursor. Below it, the text "Example: 'JohnM'" is displayed.
- Password:** A text input field.
- Password Confirm:** A text input field.
- Enabled:** A checkbox that is checked.
- Role:** A dropdown menu with "Power User" selected.
- Contact:** A dropdown menu with "None" selected.
- Buttons:** "OK" and "Cancel" buttons are located at the bottom right of the dialog.

4. Enter the following information.

- **Username:** The username can contain letters, numbers, and spaces; it cannot contain special characters. The minimum number of characters is 6; the maximum number of characters is 20.
- **Password:** The minimum number of characters is 6; the maximum number of characters is 20. The password cannot be the same as the username.
- **Password Confirm:** Re-enter the password you entered.
- **Enabled:** Select this checkbox if you want to enable this user to access Peer Management Center. You can enable or disable the user at a later date by [editing the user](#).
- **Role:** Select the [role](#) you want to assign to the user. It can be a base role or a custom role. For more details on the available roles, see [Roles for Web Client Users](#).
- **Contact:** Select the user's email address from the drop-down list. If the user's email address does not appear in the list, you can add it to **Contacts** in the [Email Configuration](#) in [Preferences](#).

5. Click **OK**.

The new user appears in the list of internal users on the User Management page.

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

Editing an Internal User

Once an internal user has been created, its user name, password, email address, and role can all be changed.

Note: The [default internal user](#) cannot be deleted, nor can its role be changed. However, you should change the password for the default internal user.

To edit an internal user from Peer Management Center:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.
3. Select the user from the list of internal users.
4. Click **Edit**.
5. Make the changes in the **Edit User Information** dialog.
6. Click **OK**.
7. Click **Apply and Close** or **Apply**.

Deleting an Internal User

Once the account of an internal user is deleted, that user can no longer access Peer Management Center through the web client.

Note: The [default internal user](#) cannot be deleted.

To delete an Active Directory user or group from Peer Management Center:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.
3. Select the user from the list of internal users.

4. Click **Delete**.
5. Click **OK** in the **Remove User** dialog.
6. Click **Apply and Close** or **Apply**.

An Active Directory (AD) user or group is one whose access to Peer Management Center is authenticated through Active Directory. Adding an Active AD user or group authenticates and authorizes that user or group members to use Peer Management Center. The AD user or group must already exist in Active Directory prior to adding the user or group to Peer Management Center. Active Directory users won't be able to access the web client until [Active Directory authentication is configured](#) in Peer Management Center.

Managing Active Directory users and groups involves:

- [Configuring Active Directory Authentication](#)
- [Adding an Active Directory User or Group](#)
- [Editing an Active Directory User or Group](#)
- [Deleting an Active Directory User or Group](#)

Configuring Active Directory Authentication

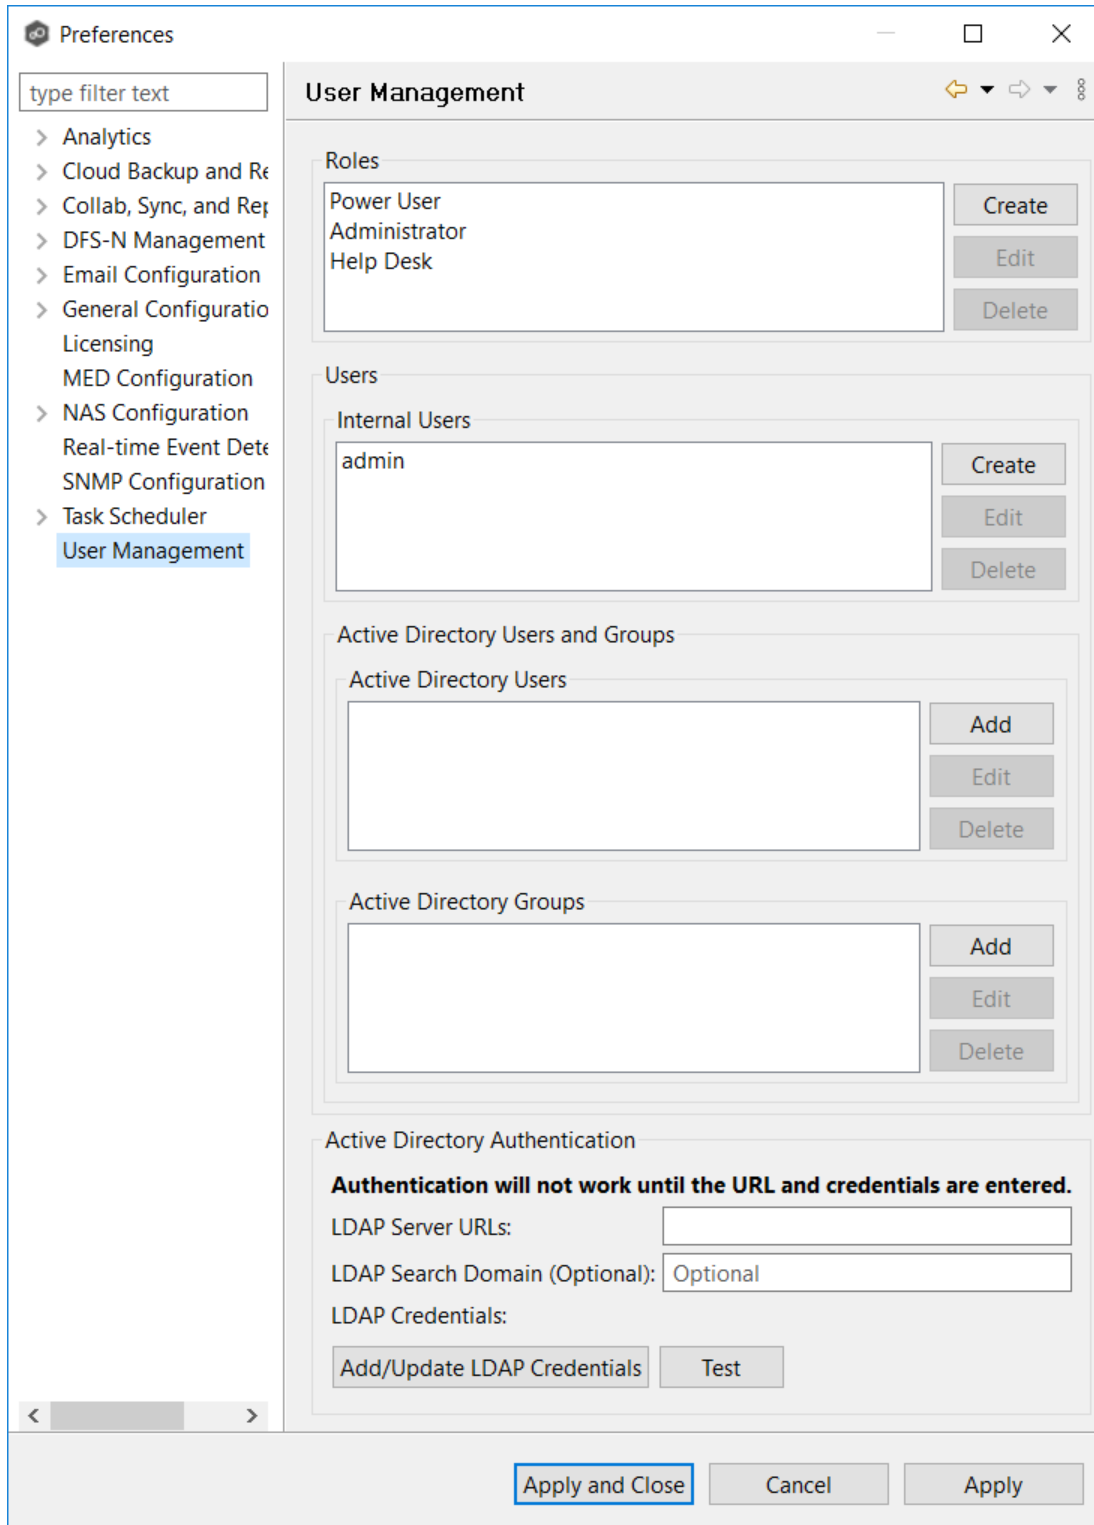
To configure Active Directory authentication, you need:

- The URL of the LDAP server
- The LDAP administrator credentials

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

To configure Active Directory authentication:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **User Management** from the navigation tree.



3. In the **LDAP Server URLs** field in the **Active Directory Authentication** section, enter the URLs of the LDAP servers on the network using one of the following formats:

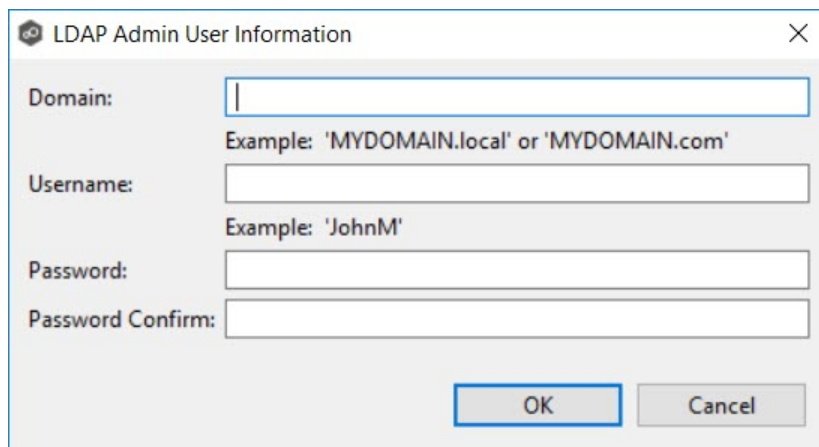
- ldap://MYDOMAIN.LOCAL
- ldaps://MYDOMAIN.LOCAL

You can enter multiple LDAP URLs separated by spaces for failover redundancy.

4. (Optional) In the **LDAP Server Domain** field, enter a domain name override to use for LDAP searches.

Enter a value when (a) the search domain is different than the domain or DNS name specified in the **LDAP Server URL** field or (b) multiple LDAP servers URLs are specified.

5. Click **Add/Update LDAP Credentials**.



The screenshot shows a dialog box titled "LDAP Admin User Information". It contains the following fields and examples:

- Domain:** [Text Input Field] Example: 'MYDOMAIN.local' or 'MYDOMAIN.com'
- Username:** [Text Input Field] Example: 'JohnM'
- Password:** [Text Input Field]
- Password Confirm:** [Text Input Field]

At the bottom of the dialog are two buttons: "OK" and "Cancel".

6. Enter the domain name, user name, and password.
7. Confirm the password.
8. Click **OK**.

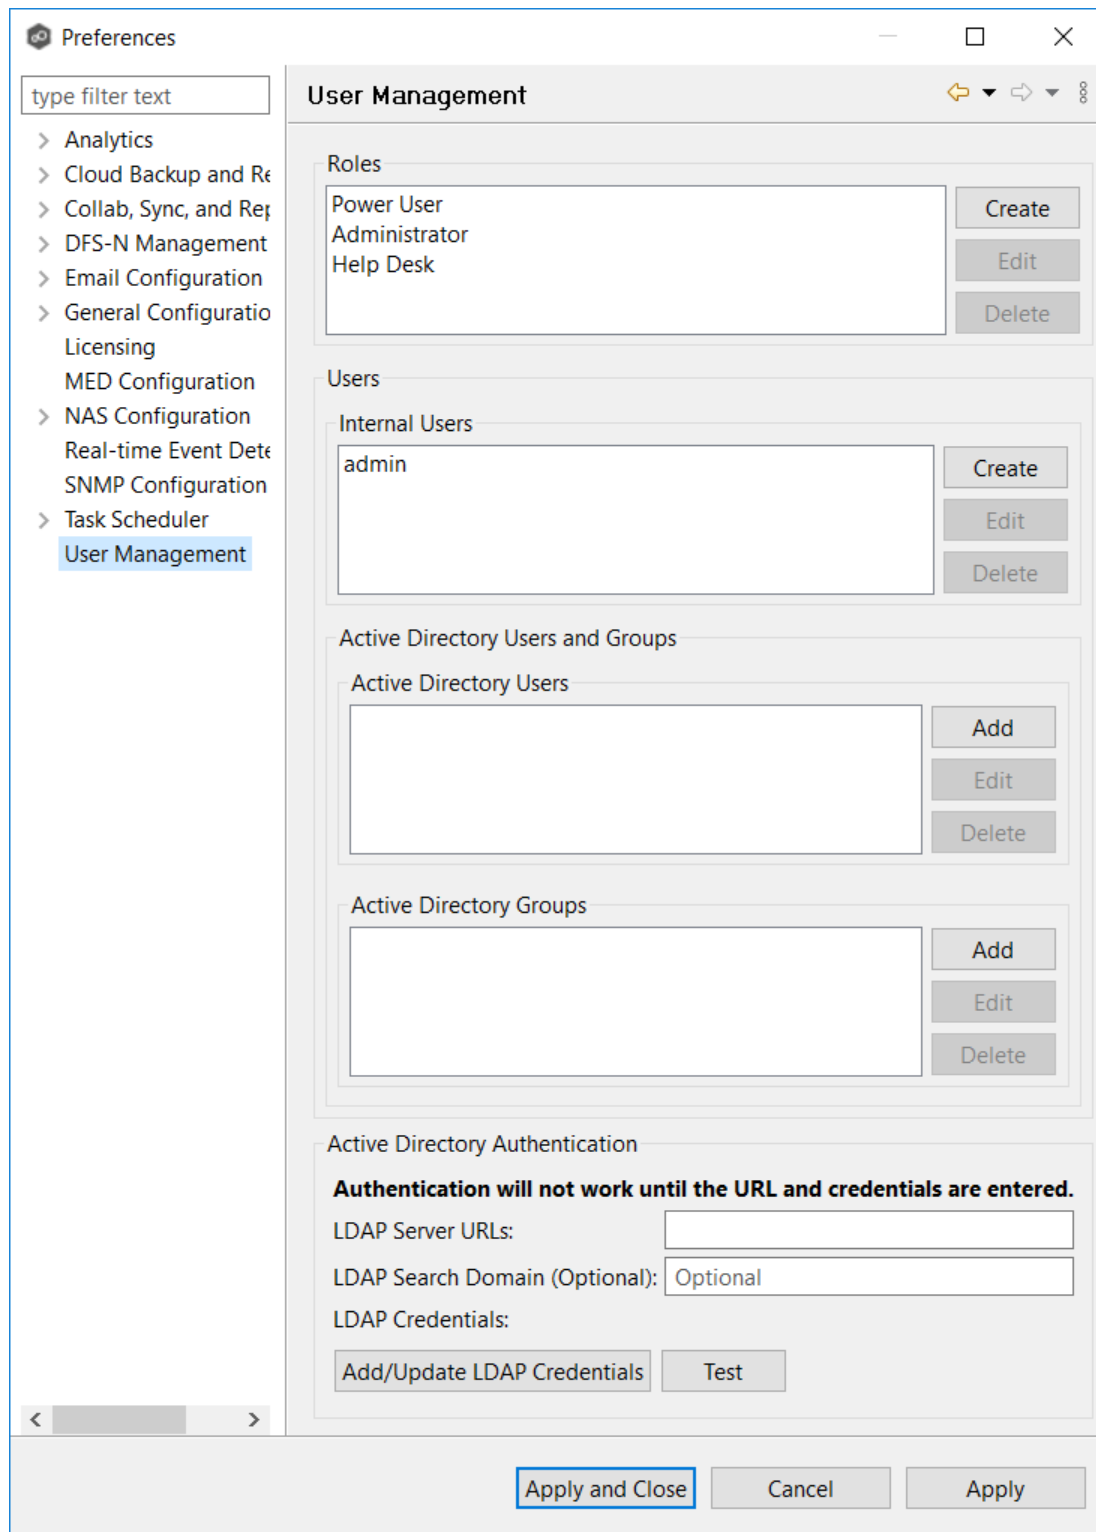
The LDAP user's information appears above the **Add/Update LDAP Admin User** button.

9. Click **Test** to verify the connection to the LDAP server.
10. Click **OK**.

Adding an Active Directory User or Group

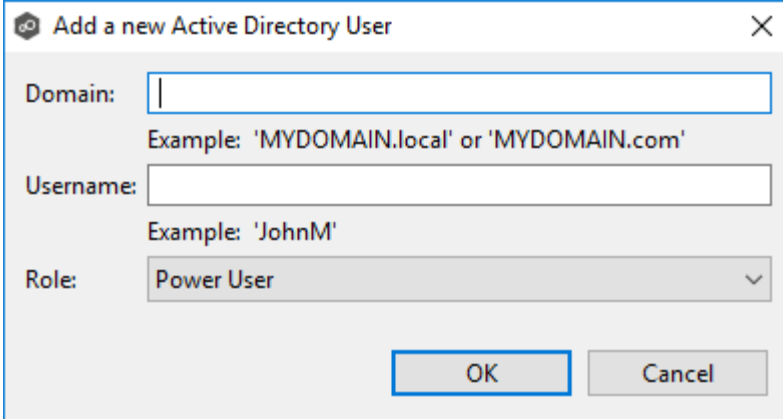
To add Active Directory users and groups to Peer Management Center, follow these steps:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **User Management** from the navigation tree.



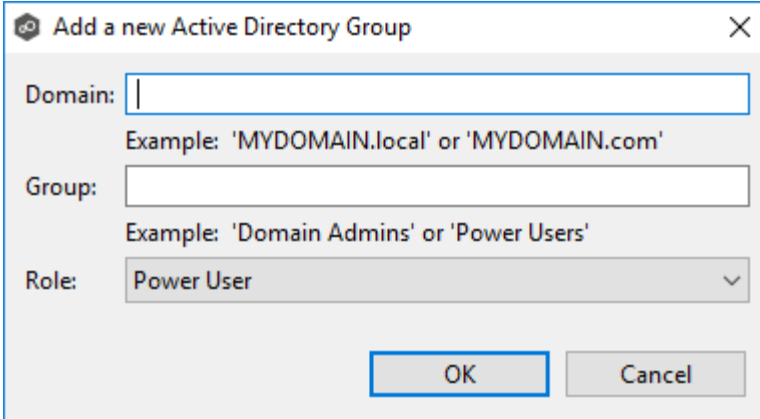
3. Add an Active Directory user or group by clicking the appropriate **Add** button.
4. Enter the information required in the dialog that appears:

- For an individual user, enter the domain name, user name, and select a role.



The screenshot shows a dialog box titled "Add a new Active Directory User". It contains three input fields: "Domain:" with a text box and an example "Example: 'MYDOMAIN.local' or 'MYDOMAIN.com'", "Username:" with a text box and an example "Example: 'JohnM'", and "Role:" with a dropdown menu currently set to "Power User". At the bottom right, there are "OK" and "Cancel" buttons.

- For a user group, enter the domain name, group name, and select a role.



The screenshot shows a dialog box titled "Add a new Active Directory Group". It contains three input fields: "Domain:" with a text box and an example "Example: 'MYDOMAIN.local' or 'MYDOMAIN.com'", "Group:" with a text box and an example "Example: 'Domain Admins' or 'Power Users'", and "Role:" with a dropdown menu currently set to "Power User". At the bottom right, there are "OK" and "Cancel" buttons.

Directory users and groups are saved in the following format:
username@mydomain.local

5. Click **OK**.

The added user or group appears in the list of Active Directory users or groups.

6. Click **OK**.

Editing an Active Directory User or Group

To edit an Active Directory user or group:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.
3. Select the AD user or group from the list of AD users or groups.
4. Click **Edit**.
5. Make the changes.
6. Click **OK**.

Deleting an Active Directory User and Group

If you delete an Active Directory user or group from Peer Management Center, that user or group will no longer have access to Peer Management Center through the web client. However, deleting the AD user or group from Peer Management Center does not delete that user or group from the Active Directory.

To delete an Active Directory user or group from Peer Management Center:

1. From the **Window** menu, select **Preferences**.
2. Select **User Management** from the navigation tree.
3. Select the AD user or group from the list of AD users or groups.
4. Click **Delete**.
5. Confirm that you want to delete the user or group.
6. Click **OK**.

Managing Roles for Web Client Users

A [role](#) determines what actions a web client user can perform and what data they can access within the web client interface.

Types of Roles

Web client users can have either a predefined (base) web role or a custom web role, providing administrators with the flexibility to manage access permissions according to their specific requirements. tailor access permissions to their specific needs:

- **Predefined (Base) Web Roles:**

- These are standard roles provided by Peer Management Center with predefined sets of permissions.

- They offer a quick and easy way to assign common access levels to users.

Base web roles cannot be modified or deleted, with one exception: [tags can be assigned to base roles](#). For a list of the permissions associated with base web roles, see [Base Web Role Permissions](#).

- **Custom Web Roles:**

- Admins can create custom web roles tailored to their specific requirements.

- This allows for fine-grained control over access permissions, accommodating unique user roles or organizational needs.

Multiple Roles Possible

A user can have multiple roles within the PMC. For example, the same user could access the PMC through different authentication methods:

- As an internal user authenticated through the PMC database with an Admin role.
- Using an Active Directory account with a Power User role.

Each authentication method may grant the user different permissions and access levels within PMC. By having multiple web roles, the user can fulfill different roles or responsibilities depending on how they access the system.

Managing Custom Roles

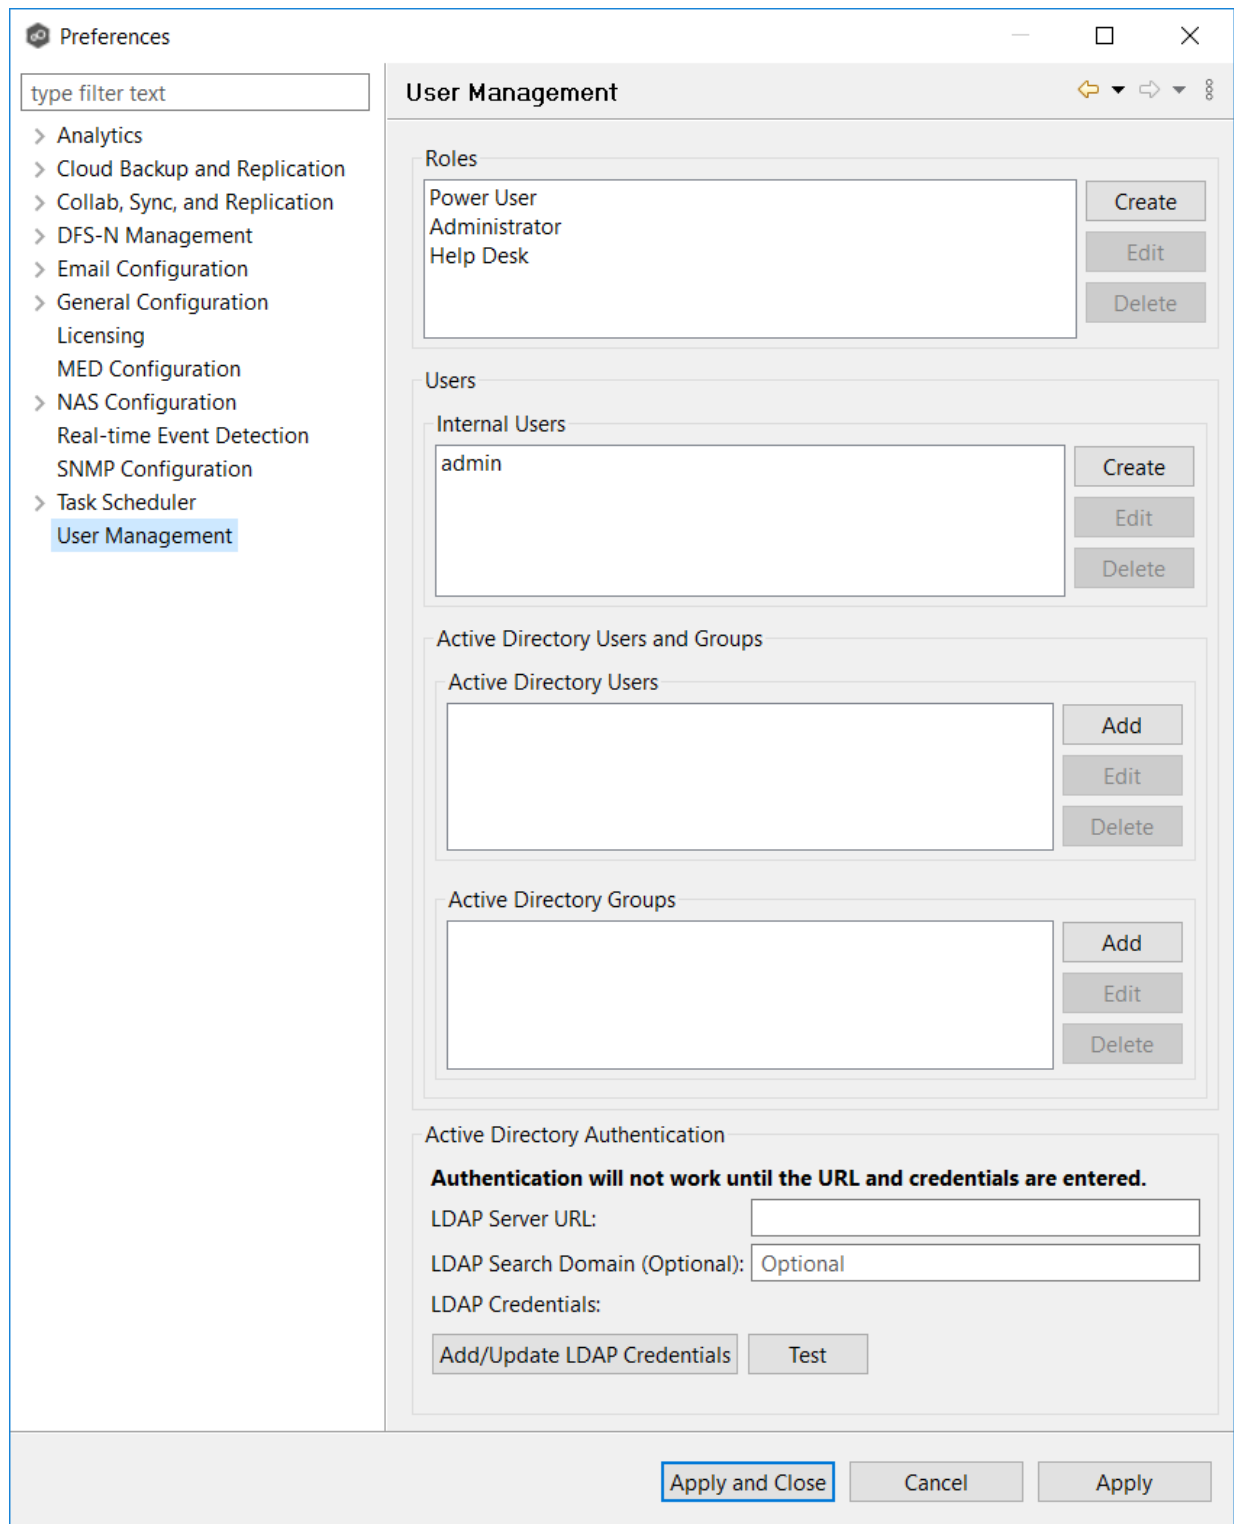
Managing web roles involves:

- [Creating a Custom Role](#)

- [Editing a Custom Role](#)
- [Deleting a Custom Role](#)
- [Assigning Tags to a Role](#)

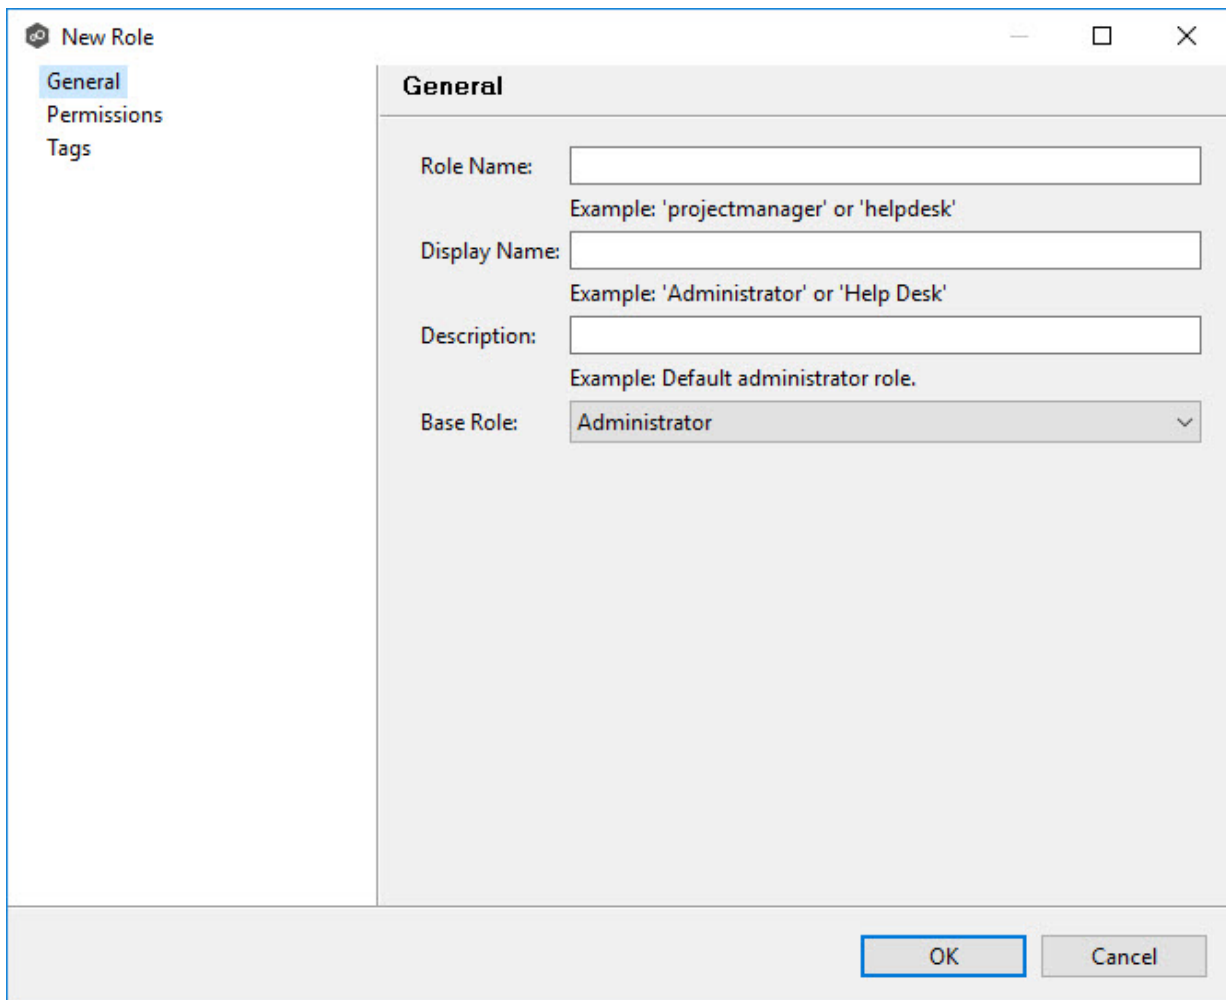
To create a custom role:

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



3. Click the **Create** button in the **Roles** section.

The **General** tab of **New Role** dialog is displayed.



The screenshot shows a 'New Role' dialog box with a sidebar on the left containing 'General', 'Permissions', and 'Tags'. The 'General' tab is selected. The main area contains the following fields:

- Role Name:** A text input field with the example: 'projectmanager' or 'helpdesk'.
- Display Name:** A text input field with the example: 'Administrator' or 'Help Desk'.
- Description:** A text input field with the example: 'Default administrator role.'
- Base Role:** A dropdown menu currently showing 'Administrator'.

At the bottom right, there are 'OK' and 'Cancel' buttons.

4. Enter the following information:

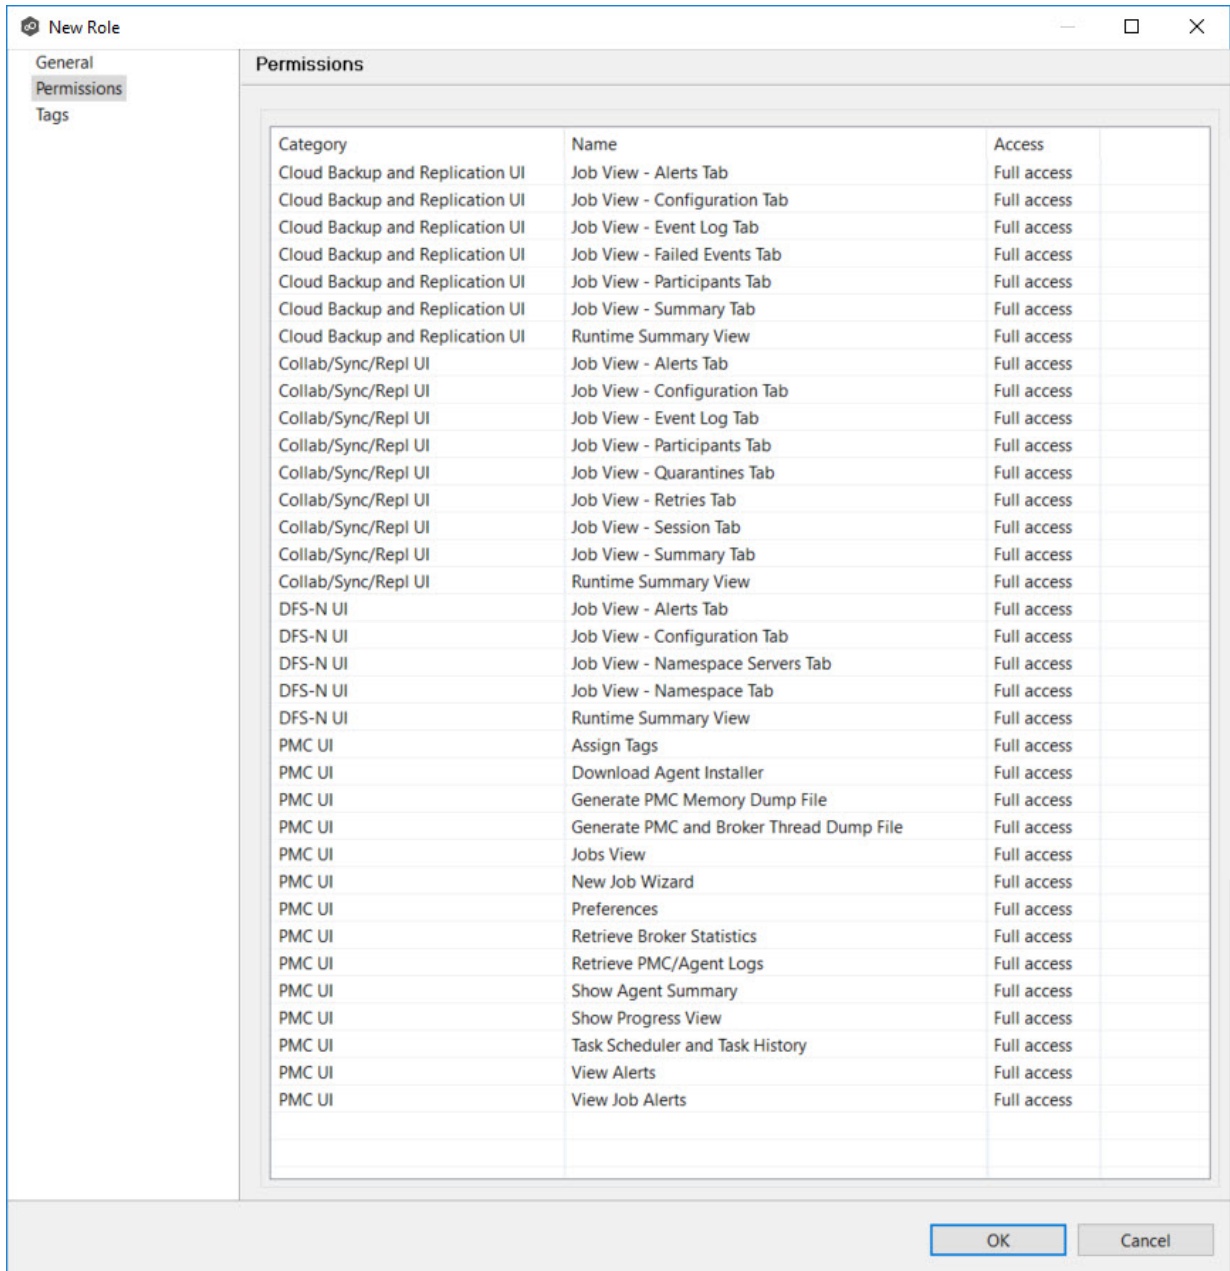
- **Role Name:** Role Name can contain only letters and numbers; it cannot contain any spaces or special characters. The maximum number of characters is 20. The Role Name is used in the internal Peer Management Center database.
- **Display Name:** Display Name can contain spaces and special characters, in addition to letters and numbers. The Display Name is displayed in the PMC user interface and reports.
- **Description:** (Optional) Use the Description to provide a brief summary of the intended use of the role.

5. Select a [base role](#) on which to base the custom role.

6. Select the **Permissions** tab.

The **Permissions** tab displays a table of [the permissions that are available to be modified for the new role](#). The **Access** column displays the current level of access that

the role has to the resource. The three levels of access are **Full access**, **View-only access**, and **No access**.



7. For each permission that you want to modify, click in the **Access** column, and then select the access level that you want for the new role.
8. (Optional) Click **Tags** to assign tags to the role.

To edit a role, select the role in the **Roles** section in the **User Management** page, and then click **Edit**.

To delete a custom web role, select the role in the **Roles** section in the **User Management** page, and then click **Delete**.

Note: You cannot delete a base web role.

For information about assigning a tag to a role, see [Assigning Tags](#).

Cloud Backup and Replication Jobs

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

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

- [Overview](#)
- [Before You Create Your First Cloud Backup and Replication Job](#)
- [Creating a Cloud Backup and Replication Job](#)
- [Running a Cloud Backup and Replication Job](#)
- [Monitoring Your Cloud Backup and Replication Jobs](#)

- [Deleting a Cloud Backup and Replication Job](#)
- [Recovering Data from the Cloud](#)

Overview

A **Cloud Backup and Replication job** brings file- to-object replication into Peer Software's capabilities for enterprise NAS environments. Leveraging the same real-time engine that powers Peer Software's multi-site, multi-vendor replication, Cloud Backup and Replication efficiently pushes data into Microsoft Azure or Amazon S3 storage in an open format that is immediately consumable by other applications and services.

Use cases for Cloud Backup and Replication include: (1) pushing exact replicas of on-premises data sets into object storage for use with burstable compute and cloud-borne services and (2) tape replacement-style backup to object with point-in-time recovery capability.

Before You Create Your First Cloud Backup and Replication Job

We strongly recommend that you configure the [Cloud Backup and Replication settings](#) (including [proxy configurations](#)), as well as other global settings such as SMTP configuration, email alerts, and before configuring your first Cloud Backup and Replication job. See [Preferences](#) for details on what and how to configure these settings.

In addition, we recommend that you set up your destination storage account before creating the job.

Creating a Cloud Backup and Replication Job

The **Create Job Wizard** walks you through the process of creating a Cloud Backup and Replication job. The process consists of the following steps:

[Step 1: Job Type and Name](#)

[Step 2: Source Storage Platform](#)

[Step 3: Management Agent](#)

[Step 4: Proxy Configuration](#)

[Step 5: Storage Information](#)

[Step 6: Source Paths](#)

[Step 7: File and Folder Filters](#)

[Step 8: Destination](#)

[Step 9: Destination Credentials](#)

[Step 10: Container or Bucket Details](#)

[Step 11: Replication and Retention Policy](#)

[Step 12: Replication Schedule](#)

[Step 13: Retention](#)

[Step 14: Source Snapshots](#)

[Step 15: Miscellaneous Options](#)

[Step 16: Email Alerts](#)


[Step 17: SNMP Notifications](#)

[Step 18: Confirmation](#)

Step 1: Job Type and Name


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

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



- Click **Cloud Backup and Replication**, and then click **Create**.
- Enter a name for the job in the dialog that appears.

The job name must be unique.



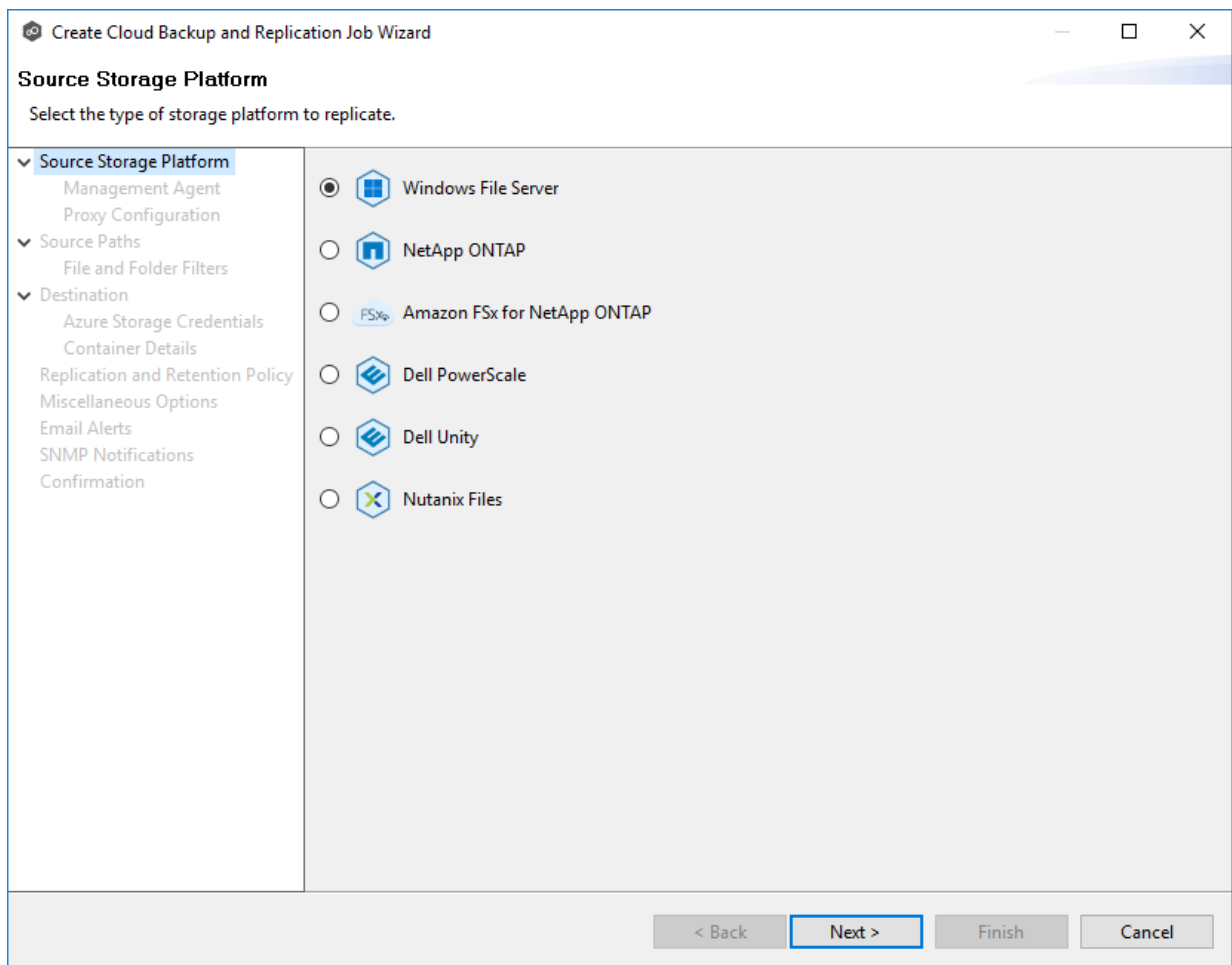
- Click **OK**.

The [Source Storage Platform](#) page appears.

Step 2: Source Storage Platform

The **Storage Platform** page lists the available storage platforms types.

1. Select the storage platform type hosting the data you want to replicate.



1. Click **Next**.

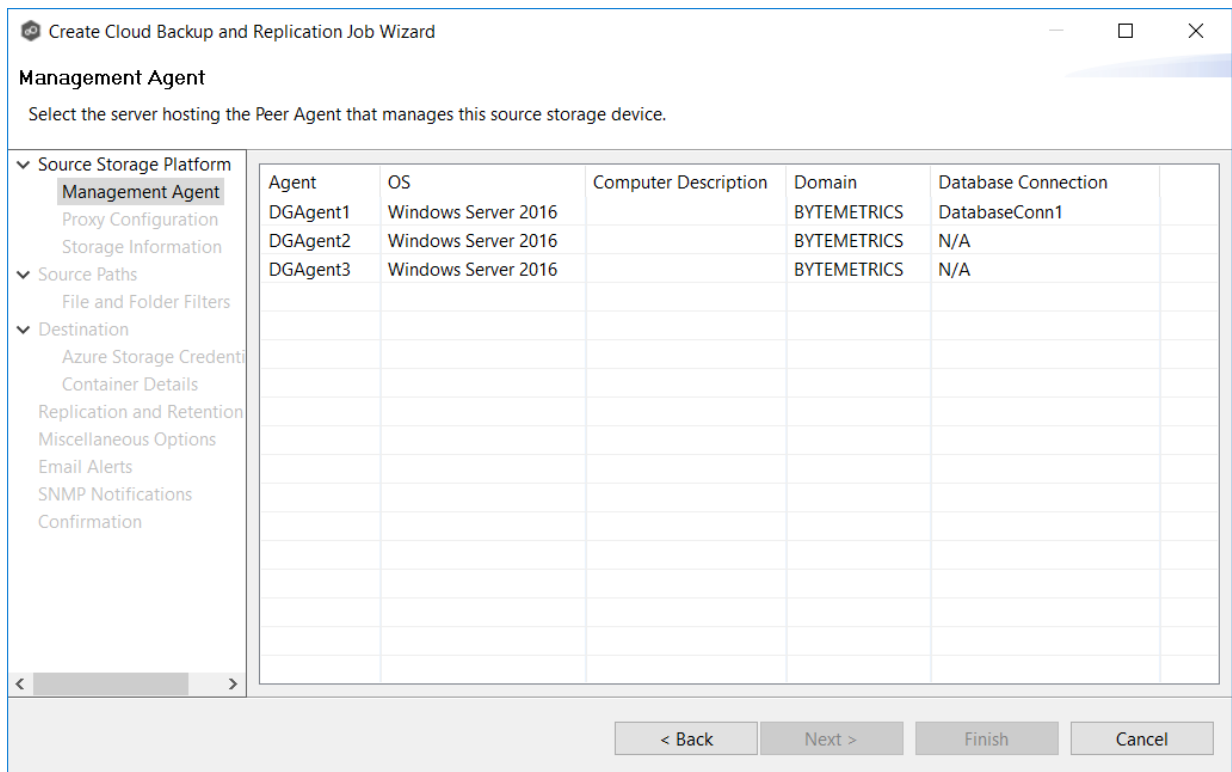
The [Management Agent](#) page appears.

Step 3: Management Agent

Each storage device that you want to replicate must have a Peer Agent that manages that device. The Peer Agent that manages a device is known as its [Management Agent](#). You can have more than one Agent managing a storage device—however, the Agents must be managing different volumes, shares, or folders on that storage device.


The **Management Agent** page lists the available Agents. In this step, you should select the Agent that manages the volumes, shares, or folders you want to replicate in this job.

1. Select the Management Agent for the volume, share, or folder you want replicated.



Tip: If the Agent you want is not listed, the Peer Agent Service may not be running on the server hosting the Agent. Try restarting the Peer Agent Service. If the service successfully connects to the [Peer Management Broker](#), then the list of available agents will be updated with that Agent.

2. If you select an Agent that does not have a database connection listed in the **Database Connection** column, a message prompts you to create the connection:



3. Click **OK**, and then configure the database connection for the selected Management Agent.

See [Database Connections](#) for instructions about creating a database connection.

4. Click **Next**.

The [Proxy Configuration](#) page appears.

Step 4: Proxy Configuration

If you do not need a proxy server to connect to outside networks, skip this step and proceed to [Step 5](#).

If you do need a proxy server to connect to outside networks, you have three options:

- Create a new proxy configuration.
- Use the existing proxy configuration. If there is an existing proxy configuration, details about the configuration will be displayed on the page.

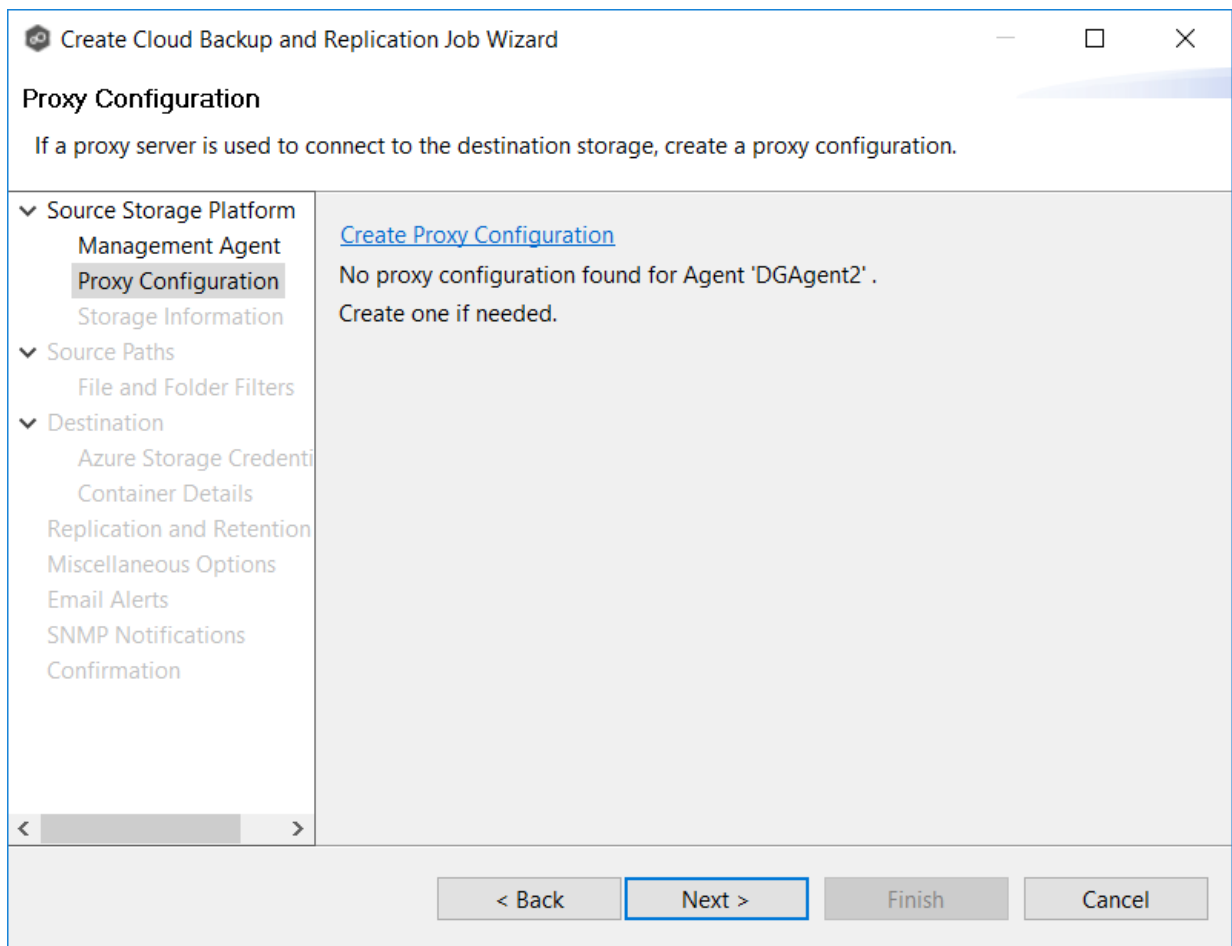
You may have created one in advance through [Cloud Back and Replication Preferences](#) or when you created another Cloud Backup and Replication job. Once a proxy configuration is created for a source storage platform, that proxy configuration is used for all Cloud Backup and Replication jobs using that agent.

- Edit an existing proxy configuration. Click the **Edit Proxy Configuration** link to edit the existing proxy.

If you edit the proxy configuration, it affects other jobs using the same Agent. Editing an existing proxy configuration has the potential to create problems with the other jobs.

If there is not an existing proxy configuration for the selected management agent, follow these steps to create a new proxy configuration:

1. Click **Create Proxy Configuration**.



The **Proxy Configuration** page is displayed. Existing proxies are listed in the Proxy Configuration table.


Field	Description
Address	Enter the IP address or fully qualified domain name of the proxy server.
Port	Enter the port number.
User Authentication	Select this checkbox if the proxy server requires authentication.

4. If your proxy server requires authentication, click the **User Authentication** checkbox, and then supply the necessary values.


Field	Description
Domain	Enter the domain name on the proxy server.
Username	Enter the user name for the proxy server.
Password	Enter the password for the proxy server.

5. Click **OK**.

If you already have jobs managed by this Agent, a message appears and identifies those jobs. They will now use the proxy as well.



After you click OK, the **Proxy Configuration** page is redisplayed. The proxy you just created now appears in the table.



11. Click **Next**.

The [Storage Information](#) page appears.

Step 5: Storage Information

On the **Storage Information** page, you will select the storage device containing data that you want to replicate and enter other information about the storage device. The contents of the **Storage Information** page varies, depending on your selection on the [Storage Platform](#) page:

- If you selected **Windows File Server**, you are prompted for configuration information relating to Windows File Server. If you selected **Windows File Server** in [Step 2](#), this page doesn't appear; skip to [Step 5: Source Paths](#)

- If you selected any other type of storage device other than Windows File Server, the **Storage Information** page requests the credentials necessary to connect to that storage device. Continue with the steps below.

For storage platforms other than Windows File Server:

1. Select **New Credentials** to enter a new set of credentials for the source storage platform or select **Existing Credentials**.
2. If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue with [Step 5: Source Paths](#).

If you selected **New Credentials**, enter the credentials for connecting to the source storage device. The information you are prompted to enter varies, depending on the type of storage platform:

[Amazon FSxN](#)

[Dell PowerScale](#)

[Dell Unity](#)

[NetApp ONTAP](#)

[Nutanix Files](#)

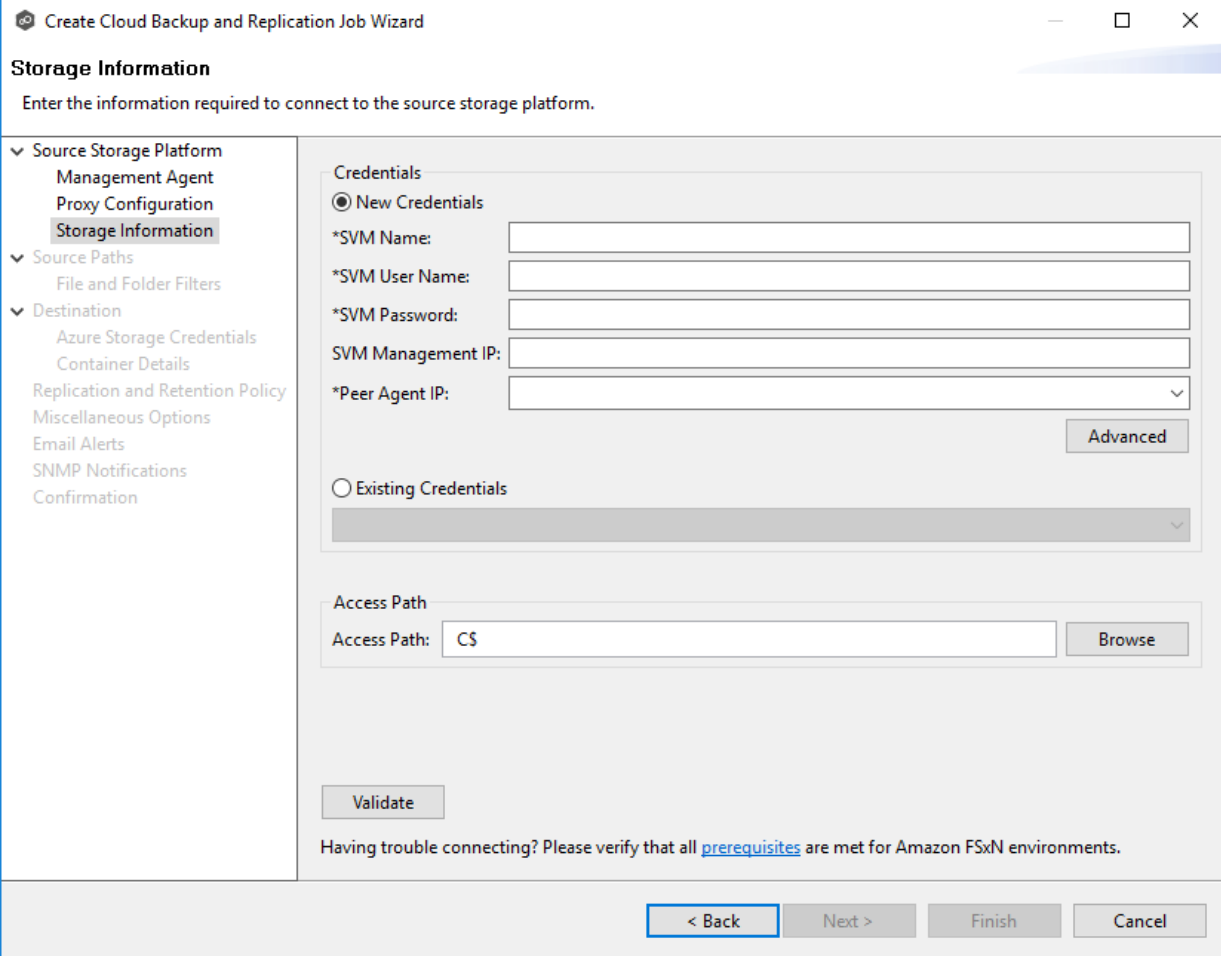
3. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

4. Click **Next**.

The [Source Paths](#) page appears.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.



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

Source Storage Platform
Management Agent
Proxy Configuration
Storage Information

Source Paths
File and Folder Filters

Destination
Azure Storage Credentials
Container Details
Replication and Retention Policy
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

Credentials

New Credentials

*SVM Name:

*SVM User Name:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

Advanced

Existing Credentials

Access Path

Access Path: Browse

Validate

Having trouble connecting? Please verify that all [prerequisites](#) are met for Amazon FSxN environments.

< Back Next > Finish Cancel

- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Source Paths](#) page.

If you selected **New Credentials**, enter the credentials for connecting to the storage device.

Option	Description
SVM Name	Enter the name, FQDN, or IP address of the Storage Virtual Machine (SVM) 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.

Option	Description
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	Optional. Enter the IP address used to access the management API of the Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already permit management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually enter the address.
Access Path	Use only when experiencing access issues. Contact the Peer Software Support team for more information.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

5. Click **Next**.

The [Source Paths](#) page is displayed.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Dell PowerScale cluster hosting the data to be replicated or select existing credentials.

The information required will vary, depending on whether you select **Syslog** or **RabbitMQ** as the connection type, due to the distinct protocols and mechanisms they employ for communication.

Syslog

Create Cloud Backup and Replication Job Wizard

Storage Information

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

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
 - Storage Information**
- Source Paths
 - File and Folder Filters
- Destination
 - Azure Storage Credentials
 - Container Details
 - Replication and Retention Policy
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation

Credentials

New Credentials

*Cluster Name:

*Cluster Management IP:

*Cluster Username:

*Cluster Password:

Cluster Access Zone:

Connection Type: Syslog RabbitMQ

Syslog

*Agent IP Address:

*Listening Port:

*SSL Certificate Path:

*SSL Private Key Path:

SSL Private Key Password:

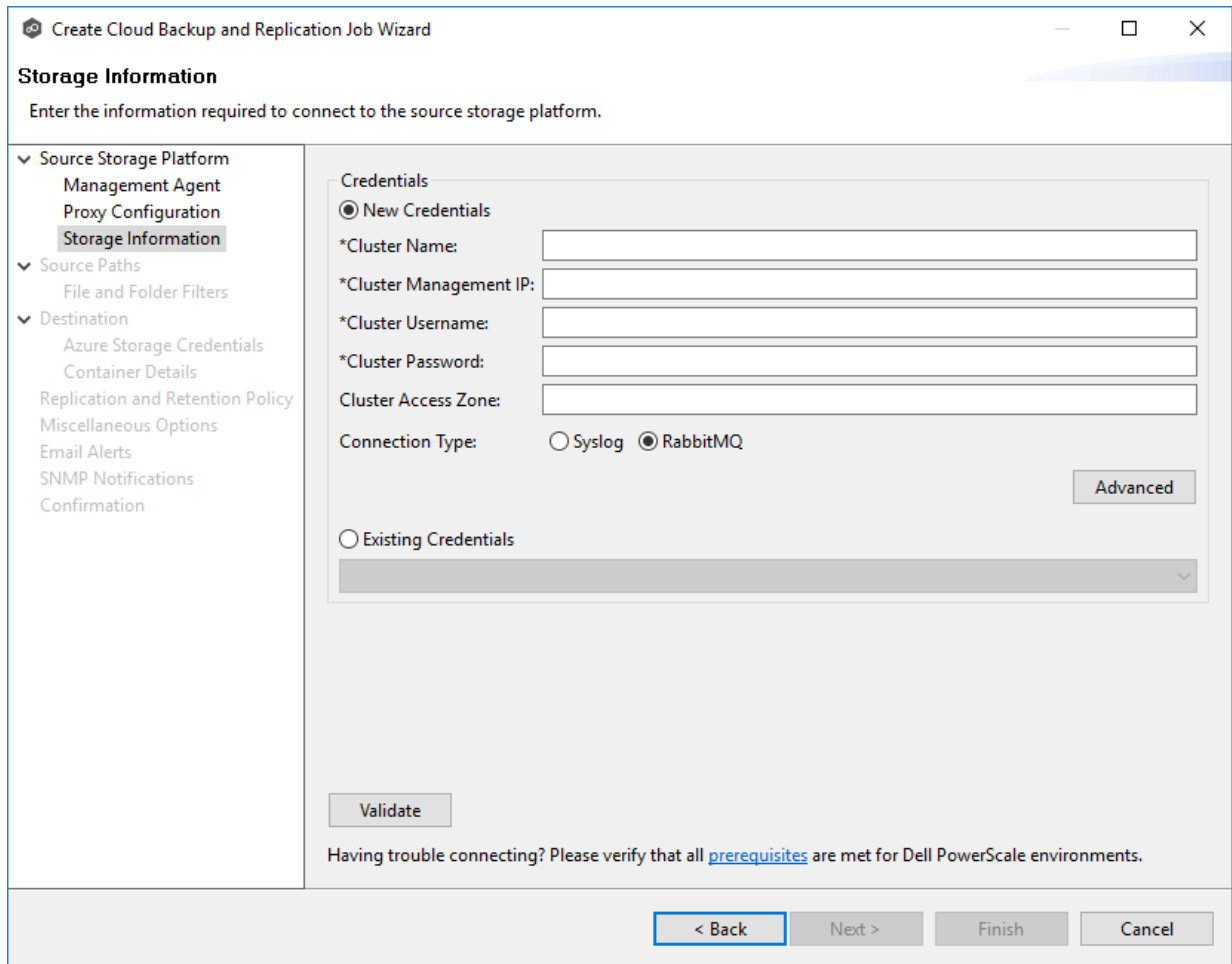
Existing Credentials

You must enter a Cluster Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments.

< Back Next > Finish Cancel

RabbitMQ



- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Source Paths](#) page.

If you selected **New Credentials**, enter the credentials for connecting to the storage device.

Field	Description
Cluster Name	Enter the name of the PowerScale cluster hosting the data to be replicated.
Cluster Management IP	Enter the IP address to use to access the REST-based API integrated into the PowerScale cluster. Required only if multiple Access Zones are in use on the cluster.

Field	Description
Cluster Username	Enter the user name for the account managing the PowerScale cluster.
Cluster Password	Enter the password for account managing the PowerScale cluster.
Cluster Access Zone	Optional. Enter the name of the access zone that is being monitored.
Connection Type	<p>Select the appropriate method for sending real-time event notifications to the Agent:</p> <ul style="list-style-type: none"> • Opt for Syslog if the storage device directly transmits notifications to the Agent. • Opt for RabbitMQ if you're utilizing the CEE framework to dispatch notifications to the Agent.

3. If you selected **Syslog**, you will need to provide values for the following fields:

Field	Description
Agent IP Address	Select the IP address of the server hosting the Agent that manages the PowerScale cluster. The cluster must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.
Listening Port	Enter the port over which the Agent will receive TLS-based syslog events from the PowerScale cluster.
SSL Certificate Path	Enter the path to the certificate to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private	Enter the path to the private key to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster.

Field	Description
Key Path	For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private Key Password	[Optional] If your private key is protected with a password, enter it here. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .

4. Click **Advanced** if you want to set [advanced options](#).

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

After the credentials are validated, a success message appears.

6. Click **Next**.

The [Source Paths](#) page is displayed.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the CIFS Server hosting the data to be replicated.

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

Source Storage Platform
Management Agent
Proxy Configuration
Storage Information

Source Paths
File and Folder Filters

Destination
Azure Storage Credentials
Container Details
Replication and Retention Policy
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

Credentials

New Credentials

*CIFS Server Name:

*Unisphere Management IP:

*Unisphere Username:

*Unisphere Password:

Advanced

Existing Credentials

Access Path

Access Path: Browse

Validate You must enter a CIFS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments.

< Back Next > Finish Cancel

- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Source Paths](#) page.

If you selected **New Credentials**, enter the credentials for connecting to the storage device.

Field	Description
CIFS Server Name	Enter the name of the NAS server hosting the data to be replicated.
Unisphere Managem	Enter the IP address of the Unisphere system used to manage the Unity storage device. The IP address should not point to the NAS

Field	Description
ent IP	server.
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.
Access Path	Use only when experiencing access issues. Contact Peer Software support for more information.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

5. Click **Next**.

The [Source Paths](#) page is displayed.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

Create Cloud Backup and Replication Job Wizard

Storage Information

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

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
 - Storage Information**
- Source Paths
 - File and Folder Filters
- Destination
 - Azure Storage Credentials
 - Container Details
 - Replication and Retention Policy
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation

Credentials

New Credentials

*SVM Name:

*SVM User Name:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

Existing Credentials

Access Path

Access Path:

Having trouble connecting? Please verify that all [prerequisites](#) are met for NetApp ONTAP environments.

- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Source Paths](#) page.

If you selected **New Credentials**, enter the credentials for connecting to the storage device.

Field	Description
SVM Name	Enter the name, FQDN, or IP address of the Storage Virtual Machine (SVM) 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.

Field	Description
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	Enter the IP address used to access the management API of the NetApp Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already allow management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.
Access Path	Use only when experiencing access issues. Contact Peer Software support for more information.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

5. Click **Next**.

The [Source Paths](#) page is displayed.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated.

Create Cloud Backup and Replication Job Wizard

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

Source Storage Platform
Management Agent
Proxy Configuration
Storage Information

Source Paths
File and Folder Filters

Destination
Azure Storage Credentials
Container Details
Replication and Retention Policy
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

Credentials

New Credentials

*Nutanix File Server Name:

*Username:

*Password:

*Peer Agent IP:

Advanced

Existing Credentials

Validate

Having trouble connecting? Please verify that all [prerequisites](#) are met for Nutanix Files environments.

< Back Next > Finish Cancel

- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Source Paths](#) page.

If you selected **New Credentials**, enter the credentials for connecting to the storage device.

Field	Description
Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.

Field	Description
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Enter the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually enter the address. The IP address should not point to the Files cluster itself.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

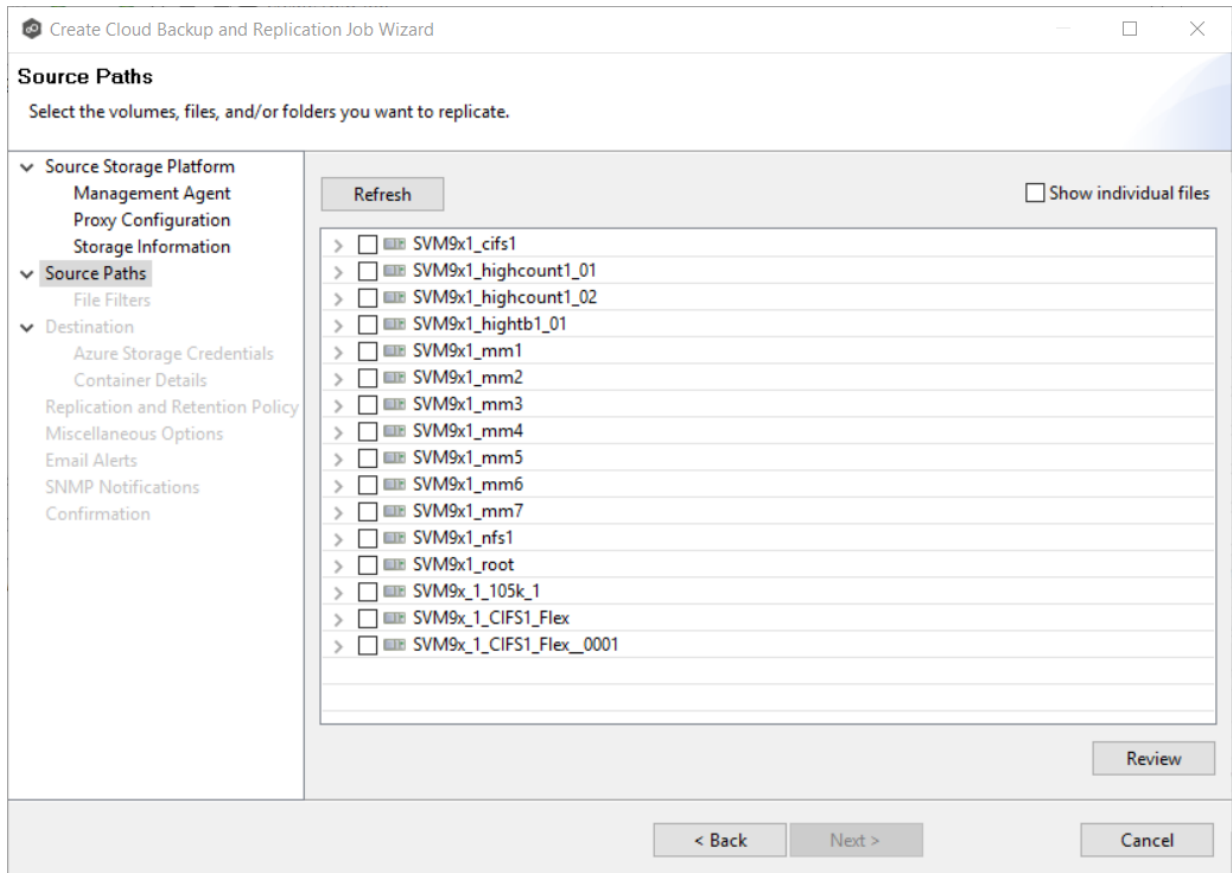
5. Click **Next**.

The [Source Paths](#) page is displayed.

Step 6: Source Paths

The **Source Paths** page displays a list of available volumes to replicate. You can choose to replicate an entire volume or selectively replicate files and folders. The files/folders/volumes selected for replication are referred to as the [watch set](#).

1. Select the paths to the files/folders/volumes you want to replicate.



Unlike other job types, you can select multiple folders to replicate:

Option	Action
The entire volume (all files and folders, including subfolders and their files)	Select the volume checkbox.
All files at the root level of the volume (but no folders)	Expand the volume, scroll to the bottom of the expanded list, and then select All Files .
A specific folder and its content (including subfolders and their files)	Expand the volume, find the desired folder, and then select its checkbox.
All files within a specific folder (but not the folder)	Expand the folder and select All Files .

Option	Action
Specific files and folders	Select the Show individual files checkbox, expand the folders, and then select the files and folders you want to replicate.

2. (Optional) Click the **Review** button to see your selections.
3. Click **Next**.

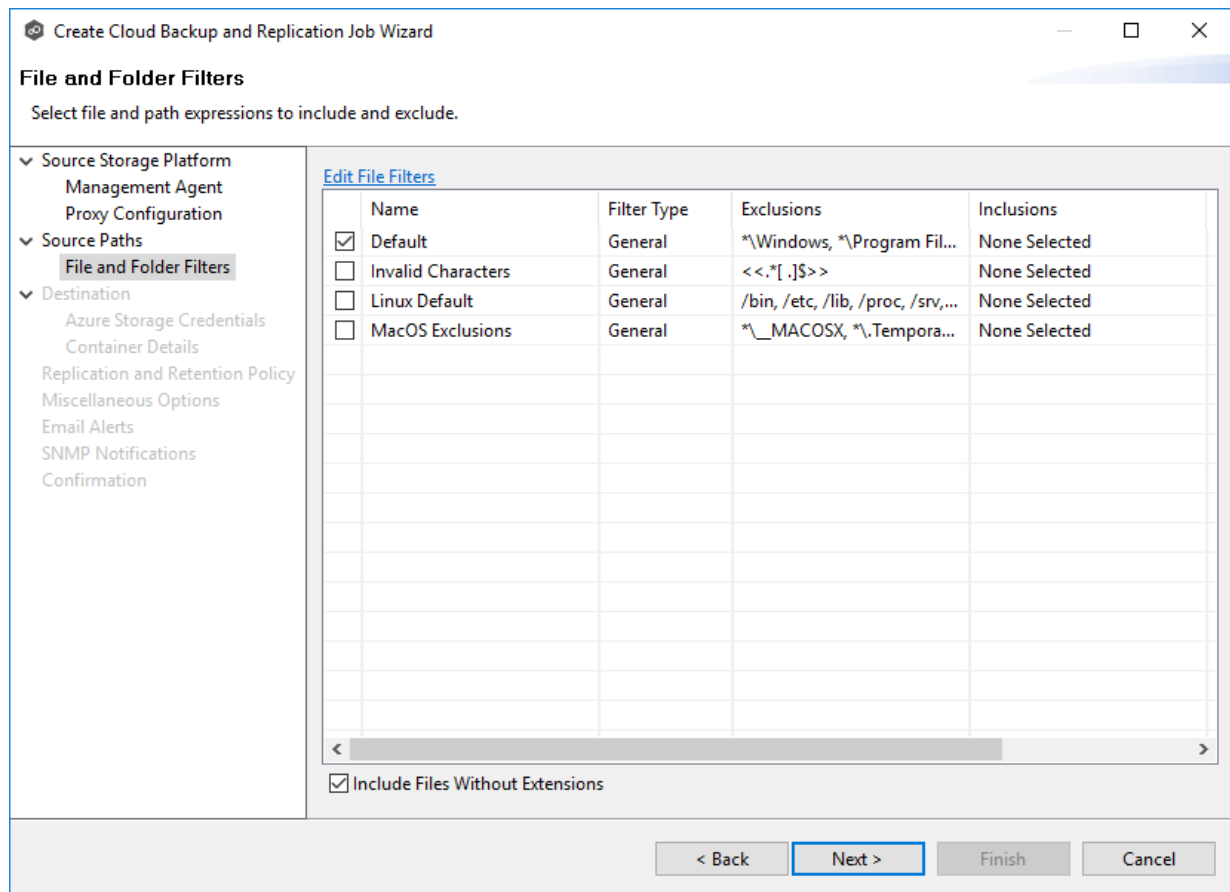
The [File and Folder Filters](#) page appears.

Step 7: File and Folder Filters

The **File and Folder Filters** page displays a list of [file and folder filters](#). By default, all files and folders selected in the **Source Paths** page will be replicated. A file or folder filter enables you to exclude and/or include files and folders from the job based on file type, extension, name, or directory path. Any file or folder that matches the filter is excluded or included from replication, depending on the filter's definition.

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

If you want to create a new file or folder filter or modify an existing one, click **Edit File Filters**. See [File and Folder Filters](#) in the [Preferences](#) section for information about creating or modifying a file filter.



2. Select the **Include Files Without Extensions** checkbox if you want to replicate files that do not have extensions.

Note: Files without extensions are ignored during replication unless you select this checkbox.

3. Click **Next**.

The [Destination](#) page appears.

Step 8: Destination

The **Destination** page displays a list of the available storage platforms to which Cloud Backup and Replication can replicate. Currently, the following platforms are supported:

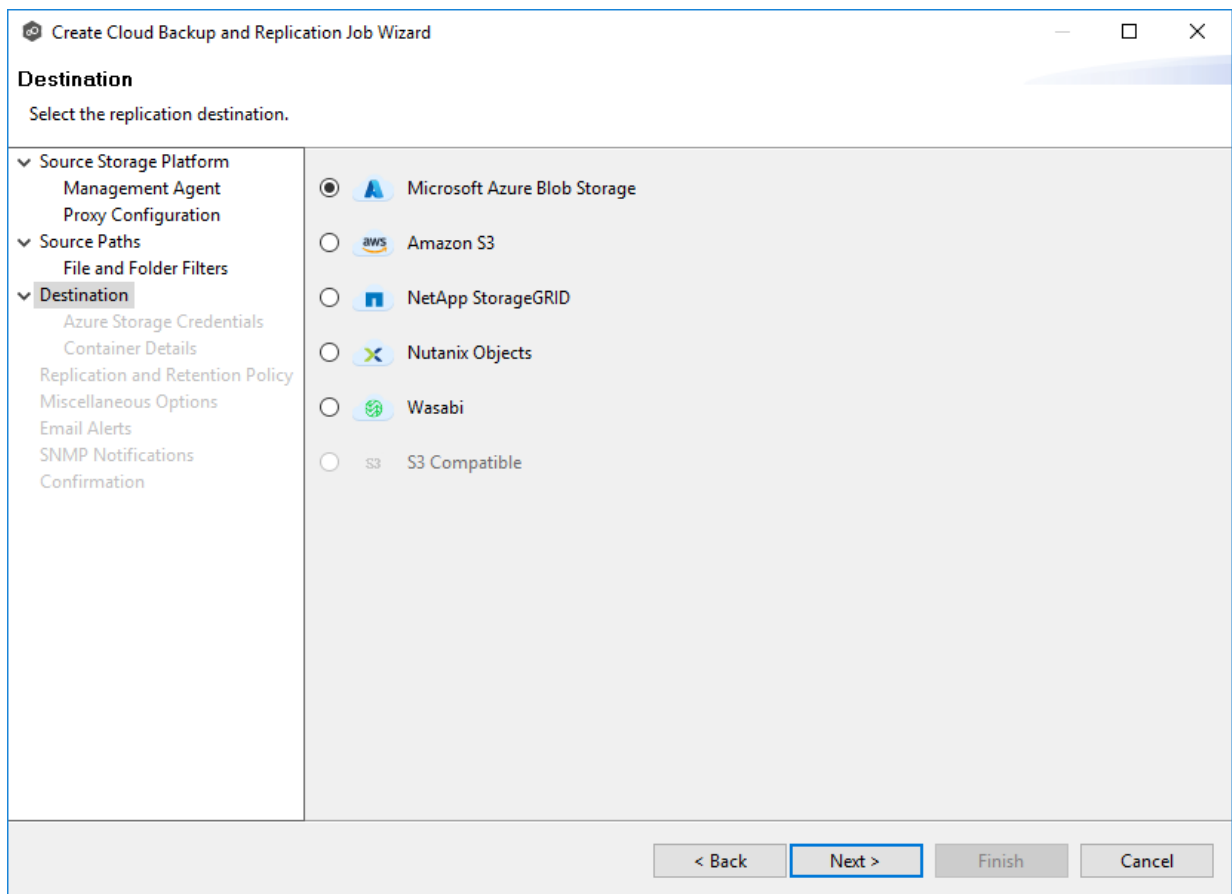
- Microsoft Azure
- Amazon S3

- NetApp StorageGRID
- Nutanix Objects
- Wasabi

In addition, some S3-compatible platforms are also supported. Contact your Peer Software Sales representative to see if the S3 compatible platform you want to use is supported.

Important: You should create the storage account before creating the Cloud Backup and Replication job.

1. Select the type of destination storage platform.



2. Click **Next**.

The [Destination Credentials](#) page appears.

Step 9: Destination Credentials

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

1. Select **New Credentials** to enter a new set of credentials for the destination storage device or select **Existing Credentials**.
2. If you selected **Existing Credentials**, select a credential from the drop-down list.

If you selected **New Credentials**, enter the credentials for connecting to the destination storage account. The information you are prompted to enter varies, depending on the type of storage platform:

[Azure Blob Storage Credentials](#)

[Amazon S3 Credentials](#)

[NetApp StorageGRID](#)

[Nutanix Objects](#)

[Wasabi Credentials](#)

3. Click **Next**.

The **Details** page for the selected destination storage account.

1. Enter the credentials to connect to a Microsoft Azure storage account. General Purpose and Blob storage accounts are supported.

Create Cloud Backup and Replication Job Wizard

Azure Storage Credentials

Enter new credentials or select existing credentials.

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
 - Storage Information
- Source Paths
 - File Filters
- Destination
 - Azure Storage Credentials**
 - Container Details
 - Replication and Retention Policy
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation

Credentials

New Credentials

*Description:

*Account:

*Shared Key: Show Key

*Endpoint Type: **Public**

Use SSL

Existing Credentials

Azure Storage

< Back
Next >
Cancel

Field	Description
Description	Enter a name for the credentials.
Account	Enter the name of the Azure storage account, which can be found in the Azure Portal.
Shared Key	Enter one of the shared keys for the Azure Storage account. The shared keys can be found in the Azure Portal.
Endpoint Type	Select the type of data center endpoint. The options are: Public , Germany , China , US Government , and Custom .
Endpoint	If you selected Custom for Endpoint Type , the Endpoint field appears. Enter the IP address of the endpoint.

Field	Description
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

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

If you are using a proxy in your environment and you get an error while trying to validate, you may want to check the [proxy configuration](#) in [Preferences](#).

3. Click **Next**.

The [Container Details](#) page appears.

1. Enter the credentials to connect to an Amazon S3 storage account.

Create Cloud Backup and Replication Job Wizard

Amazon S3 Credentials
Enter new credentials or select existing credentials.

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
 - Storage Information
- Source Paths
 - File Filters
- Destination
 - Amazon S3 Credentials**
 - Bucket Details
 - Replication and Retention Policy
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation

Credentials

New Credentials

*Description:

*Access Key:

*Secret Key: Show Key

Use SSL

Existing Credentials

Field	Description
Description	Enter a name for the credentials.
Access Key	Enter one of the shared keys of the Amazon S3 Storage account, which can be found in the Amazon AWS portal.
Secret Key	Enter the secret key of the Amazon S3 Storage account, which can be found in Amazon AWS portal.
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

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

If you are using a proxy in your environment and you get an error while trying to validate, you may want to check the [proxy configuration](#) in [Preferences](#).

3. Click **Next**.

The [Bucket Details](#) page appears.

1. Enter the credentials to connect to a NetApp StorageGRID storage account.

Field	Description
Description	Enter a name for the credentials.
Access Key	Enter one of the shared keys of the NetApp StorageGRID account, which can be found in the Tenant Manager.
Secret Key	Enter the secret key of the NetApp StorageGRID account, which can be found in the Tenant Manager.

Field	Description
Service Point	Enter the IP or name of the object store.
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

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

The [Container Details](#) page appears.

1. Enter the credentials to connect to a Nutanix Objects storage account.

Create Cloud Backup and Replication Job Wizard

Nutanix Objects Credentials
Enter new credentials or select existing credentials.

New Credentials
 *Description:
 *Access Key:
 *Secret Key: Show Key
 *Service Point:
 Use SSL
 Existing Credentials

Validate

< Back Next > Cancel

Field	Description
Description	Enter a name for the credentials.
Access Key	Enter one of the shared keys of the Nutanix Objects account, which can be found in Prism Central.
Secret Key	Enter the secret key of the Nutanix Objects account, which can be found in Prism Central.
Service Point	Enter the IP or name of the object store.

Field	Description
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

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

The [Container Details](#) page appears.

1. Enter the credentials to connect to a Wasabi storage account.

Create Cloud Backup and Replication Job Wizard

Wasabi Credentials

Enter new credentials or select existing credentials.

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
- Source Paths
 - File and Folder Filters
- Destination
 - Wasabi Credentials**
 - Bucket Details
 - Replication and Retention Policy
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation

Credentials

New Credentials

*Description:

*Access Key:

*Secret Key: Show Key

*Service Point:

Use SSL

Existing Credentials

Validate

< Back Next > Finish Cancel

Field	Description
Description	Enter a name for the credentials.
Access Key	Enter one of the shared keys of the Wasabi account.
Secret Key	Enter the secret key of the Wasabi account.
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

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

If you are using a proxy in your environment and you get an error while trying to validate, you may want to check the [proxy configuration](#) in [Preferences](#).

3. Click **Next**.

The [Bucket Details](#) page appears.

Step 10: Container or Bucket Details

The **Container Details** or **Bucket Details** page allows you to create a new storage container or bucket or choose an existing one.

1. Select **New Container/New Bucket** to create a new storage container/bucket; otherwise, select **Existing Container/Existing Bucket** to choose an existing one.
2. If you selected **Existing Container** or **Existing Bucket**, select a container or bucket from the drop-down list.

If you selected **New Container** or **New Bucket**, enter the requested information. The information you are prompted to enter varies, depending on the type of storage platform:

[Azure Blob Storage Container Details](#)

[Amazon S3 Bucket Details](#)

[NetApp StorageGRID Bucket Details](#)

[Nutanix Objects Bucket Details](#)

[Wasabi Bucket Details](#)

3. Click **Next**.

The [Replication and Retention Policy](#) page appears.

1. Select **New Container** to create a new container or select **Existing Container**.

Choose **Existing Container** if:

- You (or someone else) already created a container you want to use.
- You want to use a container that was created outside Peer Management Center.
- You don't have the permissions required to create a new container and want to use one that someone else will create.

Container Details

Create a new container or select an existing one.

Source Storage Platform

- Management Agent
- Proxy Configuration
- Storage Information

Source Paths

- File Filters

Destination

- Azure Storage Credentials
- Container Details
- Replication and Retention Policy
- Miscellaneous Options
- Email Alerts
- SNMP Notifications
- Confirmation

New Container

*Name:

dgagent2-mykirzevab

Automatically name

Existing Container

< Back Next > Cancel

2. If you selected **Existing Container**, select a container from the drop-down list. If the container does not appear in the list because the person who has the permissions to create a container has not yet created the bucket, click the **Reload** button after the container is created. The container will appear in the updated list.

If you selected **New Container**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the container; the container name must conform to the following naming rules:

- A container name must be unique.
- A container name must be a valid DNS name.
- A container name must start with a letter or number, and can contain only letters, numbers, and the dash (-) character.
- Every dash (-) character must be immediately preceded and followed by a letter or number; consecutive dashes are not permitted in container names.
- All letters in a container name must be lowercase.

- A container name must be from 3 through 63 characters long.

For more information about container names, see [Naming and referencing containers, blobs, and metadata](#).

3. Click **Next**.

The [Replication and Retention Policy](#) page appears.

1. Select **New Bucket** to create a new bucket or select **Existing Bucket**.

Choose **Existing Bucket** if:

- You (or someone else) already created a bucket you want to use.
- You want to use a bucket that was created outside Peer Management Center.
- You don't have the permissions required to create new buckets and want to use one that someone else will create.

Create Cloud Backup and Replication Job Wizard

Bucket Details

Create a new bucket or select an existing one.

Source Storage Platform
Management Agent
Proxy Configuration
Storage Information

Source Paths
File Filters

Destination
Amazon S3 Credentials
Bucket Details
Replication and Retention Policy
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

New Bucket

*Name:
dgagent2-qczkmxosg

Automatically name

*Region:
US East (N. Virginia)

Existing Bucket

< Back Next > Cancel

2. If you selected **Existing Bucket**, select a bucket from the drop-down list. If the bucket does not appear in the list because the person who has the permissions to create a bucket has not yet created the bucket, click **Reload** after the bucket is created. The bucket will appear in the updated list.

If you selected **New Bucket**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the bucket; the bucket name must conform to the following naming rules:

- A bucket name must be unique across all existing bucket names in Amazon S3 (that is, across all AWS customers). For more information, see [Bucket Restrictions and Limitations](#).
- Bucket names must comply with DNS naming conventions. For information about legacy non-DNS-compliant bucket names, see [Bucket Restrictions and Limitations](#).
- A bucket name must start with a lowercase letter or number.
- A bucket name must not contain uppercase characters or underscores.
- A bucket name must be from 3 through 63 characters long.

- A bucket name must be a series of one or more labels. Adjacent labels are separated by a single period (.). Bucket names can contain lowercase letters, numbers, and hyphens. Each label must start and end with a lowercase letter or a number.
- A bucket name must not be formatted as an IP address (for example, 192.168.5.4).
- When you use virtual hosted-style buckets with Secure Sockets Layer (SSL), the SSL wildcard certificate only matches buckets that don't contain periods. To work around this, use HTTP or write your own certificate verification logic. We recommend that you do not use periods (.) in bucket names when using virtual hosted-style buckets.
- Choose a bucket name that reflects the objects in the bucket because the bucket name is visible in the URL that points to the objects that you're going to put in your bucket. After you create the bucket, you cannot change the name, so choose wisely.

For information about naming buckets, see [Rules for Bucket Naming](#) in the Amazon Simple Storage Service Developer Guide.

3. Select the region where you want the bucket to reside.

Important: After you have created a bucket, you cannot change its region.

4. Click **Next**.

The [Replication and Retention Policy](#) page appears.

1. Select **New Bucket** to create a new bucket or select **Existing Bucket**.

Choose **Existing Bucket** if:

- You (or someone else) already created a bucket you want to use.
- You want to use a bucket that was created outside Peer Management Center.
- You don't have the permissions required to create new buckets and want to use one that someone else will create.

Bucket Details
Create a new bucket or select an existing one.

Source Storage Platform
Management Agent
Proxy Configuration
Storage Information

Source Paths
File Filters

Destination
NetApp StorageGRID Credential
Bucket Details
Replication and Retention Policy
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

New Bucket

*Name:
dgagent2-yvgjhyrdsb

Automatically name

Existing Bucket

< Back Next > Cancel

2. If you selected **Existing Bucket**, select a bucket from the drop-down list. If the bucket does not appear in the list because the person who has the permissions to create a bucket has not yet created the bucket, click **Reload** after the bucket is created. The bucket will appear in the updated list.

If you selected **New Bucket**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the bucket; the bucket name must comply with the following rules:

- Must be unique across each StorageGRID Webscale system (not just unique within the tenant account).
- Must be DNS compliant.
- Must contain between 3 and 63 characters.
- Can be a series of one or more labels, with adjacent labels separated by a period. Each label must start and end with a lowercase letter or a number and can only use lowercase letters, numbers, and hyphens.

- Must not look like a text-formatted IP address.
- Should not use periods in virtual hosted-style requests because periods will cause problems with server wildcard certificate verification.

For information about naming buckets, see [Rules for Bucket Naming](#) in the Amazon Simple Storage Service Developer Guide.

3. Click **Next**.

The [Replication and Retention Policy](#) page appears.

Choose **Existing Bucket** if:

- You (or someone else) already created a bucket you want to use.
- You want to use a bucket that was created outside Peer Management Center.
- You don't have the permissions required to create new buckets and want to use one that someone else will create.

1. Select **New Bucket** to create a new bucket or select **Existing Bucket**.

Choose **Existing Bucket** if:

- You (or someone else) already created a bucket you want to use.
- You want to use a bucket that was created outside Peer Management Center.
- You don't have the permissions required to create new buckets and want to use one that someone else will create.

Create Cloud Backup and Replication Job Wizard

Bucket Details

Create a new bucket or select an existing one.

Source Storage Platform

- Management Agent
- Proxy Configuration
- Storage Information

Source Paths

- File Filters

Destination

- Nutanix Objects Credentials
- Bucket Details**
- Replication and Retention Policy
- Miscellaneous Options
- Email Alerts
- SNMP Notifications
- Confirmation

New Bucket

*Name:

dgagent2-kolglxhhef

Automatically name

Existing Bucket

< Back Next > Cancel

2. If you selected **Existing Bucket**, select a bucket from the drop-down list. If the bucket does not appear in the list because the person who has the permissions to create a bucket has not yet created the bucket, click **Reload** after the bucket is created. The bucket will appear in the updated list.

If you selected **New Bucket**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the bucket; the bucket name must comply with the following rules:

- Must start with a number or a letter.
- Must be 3 - 255 characters long.
- Can contain lowercase letters, numbers, underscores (_), and dashes (-).
- There may be additional restrictions on bucket names in some AWS regions. We recommend that you create bucket names that are DNS-compliant, if you want to access objects using URL. For more information, see [Amazon Simple Storage Service Console user's guide](#).

3. Click **Next**.

The [Replication and Retention Policy](#) page appears.

1. Select **New Bucket** to create a new bucket or select **Existing Bucket**.

Choose **Existing Bucket** if:

- You (or someone else) already created a bucket you want to use.
- You want to use a bucket that was created outside Peer Management Center.
- You don't have the permissions required to create new buckets and want to use one that someone else will create.

Create Cloud Backup and Replication Job Wizard

Bucket Details
Create a new bucket or select an existing one.

Source Storage Platform
Management Agent
Proxy Configuration
Source Paths
File and Folder Filters
Destination
Wasabi Credentials
Bucket Details
Replication and Retention Policy
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

New Bucket
 Existing Bucket

*Name:
agent1-wniyfqbtnf
 Automatically name

< Back Next > Finish Cancel

2. If you selected **Existing Bucket**, select a bucket from the drop-down list. If the bucket does not appear in the list because the person who has the permissions to create a bucket has not yet created the bucket, click **Reload** after the bucket is created. The bucket will appear in the updated list.

If you selected **New Bucket**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the bucket.

3. Click **Next**.

The [Replication and Retention Policy](#) page appears.

Step 11: Replication and Retention Policy

Each Cloud Backup and Replication job must have a Replication and Retention policy. A Replication and Retention policy specifies:

- How often you want to scan the storage device for replication or if you want to replicate in real-time.
- Whether you want to take snapshots of the data. A **snapshot** captures the state of a file system at a point in time. There are two types of snapshots:
 - A **destination snapshot** captures an image of the data on the destination storage device immediately after replication. Destination snapshots are useful for recovering data from different periods of time. Destination snapshots track versions of the changed files and file system structure that can be used for data recovery. For more information about recovering data, see [Recovering Data](#).
 - A **source snapshot** captures an image of the data on the source storage device immediately before replication. Sources snapshots are useful for replicating open and locked files, which otherwise may not be able to be replicated. A source snapshot also ensures that the replicated data are coming from a static version of the source file system. For details about using source snapshots, see [Step 13: Source Snapshots](#).
- How long you want to retain destination snapshots.

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

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

2. If you selected **Existing Policy**, select a policy from the drop-down list, and then click **Next**. Continue with [Step 14. Miscellaneous Options](#).

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

3. Select **Enable Backup with Destination Snapshots** if you want to replicate what is on premises to the [destination storage device](#), while taking [destination snapshots](#) at specified points in times.
4. Click **Next**.

The [Replication Schedule](#) page appears.

Step 12: Replication Schedule

The **Replication Schedule** page enables you to select the frequency of the replication and when snapshots should be taken. Replication can be performed on a scheduled, batched real-time, or a continuous real-time basis.

1. Select the frequency of the replication:

- [Scheduled Scans](#) – Select this option if you want to replicate files on a scheduled basis. A scan of changes to the file system occurs on a scheduled basis, either daily or weekly, and replication of changes occurs as the scan progresses.
- [Batched Real-time](#) – Select this option if you want to continuously monitor changes to the file system but replicate changes on scheduled basis. Changes are monitored in real-time and only the latest version of changed file is replicated at scheduled times. An initial scan can be performed to establish a baseline.
- [Continuous Data Protection](#) – Select this option if you want continuously monitor changes and replicate changes in real-time. Whenever a file changes, the change is replicated in real-time.

2. Click **Next**.

The [Retention](#) page appears.

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

1. Select the **Scan at Start** checkbox if you want a baseline replication to be performed.
2. Select **Daily** or **Weekly** for the frequency of the scans:
 - Select **Daily** if you want replications performed every day. You can schedule one to four scans per day.
 - Select **Weekly** if you want to select specific days for replication. You can select one scan per day.
3. Select the day(s) and time(s) when you want the replication performed:
 - If you selected **Daily**, select the times you want the scans performed. Then, if you selected **Enable Backup with Destination Snapshots** in [Step 10](#), choose when snapshots are taken (you must take at least one snapshot). If you did not select the backup option, the **Take Destination Snapshot** options will not appear.

Create Cloud Backup and Replication Job Wizard

Replication Schedule

✘ You must select at least one replication time.

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
 - Storage Information
- Source Paths
 - File Filters
- Destination
 - Amazon S3 Credentials
 - Bucket Details
- Replication and Retention Policy
 - Replication Schedule**
 - Retention
 - Source Snapshots
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation

Scheduled Scans

Daily Weekly

Times (up to 4)

None	<input type="checkbox"/> Take Destination Snapshot
None	<input type="checkbox"/> Take Destination Snapshot
None	<input type="checkbox"/> Take Destination Snapshot
None	<input type="checkbox"/> Take Destination Snapshot

Scan at Start

Batched Real-time

Continuous Data Protection

< Back Next > Cancel

- If you selected **Weekly**, select the day(s) and time you want the replication performed. Then, if you selected **Enable Backup with Destination Snapshots** in Step 10, choose when snapshots are taken. You must take at least one snapshot. If you did not select the backup option, the **Destination Snapshot** option will not appear.

4. Click **Next**.

The [Retention](#) page appears.

If you selected **Batched Real-time** for the replication frequency:

1. Select **Scan at Start** if you want a baseline replication to be performed.
2. Select the frequency of the replications; you can schedule one to four replications per day.
3. If you selected **Enable Backup with Destination Snapshots** in [Step 10](#), choose when destination snapshots are taken (you must take at least one snapshot). The destination snapshot will be taken after the files have been replicated. If you did not select the backup option, the **Take Destination Snapshot** option will not appear.

Replication Schedule

X You must select at least one replication time.

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
 - Storage Information
- Source Paths
 - File Filters
- Destination
 - Amazon S3 Credentials
 - Bucket Details
- Replication and Retention Policy
 - Replication Schedule**
 - Retention
 - Source Snapshots
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation

Scheduled Scans

Batched Real-time

Times (up to 4)

None	<input type="checkbox"/> Take Destination Snapshot
None	<input type="checkbox"/> Take Destination Snapshot
None	<input type="checkbox"/> Take Destination Snapshot
None	<input type="checkbox"/> Take Destination Snapshot

Continuous Data Protection

< Back Next > Cancel

4. Click **Next**.

The [Retention](#) page appears.

If you selected **Continuous Data Protection** for the replication frequency:

1. Enter a value for **Processing Delay** if you want the replication to occur after a slight delay. A delay is useful to ensure that when a file or folder is created and quickly renamed, only the latest copy of the file or folder is replicated. This reduces WAN usage.
2. If you selected **Enable Backup with Destination Snapshots** in [Step 10](#), choose when snapshots are taken (you must take at least one snapshot). If you did not select the backup option, the **Take Destination Snapshot at** options will not appear.

Replication Schedule

✘ You must select at least one destination snapshot trigger time.

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
 - Storage Information
- Source Paths
 - File Filters
- Destination
 - Amazon S3 Credentials
 - Bucket Details
- Replication and Retention Policy
 - Replication Schedule**
 - Retention
 - Source Snapshots
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation

Scheduled Scans

Batched Real-time

Continuous Data Protection

Processing Delay: 5 minutes

Take Destination Snapshot at:

None

None

None

None

< Back Next > Cancel

3. Click **Next**.

The [Retention](#) page appears.

Step 13: Retention

The **Retention** page enables you to define how long you want to retain destination snapshots. You have the option to retain destination snapshots on a daily, weekly, monthly, and yearly basis. If you did not select the **Enable Backup with Destination Snapshots** in Step 10, the **Retention** page will not appear.

1. Select the **Purge all versions between snapshots** checkbox if you do not want to indefinitely retain all versions.
2. Select the retention options. The options vary according to the replication schedule you selected.

3. Click **Next**.

The [Source Snapshots](#) page appears.

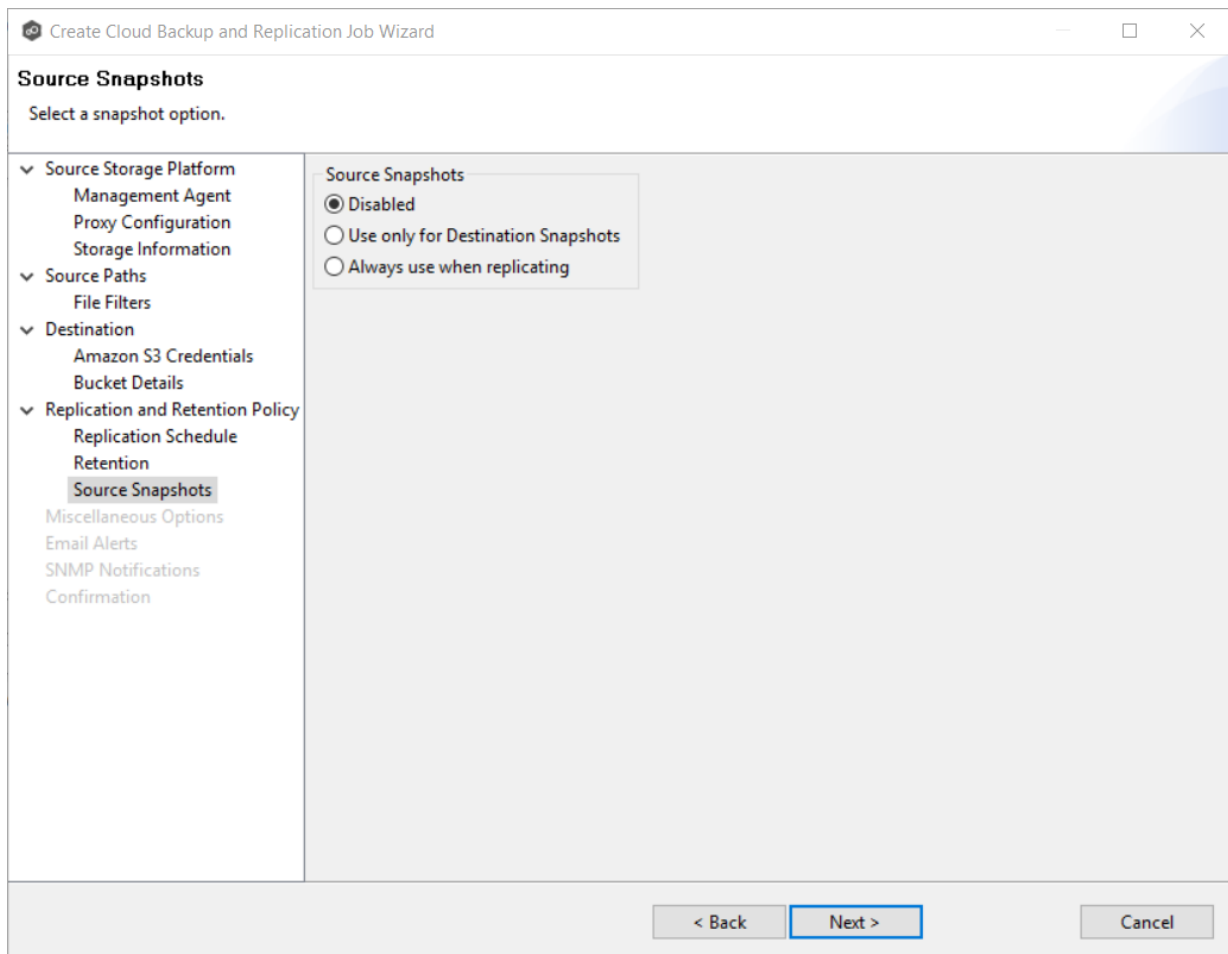
Step 14. Source Snapshots

The **Source Snapshots** page enables you to choose whether to take snapshots of the source storage before the items are replicated. A [source snapshot](#) is a read-only point-in-time version of the volume. A source snapshot allows the creation of consistent backups of a volume, ensuring that the contents do not change and are not locked while the backup is being made. It can be used to provide a consistent state of a managed file, e.g., pst files, and help with errors accessing files that are currently open.

1. Select a source snapshot option:

- Select the **Disabled** option if you do not want to take source snapshots.

- Select the **Use only for Destination Snapshots** option when you want the source snapshot to be stored on the destination storage as the destination snapshot rather than an actual destination snapshot. To use this option, you must have selected the **Enable Backup with Destination Snapshot** in Step 10.
- Select **Always use when replicating** when you want to replicate always using source snapshots.



2. Click **Next**.

The [Miscellaneous Options](#) page appears.

Step 15: Miscellaneous Options

The **Miscellaneous Options** page displays various options; the options available depend on the destination storage platform selected.

1. Select the options to apply to this job.

Option	Description
NTFS Permissions	<p>If you want NTFS permissions metadata included in the replication, select the elements to include:</p> <ul style="list-style-type: none"> • Owner – The NTFS Creator-Owner who owns the object (which is, by default, whoever 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. <p>See File Metadata Synchronization for more information about NTFS permissions metadata.</p>

Option	Description
Storage Tier/Class	<p>Only available with Azure Blob Storage with either a General Purpose v2 (GPv2) account or Blob Storage Account.</p> <p>Select a storage tier. If you do not select a tier, it will default to the tier you configured on your Azure Storage account.</p> <p>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:</p> <ul style="list-style-type: none"> • 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. <p>To read data in archive storage, Cloud Backup and Replication must first change the tier of the blob to hot or cool. This process is known as rehydration and can take up to 15 hours to complete.</p> <p>Rehydrated data remains in hot or cool storage for a specified number of days before Cloud Backup and Replication automatically returns it to archive storage.</p>
Rehydrated Data Availability (Days)	<p>Only available with Azure Blob Storage with either a General Purpose v2 (GPv2) account or Blob Storage Account.</p> <p>Rehydrated data are 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.</p>

2. Click **Next**.

The [Email Alerts](#) page appears.

Step 16: Email Alerts

This step is optional.

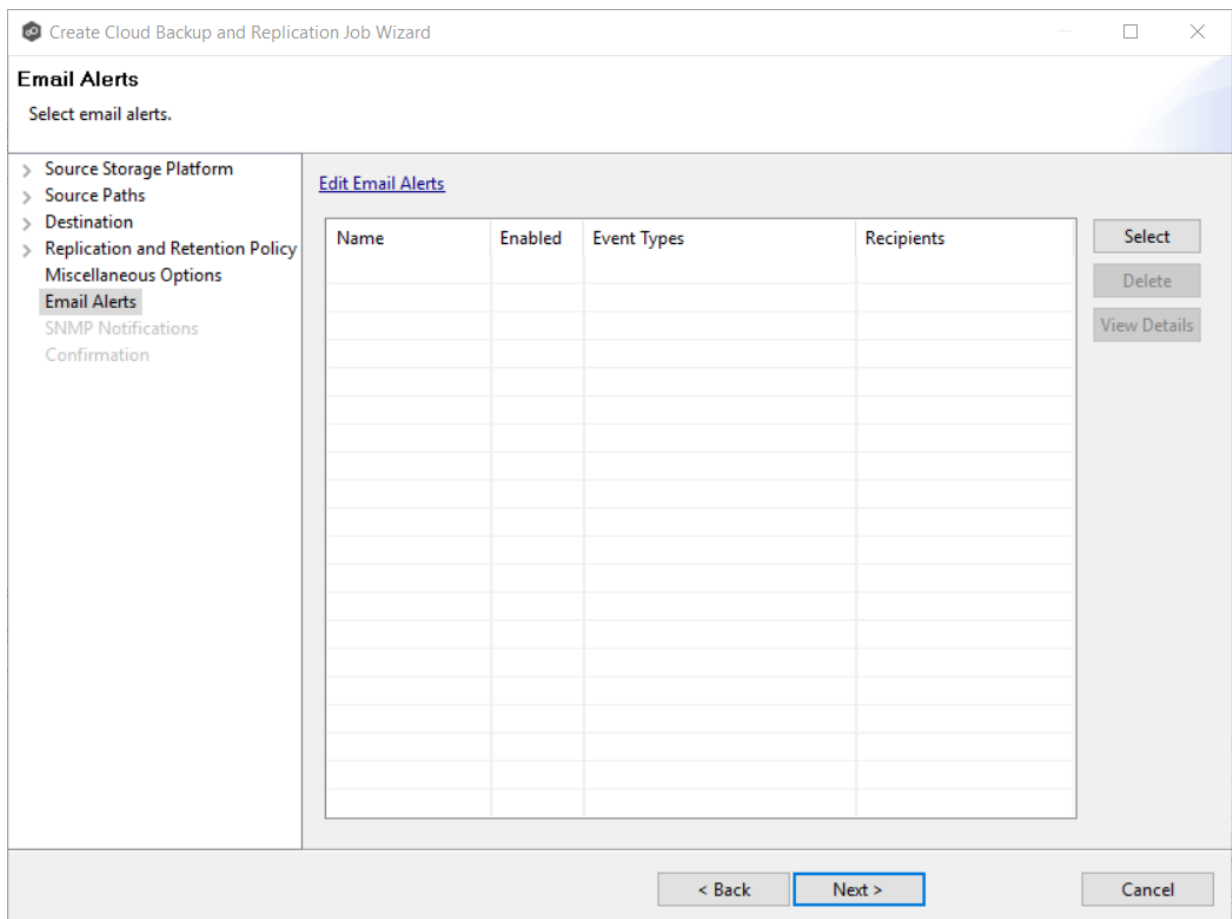
An [email alert](#) notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in [Preferences](#) and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply.

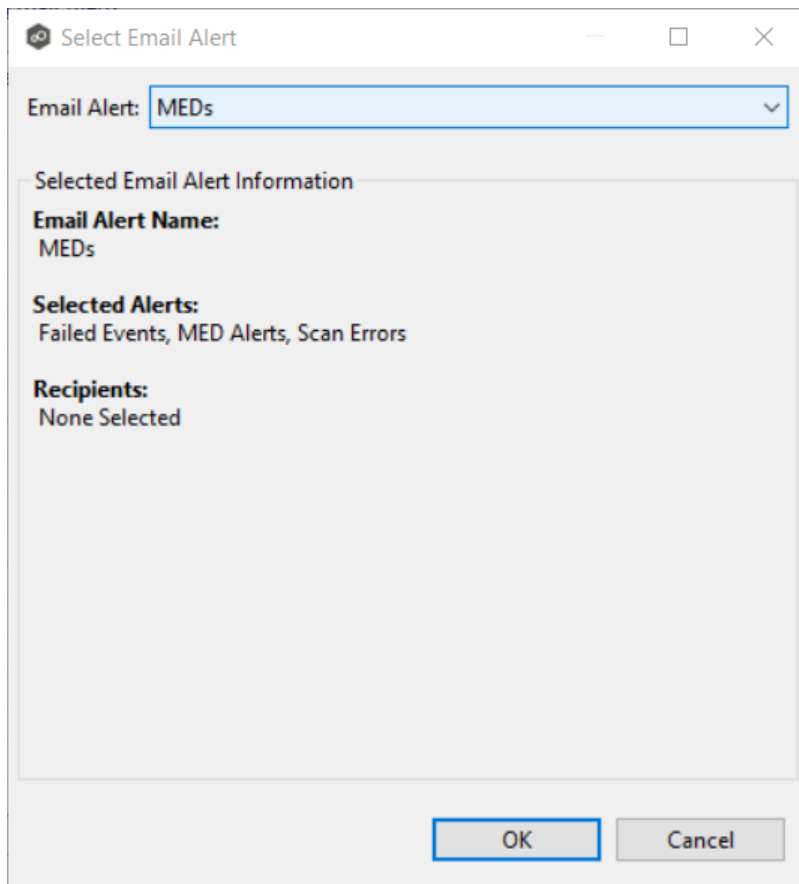
To create a new alert, see [Email Alerts](#) in the [Preferences](#) section.

To apply an existing email alert to the job:

1. Click the **Select** button.



The **Select Email Alert** dialog appears.



2. Select an alert from the **Email Alert** drop-down list, and then click **OK**.

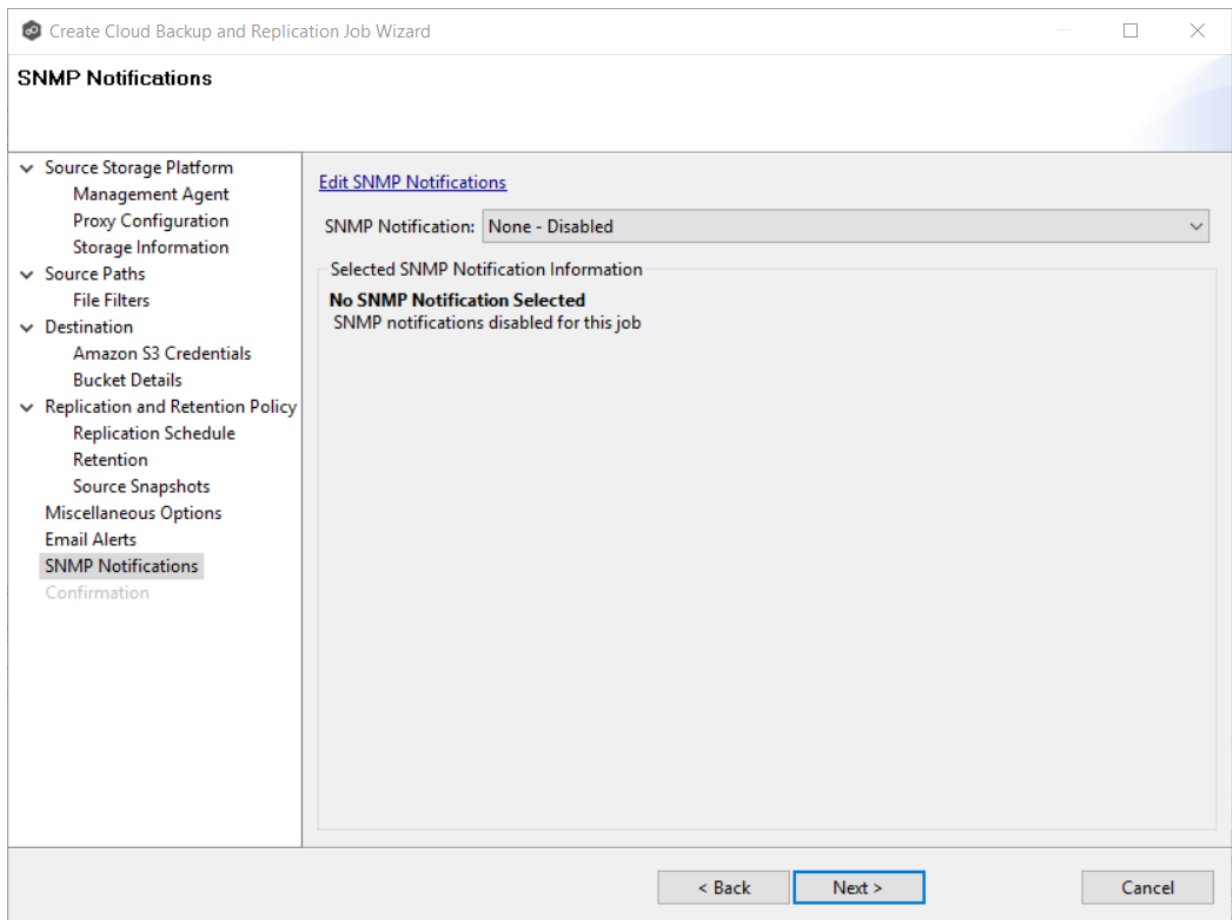
The alert is listed in the **Email Alerts** page.

Peer Software recommends that you create SNMP notifications in advance. However, from this wizard page, you can select an existing SNMP notification to apply to the job or create new SNMP notifications.

To apply an existing SNMP notification to the job or disable notifications:

1. Select an SNMP notification from the drop-down list.

To disable, select **None - Disabled**.



2. Click **Next**.

The [Confirmation](#) page appears.

Step 18: Confirmation

The **Confirmation** page displays the job configuration.

1. Review the job configuration.
2. If you need to modify the job configuration after reviewing it, click **Back** until you reach the appropriate page and make your changes.

Note: You cannot change the job name.

Create Cloud Backup and Replication Job Wizard

Confirmation
Review your job configuration.

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
 - Storage Information
- Source Paths
 - File Filters
- Destination
 - NetApp StorageGRID Credential
 - Bucket Details
 - Replication and Retention Policy
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation**

Source Storage: NetApp ONTAP | Clustered Data ONTAP

Management Agent: DGAgent2

Source Paths: One or more specific paths

Source Items:
Volume: SVM9x1_mm1
Destination: -1
Item(s) to Include:
..\
Destination: NetApp StorageGRID

Replication and Retention Policy: 3x Daily
Replication Schedule: Destination Snapshots

Replicate every day

Times:
- 05:00 (Take destination snapshot)
- 15:00 (Take destination snapshot)
- 20:00 (Take destination snapshot)

Retention Configuration (3x Daily):
Purge all versions between destination snapshots: true

Daily Retention
Keep destination snapshots taken at:
- 05:00
- 15:00
- 20:00
For: 30 day(s)

Weekly Retention
Keep destination snapshots taken on:
- Monday
Taken at: 05:00
For: 52 weeks(s)

Monthly Retention
Keep destination snapshots taken on:
First
Tuesday
Taken at: 15:00
For: 60 month(s)

Yearly Retention
Keep destination snapshots taken on:
First
- Wednesday
in month(s):
- January
Taken at: 20:00
For: 10 year(s)

Source Snapshots:
Disabled

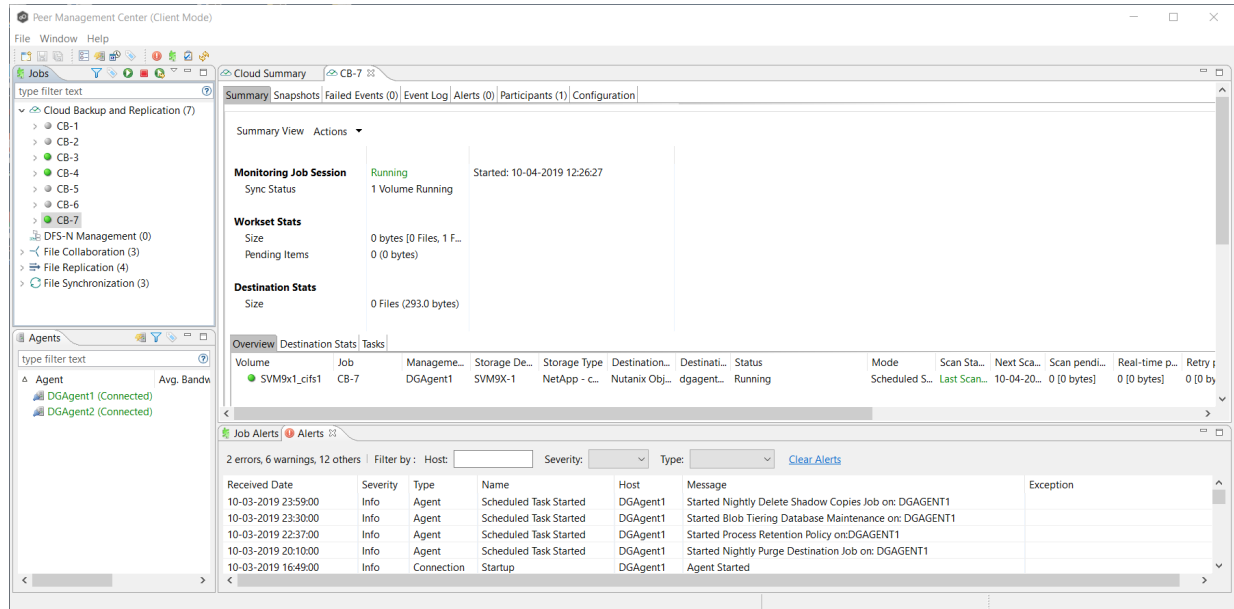
Start job after creation

< Back Next > Finish Cancel

3. Select the **Start job after creation** checkbox if you want the job to start immediately after clicking **Finish**.

4. Click **Finish**.

The **Summary** tab in the **Cloud Backup and Replication Job** runtime view is displayed.



Running a Cloud Backup and Replication Job

This section describes:

- [Starting a Cloud Backup and Replication Job](#)
- [Stopping a Cloud Backup and Replication Job](#)

Starting a Cloud Backup and Replication Job


When running a Cloud Backup and Replication job for the first time, you must manually start it. After the initial run, a job will automatically start, even when the Peer Management Center server is rebooted.

Note: You cannot run two jobs concurrently on the same volume if the [watch sets](#) contain an overlapping set of files and folders.

To manually start a job:


1. Choose one of these options:

- Right-click the job name in the **Jobs** view.
- Open a job and then click the **Start/Stop** button to the left of the **Status** field in the bottom left corner of the job's view (shown below).



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

After the job initialization has completed, the job will run. Once the job starts, the icon next to the job name in the **Jobs** view changes from gray to green.



Stopping a Cloud Backup and Replication Job

You can stop a Cloud Backup and Replication job at any time.

To stop a Cloud Backup and Replication job:

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

Or open the job and then click the **Start/Stop** button to the left of the **Status** field in the bottom left corner of the job's runtime view (shown below)

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

The icon next to the job name in the **Jobs** view changes from green to red.

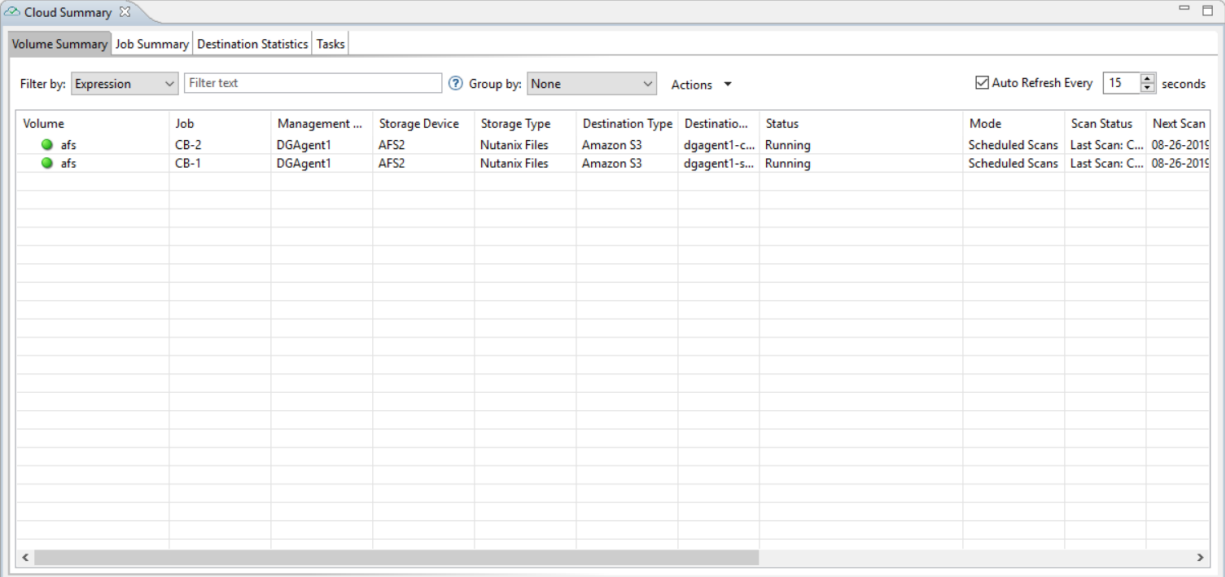
Monitoring Cloud Backup and Replication Jobs

Monitoring your Cloud Backup and Replication jobs is an important aspect of successfully replicating to the cloud. Monitoring involves checking the execution of a running job, checking the status of a job, reviewing performance statistics, making sure snapshots are created correctly, identifying problems such as a server outage, seeing how much data has been uploaded, and so forth. Cloud Backup and Replication provides several views to help you monitor the health and performance of your Cloud Backup and Replication jobs.

Many of the views are customizable tables. You can sort the columns in the view, filter by columns, add and subtract columns from the default display, and so forth.


To display a view:

- Double-click **Cloud Backup and Replication** in the **Jobs** view to display the summary view for all Cloud Backup and Replication jobs. The **Volume Summary** tab of the **Cloud Summary** view is displayed.



Volume	Job	Management ...	Storage Device	Storage Type	Destination Type	Destinatio...	Status	Mode	Scan Status	Next Scan
afs	CB-2	DGAgent1	AFS2	Nutanix Files	Amazon S3	dgagent1-c...	Running	Scheduled Scans	Last Scan: C...	08-26-2015
afs	CB-1	DGAgent1	AFS2	Nutanix Files	Amazon S3	dgagent1-s...	Running	Scheduled Scans	Last Scan: C...	08-26-2015


- Double-click the name of a Cloud Backup and Replication job in the **Jobs** view to display the runtime view associated with that job. The **Summary** tab of the runtime view is displayed.



Deleting a Cloud Backup and Replication Job


To delete a Cloud Backup and Replication job:

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



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

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



3. Click **Yes** or **No**.

If you click **Yes**, the data associated with this job will be deleted as part of a nightly clean-up process in addition to the job itself. If you click **No**, the data will not be deleted but the job will be deleted.

Recovering Data

When you need to recover data from the cloud to on-premises, you can use the **Data Recovery** wizard. To restore data, you must have an existing Cloud Backup and Replication job that has been replicating that data.

Note: You can recover data from a running job. However, if you plan to restore the data to the original location, you should stop the job first.

To recover data:

1. Open Peer Management Center.
2. In the **Jobs** view, identify the Cloud Backup and Replication job that replicated the data you want to restore.
3. Right-click the job name, and then select **Recover Volume/File(s)** from the menu.

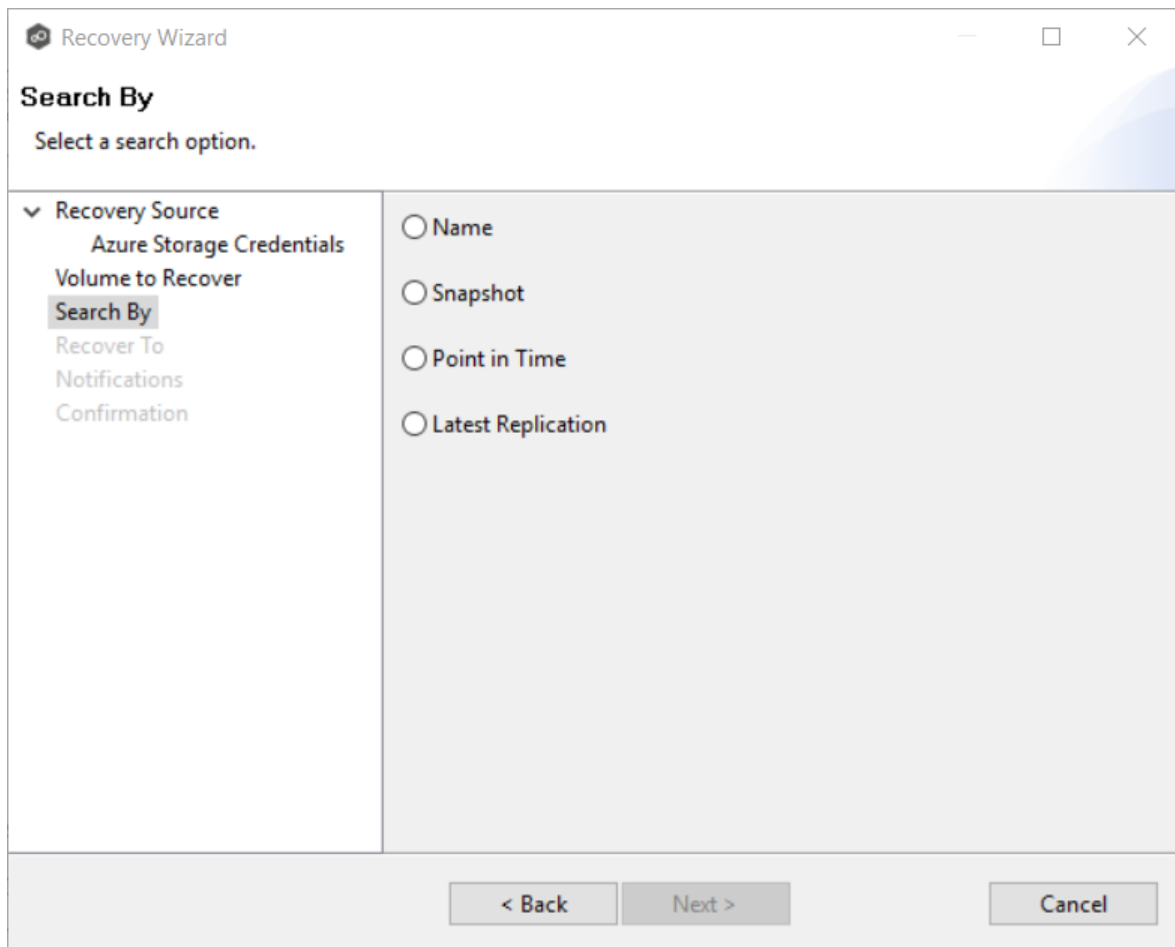
The **Recovery Wizard** opens and displays the **Volume to Recover** page. The **Storage Device** field on the page is a read-only field that displays the name of the source storage device.

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

The screenshot shows a window titled "Recovery Wizard" with the subtitle "Volume to Recover". Below the subtitle is the instruction "Select the storage device and volume that the data originated from." The window is divided into two main sections. On the left is a navigation pane with a tree view containing the following items: "Recovery Source" (expanded), "Azure Storage Credentials", "Volume to Recover" (highlighted), "Search By", "Recover To", "Notifications", and "Confirmation". On the right is the main content area with two dropdown menus: "Storage Device:" with the value "SVM9X-1" and "Volume:" with the value "SVM9x1_highcount1_01". At the bottom of the window are three buttons: "< Back", "Next >" (which is highlighted with a blue border), and "Cancel".

5. Click **Next**.

The **Search By** page is displayed.



6. Select one of the search options.

- [Name](#)
- [Snapshot](#)
- [Point in Time](#)
- [Latest Replication](#)

7. Click **Next** and continue with [Recovery Options](#).

The search pages vary according to the search option you selected.

Search Options

1. The search options are:

- [Name](#)
- [Snapshot](#)
- [Point in Time](#)
- [Latest Replication](#)

Use the **Search by Name** option if you know any part of the name of a file or folder but don't know which folder contained it on the original volume on premises.

To search by name:

1. Enter a search string in the **Name** field.

The search string can be a full or partial name and can include wildcards. If you do not enter a search string, all files and folders will be listed in the search results.

Recovery Wizard

Name
Enter a search string.

> Recovery Source
Volume to Recover
v Search By
 Name
 File/Folder Versions
 Recover To
 Notifications
 Confirmation

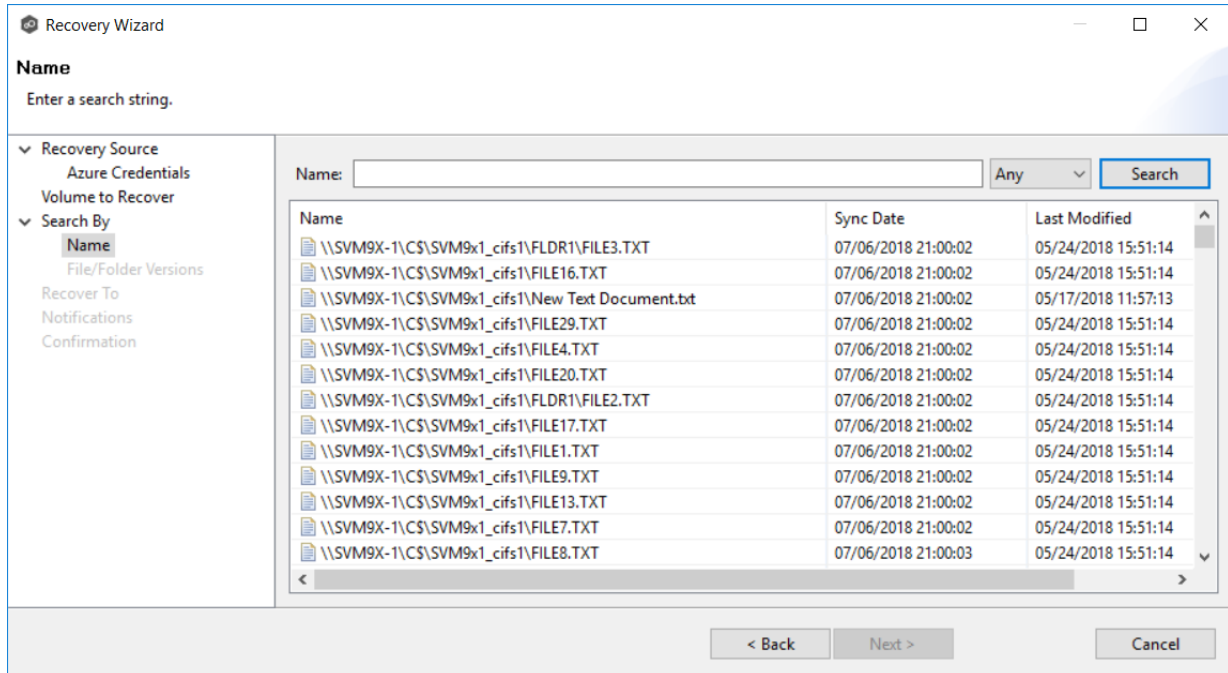
Name: Any

Name	Sync Date	Last Modified	Size

< Back Next > Cancel

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

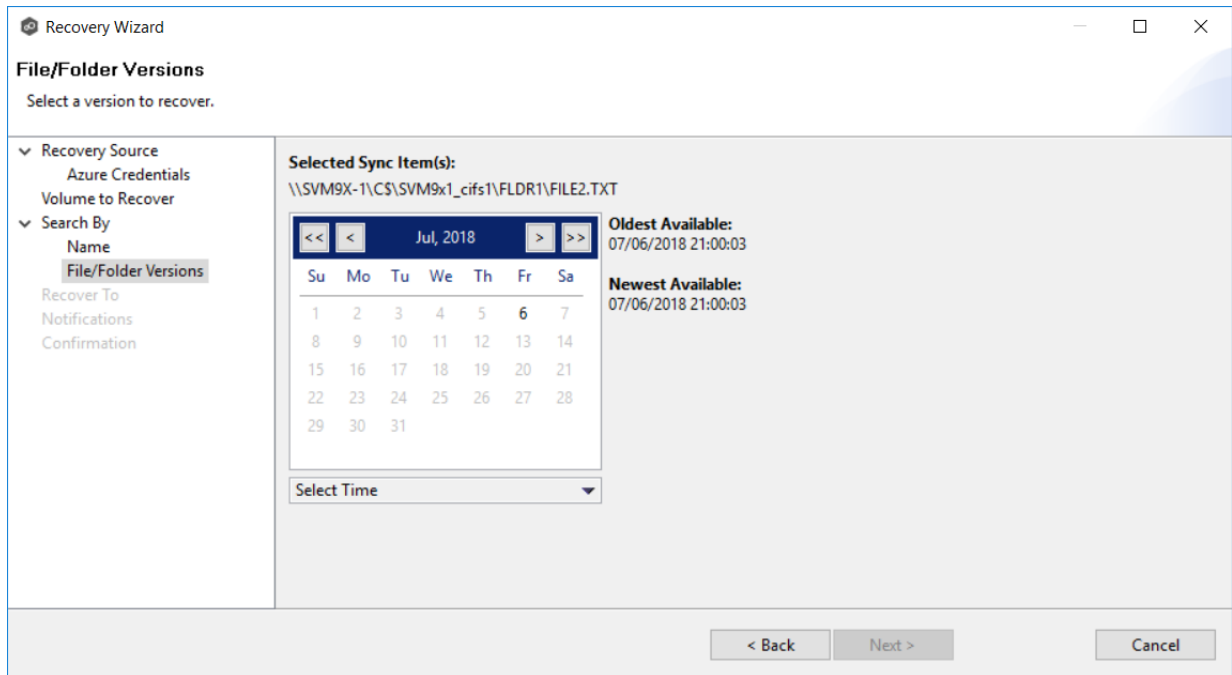
A list of matching files and/or folders appears. The **Sync Date** column shows the date the file was replicated; the **Last Modified Date** column shows the last known date and time that the file was changed on premises.



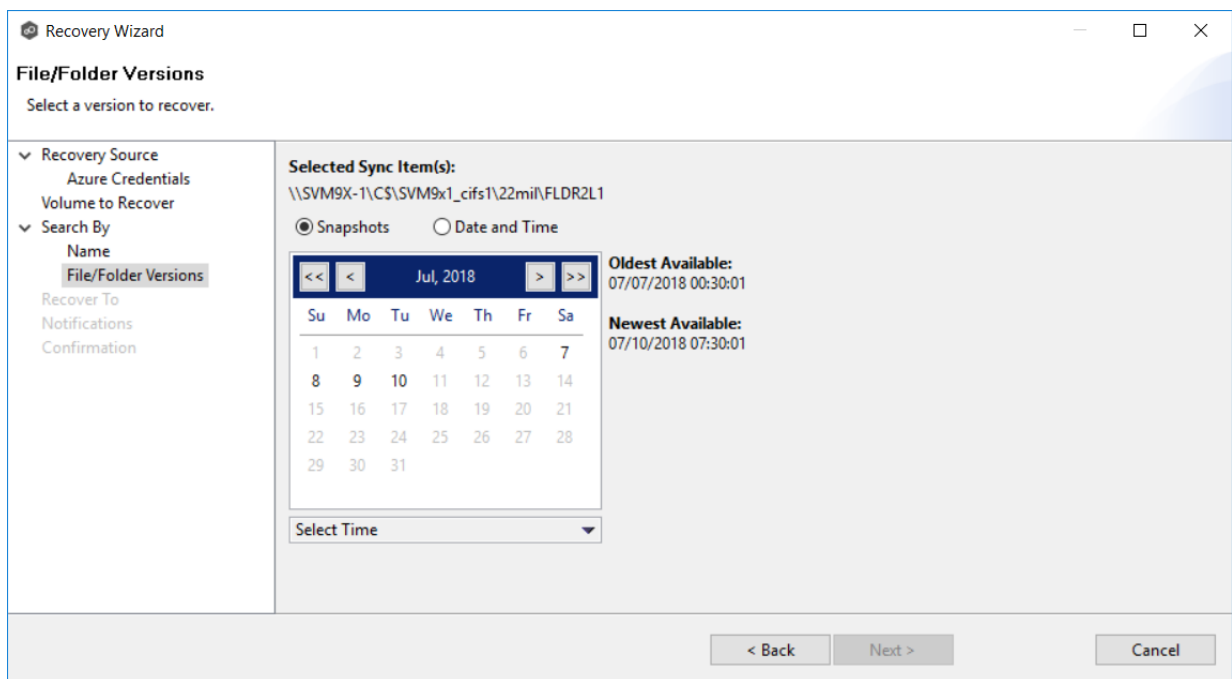
4. Select the file or folder to recover.
5. Click **Next**.

The **File/Folder Versions** page appears. Your options will vary, depending on whether you are recovering a file or folder.

6. If you selected a file to recover, all available versions of that file are presented below the calendar. Select the time of the desired version, and then click elsewhere on the page.



If you selected a folder to recover, you have two options. You can recover the contents of the folder based on a snapshot that was previously taken, or you can recover the contents of the folder as it existed at a specific point in time. Select one of the options, select a time, and then click elsewhere on the page.

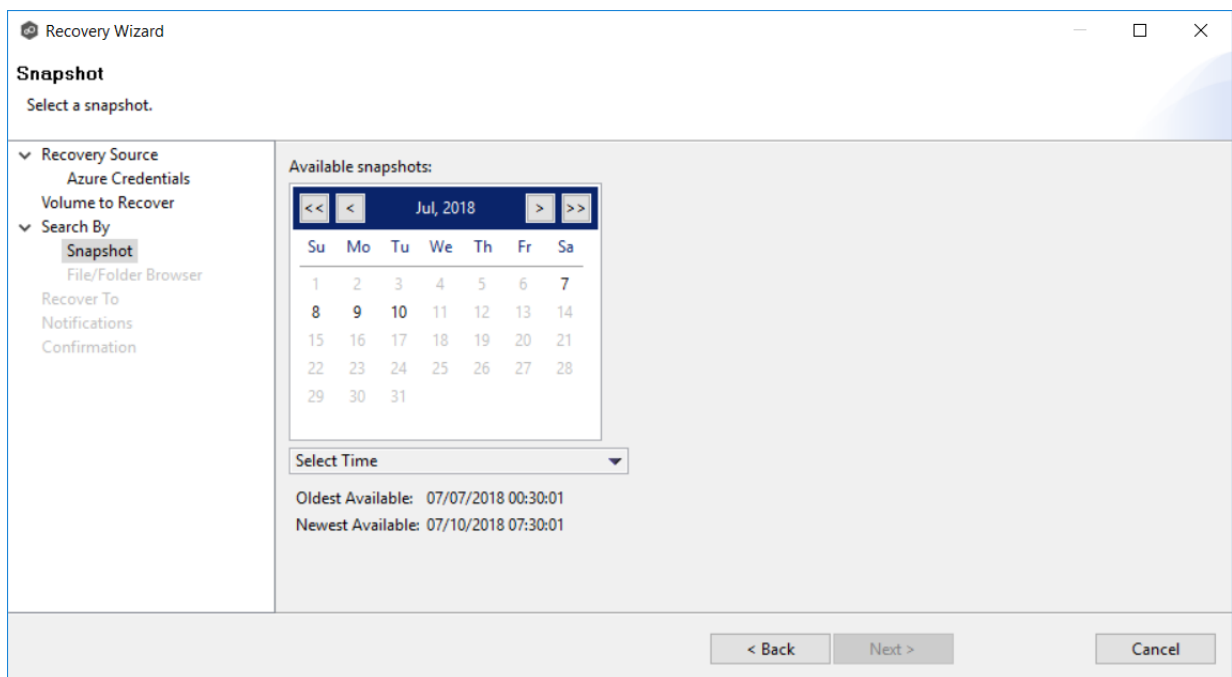


7. Click **Next** and continue with [Recovery Options](#).

Use the **Search by Snapshot** option if you want to recover data by browsing a previously taken destination snapshot. All available snapshots will be represented in the calendar widget below.

To search by snapshot:


1. Select the date of the snapshot.



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

3. Click **Next**.

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

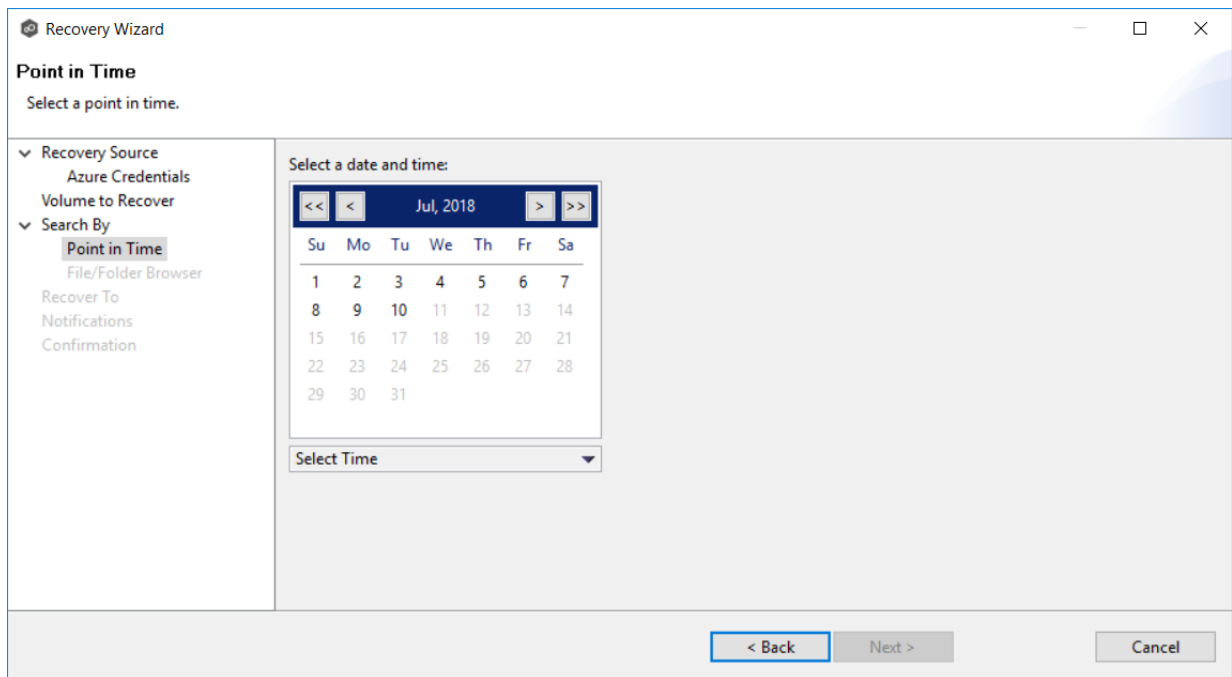


4. Select the file or folder to restore. If no snapshots are available, click **Back** and select a different search option.
5. Click **Next** and continue with [Recovery Options](#).

Use the **Search by Point in Time** option if you want to restore 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 on the page.
3. Click **Next**.

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



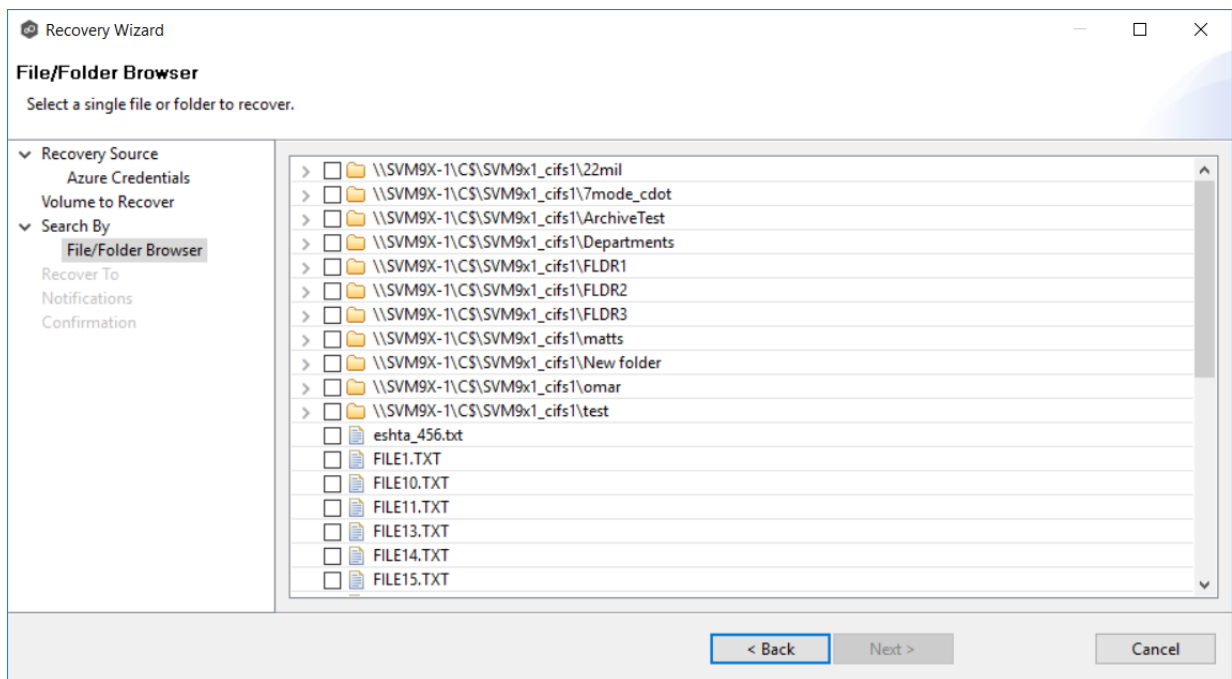
4. Select the file or folder to restore.

5. Click **Next** and continue with [Recovery Options](#).

Use the **Search by Latest Replication** option if you want to restore from the latest replication. For example, you may want to restore data from the last time that replication occurred rather than a snapshot or a point in time. This option is very useful if you selected [Continuous Data Protection](#), where replication is performed on an on-going basis.

To search by latest replication:

1. Select the file or folder to restore.



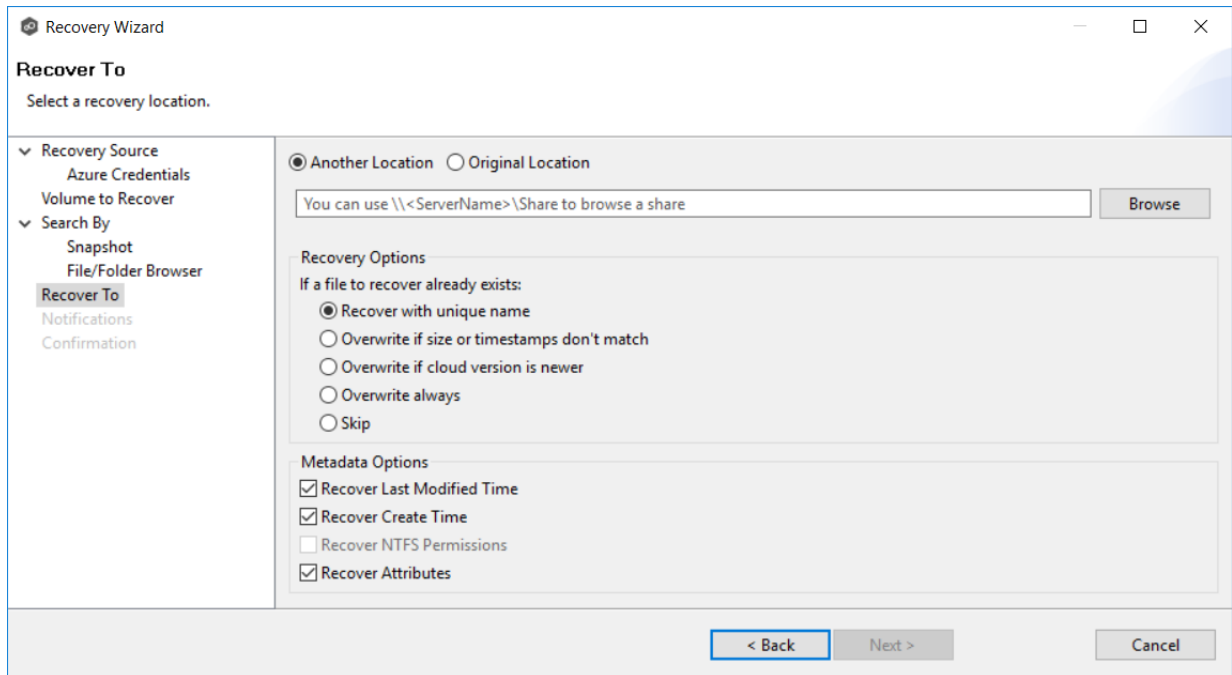
2. Click **Next** and continue with [Recovery Options](#).

Recovery Options

After you select the data to recover, the **Recover To** page appears.

1. Select the recovery location. You have two options:

- **Another Location** - Enter the UNC path to a location on another storage device.
- **Original Location** - Browse to a location on the device hosting the management agent. However, we recommend not restoring directly to the original location, especially if the job is currently running. If the version that is restored is older than the latest version in the destination storage, the restored version will not be backed up until the next scan.



2. Select the recovery options for when the file to recover already exists in the recovery location:

Recovery Option	Select this option if you want to:
Recover with unique name	Ensure that the existing file is not overwritten with the cloud version.
Overwrite if sizes or timestamps don't match	Overwrite the existing file with the cloud version if the sizes or timestamps the existing file do not match the cloud version.
Overwrite if cloud version is newer	Overwrite the existing file if the cloud version has a more recent modification date.

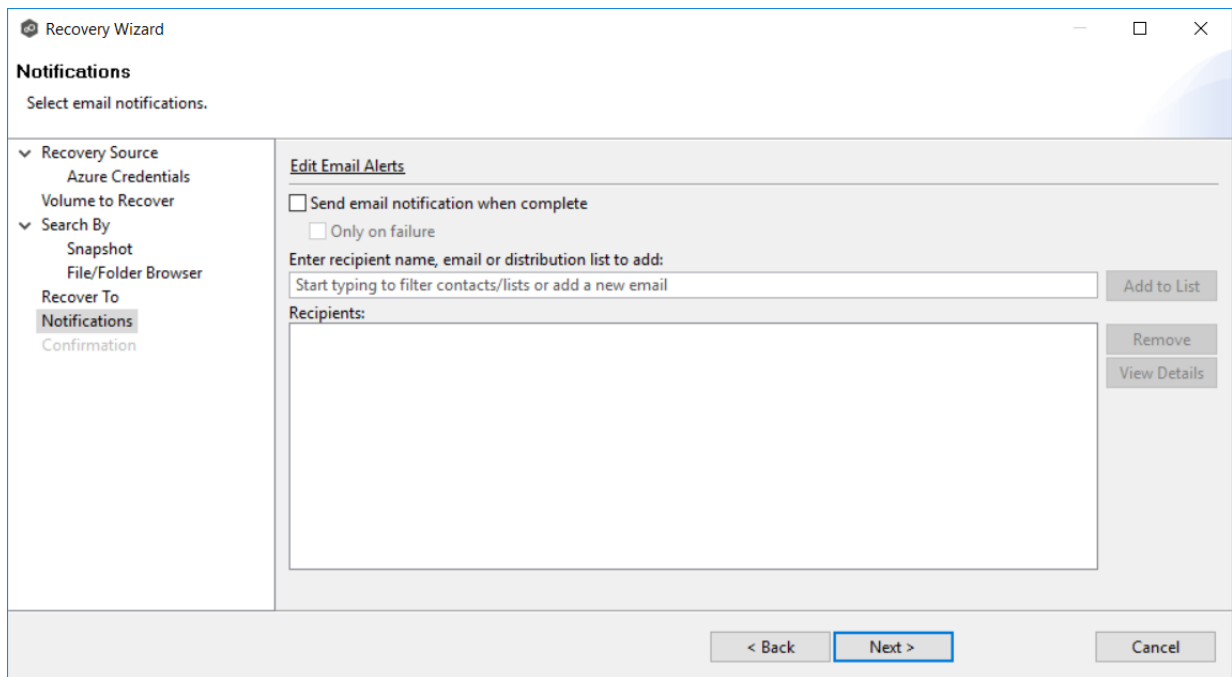
Overwrite always	Always overwrite the existing file with the cloud version.
Skip	Skip recovering a file if the file already exists.

3. Select the recovery metadata options:

Metadata Option	Select this option if you want to:
Recover Last Modified Time	Set the last modification time of a recovered file to match the last modification time stored at upload rather than the time at which it was recovered.
Recover Create Time	Set the creation time of a recovered file to match the creation time stored at upload rather than the time at which it was recovered.
Recover NTFS Permissions	Set the NTFS permissions of any recovered files and folders to match the original permissions when those files and folders were uploaded.
Recover	Set the attributes of any recovered files and folders to match the original attributes when those files and folders were uploaded.

4. (Optional) Click the **Review** button to see your selections.
5. Click **Next**.

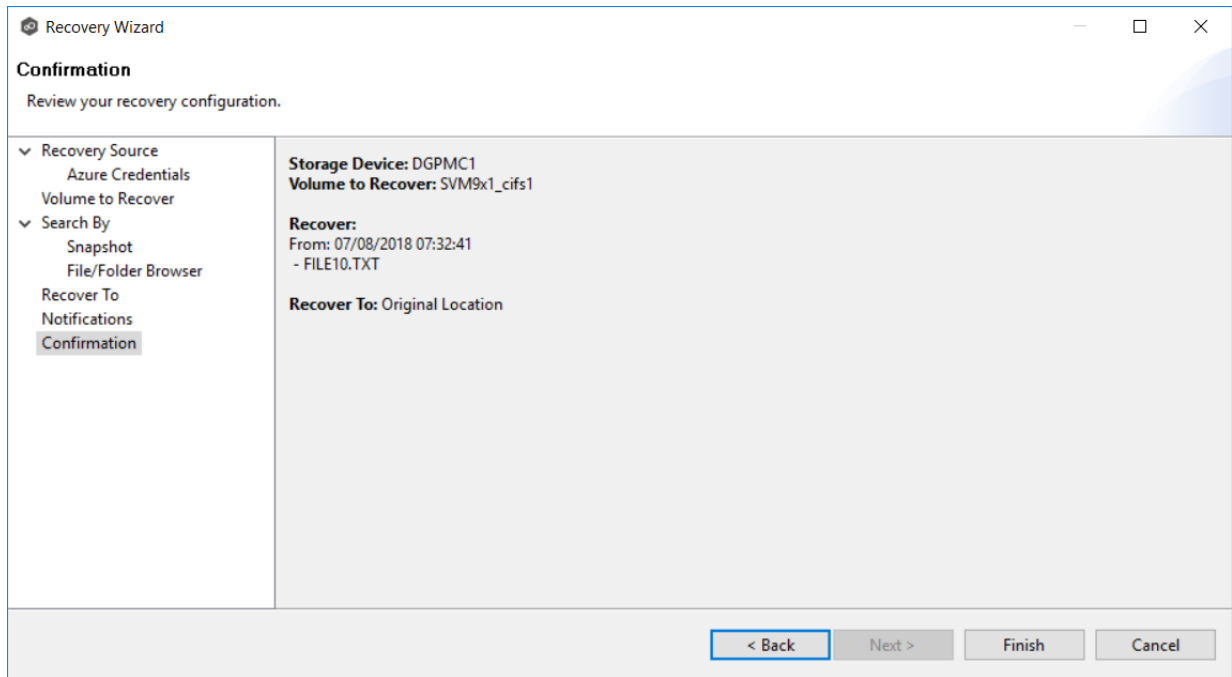
The **Notifications** page appears.



6. (Optional) Select the **Send email notification when complete** checkbox if you want notifications sent when the recovery process is complete. Select **Only on failure** if you want notifications sent only if the recovery does not successfully complete.
7. If sending notifications, enter recipients and add them to the list.
8. Click **Next**.

The **Confirmation** page is displayed.

9. Review your recovery settings.



10. Click **Finish**.

DFS-N Management Jobs

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

A [DFS namespace](#) enables you to group shared folders that are located on different servers into one or more logically structured namespaces. Each namespace appears to users as a single shared folder with a series of subfolders. However, the underlying structure of the namespace can consist of numerous file shares that are located on different servers and in multiple sites. See [DFS Namespaces](#) for more information about the benefits of using DFS namespaces in PeerGFS.

PeerGFS enables you to create a namespace and manage various activities related to it, such as creating namespace folders, adding folder targets, and linking the namespace to a File Collaboration or File Synchronization job. You could manage DFS namespace using Microsoft tools; however, you can manage DFS namespaces through a dedicated job type in Peer Management Center, the DFS-N Management job.

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

- [Creating a DFS-N Management Job](#)
- [Running a DFS-N Management Job](#)
- [Managing DFS Namespaces](#)
 - [Adding a Namespace Server](#)
 - [Adding a Namespace Folder](#)
 - [Adding a Namespace Folder Target](#)
- [Importing an Existing Namespace](#)
- [Linking a DFS Namespace to File Collaboration and File Synchronization Jobs](#)

Creating a DFS-N Management Job

The **Create Job** Wizard walks you through the process of creating a DFS-N Management job. The process consists of the following steps:

[Step 1: Job Type](#)

[Step 2: Management Agent](#)

[Step 3: Agent Verification](#)

[Step 4: Namespace Name](#)

[Step 5: Namespace Servers](#)

[Step 6: Namespace Settings](#)

[Step 7: Folders](#)

[Step 8: Email Alerts](#)

[Step 9: SNMP Notifications](#)

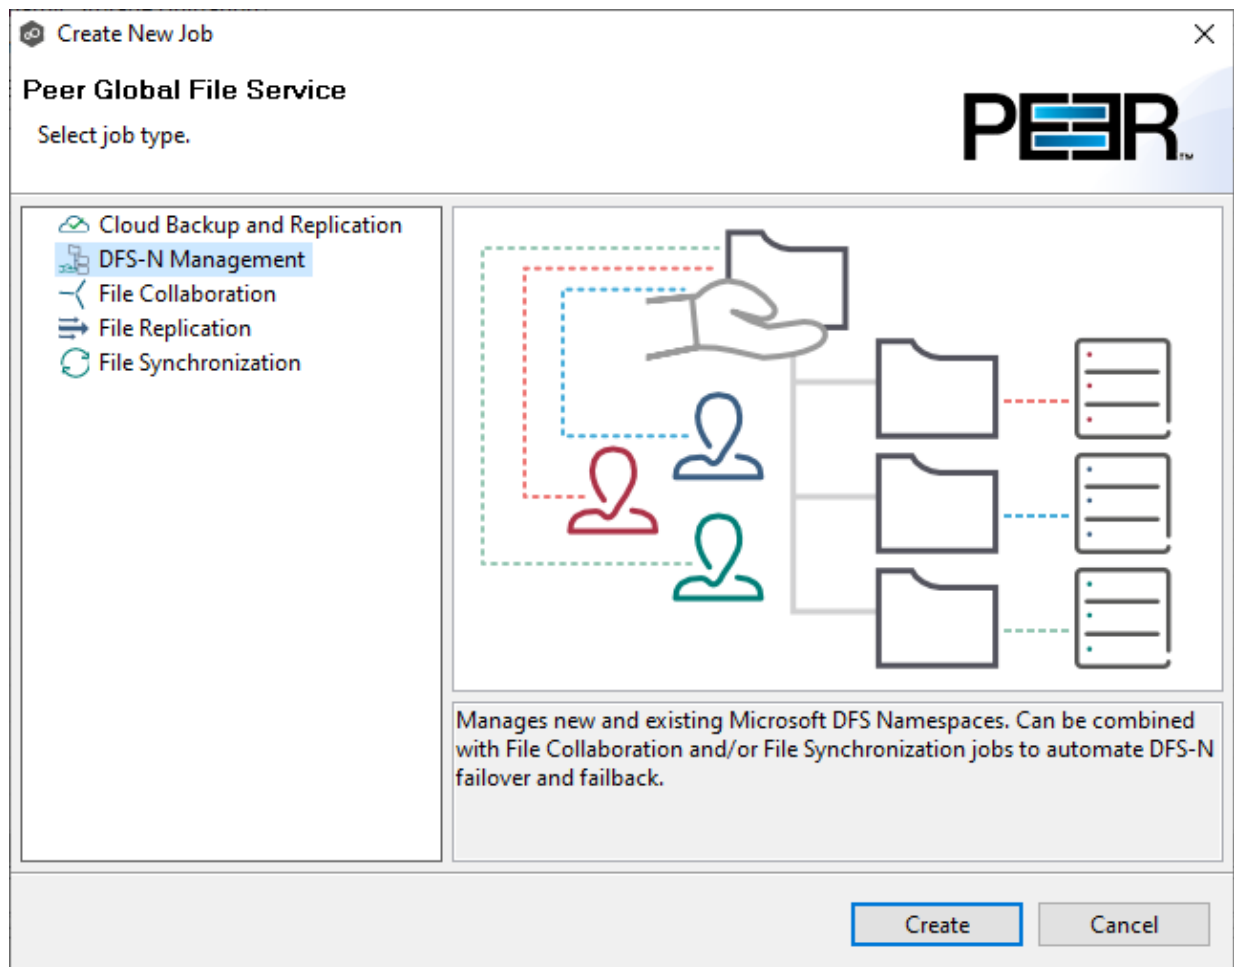
[Step 10: Review](#)

[Step 11: Results](#)**Step 1: Job Type**


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

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

3. Click **DFS-N Management**, and then click **Create**.



The following dialog appears.



4. If you have an existing namespace you want to import, click **Import Namespace**, and then follow the [steps for importing an existing namespace](#).

Otherwise, click **Create Namespace**.

The [Management Agent](#) page appears.

Step 2: Management Agent

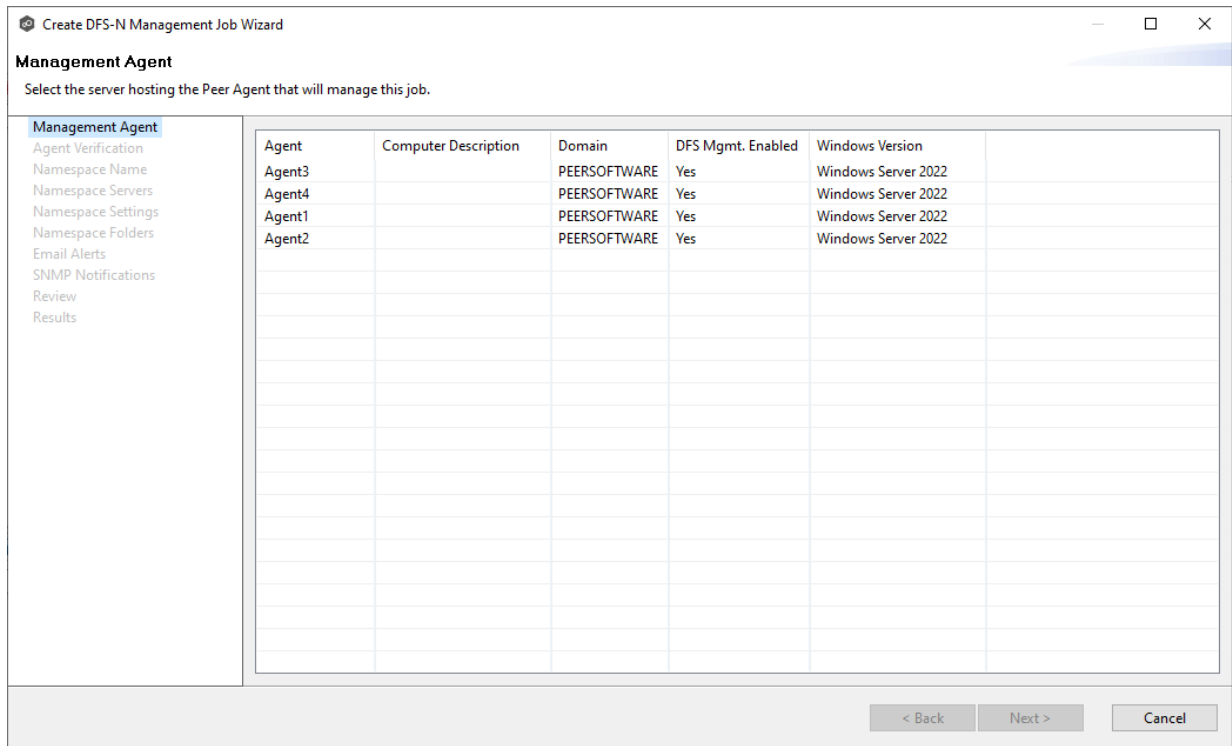
On the **Management Agent** page, you select a Management Agent for this DFS-N Management job from the list of servers that have a Peer Agent installed.

Recommendation: Select an Agent that will be dedicated to managing DFS-N Management jobs. This enables the Agent to continue managing the namespace even if other Agent servers go down. If you use a dedicated Agent for DFS-N Management jobs, this Agent will not count against the number of licensed servers.

To reduce the number of Windows servers in your environment, you can [install an Agent](#) that runs on the same server as Peer Management Center and use this Agent to manage DFS-N Management jobs. This Agent will also not count against the number of licensed servers.

1. Select an Agent that is in the domain of the DFS namespace of where you want to create the new DFS namespace.

Note: If you select an Agent that has **No** in the **DFS Mgmt. Enabled** column, the Microsoft DFS PowerShell Management toolkit will be installed in [Step 3: Agent Verification](#).



2. Click **Next**.


The [Agent Verification](#) page appears.

Step 3: Agent Verification

The purpose of the **Agent Verification** page is to verify that the environment of the selected Management Agent is set up properly to communicate with DFS Namespaces. For example, the Microsoft DFS PowerShell Management toolkit must be installed on the same system as the Management Agent and configured correctly. If the toolkit is not already installed, you will be able to install it during the verification process.

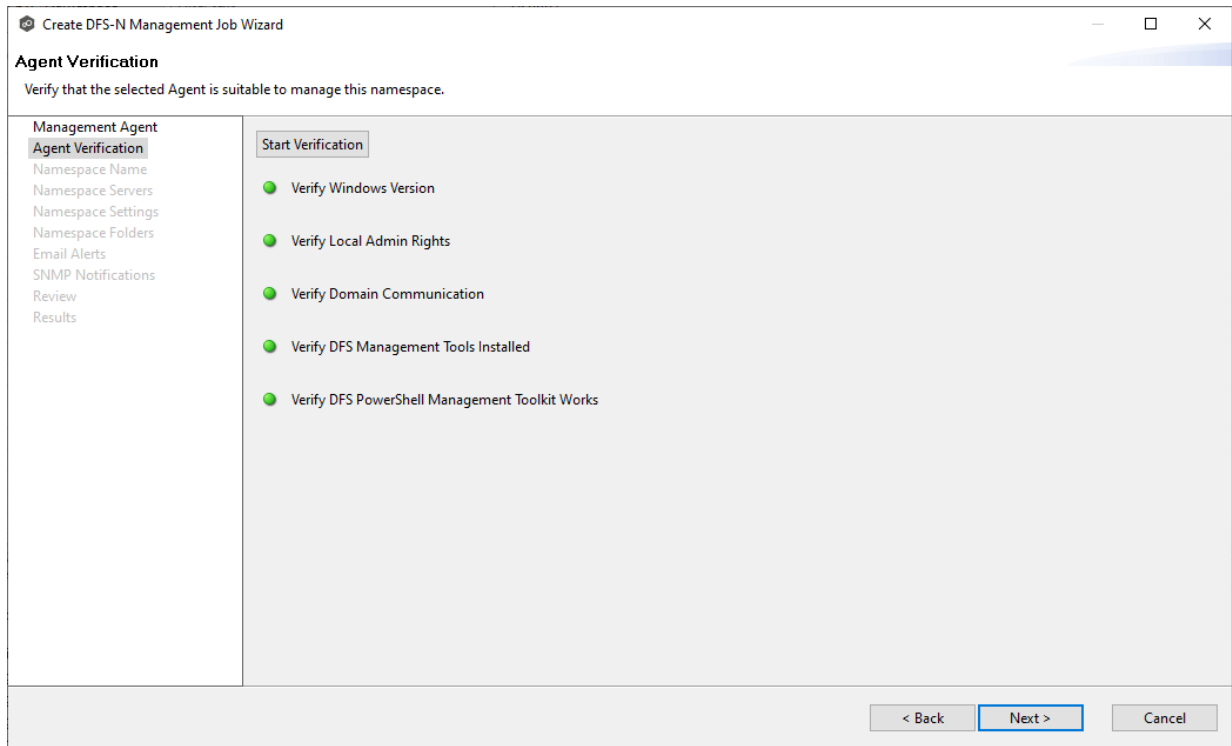
Note: The verification process does not include checking whether DFS Services is running because DFS Services doesn't have to run on the Agent server itself; it typically runs on a domain controller.

1. Click **Start Verification**.



2. If the DFS PowerShell Management toolkit is not installed, click the **Install** button that will appear next to **Verify DFS PowerShell Management Toolkit Installed**.

After the toolkit is installed, the verification continues. A green dot signifies that the verification of that element was successful.



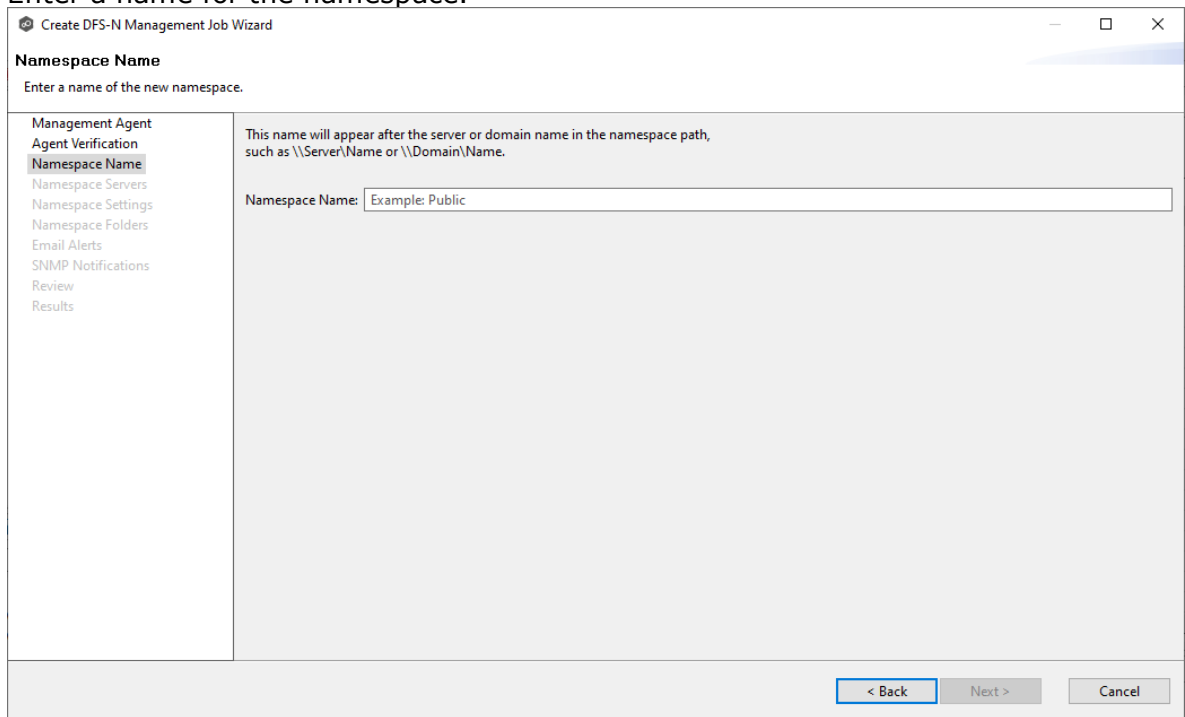
3. After the verification has successfully completed, click **Next**.

The [Namespace Name](#) page appears.

Step 4: Namespace Name

On the **Namespace Name** page, you enter a name for the new namespace. The name of the namespace will also be the name of this DFS-N Management job.

1. Enter a name for the namespace.



The screenshot shows a wizard window titled "Create DFS-N Management Job Wizard". The current step is "Namespace Name", with the instruction "Enter a name of the new namespace." A left-hand pane lists the wizard steps: Management Agent, Agent Verification, Namespace Name (highlighted), Namespace Servers, Namespace Settings, Namespace Folders, Email Alerts, SNMP Notifications, Review, and Results. The main area contains the text: "This name will appear after the server or domain name in the namespace path, such as \\Server\Name or \\Domain\Name." Below this is a text input field labeled "Namespace Name:" with the example text "Example: Public". At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

2. Click **Next**.

The [Namespace Servers](#) page appears.

Step 5: Namespace Servers

On the **Namespace Servers** page, you select one or more servers to host the namespace. A [namespace server](#) must be a member server or domain controller in the domain in which the namespace is configured. The Microsoft DFS Namespaces service must also be running on the namespace server.

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

Create 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.

Management Agent
Agent Verification
Namespace Name
Namespace Servers
Namespace Settings
Namespace Folders
Email Alerts
SNMP Notifications
Review
Results

Enter the fully qualified domain name of a server running the DFS Namespaces service.

Server Name:

Add

Servers:

Delete

< Back Next > Cancel

The server path is listed in the **Servers** area below.

Create 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.

Management Agent
Agent Verification
Namespace Name
Namespace Servers
Namespace Settings
Namespace Folders
Email Alerts
SNMP Notifications
Review
Results

Enter the fully qualified domain name of a server running the DFS Namespaces service.

Server Name:

Add

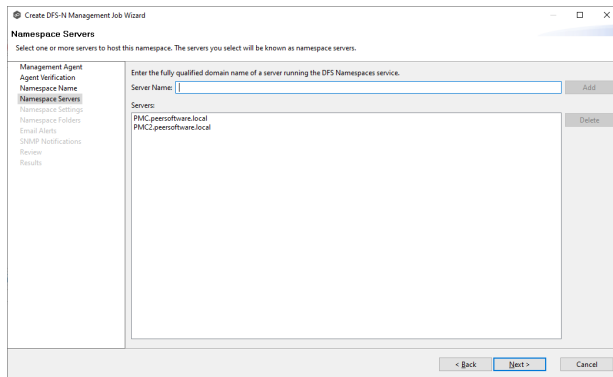
Servers:

PMC.peersoftware.local

Delete

< Back Next > Cancel

2. Add additional servers if desired.



3. Click **Next**.

The [Namespace Settings](#) page appears.

Step 6: Namespace Settings

The **Namespace Settings** page displays the namespace servers selected for the job. You can modify the following namespace server's settings if necessary:

- The local path to the DFS root share for the namespace. The default location of the DFS root share is under C:\DFSRoots\ and is specified in [DFS-N Management Job Preferences](#).
- The access permissions to the namespace server. By default, all users have full access.

To edit a server's settings:

1. In the **DFS Root Local Path** column for the server, the prepopulated path is recommended. If your DFS root share location is different than the default (C:\DFSRoots\), you can modify the first part of **DFS Root Local Path** to match. The second part of the path should be the name of the namespace.


4. Click **Next**.

The [Namespace Folders](#) page appears.

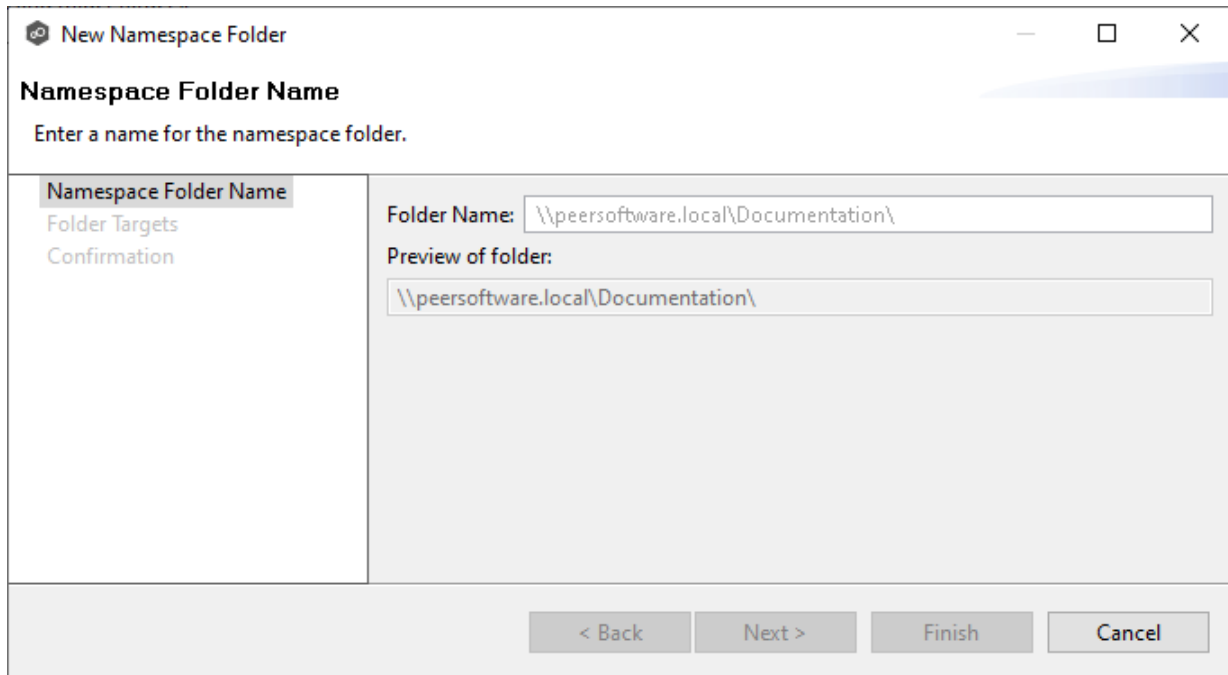
Step 7: Namespace Folders

Initially, the **Folders** page displays only the namespace path. After namespace folders and folder targets are added, they are displayed in a tree structure.

1. Select the namespace path, and then click the **Create** button.



The **Folder Name** dialog appears.



New Namespace Folder

Namespace Folder Name
Enter a name for the namespace folder.

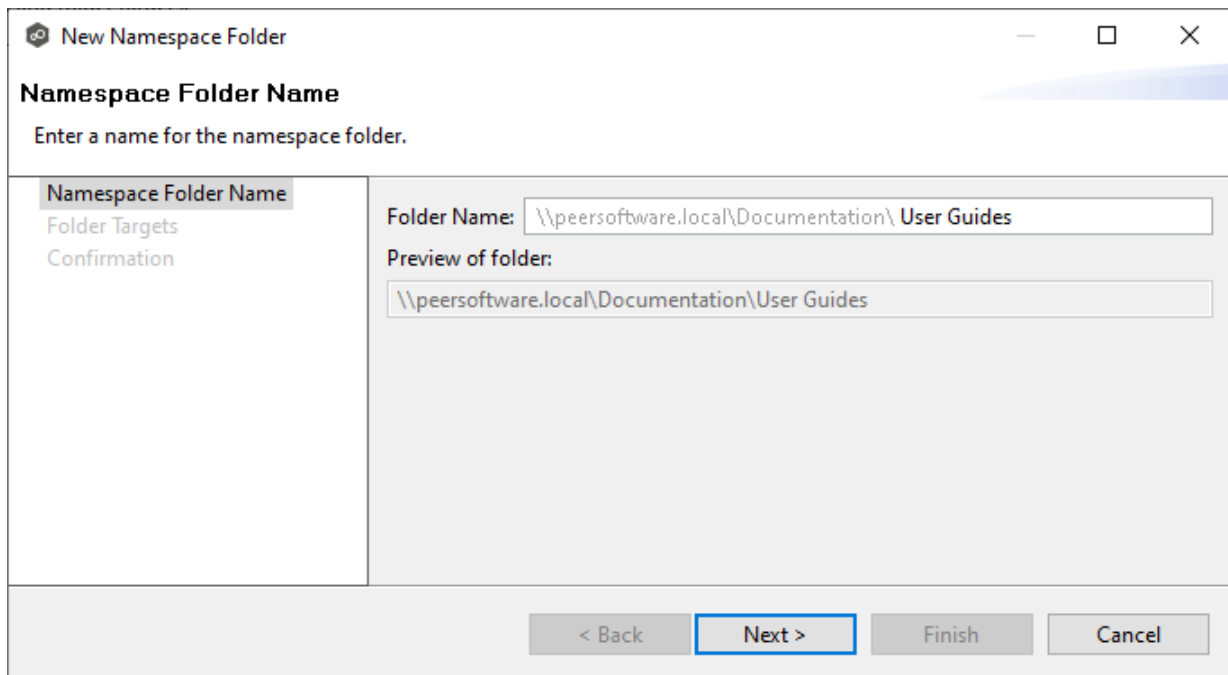
Namespace Folder Name
Folder Targets
Confirmation

Folder Name: \\peersoftware.local\Documentation\
Preview of folder:
\\peersoftware.local\Documentation\

< Back Next > Finish Cancel

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

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



New Namespace Folder

Namespace Folder Name
Enter a name for the namespace folder.

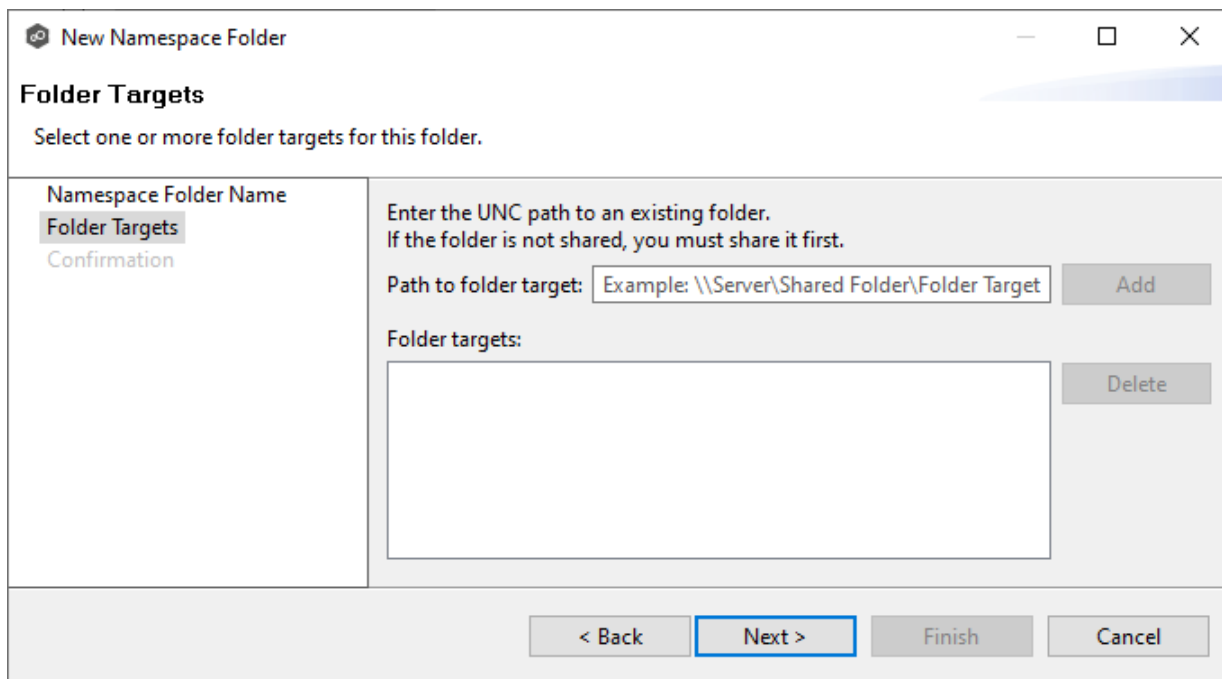
Namespace Folder Name
Folder Targets
Confirmation

Folder Name: \\peersoftware.local\Documentation\ User Guides
Preview of folder:
\\peersoftware.local\Documentation\User Guides

< Back **Next >** Finish Cancel

3. Click **Next**.

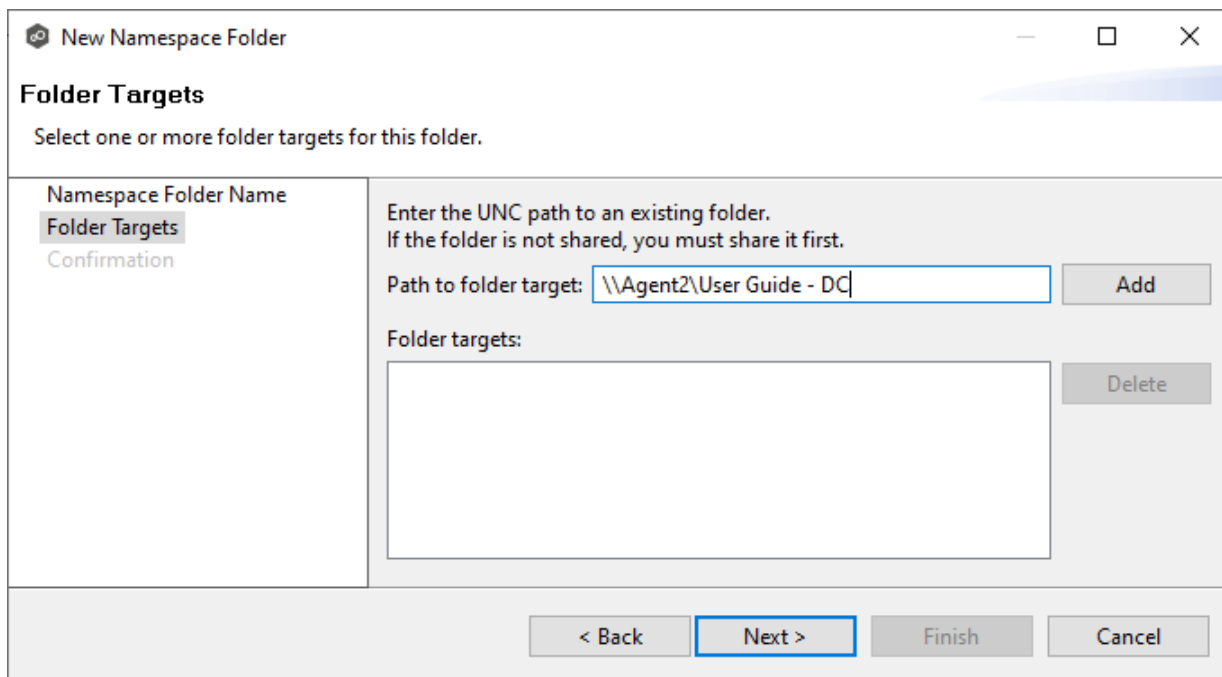
The **Folder Targets** dialog appears.



The screenshot shows a dialog box titled "New Namespace Folder" with a "Folder Targets" section. The section is divided into two columns. The left column contains "Namespace Folder Name", "Folder Targets" (highlighted), and "Confirmation". The right column contains instructions: "Enter the UNC path to an existing folder. If the folder is not shared, you must share it first." Below this is a text input field with the placeholder "Example: \\Server\Shared Folder\Folder Target" and an "Add" button. Underneath is a "Folder targets:" label and a large empty list box with a "Delete" button to its right. At the bottom of the dialog are four buttons: "< Back", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

- (Optional) To add a folder target, enter the UNC path to a shared folder, and then click **Add**.

If you haven't created your folders target yet, you can skip to Step 6 and [add folder targets to the job](#) later.



This screenshot is identical to the previous one, but the text input field now contains the UNC path "\\Agent2\User Guide - DC". The "Add" button is now active and highlighted in grey.

After the share is validated, it appears in the **Folder targets** area:


The screenshot shows the 'New Namespace Folder' dialog box with the 'Folder Targets' tab selected. The 'Folder targets' list contains one entry: '\\Agent2\User Guide - DC'. The 'Next >' button is highlighted.

5. (Optional) Add additional folder targets.

The screenshot shows the 'New Namespace Folder' dialog box with the 'Folder Targets' tab selected. The 'Folder targets' list contains three entries: '\\Agent2\User Guide - DC', '\\Agent3\User Guide - Munich', and '\\Agent4\User Guide - Cape Town'. The 'Next >' button is highlighted.


6. Click **Next**.

The **Confirmation** dialog appears.




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

The **Folders** page reappears.



- Expand the tree to view the folders and folder targets you added.



- Click **Next**.

The [Email Alerts](#) page appears.

Step 8: Email Alerts

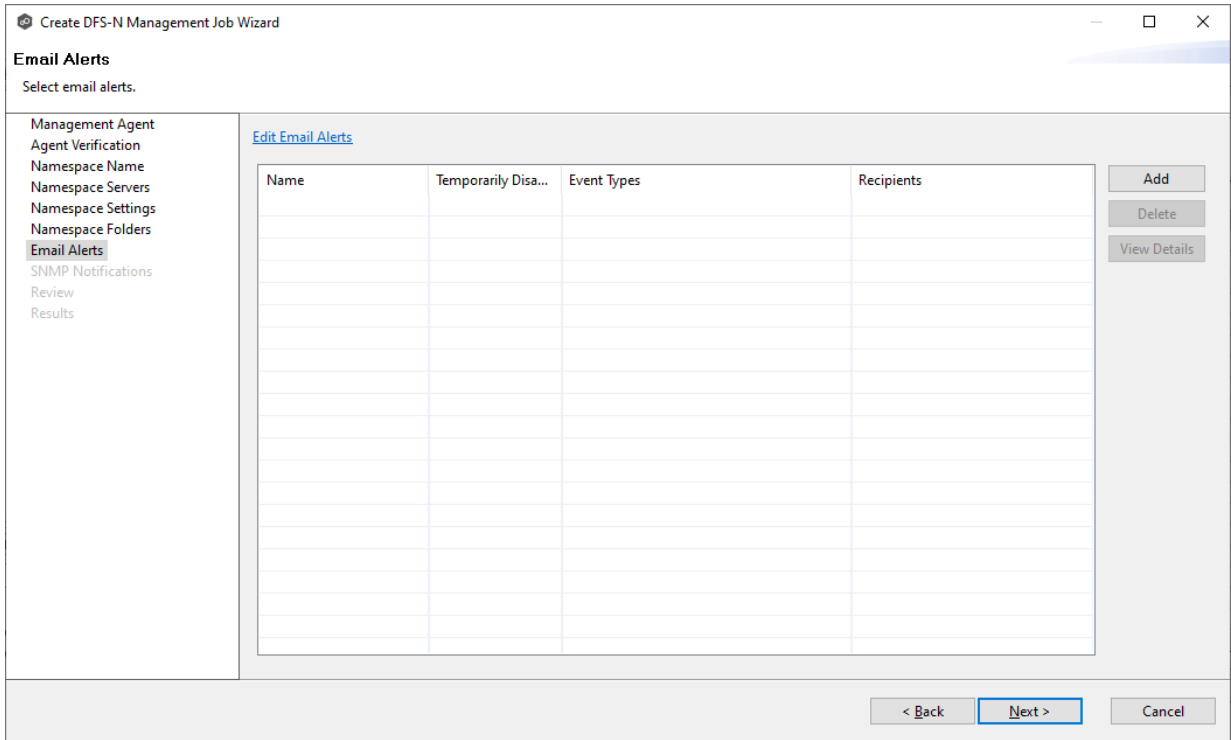
This step is optional.

An [email alert](#) notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in [Preferences](#) and can then be applied to multiple jobs of the same type.

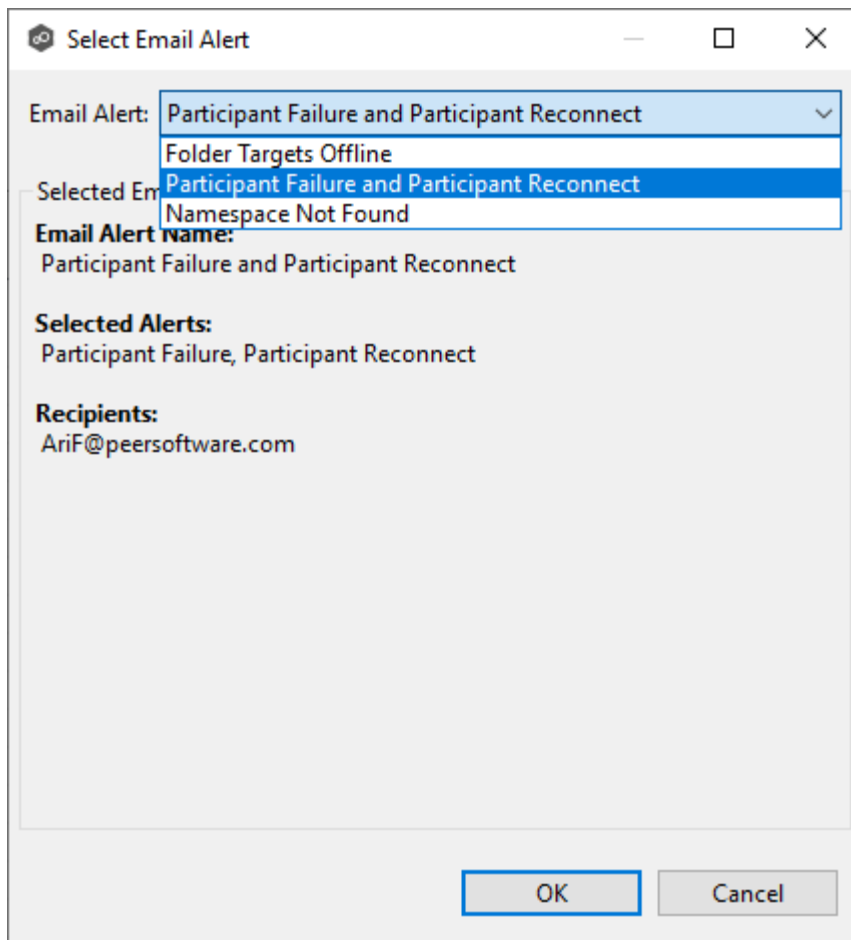
Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply. To create a new alert, see [Email Alerts](#) in the [Preferences](#) section.

To apply an existing email alert to the job:

- Click the **Add** button.




The **Select Email Alert** dialog appears.



2. Select an alert from the **Email Alert** drop-down list, and then click **OK**.

The alert is listed in the **Email Alerts** page.



3. (Optional) Repeat steps 1-3 to apply additional alerts.
4. Click **Next**.

The [SNMP Notifications](#) page appears.

Step 9: SNMP Notifications

This step is optional.

An [SNMP notification](#) notifies recipients when certain type of event occurs, for example, session abort, host failure, system alert. The **SNMP Notifications** page displays a list of notifications that have been applied to the job. When you first create a job, this list is empty. Like email alerts and file filters, an SNMP notification is defined in [Preferences](#) and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create SNMP notifications in advance. However, from this wizard page, you can select an existing SNMP notification to apply to the job or create new SNMP notifications. To create a new alert, see [SNMP Notifications](#) in the [Preferences](#) section.

To apply an existing SNMP notification to the job or disable notifications:

1. Select an SNMP notification from the drop-down list.

If you don't want to send a SNMP notification, select **None - Disabled**.

Management Agent
Agent Verification
Namespace Name
Namespace Servers
Namespace Settings
Namespace Folders
Email Alerts
SNMP Notifications
Review
Results

Select a SMNP notification.

[Edit SNMP Notifications](#)

SNMP Notification: None - Disabled

Selected SNMP Notification Information

No SNMP Notification
Namespace Not Found
SNMP notifications disabled for this job

< Back Next > Cancel

If you select a notification, details about the notification appear in the Selected SNMP Notification Information section.

Management Agent
Agent Verification
Namespace Name
Namespace Servers
Namespace Settings
Namespace Folders
Email Alerts
SNMP Notifications
Review
Results

Select a SMNP notification.

[Edit SNMP Notifications](#)

SNMP Notification: Namespace Not Found

Selected SNMP Notification Information

SNMP Notification Name: Namespace Not Found

Temporarily Disabled: No

Selected Notifications:
Namespace Not Found

< Back Next > Cancel

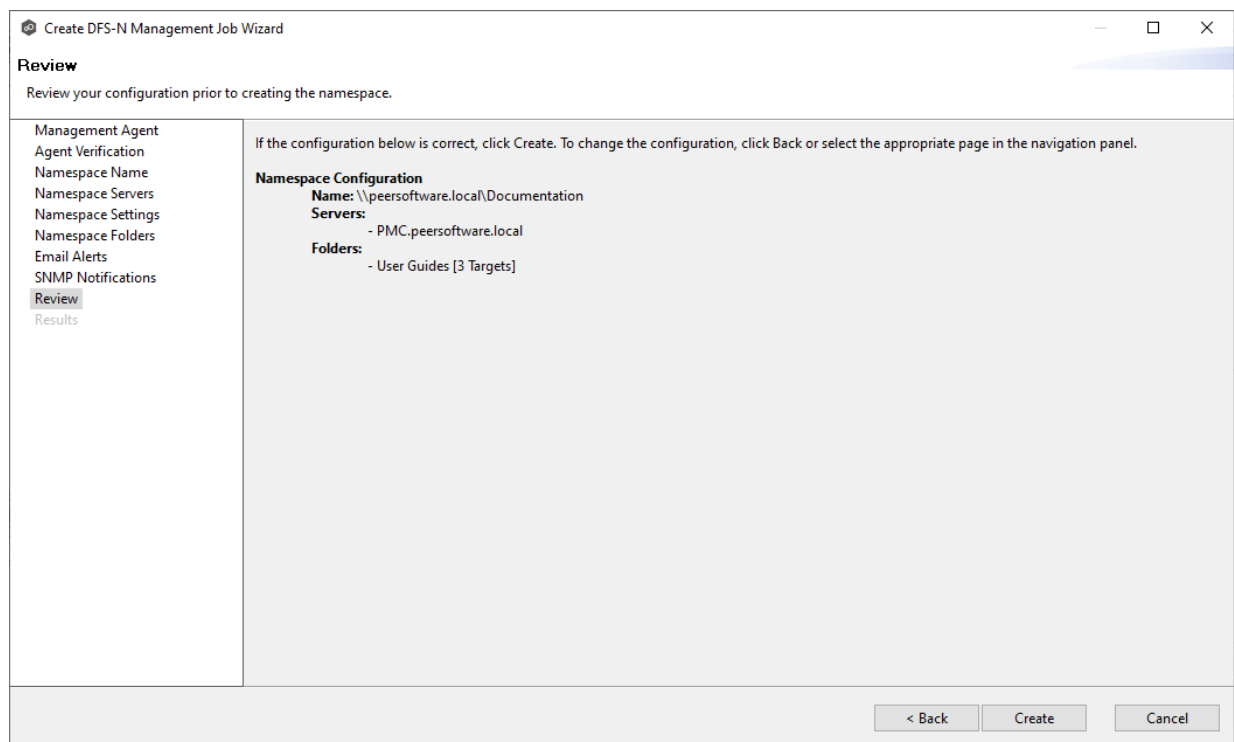
2. Click **Next**.

The [Review](#) page appears.

Step 10: Review

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

1. Review the namespace configuration.



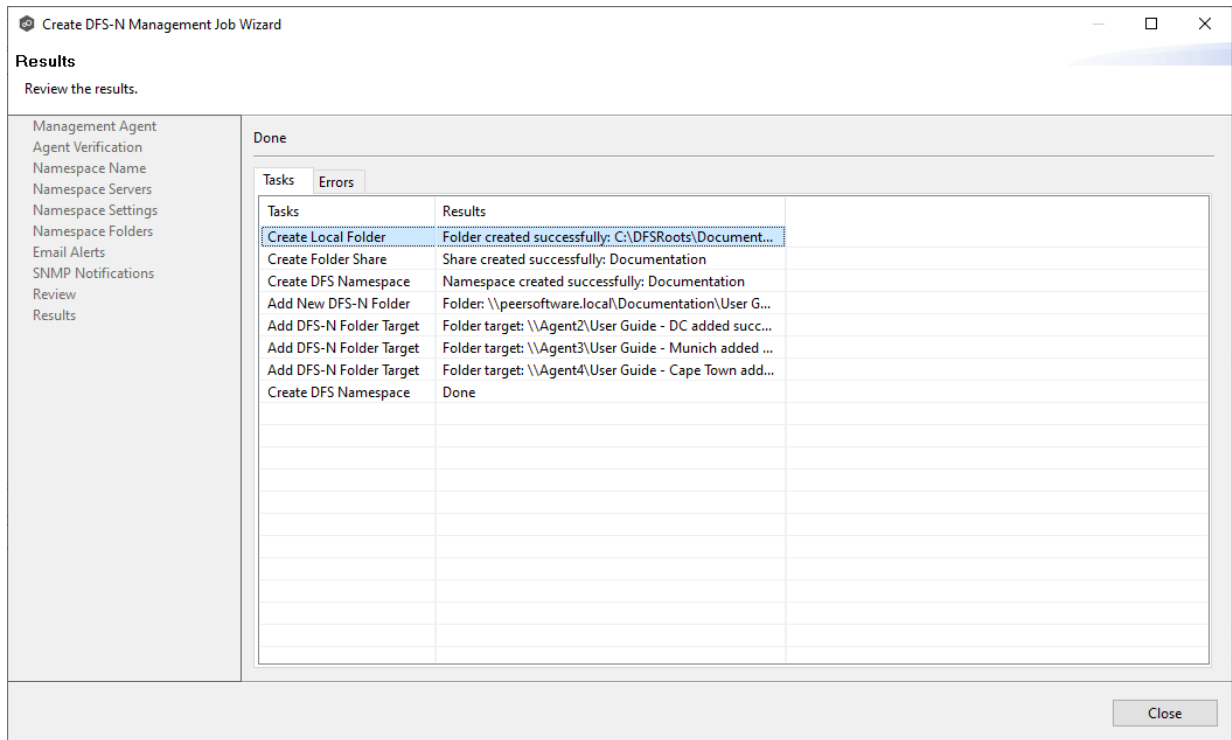
2. Click **Create** if the configuration is correct; otherwise, click **Back** and correct the configuration.

After you click **Create**, the [Results](#) page appears.

Step 11: Results

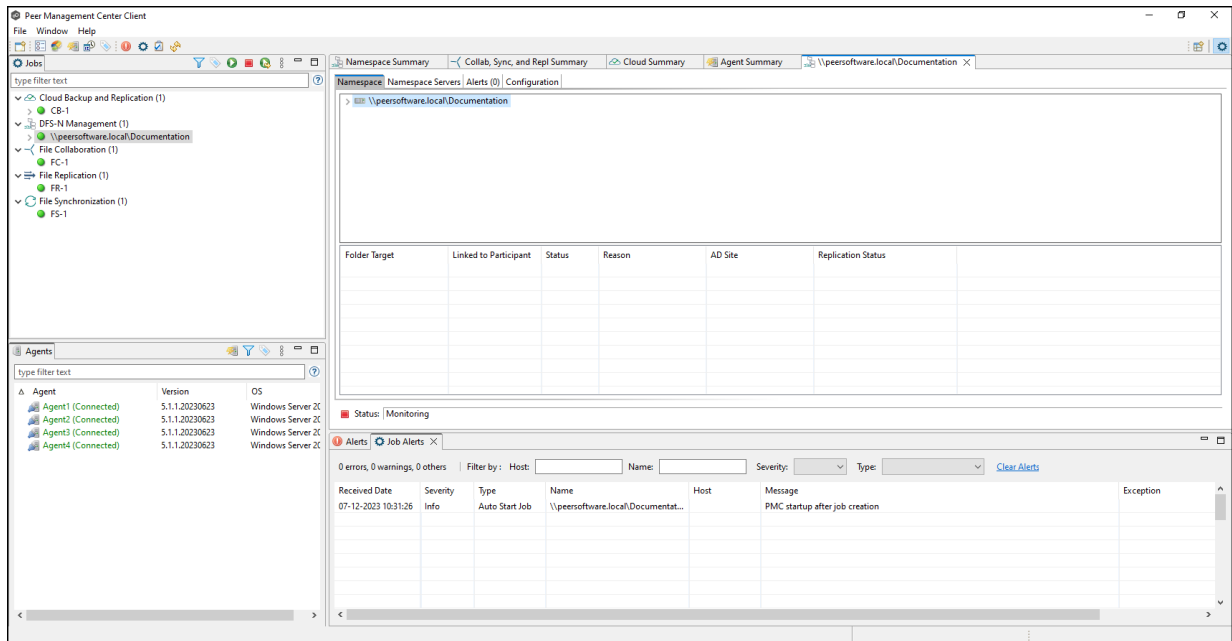
The **Results** page has two tabs: **Tasks** and **Errors**.

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




2. Review each task to verify that it was successful; if there were any errors, click the **Errors** tab to view more details about the problem.
3. After reviewing, click **Close**.

If there are no errors, the job automatically starts and the summary view for the new job is displayed.



- Select the namespace in the **Namespace** tab and then select the namespace folder to view its folder targets.

The folder targets appear in the panel below.



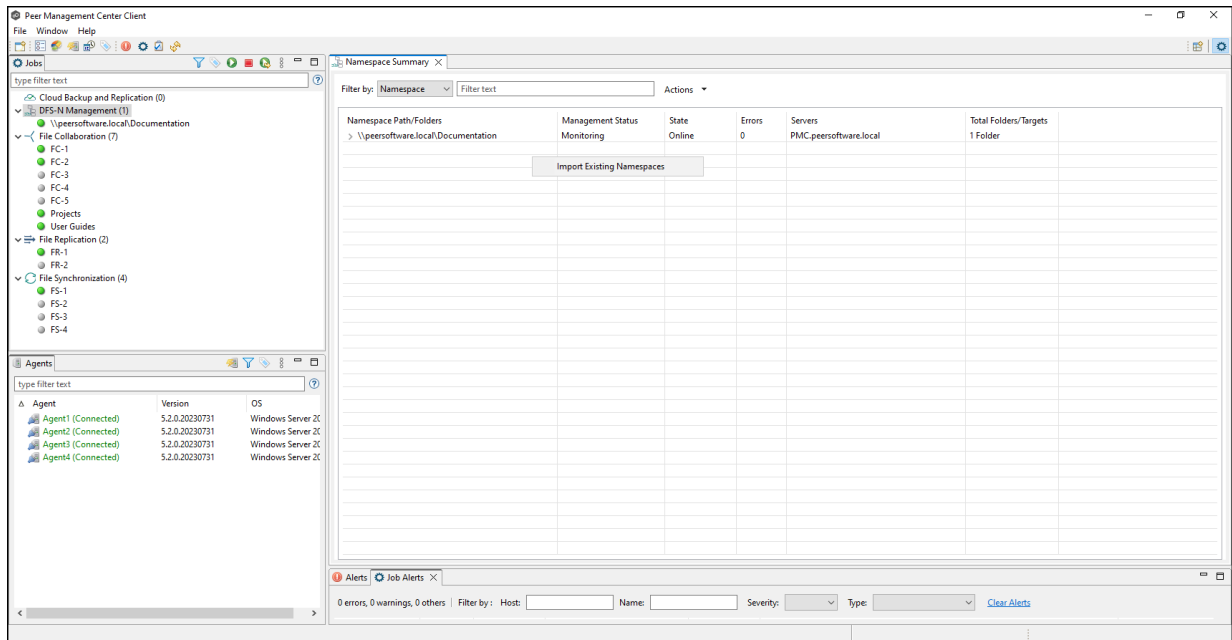
Importing an Existing Namespace

If you have an existing namespace that you want to use in in a File Collaboration or File Synchronization job, you can import the namespace. Importing the namespace automatically creates a new DFS-N Management job with the same name as the imported namespace.

You can then either [link the namespace to an existing File Collaboration or File Synchronization job](#) or [create a new File Collaboration or File Synchronization job](#) that uses the namespace.

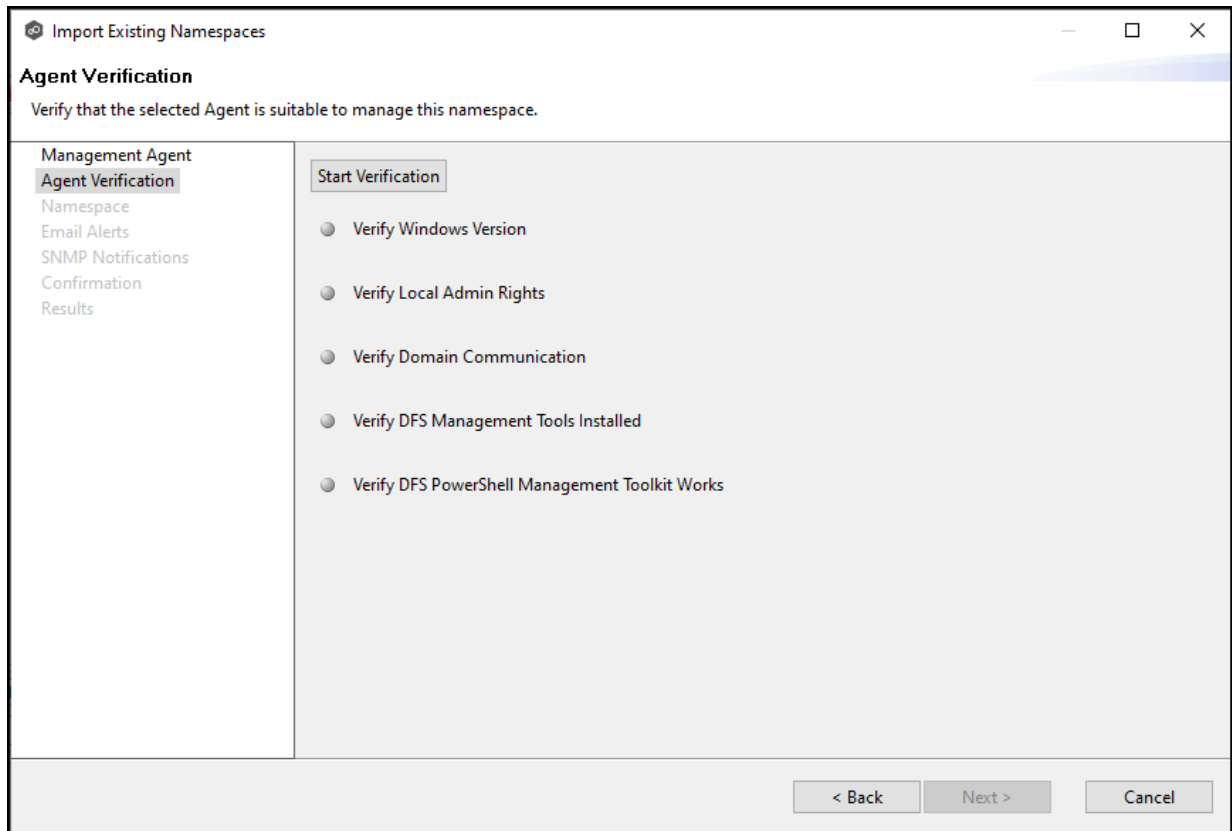
To import an existing namespace:

1. Right-click anywhere in the **Runtime Summary** tab of the **Namespace Summary** view, and then select **Import Existing Namespaces** (or right-click the DFS-N Management job type in the **Jobs** view).



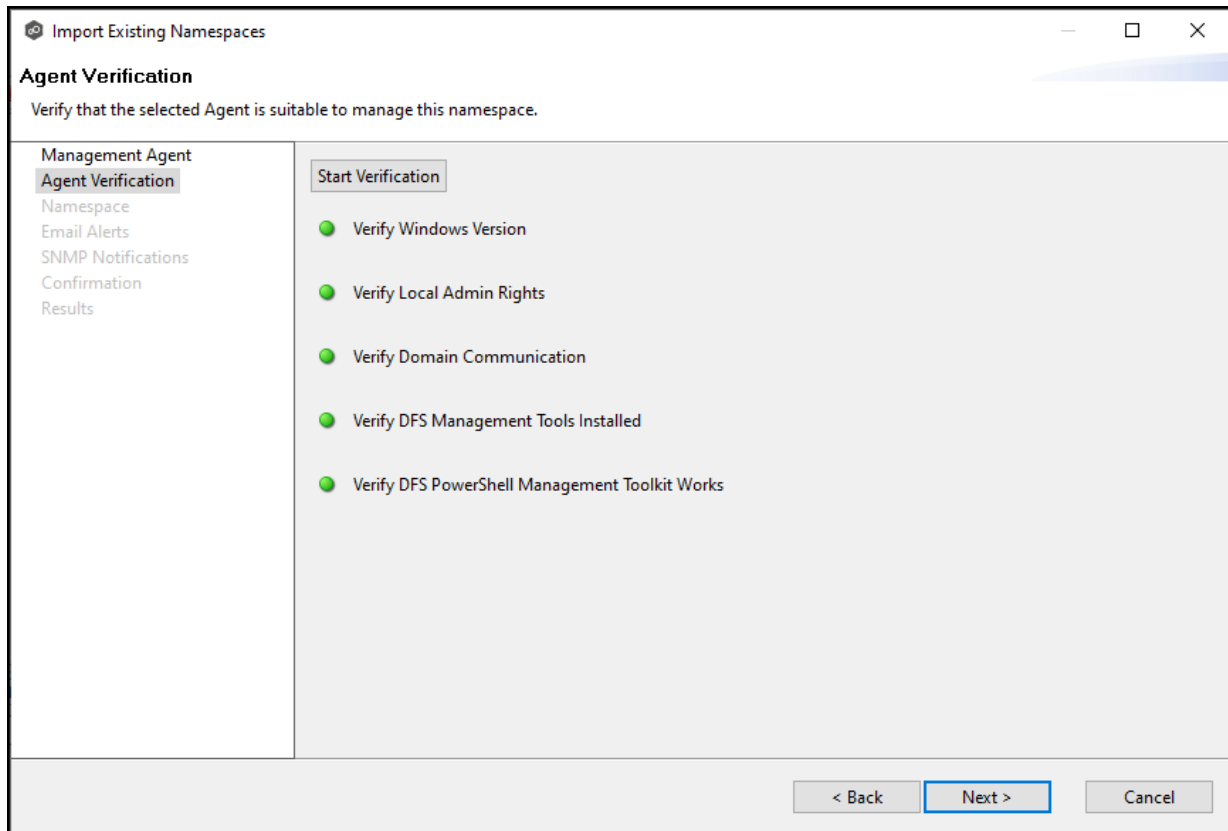
The **Import Existing Namespace** wizard appears.

2. Select a Management Agent, and then click **Next**.



4. If the DFS PowerShell Management toolkit is not installed, click the **Install** button that will appear next to **Verify DFS PowerShell Management Toolkit Installed**.

After the toolkit is installed, the verification continues. A green dot signifies that the verification of that element was successful.



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

The **Namespace** page appears. You have two options for selecting the namespace to import: either by entering its name or by selecting it from a list of namespaces.

6. If you choose **Select By Name**, enter the namespace name, and then click **Validate**. After the namespace is validated, skip to Step 8.

Import Existing Namespaces

Namespace
Select an existing namespace to import.

Management Agent
Agent Verification
Namespace
Email Alerts
SNMP Notifications
Confirmation
Results

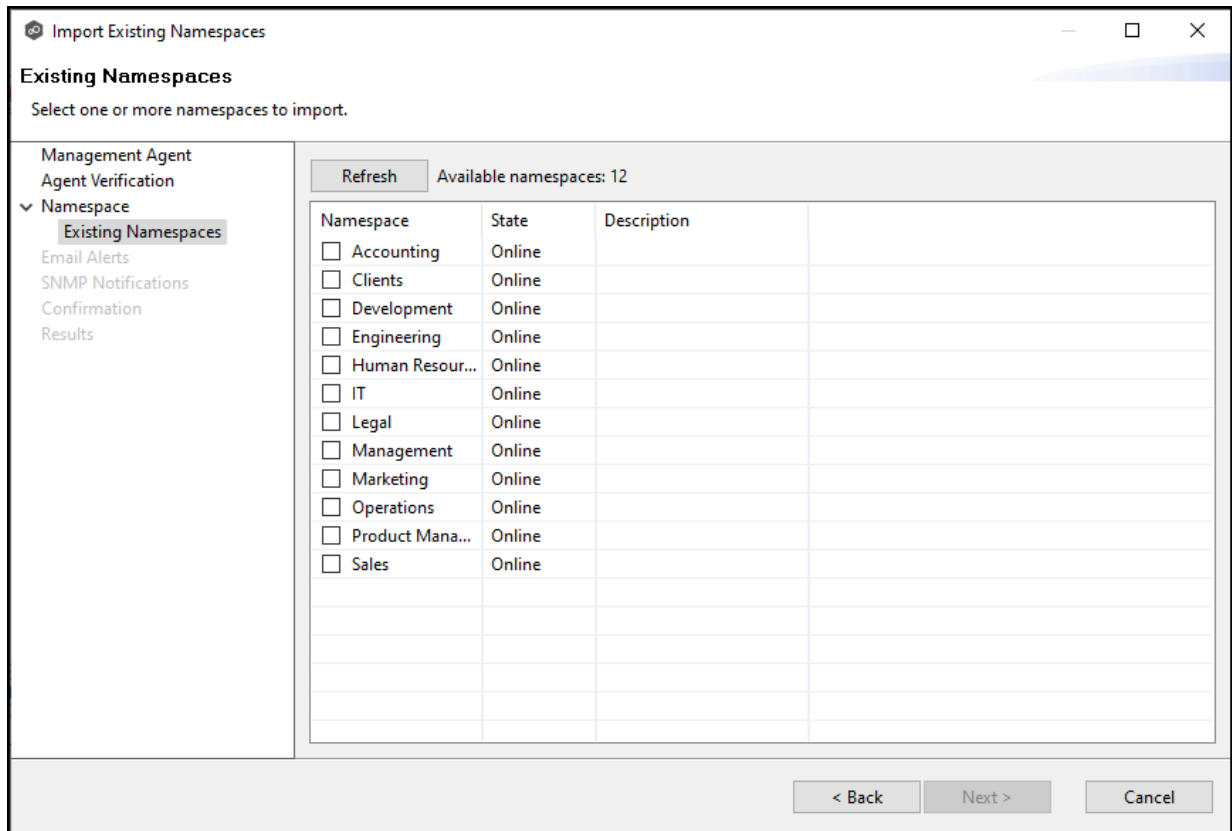
Select By Name
*Namespace Name:
 List All Namespaces

7. If you choose **List All Namespaces**, click **Next**.

The **Existing Namespace** page appears; it displays a table listing the existing namespaces.

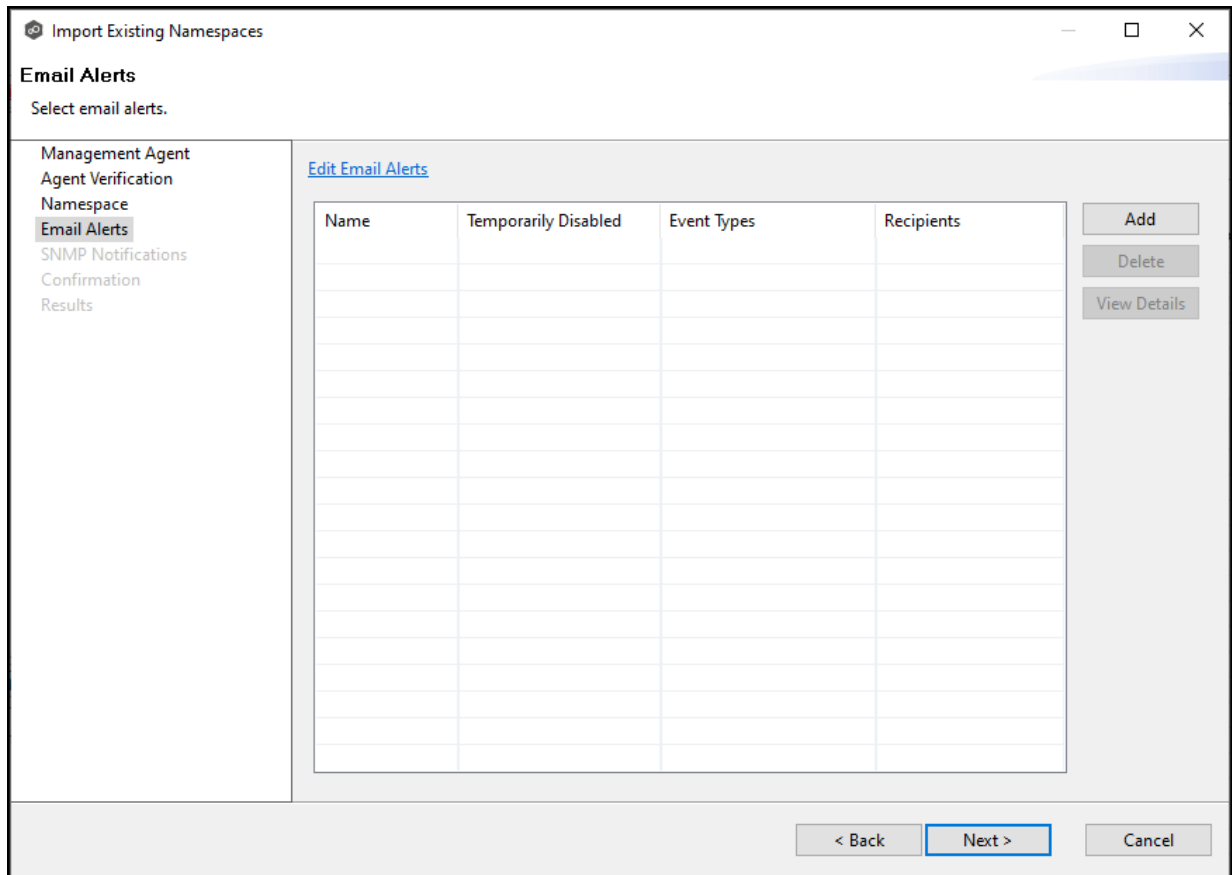
Note: It may take a few minutes for existing namespaces to appear in the table.

8. Select one or more existing namespaces from the table, and then click **Next**.




The **Email Alerts** page appears.

9. (Optional) [Select or create email alerts](#) to apply to the job, and then click **Next**.

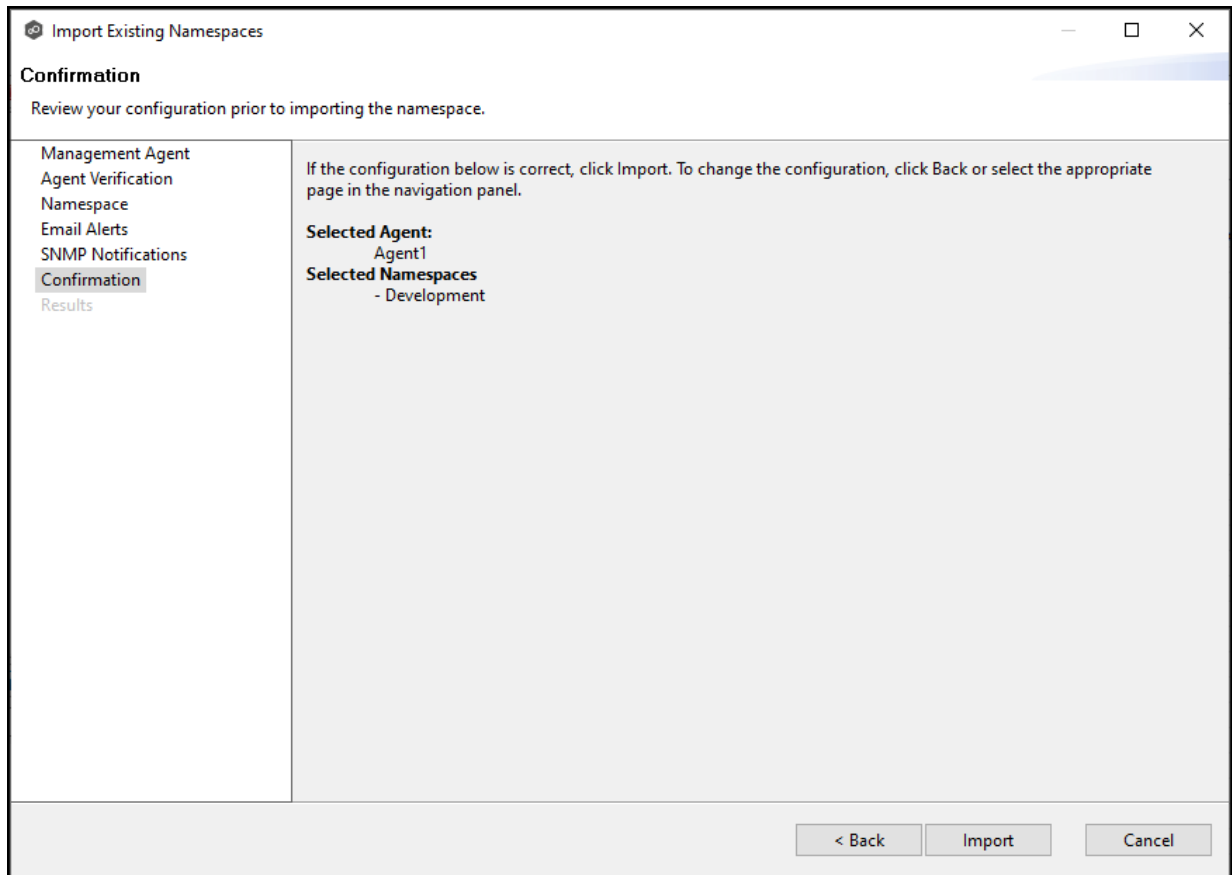


The **SNMP Notifications** page appears.

10. (Optional) [Select or create an SNMP notification](#) to apply to the job, and then click **Next**.



The **Confirmation** page appears.

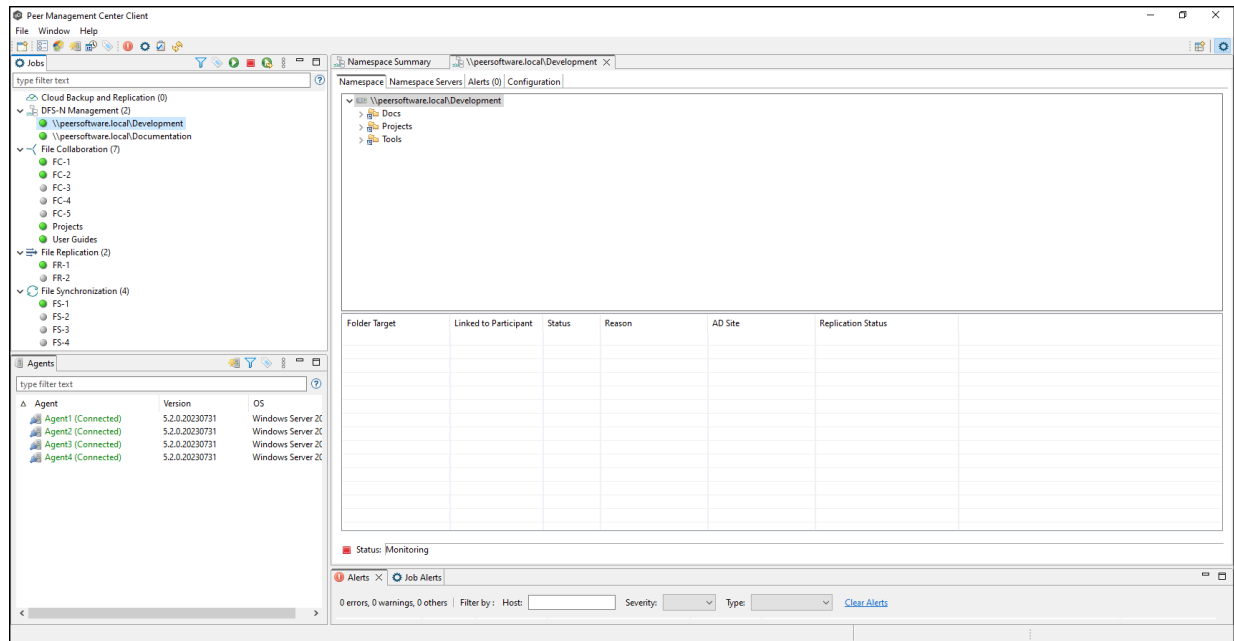


11. Review the configuration.

12. If you need to modify the job configuration after reviewing it, click **Back** until you reach the appropriate page and make your changes.

13. Once you are satisfied with the job configuration, click **Import**.

The **Results** page appears.



Running a DFS-N Management Job

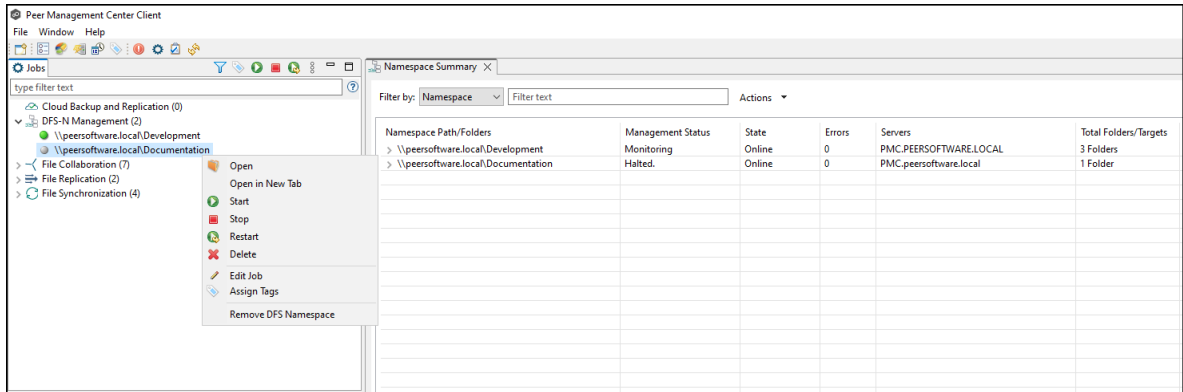
This section describes:

- [Starting a DFS-N Management Job](#)
- [Stopping a DFS-N Management Job](#)

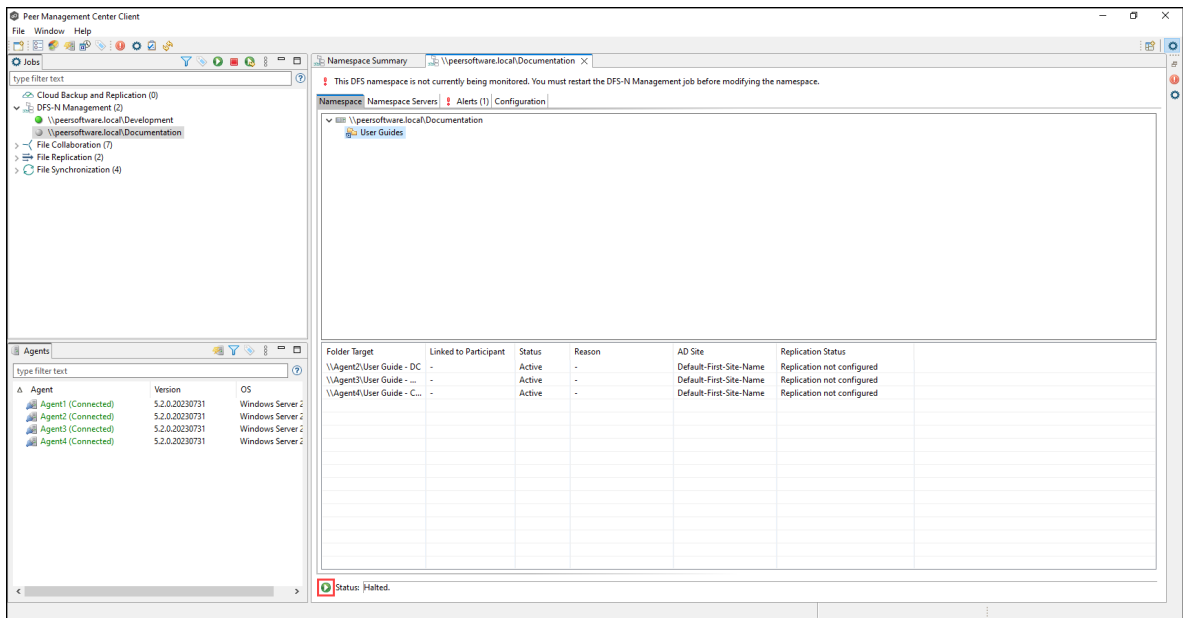
Starting a DFS-N Management Job

To manually start a DFS-N Management job:

1. Choose one of these options:
 - Right-click the job name in the **Jobs** view, and then select **Start**.

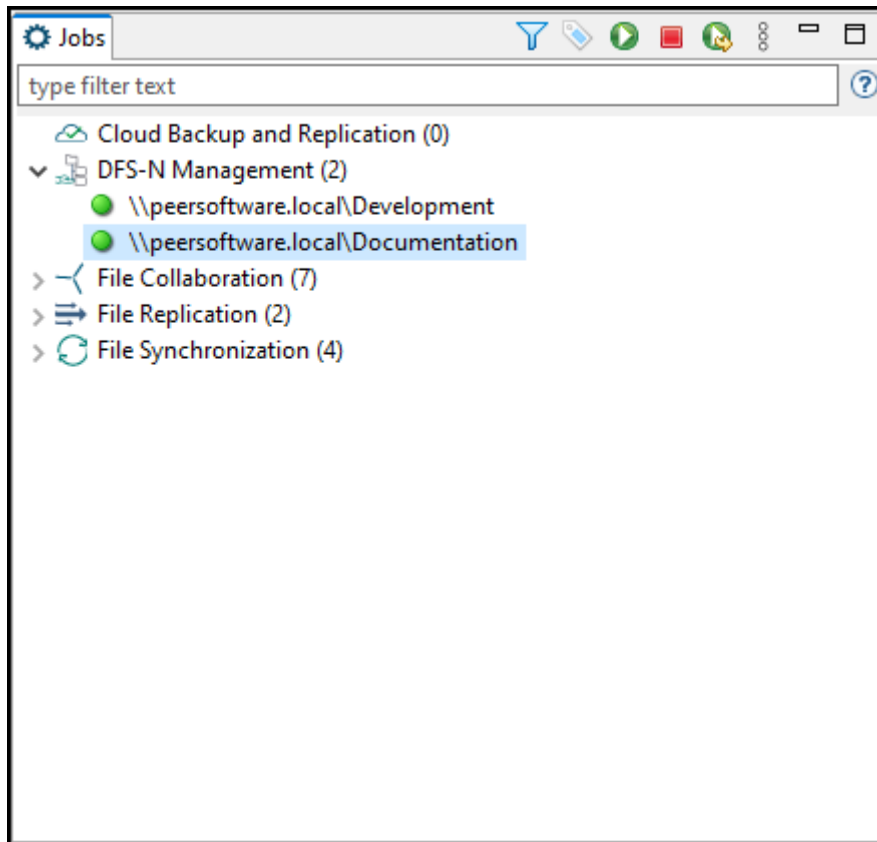


- Open the job and then click the **Start/Stop** button in the bottom left corner of the job's **Namespace** tab in the [DFS-N Management runtime view](#). You may need to scroll to the bottom of the tab to see the **Start/Stop** button.



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

After the job initialization has completed, the job will run. Once the job starts, the icon next to the job name in the **Jobs** view changes from gray to green.



Stopping a DFS-N Management Job

You can stop a DFS-N Management job at any time. Note that you cannot edit a DFS-N Management job while it is stopped.

To stop a DFS-N Management job:

1. Right-click the job name in the **Jobs** view, and then choose **Stop** from the context menu.

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

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

The icon next to the job name in the **Jobs** view changes from green to red.

Managing DFS Namespaces

This section describes:

- [Adding a Namespace Server](#)
- [Adding a Namespace Folder](#)
- [Adding a Namespace Folder Target](#)

Adding a Namespace Server

You can add a namespace server to a namespace.

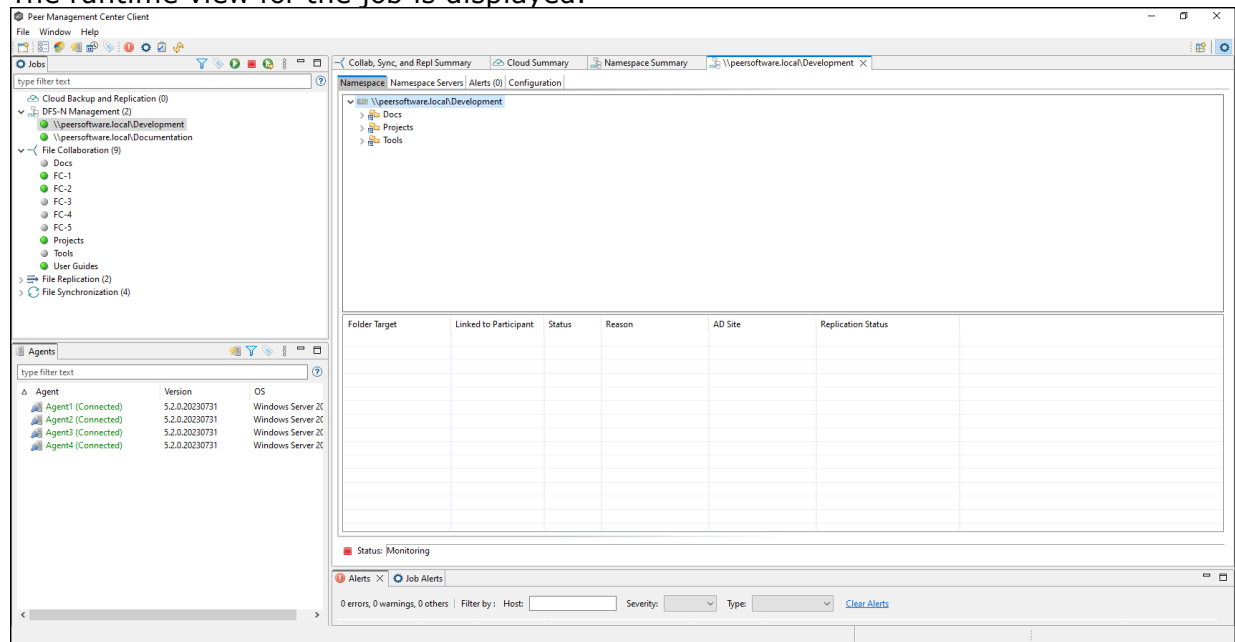
To add a namespace server:

1. Double-click the name of a DFS-N Management job in the **Jobs** view or in the **Namespace Summary** view to open the [runtime view](#) for the job.

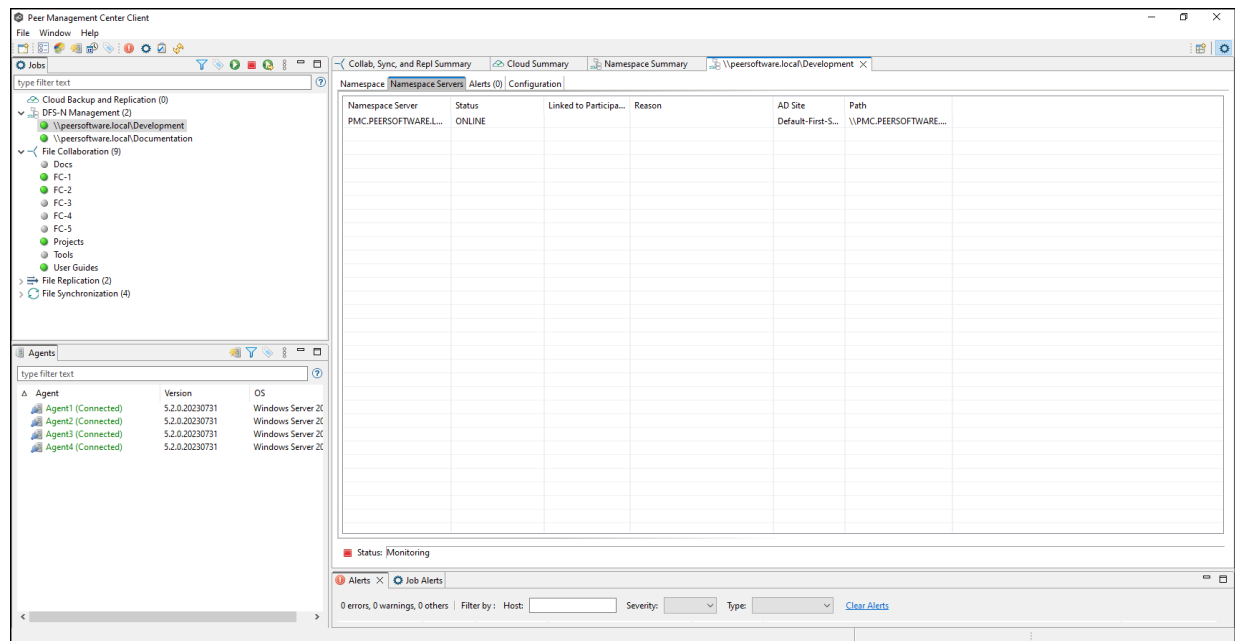
The screenshot displays the Peer Management Center Client interface. The main window is titled "Peer Management Center Client" and shows a "Jobs" view. The left sidebar contains a tree view of the namespace structure, including "DFS-N Management (2)", "Docs", "FC-1" through "FC-5", "Projects", "Tools", "User Guides", "File Replication (2)", and "File Synchronization (4)". The main pane shows a "Namespace Summary" view with a table of jobs. The table has columns for Name, Overall Status, Job Type, Failed Hosts, Quarant., Retries, Errors, Warnings, Open F..., Pending..., Queue..., Background..., Scan Status, Elapsed..., and Session Stru... The table lists various jobs such as "User Guides", "Tools", "Projects", "FS-4", "FS-3", "FS-2", "FS-1", "FR-2", "FR-1", "FC-5", "FC-4", "FC-3", "FC-2", "FC-1", and "Docs". The status of each job is indicated by a green dot (Running) or a red dot (Stopped). The bottom status bar shows "Active Jobs -> Failed Participants: 0 of 4 | Bytes Pending: 0 bytes | Bytes Transferred: 0 bytes | Opens: 0 | Initial Scans Completed: 6 of 6 | Total Size: 981.51 MB | Total Files: 100 | Total Directories: 6".

Name	Overall Status	Job Type	Failed Hosts	Quaran.	Retries	Errors	Warnin...	Open F...	Pendin...	Queue...	Backgr...	Scan Status	Elapge...	Session Stru...
User Guides	Running	File Collaboration		0	0	3	0	0	0 bytes	0	0	Completed - 00:00...	06:11:50	Size: 0 bytes,
Tools	Stopped	File Collaboration		0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes,
Projects	Running	File Collaboration		0	0	3	0	0	0 bytes	0	0	Completed - 00:00...	06:11:52	Size: 0 bytes,
FS-4	Stopped	File Synchronization		0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes,
FS-3	Stopped	File Synchronization		0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes,
FS-2	Stopped	File Synchronization		0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes,
FS-1	Running	File Synchronization		0	0	2	0	0	0 bytes	0	0	Completed - 00:00...	06:10:51	Size: 0 bytes,
FR-2	Stopped	File Replication		0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes,
FR-1	Running	File Replication		0	0	3	0	0	0 bytes	0	0	Completed - 00:00...	06:12:04	Size: 0 bytes,
FC-5	Stopped	File Collaboration		0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes,
FC-4	Stopped	File Collaboration		0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes,
FC-3	Stopped	File Collaboration		0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes,
FC-2	Running	File Collaboration		0	0	0	0	0	0 bytes	0	0	Completed - 00:00...	05:46:28	Size: 0 bytes,
FC-1	Running	File Collaboration		0	0	0	0	0	0 bytes	0	0	Completed - 00:00...	05:47:36	Size: 981.51 B,
Docs	Stopped	File Collaboration		0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 0 bytes,

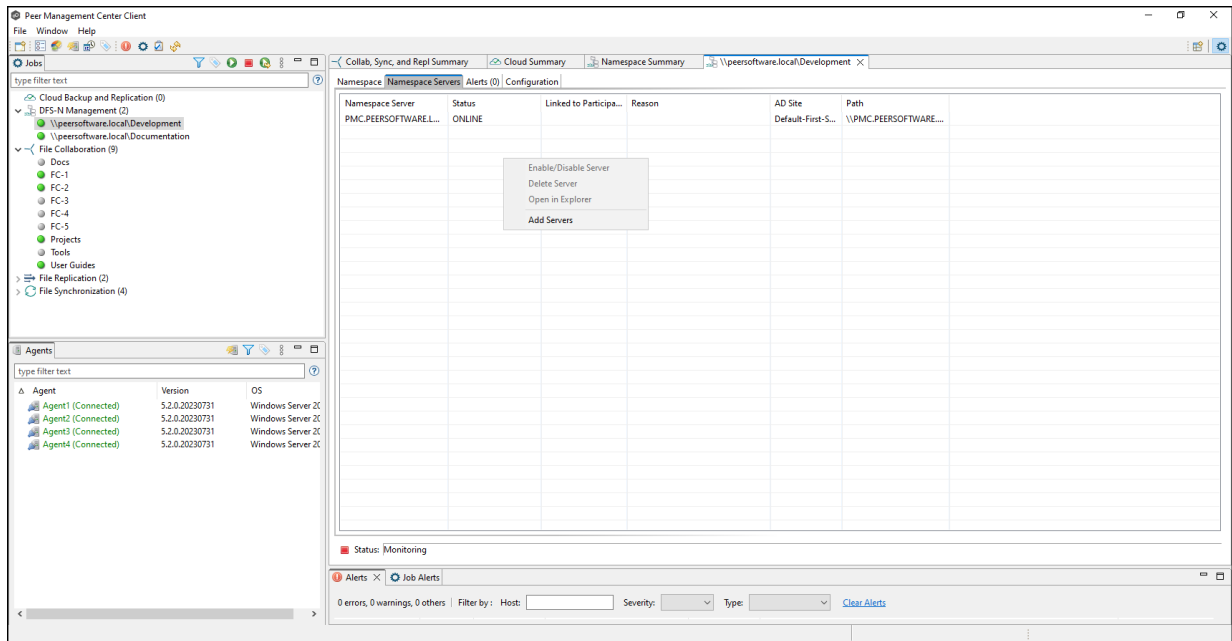
The runtime view for the job is displayed.



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

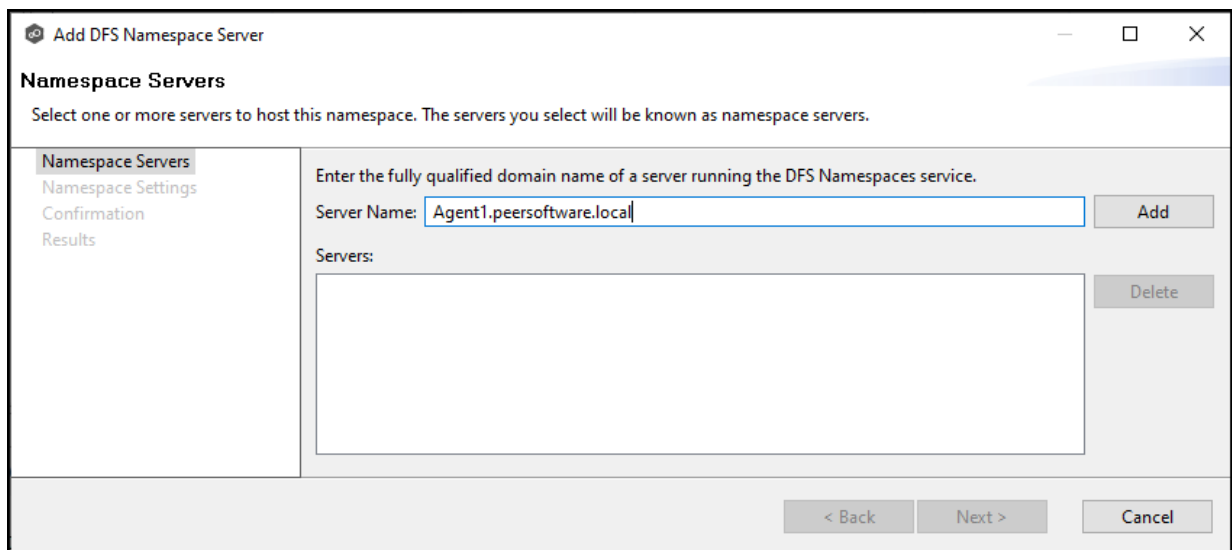


3. Right-click anywhere in the **Namespace Servers** tab, and then select **Add Servers**.

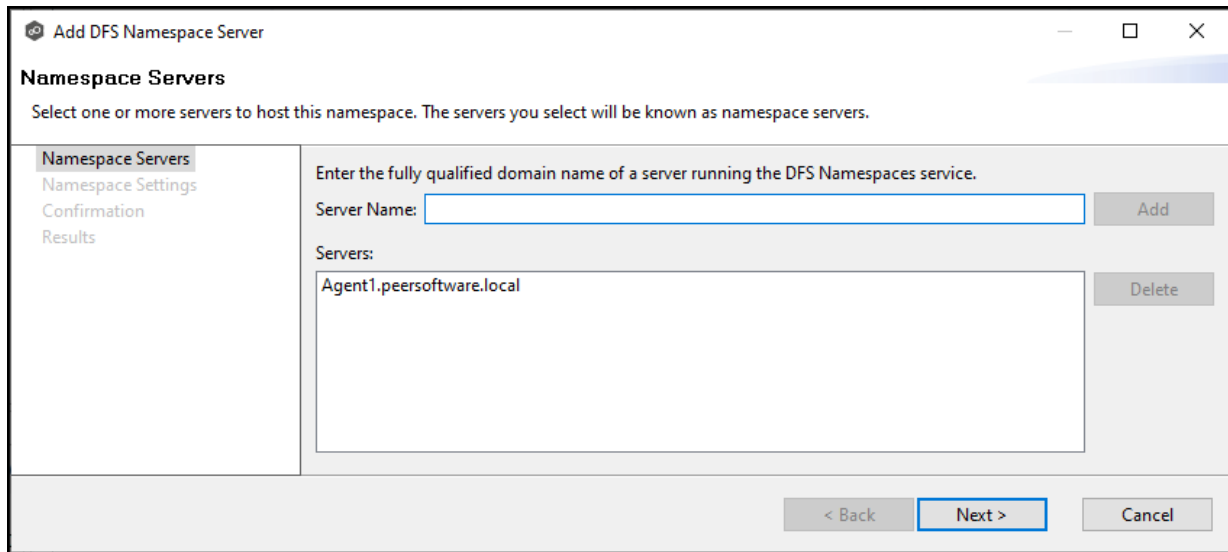


The **Add DFS Namespace Server** wizard appears.

4. Enter the fully qualified domain name (FQDN) of a namespace server in the **Server Name** field, and then click **Add**.

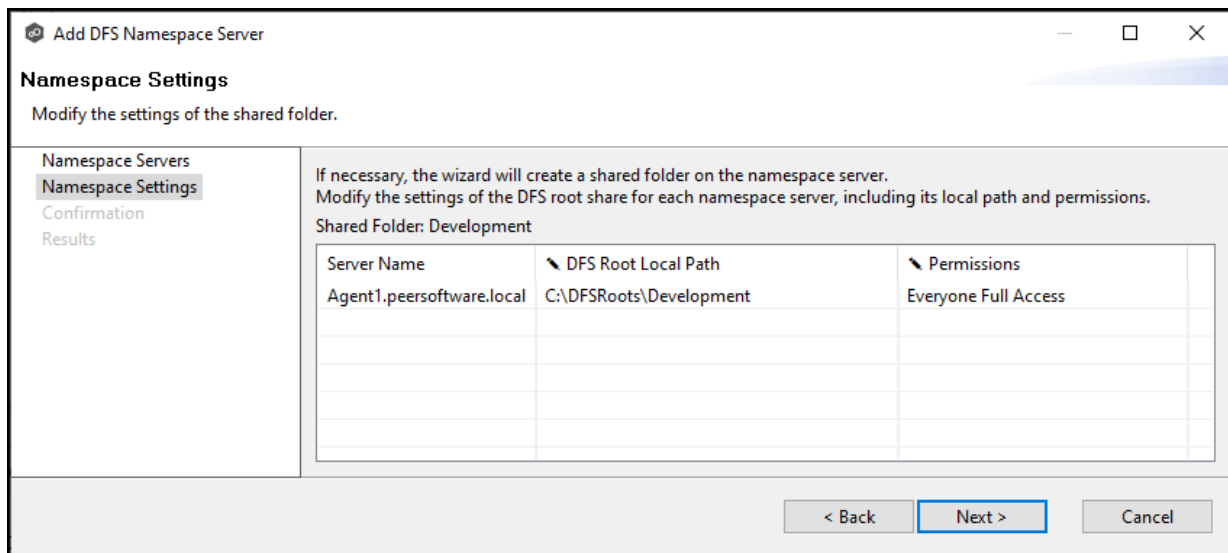


The server FQDN is listed in the area below.



5. Add additional servers if desired.
6. Click **Next**.


The **Namespace Settings** page is displayed.



7. (Optional) Edit the namespace server settings: **DFS Root Share Path** and **Permissions**.
 - To modify the local path to the DFS root share for the namespace, type a new path in the **DFS Root Local Path** column. The default location of the DFS root share is under C:\DFSRoots\ and is specified in [DFS-N Management Job Preferences](#).
 - To modify the access permissions, select a new set using the drop-down menu in the **Permissions** column.

8. Click **Next**.

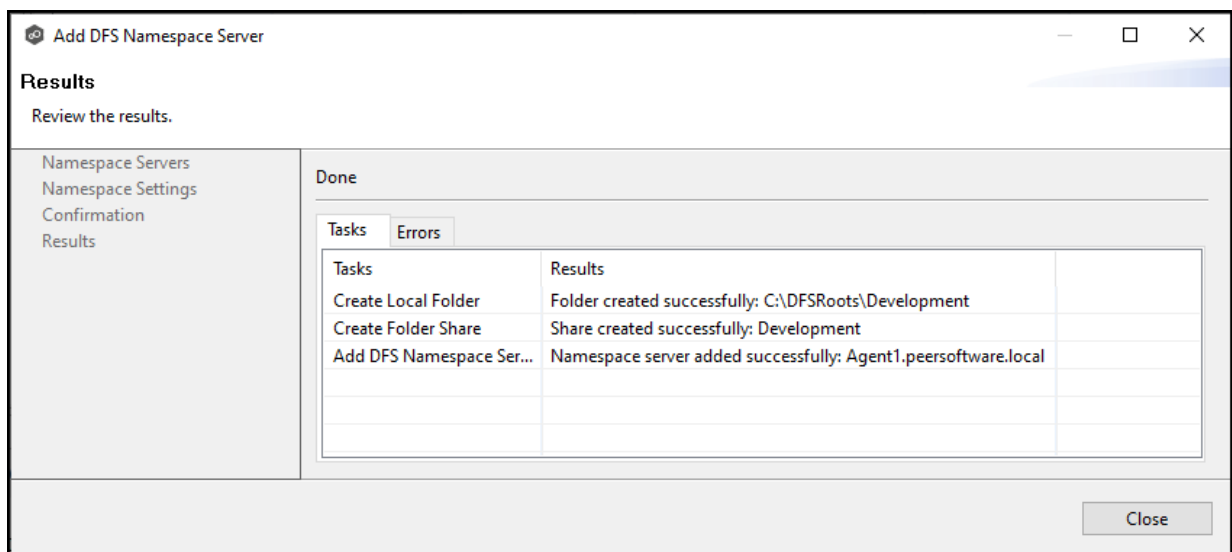
The **Confirmation** page is displayed.



9. Review the namespace server configuration.

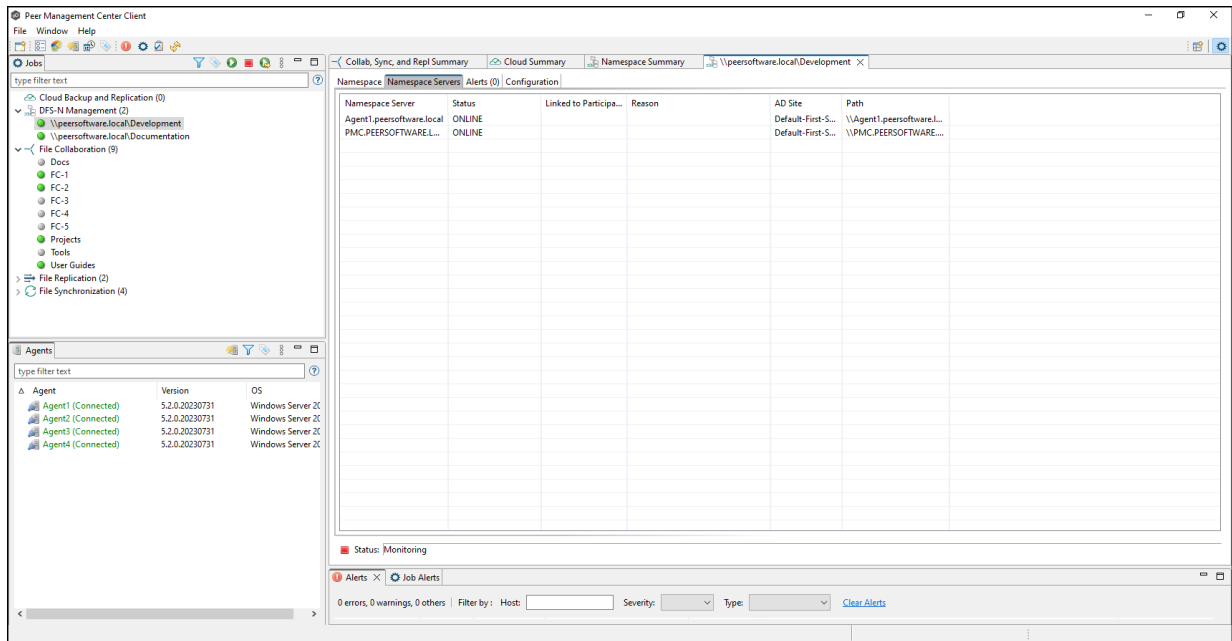
10. Click **Add** if the configuration is correct; otherwise, click **Back** and correct the configuration.

The **Results** page is displayed.



11. Click **OK**.

The newly added server is listed in the **Namespace Servers** tab.



Adding a Namespace Folder

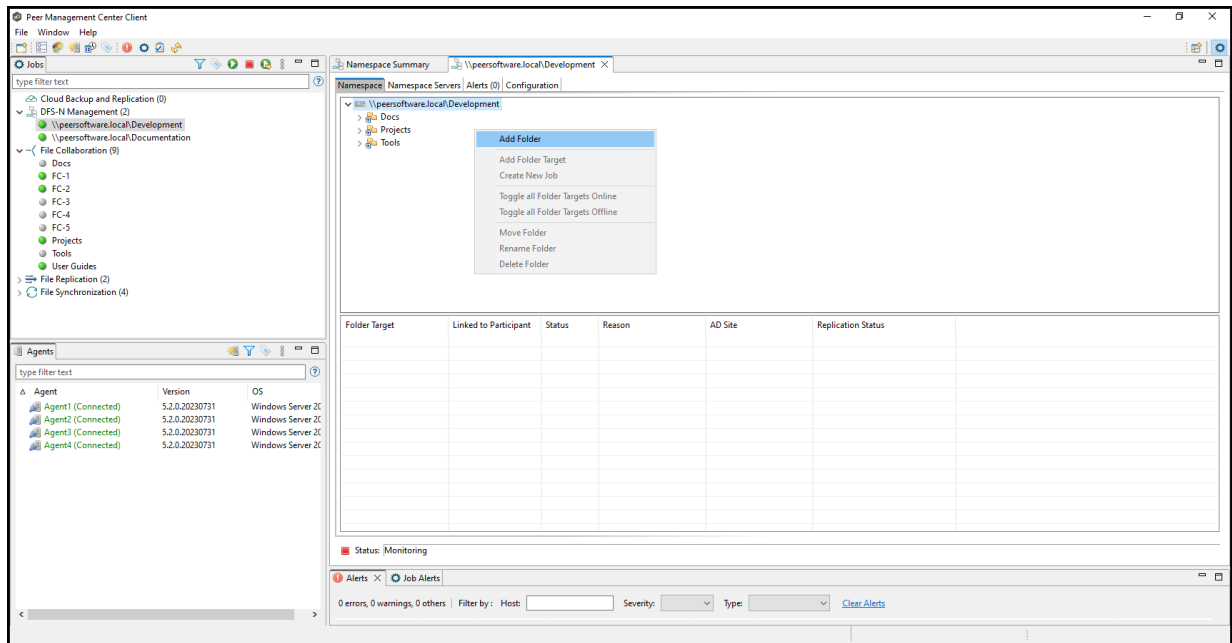
You can add a namespace folder to a namespace. At the same time, you can its folder targets or you can [add folder targets](#) later.

Note that PeerGFS does not currently support creation of nested namespace folders. In other words, you cannot add a namespace folder that is a subfolder of a namespace folder, although this can be done when using the Microsoft DFS Management Tool.

A DFS-N Namespace job must be running before you can edit it.

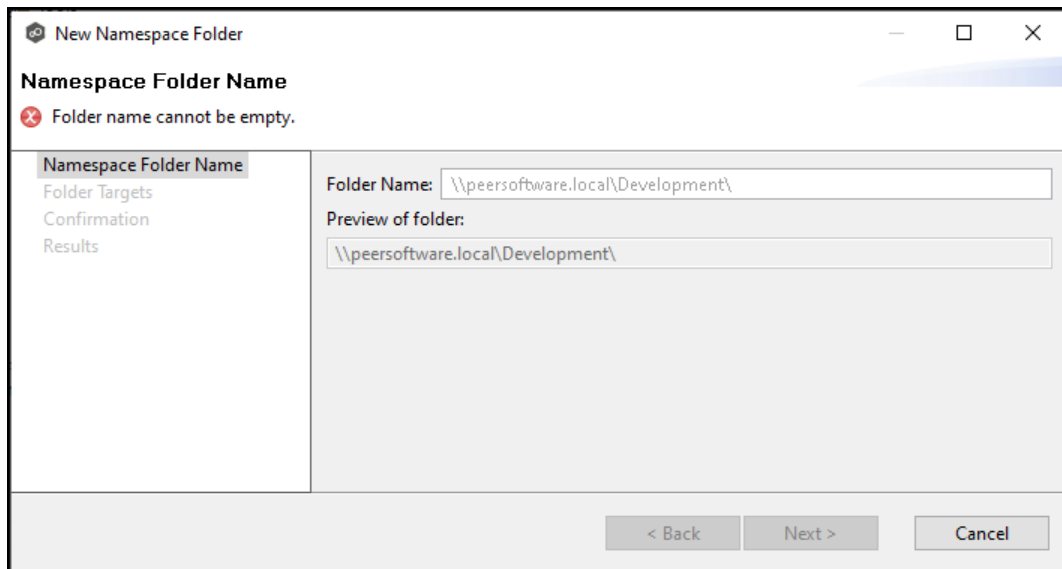
To add a namespace folder to a namespace:

1. Double-click the DFS-N Management job name in the **Jobs** view or in the [Namespace Summary view](#) to open the runtime view for the job.

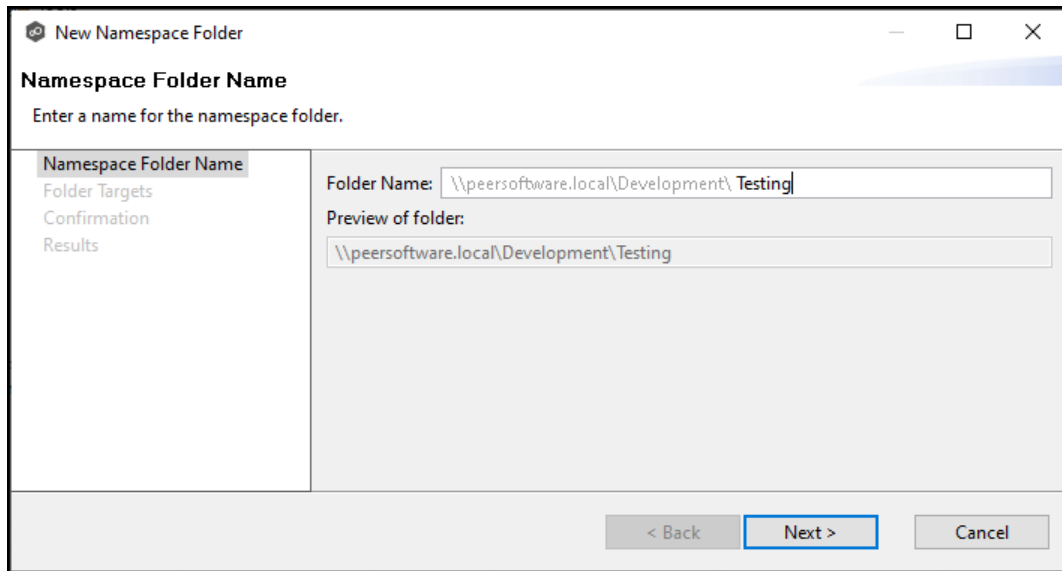


The **New Folder** wizard appears.

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



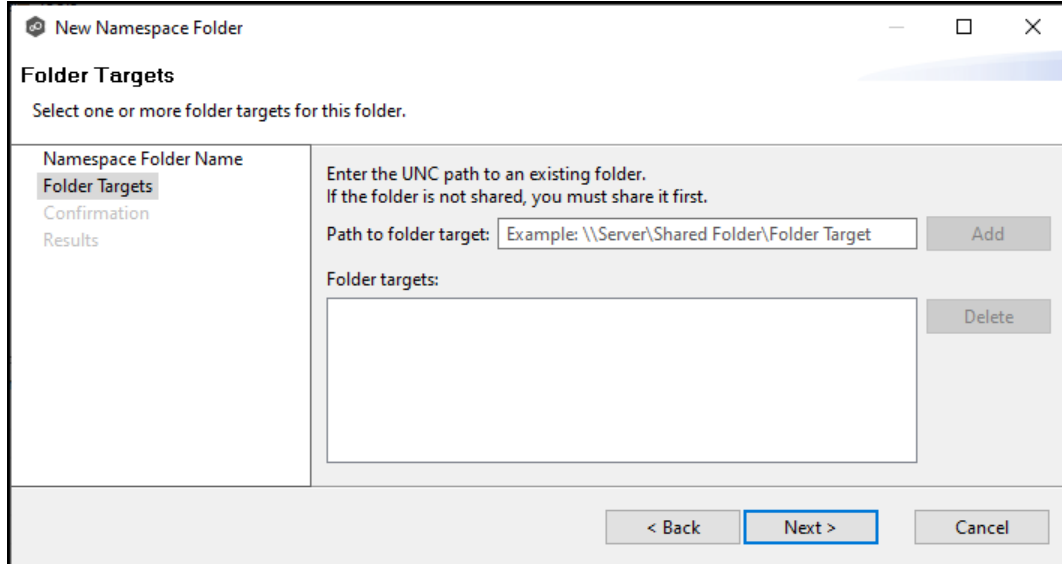
As you enter the folder name, a preview of the folder name and path appears below.



The screenshot shows a window titled "New Namespace Folder" with a sidebar on the left containing "Namespace Folder Name", "Folder Targets", "Confirmation", and "Results". The main area is titled "Namespace Folder Name" and contains the instruction "Enter a name for the namespace folder." Below this is a "Folder Name:" text box containing the path "\\peersoftware.local\Development\Testing". A "Preview of folder:" section below it shows the same path. At the bottom, there are three buttons: "< Back", "Next >" (highlighted with a blue border), and "Cancel".

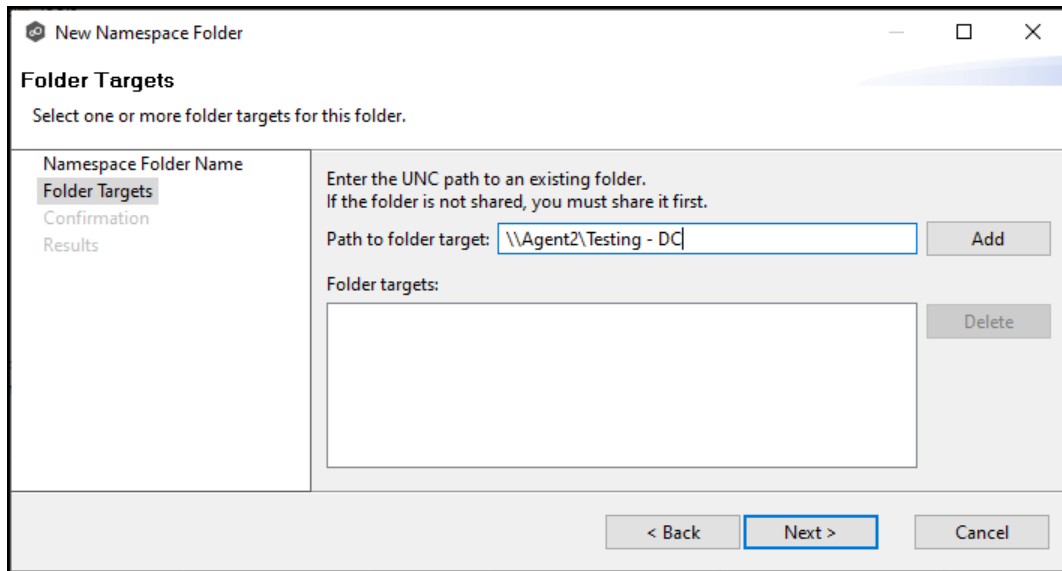
4. Click **Next**.

The **Folder Targets** page is displayed. It is optional to add folder targets for the namespace folder at this point. You can [add them later](#) if you wish. If you choose to add the folder targets now, they must already exist and be shared.



The screenshot shows the same window titled "New Namespace Folder", but the sidebar now has "Folder Targets" selected. The main area is titled "Folder Targets" and contains the instruction "Select one or more folder targets for this folder." Below this is a text box for "Path to folder target:" with the example "\\Server\Shared Folder\Folder Target" and an "Add" button. A "Folder targets:" list box is empty, with a "Delete" button to its right. At the bottom, there are three buttons: "< Back", "Next >" (highlighted with a blue border), and "Cancel".

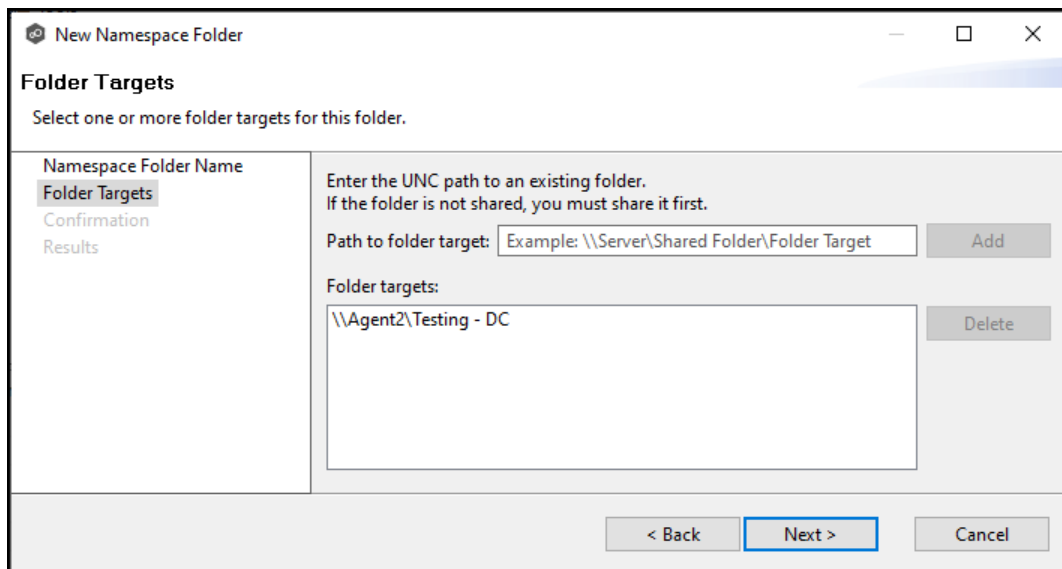
5. Click **Next** if you do not want to add folder targets at this point and continue with Step 9.
6. (Optional) Enter the UNC path to the shared folder you want to be a folder target.



The screenshot shows the 'New Namespace Folder' dialog box with the 'Folder Targets' section selected. The 'Path to folder target' field contains the UNC path '\\Agent2\Testing - DC'. The 'Folder targets' list is currently empty. The 'Add' button is visible next to the path field, and the 'Delete' button is visible next to the folder targets list. The 'Next >' button is highlighted at the bottom.

7. Click **Add**.

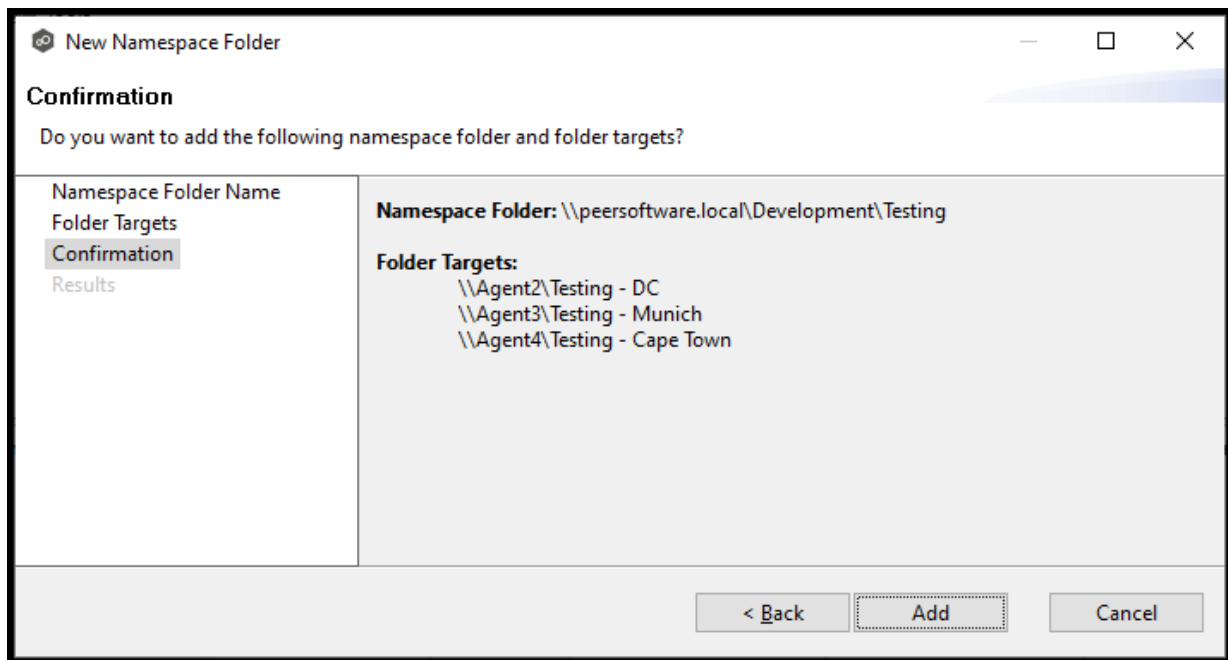
The folder target is added to the **Folder targets** section.



The screenshot shows the 'New Namespace Folder' dialog box with the 'Folder Targets' section selected. The 'Path to folder target' field now contains the example path '\\Server\Shared Folder\Folder Target'. The 'Folder targets' list now contains the path '\\Agent2\Testing - DC'. The 'Add' button is visible next to the path field, and the 'Delete' button is visible next to the folder targets list. The 'Next >' button is highlighted at the bottom.

8. Repeat Steps 6-7 to add additional folder targets if desired.
9. Click **Next**.

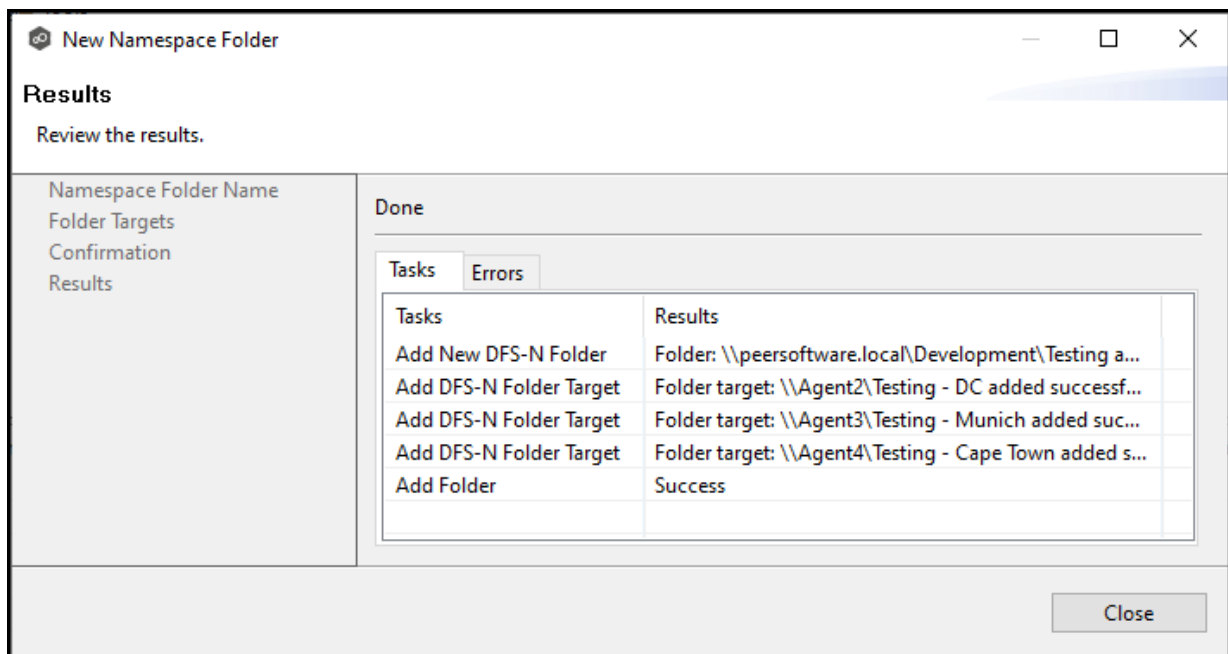
The **Confirmation** page is displayed.



10. Review the folders and folder targets.


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

The **Results** page is displayed.



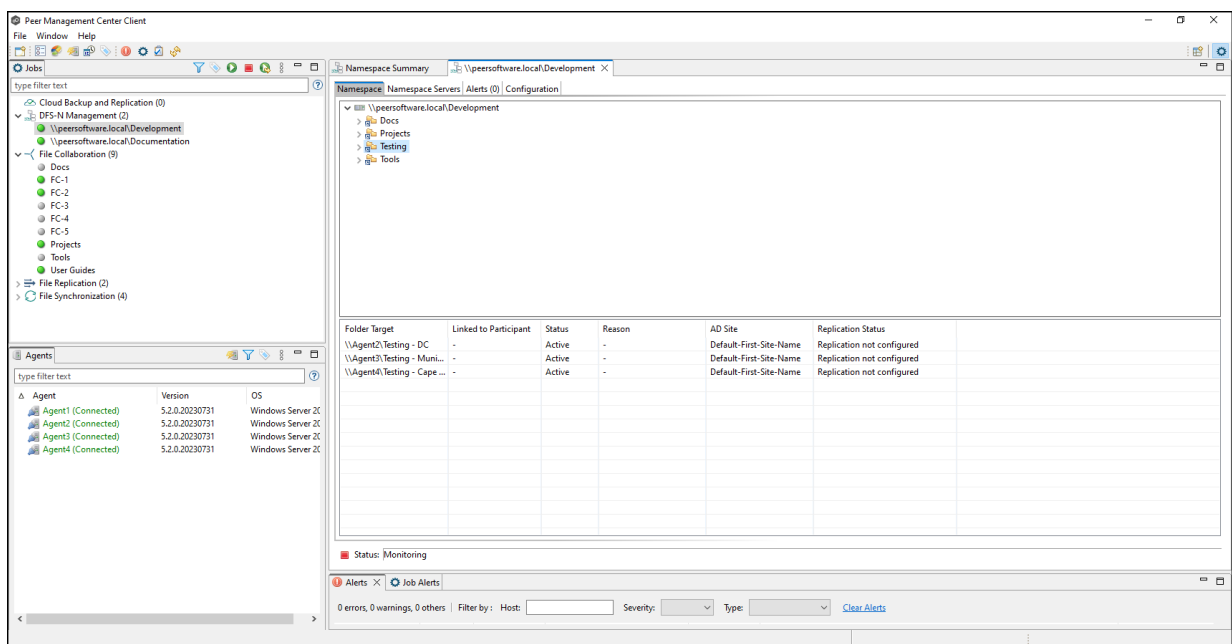
12. Click **Close**.

The runtime view for the job is displayed. The newly added folders are listed in the job's **Namespace** tab.



13. Click the folder you just added.

The newly added folder targets are listed in the **Folder Target** section of the tab. (Depending on how many namespace folders you have, you may need to scroll to view the **Folder Target** section.)



Adding a Folder Target

You can add a folder target for a namespace folder.

Note: A DFS-N Namespace job must be running before you can edit it.

To add a folder target for a namespace folder:


1. Double-click the job name in the **Jobs** view or the [Namespace Summary view](#) to open the runtime view for the job.

The screenshot displays the Peer Management Center Client interface. The main window is titled "Namespace Summary" and shows a table of namespace paths and folders. The table has columns for Namespace Path/Folders, Management Status, State, Errors, Servers, and Total Folders/Targets. The data is as follows:


Namespace Path/Folders	Management Status	State	Errors	Servers	Total Folders/Targets
\\peersoftware.local\Development	Monitoring	Online	0	Agent1.peersoftware.local, PMC.PEERSOFT...	4 Folders 3 Targets
Docs					3 Targets
Projects					3 Targets
Testing					3 Targets
Tools					3 Targets
\\peersoftware.local\Documentation	Monitoring	Online	0	PMC.peersoftware.local	1 Folder 3 Targets
User Guides					3 Targets

The interface also shows a "Jobs" view on the left with a tree structure of jobs, and an "Agents" view at the bottom left showing a list of connected agents. The "Alerts" section at the bottom right shows "0 errors, 0 warnings, 0 others" and includes filters for Host, Severity, and Type, along with a "Clear Alerts" button.

The runtime view for the job is displayed.

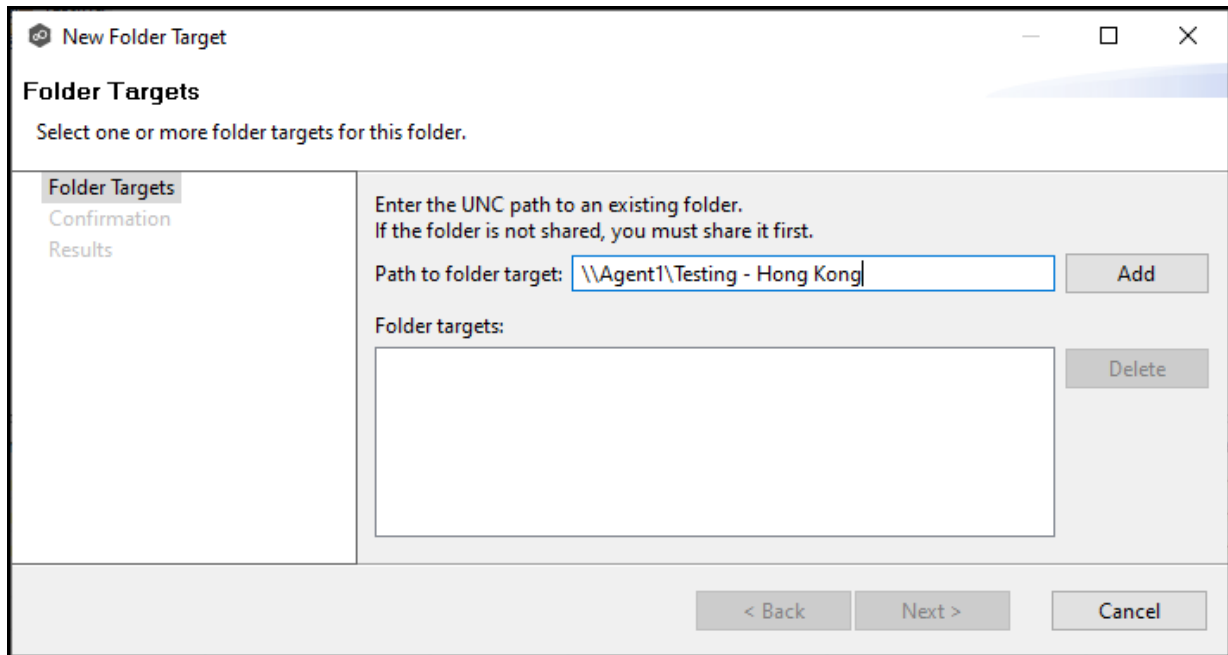


2. Right-click the folder you want to add a folder target to, and then select **Add Folder Target**.



The **New Folder Target** wizard appears.

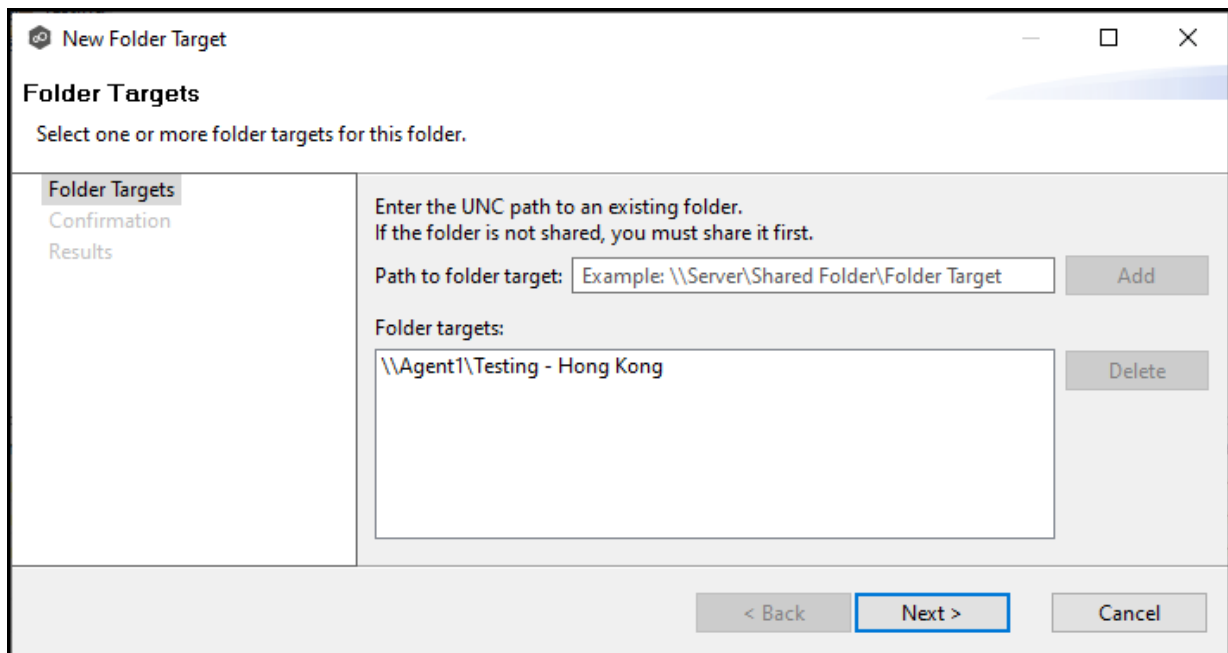
3. Enter the UNC path to a shared folder.



The screenshot shows the 'New Folder Target' dialog box with the 'Folder Targets' tab selected. The dialog has a title bar with a minimize, maximize, and close button. Below the title bar, the text 'Folder Targets' is displayed in bold, followed by the instruction 'Select one or more folder targets for this folder.' The main area is divided into two sections. The left section contains three sub-tabs: 'Folder Targets' (selected), 'Confirmation', and 'Results'. The right section contains the following elements: a text box with the placeholder text 'Enter the UNC path to an existing folder. If the folder is not shared, you must share it first.'; a text input field containing the path '\\Agent1\Testing - Hong Kong'; an 'Add' button; a 'Folder targets:' label; a list box containing the path '\\Agent1\Testing - Hong Kong'; a 'Delete' button; and a '< Back' button at the bottom left.

4. Click **Add**.

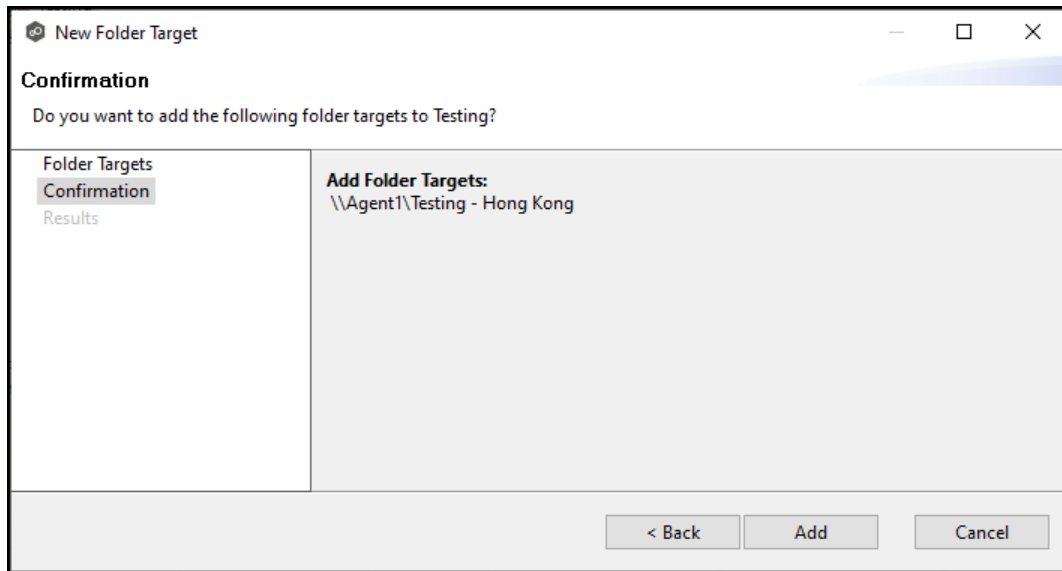
The folder target is added to the **Folder targets** section.



The screenshot shows the 'New Folder Target' dialog box with the 'Folder Targets' tab selected. The dialog has a title bar with a minimize, maximize, and close button. Below the title bar, the text 'Folder Targets' is displayed in bold, followed by the instruction 'Select one or more folder targets for this folder.' The main area is divided into two sections. The left section contains three sub-tabs: 'Folder Targets' (selected), 'Confirmation', and 'Results'. The right section contains the following elements: a text box with the placeholder text 'Enter the UNC path to an existing folder. If the folder is not shared, you must share it first.'; a text input field containing the placeholder text 'Example: \\Server\Shared Folder\Folder Target'; an 'Add' button; a 'Folder targets:' label; a list box containing the path '\\Agent1\Testing - Hong Kong'; a 'Delete' button; and a '< Back' button at the bottom left. The 'Next >' button at the bottom right is highlighted with a blue border.

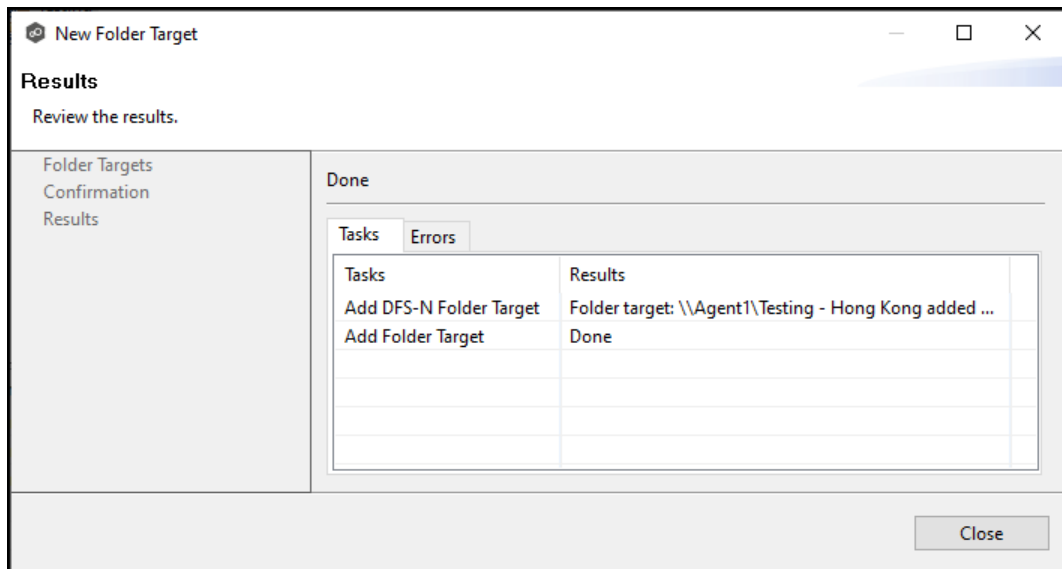
5. Repeat Steps 3-4 to add additional folder targets if desired.
6. Click **Next**.

The **Confirmation** page is displayed.




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

The **Results** page is displayed.




9. Click **Close**.

The runtime view for the job is displayed.



10. Click the folder you just added.

The newly added folder targets are listed in the **Folder Target** section of the job's **Namespace** tab. (Depending on how many namespace folders you have, you may need to scroll to view the **Folder Target** section.)



Linking a DFS Namespace to File Collaboration and File Synchronization Jobs

The primary benefit of using PeerGFS's ability to manage DFS namespaces is that PeerGFS can be configured to automatically disable and enable folder targets when they become unavailable, helping to control [failover and failback](#). This is a manual process using Microsoft DFS Namespaces, but PeerGFS can automatically disable or enable targets based on the state of a linked collaboration/synchronization job. This ensures a folder target is not available to users until a failed server has come back online and is in sync again.

To take advantage of the failover and failback capabilities, the File Collaboration or File Synchronization job must be linked to the DFS-N Management job that manages the namespace.

There are various ways to link a File Collaboration or File Synchronization job to a DFS-N Management job, including:

- If the File Collaboration or File Synchronization job does not yet exist, you can:
 - Create a File Collaboration or File Synchronization job and link it to a namespace while you are creating the job. When creating the job, you are given the option to create a new namespace, import one, or use an existing one. The DFS-N Management job is automatically created when using this method. See [Step 8: DFS Management](#) in [Creating a File Collaboration job](#) or in [Creating a File Synchronization job](#) for more information.
 - Create one from the DFS namespace folder. See [Create a File Collaboration or Synchronization Job from a Namespace Folder](#) for step-by-step instructions.
- If the File Collaboration or File Synchronization job already exists, edit the job, and link it to the DFS-N Management job. Use the [DFS-N settings page in the Edit File Collaboration Job wizard](#) and [DFS-N settings page in the Edit File Synchronization Job wizard](#) to link them. See [Linking a Namespace with an Existing File Collaboration or Synchronization Job](#) for step-by-step instructions.

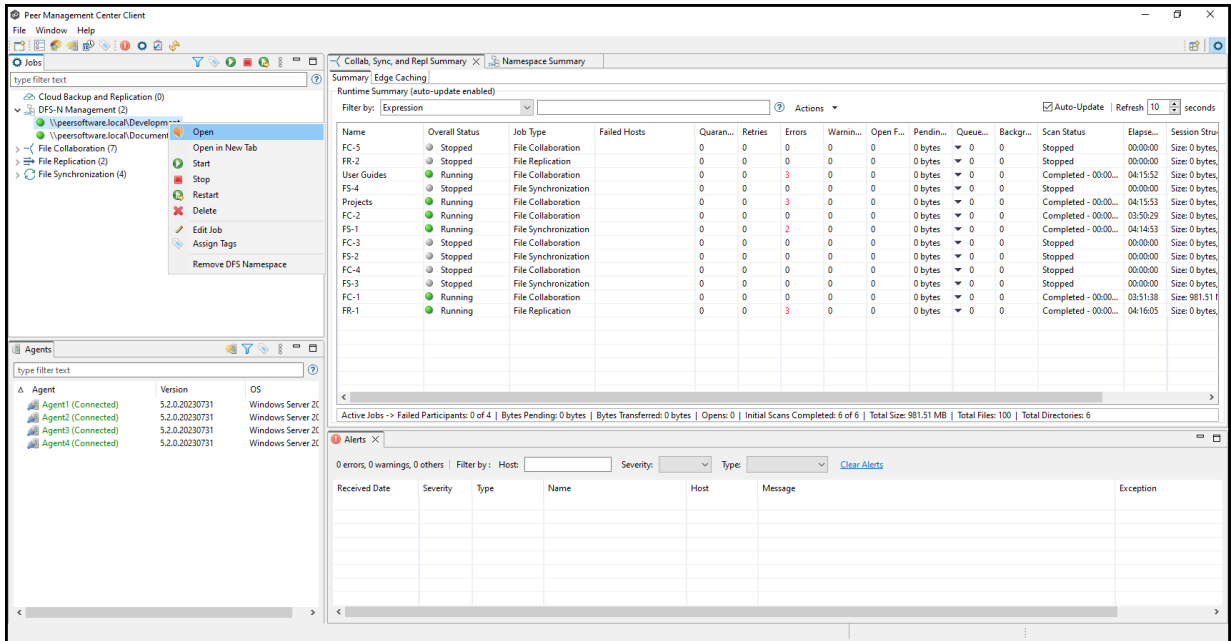
Note: Currently, only File Collaboration and File Synchronization jobs can be linked to a DFS-N Management job.

Creating a File Collaboration or File Synchronization Job from a Namespace Folder

You can create a File Collaboration or File Synchronization job from a DFS namespace folder. These steps require that the DFS namespace has been already created and is being managed by Peer Management Center.

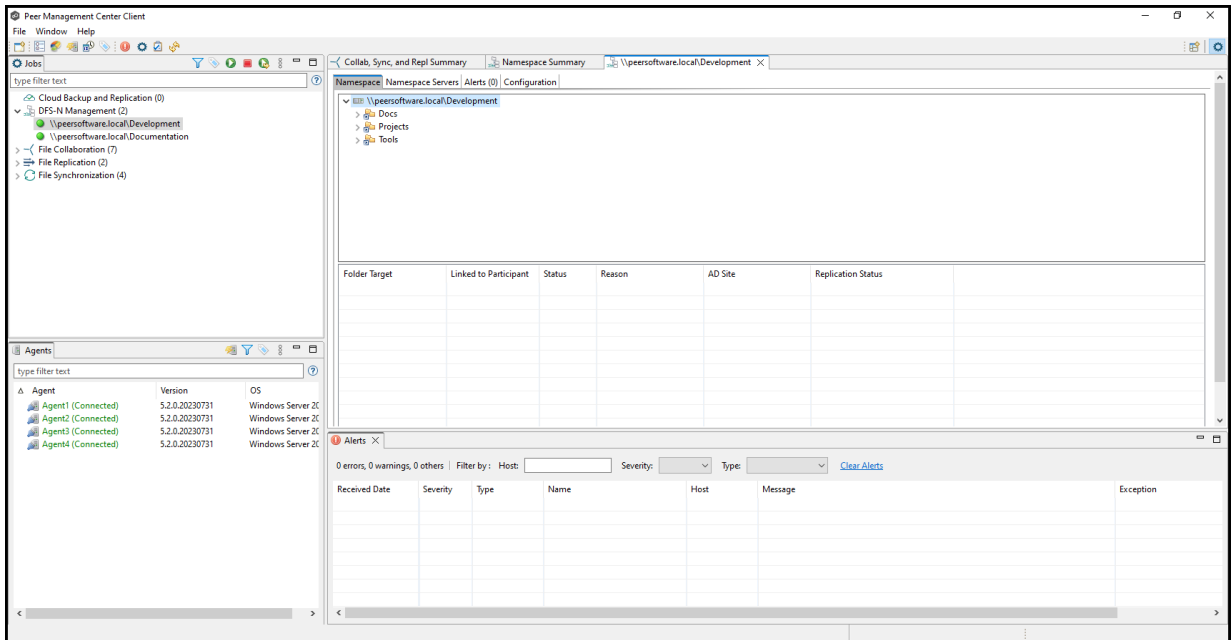
To create a File Collaboration or File Synchronization job from a namespace folder:

1. From the **Jobs** view, open the DFS-N Management job managing the namespace.




The [DFS-N Management Job runtime view](#) is displayed.

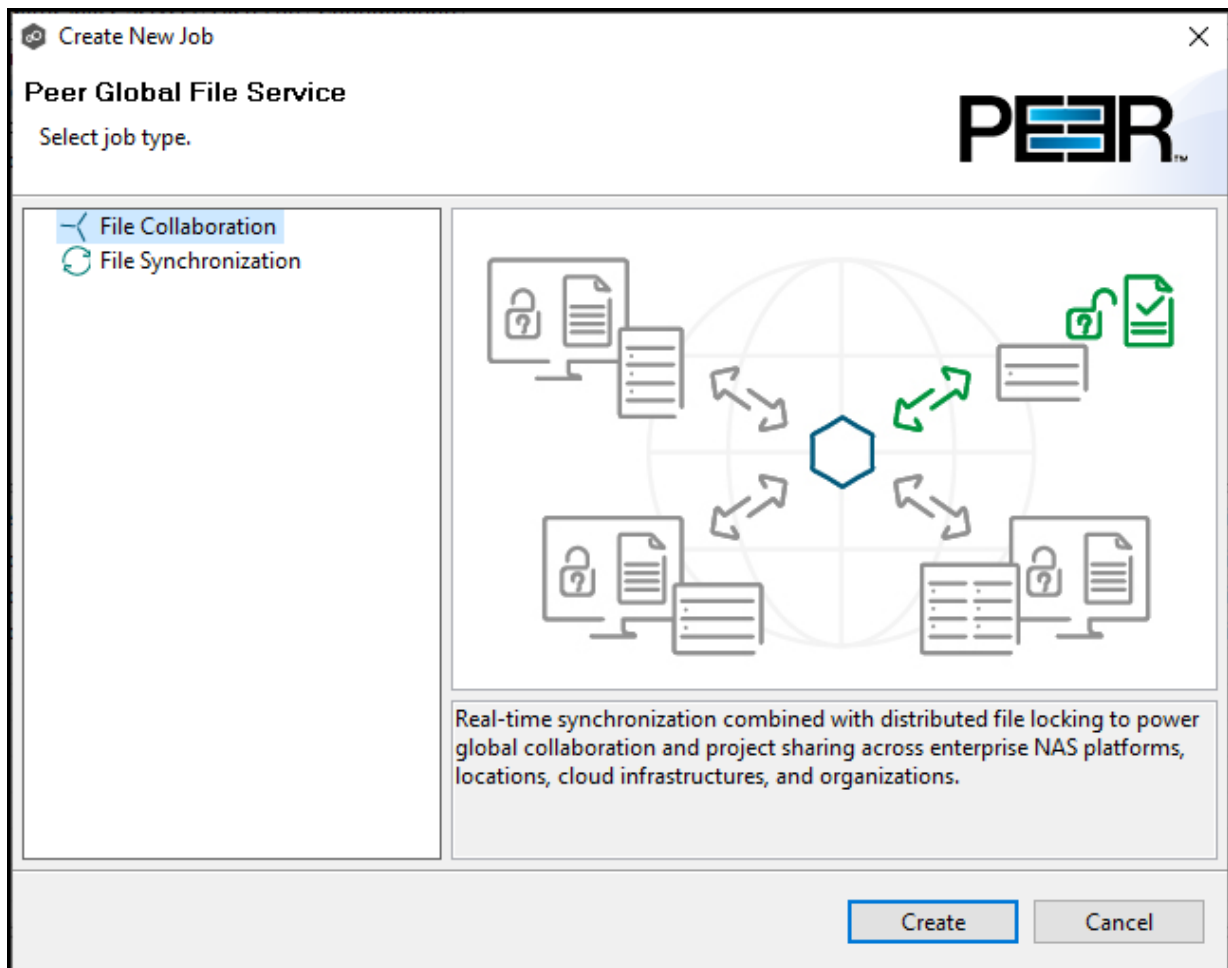
2. Open the **Namespace** tab if it is not already displayed.



3. In the **Namespace** tab, right-click the desired namespace folder and select **Create New Job**.



The **Create New Job** wizard displays a list of job types you can create: File Collaboration and File Synchronization. The other job types are not supported for use with DFS namespace management.

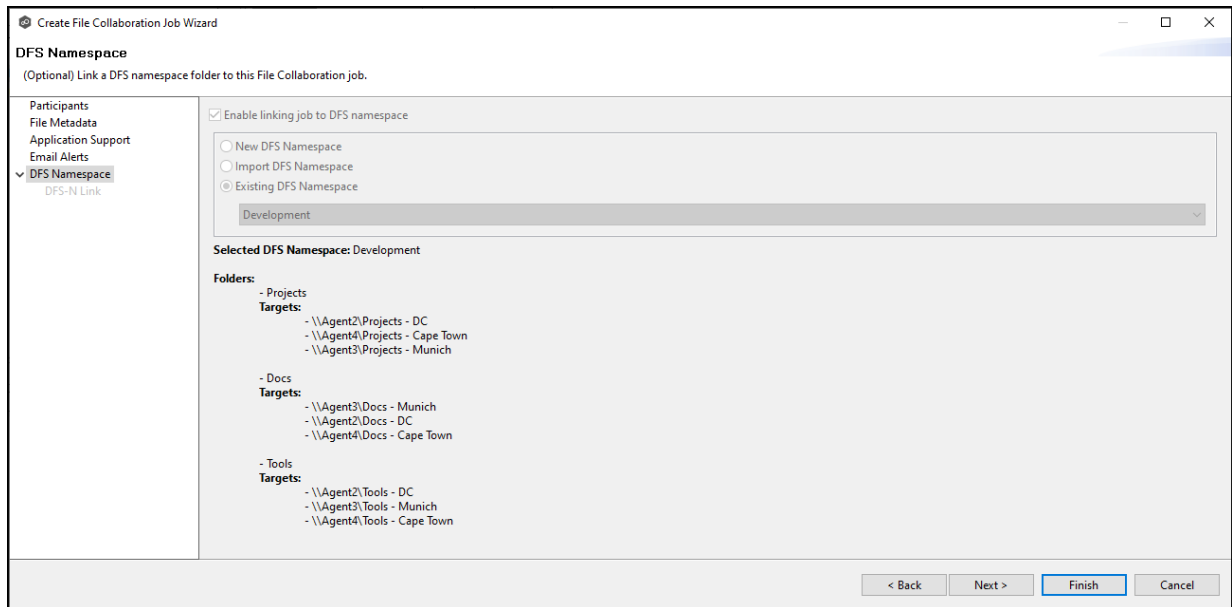


4. Select a job type and click **Create**:
 - Select **File Collaboration** if locking is required in addition to replication (for example, for data sets with shared project files).
 - Select **File Synchronization** if no locking is required (for example, with home directory and user profile datasets).
5. Follow the wizard prompts as it walks you through creating the job.

The process is the same as described in [Creating a File Collaboration Job](#) and [Creating a File Synchronization Job](#).

Once you have selected your email alerts, the **DFS Namespace** page is displayed.

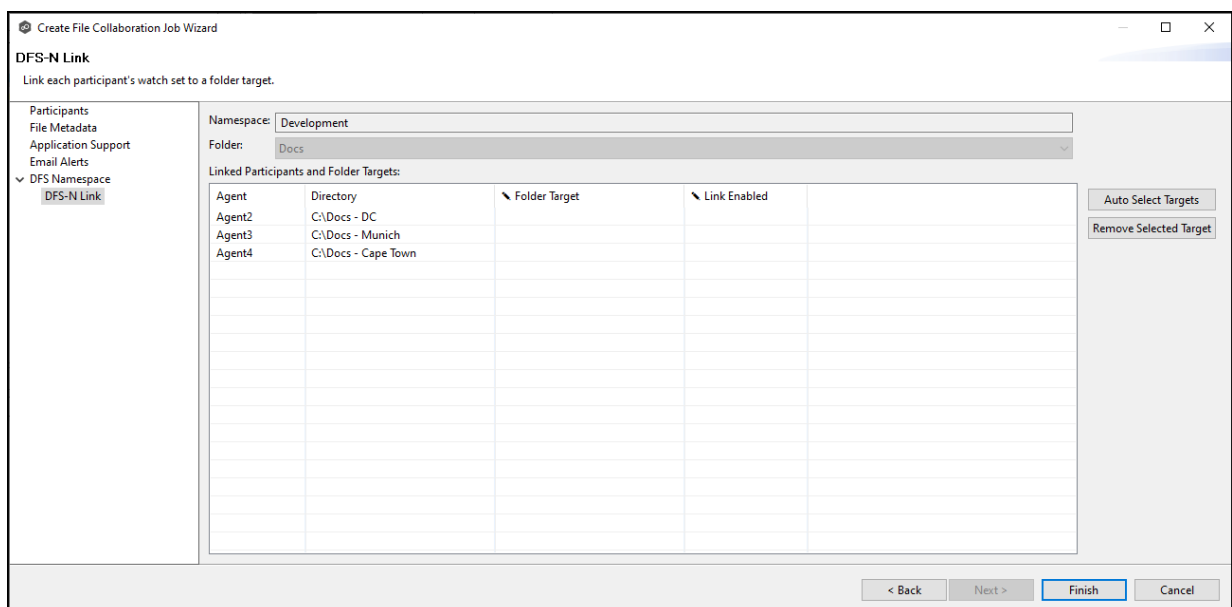
The **Enable linking job to DFS namespace** checkbox is preselected and the **Existing DFS Namespace** option is selected.

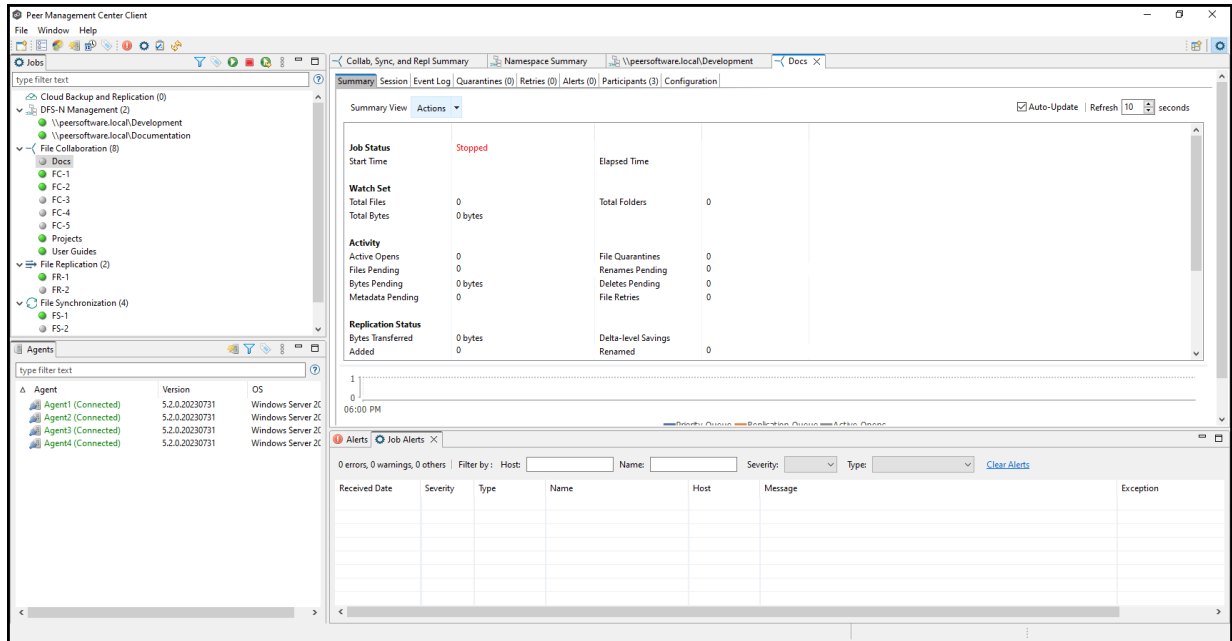


6. Click **Next** if you want to create folder targets; otherwise click **Finish** and continue with Step 9.

You can create folder targets later if you wish. See [Adding a Namespace Folder Target](#) for step-by-step instructions.

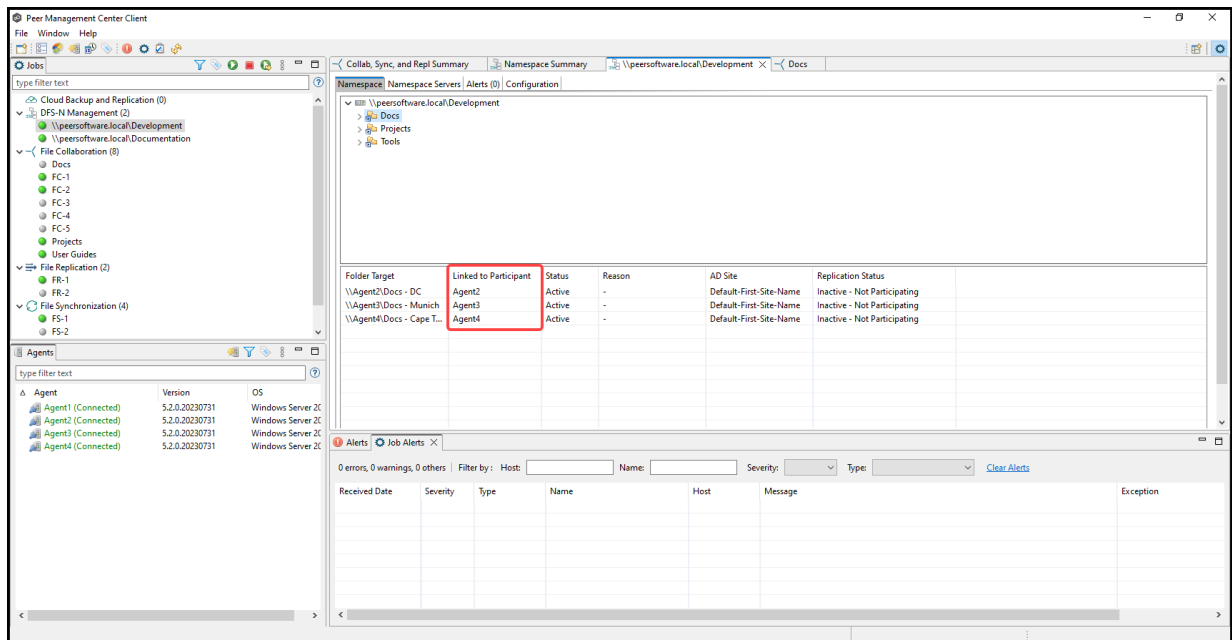
If you clicked **Next**, the **DFS Link** page appears. The purpose of this page is to link the watch path of each participant of a collaboration/synchronization job (specified in the **Path** page) to the appropriate folder target. Initially, the **Folder Target** column in the **Linked Participants and Folder Targets** table will not contain any folder targets. After linking, a folder target should be listed in this column for each participant.





10. Click the DFS-N Management job tab to display its runtime view.

The **Linked to Participant** column is now populated.

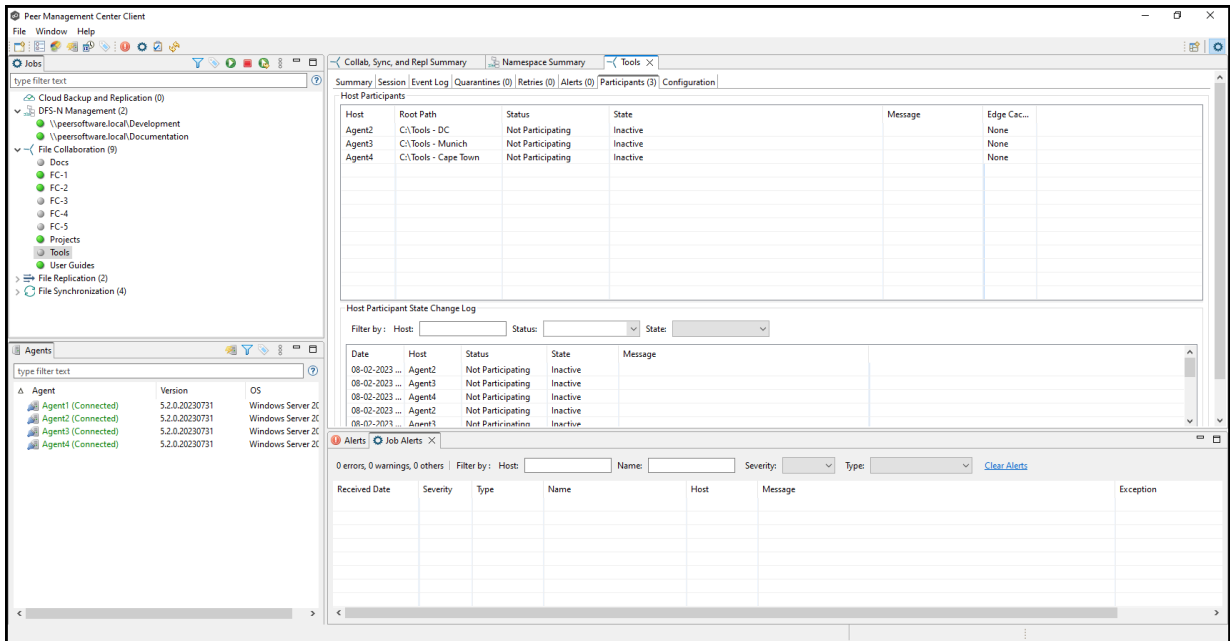


Linking an Existing Namespace Folder to an Existing File Collaboration or File Synchronization Job

You can link an existing DFS namespace folder with an existing File Collaboration or File Synchronization job. These steps require that the DFS namespace has been already created and is being managed by a DFS-N Management job.

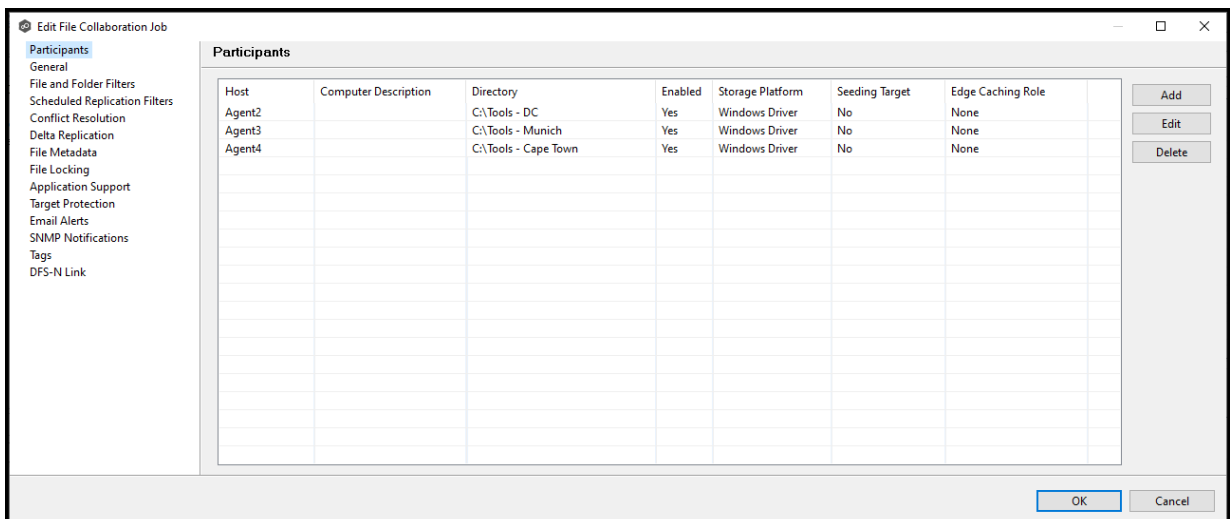
To link a namespace folder with an existing File Collaboration or File Synchronization job:

1. Select the File Collaboration or File Synchronization job in the **Jobs** view.




2. Right-click and select **Edit Job**.

The **Edit Job** wizard appears.



Tip: In most cases, clicking the **Auto Select Targets** button, PeerGFS will be able to automatically link a folder target with the appropriate participant. However, if a folder does not have the appropriate folder targets, click the **Auto Create Targets** button. The wizard that appears will use the paths configured in your File Collaboration or File Synchronization job and try to automatically create folder targets for you.



7. Review the values in the **Link Enabled** column; if **No** appears, select **Yes** from the drop-down list.
8. Once all participants are linked to the appropriate folder targets, click **OK** to save your changes.

The runtime window appears. From this point forward, if this collaboration or synchronization job is running along with its paired DFS-N Management job, Peer Management Center will automatically [failover and failback](#) folder targets.

File Collaboration Jobs

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

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 computer resources, the actual performance is as fast as possible from a physical view. At the same time File Collaboration ensures version conflicts are eliminated with file locking.

Before You Create Your First File Collaboration Job

We strongly recommend that you configure the File Collaboration settings (e.g., SMTP notifications), as well as other [global settings](#) such as SMTP email settings, email alerts, and file filters before configuring your first File Collaboration job. See [Preferences](#) for details on these settings.

Creating a File Collaboration Job

The **Create Job** wizard walks you through the process of creating a File Collaboration job. The process consists of the following steps:

[Step 1: Job Type and Name](#)

[Step 2: Participants](#)

[Step 3: File Metadata](#)

[Step 4: Application Support](#)

[Step 5: Email Alerts](#)

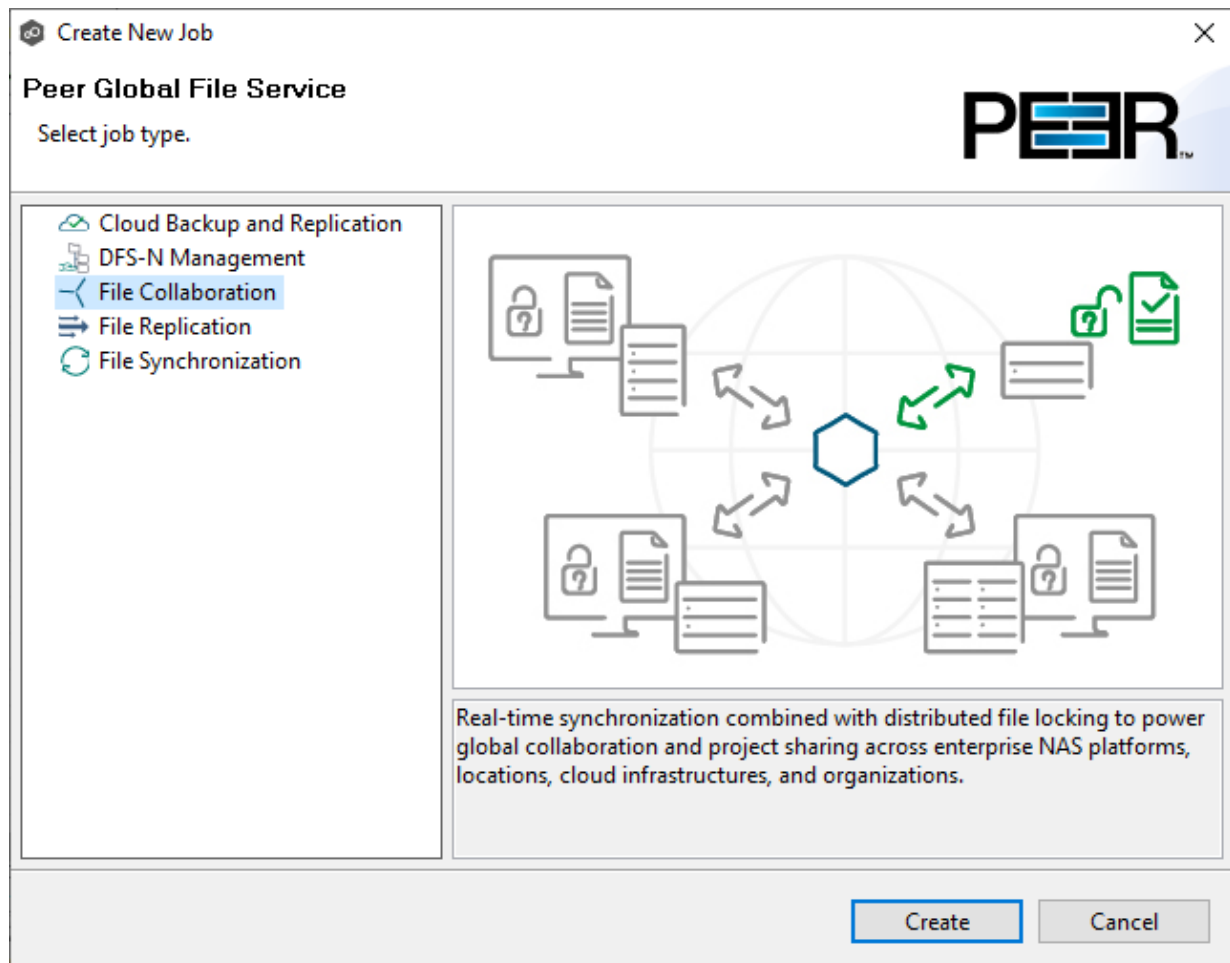
[Step 6: Save Job](#)

Additional configuration options, such as applying [file filters](#) and specifying [delta level replication](#), are available when [editing a File Collaboration job](#).

Step 1: Job Type and Name

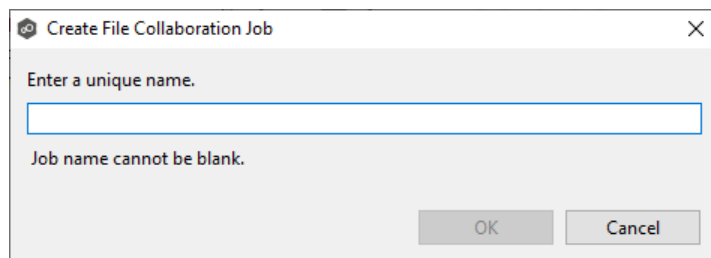
1. Open Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

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



3. Click **File Collaboration**, and then click **Create**.
4. Enter a name for the job in the dialog that appears.

The job name must be unique.



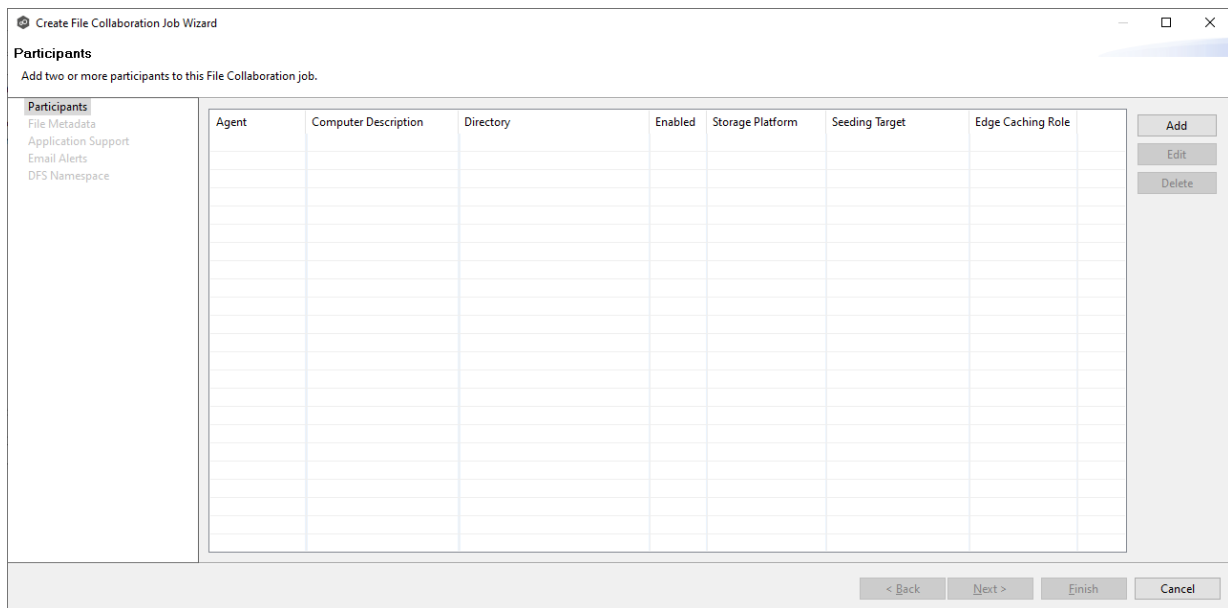
5. Click **OK**.

The [Participants](#) page appears.

Step 2: Participants

After selecting the job type and naming the job, the **Participants** page is displayed. It contains a table that will display the job [participants](#) once you have added them. A File Collaboration job must have two or more participants. A **participant** consists of an Agent and the volume, share, or folder to be replicated. A File Collaboration job synchronizes the files of participants in real-time and adds distributed locking to avoid version conflicts.

1. Click the **Add** button to start the process of adding a participant.



The **Add New Participant** wizard opens; it walks you through the steps for adding a participant:

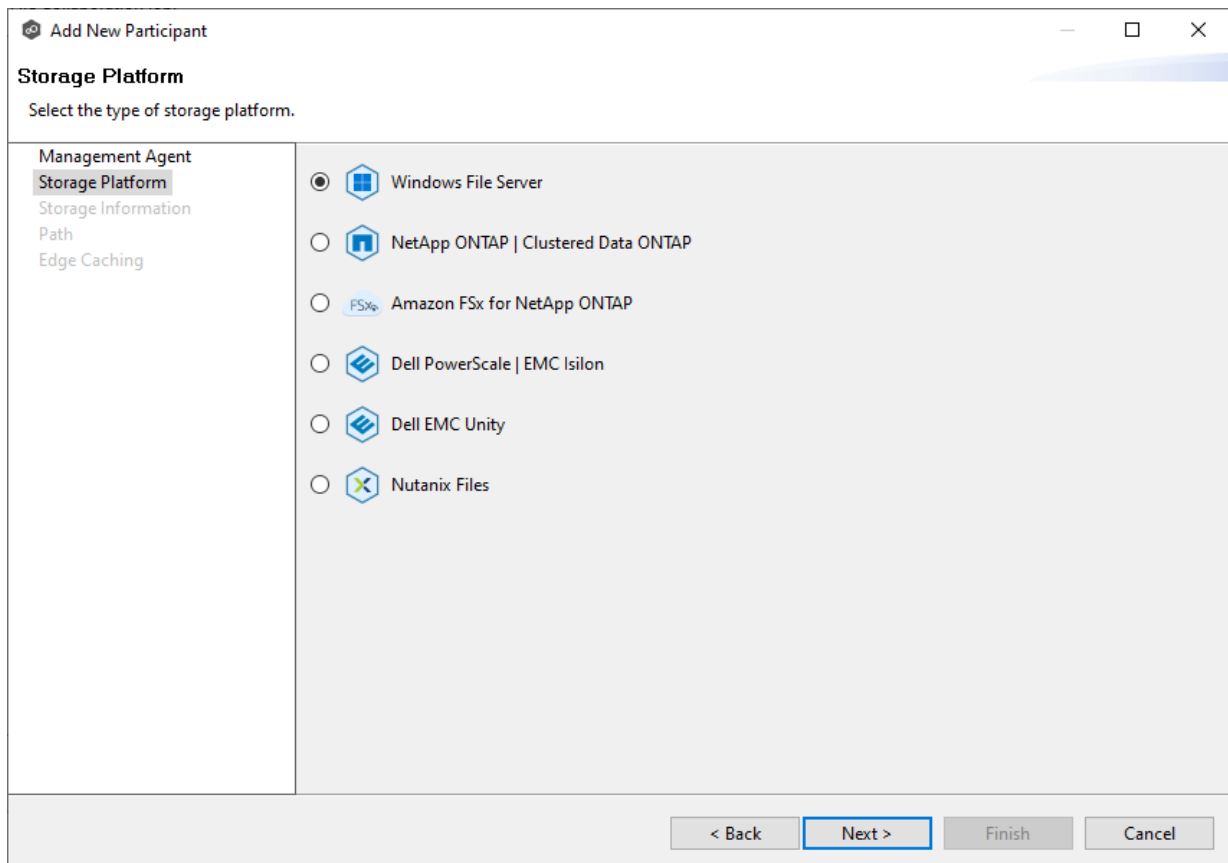
- a. [Selecting a Management Agent](#), which is the Agent that will manage the storage device that hosts data you want to replicate.
- b. [Selecting the type of storage platform](#) that hosts data you want to replicate.
- c. [Entering the credentials needed to access a specific storage device and providing other storage information.](#)
- d. [Entering the path](#) to the [watch set](#) (the data that you want to replicate) and selecting whether participant will be a [seeding target](#).
- e. (Optional) [Enabling Edge Caching](#) for the participant.

Once you have added a participant, it is listed in the **Participants** table.

The [Storage Platform](#) page is displayed.

The **Storage Platform** page lists the available storage platforms types.

1. Select the storage platform type hosting the data you want to replicate.



The screenshot shows a window titled "Add New Participant" with a sub-header "Storage Platform". Below the sub-header is the instruction "Select the type of storage platform." On the left side, there is a navigation pane with the following items: "Management Agent", "Storage Platform" (which is highlighted), "Storage Information", "Path", and "Edge Caching". The main area of the dialog displays a list of storage platform options, each with a radio button and an icon:

- Windows File Server
- NetApp ONTAP | Clustered Data ONTAP
- Amazon FSx for NetApp ONTAP
- Dell PowerScale | EMC Isilon
- Dell EMC Unity
- Nutanix Files

At the bottom of the dialog, there are four buttons: "< Back", "Next >" (which is highlighted with a blue border), "Finish", and "Cancel".

2. Click **Next**.

The [Storage Information](#) page is displayed.

On the **Storage Information** page, you will select the storage device containing data that you want to replicate and enter other information about the storage device. The contents of the **Storage Information** page varies, depending on your selection on the [Storage Platform](#) page:

- If you selected **Windows File Server**, you are prompted for configuration information relating to Windows File Server. Continue with the [Windows File Server](#) page.
- If you selected any other type of storage device other than Windows File Server, the **Storage Information** page requests the credentials necessary to connect to that storage device. Continue with the steps below.

For storage platforms other than Windows File Server:

1. Select **New Credentials** or **Existing Credentials**.
2. If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Path](#) page.

If you selected **New Credentials**, enter the credentials for connecting to the storage device. The information you are prompted to enter varies, depending on the type of storage platform:

[Amazon FSxN](#)

[Dell PowerScale](#)

[Dell Unity](#)

[NetApp ONTAP](#)

[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.

Amazon FSxN

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

The screenshot shows a window titled "Add New Participant" with a "Storage Information" tab selected. The window contains a sidebar with navigation options: "Management Agent", "Storage Platform", "Storage Information" (highlighted), "Path", and "Edge Caching". The main area is titled "Storage Information" and contains the instruction "Enter the information required to connect to the storage device." Under the "Credentials" section, there are two radio buttons: "New Credentials" (selected) and "Existing Credentials". The "New Credentials" section includes five input fields: "*SVM Name:", "*SVM User Name:", "*SVM Password:", "SVM Management IP:", and "*Peer Agent IP:". An "Advanced" button is located to the right of the "*Peer Agent IP" field. Below the "Existing Credentials" radio button is a drop-down menu. At the bottom of the main area is a "Validate" button. A footer message reads: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Amazon FSxN environments." At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

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**, supply the required information.

Field	Description
SVM Name	Enter the name, FQDN, or IP address of the Storage Virtual Machine (SVM) hosting the data to be replicated.
SVM User Name	Enter the user name for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	Optional. Enter the IP address used to access the management API of the Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already permit management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually enter the address.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

5. Click **Next**.

The [Path](#) page is displayed.

Dell PowerScale

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect the PowerScale cluster hosting the data to be replicated.

The information required will vary, depending on whether you select **Syslog** or **RabbitMQ** as the connection type, due to the distinct protocols and mechanisms they employ for communication.

Syslog

Add New Participant

Storage Information

Enter the information required to connect to the storage device.

Management Agent
Storage Platform
Storage Information
Path
Edge Caching

Credentials

New Credentials

*Cluster Name:

*Cluster Management IP:

*Cluster Username:

*Cluster Password:

Cluster Access Zone:

Connection Type: Syslog RabbitMQ

Syslog

*Agent IP Address:

*Listening Port:

*SSL Certificate Path:

*SSL Private Key Path:

SSL Private Key Password:

Existing Credentials

You must enter a Cluster Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments.

RabbitMQ

The screenshot shows a window titled "Add New Participant" with a sidebar on the left containing "Management Agent", "Storage Platform", "Storage Information" (selected), "Path", and "Edge Caching". The main area is titled "Storage Information" and contains the instruction "Enter the information required to connect to the storage device." Under the "Credentials" section, there are two radio buttons: "New Credentials" (selected) and "Existing Credentials". The "New Credentials" section includes five text input fields: "*Cluster Name:", "*Cluster Management IP:", "*Cluster Username:", "*Cluster Password:", and "Cluster Access Zone:". Below these is a "Connection Type" section with radio buttons for "Syslog" and "RabbitMQ" (selected). An "Advanced" button is located to the right of the "Connection Type" section. The "Existing Credentials" section has a drop-down menu. A "Validate" button is at the bottom left of the form area. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel". A note at the bottom of the form area reads: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments."

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**, supply the required information.

Field	Description
Cluster Name	Enter the name of the PowerScale cluster hosting the data to be replicated.
Cluster Management IP	Enter the IP address to use to access the REST-based API integrated into the PowerScale cluster. Required only if multiple Access Zones are in use on the cluster.
Cluster Username	Enter the user name for the account managing the PowerScale cluster.
Cluster Password	Enter the password for account managing the PowerScale cluster.
Cluster Access Zone	Optional. The name of the access zone that is being monitored.
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: <ul style="list-style-type: none"> • Opt for Syslog if the storage device directly transmits notifications to the Agent. • Opt for RabbitMQ if you're utilizing the CEE framework to dispatch notifications to the Agent.

3. If you selected **Syslog**, you will need to provide values for the following fields:

Field	Description
Agent IP Address	Select the IP address of the server hosting the Agent that manages the PowerScale cluster. The cluster must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.
Listening Port	Enter the port over which the Agent will receive TLS-based syslog events from the PowerScale cluster.

Field	Description
SSL Certificate Path	Enter the path to the certificate to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private Key Path	Enter the path to the private key to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private Key Password	[Optional] If your private key is protected with a password, enter it here. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .

4. Click **Advanced** if you want to set [advanced options](#).
5. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

6. Click **Next**.

The [Path](#) page is displayed.

Dell Unity

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the NAS server hosting the data to be replicated.

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. The main heading is "Storage Information" with a sub-heading "Enter the information required to connect to the storage device." On the left is a navigation pane with "Storage Information" selected. The main area has two radio buttons: "New Credentials" (selected) and "Existing Credentials". Under "New Credentials" are four text boxes for "*CIFS Server Name:", "*Unisphere Management IP:", "*Unisphere Username:", and "*Unisphere Password:". An "Advanced" button is to the right. Below the "Existing Credentials" radio button is a drop-down menu. At the bottom, a "Validate" button is followed by the error message "You must enter a CIFS Server Name." Below that is a link: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments." At the very bottom are four buttons: "< Back", "Next >", "Finish", and "Cancel".

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**, supply the required information.

Field	Description
CIFS Server Name	Enter the name of the NAS server hosting the data to be replicated.
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.
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.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

5. Click **Next**.

The [Path](#) page is displayed.

NetApp ONTAP

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. The main heading is "Storage Information" with a sub-heading "Enter the information required to connect to the storage device." On the left is a navigation pane with "Storage Information" selected. The main area has two radio buttons: "New Credentials" (selected) and "Existing Credentials". Under "New Credentials" are five input fields: "*SVM Name:", "*SVM User Name:", "*SVM Password:", "SVM Management IP:", and "*Peer Agent IP:". An "Advanced" button is to the right of the last field. Under "Existing Credentials" is a drop-down menu. A "Validate" button is at the bottom left. A footer contains "< Back", "Next >", "Finish", and "Cancel" buttons. A note at the bottom says: "Having trouble connecting? Please verify that all [prerequisites](#) are met for NetApp ONTAP environments."

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**, supply the required information.

Field	Description
SVM Name	Enter the name, FQDN, or IP address of the Storage Virtual Machine (SVM) hosting the data to be replicated.
SVM User Name	Enter the user name for the account managing the Storage Virtual Machine. The account must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. The account must not be a cluster management account.
SVM Management IP	Optional. Enter the IP address used to access the management API of the Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already permit management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

5. Click **Next**.

The [Path](#) page is displayed.

Nutanix Files

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated.

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. The main heading is "Storage Information" with a sub-heading "Enter the information required to connect to the storage device." On the left is a navigation pane with "Storage Information" selected. The main area has two radio buttons: "New Credentials" (selected) and "Existing Credentials". Under "New Credentials" are four input fields: "*Nutanix File Server Name:", "*Username:", "*Password:", and "*Peer Agent IP:". There is an "Advanced" button to the right of the "Peer Agent IP" field. Under "Existing Credentials" is a drop-down menu. At the bottom left is a "Validate" button. At the bottom right are four buttons: "< Back", "Next >", "Finish", and "Cancel". A note at the bottom reads: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Nutanix Files environments."

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**, supply the required information.

Field	Description
Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Enter the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually enter the address. The IP address should not point to the Files cluster itself.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials.

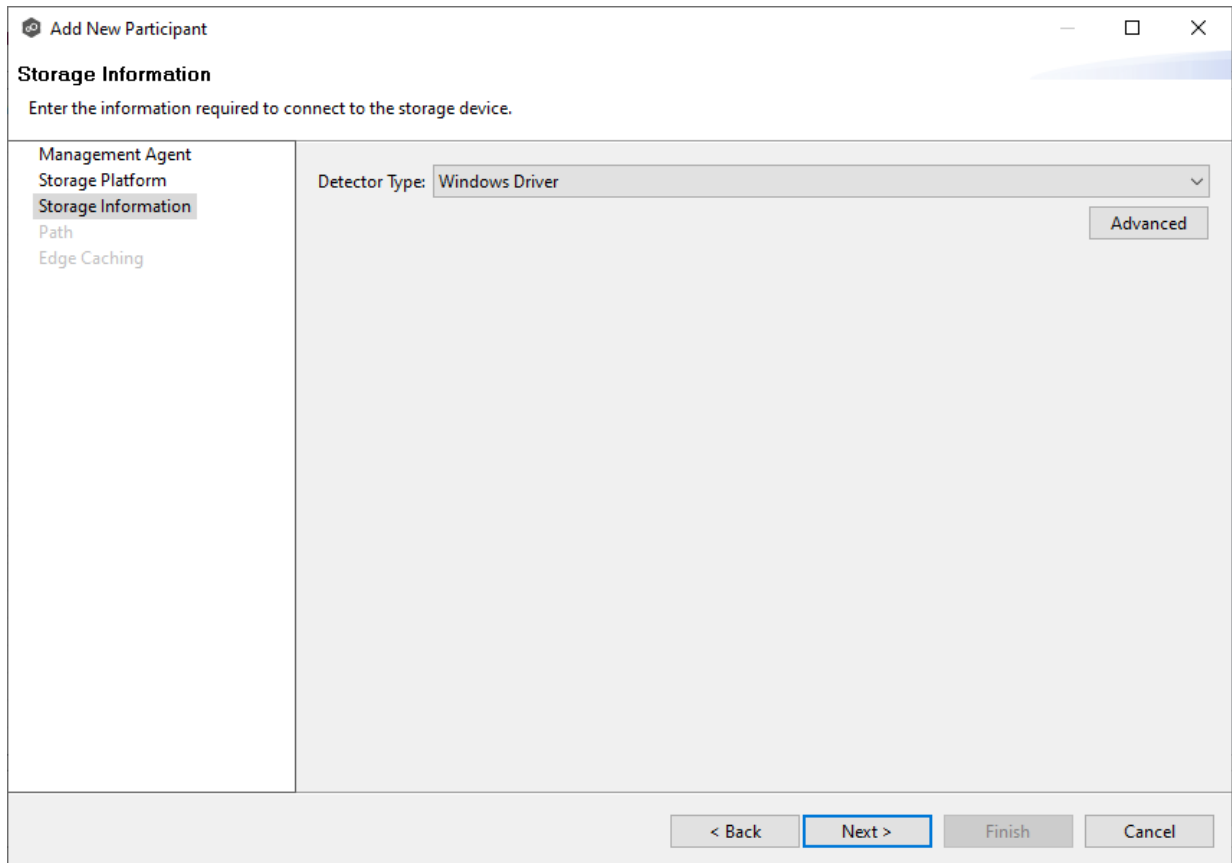
After the credentials are validated, a success message appears.

5. Click **Next**.

The [Path](#) page is displayed.

Windows File Server

1. Select the **Detector Type**:
 - Select **Windows Driver** for more robust logging and better performance (Recommended).
 - Select **Windows** if suggested by Peer Technical Support.




2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Next**.

The [Path](#) page is displayed.

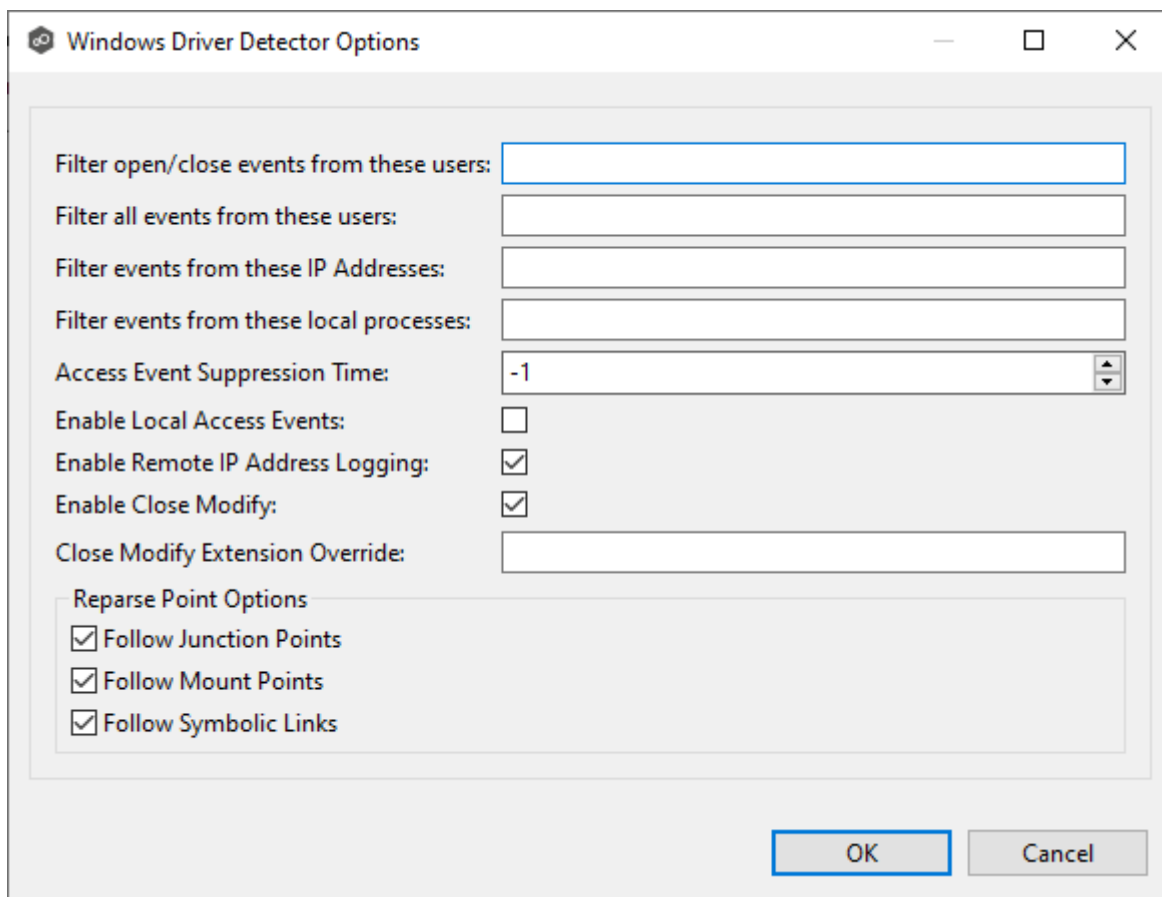
1. Modify the options as desired.

The available options depend on the detector type selected: **Windows** or **Windows Driver**.

Windows



Windows Driver



Option	Description
Filter open/close events from these users	Enter a comma-separated list of user account names from which all file opens and closes will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user name under which a backup and/or archival service is running.
Filter all events from these users	Enter a comma-separated list of user account names from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user name under which a backup and/or archival service is running.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the IP addresses on which a backup and/or archival service is running.
Filter events from these local processes	Enter a comma-separated list of local process names on the Agent server from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the specific process names under which a backup and/or archival service is running.
Access Event Suppression Time	Enter the number of seconds to delay an open event before being processed. Use this option to help reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time an open event will be delayed based on the load of the system.
Enable Local Access Events	Enable tracking of opens and closes that are performed locally on the Agent server.
Enable Remote IP Logging	Enable logging of client IP addresses for all real-time activity.
Enable Close Modify	When enabled, no modify or write events will be detected. Instead, replication of a modified file will be performed when the file is closed.
Close Modify Extension Override	Enter a comma-separated list of exclusions for the Enable Close Modify option. All modify/write events will be detected for these files. This is important for those who rely on sync-on-save functionality.

For more information about junction points or symbolic links, contact [<% SUPPORT_EMAIL%](#)

2. Click **OK**.

The **Path** page is where you specify the path to the volume, share, or folder you want to replicate. This volume, share, or folder is referred to as the [watch set](#). The watch set can contain only a single volume, share, or folder. If you want to replicate multiple volumes, shares, or folders, you need to create a separate job for each one.

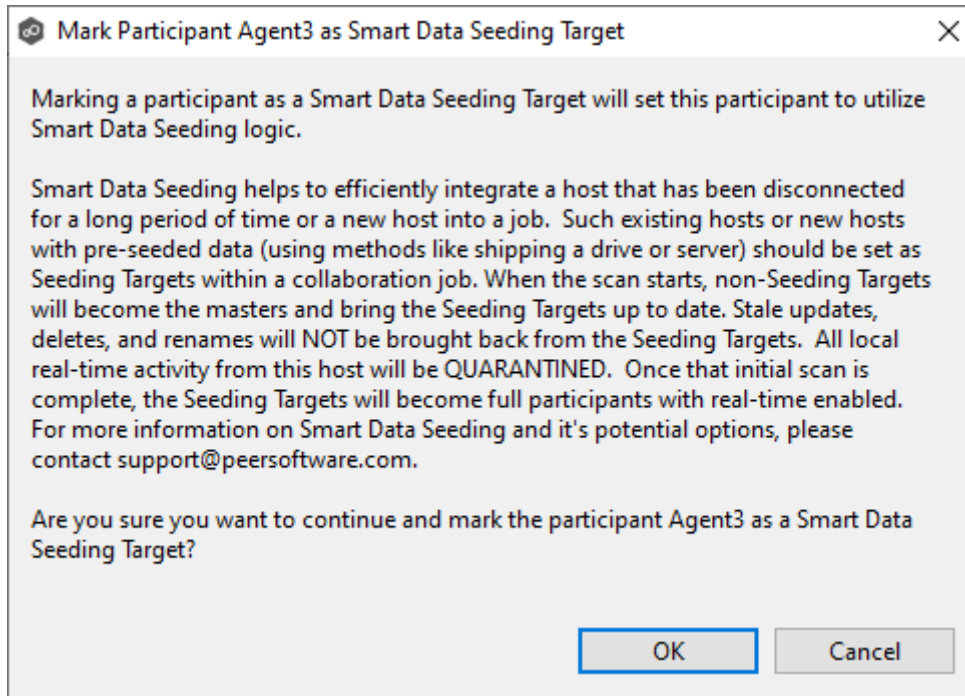
1. Browse to or enter the path to the watch set.

The screenshot shows a dialog box titled "Add New Participant" with a close button (X) in the top right corner. The dialog is divided into two main sections. On the left is a vertical navigation pane with the following items: "Management Agent", "Storage Platform", "Storage Information", "Path" (which is highlighted with a blue background), and "Edge Caching". The main area on the right is titled "Path" and contains a red error message: "You must enter a valid path." Below this message is a text input field labeled "Enter Path" and a "Browse" button to its right. Underneath the input field is a checkbox labeled "Seeding Target". At the bottom of the dialog, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. (Optional) Select the **Seeding Target** checkbox, and then click **OK** in the dialog that appears.

If you select this option, a message describing seeding behavior is displayed. Multiple participants in a File Collaboration job can be set as smart data seeding targets;

however, at least one participant should not be set as a smart data seeding target. All participants that are not set as seeding targets will become sources for the smart data seeding targets. For more information about smart data seeding, see [Smart Data Seeding](#) or contact [Peer Support](#).



3. Click **Next** if you want this participant to use [Edge Caching](#); otherwise, click **Finish** to complete the wizard for this participant.

If you click **Finish**, return to [Step 2: Participants](#) to add more participants, if applicable. A File Collaboration job must have at least two participants.

Edge Caching is a method for conserving space on storage devices by caching files until needed. Edge Caching saves space by stubbing files and rehydrating them as needed. Edge Caching is optional; if you don't need to conserve space on the storage device managed by the Agent, then you do not need to select this option.

If you enable [Edge Caching](#) for a participant, you must designate the participant as either a **master** or **edge** participant.

- **Master participant** - A master participant always has a complete set of files for that job. None of the files are stubbed; they are stored physically on that device.

- **Edge participant** - A subset of the files stored on a master participant are physically stored on an edge participant; the rest of the files on an edge participant are stub files that take up minimal space. Edge Caching allows users to seamlessly retrieve stubbed files directly from a master participant as needed; when retrieved, the local stub file is rehydrated so that the full file is stored locally on the edge participant.

A job can have master and edge participants, as well as participants that don't have either role. If you do not choose to enable Edge Caching for a participant, it will always have a full set of files like a master participant but will not be used to serve file content to any edge participants.

Notes:

- A participant can be a master participant for some jobs and an edge participant for other jobs.
- A job needs at least one master participant that isn't a seeding target. If there is only one master participant for the job, it should not be a seeding target.

For more information about Edge Caching, see [Edge Caching](#) in [Advanced Topics](#).

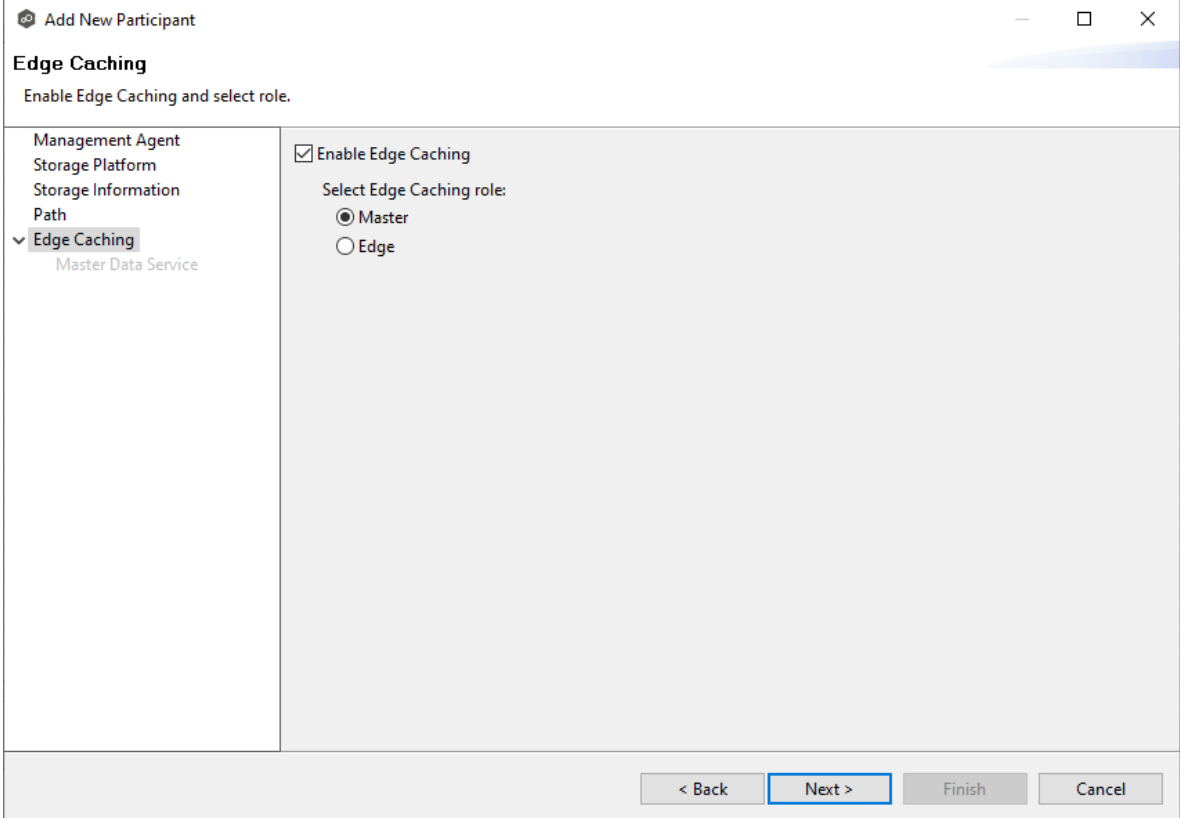
1. Select the **Enable Edge Caching** checkbox if you want this participant to be able to use Edge Caching; otherwise, click **Finish**.

The screenshot shows a window titled "Add New Participant" with a sub-header "Edge Caching" and the instruction "Enable Edge Caching and select role." On the left, a list of options includes "Management Agent", "Storage Platform", "Storage Information", "Path", and "Edge Caching", with "Edge Caching" highlighted. The main content area contains a checkbox labeled "Enable Edge Caching" which is currently unchecked. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

If you enable Edge Caching, the Edge Caching role options are displayed; the **Master** role is selected by default.

2. Choose an Edge Caching role for the participant:

- Choose **Master** if the storage device managed by the Agent will contain complete copies of all files for this job. Any type of storage platform can be a master participant.



The screenshot shows a window titled "Add New Participant" with a close button (X) in the top right corner. The main heading is "Edge Caching" with the instruction "Enable Edge Caching and select role." Below this, there is a tree view on the left with the following items: "Management Agent", "Storage Platform", "Storage Information", "Path", "Edge Caching" (expanded), and "Master Data Service". The "Edge Caching" section is active, showing a checked box for "Enable Edge Caching" and a "Select Edge Caching role:" section with two radio buttons: "Master" (selected) and "Edge". At the bottom of the window, there are four buttons: "< Back", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

- Choose **Edge** if you want to conserve space on the storage device managed by the Agent. Only Windows File Servers can be an edge participant.

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. The main heading is "Edge Caching" with the instruction "Enable Edge Caching and select role." On the left, a navigation pane lists several options: "Management Agent", "Storage Platform", "Storage Information", "Path", "Edge Caching" (which is expanded to show "Volume Policy", "Utilization Policy", and "Pinning Filter"), and "Edge Caching" (which is selected). The main area contains a checked checkbox for "Enable Edge Caching" and a section titled "Select Edge Caching role:" with two radio buttons: "Master" (unselected) and "Edge" (selected). At the bottom right, there are four buttons: "< Back", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

3. Click **Next**.

- If you selected **Master**, continue with the [Master Data Service](#) page.
- If you selected **Edge**, continue with the [Volume Policy](#) page.

Master Data Service

The **Master Data Service** page appears if you chose the master role for the participant. The Master Data Service handles requests from edge participants for files on a master participant. The Master Data Service is installed on the Agent server as part of the Peer Agent installation process.

The first two fields on this page are automatically populated:

- **Protocol:** This field lists the protocol that will be used to transfer file content between master participants and edge participants. HTTPS is currently the only available option as it requires only a single open firewall port on each master participant and does a better job of handling latency over WAN-based connections.

- **Agent Name:** This field lists the name of the Management Agent that you selected at the beginning of Step 2.

1. (Optional) Enter a value for **Agent Alias**. The value can be a hostname, FDQN, or IP address.

A value for this field is required only if the name of the Agent cannot be converted to an IP address via DNS. If an alias is entered, it will be used by the Edge Service on edge participants to connect with the Master Data Service. If no alias is entered, the Agent's name will be used.

2. (Optional) Modify the port number that the Master Data Service will listen on for this master participant.

A default value for the port number, 8446, is set when the Agent is installed. If you modify the port number, the Master Data Service is started with the new port number.

Add New Participant

Master Data Service

Configure access to the Master Data Service.

Management Agent
Storage Platform
Storage Information
Path
Edge Caching
Master Data Service

Protocol: HTTPS

Agent Name: Agent2

Agent Alias:

Port: 8446

< Back Next > Finish Cancel

Note: If the Agent you selected is already being used as a master participant in another job utilizing Edge Caching, then the existing Master Data Service parameters will be displayed. You can edit the values by clicking the **Edit Master Data Service** link. If you modify the port number, the Master Data Service will be restarted, and the new port number will take effect immediately. Any modifications you apply will be applied to every other job that uses this Agent as a master participant.

4. Continue adding more participants if applicable or continue with [Step 4: Master-Edge Assignment](#).

Volume Policy

The **Volume Policy** page appears if you chose the edge role for the participant.

A volume policy is applied when a caching scan is run. The primary purpose of a **volume policy** is to specify how much space is available to Edge Caching on a specific volume (or drive letter), i.e., to define the **cache size**. The cache size specifies the maximum amount of disk space you want to allocate to Edge Caching for fully hydrated files on the volume specified by the path on the **Path** page. For example, if the participant is configured to monitor D:\Data, the volume policy for this participant would apply to the D volume.

The cache size can be specified as a percentage of the volume disk space or as a fixed size. For example, if an edge participant is configured to monitor a volume that has 1 TB of disk space, and you tell Edge Caching to use 75% of that volume, then up to 750 GB of files could be locally available on the volume monitored by that edge participant. For optimal performance, we recommend that this cache be dedicated to Edge Caching's use on this volume.

A volume policy applies to each job where the following three elements are true:

- Edge Caching is enabled for the job.
- The participant is an edge participant.
- The paths specified for each job share the same volume.

To create a volume policy:

1. In the **Cache Size** section, choose an option for setting the cache size:
 - Use up to X % of this volume
 - Use up to X size of this volume

Add New Participant

Volume Policy

Please enter a valid path

Management Agent
Storage Platform
Storage Information
Path
Edge Caching
Volume Policy
Utilization Policy
Pinning Filter

Cache Size

Use up to 75 % of this volume

Use up to 10 GB of this volume

Cache Threshold Alerts

Send alerts when:

Disk space is less than 512 MB

Cache usage exceeds 80 % of the cache size

Caching Scan Schedule

Scan every 1 day(s) at 10:00:00 PM

Define schedule

Run caching scan after job start

*Temporary Storage Path: Browse

\PeerTempPath

< Back Next > Finish Cancel

- In the **Cache Threshold Alerts** section, set threshold values for automatic alerts about free disk space and cache usage.

Peer Management Center will automatically display alerts in the **Alerts** tab:

- The amount of free disk space on the volume falls below the specified value. For example, if a 1 TB volume has 500 MB of free space and the threshold is set to 512 MB, an alert will be sent.
- Cache usage on the volume exceeds the specified percentage of the cache size. For example, if the cache size is set to 80%, equating to 750 GB, Edge Caching will start sending alerts when it has used 600 GB.

You can also send cache threshold alerts via [email alerts](#) and [SNMP notifications](#). You configure these in [Edge Caching](#) preferences for File Collaboration and File Synchronization jobs.

- In the **Caching Scan Schedule**, set the frequency that the volume should be scanned to optimize the balance of fully hydrated vs stubbed files.

This scan can be run daily at a specified time, or you can define a more customized schedule.

You can also set whether or not to run a caching scan after a job is started by checking or unchecking **Run caching scan after job start**.

4. Select **Run caching scan** after job start.
5. In the **Temporary Storage Path** field, enter a path or browse to the location you want to be used as temporary storage space.

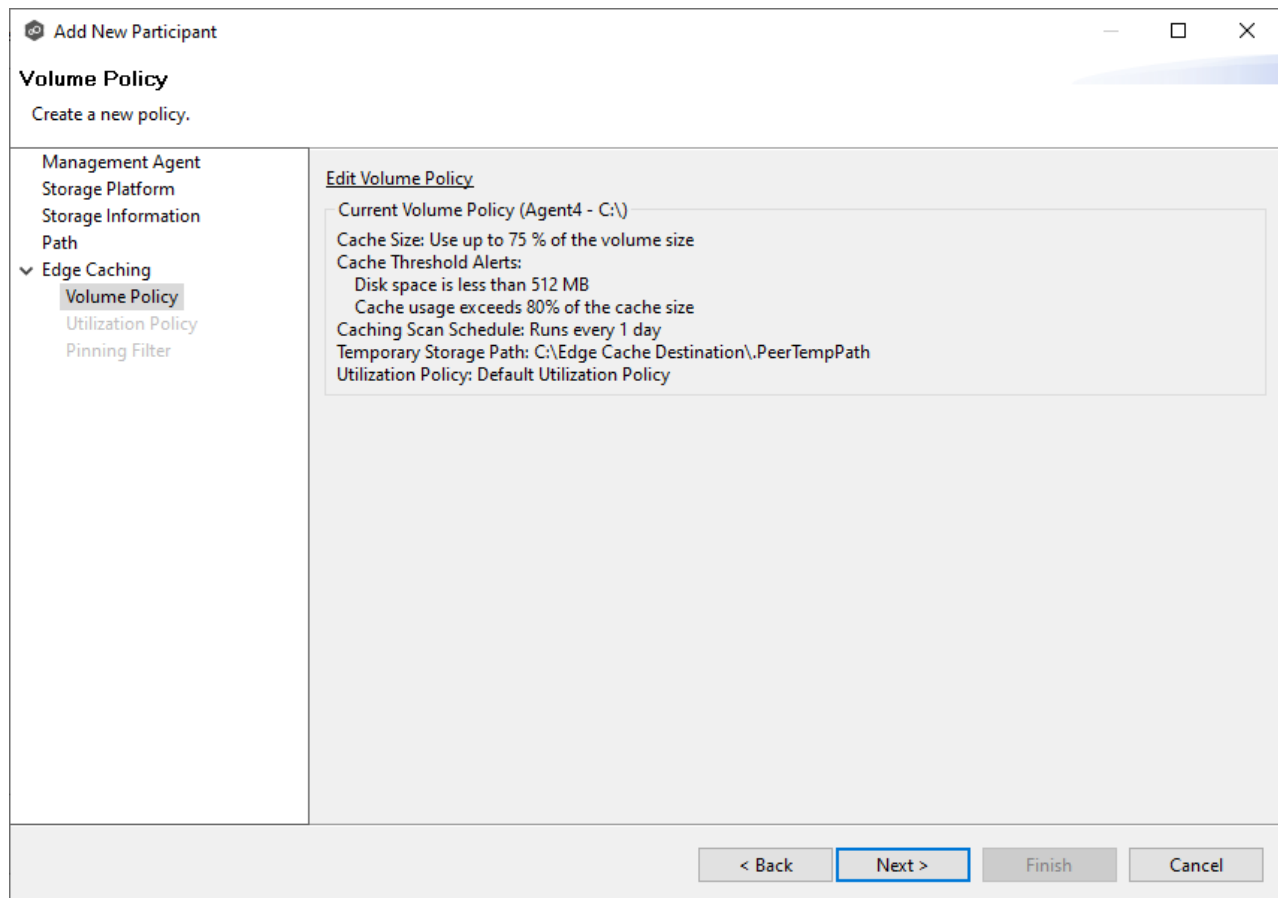
The temporary storage space will be used to store the content of stub files as they are being rehydrated. The content of files undergoing rehydration are referred to as **file blocks**. File blocks are fixed-length chunks of data that are read into memory when requested by an application. Edge Caching will create a subfolder named **.PeerTempPath** under the location that you specify and temporarily store the file blocks in that subfolder.

For optimal performance, we recommend that the temporary storage space be on the same volume as the watch set. If that is not possible, it should be on a high-performance disk.

6. Click **Next**.

The [Utilization Policy](#) page appears.

Note: If the Agent you selected is already being used as an edge participant in another job utilizing Edge Caching, the existing volume policy will be displayed on this page. You can edit the existing volume policy; however, be aware that any modifications to the volume policy will be applied to every other job that uses this Agent as an edge participant and "touches" the same volume.



Utilization Policy

The **Utilization Policy** page appears if you chose the edge role for the participant. The primary purpose of a **utilization policy** is to specify the parameters that govern when files on this edge participant should be stubbed or fully hydrated. Whereas the volume policy controls how much space is available to Edge Caching on a specific volume (or drive letter), the utilization policy controls whether to stub or hydrate a file.

Utilization policy parameters are based on the size of the files to be potentially stubbed and when they were last accessed and modified. A utilization policy enables you to balance getting the best performance while keeping the cache as full as possible.

You can select an existing utilization policy to apply to the job or create a new utilization policy. Whereas a volume policy is specific to a volume, a utilization policy can be reused for multiple jobs.

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

2. If you selected **Existing Policy**, select the policy, and then click **Next**.

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

3. (Optional) In the **File Size** section, select one or both options:

Field	Description
Keep files local if less than X size	Select this option if you want files under a specified size to remain local.
Stub files if greater than X size	Select this option if you want files over a specified size to be stubbed.

4. (Optional) In the **Time Period** section, select one of the options:

Field	Description
Keep recently used files local based on a dynamic set of rules	Select this option if you want Edge Caching to control when to stub files based on last accessed and last modified times. Edge Caching dynamically adjusts the accessed and modified time periods it uses based on the total amount of space that Edge Caching is actively using on a volume.
Keep recently used files local based on the following rules	Select this option if you want to specify the dates used when stubbing files based on the last accessed and last modified.

5. If you selected the **Keep recently used files local based on the following rules** option in **Time Period**, two additional options are enabled; select the options to be used and specify the time periods to be used:

Field	Description
Stub files if not modified within the past X time period	Select this option and specify a time range to use the last modified date of a file to determine if it should be kept local or stubbed.
Stub files is not accessed within the past X time period	Select this option and specify a time range to use the last accessed date of a file to determine if it should be kept local or stubbed.

6. (Optional) In the **Stubbing Override** section, select the checkbox and a time period if you want to use the last accessed date of all recently hydrated files to determine if they should be kept local or re-stubbed.
7. Click **Next** or **Finish**.

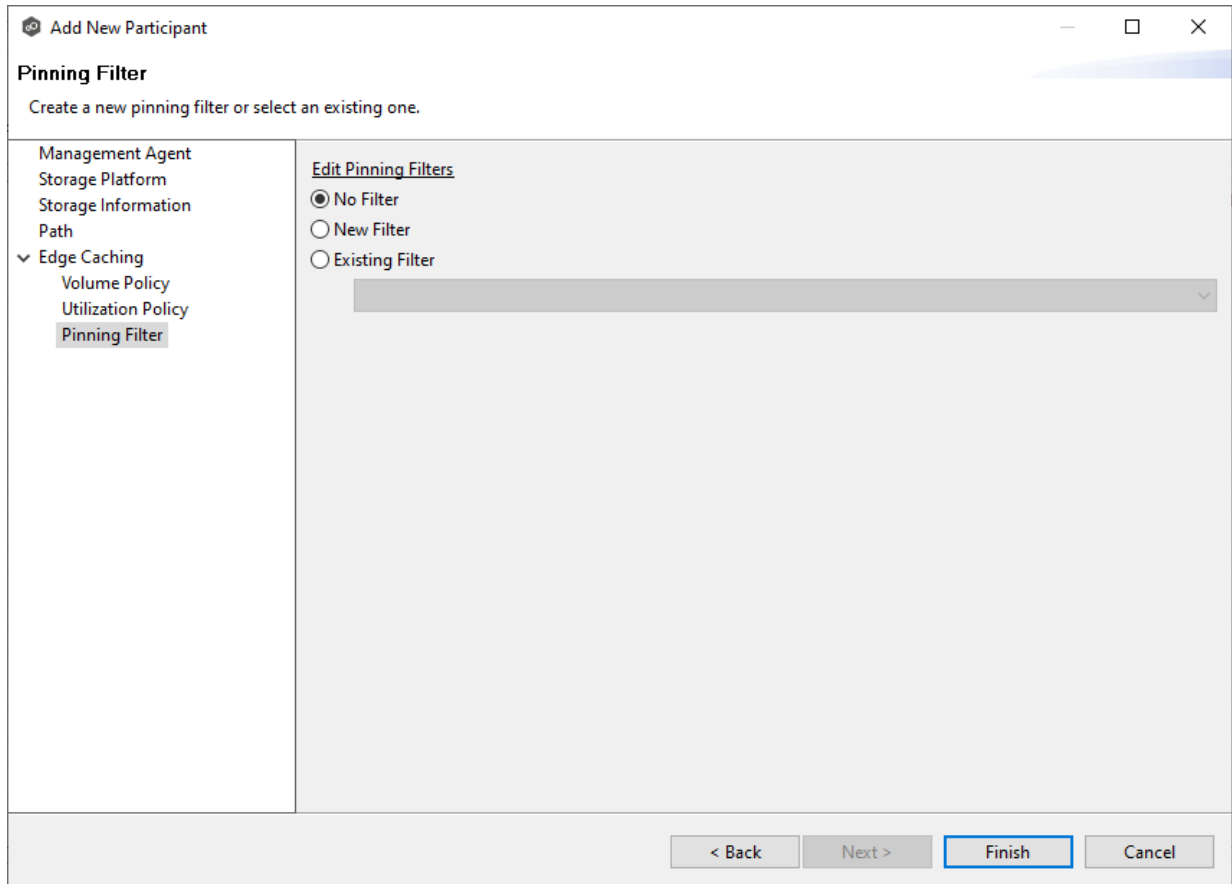
If you click **Next**, the [Pinning Filter](#) page appears.

Pinning Filter

The **Pinning Filter** page allows you to create a new pinning filter or select an existing pinning filter to apply to the job. A **pinning filter** specifies whether specific files or files in a particular

directory are always stubbed or always local on an edge participant. A pinning filter similar is to a utilization policy—it can be applied to multiple jobs. If there is a conflict between a pinning filter and utilization policy (where, for example, you might have something set to be always stubbed), the pinning filter will take precedence. Pinning filters are optional.

1. Select one of the options: **No Filter**, **New Filter**, or **Existing Filter**.



The screenshot shows a window titled "Add New Participant" with a "Pinning Filter" section. The section contains the instruction "Create a new pinning filter or select an existing one." Below this, there is a left-hand navigation pane with a tree view showing categories like "Management Agent", "Storage Platform", "Storage Information", "Path", "Edge Caching", "Volume Policy", "Utilization Policy", and "Pinning Filter". The "Pinning Filter" option is selected. To the right of the navigation pane, there is a section titled "Edit Pinning Filters" with three radio button options: "No Filter" (which is selected), "New Filter", and "Existing Filter". Below these options is a dropdown menu. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish" (which is highlighted with a blue border), and "Cancel".

2. If you selected **No Filter**, click **Finish**; if you selected **Existing Filter**, select the filter, and then click **Finish**.

If you selected **New Filter**, enter a name for the filter.

The screenshot shows the 'Add New Participant' dialog box with the 'Pinning Filter' section active. A red error message at the top states 'Name should not be empty.' The left sidebar contains a tree view with 'Pinning Filter' selected. The main area is titled 'Edit Pinning Filters' and has two radio buttons: 'No Filter' and 'New Filter' (which is selected). Below the radio buttons is a text field for '*Name:'. Underneath is a table titled 'Pinning Rules:' with two columns: 'Path' and 'Pinning State'. The table is currently empty. To the right of the table are three buttons: 'Create', 'Edit', and 'Delete'. Below the table is another radio button for 'Existing Filter' and a dropdown menu. At the bottom of the dialog are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

3. Enter a name for the filter.
4. Click **Create**.

The **Create Pinning Rule** dialog appears.

The screenshot shows the 'Create Pinning Rule' dialog box. It has a title bar with a close button. The main text says 'Enter a file name or path and choose pinning state.' There are two input fields: 'Pattern:' with an empty text box, and 'Pinning State:' with a dropdown menu showing 'Local at Edge'. At the bottom are two buttons: 'OK' and 'Cancel'.

5. Enter a file name or path in the **Pattern** field and then choose a pinning state: **Local at Edge** or **Stubbed at Edge**.

- Click **OK**.

The rule appears in the filter table.

Add New Participant

Pinning Filter

Create a new pinning filter or select an existing one.

Management Agent
Storage Platform
Storage Information
Path
Edge Caching
Volume Policy
Utilization Policy
Pinning Filter

Edit Pinning Filters

No Filter
 New Filter

*Name: Outlook Files

Pinning Rules:

Path	Pinning State
*.pst	Local at Edge

Existing Filter

Create
Edit
Delete

< Back Next > **Finish** Cancel

- (Optional) Create additional pinning rules.
- Click **Finish**.

The **Participants** page reappears. The participant is listed in the **Participants** table with the **Edge** role.

Create File Collaboration Job Wizard

Participants
Add two or more participants to this File Collaboration job.

Participants	Agent	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Cachin...
Master-Edge Assignment	Agent2		C:\FC-5 Folder A	Yes	Windows Driver	No	Master
File Metadata	Agent3		C:\FC-5 Folder B	Yes	Windows Driver	No	Master
Application Support	Agent4		C:\FC-5 Folder C	Yes	Windows Driver	No	Edge
Email Alerts							
DFS Namespace							

Buttons: Add, Edit, Delete

Navigation: < Back, Next >, Finish, Cancel

- Continue adding more participants if applicable or continue with [Step 3: Master-Edge Assignment](#)

Step 3: Master-Edge Assignment

This step is optional.

The **Master-Edge Assignment** page appears only if you enabled Edge Caching for one or more participants in Step 2. The purpose of this page is to allow you to assign one or more master participants to each edge participant.

Every edge participant must have at least one master participant assigned to it, so that Edge Caching will know which master participants to use when reading or rehydrating a stub file on an edge participant. For each edge participant, you want to assign the master participant that is the fastest and closest to it. This will increase the speed of rehydrating a stub file.

It is highly recommended that you assign, if possible, more than one master participant to an edge participant, so that if one master participant is not available, another master participant is able to provide the file content to the edge participant. You can set the order in which the master participants are consulted.

If you have only one edge participant and one master participant, the master participant is automatically assigned to the edge participant and listed in the Master Participants column. If you have multiple master participants, you need to explicitly assign master participants to each edge participant and then set the failover order.

- Select an edge participant in the **Assignment** table.

to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized to all participants as they occur.

- **Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan** - Select this option if you want the metadata synchronized during the initial scan. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized during the initial scan.

2. Click **OK** in the message that appears after selecting a metadata option.
3. If you selected either of the first two options in the **Synchronize Security Descriptor Options** section, select the security descriptor components (Owner, DACL, and SACL) to be synchronized.
4. If you selected the option for metadata synchronization during the initial scan, select the host to be used as the [master host](#) in case of metadata conflict.

If a master host is not selected, then no metadata synchronization will be performed during the initial scan. If one or more security descriptors do not match across participants during the initial scan, [conflict resolution](#) will use permissions from the designated master host as the winner. If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

5. Click **Next**.

The [Application Support](#) page is displayed.

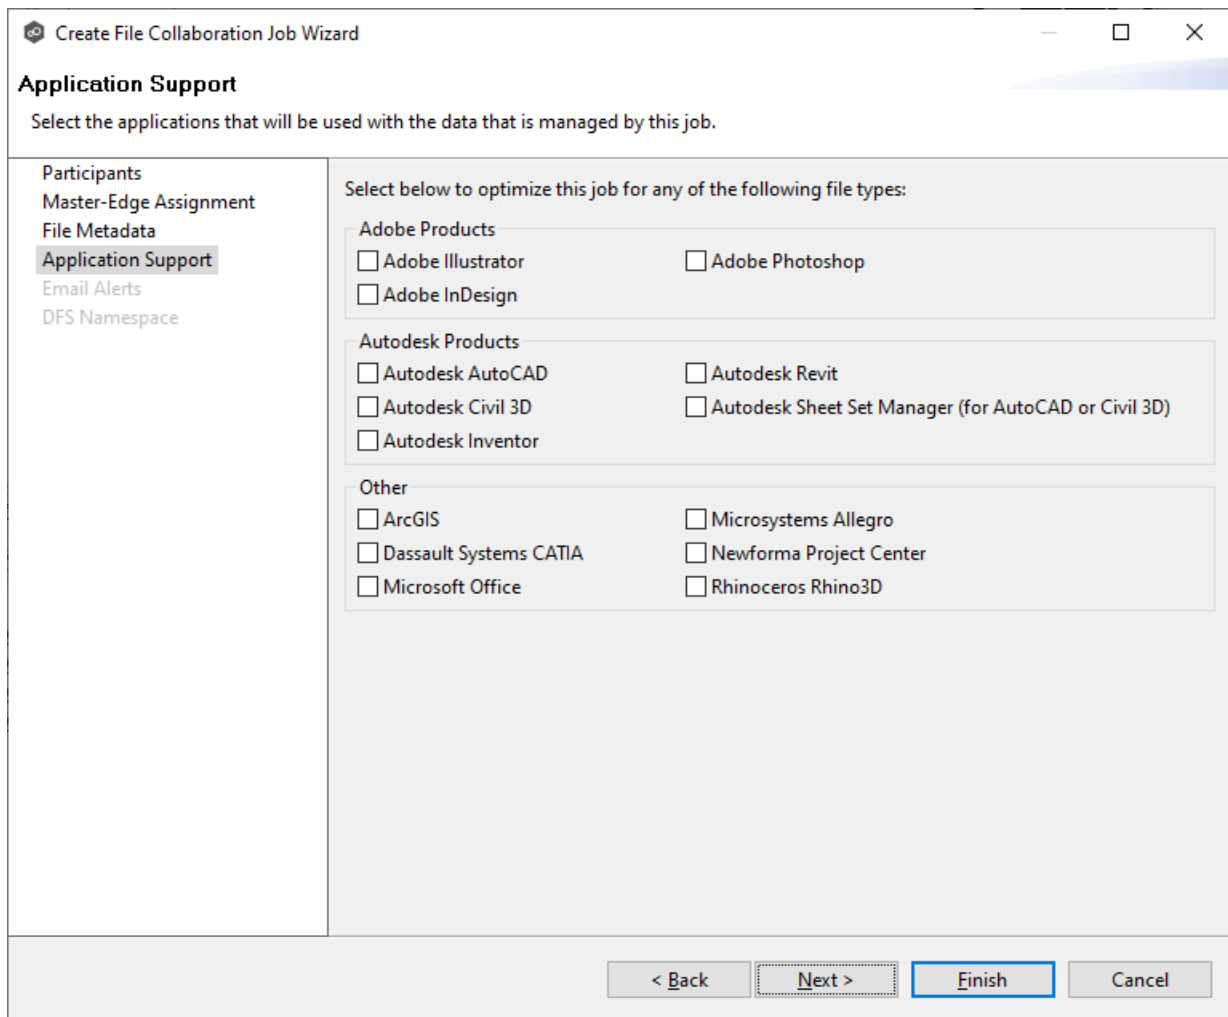
Step 5: Application Support

This step is optional.

Application Support enables automatic optimization of a file collaboration job for files created by certain 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 [Peer Support](#).

1. Select the applications that have files in the job's watch set.



2. Click **Next**.

The [Email Alerts](#) page is displayed.

Step 6: Email Alerts

This step is optional.


An email alert notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in [Preferences](#) and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply.

To create a new alert, see [Email Alerts](#) in the [Preferences](#) section.

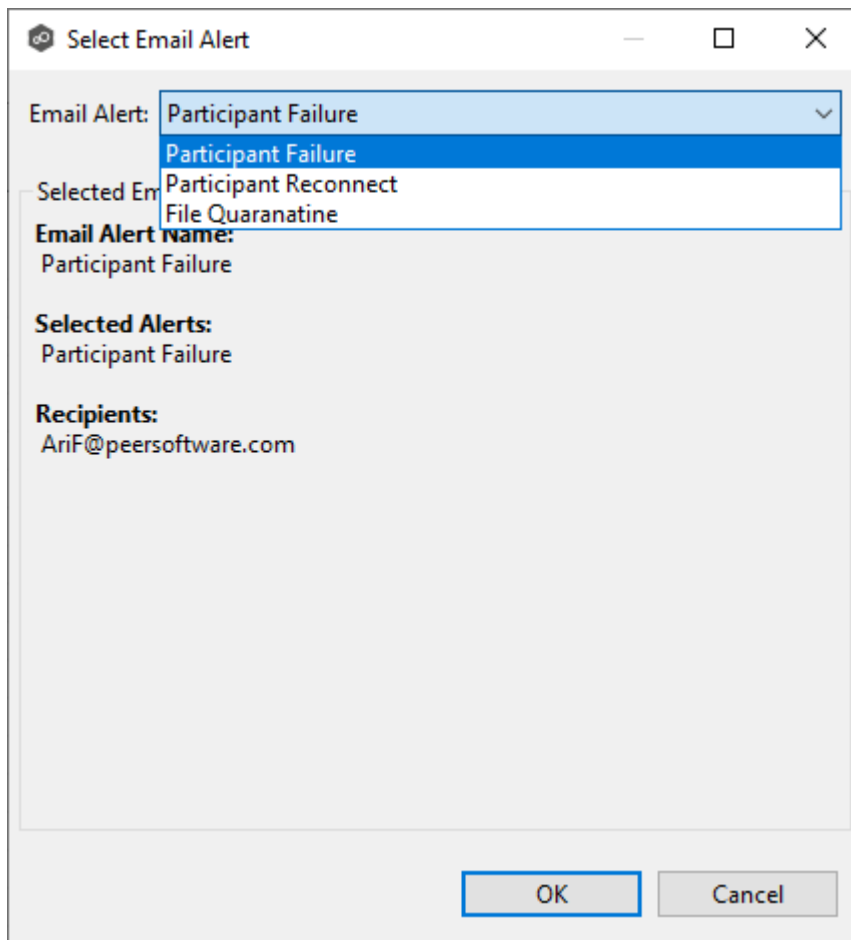
To apply an existing email alert to the job.

1. Click the **Select** button.




The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.

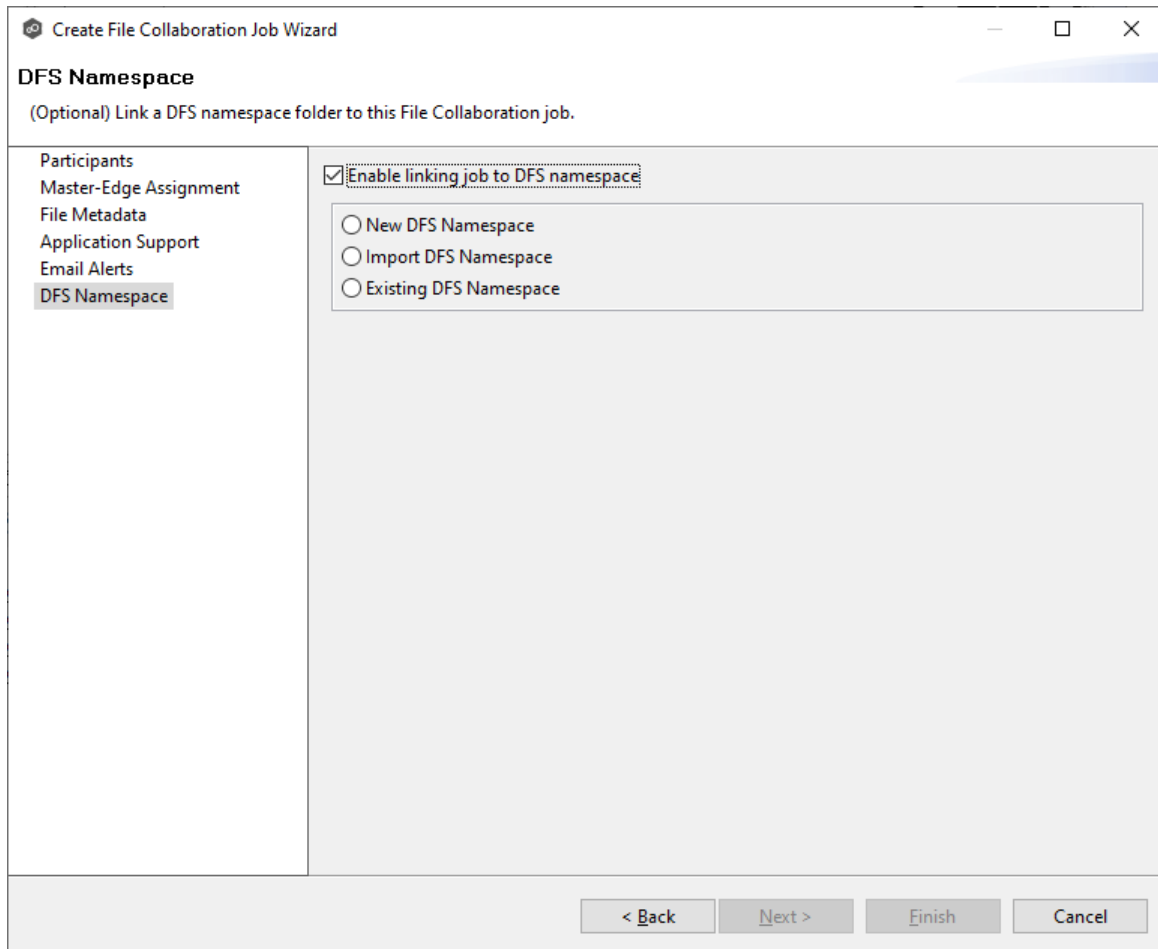


3. Click **OK**.

The alert is listed in the **Email Alerts** page.



The three options are enabled.



2. Select one of the three options:

- **New DFS Namespace** - Select this option if you want to create a new namespace. If you select this option, the **Create DFS-N Management Job Wizard** opens. Follow these steps to [create a new namespace](#).
- **Import DFS Namespace** - Select this option if you have a namespace that was created using the Microsoft DFS Management tool and is not currently being managed by a DFS-N Management job. If you select this option, the **Import Existing Namespaces** Wizard opens. For detailed instructions, follow these steps to [import an existing namespace](#).
- **Existing DFS Namespace** - Select this option if you want to use an existing namespace that is being managed by a DFS-N Management job. If you select this option, it will display the namespace folder and folders associated with namespace.

Create File Collaboration Job Wizard

DFS Namespace
(Optional) Link a DFS namespace folder to this File Collaboration job.

Participants
Master-Edge Assignment
File Metadata
Application Support
Email Alerts
▼ **DFS Namespace**
DFS-N Link

Enable linking job to DFS namespace

New DFS Namespace
 Import DFS Namespace
 Existing DFS Namespace

Documentation

Selected DFS Namespace: Documentation

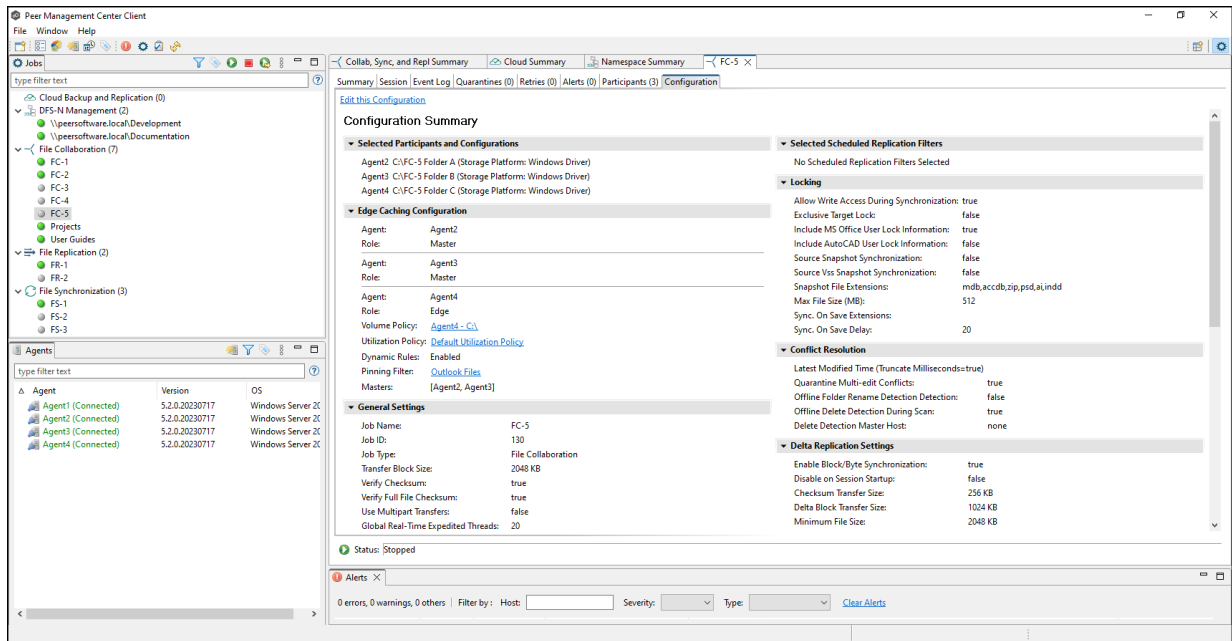
Folders:
- User Guides

Targets:
- \\Agent4\User Guide - Cape Town
- \\Agent3\User Guide - Munich
- \\Agent2\User Guide - DC

< Back Next > **Finish** Cancel

Click **Next** if you want to link participants with folder targets on the **DFS Link** page; otherwise continue with Step 3.

For more information about linking participants to folder targets, see [Linking an Existing Namespace Folder to an Existing File Collaboration or File Synchronization Job](#).



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 Link](#)
- [File and Folder Filters](#)
- [File Metadata](#) - Some options are available only when editing the job.
- [File Locking](#)
- [General](#)

- [Scheduled Replication Filters](#)
 - [Conflict Resolution](#)
 - [Delta Replication](#)
 - [File Metadata](#)
 - [File Locking](#)
 - [Application Support](#)
 - [Target Protection](#)
 - [Email Alerts](#)
 - [SNMP Notifications](#)
 - [Tags](#)
 - [DFS-N Link](#)
4. Click **OK** when finished.

Participants

The **Participants** page in the **Edit File Collaboration Configuration** dialog allows you to:

- [Add and delete participants from a job.](#)
- [Edit a participant.](#)




This topic describes [adding](#) and [deleting](#) participants in a File Collaboration job.

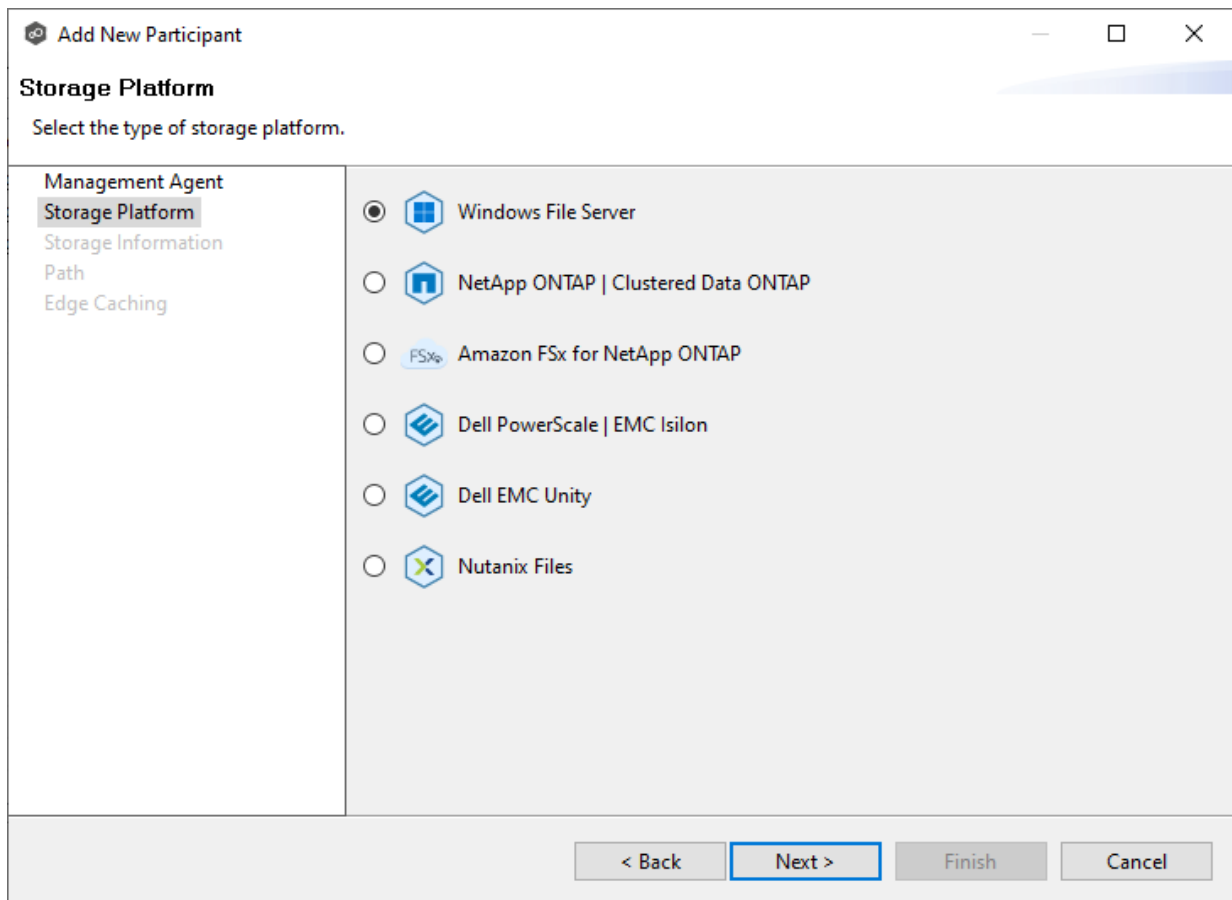
Adding a Participant to a File Collaboration Job

To add a participant to a file collaboration job:

1. Select the job in the **Jobs** view; right-click and select **Edit Job**.

The **Edit File Collaboration** dialog opens; the **Participants** page displays the current job participants.





4. Select the type of storage platform that hosts the data you want to collaborate on, and then click **Next**.

The **Storage Information** page appears; the choices available depend on your selection on the **Storage Platform** page.

5. Enter the requested information for your storage platform:

[Windows File Server](#)

[NetApp ONTAP](#)

[Amazon FSxN](#)

[Dell PowerScale](#)

[Dell EMC Unity](#)

[Nutanix Files](#)

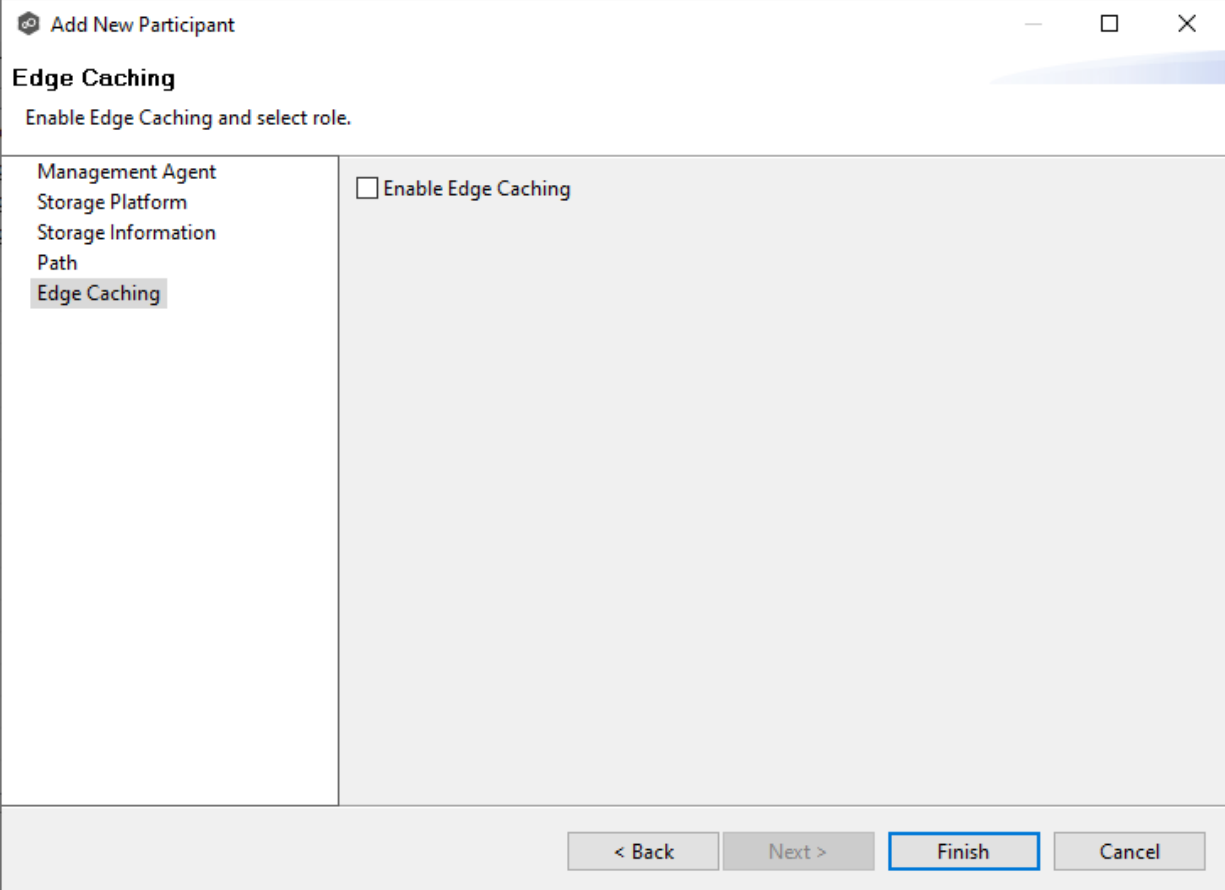
6. Click **Next**.

The **Path** page appears.

The screenshot shows a window titled "Add New Participant" with a "Path" sub-header. A red error message states "You must enter a valid path." The main content area is divided into two sections. On the left is a vertical navigation menu with the following items: "Management Agent", "Storage Platform", "Storage Information", "Path" (which is highlighted), and "Edge Caching". The right section contains an "Enter Path" text input field, a "Browse" button to its right, and a "Seeding Target" checkbox below the input field. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

7. Browse to or enter the path to the [watch set](#).
8. (Optional) Select the **Seeding Target** checkbox, and then click **OK** in the dialog that appears.
9. Click **Next**.

The **Edge Caching** page appears.



The screenshot shows a dialog box titled "Add New Participant" with a close button in the top right corner. The main heading is "Edge Caching" with the instruction "Enable Edge Caching and select role." Below this, there is a list of roles on the left: "Management Agent", "Storage Platform", "Storage Information", "Path", and "Edge Caching" (which is highlighted). To the right of this list is a checkbox labeled "Enable Edge Caching". At the bottom of the dialog, there are four buttons: "< Back", "Next >", "Finish" (which is highlighted with a blue border), and "Cancel".

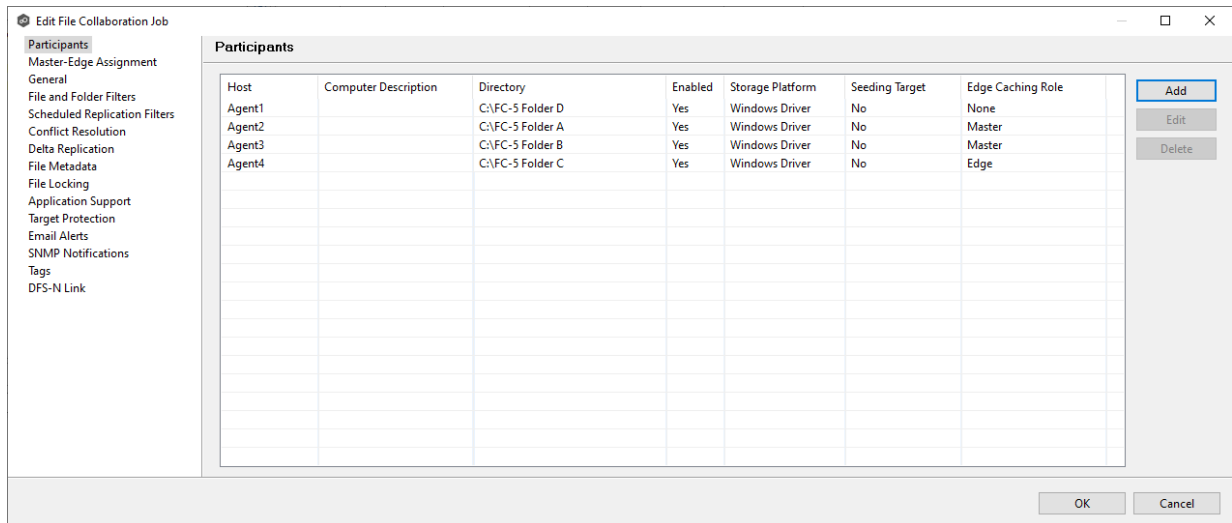
10. (Optional) Select the **Enable Edge Caching** checkbox if you want this participant to be able to use Edge Caching; otherwise, click **Finish**.

11. If you enabled Edge Caching, follow the steps outlined in [Step 2: Edge Caching](#) in [Creating a File Collaboration Job](#).

For more information about Edge Caching, see [Edge Caching](#) in [Advanced Topics](#).

12. Click **Finish**.

The new participant appears in the **Participants** table.

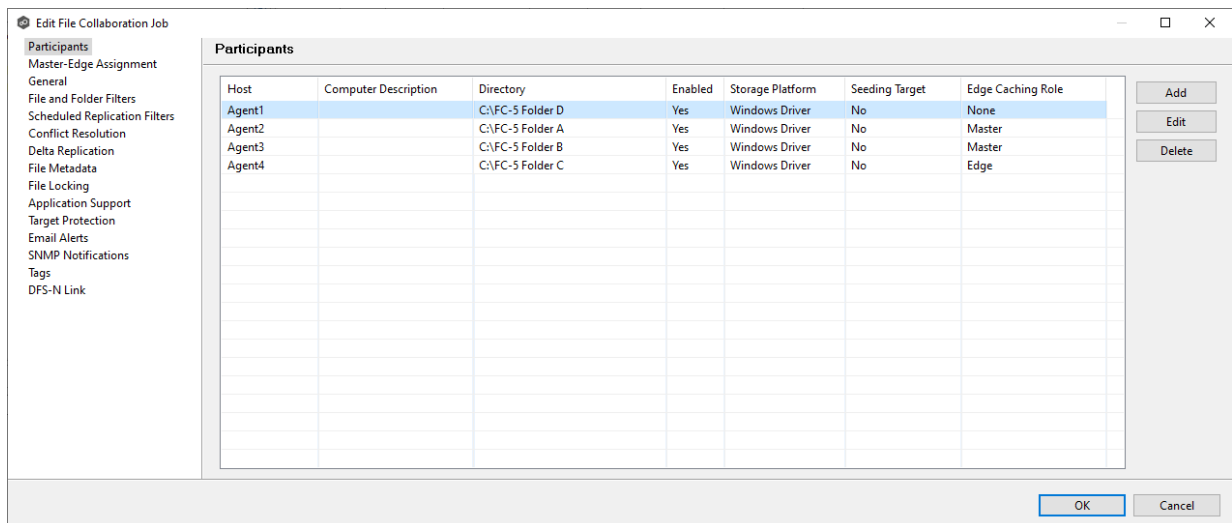


13. Click **OK** to close the Edit wizard or select another configuration item to modify.

Deleting a Participant from a File Collaboration Job

To delete a participant from a File Collaboration job:

1. In the **Edit File Collaboration** dialog, select the participant in the **Participants** table you want to remove from the job.

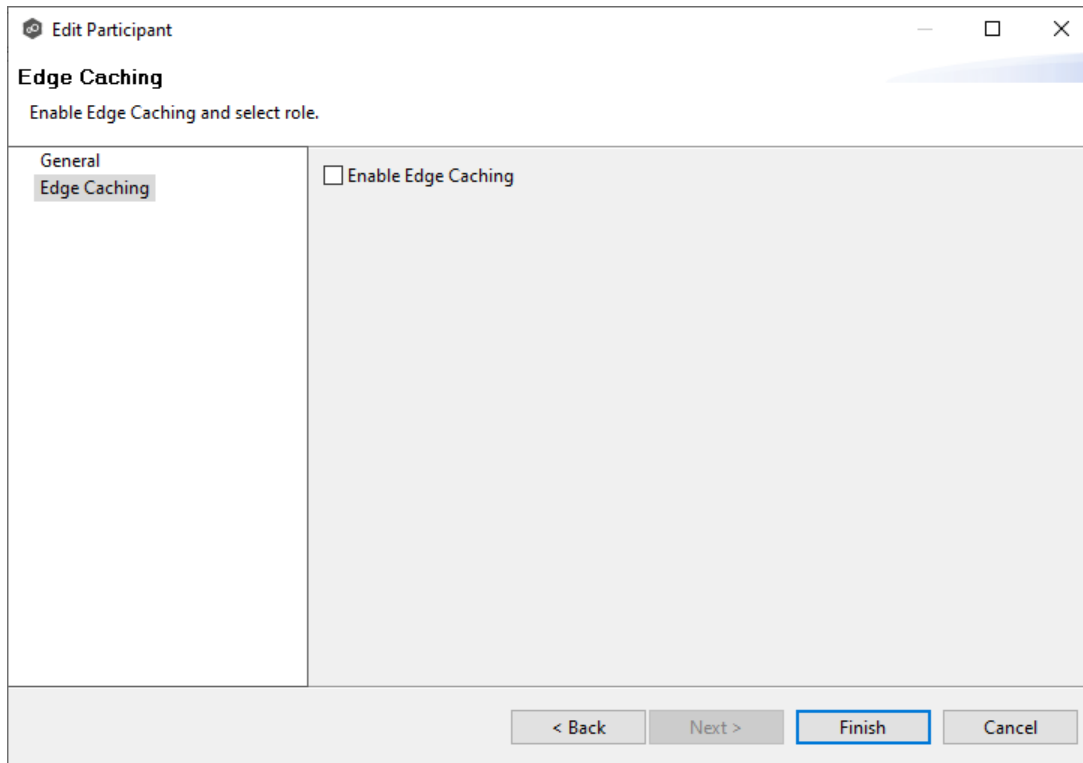


2. Click the **Delete** button.

The screenshot shows the 'Edit Participant' dialog box with the 'General' tab selected. The dialog has a title bar with a close button and a maximize button. Below the title bar, the text 'General' is followed by 'Modify these settings.' The main area is divided into two panes. The left pane is titled 'General' and contains 'Edge Caching'. The right pane contains the following settings: 'Enabled:' with a checked checkbox; 'Host:' with a text box containing 'Agent1'; 'Event Detector:' with a dropdown menu showing 'Windows Driver' and a link 'Edit Detector Configuration'; 'Directory:' with a text box containing 'C:\FC-5 Folder D' and a 'Browse' button; and 'Seeding Target' with an unchecked checkbox. At the bottom of the dialog are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

3. To enable or disable the Agent, select or deselect the **Enabled** checkbox.
4. To change the directory/folder/share that is replicated, enter the path or browse to the new watch set in the **Directory** field.
5. If the settings required to connect to the storage device have changed, click **Edit Detector Configuration**, and then make the necessary modifications.
6. To change whether the participant is a seeding target, select or deselect the **Seeding Target** checkbox.
7. Click **Next** to change Edge Caching options; otherwise, click **Finish**, and continue with Step 10.

If you clicked **Next**, the Edge Caching page appears.



8. (Optional) Select the **Enable Edge Caching** checkbox if you want this participant to be able to use Edge Caching.
9. If you enabled Edge Caching, follow the steps outlined in [Step 2: Edge Caching in Creating a File Collaboration Job](#).
10. Click **OK** to close the **Edit Participant** wizard.

Master-Edge Assignment

This page appears only when Edge Caching is enabled for the job.

Every edge participant must have at least one master participant assigned to it, so that Edge Caching will know which master participants to use when reading or rehydrating a stub file on an edge participant. For each edge participant, you want to assign the master participant that is the fastest and closest to it. This will increase the speed of rehydrating a stub file.

It is highly recommended that you assign, if possible, more than one master participant to an edge participant, so that if one master participant is not available, another master participant is able to provide the file content to the edge participant. You can set the order in which the master participants are consulted.

General

The **General** page in the **Edit File Collaboration Job** dialog presents miscellaneous settings pertaining to a File Collaboration job. You may want to consult with Peer Software's Technical Support team before modifying these values.

To modify these settings:

1. Enter the values recommended by Peer Software Support.

Option	Description
Job ID	Unique, system-generated job identifier that cannot be edited.
Job Type	Identifies the job type. This cannot be modified.
Job Name	The name of this File Collaboration job. This name must be unique.

Option	Description
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 .
Verify Block Checksums	If selected, each block sent will be checksummed at both the source and target(s) Agents.
Verify Full File Checksums	If selected, the entire file will be checksummed <i>after</i> it has been sent from the source to all target Agents. If temp files are enabled, the checksum on the target will be calculated prior to renaming the temp file to the base file's name. If the checksums between source and target(s) do not match, the file will be sent to the retry engine to reattempt the transfer.
Enable Multipart Transfers	If selected, larger files will be broken into smaller chunks and these chunks will be sent over the wire in a parallel fashion. This feature will automatically throttle itself if many other transfers are also waiting to be processed.
Synchronization 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 are NOT seeded, the FOLDER_BY_FOLDER method would be best. For environments where most data are seeded, the BULK_CHECKSUM method will result in a faster first scan.
Remove Filtered Files On Folder Delete	If selected, then all child files on target hosts will be deleted when its parent folder is deleted on another source host. Otherwise, filtered files will be left intact on targets when a parent folder is deleted on another source host.
Require All Hosts At	If selected, requires all participating hosts to be online and available at the start of the File Collaboration job in order for the

Option	Description
Start	job to successfully start.
Auto Start	If selected, then this file collaboration session will automatically be started when the Peer Management Center Service is started.

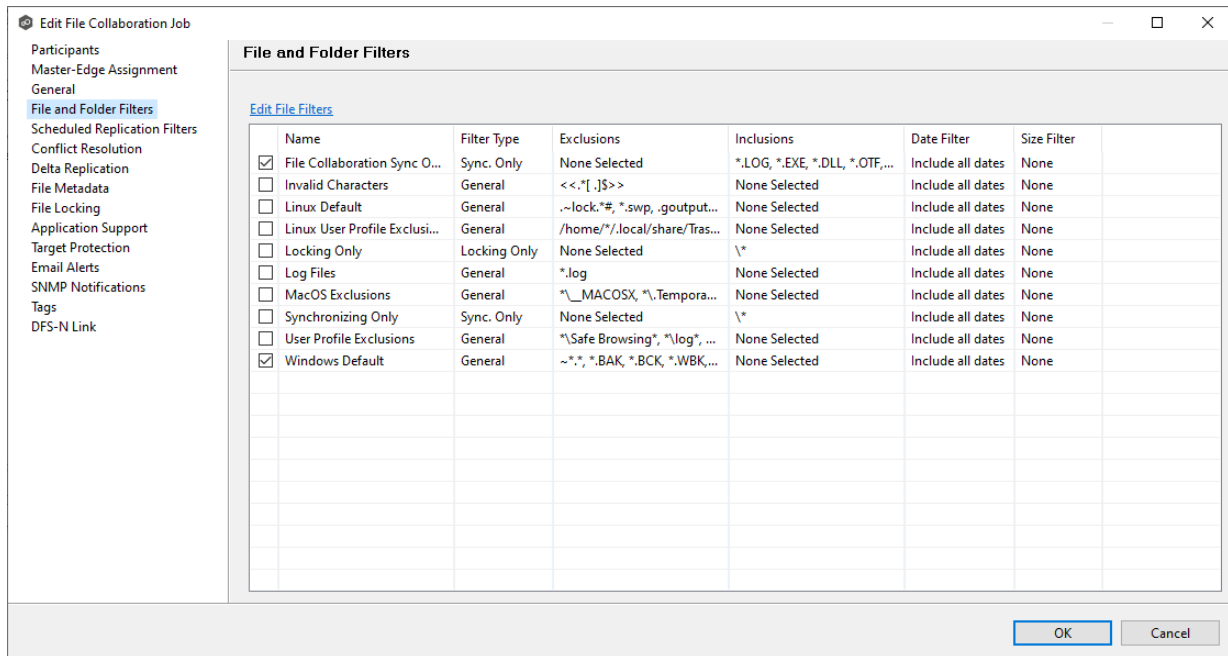
2. Click **OK** to close the Edit wizard or select another configuration item to modify.

File and Folder Filters

The **File and Folder Filters** page in the **Edit File Collaboration Job** dialog displays a list of [file and folder filters](#). A file or folder filter enables you to exclude and/or include files and folders from the job based on file type, extension, name, or directory path. Any file or folder that matches the filter is excluded or included from replication, depending on the filter's definition. By default, all files and folders selected in the **Source Paths** page will be replicated.

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

If you want to create a new file or folder filter or modify an existing one, click **Edit File Filters**. See [File and Folder Filters](#) in the [Preferences](#) section for information about creating or modifying a file filter.



2. Click **OK** to close the Edit wizard or select another configuration item to modify.

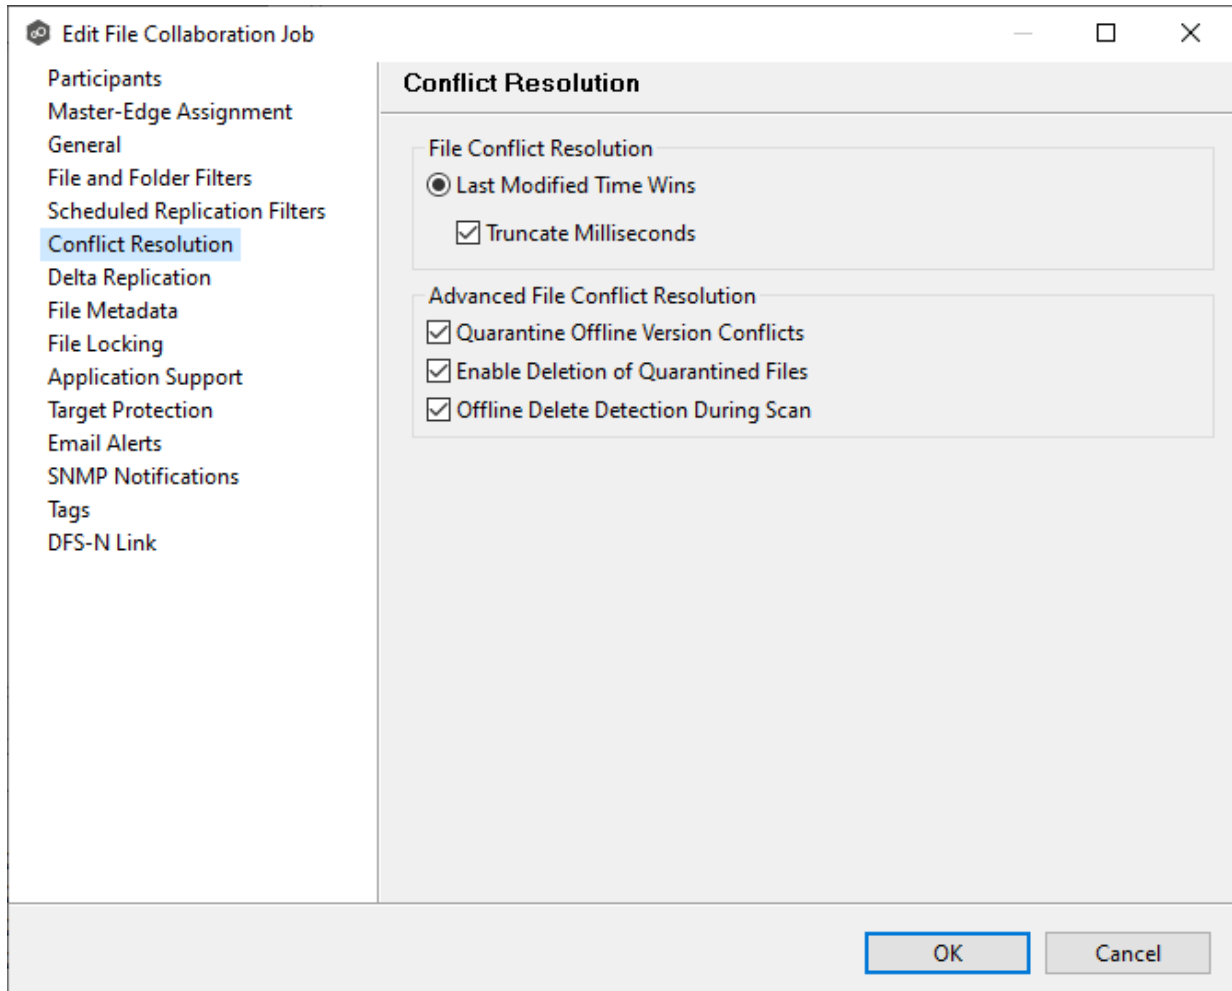
Scheduled Replication Filters

The **Scheduled Replication Filters** page in the **Edit File Collaboration Job** dialog displays a list of [scheduled replication filters](#). A scheduled replication filter enables you to identify files and folders that you don't want replicated in real-time.

1. Select the scheduled replication filters you want to apply to the job.

If you want to create a new filter or modify an existing one, click **Edit Scheduled Replication Filters**. See [Scheduled Replication Filters](#) in the [Preferences](#) section for information about creating or modifying a scheduled replication filter.

As some storage platforms and applications track milliseconds slightly differently, selecting this option will prevent subtle millisecond differences from causing an otherwise in-sync file to be replicated or quarantined.



2. Select the **Advanced File Conflict Resolution** options you want applied:

Option	Description
Quarantine Offline Version Conflicts	Select this option if you want Peer Management Center to quarantine a file that was updated in two or more locations while the collaboration session was not running. If it is not selected, the file with the most recent last modified time will be replicated to all other participants.
Enable Deletion of	Select this option if you want Peer Management Center to process a delete event for a quarantined file. If it is not selected, the

Option	Description
Quarantined Files	quarantined file is not deleted and remains quarantined.
Offline Delete Detection During Scan	Select this option (and enable target protection), if you want any files or folders that were deleted while the job was stopped to be deleted from all participants when the job is restarted. If it is not selected, then the deleted file or folder is restored from a participant with the file or folder to any participant where it no longer exists.

3. Click **OK** to close the Edit wizard or select another configuration item to modify.

Delta Replication

The **Delta Replication** page in the **Edit File Collaboration Job** dialog allows you to specify the delta-replication options to use for the selected File Collaboration job. Delta-level replication is a byte replication technology that enables block/byte level synchronization for a File Collaboration job. Through this feature, Peer Management Center is able to transmit only the bytes/blocks of a file that have changed instead of transferring the entire file. This results in much lower network bandwidth utilization, which can be an enormous benefit if you are transferring files across a slow WAN or VPN, as well as across a high-volume LAN.

Delta-level replication is enabled on a per File Collaboration job basis and generally affects all files in the [watch set](#). You will only benefit from delta-level replication for files that do not change much between file modifications, which includes most document editing programs.

To modify delta-level replication options:

1. Modify the following the fields as necessary.

Edit File Collaboration Job

Participants
Master-Edge Assignment
General
File and Folder Filters
Scheduled Replication Filters
Conflict Resolution
Delta Replication
File Metadata
File Locking
Application Support
Target Protection
Email Alerts
SNMP Notifications
Tags
DFS-N Link

Delta Replication

Enable Delta-level Replication:

Checksum Transfer Size (KB): 256

Delta Block Transfer Size (KB): 1024

Minimum File Size (KB): 2048

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

Excluded File Name Patterns

OK Cancel

Field	Description
Enable Delta-Level Replication	Select to enable delta encoded file transfers which only sends the file blocks that are different between source and target(s). If this is disabled, the standard file copy method will be used to synchronize files.
Checksum Transfer Size (KB)	Enter the block in kilobytes used to transfer checksums from target to source at one time. Larger sizes will result in faster checksum transfer, but will consume more memory on the Peer Agents
Delta Block Transfer Size (KB)	Enter the block size in kilobytes used to transfer delta encoded data from source to target at one time. Larger sizes will result in faster overall file transfers but will consume more memory on the Peer Agents.
Minimum File Size (KB)	Enter the minimum size of files in kilobytes to perform delta encoding for. If a file is less than this size, then delta encoding will not be performed.
Minimum File Size Percentage Target/Source	Enter the minimum allowed file size difference between source and target, as a percentage, to perform delta encoding. If the target file size is less than this percentage of the source file size, then delta encoding will not be performed.

Field	Description
Excluded File Extensions	Enter a comma-separated list of file extension patterns to be excluded from delta encoding, e.g., zip, jpg, png. In general, compressed files should be excluded from delta encoding and the most popular compressed file formats are excluded by default.
Excluded File Name Patterns	Enter a list of file name patterns to be excluded from delta encoding. If a file name matches any pattern in this list, then it will be excluded from delta encoding transfers and a regular file transfer will be performed. See File and Folder Filters for more information on specifying wildcard expressions.

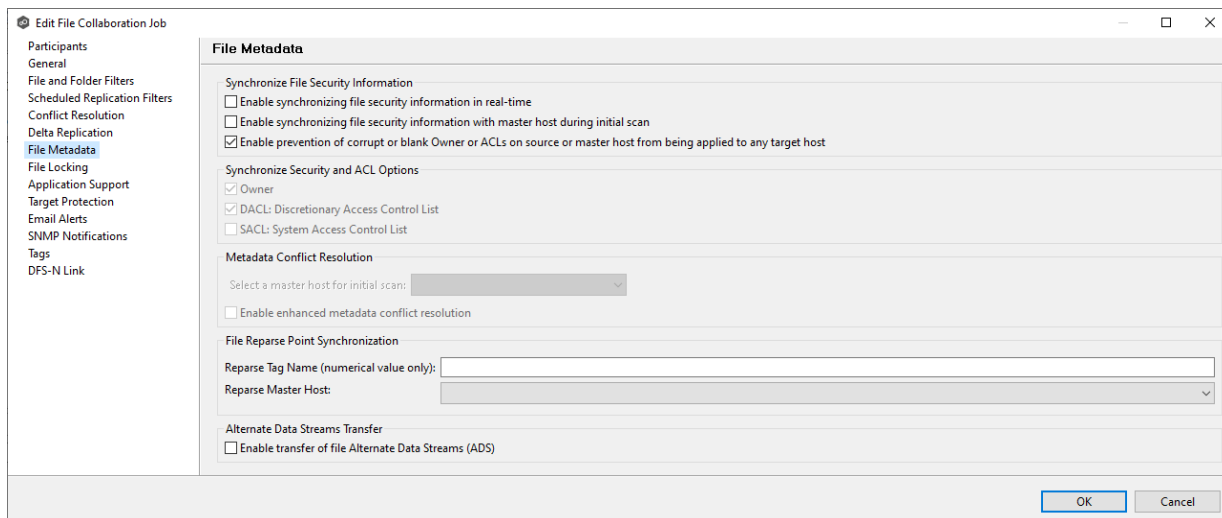
2. Click **OK** to close the Edit wizard or select another configuration item to modify.

File Metadata

The **File Metadata** page allows you to specify whether you wish to replicate file security metadata and the types of metadata for synchronization. Additionally, it provides the ability to designate the metadata source (volume, share, export, or folder) to resolve conflicts during the initial synchronization. This designated source, utilized in case of conflicts, is referred to as the [master host](#). This page also provides some additional options not available when creating the job. See [File Metadata Synchronization](#) in [Advanced Topics](#) for more information about file metadata replication.

To modify file metadata synchronization settings:

1. Select when you want the metadata synchronized (you can select one or both options):
 - **Enable synchronizing NTFS security descriptors (ACLs) in real-time** - Select this option if you want the metadata replicated in real-time. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be transferred to the target host file(s) as they occur.
 - **Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan** - Select this option if you want the metadata replicated during the initial scan. If enabled, changes to the selected security descriptor components (Owner, DACL, and SACL) will be synchronized during the initial scan.



2. Click **OK** in the message that appears after selecting a metadata option.
3. If you selected either of the first two options in the **Synchronize Security Descriptor (ACLs)** section, select the security descriptor components (**Owner**, **DACL**, and **SACL**) to be synchronized.
4. If you selected the option for metadata synchronization during the initial scan (second option in the **Synchronize Security Descriptor (ACLs)** section, select the host to be used as the [master host](#) in case of file metadata conflict.

If a master host is not selected, then no metadata synchronization will be performed during the initial scan. If one or more security descriptors do not match across participants during the initial scan, [conflict resolution](#) will use permissions from the designated master host as the winner. If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

5. (Optional) Select the **Enable enhanced metadata conflict resolution** checkbox.

This option is only available when both of the first two options in the **Synchronize Security Descriptor (ACLs)** section are enabled, and **Owner** is selected under **Synchronize Security Descriptor Options**.

If you select **Enable enhanced metadata conflict resolution**, this will prevent the Peer Agent service account from being assigned as the owner of that file or folder when a metadata conflict occurs, and a file or folder is written to a target. If the Peer Agent service account were to be assigned as the owner, the user may not have permission to access the file or folder.

Note: The Peer Agent service account cannot be a local or system administrator. As described in [Peer Global File Service - Environmental Requirements](#), the Peer agent service account should be an actual user.

6. (Optional) Enter values for one or both file reparse point data synchronization options:

- **Reparse Tag Name** - Enter a single numerical value. Must be either blank (if blank, reparse synchronization will be disabled) or greater than or equal to 0. The default for Symantec Enterprise Vault is 16. A value of 0 enables reparse point synchronization for all reparse file types. If you are unsure as to what value to use, then contact Peer Software technical support, or you can use a value of 0 if you are sure that you are only utilizing one vendor's reparse point functionality.
- **Reparse Master Host** - Select a master host. If a master host is selected, then when the last modified times and file sizes match on all hosts, but the file reparse attribute differs (e.g., archived/offline versus unarchived on file server), then the file reparse data will be synchronized to match the file located on the master host. For Enterprise Vault, this should be the server where you run the archiving task on. If the value is left blank, then no reparse data synchronization will be performed, and the files will be left in their current state.

Note: Use this option only if you are utilizing archiving or hierarchical storage solutions that make use of NTFS file reparse points to access data in a remote location, such as Symantec's Enterprise Vault. Enabling this option allows synchronization of a file's reparse data, and not the actual offline content, to target hosts, and prevents the offline file from being recalled from the remote storage device.

7. (Optional) Select the **Enable transfer of file Alternate Data Stream (ADS)** checkbox.

If enabled, Alternate Data Streams (ADS) of updated files will be transferred to the corresponding files on target participants as a post process of the normal file synchronization.

Known limitation: ADS information is transferred only when a modification on the actual file itself is detected. ADS will not be compared between participants. The updated file's ADS will be applied to the corresponding files on target participants.

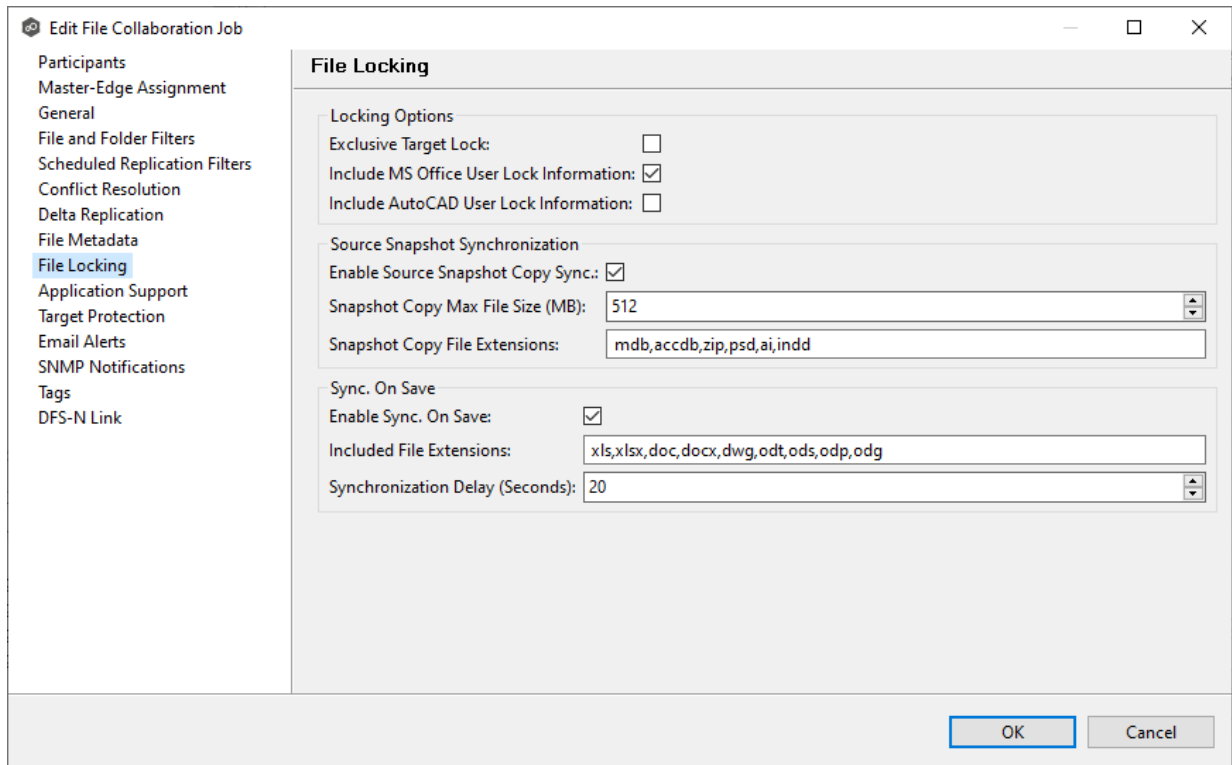
8. Click **OK** to close the Edit wizard or select another configuration item to modify.

File Locking

The **File Locking** page in the **Edit File Collaboration Job** dialog presents options pertaining to how source and target files are locked by Peer Management Center.

To modify file locking options:

1. Modify these fields as needed:



- Click **OK** to close the Edit wizard or select another configuration item to modify.

Locking Options

Option	Description
Exclusive Target Lock	If enabled, then whenever possible, an exclusive lock will be obtained on target file handles, which will prevent users from opening the file (even in read-only mode) while a user has the file opened on the source host. When this option is disabled, then users will be allowed to open files for read-only if the application allows for this.
Include MS Office User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported Microsoft Office files (e.g., Word, Excel, and PowerPoint).
Include AutoCAD User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported AutoCAD files.

Source Snapshot Synchronization Option

Option	Description
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.

Sync On Save Options

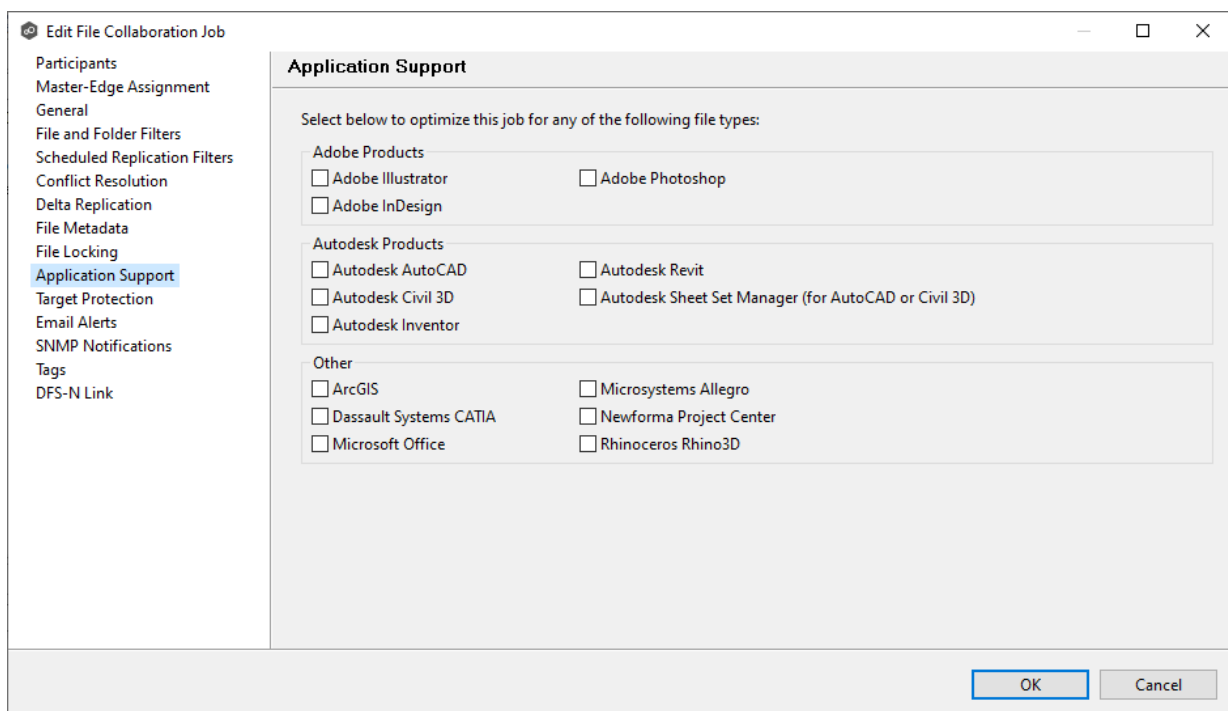
Option	Description
Exclusive Target Lock	If enabled, then whenever possible, an exclusive lock will be obtained on target file handles, which will prevent users from opening the file (even in read-only mode) while a user has the file opened on the source host. When this option is disabled, then users will be allowed to open files for read-only if the application allows for this.
Include MS Office User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported Microsoft Office files (e.g., Word, Excel, and PowerPoint).
Include AutoCAD User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported AutoCAD files.
Enable Source Snapshot Copy Sync.	If enabled, a snapshot copy of the source file will be created for files that meet the snapshot configuration criteria below, and this copy will be used for synchronization purposes. In addition, no file handle will be held on the source file except while making a copy of the file.
Snapshot Copy Max File Size (MB)	The maximum file size for which source snapshot synchronization will be utilized.
Snapshot Copy File Extensions	A comma-separated list of file extensions for which source snapshot synchronization will be utilized.
Enable Sync. On Save	If enabled, this feature will allow supported file types to be synchronized after a user saves a file, rather than waiting for the file to close.
Included File Extensions	A comma-separated list of file extensions for which to enable the Sync. On Save feature.
Synchronization Delay (Seconds)	The number of seconds to wait after a file has been saved before initiating a synchronization of the file.

Application Support

When you create a File Collaboration job, you have the option of [selecting applications to be automatically optimized](#). When editing the job, you can modify your selections in the **Application Support** page in the **Edit File Collaboration Job** dialog.

To modify which applications are optimized:

1. Select the applications to be optimized.



2. Click **OK** to close the Edit wizard or select another configuration item to modify.

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 a job has target protection enabled, then whenever a file is deleted or modified on the source host but before the changes are propagated to the targets, a copy of the existing file on the target is moved to the Peer Management Center trash bin.

The trash bin is located in a hidden folder named **.pc-trash_bin** found in the root directory of the [watch set](#) of the target host. A backup file is placed in the same directory hierarchy location as the source folder in the watch set within the recycle bin folder. If you need to restore a previous version of a file, you can copy the file from the trash bin into the corresponding location in the watch set and the changes will be propagated to all other collaboration hosts.

You can configure target protection in the **Target Protection** page in the **Edit File Collaboration Job** dialog.

Edit File Collaboration Job

- Participants
- Master-Edge Assignment
- General
- File and Folder Filters
- Scheduled Replication Filters
- Conflict Resolution
- Delta Replication
- File Metadata
- File Locking
- Application Support
- Target Protection**
- Email Alerts
- SNMP Notifications
- Tags
- DFS-N Link

Target Protection

Enabled:

of Backup Files to Keep: 3

of Days to Keep: 30

Trash Bin: .pc-trash_bin

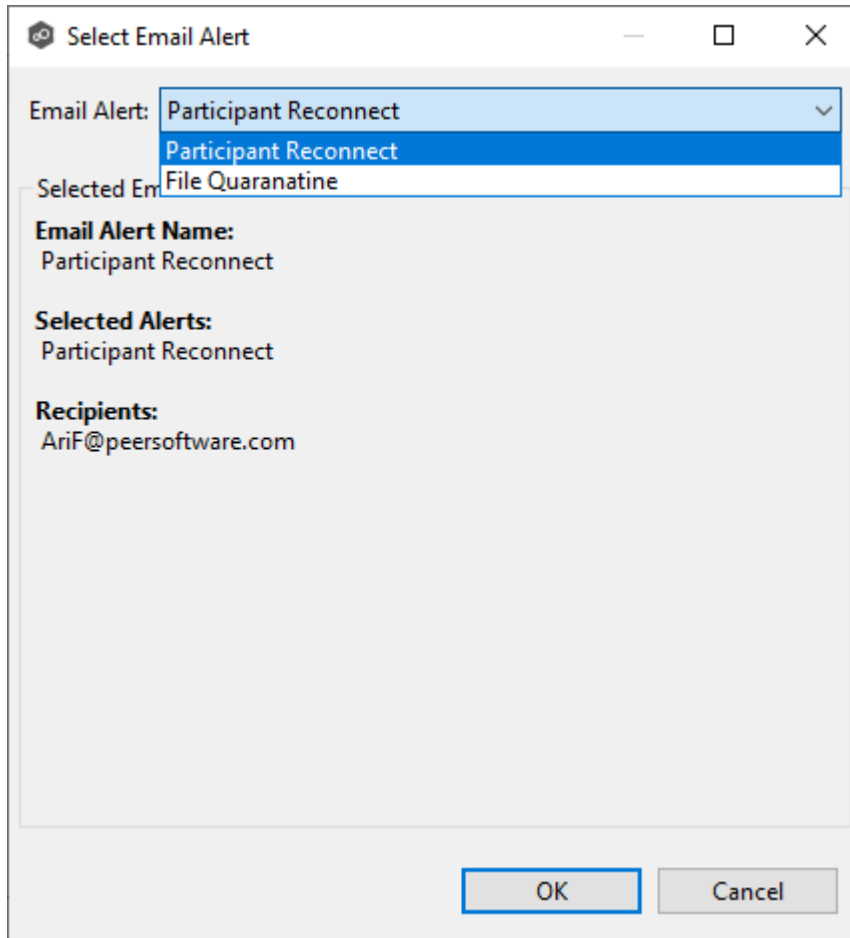
OK Cancel

Modify the fields as needed:


Field	Description
Enabled	Enables target protection.

The **Select Email Alert** dialog opens.

2. Select the email alert from the drop-down list, and then click **OK**.



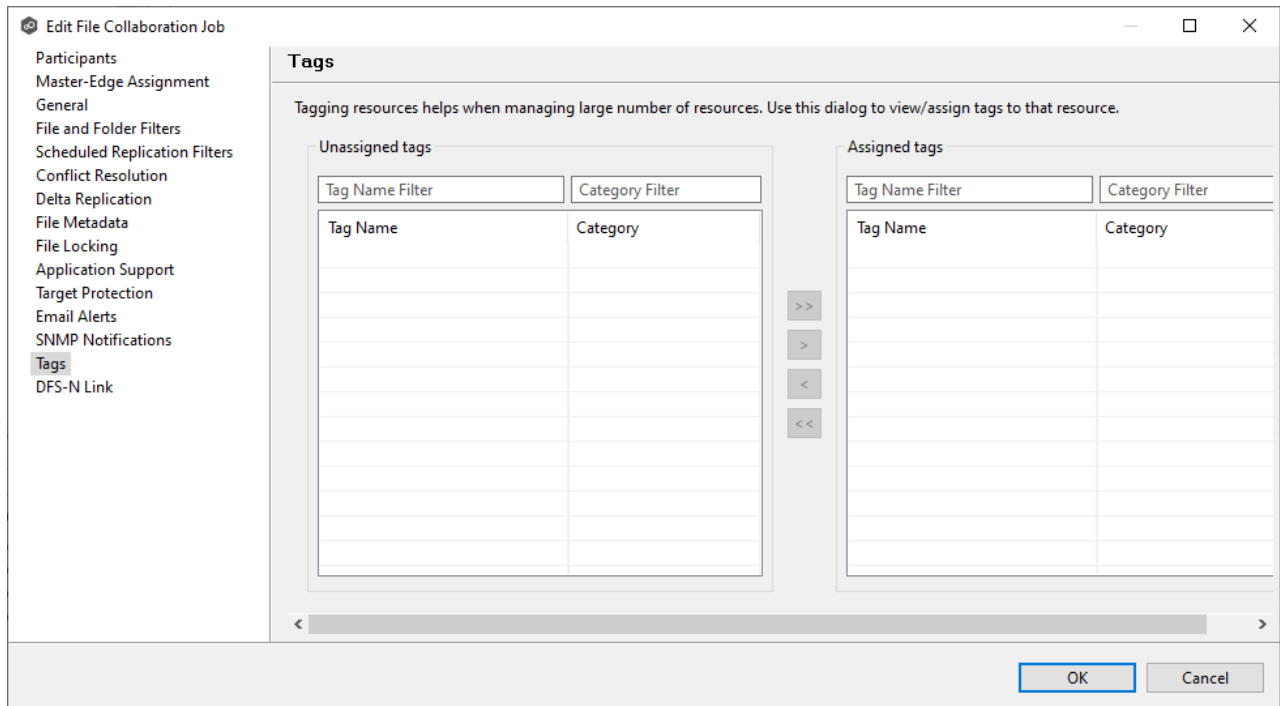
The newly added email alert appears in the **Email Alerts** table.



2. Click **OK** to close the Edit wizard or select another configuration item to modify.

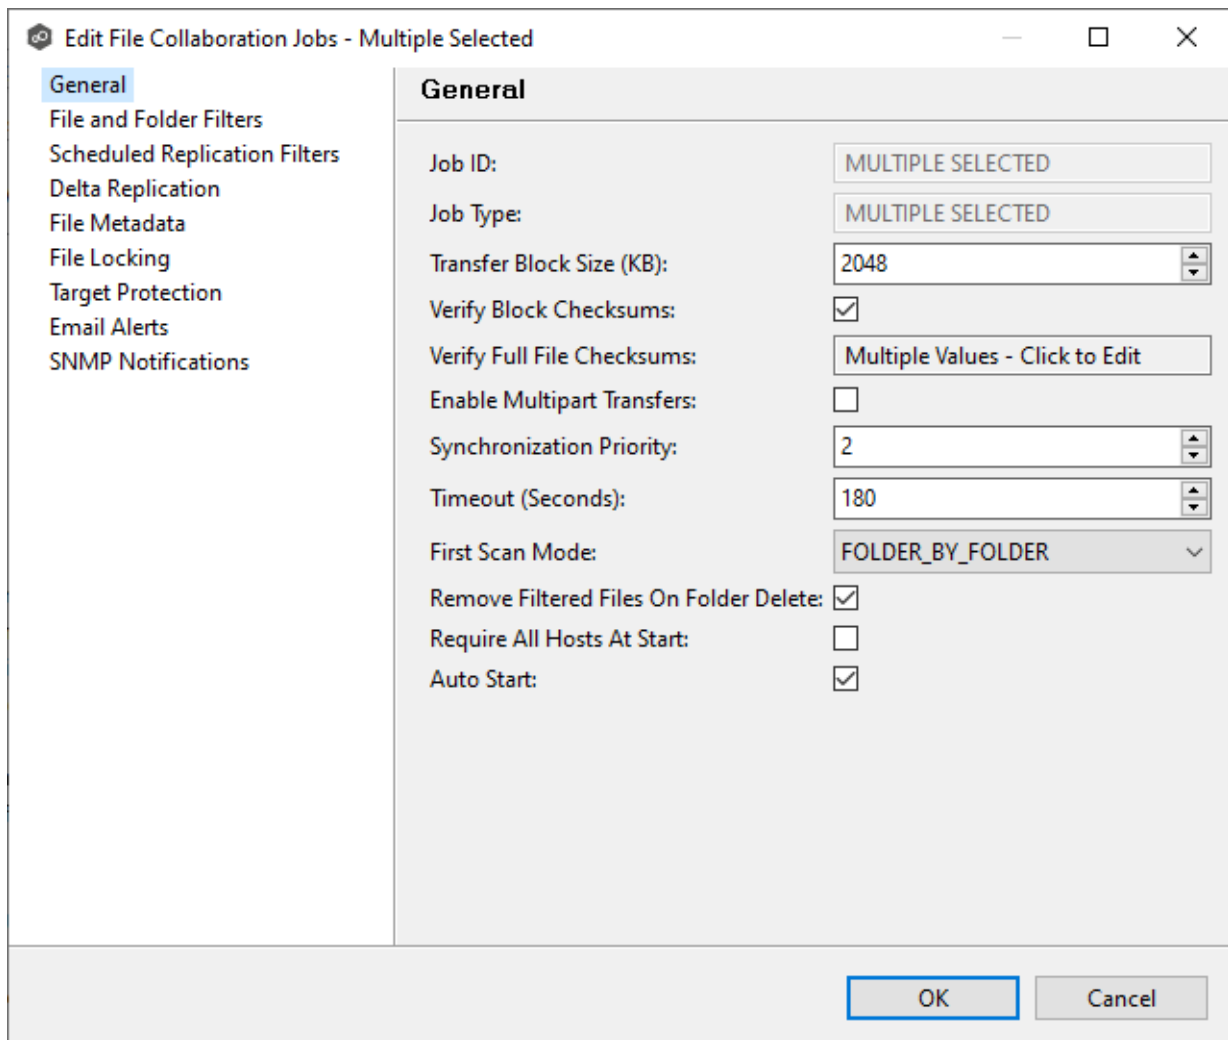
Tags

The **Tags** page in the **Edit File Collaboration Job** dialog allows you to assign existing tags and categories to the selected job. This page is not available in [Multi-Job Editing](#) mode. For more information about tags, see [Tags](#) in the [Basic Concepts](#) section.

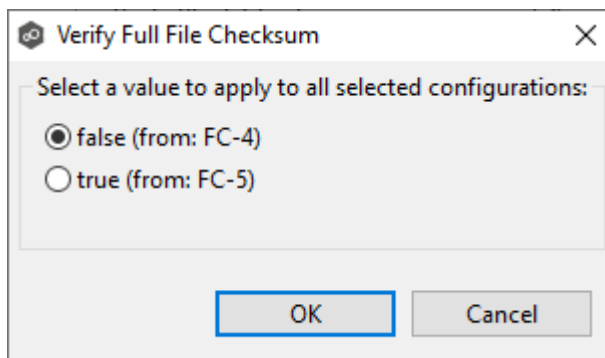


DFS-N Link

The **DFS-N** page in the **Edit File Collaboration Job** dialog presents options for linking a DFS namespace folder to this job. See [Linking a Namespace Folder to an Existing File Collaboration or Synchronization Job](#) for more information.



In this dialog, any discrepancies between multiple selected jobs will generally be illustrated by a read-only text field with the caption **Multiple Values - Click to Edit**. Clicking this field will bring up a dialog similar to the following:



This dialog gives you the option of choosing a value that is already used by one or more selected File Collaboration jobs, in addition to the ability to use your own value. Notice

that variances in the look and feel of this pop-up dialog above depend on the type of information it is trying to represent (for example, text vs. a checkbox vs. a list of items).

Upon clicking **OK**, the read-only text field you originally clicked will be updated to reflect the new value. Any fields that have changed will be marked by a small caution sign. On saving in this multi-job edit dialog, the changed values will be applied to all selected jobs.

Note: Read all information on each configuration page carefully when using the multi-job edit dialog. A few pages operate in a slightly different manner than mentioned above. All the necessary information is provided at the top of these pages in bold text.

Running and Managing a File Collaboration Job

The topics in this section provide some basic information about starting, stopping, and managing File Collaboration jobs:

- [Overview](#)
- [Starting a File Collaboration Job](#)
- [Stopping a File Collaboration Job](#)
- [Auto-Restarting a File Collaboration Job](#)
- [Host Connectivity Issues](#)
- [Removing a File from Quarantine](#)
- [Manual Retries](#)

Overview

This topic describes:

- The [initialization process](#) for a File Collaboration job: What occurs the first time you run a File Collaboration job.
- The [initial synchronization process](#): How files are synchronized the first time you run a File Collaboration job.

The initialization process for a File Collaboration job consists of the following steps:

1. All participating hosts are contacted to make sure they are online and properly configured.
2. [Real-time event detection](#) is initialized on all participating hosts where file locks and changes will be propagated in real-time to all participating hosts. You can view real-time activity and history via the various [Runtime Job views](#) for the open job.
3. The [initial synchronization process](#) is started; all the configured root folders on the participating hosts are scanned in the background, and a listing of all folders and files are sent back to the running job.
4. The background directory scan results are analyzed, and directory structures compared to see which files are missing from which hosts. In addition, file conflict resolution is performed to decide which copy to use as the master for any detected file conflicts based on the [File Conflict Resolution](#) settings.
5. After the analysis is performed, all files that need to be synchronized are copied to the pertinent host(s).

Before you start a File Collaboration job for the first time, you need to decide how you would like the [initial synchronization](#) to be performed.

During the initial synchronization process:

- The watch set is recursively scanned on all participating hosts.
- File conflict resolution is performed.
- Any files that require updating are synchronized with the most current copy of the file.

The two primary options are:

- Have the File Collaboration job perform the initial synchronization based on the [Conflict Resolution](#) settings.

- [Pre-seed](#) all [participating hosts](#) with the correct folder and file hierarchy for the configured root folders before starting the session.

If you have a large data set, we strongly recommend that you perform the initial synchronization manually by copying the data from a host with the most current copy to all other participating hosts. This needs to be done only once--before the first time that you run the File Collaboration job.

If you choose the first option, click the **Start** button to begin [collaboration session initialization](#). Otherwise, pre-seed each participating host with the necessary data, then click the **Start** button.

Starting a File Collaboration Job


Before starting a File Collaboration job for the first time, make sure that you have decided how you want the [initial synchronization](#) to be performed.

When running a File Collaboration job for the first time, you must manually start it. After the initial run, a job will automatically start, even when the Peer Management Center server is rebooted.

Note: You cannot run two jobs concurrently on the same volume if the [watch sets](#) contain an overlapping set of files and folders.


To manually start a job:

1. Choose one of three options:
 - Right-click the job name in the **Jobs** view.
 - Right-click the job name in the **Summary** tab of the **Collab, Sync, and Repl Summary** summary view, and then choose **Start** from the context menu.
 - Open a job and then click the **Start/Stop** button to the left of the **Status** field in the bottom left corner of the job's runtime view (shown below).



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

After the job initialization has completed, the job will run. Once the job starts, the icon next to the job name in the **Jobs** view changes from gray to green.



Stopping a File Collaboration Job

You can stop a File Collaboration job at any time by selecting the job in the **Jobs** view and clicking the **Stop** button. Doing this shuts down the real-time file event detection and closes 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.

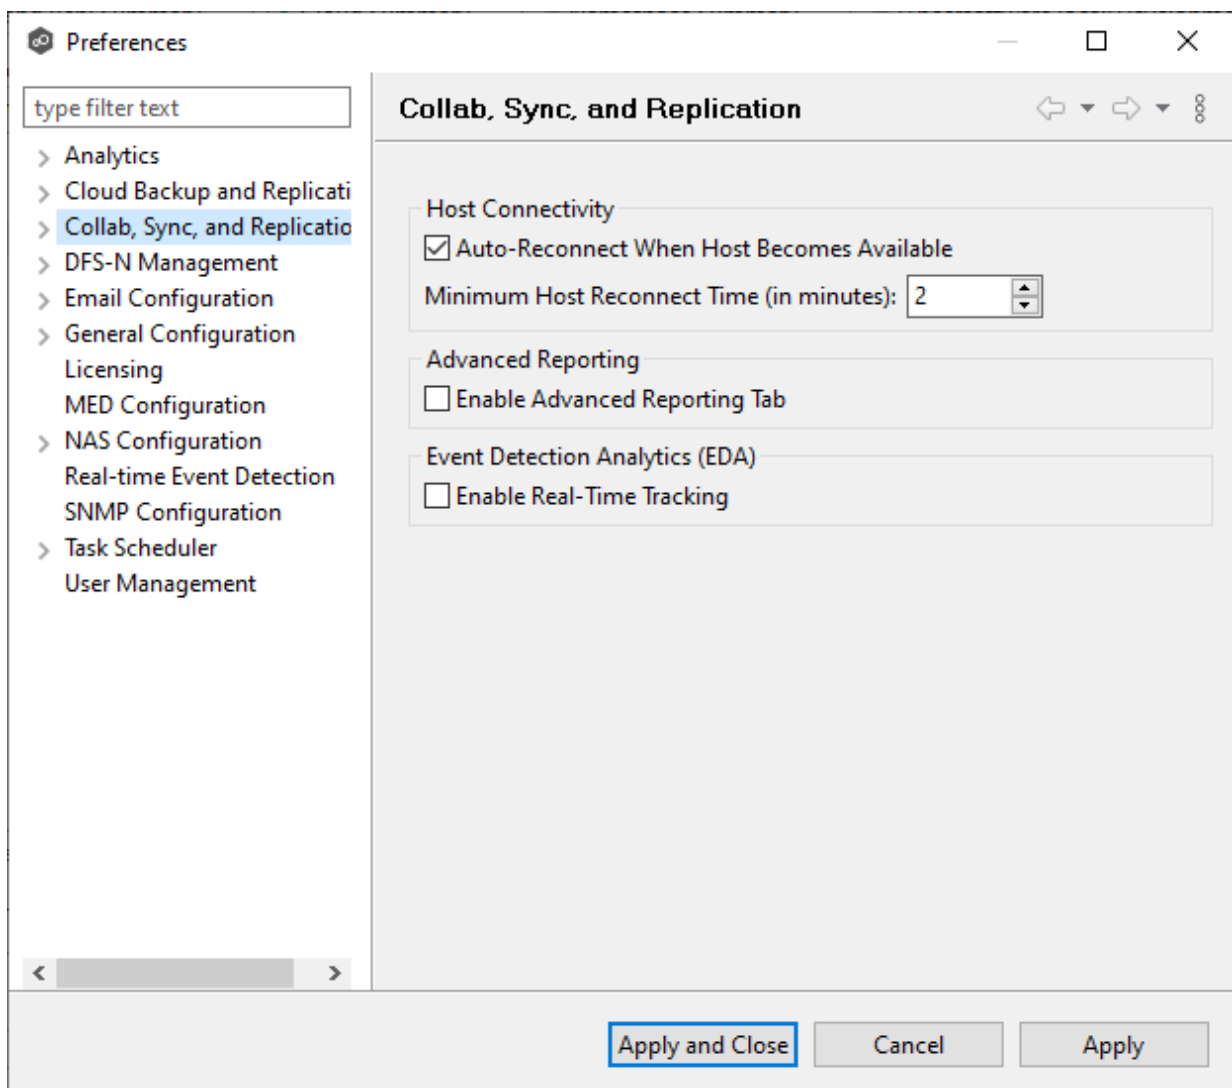
After a host becomes unavailable and the [quorum](#) is lost on a running File Collaboration job, the job automatically stops running and enters a pending state, waiting for one or more hosts to become available again so that the quorum can be met. Once the quorum is met, the pending job will automatically be restarted, beginning with a scan of all root folders.

In a job where a host becomes unavailable, but the quorum is not lost, the remaining hosts will continue collaborating. If the unavailable host becomes available once again, it is brought back into the running job and a background scan begins on all participating hosts, similar in fashion to the initial background scan at the start of a job.

You can enable all File Collaboration jobs to auto-restart. You can also disable auto-restart File Collaboration jobs on a per-job and per-host instance. For more information on disabling auto-restart at the job level, see [Participants Tab](#).

To enable all File Collaboration jobs to auto-restart:

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



3. Select the **Auto Reconnect when Host Becomes Available** checkbox.

4. Enter the minimum number of minutes to wait after an Agent reconnects before re-enabling it in any associated jobs in the **Minimum Host Reconnect Time** field.
5. Click **OK**.

Host Connectivity Issues

Peer Management Center is designed to be run in an environment where all [participating hosts](#) are highly available and on highly available networks. The two primary connectivity issues result from:

- [Unavailable Hosts](#)
- [Quorum Not Met](#)

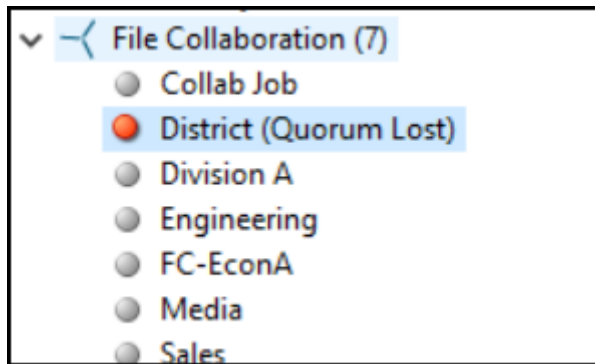
Unavailable Hosts

If a host becomes unavailable while a File Collaboration job is running and is unreachable within the configured timeout period (specified in the job's [General settings](#)), it may be removed from collaboration. If no response is received while performing a file collaboration operation within the timeout period, Peer Management Center pings the host; if still no response, the host is taken out of the running session, a FATAL event is logged, and the [Participants tab](#) for the job is updated to indicate that the host has failed. In addition, if [email alerts](#) and/or [SNMP notifications](#) are configured and enabled for **Host Timeouts**, then the appropriate message(s) are sent.

If [auto-restart](#) not enabled, you must stop and start the File Collaboration job to bring any failed hosts back into the session. As a result, all root folders on all hosts will need to be scanned again to detect any inconsistencies. Therefore, if you are operating over a WAN with low bandwidth, you will want to set the timeout to a higher value on each related job.

Quorum Not Met

For a File Collaboration job to run correctly, a quorum of available hosts must be met. When a quorum is lost, a message appears after the job name in the **Jobs** view.



The quorum is currently set to at least two hosts, and if quorum is not met, then the collaboration session is automatically terminated. If [email alerts](#) and/or [SNMP notifications](#) are configured and enabled for **Session Aborts**, then the appropriate message(s) are sent.

Removing a File from Quarantine

Quarantines are a key feature of Peer Global File Service, used for resolving version conflicts. For more detailed information on how quarantines work, see [Conflicts, Retries, and Quarantines](#) in the [Advanced Topics](#) section.

You must explicitly remove a file from quarantine in order to have it participate in the collaboration session once again.

You may also choose to perform no action, in which case, the file is removed from the **Quarantines** table but none of the file versions are modified. Therefore, if the files are not currently in-sync, then the next time the file is modified, changes will be propagated to the other hosts. If an error occurs while removing the quarantine, then the **Status** field in the **Quarantines** table is updated to reflect the error.

To remove a file (or multiple files) from quarantine:

1. Open the job.
2. The runtime view for the job appears.
3. Click the [Quarantines tab](#).
4. Select the file(s) in the **Quarantines** table.
5. Select the host with the correct version.
6. Click the **Release Conflict** button.

After doing this, all hosts are checked to make sure the file is not currently locked by anyone. If no locks are found, then locks are obtained on all versions of the file and the targets that are out-of-date are synchronized with the selected source host.

Manual Retries

Retries are a key feature of Peer Global File Service, used for automatically handling errors in the collaboration environment that would have otherwise led to a quarantine. For more detailed information on how retries work, see [Conflicts, Retries, and Quarantines](#) in the [Advanced Topics](#) section.

When a file is put in the retry list, it will be automatically retried based on the settings defined in [File Retries](#) in [Preferences](#). If need be, you can also manually force the retry of a file. This can be done from the Retries list of a specific File Collaboration job.

You may also choose to perform no action, in which case, the file is removed from the Retries list but none of the file versions are modified. Therefore, if the files are not currently in-sync, then the next time the file is modified, changes will be propagated to the other hosts. If an error occurs while forcing the retry, then the **Status** field in the **Retries** table is updated to reflect the error.

To manually force the retry of a file (or multiple files):

1. Select the file(s) in the **Retries** list.
2. Select the host with the correct version.
3. Click the **Release Conflict** button.

After doing this, all hosts are checked to make sure the file is not currently locked by anyone. If no locks are found, then locks are obtained on all versions of the file and the targets that are out-of-date are synchronized with the selected source host.

File Replication Jobs

This section provides information about creating a File Replication job.

- [Overview](#)
- [Before You Create Your First File Replication Job](#)

- [Creating a File Replication Job](#)

Overview

A **File Replication** job is designed to push files one way from a single file server (known as the **source storage device**) to one or more file servers (known as the destinations or **target storage devices**). This job type requires an Agent for the source storage device and an Agent for each target storage device. However, only the Agent for the source storage device will register with its local storage platform for real-time activity. The destination Agents will simply act as an intermediary to the destination file server.

Before You Create Your First File Replication Job

We strongly recommend that you configure global settings such as SMTP configuration and email alerts before configuring your first File Replication job. See [Preferences](#) for details on what and how to configure these settings.

Creating a File Replication Job

The **Create Job Wizard** walks you through the process of creating a File Replication job:

[Step 1: Job Type and Name](#)

[Step 2: Source Platform](#)

[Step 3: Source Agent](#)

[Step 4: Storage Information](#)

[Step 5: Source Path](#)

[Step 6: Destination Agent](#)

[Step 7: Destination Path](#)

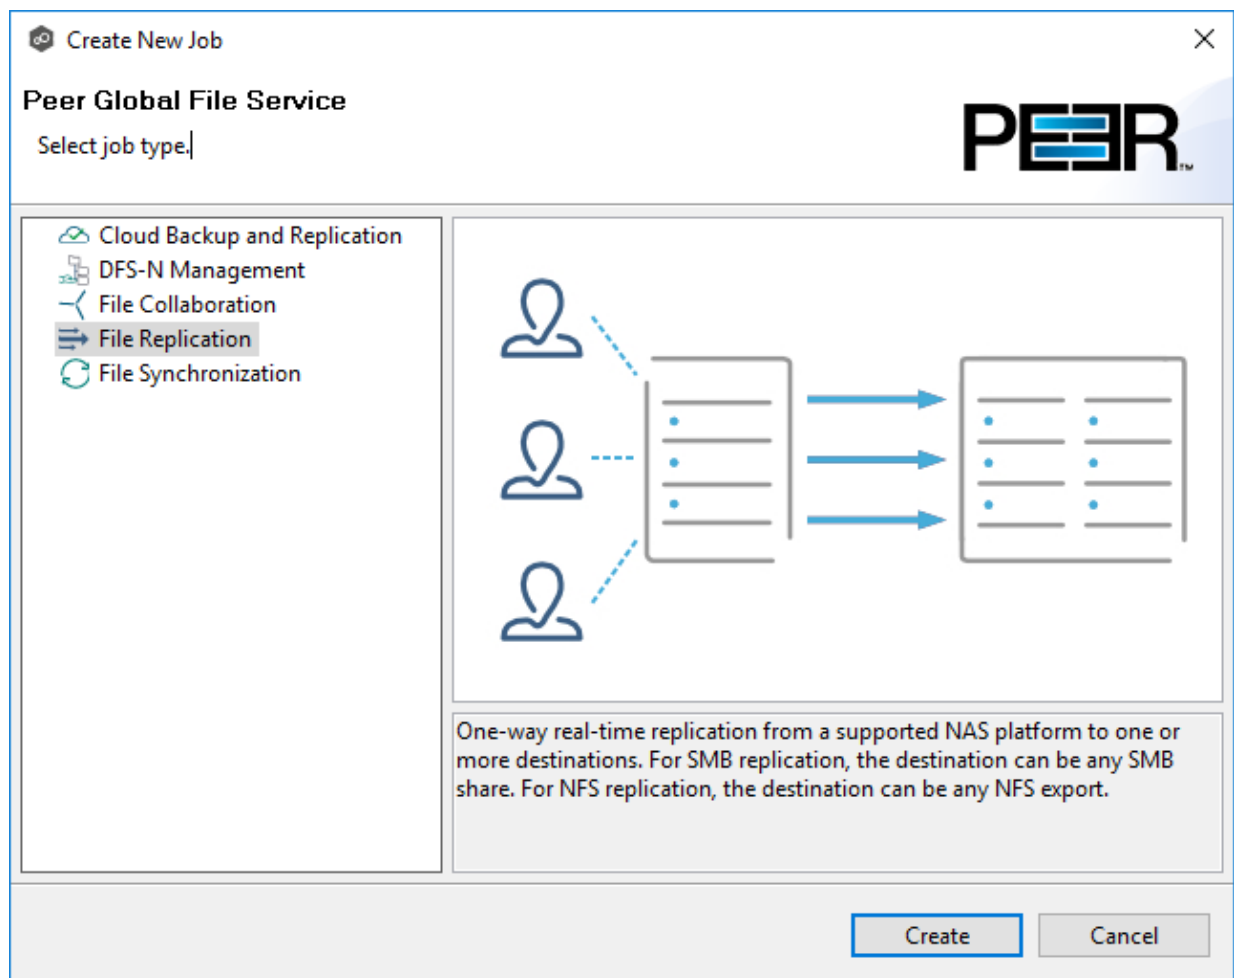
[Step 8: File Metadata](#)

[Step 9: Email Alerts](#)

[Step 10: Save Job](#)**Step 1: Job Type and Name**

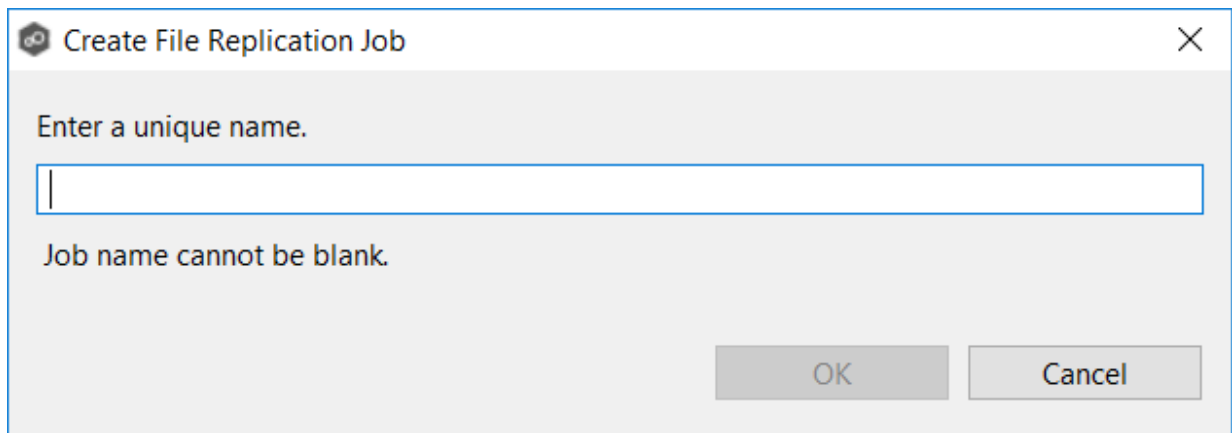
1. Open Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

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



3. Click **File Replication**, and then click **Create**.
4. Enter a name for the job in the dialog that appears.

The job name must be unique.



Create File Replication Job

Enter a unique name.

Job name cannot be blank.

OK Cancel

5. Click **OK** to proceed to the [Source Agent](#) page.

Step 2: Source Agent

Overview

Use the **Source Agent** page to select the Source Agent and file-sharing protocol for the job:

- Choose a file-sharing protocol: **SMB**, **NFS**, or **Multi-Protocol**, based on the workload for this job

Note: Not all licenses support multi-protocol configurations. Verify your license capabilities before configuring jobs that require both SMB and NFS.

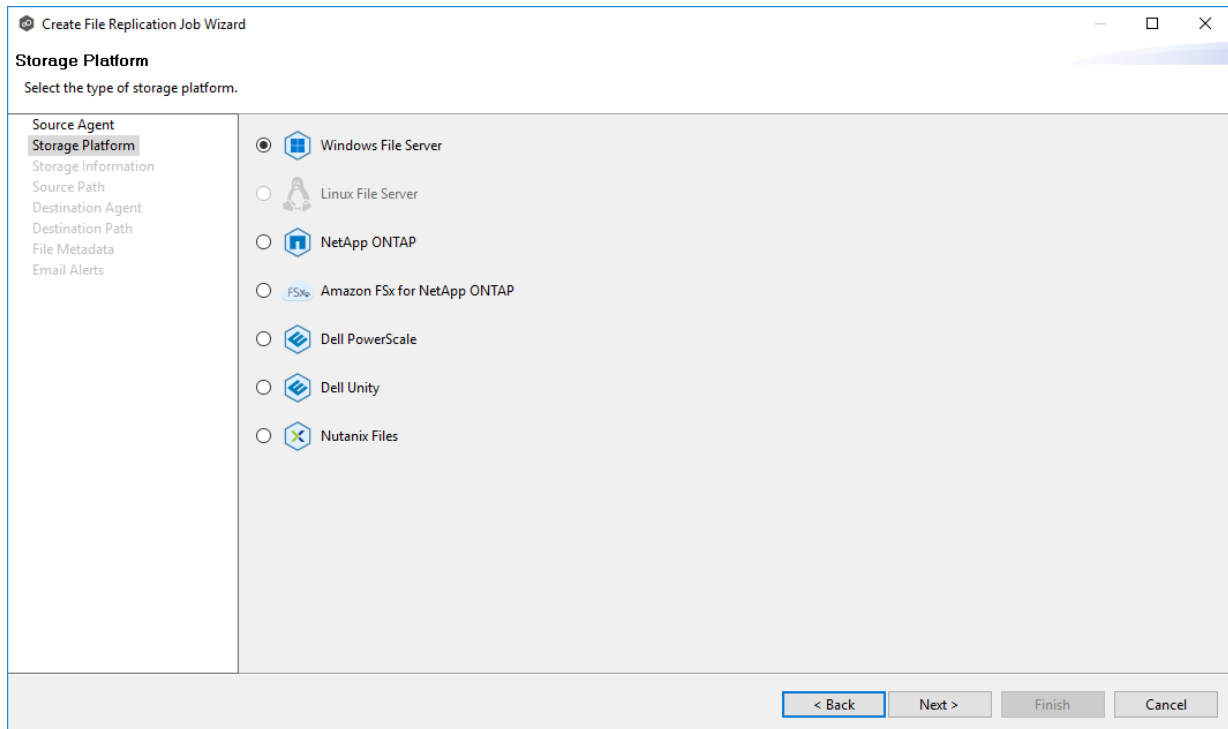
- Select an Agent to manage the data you want to replicate. After selecting a protocol, only compatible Agents will be available:
 - **SMB** requires Windows-based Agents.
 - **NFS** or **Multi-Protocol** requires Linux-based Agents.
- Both Source and Destination Agents must use the same protocol—Windows and Linux Agents cannot be mixed in the same job.
- The wizard will guide you through configuring protocol-specific settings.

Note: You can assign multiple Agents to a storage device, but each must manage different volumes, shares, exports, or folders.

Step 3: Storage Platform

The **Storage Platform** page lists the available storage platforms types. The available storage platform types depends on the protocol and Agent selected on the [Source Agent](#) page.

1. Select the storage platform type that hosts the data you want to replicate.



These are the available platforms:

Protocol/Agent Selection	Available Storage Platform Types
SMB Protocol and Windows-based Agent	<ul style="list-style-type: none"> • Windows File Server • NetApp ONTAP • Amazon FSxN • Dell PowerScale • Dell Unity • Nutanix Files
NFS Protocol and Linux-based Agent	<ul style="list-style-type: none"> • Linux File Server • NetApp ONTAP • Amazon FSxN • Dell PowerScale
Multi-Protocol and Linux-based Agent	<ul style="list-style-type: none"> • NetApp ONTAP • Amazon FSxN • Dell PowerScale

2. Click **Next** to proceed to the the [Storage Information](#) page.

Step 4: Storage Information

On the **Storage Information** page, enter the required details for the storage platform selected on the [Storage Platform](#) page. The fields displayed vary based on the selected platform type.

1. If you selected **Windows File Server** or **Linux File Server**, provide the required configuration details. See [Windows File Server](#) or [Linux File Server](#) for more information.

If you selected any other storage platform type, skip to Step 2.

2. Select **New Credentials** or **Existing Credentials**.

3. Enter credentials:

- If you select **Existing Credentials**, select a entry from the drop-down list, and then click **Next** to proceed to [Step 5. Source Path](#).
- If you select **New Credentials**, enter the required connection details. The fields vary depending on the storage platform type. For details, see:

[Amazon FSxN](#)

[Dell PowerScale](#)

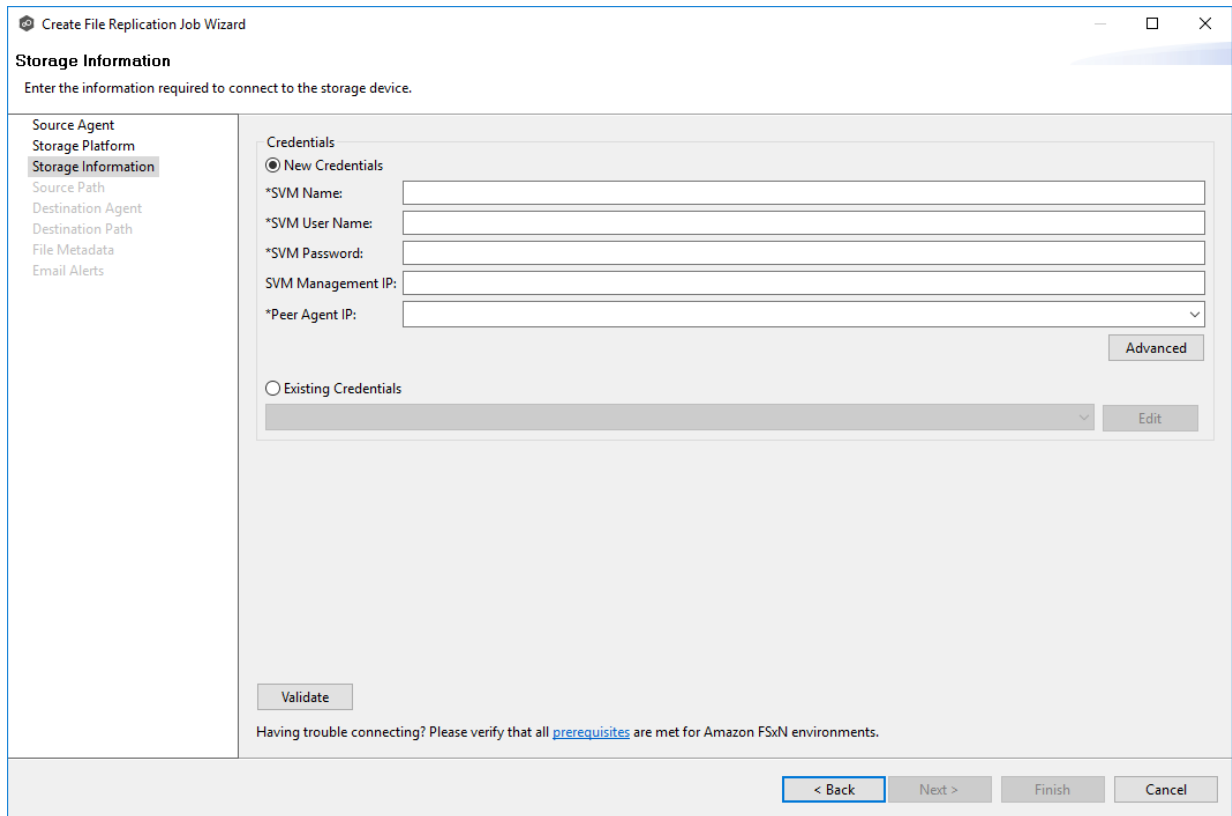
[Dell Unity](#)

[NetApp ONTAP](#)

[Nutanix Files](#)

4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
5. Click **Next** to proceed to the [Source Path](#) page.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine (SVM) hosting the data to be replicated.



Create File Replication Job Wizard

Storage Information
Enter the information required to connect to the storage device.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Credentials

New Credentials

*SVM Name:

*SVM User Name:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

Advanced

Existing Credentials

Edit

Validate

Having trouble connecting? Please verify that all [prerequisites](#) are met for Amazon FSxN environments.

< Back Next > Finish Cancel

2. Enter credentials:

- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Source Path](#) page.
- If you selected **New Credentials**, supply the required information:

Field	Description
SVM Name	Enter the name, FQDN (Fully Qualified Domain Name), or IP address of the Storage Virtual Machine (SVM) hosting the data to be replicated.
SVM User Name	Enter the username for the account managing the SVM. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the SVM. This must not be a cluster management account.
SVM Management IP	Optional. Enter the IP address for accessing the SVM's management API. This field is not required if the data LIFs (Logical Interfaces) associated with the specified SVM Name already permit management access,
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the SVM. The SVM must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter it.

3. Click **Advanced** if you want to set [advanced options](#).
 4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
 5. Click **Next** to proceed to the [Source Path](#) page.
1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the PowerScale cluster hosting the data to be replicated.
 2. Enter credentials:
 - If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Source Path](#) page.
 - If you selected **New Credentials**, supply the required information.

The information required will vary depending on whether you select **RabbitMQ** or **Syslog** as the connection type, due to the distinct protocols and mechanisms they employ for communication.

If you selected **Rabbit MQ**, this page appears:

Create File Replication Job Wizard

Storage Information
Enter the information required to connect to the storage device.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Credentials

New Credentials

*SmartConnect Zone Name:

*Cluster Management IP:

*Cluster Username:

*Cluster Password:

Access Zone:

Connection Type: Syslog RabbitMQ

Existing Credentials

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments.

If you selected **Syslog**, this page appears:

Create File Replication Job Wizard

Storage Information
Enter the information required to connect to the storage device.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Credentials

New Credentials

*SmartConnect Zone Name:

*Cluster Management IP:

*Cluster Username:

*Cluster Password:

Access Zones:

Connection Type: Syslog RabbitMQ

Syslog

*Agent IP Address:

*Listening Port:

*SSL Certificate Path:

*SSL Private Key Path:

SSL Private Key Password:

Existing Credentials

You must enter a Cluster Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments.

Field	Description
SmartConnect Zone Name	Enter the name of the PowerScale cluster hosting the data to be replicated.
Cluster Management IP	Enter the IP address to access the REST-based API integrated into the PowerScale cluster. This field is required only if multiple access zones are in use on the cluster.
Cluster Username	Enter the username for the account managing the PowerScale cluster.
Cluster Password	Enter the password for account managing the PowerScale cluster.
Access Zone	Optional. Enter the name of the access zone being monitored.
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: <ul style="list-style-type: none"> • Syslog: Choose if the storage device directly transmits notifications to the Agent. • RabbitMQ: Choose if you are utilizing the CEE framework to dispatch notifications to the Agent.

3. If you selected **Syslog**, you must also enter values for the following fields:

Field	Description
Agent IP Address	Select the IP address of the server hosting the Agent that manages the PowerScale cluster. The cluster must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.
Listening Port	Enter the port over which the Agent will receive TLS-based syslog events from the PowerScale cluster.

Field	Description
SSL Certificate Path	Enter the path to the certificate to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see Dell PowerScale Syslog Configuration Guide .
SSL Private Key Path	Enter the path to the private key to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see Dell PowerScale Syslog Configuration Guide .
SSL Private Key Password	[Optional] If your private key is protected with a password, enter it here. For more information, see Dell PowerScale Syslog Configuration Guide .

4. Click **Advanced** if you want to set [advanced options](#).
 5. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
 6. Click **Next** to proceed to the [Source Path](#) page.
-
1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the NAS Server hosting the data to be replicated.

The screenshot shows a window titled "Create File Replication Job Wizard" with a sidebar on the left containing the following items: Source Agent, Storage Platform, Storage Information (highlighted), Source Path, Destination Agent, Destination Path, File Metadata, and Email Alerts. The main area is titled "Storage Information" and contains the instruction "Enter the information required to connect to the storage device." Under the "Credentials" section, the "New Credentials" radio button is selected. Below it are four text input fields labeled: "*NAS Server Name:", "*Unisphere Management IP:", "*Unisphere Username:", and "*Unisphere Password:". To the right of these fields is an "Advanced" button. Below the "New Credentials" section is the "Existing Credentials" radio button, which is unselected. Below it is a greyed-out dropdown menu and an "Edit" button. At the bottom left of the main area is a "Validate" button. To its right is a validation message: "You must enter a NAS Server Name." Below this message is a link: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments." At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. Enter credentials:

- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to proceed to the [Source Path](#) page.
- If you selected **New Credentials**, supply the required information.

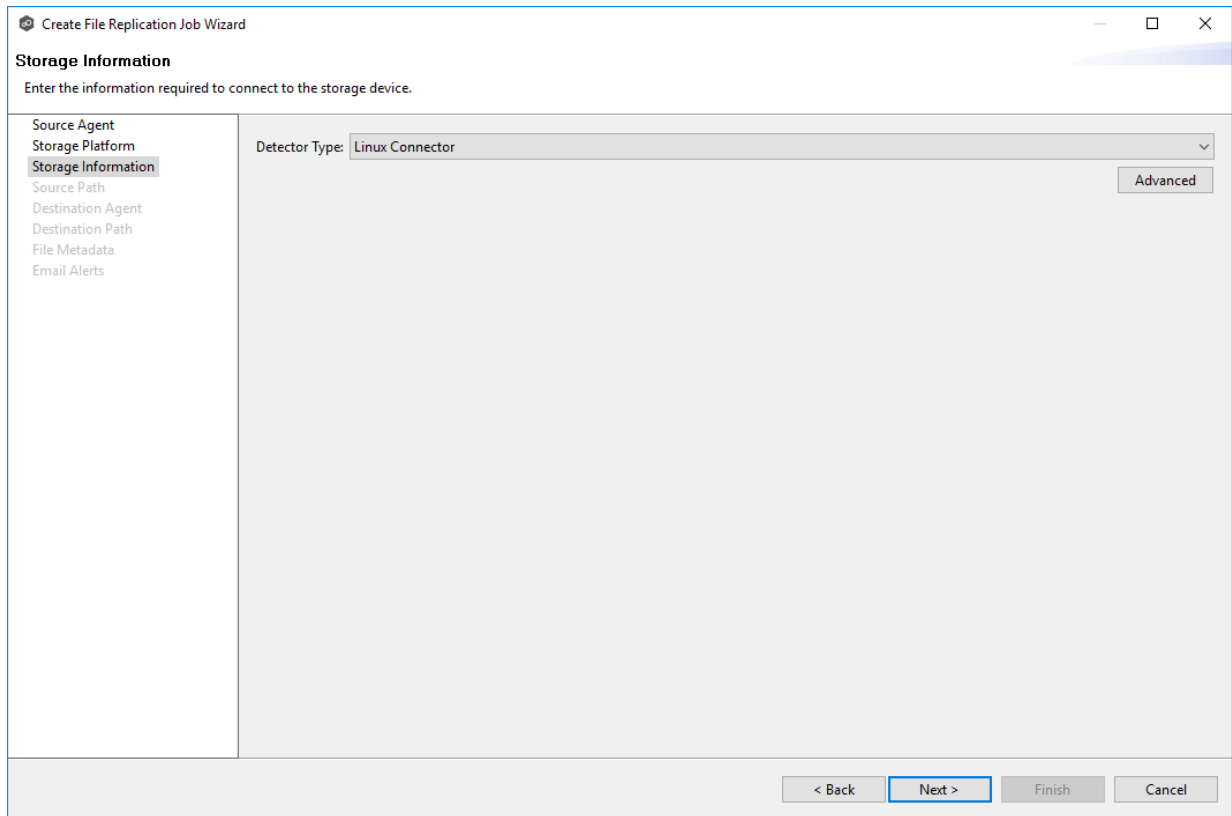
Field	Description
NAS Server Name	Enter the name of the NAS server hosting the data to be replicated.
Unisphere Management IP	Enter the IP address of the Unisphere system that manage the Unity storage devices. The IP address should not point to the NAS server.
Unisphere Username	Enter the username for the Unisphere account managing the Unity storage device.
Unisphere Password	Enter the password for the Unisphere account managing the Unity storage device.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
5. Click **Next**.

The [Source Path](#) page is displayed.

1. Select the **Detector Type**.

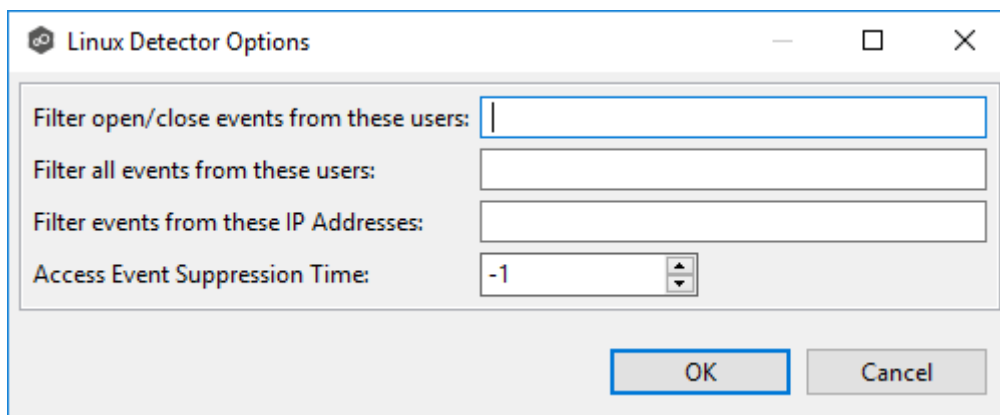
Detector type refers to the type of storage platform that is being monitored for real-time file activity.



2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Next**.

Linux File Server Advanced Options

1. Modify the options as desired.



Option	Description
Filter open/close events from these users	Enter a comma-separated list of user account names from which all file open and close events will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user account under which a backup or archival service is running.
Filter all events from these users	Enter a comma-separated list of user account names from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user account under which a backup or archival service is running.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the IP addresses on which a backup or archival service is running.
Filter events from these local processes	Enter a comma-separated list of local process names on the Agent server from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the specific process names under which a backup or archival service is running.
Access Event Suppression Time	Enter the number of seconds to delay an open event before being processed. Use this option to reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time an open event will be delayed based on the load of the system.

2. Click **OK**.

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine hosting the data to be replicated.

The screenshot shows a wizard window titled "Create File Replication Job Wizard" with a sidebar on the left containing the following steps: Source Agent, Storage Platform, **Storage Information** (highlighted), Source Path, Destination Agent, Destination Path, File Metadata, and Email Alerts. The main area is titled "Storage Information" and contains the instruction "Enter the information required to connect to the storage device." Under the "Credentials" section, there are two radio buttons: "New Credentials" (selected) and "Existing Credentials". The "New Credentials" section includes five input fields: "*SVM Name:", "*SVM User Name:", "*SVM Password:", "SVM Management IP:", and "*Peer Agent IP:". The "Existing Credentials" section has a drop-down menu and an "Edit" button. A "Validate" button is located at the bottom left of the main area. Below the "Validate" button is a note: "Having trouble connecting? Please verify that all [prerequisites](#) are met for NetApp ONTAP environments." At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. Enter credentials:

- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Source Path](#) page.
- If you selected **New Credentials**, supply the required information.

Field	Description
SVM Name	Enter the name, FQDN (Fully Qualified Domain Name), or IP address of the Storage Virtual Machine (SVM) hosting the data to be replicated.
SVM User Name	Enter the username for the account managing the SVM. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the SVM. This must not be a cluster management account.
SVM Management IP	Optional. Enter the IP address for accessing the SVM's management API. This field is not required if the data LIFs (Logical Interfaces) associated with the specified SVM Name already permit management access,
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the SVM. The SVM must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter it.

3. Click **Advanced** if you want to set [advanced options](#).
 4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
 5. Click **Next** to proceed to the [Source Path](#) page.
1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Nutanix Files cluster hosting the data to be replicated.

Create File Replication Job Wizard

Storage Information
Enter the information required to connect to the storage device.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Credentials
 New Credentials
*Nutanix File Server Name:
*Username:
*Password:
*Peer Agent IP:
 Existing Credentials

Having trouble connecting? Please verify that all [prerequisites](#) are met for Nutanix Files environments.

2. Enter credentials:

- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next** to continue to the [Source Path](#) page.
- If you selected **New Credentials**, supply the required information.


Field	Description
Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the username for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Enter the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually it. The IP address should not point to the Files cluster itself.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
5. Click **Next** to proceed to the [Source Path](#) page.

1. Select the **Detector Type**.

Detector type refers to the type of storage platform that is being monitored for real-time file activity.

- Select **Windows Driver** for more robust logging and better performance (Recommended).
- Select **Windows** if suggested by Peer Technical Support.



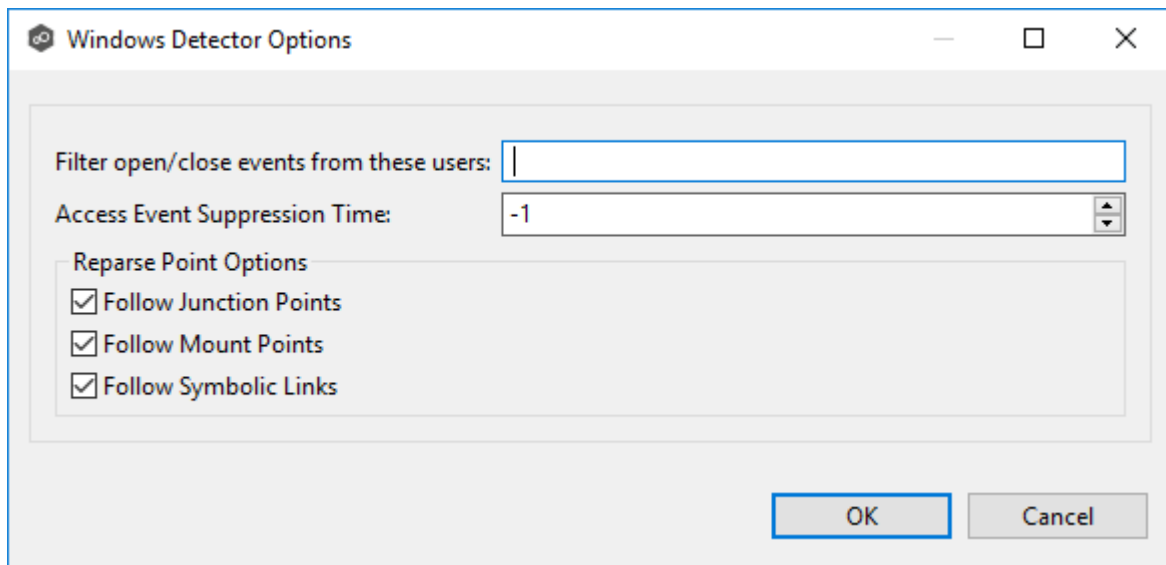
2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Next**.

Windows File Server Advanced Options

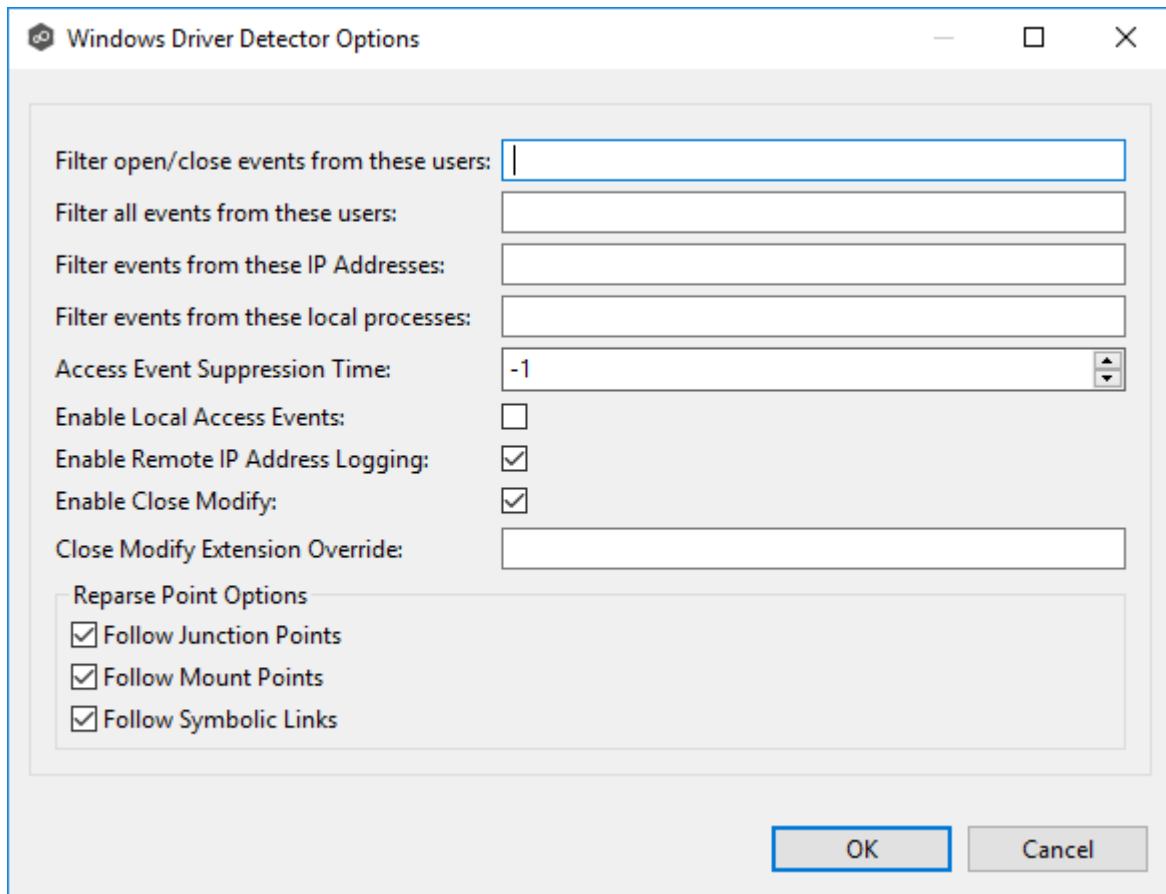
1. Modify the options as desired.

The available options depend on the detector type selected: **Windows** or **Windows Driver**.

Windows



Windows Driver



Option	Description
Filter open/close events from these users	Enter a comma-separated list of user account names from which all file open and close events will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user account under which a backup or archival service is running.
Filter all events from these users	Enter a comma-separated list of user account names from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user account under which a backup or archival service is running.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the IP addresses on which a backup or archival service is running.
Filter events from these local processes	Enter a comma-separated list of local process names on the Agent server from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the specific process names under which a backup or archival service is running.
Access Event Suppression Time	Enter the number of seconds to delay an open event before being processed. Use this option to reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time an open event will be delayed based on the load of the system.
Enable Local Access Events	Enable tracking of opens and closes that are performed locally on the Agent server.
Enable Remote IP Logging	Enable logging of client IP addresses for all real-time activity.
Enable Close Modify	When enabled, no modify or write events will be detected. Instead, replication of a modified file will be triggered when the file is closed.
Close Modify Extension Override	Enter a comma-separated list of exclusions for the Enable Close Modify option. All modify/write events will be detected for these files. This is important for those who rely on sync-on-save functionality.

For more information about junction points or symbolic links, contact [<% SUPPORT_EMAIL%](#)

2. Click **OK**.

Step 5: Source Path

The **Source Path** page allows you to specify the path to the volume, share, export, or folder you want to replicate. This path defines the [watch set](#), which can include only one volume, share, export, or folder. If you want to replicate multiple volumes, shares, exports, or folders, you need to create a separate job for each one.

1. Enter or browse to the watch set:

Create File Replication Job Wizard

Source Path

You must enter a valid path.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Enter Path

< Back

For SMB:

If the source storage device is:

- A Windows File Server, enter the local path to the watch set in the following format:
C:\data
- A NAS device, enter a path to the watch set in the following UNC format: \
\server_name_or_ip\shared_folder\file_or_directory_path

For NFS:

If the source storage device is:

- A Linux File Server, enter the local path to the watch set in the following format: `/mnt/data`
Note that the path is case-sensitive.
- A NAS device, enter the export path to the watch set in the following format:
`host:/volume`

For Multi-Protocol:

- Enter the export path to the watch set in the following format: `host:/volume`
2. If you are creating a multi-protocol job, you may need to specify SMB access. If not, skip to Step 4.

Create File Replication Job Wizard

Source Path

Path is missing the export name.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

PTSV93A:/ Enter Path Browse

SMB Access

⚠ Required only when synchronizing Windows-style permissions.

[Edit SMB Access Configuration](#)

SMB Access Configuration: ▼

SMB Access Share:

Use Admin Share (CS)

Specify Share Name:

< Back Next > Finish Cancel

SMB access configuration enables the Linux Agent to access data on the file server using the SMB file sharing protocol. It is required in the following situations when creating a multi-protocol job:

- Synchronizing Windows-style permissions
- Synchronizing both Windows-style permissions and Linux-styled permission (mixed permissions)

- Using Nutanix Files.

You will select how you want to synchronize permissions when you get to the **File Metadata** page in this wizard.

Recommended: Create SMB access configurations in advance in the **Storage Platforms** section of **Preferences**.

3. Specify the SMB access configuration:

- To use an **existing** SMB access configuration:

- a. Select it from the **SMB Access Configuration** drop-down.
- b. Select an **SMB Access Share** option:

This share will be used by the Agent to access your data via the SMB protocol. Any Windows-style permissions will be read and written to via the share.

- **Use Admin Share:** Select this option if you have an admin share that you would like to use. The same admin share can often be used for multiple jobs.
 - **Specify Share Name:** If you want to or need a dedicated share for the job, select this option and enter the share name.
- To create a **new** SMB access configuration:
 - a. Click **Edit SMB Access Configuration**.
 - b. In the dialog that appears, click **Create**.


Field	Description
Domain Username	Enter the domain username in the format: Domain\Username
Domain Password	Enter the password for the domain account.

- d. Click **OK** to save the access configuration.
 - e. Click **Apply and Close** to apply the configuration to the job and close the dialog.
4. Click **Next** to proceed to the [Destination Agent](#) page.

Step 6: Destination Agent

The **Destination Agent** page lists available Agents, excluding the designated Source Agent. This Destination Agent is responsible for writing files and metadata to the destination storage device. No credentials are required since it does not perform real-time monitoring.

1. Select the Agent that manages the destination storage device. If the destination is a Windows File Server, the Agent must be installed on it.



For SMB:

If the destination storage device is:

- A Windows File Server, enter the local path in the following format: `C:\data`
- A NAS device, enter a path in the following UNC format: `\server_name_or_ip\shared_folder\file_or_directory_path`

For NFS:

If the destination storage device:

- A Linux File Server, enter the local path in the following format: `/mnt/data`
Note that the path is case-sensitive.
- A NAS device, enter the export path in the following format: `host:/volume`

For Multi-Protocol:

- Enter the export path to the watch set in the following format: `host:/volume`

Note:

- 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.
 - If you enter the start of an NFS path and click **Browse**, the **Folder Browser** dialog will attempt to present a list of the available exports on the specified NFS server. However, you may need to enter the full NFS path manually.
2. If you are creating a multi-protocol job, you may need to specify SMB access. If not, skip to Step 5.

SMB access configuration enables the Linux Agent to access data on the file server using the SMB file sharing protocol. It is required in the following situations when creating a multi-protocol job:

- Synchronizing Windows-style permissions
- Synchronizing both Windows-style permissions and Linux-styled permission (mixed permissions)
- Using Nutanix Files.

You will select how you want to synchronize permissions when you get to the **File Metadata** page in this wizard.

Recommended: Create SMB access configurations in advance in the **Storage Platforms** section of **Preferences**.

- c. In the **New SMB Access Configuration** dialog, supply the requested information.

Field	Description
Description	Enter a description that identifies this configuration.
Domain Username	Enter the domain username in the format: <code>Domain\Username</code>
Domain Password	Enter the password for the domain account.

- d. Click **OK** to save the access configuration.
- e. Click **Apply and Close** to apply the configuration to the job and close the dialog.
4. Enter a value for **SMB Path**.

The **SMB Path** field is a UNC path that points to the same data on the destination file server as the **NFS Path**, but uses the SMB file-sharing protocol. The file server name is automatically populated based on your destination file server in the **NFS Path**, and the file server share is prefilled based on your **SMB Access Share**. Depending on your file server's configuration, you may need to add additional text to the prefilled server name and share. See [SMB Path](#) for details and an example.

5. Click **Next** to proceed to the [File Metadata](#) page.

The **SMB Path** field is a UNC path that points to the same data on the destination file server as the **NFS Path**, but uses the SMB file-sharing protocol. It combines the following elements:

- **File Server Name or IP:** Automatically populated with the file server name from the **NFS Path** field.
- **SMB Access Share:** Automatically populated with your selected **SMB Access Share**.
- **Relative Path under the SMB Access Share:** Depending on your file server's configuration, this relative path may not be necessary.

To set the relative path, use the following example:

- **File Server Name:** `fs1`
- **NFS-formatted Destination Path:** `fs1:/myexport/data`
- **SMB Access Share:** `myshare` (mapped to `/myexport` on the server)

Solution: Add `/data` to the end of **SMB Path**.

This results in:

- **SMB Path:** `\\fs1\myshare\data`
- **NFS Path:** `fs1:/myexport/data`

Both paths now point to the same directory on the file server.

Step 8: File Metadata

This step is optional.

The **File Metadata** page allows you to specify whether you wish to replicate file security metadata and the types of metadata for synchronization. Additionally, it provides the ability to designate the metadata source (volume, share, export, or folder) to resolve conflicts during the initial synchronization. This designated source, utilized in case of conflicts, is referred to as the [master host](#). For more information on synchronizing file security metadata, see [File Metadata Synchronization](#) in [Advanced Topics](#).

The contents of the File Metadata page vary depending on the file-sharing protocol you selected and whether you selected Windows-based or Linux-based Agents for your job:

- If you selected SMB protocol and a Windows-based Agent, proceed with [SMB File Metadata](#).

- If you selected a NFS protocol and a Linux-based Agent, proceed with [NFS File Metadata](#).
- If you selected a Multi-Protocol and a Linux-based Agent, proceed with [Multi-Protocol File Metadata](#).

1. **Select when to replicate metadata** (one or both options):

- **Real-time synchronization:** Metadata changes (Owner, DACL, SACL) will be synchronized with the target host as they occur.
- **Initial scan synchronization:** Metadata will be synchronized during the initial scan.

Create File Replication Job Wizard

File Metadata
Configure the replication of security permissions.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Synchronize File Security Information

Enable synchronizing file security information in real-time

Enable synchronizing file security information with master host during initial scan

Synchronize Security Options

Owner

DACL: Discretionary Access Control List

SACL: System Access Control List

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

2. **Confirm metadata synchronization.**

If you selected either option, click **OK** in the confirmation message.

3. **Select security descriptor components to be synchronized:**

- Choose which components (Owner, DACL, SACL) to synchronize.
4. **Select the master host (if applicable):**
 - If synchronizing metadata during the initial scan, choose a [master host](#) to resolve conflicts.
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If security descriptor components differ across participants, [conflict resolution](#) will use the master host's permissions.
 - If the file does not exist on the master host, a participant will be chosen randomly.
 5. Click **Next** to proceed to the [Email Alerts](#) page.
-
1. **Select when to replicate metadata** (one or both options):
 - **Real-time synchronization:** Metadata changes (Owner, Group, Linux Permissions, ACL) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.
- Note:** Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.

File Metadata

You must select a master host when enabling security or attributes synchronization during initial scan.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Synchronize File Security Information

- Enable synchronizing file security information in real-time
- Enable synchronizing file security information with master host during initial scan

Synchronize Security Options

- Owner
- Group
- Linux Permissions
- ACL (NFSv4 or POSIX)

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

2. Confirm metadata synchronization.

If you selected either option, click **OK** in the confirmation message.

3. Select access controls to be synchronized:

- Choose which components (Owner, Group, Linux Permissions, ACL) to synchronize.

4. Select the master host (if applicable):

- If synchronizing metadata during the initial scan, choose a [master host](#) to resolve conflicts.
- If no master host is selected, metadata synchronization will not occur during the initial scan.
- If access controls differ across participants, [conflict resolution](#) will use the master host's permissions.
- If the file does not exist on the master host, a participant will be chosen randomly.

5. Click **Next** to proceed to the [Email Alerts](#) page.

1. **Select when to replicate metadata** (one or both options):

- **Real-time synchronization:** Metadata changes are synchronized with the target host as they occur.
- **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.

2. **Confirm metadata synchronization.**

If you selected either option, click **OK** in the confirmation message.

3. **Select the security style:**

- Windows-style permissions
- Linux-style permissions
- Mixed permissions (available only with Dell PowerScale)

4. **Select permissions to be synchronized:**

For a general discussion of synchronizing permissions in a multi-protocol job, see the Peer Software knowledge base article [How Do Permissions Work with Multi-Protocol?](#)

- Windows-style permissions

Create File Replication Job Wizard

File Metadata

You must select a master host when enabling security or attributes synchronization during initial scan.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Synchronize File Security Information

Enable synchronizing file security information in real-time

Enable synchronizing file security information with master host during initial scan

Synchronize Security Options

Security Style: **Windows-style permissions** [How do permissions work with Multi-Protocol?](#)

Owner

DACL: Discretionary Access Control List

Validate

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

- Linux-style permissions

Create File Replication Job Wizard

File Metadata

You must select a master host when enabling security or attributes synchronization during initial scan.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Synchronize File Security Information

Enable synchronizing file security information in real-time

Enable synchronizing file security information with master host during initial scan

Synchronize Security Options

Security Style: **Linux-style permissions** [How do permissions work with Multi-Protocol?](#)

Owner

Group

Linux Permissions

ACL (NFSv4 or POSIX)

Validate

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

- Mixed permissions (available only with Dell PowerScale)

5. Select the master host (if applicable):

- If synchronizing metadata during the initial scan, choose a [master host](#) to resolve conflicts.
- If no master host is selected, metadata synchronization will not occur during the initial scan.
- If access controls differ across participants, [conflict resolution](#) will use the master host's permissions.
- If the file does not exist on the master host, a participant will be chosen randomly.

6. Click **Next** to proceed to the [Email Alerts](#) page.

Step 9: Email Alerts

This step is optional.

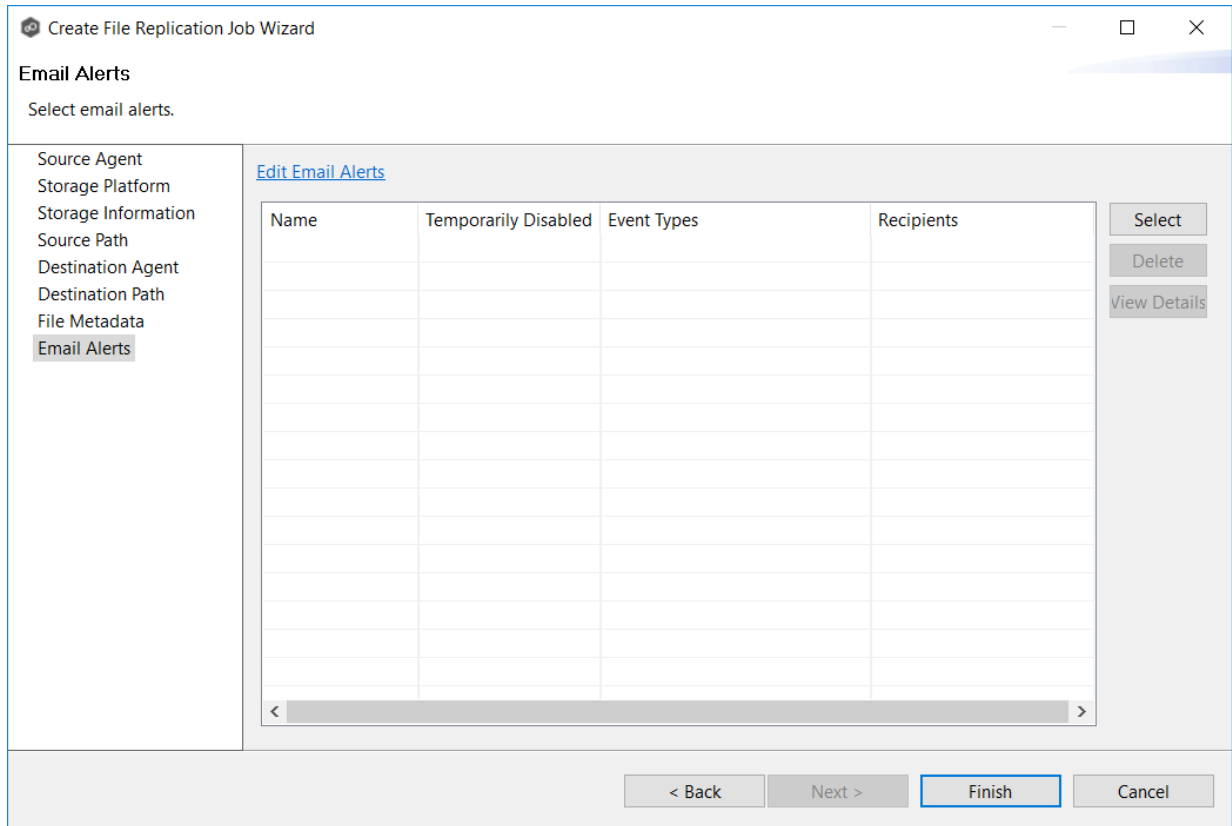
An [email alert](#) notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in [Preferences](#) and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply.

To create a new alert, see [Email Alerts](#) in the [Preferences](#) section.

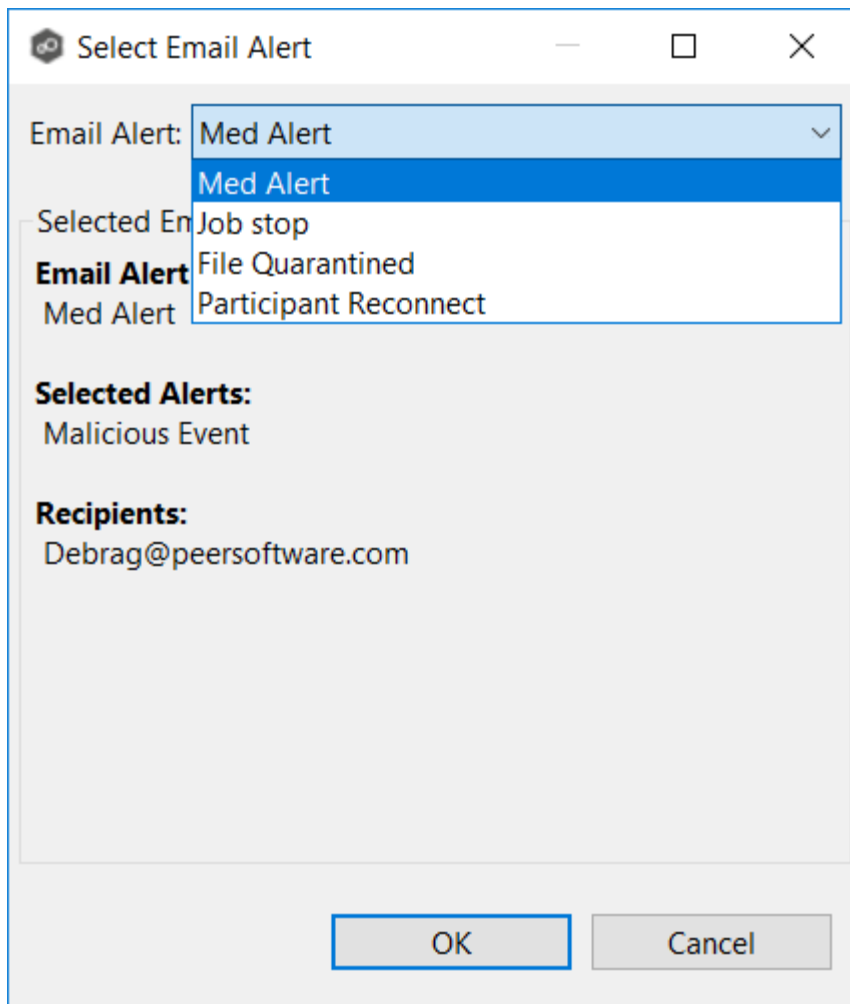
To apply an existing email alert to the job.

1. Click the **Select** button.



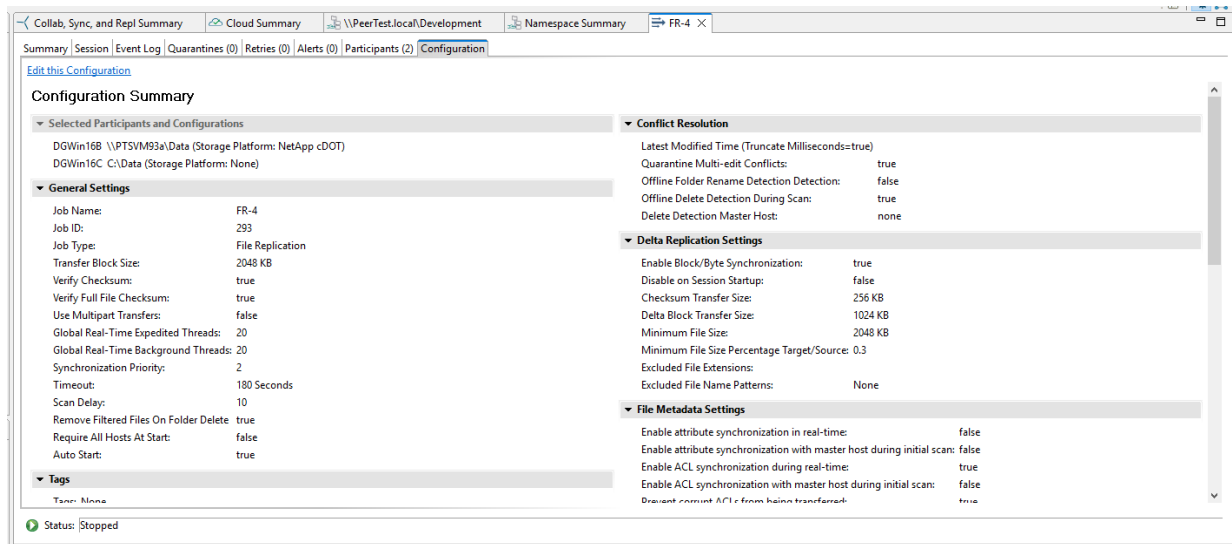
The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.



3. Click **OK**.

The alert is listed on the **Email Alerts** page.



Editing a File Replication Job

You can edit a File Replication job while it is running; however, any changes will not take effect until the job is restarted.

Overview

When you create a File Replication job, the **Create Job** wizard guides you through the process, presenting the most common options for configuration. When editing a job, you have access to all options, allowing you to fine-tune the job configuration. Options not included in the initial job creation include:

- [Application Support](#)
- [Conflict Resolution](#)
- [Delta Replication](#)
- [File and Folder Filters](#)
- [File Locking](#)
- [General](#)
- [Scheduled Replication Filters](#)
- [SNMP Notifications](#)

- [Delta Replication](#)
- [File Metadata](#)
- [File Locking](#)
- [Application Support](#)
- [Target Protection](#)
- [Email Alerts](#)
- [SNMP Notifications](#)
- [Tags](#)

4. Click **OK** when finished.

Participants

The **Participants** page in the **Edit File Replication Job** dialog allows you to:

- [Add and remove participants from a job.](#)
- [Edit a participant.](#)

This topic describes [adding](#) and [deleting](#) participants in a File Replication job.

Adding a Participant to a File Replication Job

To add a participant to a File Replication job:

1. Select the job in the **Jobs** view; right click and select **Edit Job**.

The **Edit File Replication Job** dialog opens; the **Participants** page displays the current job participants.

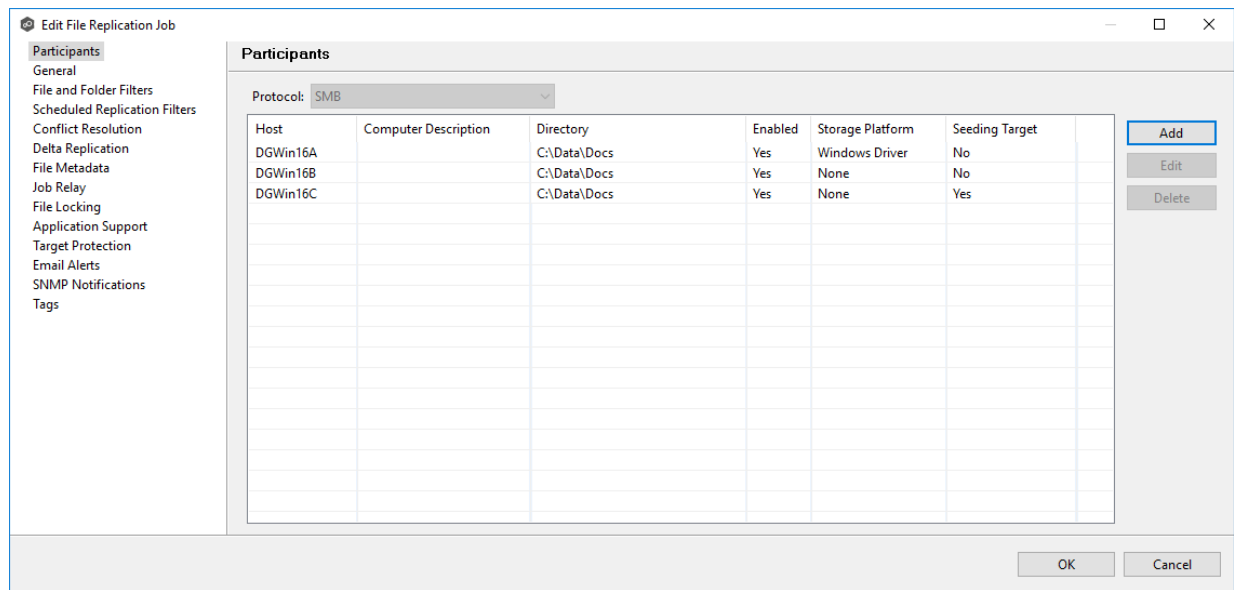
The screenshot shows a window titled "Add New Participant" with a "Destination Path" section. A red error icon and message state "You must enter a valid path." The "Destination Path" section has a "Destination Agent" dropdown menu and a "Destination Path" input field. The input field contains the placeholder text "Enter the path in Windows or UNC format" and a "Browse" button. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

4. Enter or browse to the path to the [watch set](#):

- If the destination storage device is a Windows file server, this path should be a local path such as D:\Data, or it can be the UNC path to any SMB-capable file server.
- If the destination storage device is a Linux-based device, the path should be in NFS format. It should point to an NFS export or a subfolder under the export (e.g. server:/export-name/subfolder).

5. Click **Finish**.

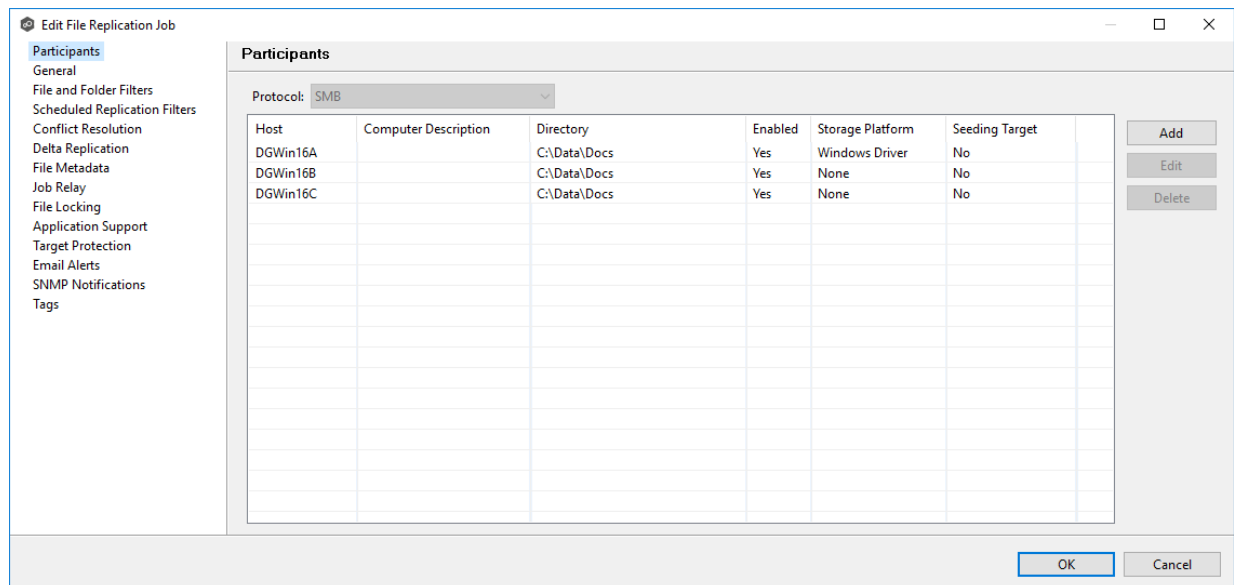
The new participant appears in the **Participants** table.



Deleting a Participant from a File Replication Job

To delete a participant from a File Replication job:

1. In the **Edit File Replication Job** dialog, select the participant in the **Participants** table you want to remove from the job. You cannot remove the Source Agent.



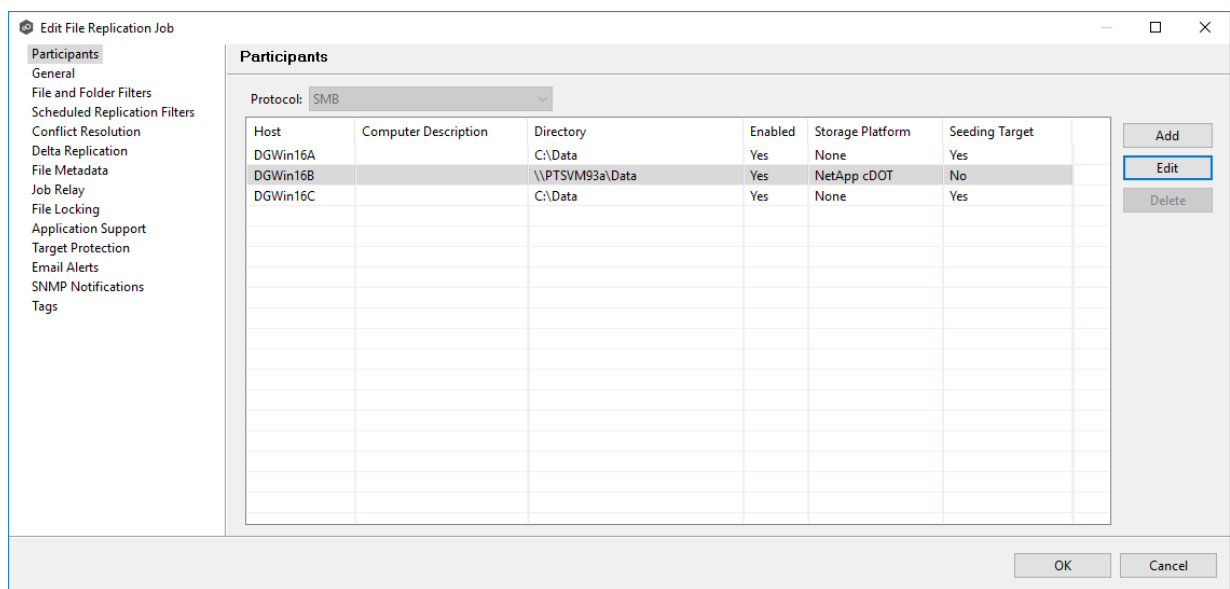
2. Click the **Delete** button.
3. Click **OK** in the **Delete Confirmation** dialog.

The participant is removed from the **Participants** table.

Note: A File Replication job must have at least two participants, so if after deleting a participant, there is only a single participant, you must add another participant to the job.

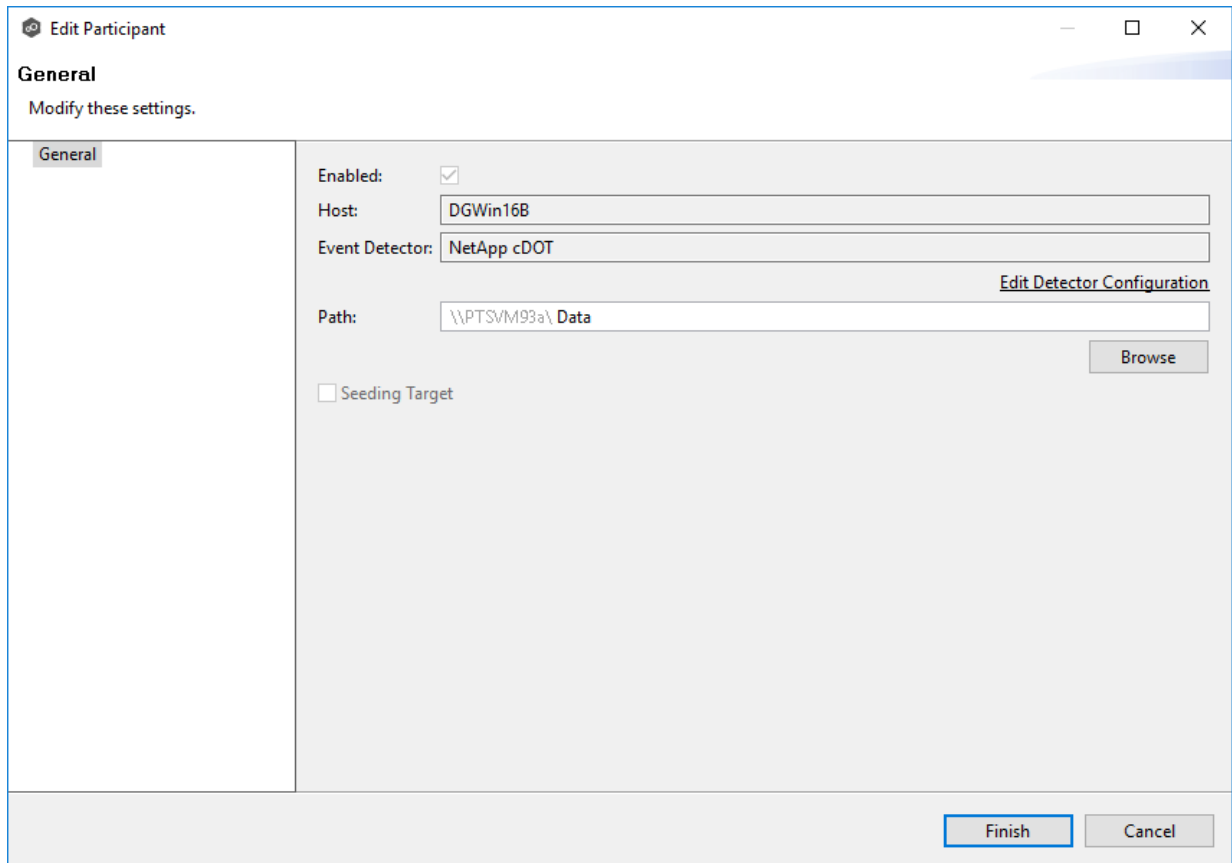
To edit a participant:

1. In the **Edit File Replication Job** dialog, select the participant in the **Participants** table you want to edit.



2. Click **Edit**.

The **Edit Participant** dialog appears.



The screenshot shows the 'Edit Participant' dialog box with the 'General' tab selected. The dialog has a title bar with a minimize, maximize, and close button. Below the title bar, the text 'General' is followed by 'Modify these settings.' The main area is divided into two panes. The left pane is labeled 'General'. The right pane contains the following fields and controls:

- Enabled:** A checked checkbox.
- Host:** A text box containing 'DGWin16B'.
- Event Detector:** A text box containing 'NetApp cDOT'. To the right of this field is a link labeled 'Edit Detector Configuration'.
- Path:** A text box containing '\\PTSVM93a\Data'. To the right of this field is a 'Browse' button.
- Seeding Target:** An unchecked checkbox.

At the bottom right of the dialog are two buttons: 'Finish' and 'Cancel'.

- To change the volume, share, export, or folder that is replicated, enter the path to the new watch set in the **Directory** field or browse to it.
 - If the destination storage device is a Windows file server, this path should be a local path such as D:\Data, or it can be the UNC path to any SMB-capable file server.
 - If the destination storage device is a Linux-based device, the path should be in NFS format. It should point to an NFS export or a subfolder under the export (e.g. server:/export-name/subfolder).
- If you are editing a Source Agent, you can edit the Detector Configuration. If it is a Destination Agent, the option is disabled.
- Click **OK** to close the Edit wizard or select another configuration item to modify.

General

The **General** page in the **Edit File Replication Job** dialog presents miscellaneous settings pertaining to a File Replication 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.

Edit File Replication Job

Participants

- General
- File and Folder Filters
- Scheduled Replication Filters
- Conflict Resolution
- Delta Replication
- File Metadata
- Job Relay
- File Locking
- Target Protection
- Email Alerts
- SNMP Notifications
- Tags

General

Job ID: 351

Job Type: File Replication

Job Name: FR-7

Job Description:

Transfer Block Size (KB): 2048

Verify Block Checksums:

Verify Full File Checksums:

Enable Multipart Transfers:

Synchronization Priority: 2

Timeout (Seconds): 180

First Scan Mode: FOLDER_BY_FOLDER

Remove Filtered Files On Folder Delete:

Delete Source File After Successful Transfer:

Require All Hosts At Start:

Auto Start:

Auto Start Priority: 1

OK Cancel

Option	Description
Job ID	Unique, system-generated job identifier that cannot be edited.
Job Type	Identifies the job type. This cannot be modified.
Job Name	The name of this File Replication 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 .
Verify Block Checksums	If selected, each block sent will be checksummed at both the source and target(s) Agents.

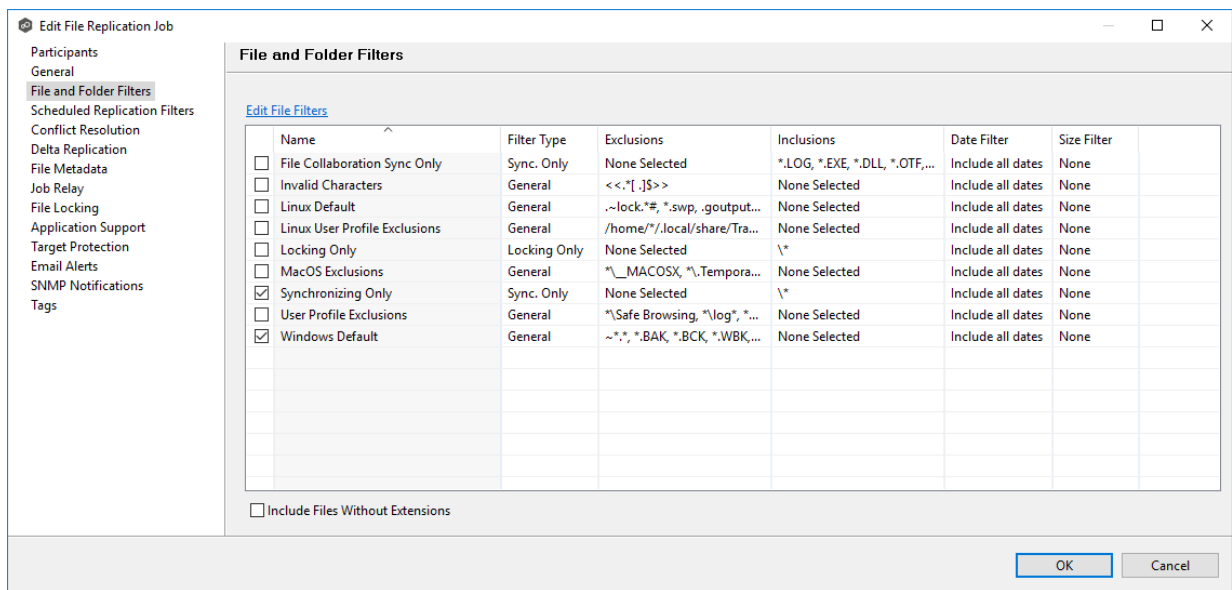
Option	Description
Verify Full File Checksums	If selected, the entire file will be checksummed after it has been sent from the source to target Agents. If temp files are enabled, the checksum on the target will be calculated prior to renaming the temp file to the base file's name. If the checksums between source and target(s) do not match, the file will be sent to the retry engine to reattempt the transfer.
Enable Multipart Transfers	If selected, larger files will be broken into smaller chunks and these chunks will be sent over the wire in a parallel fashion. This feature will automatically throttle itself if many other transfers are also waiting to be processed.
Synchronization 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 are NOT seeded, the FOLDER_BY_FOLDER method would be best. For environments where most data are seeded, the BULK_CHECKSUM method will result in a faster first scan.
Remove Filtered Files On Folder Delete	If selected, then all child files on target hosts will be deleted when its parent folder is deleted on another source host. Otherwise, filtered files will be left intact on targets when a parent folder is deleted on another source host.
Require All Hosts At Start	If selected, requires all participating hosts to be online and available at the start of the File Replication job in order for the job to successfully start.
Auto Start	If selected, then this file synchronization session will automatically be started when the Peer Management Center Service is started.

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

File and Folder Filters

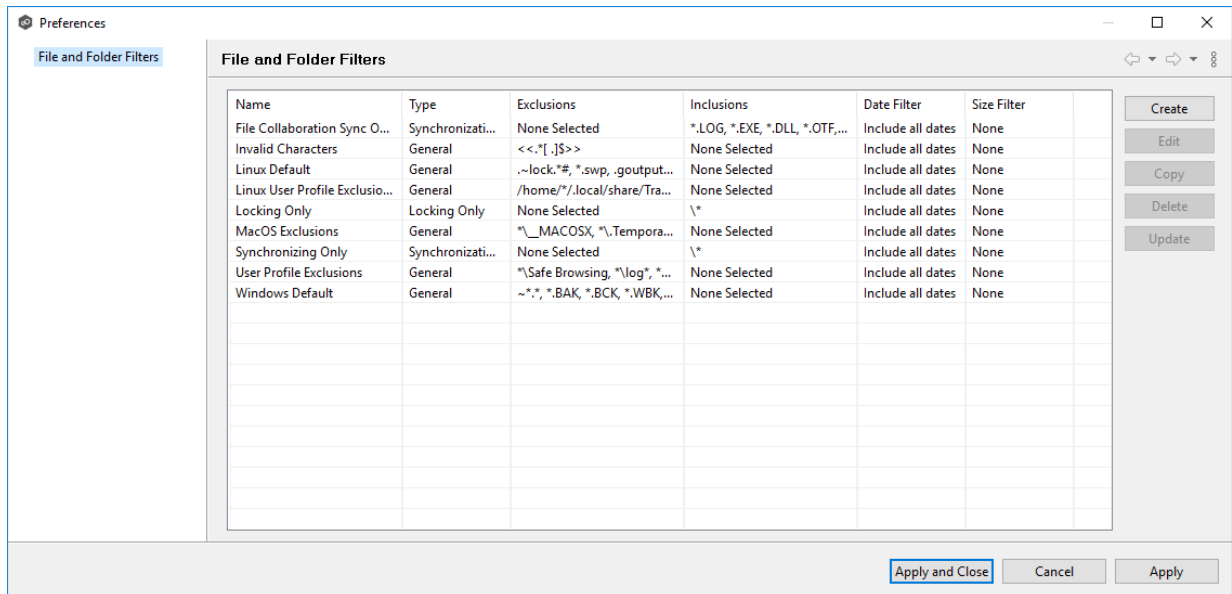
The **File and Folder Filters** page in the **Edit File Synchronization Job** dialog displays a list of [file and folder filters](#). A file or folder filter enables you to exclude and/or include files and folders from the job based on file type, extension, name, or directory path. Any file or folder that matches the filter is excluded or included from replication, depending on the filter's definition. By default, all files and folders selected in the **Source Paths** page will be replicated.

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



2. If you want to create a new file or folder filter, modify an existing one, or update a filter, click **Edit File Filters**.

The **File and Folder Filters** dialog appears. You cannot edit predefined filters. See [File Filters](#) in the [Preferences](#) section for information about creating or modifying a file filter.



3. Select the **Include Files Without Extensions** checkbox if you want to replicate files that do not have extensions.

Note: Files without extensions are ignored during replication unless you select this checkbox.

4. Click **OK** to close the Edit wizard or select another configuration item to modify.

Scheduled Replication Filters

The **Scheduled Replication** page in the **Edit File Synchronization Job** dialog displays a list of [scheduled replication filters](#). A scheduled replication filter enables you to identify files and folders that you don't want replicated in real-time.

1. Select the scheduled replication filters you want to apply to the job.

3. Click **OK** to close the Edit wizard or select another configuration item to modify.

Conflict Resolution

By default, any file conflicts that are encountered during the [initial synchronization process](#) are automatically resolved by Peer Management Center. Peer Management Center resolves the conflict by selecting the file with the most recent modification time. Conflicts that cannot be automatically resolved result in the files being quarantined. The **Conflict Resolution** page in the **Edit File Replication Job** allows you to select options for resolving file conflicts and quarantines.

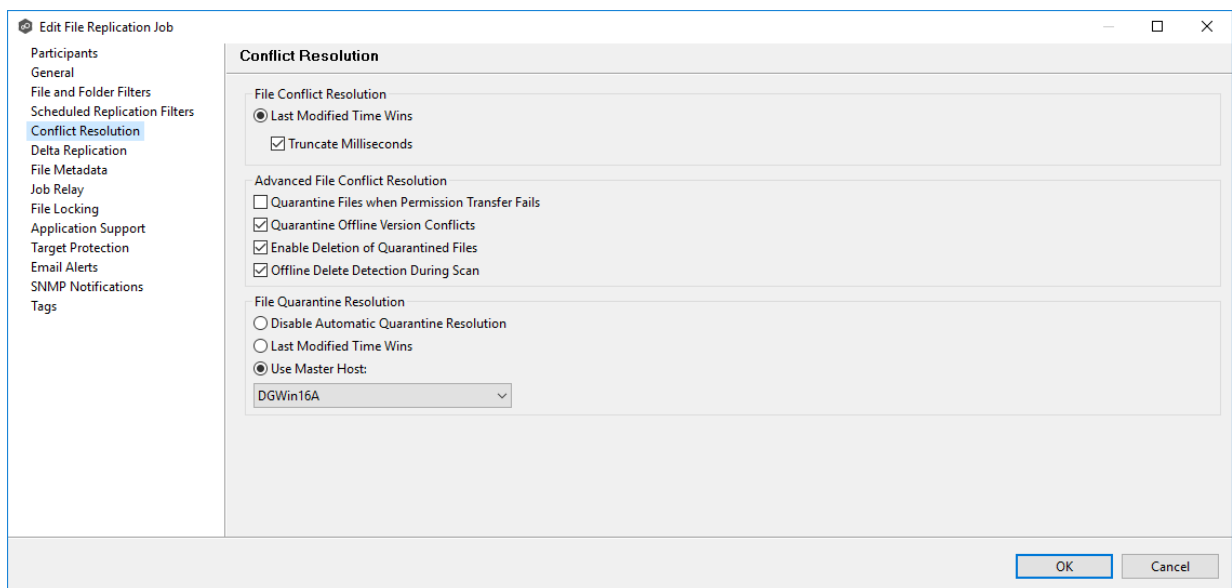
However, if you want to resolve the conflicts yourself, you can contact Peer Software to enable manual resolution. With manual resolution, you can select the host with the correct version of the file.

For more information about the cause of file conflicts, see [Conflicts, Retries, and Quarantines](#).

To modify conflict resolution settings for the File Synchronization job:

1. In the **File Conflict Resolution** section, select the **Truncate Milliseconds** option if you want the millisecond value truncated from each time stamp when comparing the time stamps of a file on two or more hosts.

As some storage platforms and applications track milliseconds slightly differently, selecting this option will prevent subtle millisecond differences from causing an otherwise in-sync file to be replicated or quarantined.



2. Select the **Advanced File Conflict Resolution** options you want applied:

Option	Description
Quarantine Offline Version Conflicts	Select this option if you want Peer Management Center to quarantine a file that was updated in two or more locations while the collaboration session was not running. If it is not selected, the file with the most recent last modified time will be replicated to all other participants.
Enable Deletion of Quarantined Files	Select this option if you want Peer Management Center to process a delete event for a quarantined file. If it is not selected, the quarantined file is not deleted and remains quarantined.
Offline Delete Detection During Scan	Select this option (and enabled target protection), if you want any files or folders that were deleted while the job was stopped to be deleted from all participants when the job is restarted. If it is not selected, then the deleted file or folder is restored from a participant with the file or folder to any participant where it no longer exists.

3. Select an option for automatically resolving quarantines (this option is intended to be used in environments where a single file server is active for a job):

Option	Description
Disable Automatic Resolution of Quarantines	Select this option if you want to manually resolve quarantines.
Last Modified Time	Select this option if you want quarantines automatically resolved by selecting the file with the latest modification time.
Use Master Host	Select this option if you want quarantines automatically resolved by selecting the file on the Master Host.

4. Click **OK** to close the Edit wizard or select another configuration item to modify.

Delta Replication

The **Delta Replication** page in the **Edit File Replication Job** dialog allows you to specify the delta-replication options to use for the selected File Replication job. Delta-level replication is a byte replication technology that enables block/byte level synchronization for a File Replication job. Through this feature, Peer Management Center is able to transmit only the bytes/blocks of a file that have changed instead of transferring the entire file. This results in much lower network bandwidth utilization, which can be an enormous benefit if you are transferring files across a slow WAN or VPN, as well as across a high-volume LAN.

Delta-level replication is enabled on a per File Replication job basis and generally affects all files in the [watch set](#). You will only benefit from delta-level replication for files that do not change much between file modifications, which includes most document editing programs.

To modify delta-level replication options:

1. Modify the following the fields as necessary.

Field	Description
Enable Delta-Level Replication	Select to enable delta encoded file transfers which only sends the file blocks that are different between source and target(s). If this is disabled, the standard file copy method will be used to synchronize files.

Field	Description
Checksum Transfer Size (KB)	Enter the block in kilobytes used to transfer checksums from target to source at one time. Larger sizes will result in faster checksum transfer, but will consume more memory on the Peer Agents
Delta Block Transfer Size (KB)	Enter the block size in kilobytes used to transfer delta encoded data from source to target at one time. Larger sizes will result in faster overall file transfers but will consume more memory on the Peer Agents.
Minimum File Size (KB)	Enter the minimum size of files in kilobytes to perform delta encoding for. If a file is less than this size, then delta encoding will not be performed.
Minimum File Size Percentage Target/Source	Enter the minimum allowed file size difference between source and target, as a percentage, to perform delta encoding. If the target file size is less than this percentage of the source file size, then delta encoding will not be performed.
Excluded File Extensions	Enter a comma-separated list of file extension patterns to be excluded from delta encoding, e.g., zip, jpg, png. In general, compressed files should be excluded from delta encoding and the most popular compressed file formats are excluded by default.
Excluded File Name Patterns	Enter a list of file name patterns to be excluded from delta encoding. If a file name matches any pattern in this list, then it will be excluded from delta encoding transfers and a regular file transfer will be performed. See File and Folder Filters for more information on specifying wildcard expressions.

- Click **OK** to close the Edit wizard or select another configuration item to modify.

File Metadata

The **File Metadata** page allows you to specify whether you wish to replicate file security metadata and the types of metadata for synchronization. Additionally, it provides the ability to

designate the metadata source (volume, share, export, or folder) to resolve conflicts during the initial synchronization. This designated source, utilized in case of conflicts, is referred to as the [master host](#). This page also provides some additional options not available when creating the job. See [File Metadata Synchronization](#) in [Advanced Topics](#) for more information about file metadata synchronization.

The contents of the File Metadata page vary depending on whether you are using Windows-based or Linux-based Agents for your job:

- If you selected Windows-based, proceed with [SMB File Metadata](#).
- If you selected Linux-based, proceed with [NFS File Metadata](#).
- If you selected multi-protocol (Windows and Linux), proceed with [Multi-Protocol Metadata](#).

To modify file metadata synchronization settings:

1. **Select when to replicate metadata** (one or both options):
 - **Real-time synchronization:** Metadata changes (Owner, DACL, SACL) will be synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata will be synchronized during the initial scan.

The screenshot shows the 'Edit File Replication Job' dialog box with the 'File Metadata' tab selected. The left sidebar lists various configuration categories, with 'File Metadata' highlighted. The main area contains the following settings:

- Synchronize File Security Information:**
 - Enable synchronizing file security information in real-time
 - Enable synchronizing file security information with master host during initial scan
 - Enable prevention of corrupt or blank Owner or ACLs on source or master host from being applied to any target host
- Synchronize Security Options:**
 - Owner
 - DACL: Discretionary Access Control List
 - SACL: System Access Control List
- Metadata Conflict Resolution:**
 - Select a master host for initial scan:
 - Enable enhanced metadata conflict resolution
- 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)
- File Timestamps Transfer:**
 - Use source create time
 - Use source last accessed time

At the bottom right, there are 'OK' and 'Cancel' buttons.

2. Confirm metadata synchronization.

If you selected either option, click **OK** in the confirmation message.

3. Prevent corrupted or blank metadata from propagating (optional):

- Select **Enable prevention of corrupt or blank Owner or ACLs** to prevent invalid ownership or ACL entries on the source or master host from being applied to any target host. The replication process will skip or correct these entries to maintain data integrity and security.

4. Select security descriptor components to be synchronized:

- Choose which components (Owner, DACL, SACL) to synchronize.

Note: To synchronize Owner or SACLs, the user that a Peer Agent service is run under on each participating host must have permission to read and write Owner and SACLs.

5. Select the master host (if applicable):

- If synchronizing metadata during the initial scan, choose a master host to resolve conflicts.
- If no master host is selected, metadata synchronization will not occur during the initial scan.
- If security descriptor components differ across participants, conflict resolution will use the master host's permissions.
- If the file does not exist on the master host, a participant will be chosen randomly.

6. Select **Enable enhanced metadata conflict resolution** if you want to prevent the Peer Agent service account from being assigned as the owner of that file or folder when a metadata conflict occurs, and a file or folder is written to a target. If the Peer Agent service account were to be assigned as the owner, the user may not have permission to access the file or folder.

Note: The Peer Agent service account cannot be a local or system administrator. As described in [Peer Global File Service - Environmental Requirements](#), the Peer agent service account should be an actual user.

7. (Optional) Enter values for one or both file reparse point data synchronization options:

- **Reparse Tag Name** - Enter a single numerical value. Must be either blank (if blank, reparse synchronization will be disabled) or greater than or equal to 0. The default for Symantec Enterprise Vault is 16. A value of 0 enables reparse point synchronization for all reparse file types. If you are unsure as to what value to use, then contact Peer Software Support, or you can use a value of 0 if you are sure that you are only utilizing one vendor's reparse point functionality.
- **Reparse Master Host** - Select a master host. If a master host is selected, then when the last modified times and file sizes match on all hosts, but the file reparse attribute differs (e.g., archived/offline versus unarchived on file server), then the file reparse data will be synchronized to match the file located on the master host. For

Enterprise Vault, this should be the server where you run the archiving task on. If the value is left blank, then no reparse data synchronization will be performed, and the files will be left in their current state.

Note: Use this option only if you are utilizing archiving or hierarchical storage solutions that make use of NTFS file reparse points to access data in a remote location, such as Symantec's Enterprise Vault. Enabling this option allows synchronization of a file's reparse data, and not the actual offline content, to target hosts, and prevents the offline file from being recalled from the remote storage device.

8. (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.

Note: The **Alternate Data Stream Transfer** option is not available if the job is using Edge Caching.

9. In the **File Timestamps Transfer** section, select one or both options:

- **Use source create time:** Enable this option if you want to preserve the create time (or creation timestamp) of the source file during the transfer process. Enabling this option ensures that the destination file retains the same create time as the source file after the transfer is complete.
- **Use source last accessed time:** Enable this option if you want to preserve the last accessed time (or access timestamp) of the source file during the transfer process. Enabling this option ensures that the destination file inherits the same last accessed time as the source file after the transfer is complete.

These timestamps are only applied during transfer operations. If the timestamps on the source change but the contents of the file remain the same, they will not be applied to the destination.

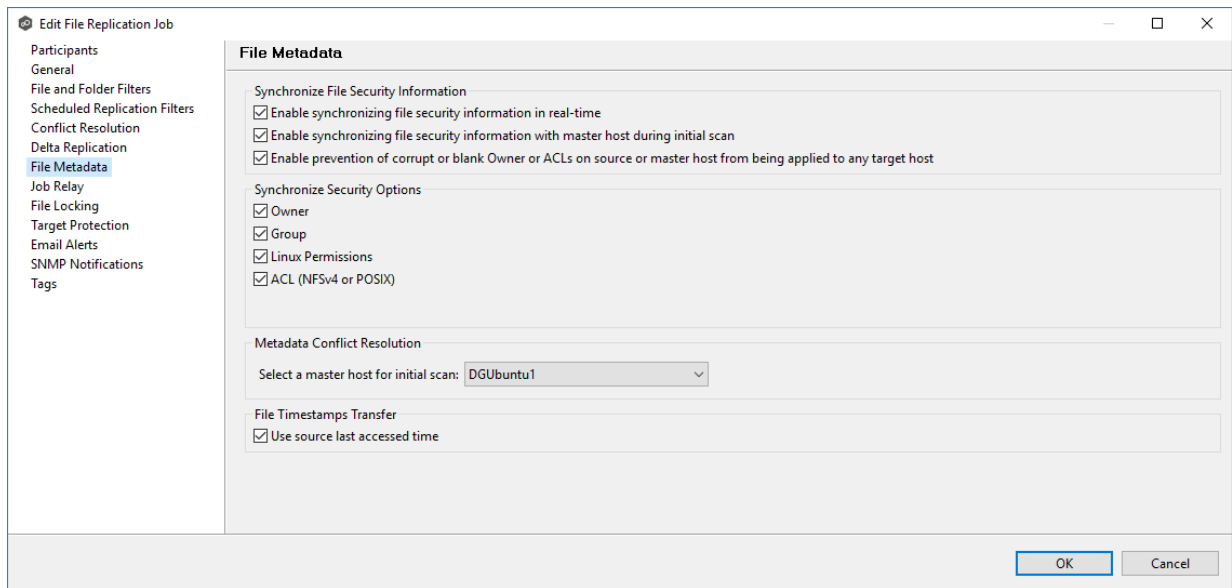
10. Click **OK** to close the Edit wizard or select another configuration item to modify.

To modify file metadata synchronization settings:

1. **Select when to replicate metadata** (first or second options):

- **Real-time synchronization:** Metadata changes (Owner, Group, Linux Permissions, ACL) are synchronized with the target host as they occur.
- **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.



2. Confirm metadata synchronization.

If you selected either of the first two options, click **OK** in the confirmation message.

3. Prevent corrupted or blank metadata from propagating (optional):

- Select **Enable prevention of corrupt or blank Owner or ACLs** to prevent invalid ownership or ACL entries on the source or master host from being applied to any target host. The replication process will skip or correct these entries to maintain data integrity and security.

4. Select access controls to be synchronized:

- Choose which components (Owner, Group, Linux Permissions, ACL) to synchronize.

Note: To synchronize Owner or ACLs, the user that a [Peer Agent](#) service is run under on each participating host must have permission to read and write Owner and ACLs.

5. Select the master host (if applicable):

- If synchronizing metadata during the initial scan, choose a [master host](#) to resolve conflicts.

This option is only available when both of the first two options in the **Synchronize File Security Information** section are enabled, and **Owner** is selected under **Synchronize Security and ACL Options**.

- If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If access controls differ across participants, [conflict resolution](#) will use the master host's permissions.
 - If the file does not exist on the master host, a participant will be chosen randomly.
6. Select **Enable enhanced metadata conflict resolution** if you want to prevent the Peer Agent service account from being assigned as the owner of that file or folder when a metadata conflict occurs, and a file or folder is written to a target. If the Peer Agent service account were to be assigned as the owner, the user may not have permission to access the file or folder.

Note: The Peer Agent service account cannot be a local or system administrator. As described in [Peer Global File Service - Environmental Requirements](#), the Peer agent service account should be an actual user.

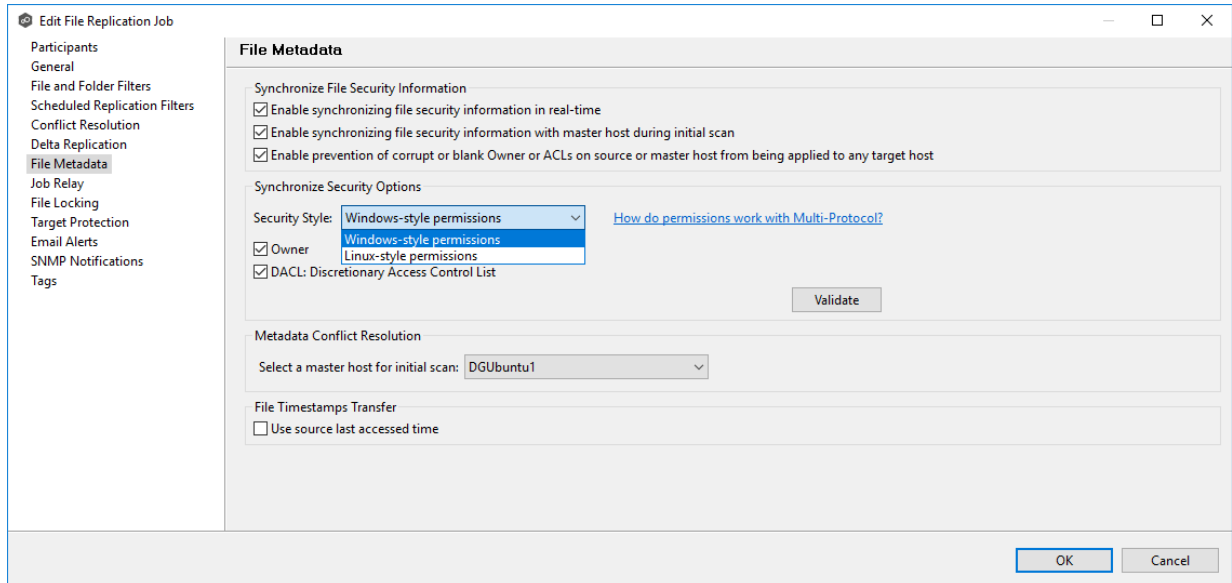
7. In the **File Timestamps Transfer** section, select **Use source last accessed time** if you want to preserve the last accessed time of the source file during the transfer process. This option ensures that the destination file inherits the same last accessed time as the source file after the transfer is complete.
8. Click **OK** to close the Edit wizard or select another configuration item to modify.

To modify file metadata synchronization settings:

1. **Select when to replicate metadata** (one or both options):
 - **Real-time synchronization:** Metadata changes are synchronized with the target host as they occur.

- **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.



2. Confirm metadata synchronization.

If you selected either one of the first two options, click **OK** in the confirmation message.

3. Prevent corrupted or blank metadata from propagating (optional):

- Select **Enable prevention of corrupt or blank Owner or ACLs** to prevent invalid ownership or ACL entries on the source or master host from being applied to any target host. The replication process will skip or correct these entries to maintain data integrity and security.

4. Select the security style:

- Windows-style permissions
- Linux-style permissions
- Mixed permissions (available only with Dell PowerScale)

5. Select permissions to be synchronized:

For a general discussion of synchronizing permissions in a multi-protocol job, see the Peer Software knowledge base article [How Do Permissions Work with Multi-Protocol?](#)

6. Select the master host (if applicable):

- If synchronizing metadata during the initial scan, choose a [master host](#) to resolve conflicts.
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If access controls differ across participants, [conflict resolution](#) will use the master host's permissions.
 - If the file does not exist on the master host, a participant will be chosen randomly.
7. Click **OK** to close the Edit wizard or select another configuration item to modify.

Job Relay

The **Job Relay** page allows you to set up a job to replicate data from the target of one job to other destinations. This is called a **relay job**. A relay job is useful when you want to replicate the same set (or subset) of data to multiple destinations but want to minimize data traffic from the initial source. Normally, you cannot have jobs that have overlapping watch sets running at the same time; however, turning on job relay allows that.

For example, if you want to replicate data from Source A to Target B, and then replicate that data from Source B to Targets C, D, and E, you could do this by setting up two File Replication jobs:

- **Replication Job 1** replicates from Source A to Target B.
- **Replication Job 2** replicates from Source B to Targets C, D, and E.

In this example, **Replication Job 2** is the **relay job**. It replicates (or relays) the data it received from **Replication Job 1** to the other destination targets (C, D, and E).

Setting Up a Job Relay

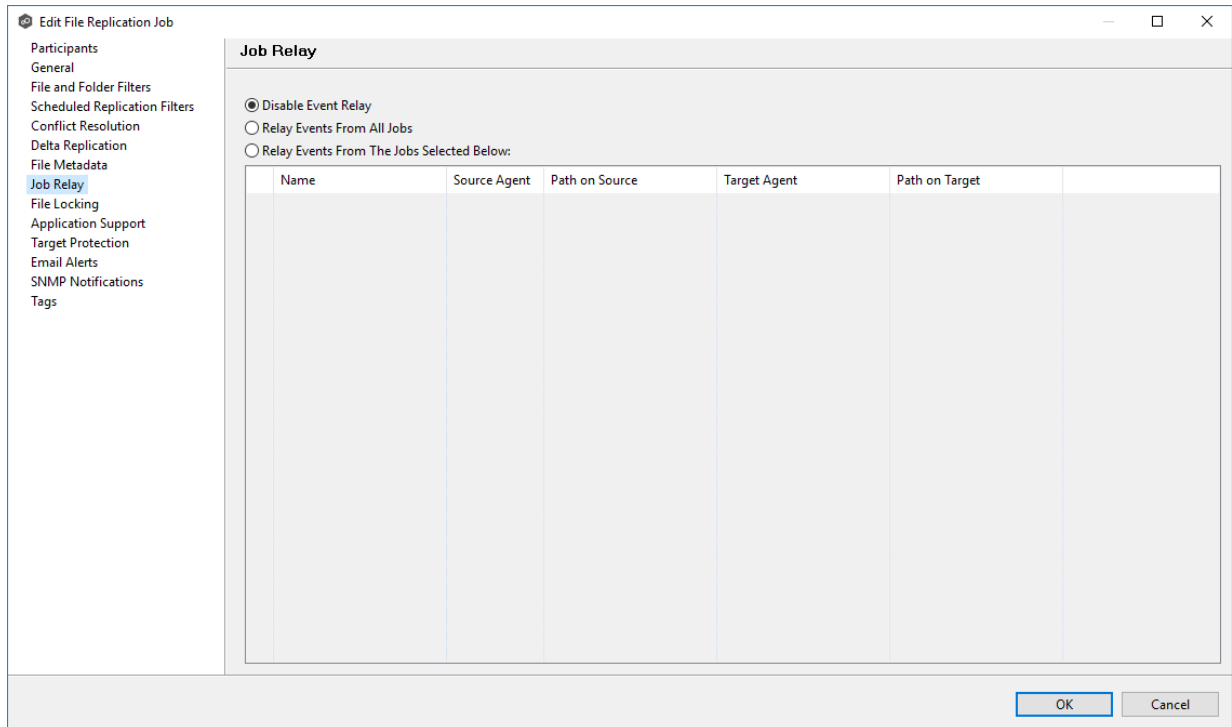
By default, job relay is disabled. To set up a job relay situation, you must:

1. Create the initial job that replicates data from Source A to Target B.
2. Create the second job, which will be the **relay job**. When creating a replication job, you can add only one target participants during creation; you can add additional target participants when you edit the job.

- (Optional) Edit the second job to add additional target participants.

Now you are ready to set up the relay job.

- Continue editing the second job to set up the job relay. (If you clicked Finish in Step 3, select the second job again and edit it.)
- Click **Job Relay** in the navigation tree.



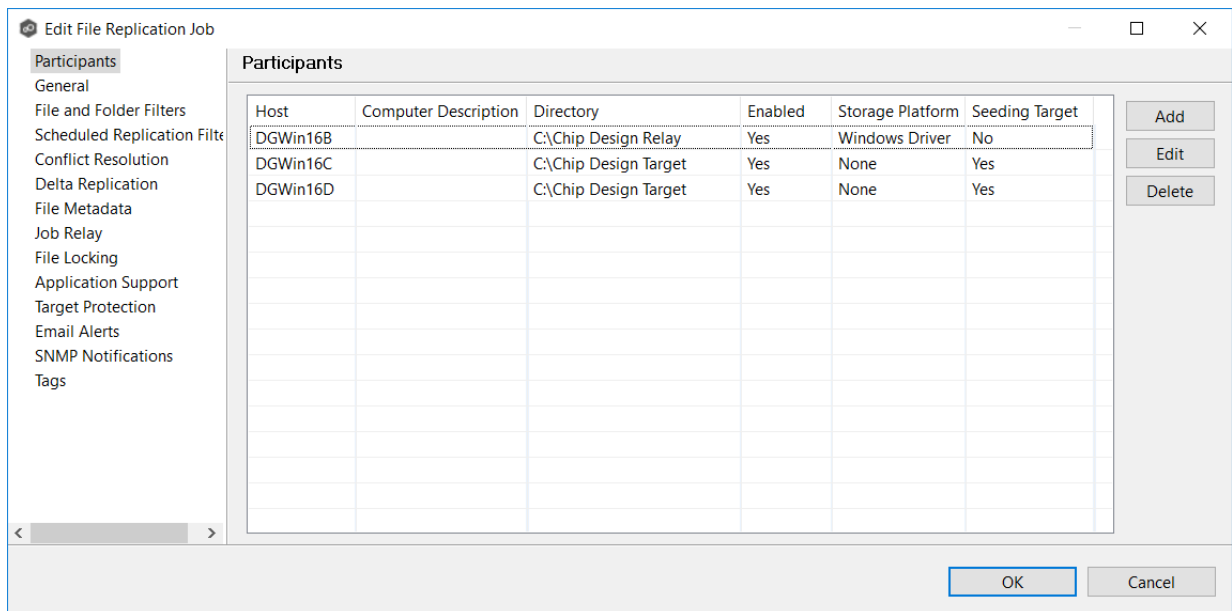
- Select one of the relay options:

- **Relay Events from All Jobs:** Select this option if you want to relay data from all jobs that point to the source path of this relay job. These jobs can also point to a child folder under the source path of this relay job.
- **Relay Events From the Jobs Selected Below:** Select this option if you want to explicitly choose which jobs can be relayed. See **Relaying Events from Selected Jobs** for more information.

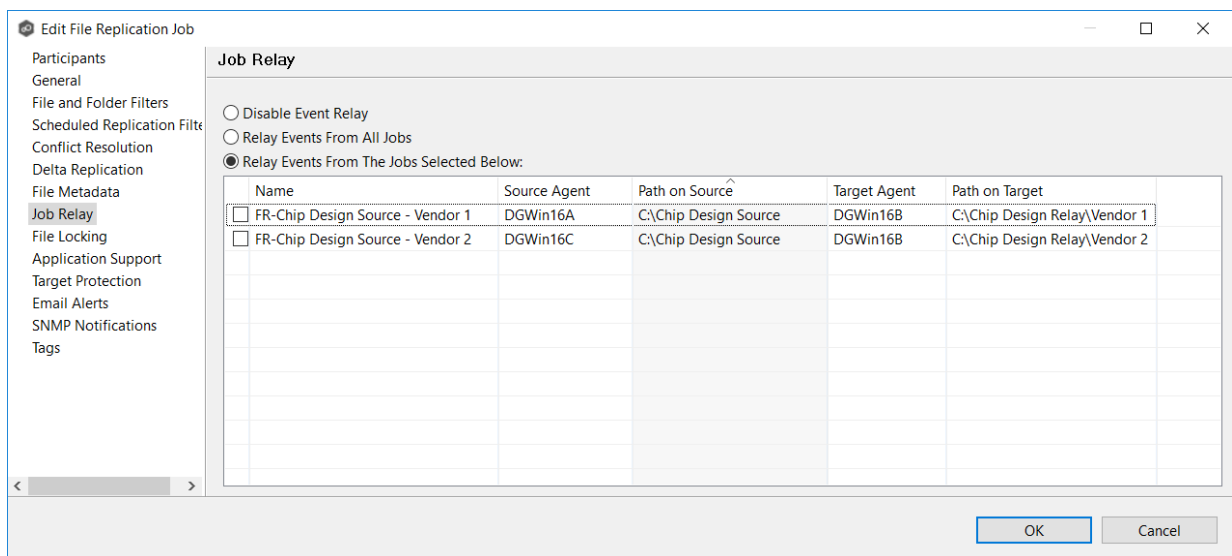
- Click **OK**.

Relaying Events from Selected Jobs

In the following example, the relay job has two target participants: DGWin16C and DGWin16D.



In the Job Relay page, you select which jobs are replicated to the target participants:

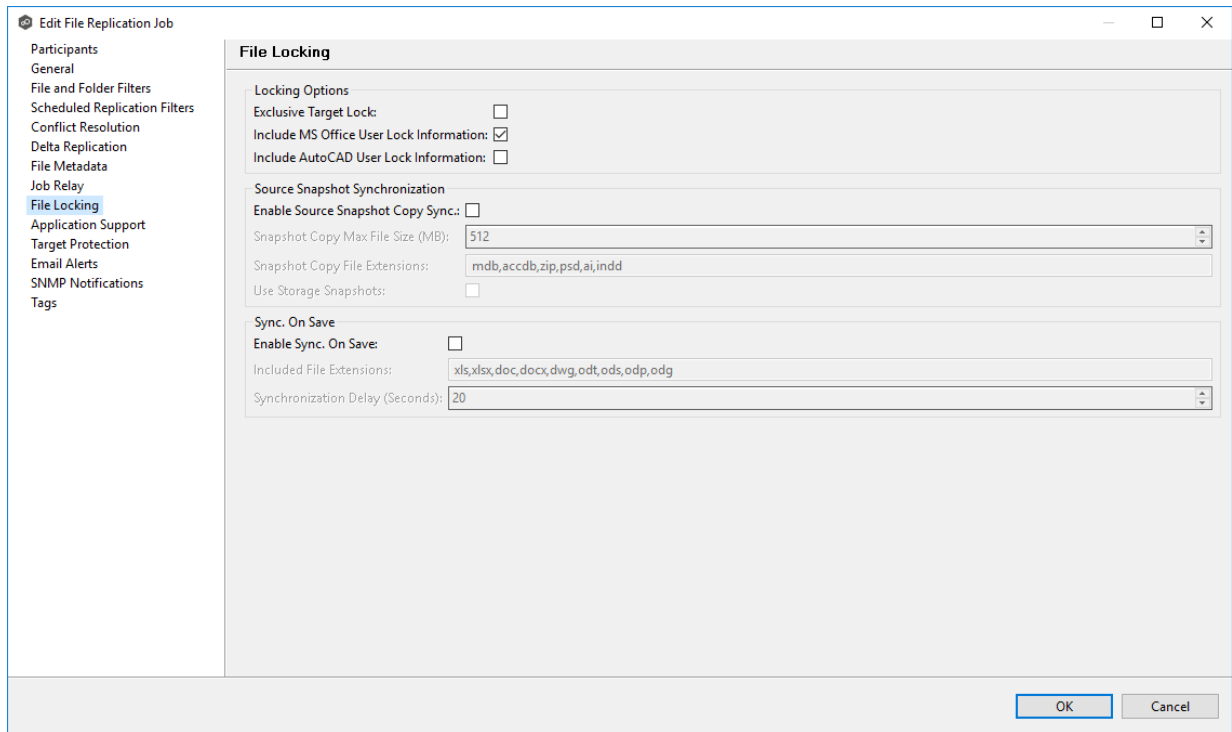


File Locking

The **File Locking** page in the **Edit File Replication Job** dialog presents options pertaining to how source and target files are locked by Peer Management Center.

To modify file locking options:

1. Modify these fields as needed:



- Click **OK** to close the Edit wizard or select another configuration item to modify.

Locking Options

Option	Description
Exclusive Target Lock	If enabled, then whenever possible, an exclusive lock will be obtained on target file handles, which will prevent users from opening the file (even in read-only mode) while a user has the file opened on the source host. When this option is disabled, then users will be allowed to open files for read-only if the application allows for this.
Include MS Office User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported Microsoft Office files (e.g., Word, Excel, and PowerPoint).
Include AutoCAD User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported AutoCAD files.

Source Snapshot Synchronization Option

Option	Description
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.

Sync On Save Options

Option	Description
Exclusive Target Lock	If enabled, then whenever possible, an exclusive lock will be obtained on target file handles, which will prevent users from opening the file (even in read-only mode) while a user has the file opened on the source host. When this option is disabled, then users will be allowed to open files for read-only if the application allows for this.
Include MS Office User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported Microsoft Office files (e.g., Word, Excel, and PowerPoint).
Include AutoCAD User Lock Information	If enabled, user lock information (if available) will be propagated to target locks for supported AutoCAD files.
Enable Source Snapshot Copy Sync.	If enabled, a snapshot copy of the source file will be created for files that meet the snapshot configuration criteria below, and this copy will be used for synchronization purposes. In addition, no file handle will be held on the source file except while making a copy of the file.
Snapshot Copy Max File Size (MB)	The maximum file size for which source snapshot synchronization will be utilized.
Snapshot Copy File Extensions	A comma-separated list of file extensions for which source snapshot synchronization will be utilized.
Enable Sync. On Save	If enabled, this feature will allow supported file types to be synchronized after a user saves a file, rather than waiting for the file to close.
Included File Extensions	A comma-separated list of file extensions for which to enable the Sync. On Save feature.
Synchronization Delay (Seconds)	The number of seconds to wait after a file has been saved before initiating a synchronization of the file.

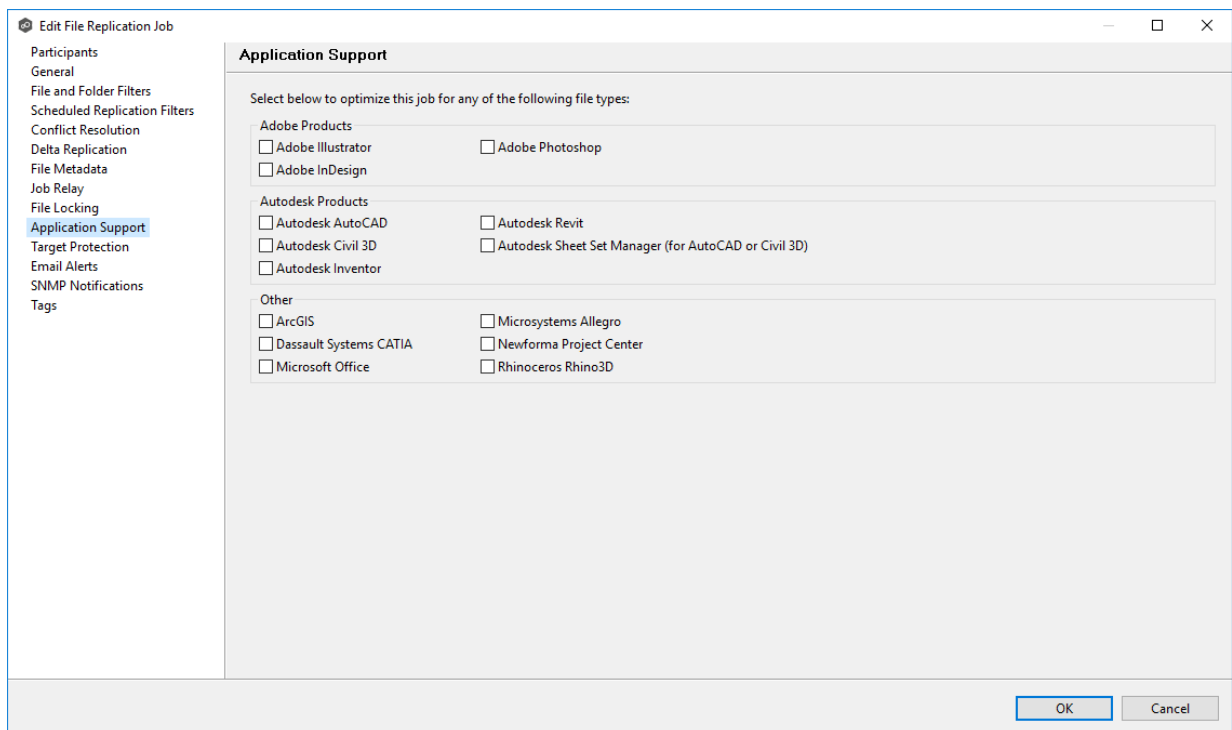
Application Support

Application Support enables automatic optimization of a file replication job for files created by certain 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 [Peer Support](#).

To modify which applications are optimized:

1. Select the applications to be optimized.



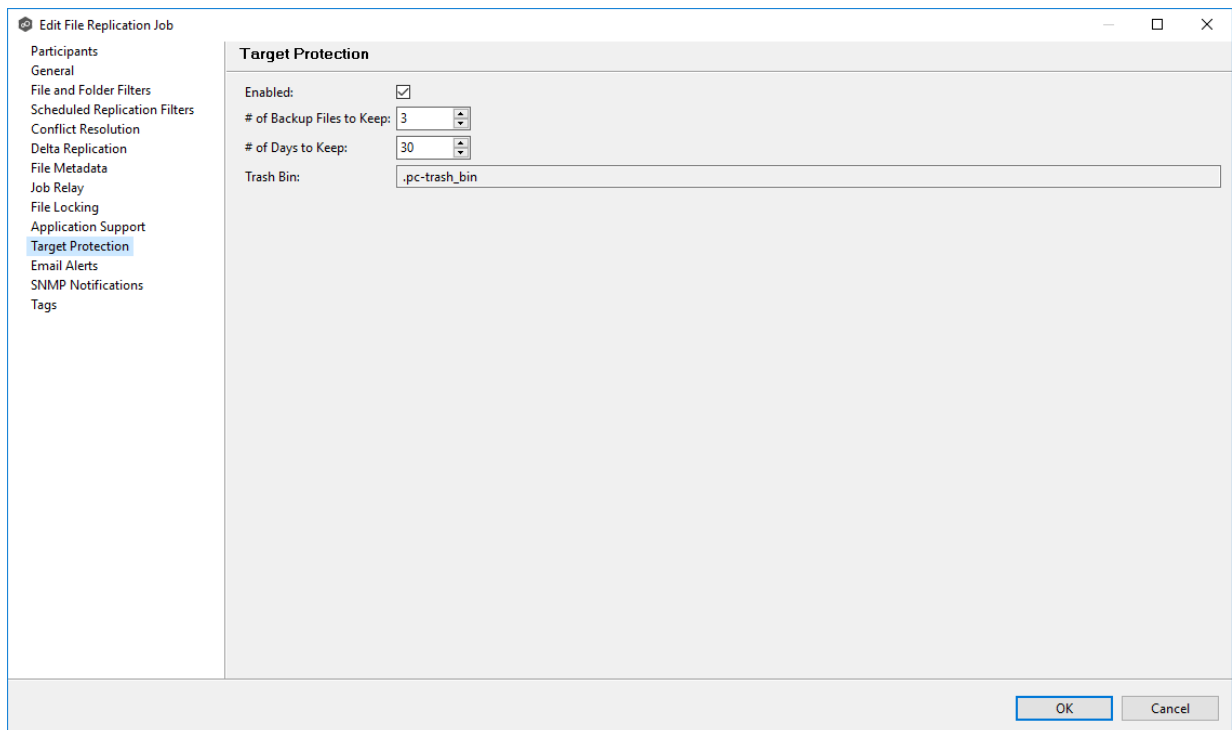
2. Click **OK** to close the Edit wizard or select another configuration item to modify.

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 a job has target protection enabled, then whenever a file is deleted or modified on the source host but before the changes are propagated to the targets, a copy of the existing file on the target is moved to the Peer Management Center trash bin.

The trash bin is located in a hidden folder named **.pc-trash_bin** found in the root directory of the [watch set](#) of the target host. A backup file is placed in the same directory hierarchy location as the source folder in the watch set within the recycle bin folder. If you need to restore a previous version of a file, you can copy the file from the trash bin into the corresponding location in the watch set and the changes will be propagated to all other collaboration hosts.

Target protection configuration is available by selecting **Target Protection** from the tree node within the Edit File Synchronization Configuration dialog.



Modify the fields as needed:

Field	Description
Enabled	Enables target protection.
# of Backup	The maximum number of backup copies of an individual file to keep in the trash bin before purging the oldest copy.

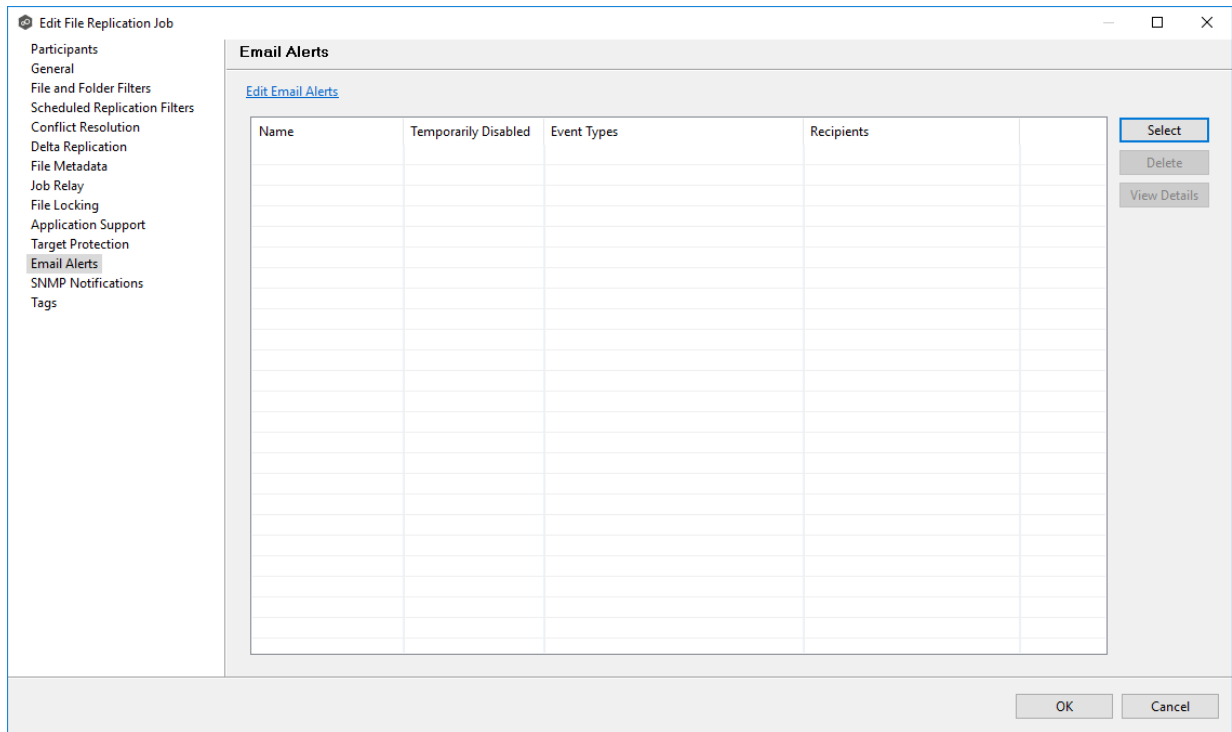
Field	Description
Files to Keep	
# of Days to Keep	The number of days to keep a backup archive copy around before deleting from disk. A value of 0 will disable purging any files from archive.
Trash Bin	The trash bin folder name located in the root directory of the watch set. This is a hidden folder and the name cannot be changed by the end user.

Email Alerts

The **Email Alerts** page in the **Edit File Replication Job** dialog allows you to select which email alerts to apply to a File Replication 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 Replication job.

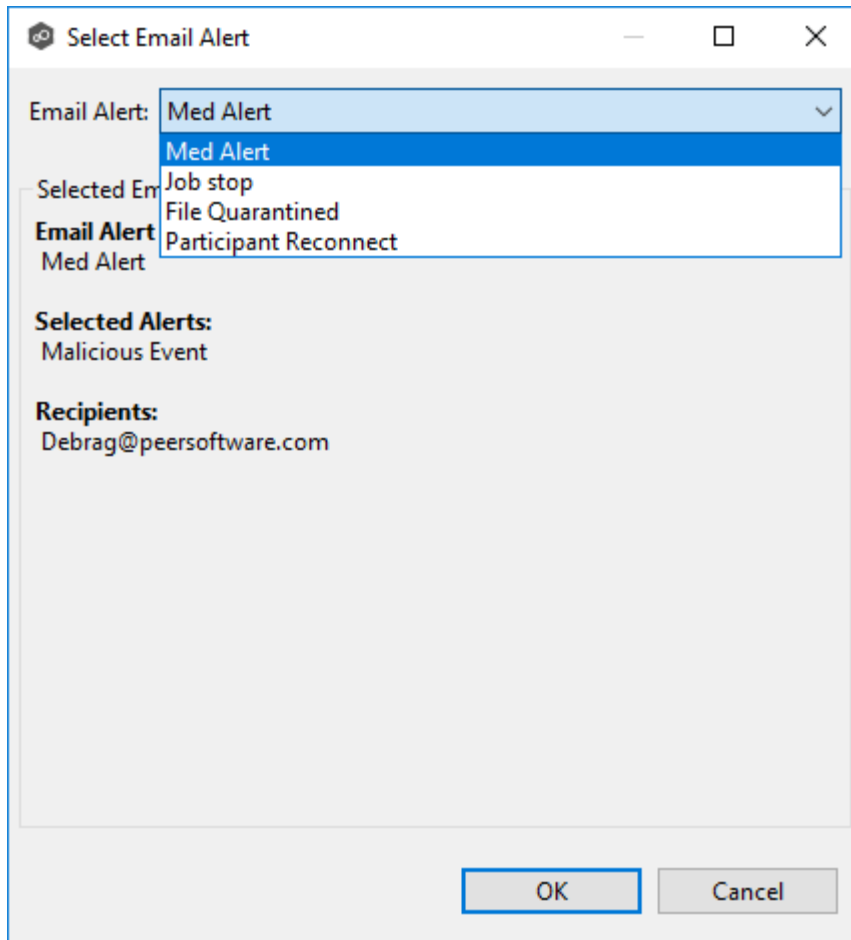
To apply email alerts to a File Replication job while editing the job:

1. Click the **Select** button.

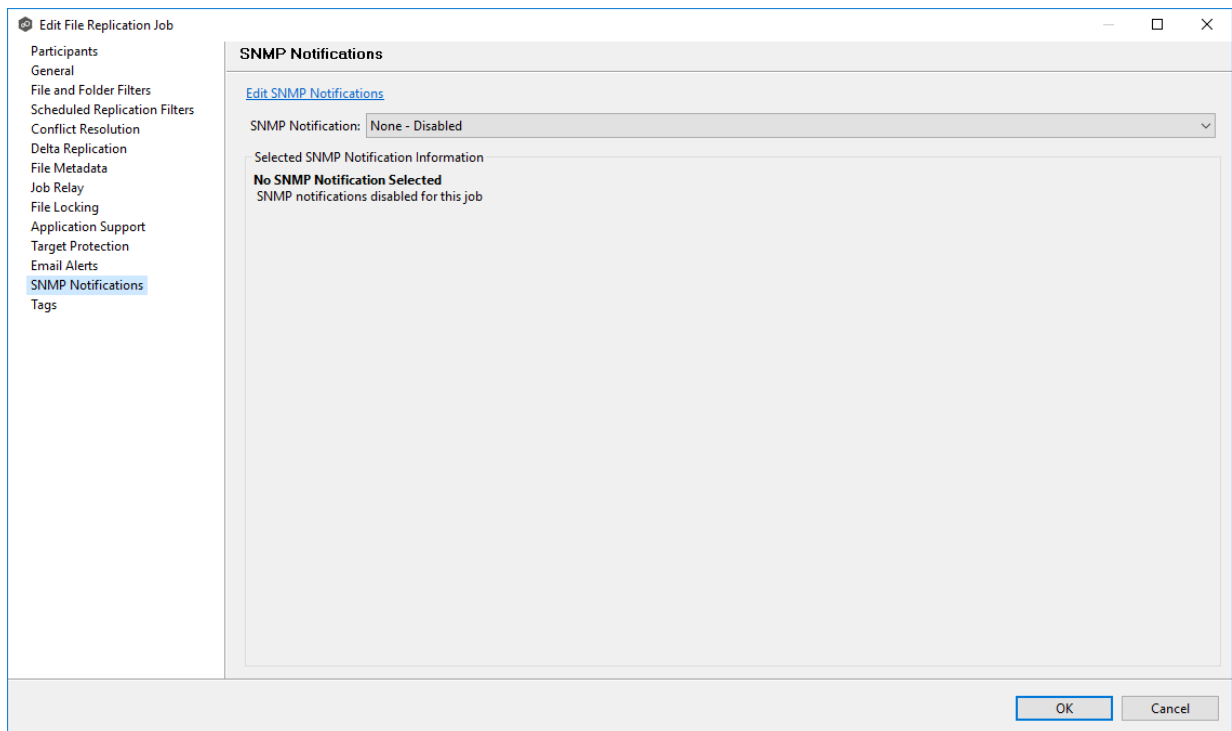


The **Select Email Alert** dialog opens.

2. Select the email alert from the drop-down list, and then click **OK**.



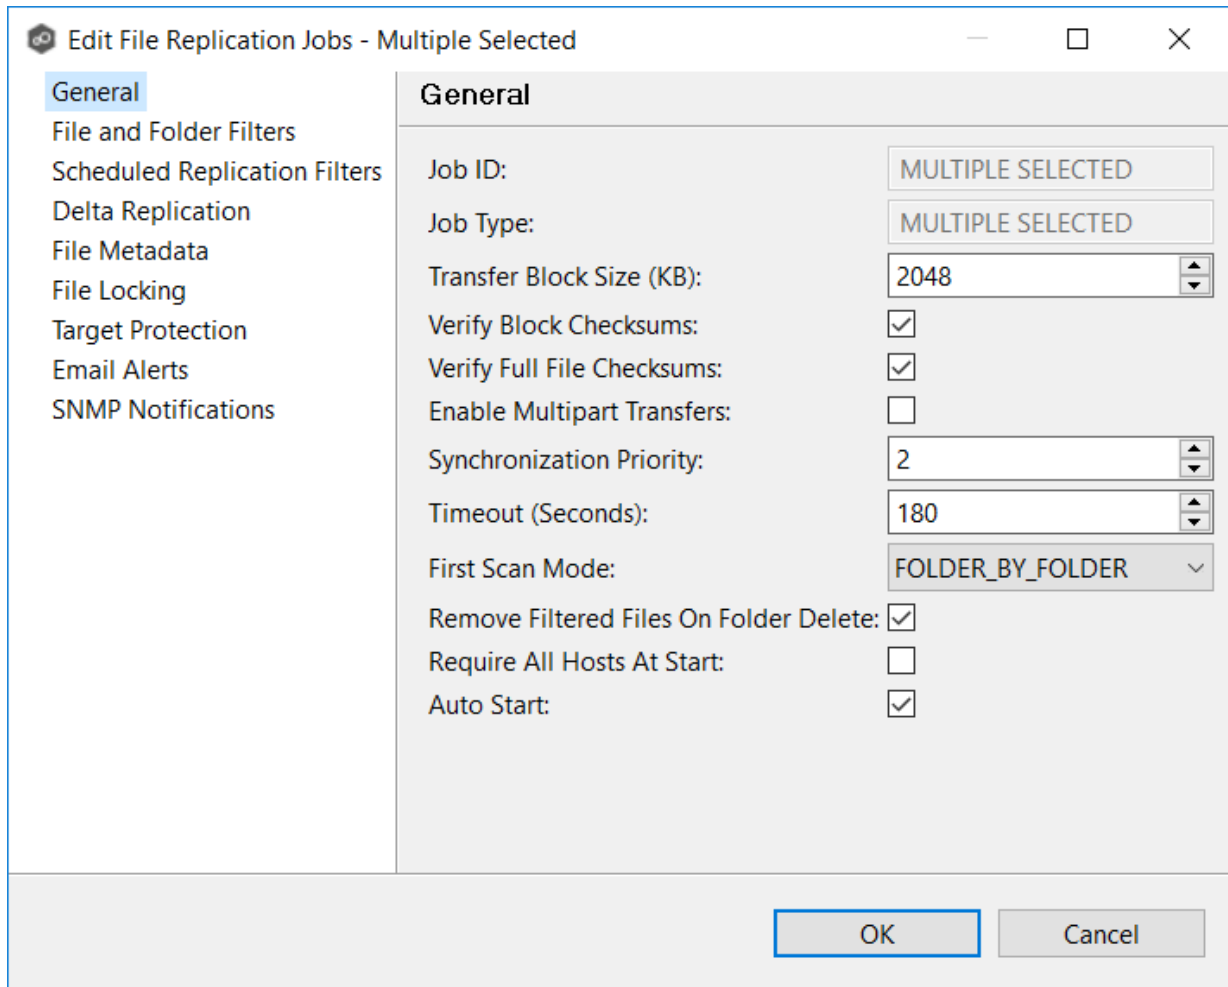
The newly added email alert appears in the **Email Alerts** table.



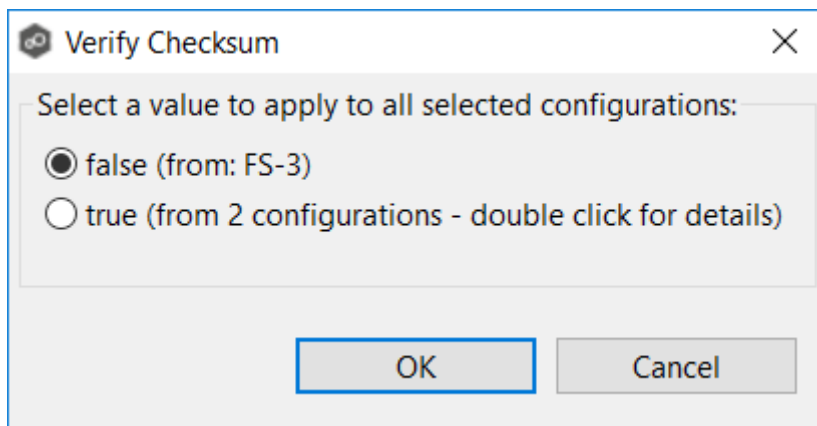
2. Click **OK** to close the Edit wizard or select another configuration item to modify.

Tags

The **Tags** page in the **Edit File Replication Job** dialog allows you to assign existing tags and categories to the selected job. This page is not available in [Multi-Job Editing](#) mode. For more information about tags, see [Tags](#) in the [Basic Concepts](#) section.



In this dialog, any discrepancies between multiple selected jobs will generally be illustrated by a read-only text field with the caption **Multiple Values - Click to Edit**. Clicking this field will bring up a dialog similar to the following:



This dialog gives you the option of choosing a value that is already used by one or more selected File Synchronization jobs, in addition to the ability to use your own value.

Notice that variances in the look and feel of this pop-up dialog above depend on the type of information it is trying to represent (for example, text vs. a checkbox vs. a list of items).

Upon clicking OK, the read-only text field you originally clicked will be updated to reflect the new value. Any fields that have changed will be marked by a small caution sign. On saving in this multi-job edit dialog, the changed values will be applied to all selected jobs.

Note: Read all information on each configuration page carefully when using the multi-job edit dialog. A few pages operate in a slightly different manner than mentioned above. All of the necessary information is provided at the top of these pages in bold text.

File Synchronization Jobs

A File Synchronization job requires at least two participants. Each participant consists of an Agent and the volume, share, export, or folder to be synchronized. A File Synchronization job synchronizes the files of participants in real-time.

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 enables real-time, multi-directional synchronization across various storage platforms and locations. It is designed for non-collaborative workloads where files must remain in sync in real time across multiple locations without persistent locking. While File Synchronization jobs do temporarily lock files during synchronization to ensure data integrity when reading and writing, the lock is released once synchronization is complete. This job type is specifically optimized for user home directories and profiles.

A File Synchronization job requires at least two participants. Each participant consists of an Agent and the volume, share, export, or folder to be synchronized. A File Synchronization job synchronizes the files of participants in real-time.

Before You Create Your First File Synchronization Job

We strongly recommend that you configure global settings such as SMTP configuration and email alerts before configuring your first File Synchronization job. See [Preferences](#) for details on what and how to configure these settings.

Creating a File Synchronization Job

The **Create Job Wizard** walks you through the process of creating a File Synchronization job:

[Step 1: Job Type and Name](#)

[Step 2: Participants](#)

[Step 3: File Metadata](#)

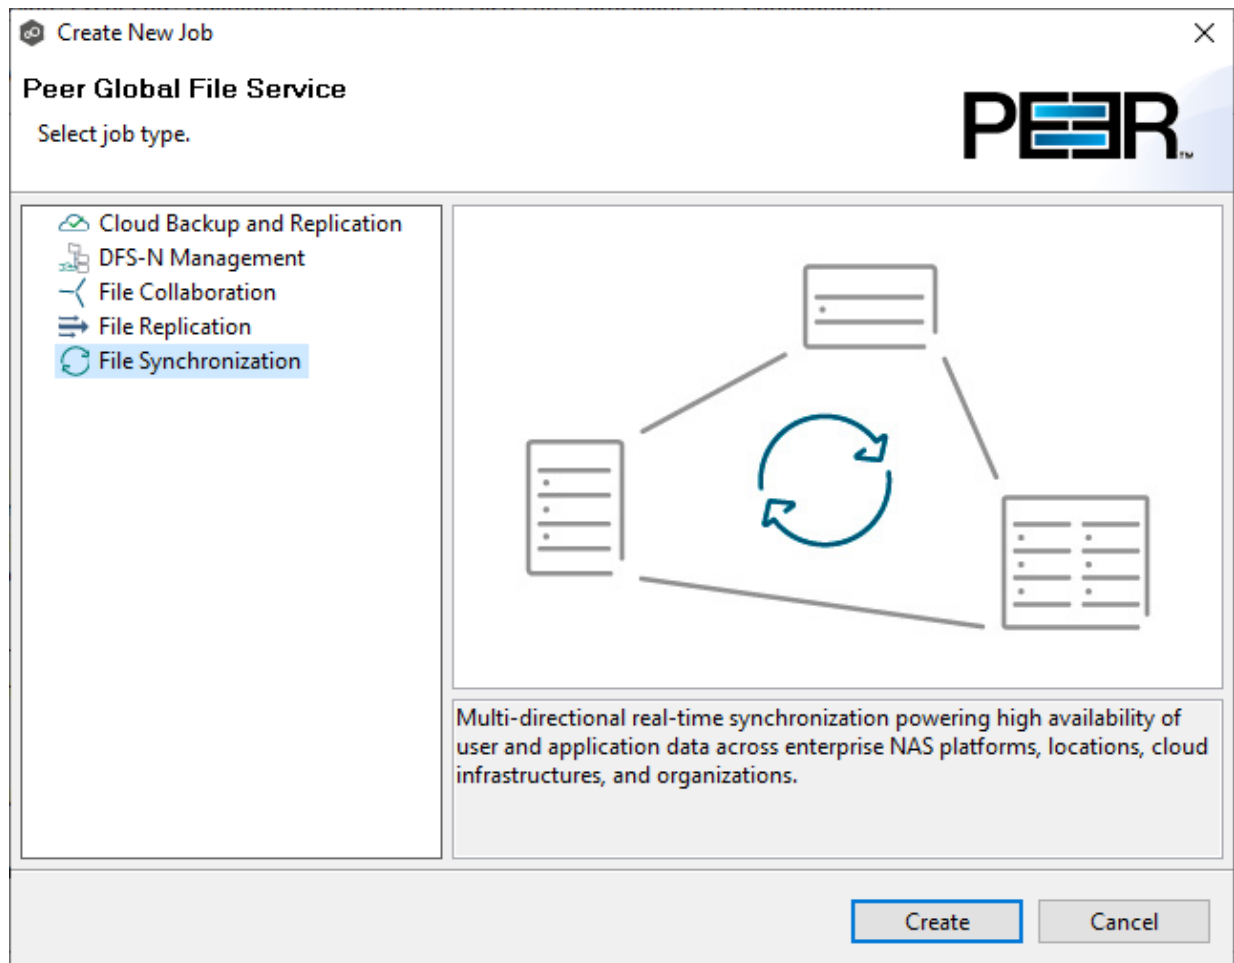
[Step 4: Email Alerts](#)

[Step 5: Save Job](#)

Step 1: Job Type and Name

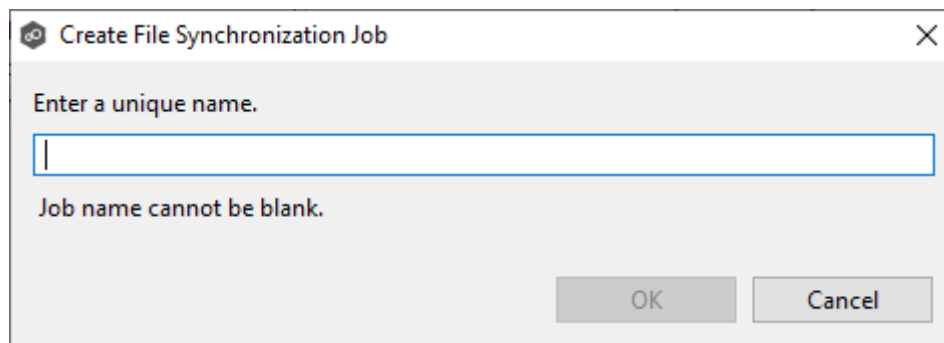
1. Open Peer Management Center.
2. From the **File** menu, select **New Job** (or click the **New Job** button on the toolbar).

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



- Click **File Synchronization**, and then click **Create**.
- Enter a name for the job in the dialog that appears.

The job name must be unique.



- Click **OK** to proceed to the [Participants](#) page.

Step 2: Participants

Use the **Participants** page to select the file-sharing protocol for the job and add participants to the job.

Overview of Adding Participants

1. Choose a file-sharing protocol: **SMB**, **NFS**, or **Multi-Protocol**, based on the workload for this job
2. Click the **Add** button to start the process of adding a participant.

The **Add New Participant** wizard opens; it walks you through the steps for adding a participant:

- a. [Selecting a Management Agent](#), which is the Agent that will manage the storage device that hosts data you want to replicate.
- b. [Selecting the type of storage platform](#) that hosts data you want to replicate.
- c. [Entering the credentials needed to access a specific storage device and providing other storage information](#).
- d. [Entering the path](#) to the [watch set](#) (the data that you want to replicate) and selecting whether participant will be a [seeding target](#).
- e. (Optional) [Enabling Edge Caching](#) for the participant.

Once you have added a participant, it is listed in the **Participants** table.

3. To add more participants, click **Add** and repeat the steps for each participant you want to add to the job.
4. Once you have added all the participants, click **Next**.

The page that appears next depends on options you selected when adding a participant:

- if you have enabled Edge Caching for a participant, the [Edge Caching](#) page appears.
- Otherwise, the [File Metadata](#) page appears.

Overview

The **Management Agent** page lists available Agents. Select an Agent to be responsible for managing the volumes, shares, exports, or folders to be synchronized in this job.

- Choose a file-sharing protocol: **SMB**, **NFS**, or **Multi-Protocol**, based on the workload for this job

Note: Not all licenses support multi-protocol configurations. Verify your license before configuring jobs that require both SMB and NFS.

- After selecting a protocol, only compatible Agents will be available:
 - **SMB** requires Windows-based Agents.
 - **NFS** or **Multi-Protocol** requires Linux-based Agents.
 - All Agents in the job must use the same protocol—Windows and Linux Agents cannot be mixed in the same job.

The wizard will guide you through configuring protocol-specific settings.

Note: Multiple Agents can be assigned to a storage device, but each must manage different volumes, shares, exports, or folders.

Adding Participants to a Job

To add participants to a job:

1. Begin by selecting the file-sharing protocol: **SMB**, **NFS**, or **Multi-Protocol**.

The **default protocol** is determined by the available Agents and licensing:

- If all Agents are Windows-based, SMB will be the default and the only option.
- If all Agents are Linux-based and Multi-Protocol is not allowed by the user license, NFS will be the default and the only option.
- If all Agents are Linux-based and Multi-Protocol is allowed by the user license, no default option will be preselected.

3. Select an Agent to manage the storage device.

Tip: If the Agent you want is not listed, the Peer Agent Service may not be running on the server hosting the Agent. Try restarting the Peer Agent Service. If the service successfully connects to the [Peer Management Broker](#), then the list of available agents will be updated with that Agent.

4. Click **Next** to proceed to the [Storage Platform](#) page.

The **Storage Platform** page lists the available storage platforms types. The available storage platform types depends on the protocol and Agent selected on the [Management Agent](#) page.

1. Select the storage platform type hosting the data you want to replicate.

The screenshot shows a window titled "Add New Participant" with a sub-header "Storage Platform". Below the sub-header is the instruction "Select the type of storage platform." On the left side, there is a navigation pane with the following items: "Management Agent", "Storage Platform" (which is highlighted), "Storage Information", "Path", and "Edge Caching". The main area of the window displays a list of storage platform options, each with a radio button and an icon:

- Windows File Server
- Linux File Server
- NetApp ONTAP
- Amazon FSx for NetApp ONTAP
- Dell PowerScale
- Dell Unity
- Nutanix Files

At the bottom of the window, there are four buttons: "< Back", "Next >" (which is highlighted with a blue border), "Finish", and "Cancel".

These are the available platforms:

Protocol/Agent Selection	Available Storage Platform Types
SMB Protocol and Windows-based Agent	<ul style="list-style-type: none">• Windows File Server• NetApp ONTAP• Amazon FSxN• Dell PowerScale• Dell Unity• Nutanix Files
NFS Protocol and Linux-based Agent	<ul style="list-style-type: none">• Linux File Server• NetApp ONTAP• Amazon FSxN• Dell PowerScale
Multi-Protocol and Linux-based Agent	<ul style="list-style-type: none">• NetApp ONTAP• Amazon FSxN• Dell PowerScale

2. Click **Next** to proceed to the [Storage Information](#) page.

On the **Storage Information** page, enter the required details for the storage platform selected on the [Storage Platform](#) page. The fields displayed vary based on the selected platform type.

1. If you selected **Windows File Server** or **Linux File Server**, provide the required configuration details. See [Windows File Server](#) or [Linux File Server](#) for more information.

If you selected any other storage platform type, skip to Step 2.

2. Select **New Credentials** or **Existing Credentials**.
3. Enter credentials:
 - If you select **Existing Credentials**, select a entry from the drop-down list, and then click **Next** to proceed to the [Path](#) page.
 - If you select **New Credentials**, enter the required connection details. The fields vary depending on the storage platform type. For details, see:
 - [Amazon FSxN](#)
 - [Dell PowerScale](#)
 - [Dell Unity](#)
 - [NetApp ONTAP](#)
 - [Nutanix Files](#)
4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
5. Click **Next** to proceed to the [Path](#) page.

Amazon FSxN

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Storage Virtual Machine hosting the data to be synchronized.

The screenshot shows a window titled "Add New Participant" with a sidebar on the left containing "Management Agent", "Storage Platform", "Storage Information" (selected), and "Path". The main area is titled "Storage Information" and contains the instruction "Enter the information required to connect to the storage device." Under the "Credentials" section, there are two radio buttons: "New Credentials" (selected) and "Existing Credentials". The "New Credentials" section includes five input fields: "*SVM Name:", "*SVM User Name:", "*SVM Password:", "SVM Management IP:", and "*Peer Agent IP:". There is an "Advanced" button to the right of the "*Peer Agent IP" field. The "Existing Credentials" section has a drop-down menu and an "Edit" button. At the bottom left of the main area is a "Validate" button. At the bottom right of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. Enter credentials:

- 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**, supply the required information.

Field	Description
SVM Name	Enter the name, FQDN, or IP address of the Storage Virtual Machine (SVM) hosting the data to be replicated.
SVM User Name	Enter the user name for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	Optional. Enter the IP address used to access the management API of the Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already permit management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually enter the address.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
5. Click **Next** to proceed to the [Path](#) page.

Dell PowerScale

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect the PowerScale cluster hosting the data to be synchronized.
2. Enter credentials:
 - 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**, supply the required information.

The information required will vary, depending on whether you select **RabbitMQ** or **Syslog** as the connection type, due to the distinct protocols and mechanisms they employ for communication.

If you selected **RabbitMQ**, this page appears:

The screenshot shows a window titled "Add New Participant" with a "Storage Information" tab selected. The window contains a sidebar with navigation options: "Management Agent", "Storage Platform", "Storage Information" (highlighted), "Path", and "Edge Caching". The main area is titled "Storage Information" and contains the instruction "Enter the information required to connect to the storage device." Below this, there are two sections for "Credentials". The "New Credentials" section is active, featuring radio buttons and input fields for: "*Cluster Name:", "*Cluster Management IP:", "*Cluster Username:", "*Cluster Password:", and "Cluster Access Zone:". The "Connection Type" is set to "RabbitMQ" (selected) with "Syslog" as an alternative. An "Advanced" button is located to the right of the "Cluster Access Zone" field. Below the "New Credentials" section is the "Existing Credentials" section, which is currently empty. At the bottom left of the main area is a "Validate" button. A footer message reads: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments." At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

If you selected **Syslog**, this page appears:

Add New Participant

Storage Information

Enter the information required to connect to the storage device.

- Management Agent
- Storage Platform
- Storage Information**
- Path
- Edge Caching

Credentials

New Credentials

*Cluster Name:

*Cluster Management IP:

*Cluster Username:

*Cluster Password:

Cluster Access Zone:

Connection Type: Syslog RabbitMQ

Syslog

*Agent IP Address:

*Listening Port:

*SSL Certificate Path:

*SSL Private Key Path:

SSL Private Key Password:

Existing Credentials

You must enter a Cluster Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments.

Field	Description
Cluster Name	Enter the name of the PowerScale cluster hosting the data to be replicated.
Cluster Management IP	Enter the IP address to use to access the REST-based API integrated into the PowerScale cluster. Required only if multiple Access Zones are in use on the cluster.
Cluster Username	Enter the user name for the account managing the PowerScale cluster.
Cluster Password	Enter the password for account managing the PowerScale cluster.
Cluster Access Zone	Optional. The name of the access zone that is being monitored.
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: <ul style="list-style-type: none"> • Opt for Syslog if the storage device directly transmits notifications to the Agent. • Opt for RabbitMQ if you're utilizing the CEE framework to dispatch notifications to the Agent.

3. If you selected **Syslog**, you must also enter values for the following fields:

Field	Description
Agent IP Address	Select the IP address of the server hosting the Agent that manages the PowerScale cluster. The cluster must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.
Listening Port	Enter the port over which the Agent will receive TLS-based syslog events from the PowerScale cluster.

Field	Description
SSL Certificate Path	Enter the path to the certificate to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private Key Path	Enter the path to the private key to be used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .
SSL Private Key Password	[Optional] If your private key is protected with a password, enter it here. For more information, see https://kb.peersoftware.com/kb/dell-powerscale-syslog-configuration-guide .

4. Click **Advanced** if you want to set [advanced options](#).
5. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
6. Click **Next** to proceed to the [Path](#) page.

Dell Unity

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the CIFS Server hosting the data to be synchronized.

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. The main heading is "Storage Information" with a sub-heading "Enter the information required to connect to the storage device." On the left is a sidebar with a tree view containing "Management Agent", "Storage Platform", "Storage Information" (highlighted), "Path", and "Edge Caching". The main area is titled "Credentials" and has two radio buttons: "New Credentials" (selected) and "Existing Credentials". Under "New Credentials" are four text input fields: "*NAS Server Name:", "*Unisphere Management IP:", "*Unisphere Username:", and "*Unisphere Password:". An "Advanced" button is to the right of the password field. Under "Existing Credentials" is a drop-down menu and an "Edit" button. At the bottom left is a "Validate" button with the message "You must enter a NAS Server Name." Below that is a link: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments." At the bottom right are four buttons: "< Back" (highlighted), "Next >", "Finish", and "Cancel".

2. Enter credentials:

- 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**, supply the required information.

Field	Description
CIFS Server Name	Enter the name of the NAS server hosting the data to be replicated.
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.
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.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
5. Click **Next** to proceed to the [Path](#) page.

Linux File Server

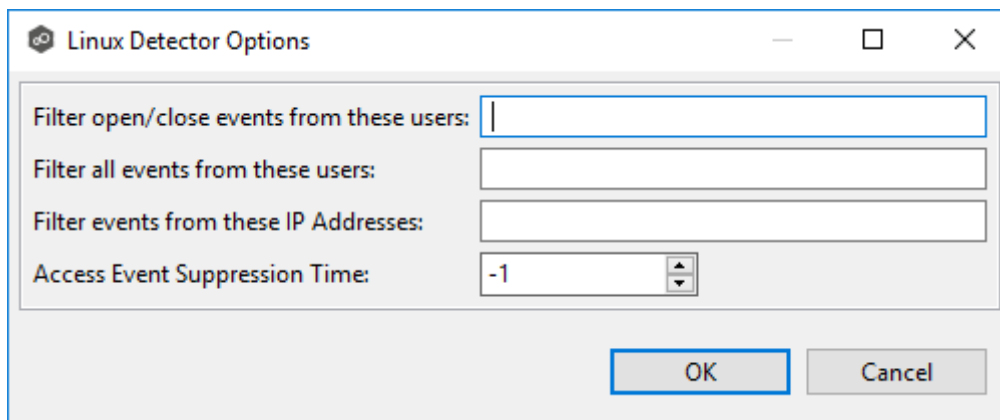
1. Select the **Detector Type**.

Detector type refers to the type of storage platform that is being monitored for real-time file activity.

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. Below the title bar, the text "Storage Information" is displayed, followed by the instruction "Enter the information required to connect to the storage device." On the left side, there is a vertical list of steps: "Management Agent", "Storage Platform", "Storage Information" (which is highlighted with a grey background), and "Path". The main area of the dialog is a large, empty grey rectangle. At the top of this area, there is a label "Detector Type:" followed by a dropdown menu showing "Linux Connector" and a downward arrow. To the right of the main area, there is a button labeled "Advanced". At the bottom of the dialog, there are four buttons: "< Back", "Next >" (which is highlighted with a blue border), "Finish", and "Cancel".

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Next** to proceed to the [Path](#) page.

1. Modify the options as desired.



The image shows a dialog box titled "Linux Detector Options". It contains four input fields and two buttons. The first field is labeled "Filter open/close events from these users:" and is empty. The second field is labeled "Filter all events from these users:" and is empty. The third field is labeled "Filter events from these IP Addresses:" and is empty. The fourth field is labeled "Access Event Suppression Time:" and contains the value "-1". The "OK" button is highlighted with a blue border, and the "Cancel" button is disabled.

Filter open/close events from these users:	<input type="text"/>
Filter all events from these users:	<input type="text"/>
Filter events from these IP Addresses:	<input type="text"/>
Access Event Suppression Time:	<input type="text" value="-1"/>

Option	Description
Filter open/close events from these users	Enter a comma-separated list of user account names from which all file open and close events will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user account under which a backup or archival service is running.
Filter all events from these users	Enter a comma-separated list of user account names from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the user account under which a backup or archival service is running.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the IP addresses on which a backup or archival service is running.
Filter events from these local processes	Enter a comma-separated list of local process names on the Agent server from which all file activities will be ignored. This option can be used to filter out events from backup and/or archival services by filtering on the specific process names under which a backup or archival service is running.
Access Event Suppression Time	Enter the number of seconds to delay an open event before being processed. Use this option to reduce the amount of chatter generated by Windows clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time an open event will be delayed based on the load of the system.

2. Click **OK**.

NetApp ONTAP

1. Select **New Credentials** or **Existing Credentials** to provide the credentials to connect to the Storage Virtual Machine hosting the data to be synchronized:

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. The window has a sidebar on the left with the following items: "Management Agent", "Storage Platform", "Storage Information" (which is highlighted), and "Path". The main area of the window is titled "Storage Information" and contains the instruction "Enter the information required to connect to the storage device." Below this instruction, there are two sections for "Credentials". The first section, "New Credentials", is selected with a radio button and contains five input fields: "*SVM Name:", "*SVM User Name:", "*SVM Password:", "SVM Management IP:", and "*Peer Agent IP:". An "Advanced" button is located to the right of the "*Peer Agent IP:" field. The second section, "Existing Credentials", is unselected and contains a drop-down menu and an "Edit" button. At the bottom of the main area, there is a "Validate" button followed by the text "You must enter the name." Below this is a link: "Having trouble connecting? Please verify that all [prerequisites](#) are met for NetApp ONTAP environments." At the very bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. Enter credentials:

- 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**, supply the required information.

Field	Description
SVM Name	Enter the name, FDQN, or IP address of the Storage Virtual Machine (SVM) hosting the data to be replicated.
SVM User Name	Enter the user name for the account managing the Storage Virtual Machine. The account must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. The account must not be a cluster management account.
SVM Management IP	Optional. Enter the IP address used to access the management API of the Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already permit management access, this field is not required.
Peer Agent IP	Select the IP address of the server hosting the Agent that manages the Storage Virtual Machine. The Storage Virtual Machine must be able to route traffic to this IP address. If the desired IP address does not appear, manually enter the address.

3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
5. Click **Next** to proceed to the [Path](#) page.

Nutanix Files

1. Select **New Credentials** or **Existing Credentials** to enter the credentials to connect to the Nutanix Files cluster hosting the data to be synchronized.

The screenshot shows a window titled "Add New Participant" with a "Storage Information" tab selected. The window contains a sidebar with navigation options: "Management Agent", "Storage Platform", "Storage Information" (highlighted), and "Path". The main area is titled "Storage Information" and contains the instruction "Enter the information required to connect to the storage device." Below this, there are two sections for "Credentials". The "New Credentials" section is selected and contains four input fields: "*Nutanix File Server Name:", "*Username:", "*Password:", and "*Peer Agent IP:". There is an "Advanced" button to the right of the "*Peer Agent IP:" field. The "Existing Credentials" section is unselected and contains a drop-down menu and an "Edit" button. At the bottom left of the main area is a "Validate" button. At the bottom right of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. Enter credentials:

- 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**, supply the required information.

Field	Description
Nutanix File Server Name	Enter the name of the Nutanix Files cluster hosting the data to be replicated.
Username	Enter the user name for the account managing the Nutanix Files cluster via its management APIs.
Password	Enter the password for the account managing the Nutanix Files cluster via its management APIs.
Peer Agent IP	Enter the IP address of the server hosting the Agent that manages the Nutanix Files cluster. The Files cluster must be able to route traffic to the specified IP address. If the desired IP address does not appear, manually enter the address. The IP address should not point to the Files cluster itself.

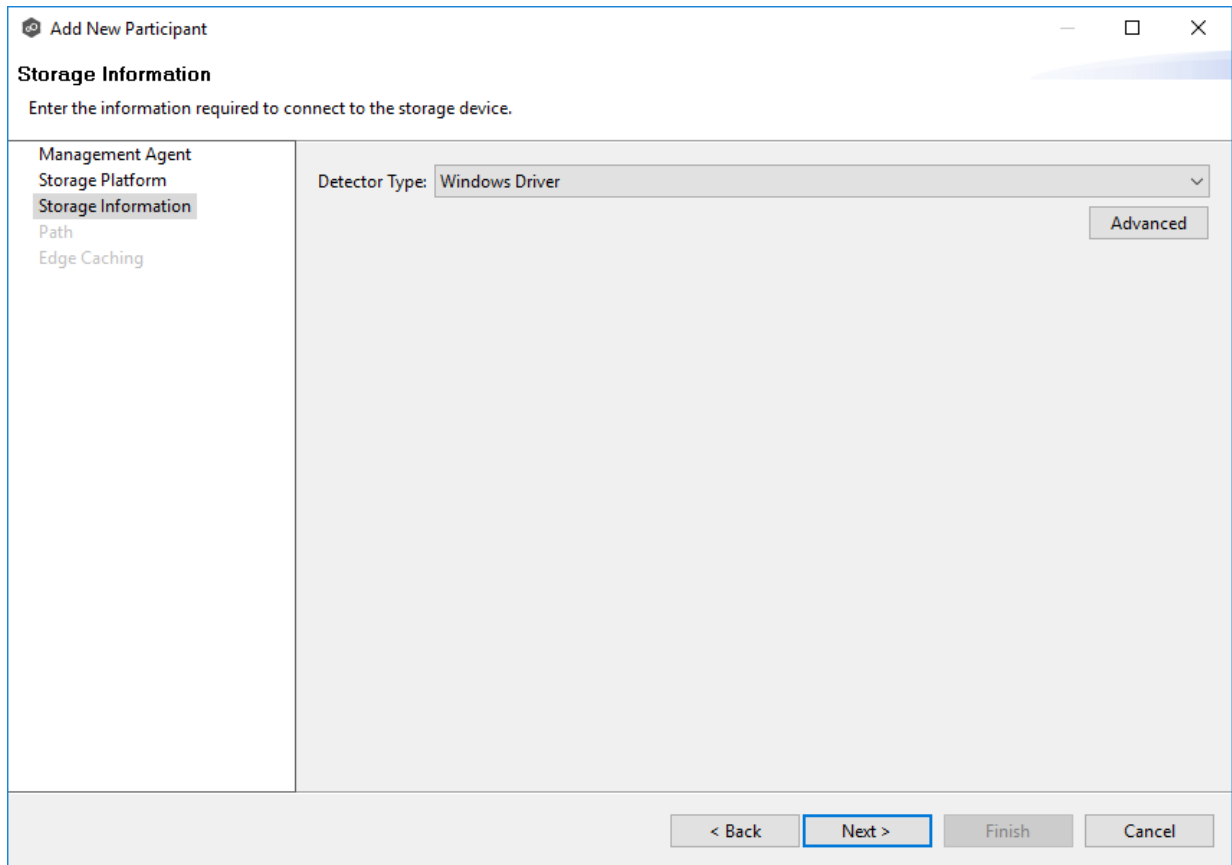
3. Click **Advanced** if you want to set [advanced options](#).
4. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
5. Click **Next** to proceed to the [Path](#) page.

Windows File Server

1. Select the detector type:

Detector type refers to the type of storage platform that is being monitored for real-time file activity.

- Select **Windows Driver** for more robust logging and better performance (Recommended).
- Select **Windows** if suggested by Peer Technical Support.

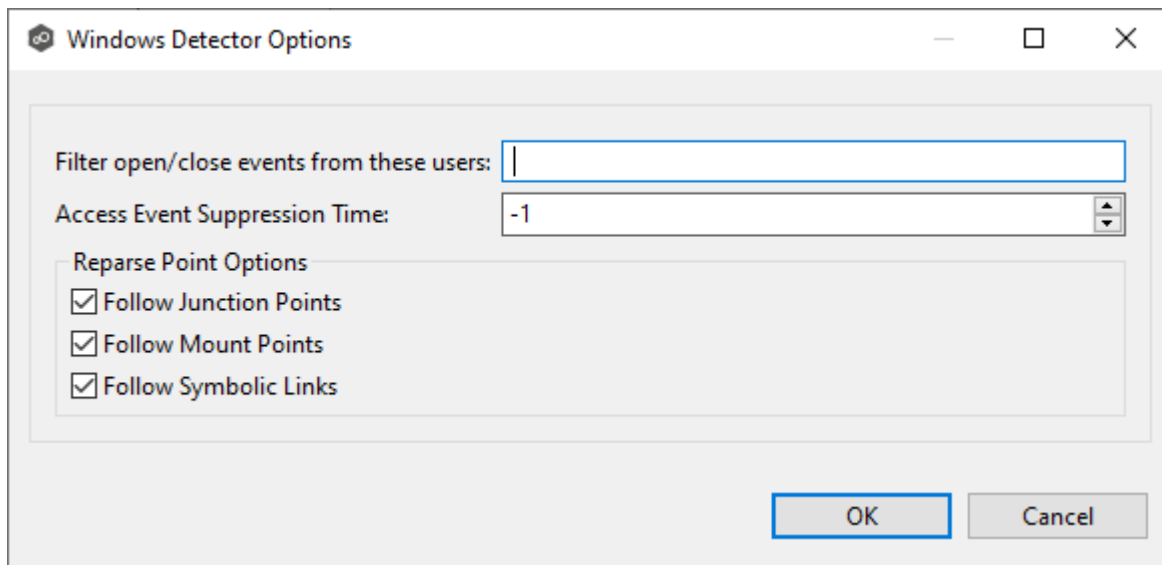


2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Next** to proceed to the [Path](#) page.

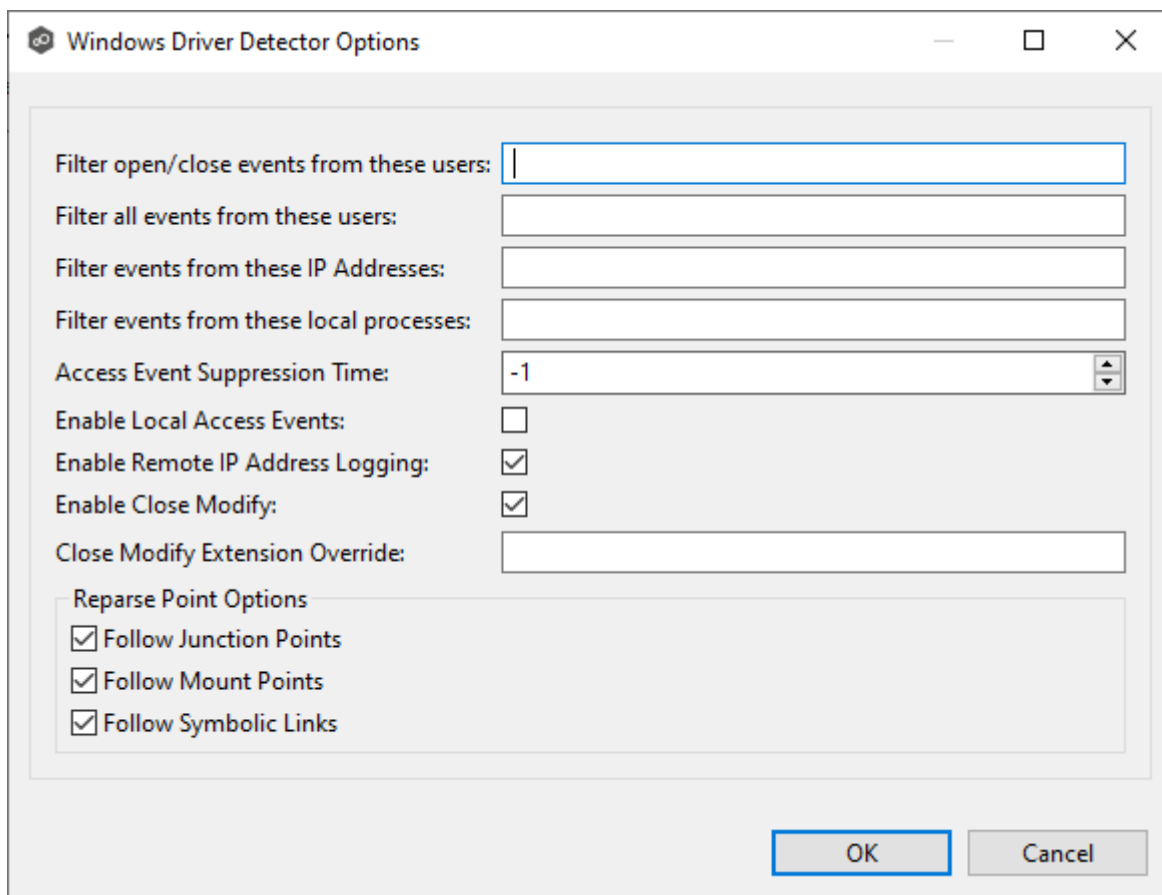
1. Modify the options as desired.

The available options depend on the detector type selected: **Windows** or **Windows Driver**.

Windows



Windows Driver



Option	Description
Filter open/close events from these users	A comma-separated list of user account names from which all opens and closes will be ignored. Ideal for filtering out events from backup and/or archival services by filtering on the username under which a backup and/or archival service is running.
Access Event Suppression Time	Represents number of seconds an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time an open event will be delayed based on the load of the system.
Follow Junction Points	Enables junction point support for the selected Windows File Server.
Follow Mount Points	Enables mount point support for the selected Windows File Server.
Follow Symbolic Links	Enables symbolic link support for the selected Windows File Server.

For more information about junction points or symbolic links, contact [<% SUPPORT_EMAIL%](mailto:SUPPORT_EMAIL%)

2. Click **OK**.

The **Path** page allows you to specify the path to the volume, share, export, or folder you want to synchronize. The path defines the [watch set](#), which can contain only one volume, share, export, or folder. To synchronize multiple watch sets, create a separate job for each.

1. Enter the path or browse to the watch set:

The screenshot shows a window titled "Add New Participant" with a "Path" step selected in a sidebar. The main area contains a text input field labeled "Enter Path" with a "Browse" button to its right. Below the input field is a checkbox labeled "Seeding Target". At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel". A red error message at the top left of the main area reads "You must enter a valid path."

For SMB:

If the source storage device is:

- A Windows File Server, enter the local path to the watch set in the following format:
C:\data
- A NAS device, enter a path to the watch set in the following UNC format: \
\server_name_or_ip\shared_folder\file_or_directory_path

For NFS:

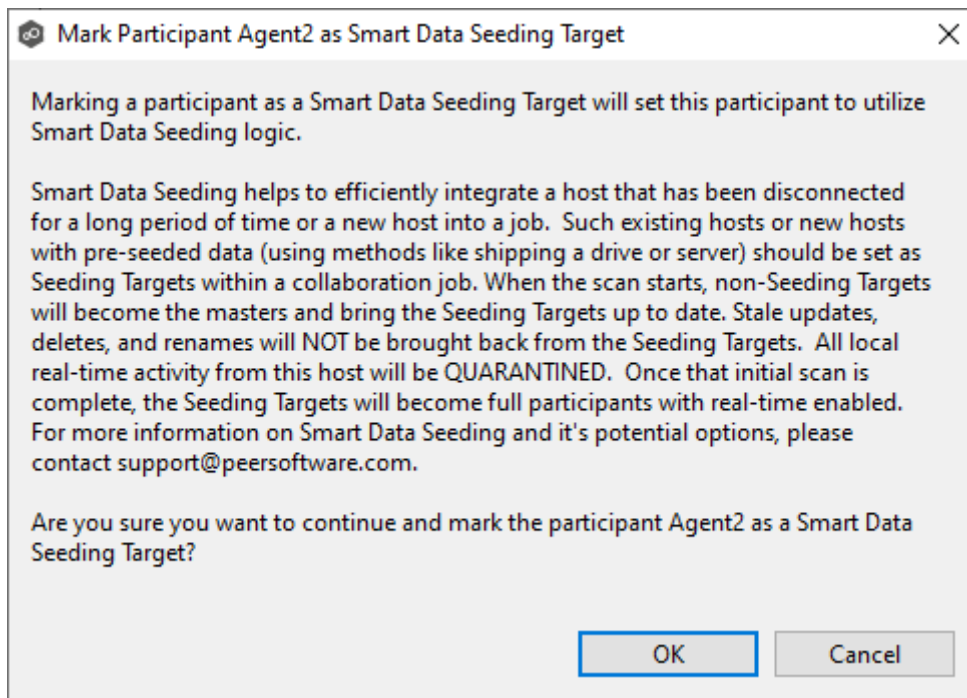
If the source storage device is:

- A Linux File Server, enter the local path to the watch set in the following format: /mnt/data
Note that the path is case-sensitive.
- A NAS device, enter the export path to the watch set in the following format:
host:/volume

For Multi-Protocol:

- Enter the export path to the watch set in the following format: `host:/volume`
2. (Optional) Select the **Seeding Target** checkbox, and then click **OK** in the dialog that appears.

If you select this option, a message describing seeding behavior is displayed. Multiple participants in a File Synchronization job can be set as smart data seeding targets; however, at least one participant should not be set as a smart data seeding target. This participant will be acting as the "master" source for the smart data seeding targets. For more information about smart data seeding, see [Smart Data Seeding](#) or contact [Peer Support](#).



3. If you are creating a multi-protocol job, you may need to specify SMB access. If not, skip to Step 5.

The screenshot shows a window titled "Add New Participant" with a sidebar on the left containing the following menu items: Management Agent, Storage Platform, Storage Information, and Path (which is highlighted). The main content area is titled "Path" and displays an error message: "Path is missing the export name." Below the error message is a text input field containing "PTSV93A:/ Enter Path" and a "Browse" button. There is a checkbox for "Seeding Target" which is currently unchecked. A section titled "SMB Access" contains a warning icon and the text "Required only when synchronizing Windows-style permissions." Below this is a link "Edit SMB Access Configuration". The "SMB Access Configuration:" label is followed by a dropdown menu. The "SMB Access Share:" label has two radio button options: "Use Admin Share (CS)" (which is selected) and "Specify Share Name:" followed by an empty text input field. At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

SMB access configuration enables the Linux Agent to access data on the file server using the SMB file sharing protocol. It is required in the following situations when creating a multi-protocol job:

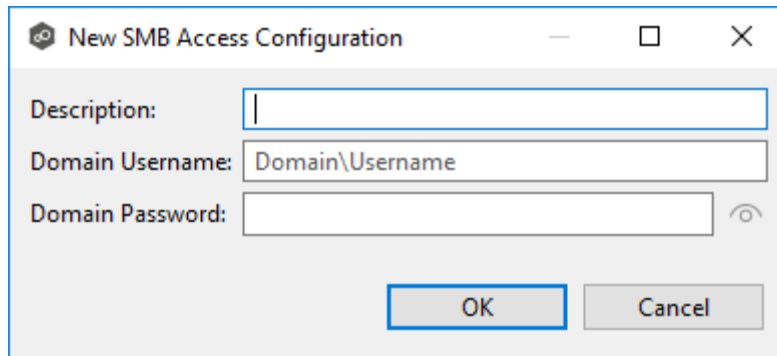
- Synchronizing Windows-style permissions
- Synchronizing both Windows-style permissions and Linux-styled permission (mixed permissions)
- Using Nutanix Files.

You will select how you want to synchronize permissions when you get to the **File Metadata** page in this wizard.

Recommended: Create SMB access configurations in advance in the **Storage Platforms** section of **Preferences**.

4. Specify the SMB access configuration:

- To use an **existing** SMB access configuration:
 - a. Select it from the **SMB Access Configuration** drop-down.



The screenshot shows a dialog box titled "New SMB Access Configuration". It has three input fields: "Description" (empty), "Domain Username" (containing "Domain\Username"), and "Domain Password" (empty with a visibility icon). At the bottom, there are "OK" and "Cancel" buttons.

Field	Description
Description	Enter a description that identifies this configuration.
Domain Username	Enter the domain username in the format: <code>Domain\Username</code>
Domain Password	Enter the password for the domain account.

- d. Click **OK** to save the access configuration.
 - e. Click **Apply and Close** to apply the configuration to the job and close the dialog.
5. If edge caching is enabled for this job, click **Next**; otherwise Click **Finish** to complete the wizard for this participant.
 6. Return to [Step 2: Participants](#) to add more participants, if applicable. A File Synchronization job must have at least two participants. If you have added all the participants, continue with [Step 4: File Metadata](#).

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

Edge Caching is a method for conserving space on storage devices by caching files until needed. Edge Caching saves space by stubbing files and rehydrating them as needed. Edge Caching is optional; if you don't need to conserve space on the storage device managed by the Agent, then you do not need to select this option.

If you enable [Edge Caching](#) for a participant, you must designate the participant as either a **master** or **edge** participant.

- **Master participant** - A master participant always has a complete set of files for that job. None of the files are stubbed; they are stored physically on that device.
- **Edge participant** - A subset of the files stored on a master participant are physically stored on an edge participant; the rest of the files on an edge participant are stub files that take up minimal space. Edge Caching allows users to seamlessly retrieve stubbed files directly from a master participant as needed; when retrieved, the local stub file is rehydrated so that the full file is stored locally on the edge participant.

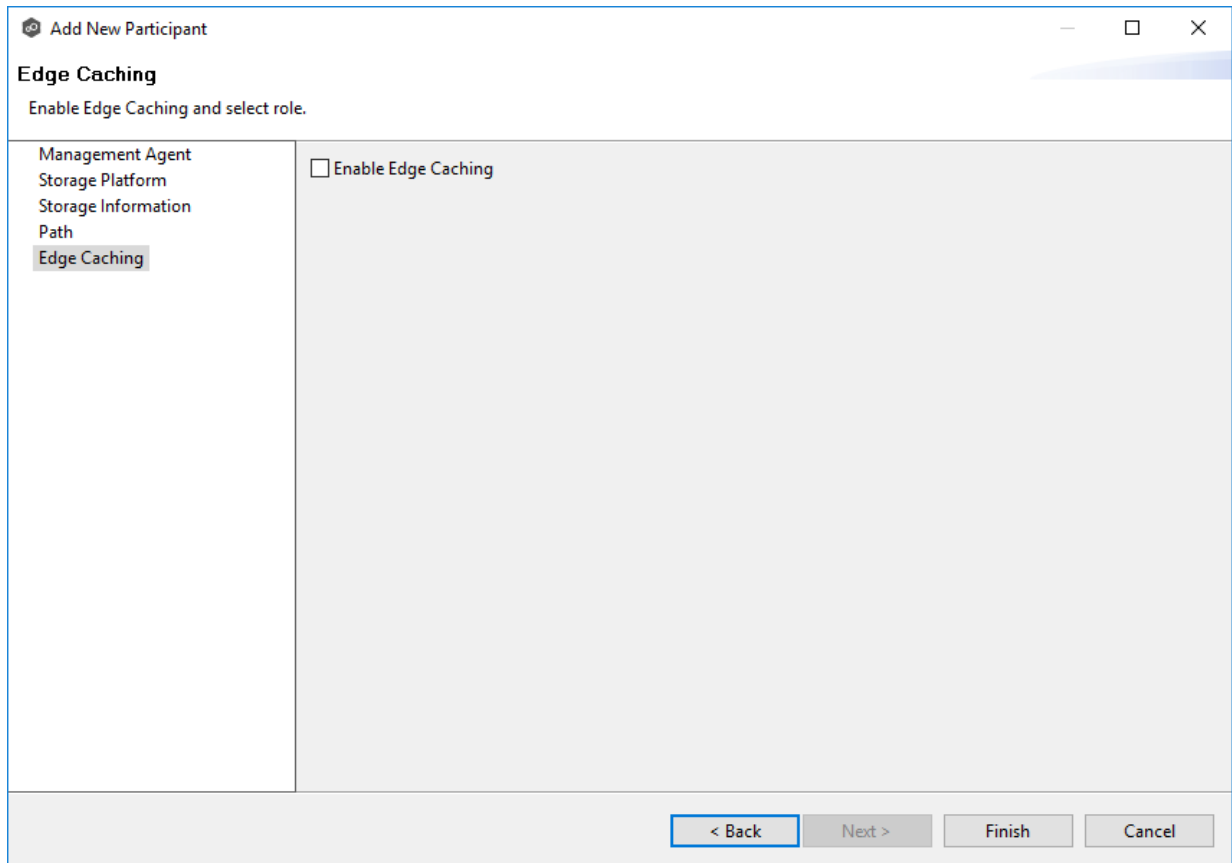
A job can have master and edge participants, as well as participants that don't have either role. If you do not choose to enable Edge Caching for a participant, it will always have a full set of files like a master participant but will not be used to serve file content to any edge participants.

Notes:

- A participant can be a master participant for some jobs and an edge participant for other jobs.
- A job needs at least one master participant that isn't a seeding target. If there is only one master participant for the job, it should not be a seeding target.

For more information about Edge Caching, see [Edge Caching](#) in [Advanced Topics](#).

1. Select the **Enable Edge Caching** checkbox if you want this participant to be able to use Edge Caching; otherwise, click **Finish**.



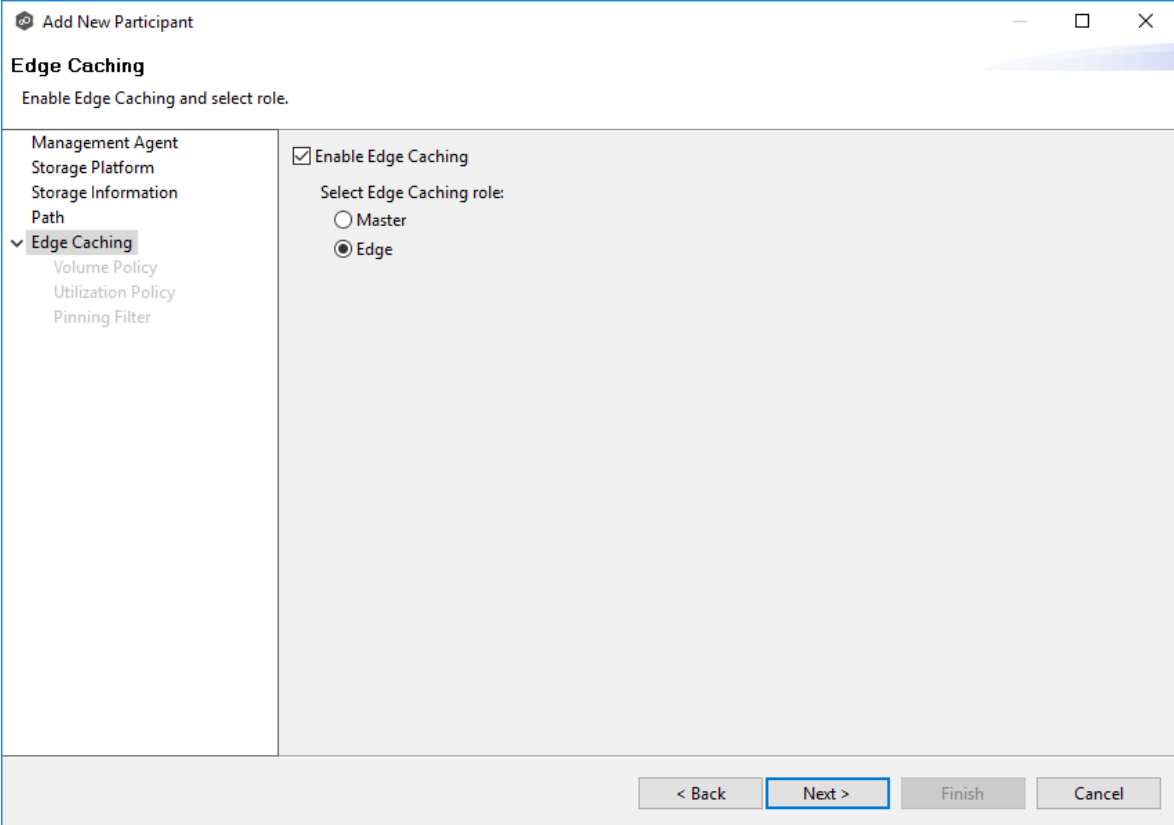
The screenshot shows a window titled "Add New Participant" with a close button (X) in the top right corner. Below the title bar, the section is labeled "Edge Caching" with the instruction "Enable Edge Caching and select role." The main area is split into two panes. The left pane contains a list of options: "Management Agent", "Storage Platform", "Storage Information", "Path", and "Edge Caching", with "Edge Caching" highlighted. The right pane contains a checkbox labeled "Enable Edge Caching" which is currently unchecked. At the bottom of the window, there are four buttons: "< Back" (highlighted with a blue border), "Next >", "Finish", and "Cancel".

If you enable Edge Caching, the Edge Caching role options are displayed; the **Master** role is selected by default.

2. Choose an Edge Caching role for the participant:
 - Choose **Master** if the storage device managed by the Agent will contain complete copies of all files for this job. Any type of storage platform can be a master participant.

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. The main heading is "Edge Caching" with the instruction "Enable Edge Caching and select role." Below this, there is a left-hand navigation pane with the following items: "Management Agent", "Storage Platform", "Storage Information", "Path", "Edge Caching" (which is expanded to show "Master Data Service"), and "Master Data Service". The main area of the dialog contains a checked checkbox for "Enable Edge Caching" and a section titled "Select Edge Caching role:" with two radio button options: "Master" (which is selected) and "Edge". At the bottom of the dialog, there are four buttons: "< Back", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

- Choose **Edge** if you want to conserve space on the storage device managed by the Agent. Only Windows File Servers can be an edge participant.



The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. The main heading is "Edge Caching" with the instruction "Enable Edge Caching and select role." On the left, a navigation pane lists several categories: "Management Agent", "Storage Platform", "Storage Information", "Path", "Edge Caching" (which is expanded to show "Volume Policy", "Utilization Policy", and "Pinning Filter"), and "Edge Caching" (which is selected). The main area contains a checkbox labeled "Enable Edge Caching" which is checked. Below it, the text "Select Edge Caching role:" is followed by two radio button options: "Master" (unselected) and "Edge" (selected). At the bottom of the window, there are four buttons: "< Back", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

3. Click **Next**.

- If you selected **Master**, continue with the [Master Data Service](#) page.
- If you selected **Edge**, continue with the [Volume Policy](#) page.

Master Data Service

The **Master Data Service** page appears if you chose the master role for the participant. The Master Data Service handles requests from edge participants for files on a master participant. The Master Data Service is installed on the Agent server as part of the Peer Agent installation process.

The first two fields on this page are automatically populated:

- **Protocol:** This field lists the protocol that will be used to transfer file content between master participants and edge participants. HTTPS is currently the only available option as it requires only a single open firewall port on each master participant and does a better job of handling latency over WAN-based connections.

- **Agent Name:** This field lists the name of the Management Agent that you selected at the beginning of Step 2.

1. (Optional) Enter a value for **Agent Alias**. The value can be a hostname, FDQN, or IP address.

A value for this field is required only if the name of the Agent cannot be converted to an IP address via DNS. If an alias is entered, it will be used by the Edge Service on edge participants to connect with the Master Data Service. If no alias is entered, the name in the Agent Name will be used.

2. (Optional) Modify the port number that the Master Data Service will listen on for this master participant.

A default value for the port number, 8446, is set when the Agent is installed. If you modify the port number, the Master Data Service is started with the new port number.

The screenshot shows a window titled "Add New Participant" with a sub-header "Master Data Service" and the instruction "Configure access to the Master Data Service." On the left is a sidebar with a tree view containing "Management Agent", "Storage Platform", "Storage Information", "Path", "Edge Caching" (expanded), and "Master Data Service" (selected). The main content area has four input fields: "Protocol:" with "HTTPS", "Agent Name:" with "Agent2", "Agent Alias:" (empty), and "Port:" with "8446". At the bottom right are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Note: If the Agent you selected is already being used as a master participant in another job utilizing Edge Caching, then the existing Master Data Service parameters will be displayed. You can edit the values by clicking the **Edit Master Data Service** link. If you modify the port number, the Master Data Service will be restarted, and the new port number will take effect immediately. Any modifications you apply will be applied to every other job that uses this Agent as a master participant.

4. Continue adding more participants if applicable or continue with [Step 3: Master-Edge Assignment](#).

Volume Policy

The **Volume Policy** page appears if you chose the edge role for the participant.

A volume policy is applied when a caching scan is run. The primary purpose of a **volume policy** is to specify how much space is available to Edge Caching on a specific volume (or drive letter), i.e., to define the **cache size**. The cache size specifies the maximum amount of disk space you want to allocate to Edge Caching for fully hydrated files on the volume specified by the path on the **Path** page. For example, if the participant is configured to monitor D:\Data, the volume policy for this participant would apply to the D volume.

The cache size can be specified as a percentage of the volume disk space or as a fixed size. For example, if an edge participant is configured to monitor a volume that has 1 TB of disk space, and you tell Edge Caching to use 75% of that volume, then up to 750 GB of files could be locally available on the volume monitored by that edge participant. For optimal performance, we recommend that this cache be dedicated to Edge Caching's use on this volume.

A volume policy applies to each job where the following three elements are true:

- Edge Caching is enabled for the job.
- The participant is an edge participant.
- The paths specified for each job share the same volume.

To create a volume policy:

1. In the **Cache Size** section, choose an option for setting the cache size:
 - Use up to X % of this volume
 - Use up to X size of this volume

Add New Participant

Volume Policy

Please enter a valid path

Management Agent
Storage Platform
Storage Information
Path
Edge Caching
Volume Policy
Utilization Policy
Pinning Filter

Cache Size

Use up to 75 % of this volume

Use up to 10 GB of this volume

Cache Threshold Alerts

Send alerts when:

Disk space is less than 512 MB

Cache usage exceeds 80 % of the cache size

Caching Scan Schedule

Scan every 1 day(s) at 10:00:00 PM

Define schedule

Run caching scan after job start

*Temporary Storage Path: Browse

\PeerTempPath

< Back Next > Finish Cancel

- In the **Cache Threshold Alerts** section, set threshold values for automatic alerts about free disk space and cache usage.

Peer Management Center will automatically display alerts in the **Alerts** tab when:

- The amount of free disk space on the volume falls below the specified value. For example, if a 1 TB volume has 500 MB of free space and the threshold is set to 512 MB, an alert will be sent.
- Cache usage on the volume exceeds the specified percentage of the cache size. For example, if the cache size is set to 80%, equating to 750 GB, Edge Caching will start sending alerts when it has used 600 GB.

You can also send cache threshold alerts via [email alerts](#) and [SNMP notifications](#). You configure these in [Edge Caching](#) preferences for File Collaboration and File Synchronization jobs.

- In the **Caching Scan Schedule**, set the frequency that the volume should be scanned to optimize the balance of fully hydrated vs stubbed files.

This scan can be run daily at a specified time, or you can define a more customized schedule.

4. In the **Temporary Storage Path** field, enter a path or browse to the location you want to be used as temporary storage space.

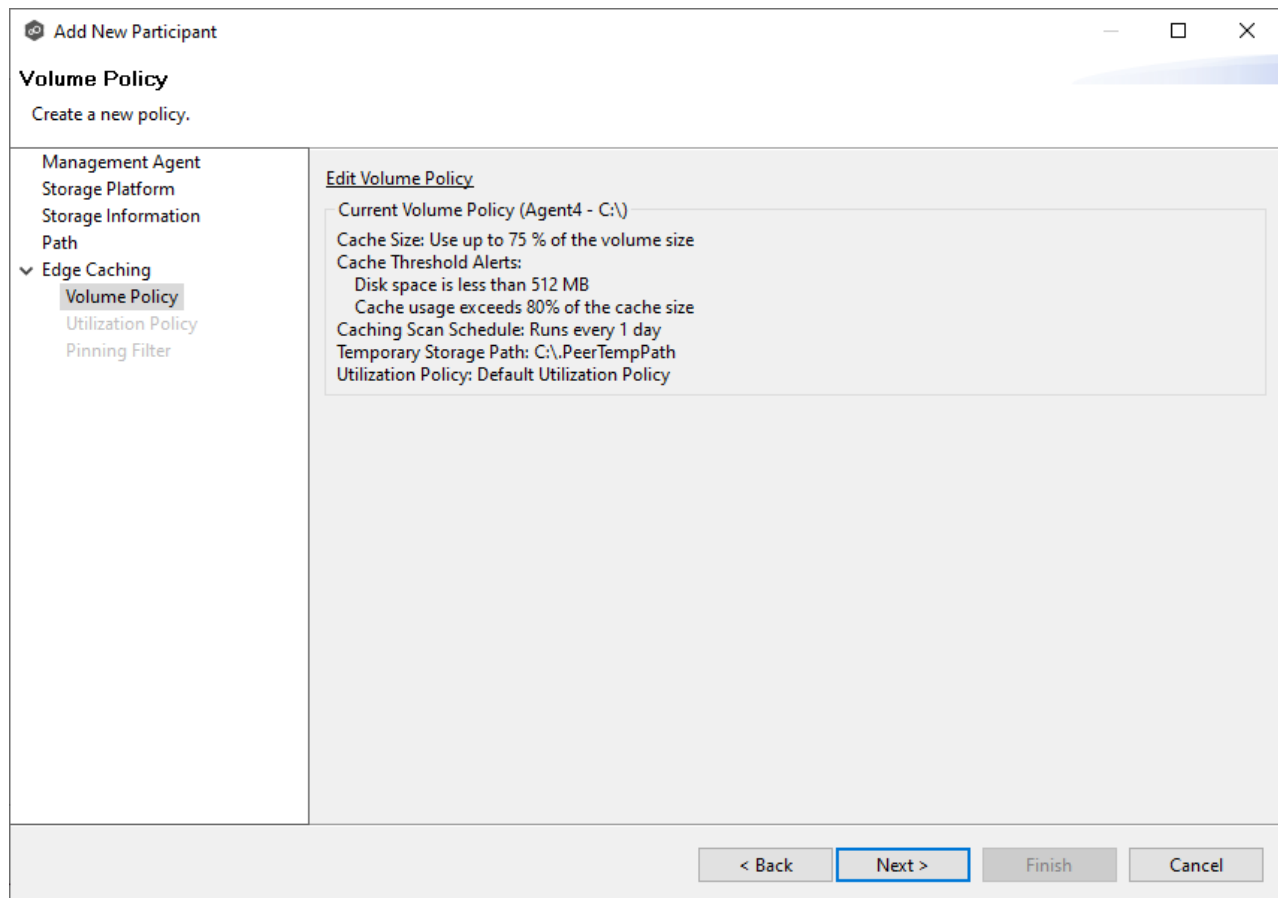
The temporary storage space will be used to store the content of stub files as they are being rehydrated. The content of files undergoing rehydration are referred to as **file blocks**. File blocks are fixed-length chunks of data that are read into memory when requested by an application. Edge Caching will create a subfolder named **.PeerTempPath** under the location that you specify and temporarily store the file blocks in that subfolder.

For optimal performance, we recommend that the temporary storage space be on the same volume as the watch set. If that is not possible, it should be on a high-performance disk.

5. Click **Next**.

The [Utilization Policy](#) page appears.

Note: If the Agent you selected is already being used as an edge participant in another job utilizing Edge Caching, the existing volume policy will be displayed on this page. You can edit the existing volume policy; however, be aware that any modifications to the volume policy will be applied to every other job that uses this Agent as an edge participant and "touches" the same volume.



Utilization Policy

The **Utilization Policy** page appears if you chose the edge role for the participant. The primary purpose of a **utilization policy** is to specify the parameters that govern when files on this edge participant should be stubbed or fully hydrated. Whereas the volume policy controls how much space is available to Edge Caching on a specific volume (or drive letter), the utilization policy controls whether to stub or hydrate a file.

Utilization policy parameters are based on the size of the files to be potentially stubbed and when they were last accessed and modified. A utilization policy enables you to balance getting the best performance while keeping the cache as full as possible. You can select an existing utilization policy to apply to the job or create a new utilization policy. Whereas a volume policy is specific to a volume, a utilization policy can be reused for multiple jobs.

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

- If you selected **Existing Policy**, select the policy, and then click **Next**.

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

- (Optional) In the **File Size** section, select one or both options:

Field	Description
Keep files local if less than X size	Select this option if you want files under a specified size to remain local.
Stub files if greater than X size	Select this option if you want files over a specified size to be stubbed.

- (Optional) In the **Time Period** section, select one of the options:

Field	Description
Keep recently used files local based on a dynamic set of rules	Select this option if you want Edge Caching to control when to stub files based on last accessed and last modified times. Edge Caching dynamically adjusts the accessed and modified time periods it uses based on the total amount of space that Edge Caching is actively using on a volume.
Keep recently used files local based on the following rules	Select this option if you want to specify the dates used when stubbing files based on the last accessed and last modified.

5. If you selected the **Keep recently used files local based on the following rules** option in **Time Period**, two additional options are enabled; select the options to be used and specify the time periods to be used:

Field	Description
Stub files if not modified within the past X time period	Select this option and specify a time range to use the last modified date of a file to determine if it should be kept local or stubbed.
Stub files is not accessed within the past X time period	Select this option and specify a time range to use the last accessed date of a file to determine if it should be kept local or stubbed.

6. (Optional) In the **Stubbing Override** section, select the checkbox and a time period if you want to use the last accessed date of all recently hydrated files to determine if they should be kept local or re-stubbed.
7. Click **Next** or **Finish**.

If you click **Next**, the [Pinning Filter](#) page appears.

Pinning Filter

The **Pinning Filter** page allows you to create a new pinning filter or select an existing pinning filter to apply to the job. A **pinning filter** specifies whether specific files or files in a particular

directory are always stubbed or always local on an edge participant. A pinning filter similar is to a utilization policy—it can be applied to multiple jobs. If there is a conflict between a pinning filter and utilization policy (where, for example, you might have something set to be always stubbed), the pinning filter will take precedence. Pinning filters are optional.

1. Select one of the options: **No Filter**, **New Filter**, or **Existing Filter**.

The screenshot shows a window titled "Add New Participant" with a "Pinning Filter" section. The section contains the instruction "Create a new pinning filter or select an existing one." Below this, there is a tree view on the left with "Pinning Filter" selected. The main area shows "Edit Pinning Filters" with three radio button options: "No Filter" (selected), "New Filter", and "Existing Filter". Below the radio buttons is a dropdown menu. At the bottom of the window are four buttons: "< Back", "Next >", "Finish" (highlighted with a blue border), and "Cancel".

2. If you selected **No Filter**, click **Finish**; if you selected **Existing Filter**, select the filter, and then click **Finish**.

If you selected **New Filter**, enter a name for the filter.

The screenshot shows the 'Add New Participant' dialog box with the 'Pinning Filter' section selected in the left-hand navigation pane. The main area is titled 'Pinning Filter' and contains a red error message: 'Name should not be empty.' Below this, there are two radio buttons: 'No Filter' and 'New Filter', with 'New Filter' selected. A text field labeled '*Name:' is empty. Underneath is a table titled 'Pinning Rules:' with two columns: 'Path' and 'Pinning State'. The table is currently empty. To the right of the table are three buttons: 'Create', 'Edit', and 'Delete'. Below the table, there is an 'Existing Filter' radio button and a dropdown menu. At the bottom of the dialog are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

3. Enter a name for the filter.
4. Click **Create**.

The **Create Pinning Rule** dialog appears.

The screenshot shows the 'Create Pinning Rule' dialog box. It has a title bar with a gear icon and the text 'Create Pinning Rule'. Below the title bar is a light gray instruction box: 'Enter a file name or path and choose pinning state.' There are two input fields: 'Pattern:' with an empty text box, and 'Pinning State:' with a dropdown menu showing 'Local at Edge'. At the bottom right are two buttons: 'OK' and 'Cancel'.

5. Enter a file name or path in the **Pattern** field and then choose a pinning state: **Local at Edge** or **Stubbed at Edge**.
6. Click **OK**.

The rule appears in the filter table.

Add New Participant

Pinning Filter
Create a new pinning filter or select an existing one.

Management Agent
Storage Platform
Storage Information
Path
Edge Caching
Volume Policy
Utilization Policy
Pinning Filter

Edit Pinning Filters
 No Filter
 New Filter

*Name: Outlook Files

Pinning Rules:

Path	Pinning State
*.pst	Local at Edge

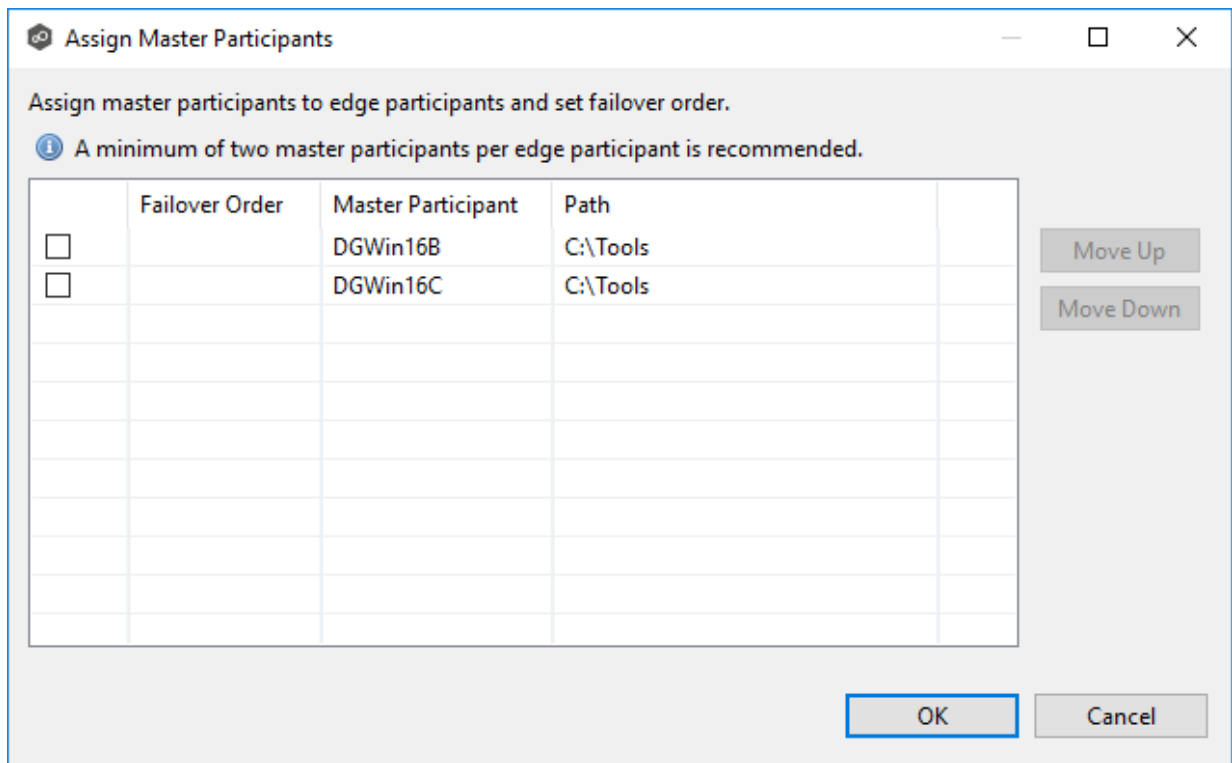
Existing Filter

Create
Edit
Delete

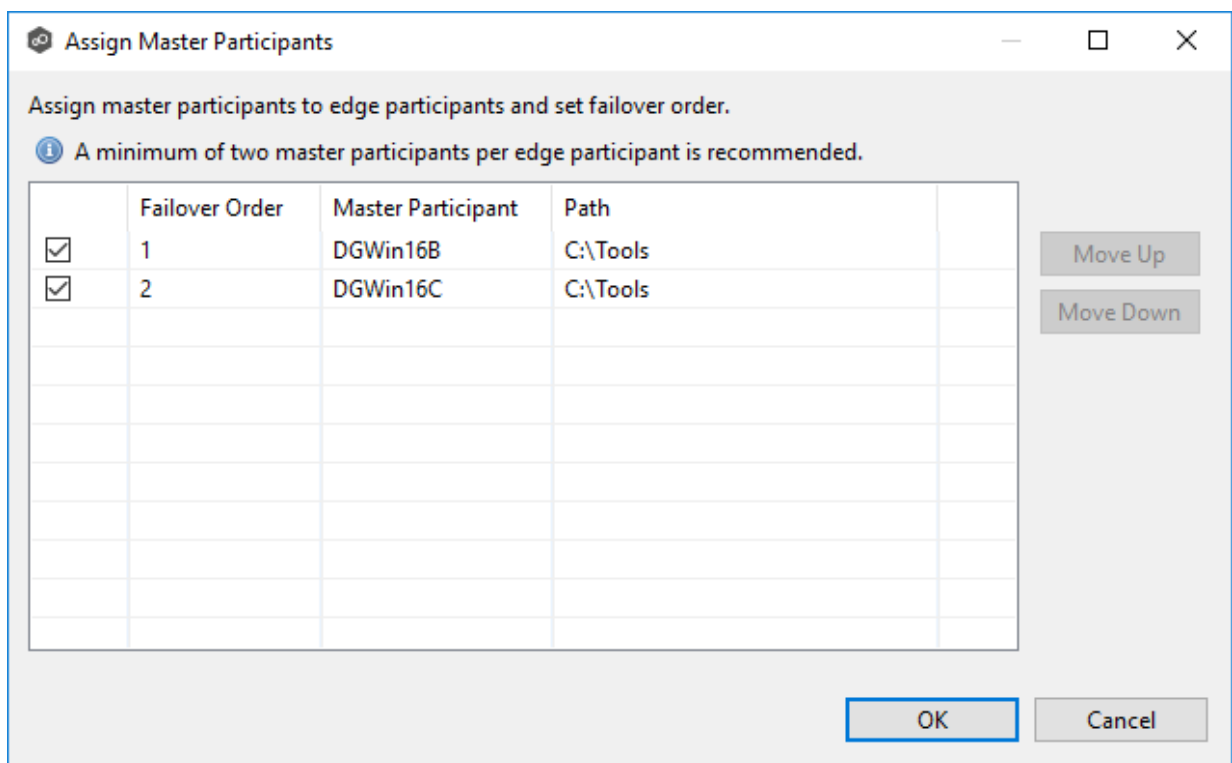
< Back Next > **Finish** Cancel

- (Optional) Create additional pinning rules.
- Click **Finish**.

The **Participants** page reappears. The participant is listed in the **Participants** table with the **Edge** role.



3. Select the master participants you want to assign to the edge participant.



7. Click **Next**.

The [File Metadata](#) page appears.

Step 4: File Metadata

This step is optional.

The **File Metadata** page allows you to specify whether you wish to replicate file security metadata and the types of metadata for synchronization. Additionally, it provides the ability to designate the metadata source (volume, share, export, or folder) to resolve conflicts during the initial synchronization. This designated source, utilized in case of conflicts, is referred to as the [master host](#).

The contents of the File Metadata page vary depending on whether you are using Windows-based or Linux-based Agents for your job:

- If you selected Windows-based, proceed with [File Metadata](#).
- If you selected Linux-based, proceed with [NFS File Metadata](#).
- If you selected Mixed Permissions, proceed with [Mixed Permissions](#).

For more information on synchronizing NTFS metadata, see [File Metadata Synchronization](#) in the [Advanced Topics](#) section.

1. Select Synchronization Options:

In the **Synchronize File Security Information** section, choose when to replicate metadata (you can select one or both options):

- **Enable synchronizing file security information in real-time:** Synchronizes metadata in real-time. Changes to the selected security descriptor components (Owner, DACL, and SACL) are synchronized to all participants as they occur.

- **Enable synchronizing file security information with master host during initial scan:** Synchronizes metadata during the initial scan. Changes to the selected security descriptor components are replicated during the initial scan.

2. Acknowledge the Message:

If you selected an option in **Synchronize File Security Information**, click **OK** in the message that appears.

3. Select Security Descriptor Components:

Choose the security descriptor components (Owner, DACL, and SACL) to synchronize.

4. Select a Master Host (If Applicable):

If you selected the option for metadata synchronization with the master host during the initial scan, choose the host to act as the master host in case of file metadata conflicts.

- If a master host is not selected, metadata synchronization will not be performed during the initial scan.
- If security descriptors do not match across participants during the initial scan, conflict resolution will use permissions from the designated master host as the winner.
- If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

5. Proceed:

Click **Next** to display the **Email Alerts** page.

Create File Synchronization Job Wizard

File Metadata

Configure the replication of security permissions.

Participants

- File Metadata
- Email Alerts

Synchronize File Security Information

- Enable synchronizing file security information in real-time
- Enable synchronizing file security information with master host during initial scan

Synchronize Security and ACL Options

- Owner
- Group
- Linux Permissions
- ACL (NFSv4 or POSIX)

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

Create File Synchronization Job Wizard

File Metadata

✖ You must select a master host when enabling security or attributes synchronization during initial scan.

Participants
File Metadata
Email Alerts

Synchronize File Security Information

- Enable synchronizing file security information in real-time
- Enable synchronizing file security information with master host during initial scan

Synchronize Security Options

Security Type: [How do permissions work with Multi-Protocol?](#)

- Owner
- Group
- Linux Permissions
- ACL (NFSv4 or POSIX)

Metadata Conflict Resolution

Select master host for initial scan:

To modify file metadata synchronization settings:

1. In the **Synchronize File Security Information** section, select when you want the metadata replicated (you can select one or both options):
2. By default both are selected.
 - **Enable synchronizing file security information in real-time** - Select this option if you want the metadata synchronized in real-time. If enabled, changes to the selected access controls (Owner, Group, Linux Permissions, and ACL) will be synchronized to all participants as they occur.
 - **Enable synchronizing file security information with master host during initial scan** - Select this option if you want the metadata replicated during the initial scan. If enabled, changes to the selected security descriptor components (Owner, Group, Linux Permissions, and ACL) will be synchronized during the initial scan.

Note: Nutanix does not support Access Control Lists (ACL) in the Network File System (NFS) version 4 (NFSv4) or POSIX formats.

3. If you selected an option in **Synchronize File Security Information**, click **OK** in the message that appears.
4. Select the access controls (Owner, Group, Linux Permissions, and ACL) that you want to synchronize.
5. If you selected the option for metadata synchronization with the master host during the initial scan, select the host to be used as the [master host](#) in case of file metadata conflict.

If a master host is not selected, then no metadata synchronization will be performed during the initial scan. If one or more access controls do not match across participants during the initial scan, [conflict resolution](#) will use permissions from the designated master host as the winner. If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

6. Click **Next**.

The [Email Alerts](#) page is displayed.

Edit File Synchronization Job

Participants
General
File and Folder Filters
Scheduled Replication Filters
Conflict Resolution
Delta Replication
File Metadata
File Locking
Target Protection
Email Alerts
SNMP Notifications
Tags

File Metadata

Synchronize File Attributes

- Enable synchronizing Linux file attributes in real-time
- Enable synchronizing Linux file attributes with master host during initial scan

Synchronize File Security Information

- Enable synchronizing file security information in real-time
- Enable synchronizing file security information with master host during initial scan
- Enable prevention of corrupt or blank Owner or ACLs on source or master host from being applied to any target host

Synchronize Security Options

Security Type: **Mixed Permissions** [How do Multi-Protocol and permissions work?](#)

Linux	Windows
<input checked="" type="checkbox"/> Owner	<input checked="" type="checkbox"/> Owner
<input checked="" type="checkbox"/> Group	<input checked="" type="checkbox"/> DACL: Discretionary Access Control List
<input checked="" type="checkbox"/> Linux Permissions	<input type="checkbox"/> SACL: System Access Control List
<input type="checkbox"/> ACL (NFSv4 or POSIX)	

Validate

Metadata Conflict Resolution

Select a master host for initial scan: **TEF1-UBUNTU-7**

File Timestamps Transfer

- Use source last accessed time

OK Cancel

Create File Synchronization Job Wizard

File Metadata

✘ You must select a master host when enabling security or attributes synchronization during initial scan.

Participants

- File Metadata
- Email Alerts

Synchronize File Security Information

- Enable synchronizing file security information in real-time
- Enable synchronizing file security information with master host during initial scan

Synchronize Security Options

Security Type: Mixed Permissions [How do permissions work with Multi-Protocol?](#)

Windows Permissions	Linux Permissions
<input checked="" type="checkbox"/> Owner	<input checked="" type="checkbox"/> Owner
<input checked="" type="checkbox"/> DACL: Discretionary Access Control List	<input checked="" type="checkbox"/> Group
	<input checked="" type="checkbox"/> Linux Permissions
	<input type="checkbox"/> ACL (NFSv4 or POSIX)

Validate

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

To modify file metadata synchronization settings:

- In the **Synchronize File Security Information** section, select when you want the metadata replicated (you can select one or both options):
 - Enable synchronizing file security information in real-time** - Select this option if you want the metadata synchronized in real-time. If enabled, changes to the selected access controls (Owner, Group, Linux Permissions, and ACL) will be synchronized to all participants as they occur.
 - Enable synchronizing file security information with master host during initial scan** - Select this option if you want the metadata replicated during the initial scan. If enabled, changes to the selected security descriptor components (Owner, Group, Linux Permissions, and ACL) will be synchronized during the initial scan.

Note: Nutanix does not support Access Control Lists (ACL) in the Network File System (NFS) version 4 (NFSv4) or POSIX formats.

2. If you selected an option in **Synchronize File Security Information**, click **OK** in the message that appears.
3. Select the access controls (Owner, Group, Linux Permissions, and ACL) that you want to synchronize.
4. If you selected the option for metadata synchronization with the master host during the initial scan, select the host to be used as the [master host](#) in case of file metadata conflict.

If a master host is not selected, then no metadata synchronization will be performed during the initial scan. If one or more access controls do not match across participants during the initial scan, [conflict resolution](#) will use permissions from the designated master host as the winner. If the file does not exist on the designated master host, a winner will be randomly picked from the other participants.

5. Click **Next**.

The [Email Alerts](#) page is displayed.

Step 5: Email Alerts

This step is optional.

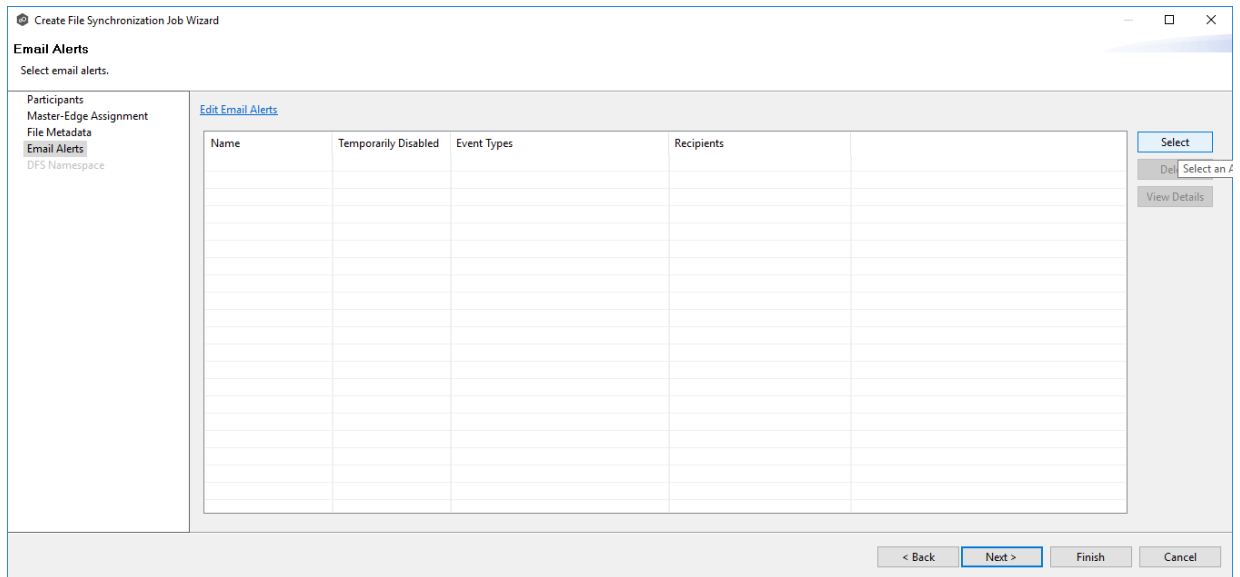
An email alert notifies recipients when a certain type of event occurs, for example, session abort, host failure, or system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in [Preferences](#) and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply.

To create a new alert, see [Email Alerts](#) in the [Preferences](#) section.

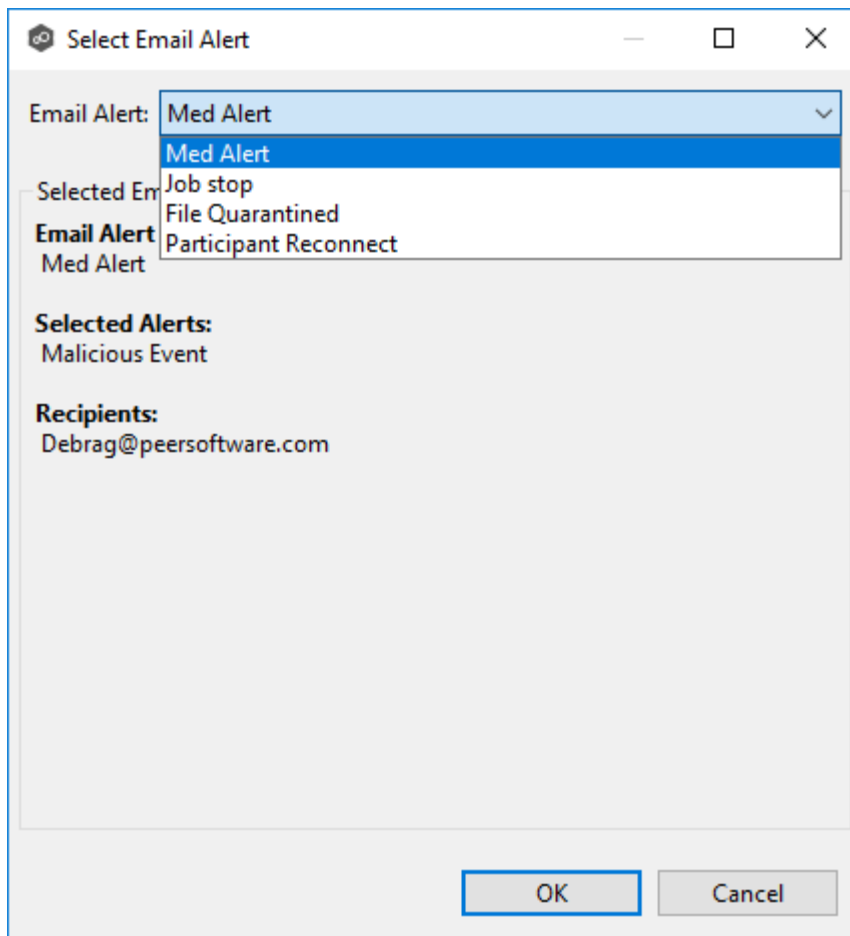
To apply an existing email alert to the job.

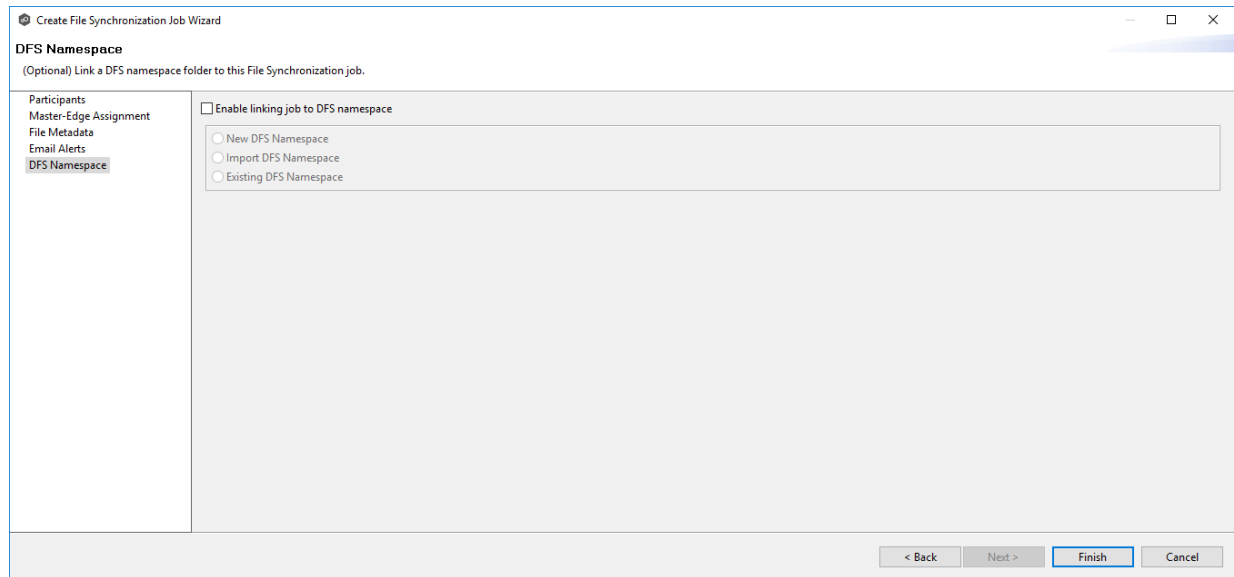
1. Click the **Select** button.



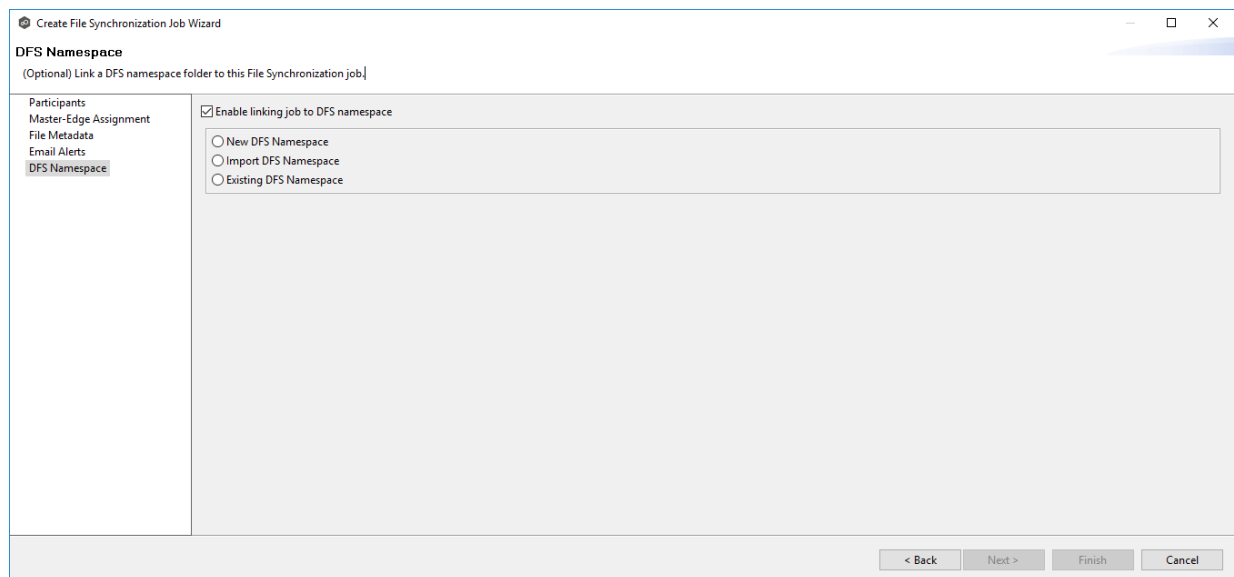
The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.





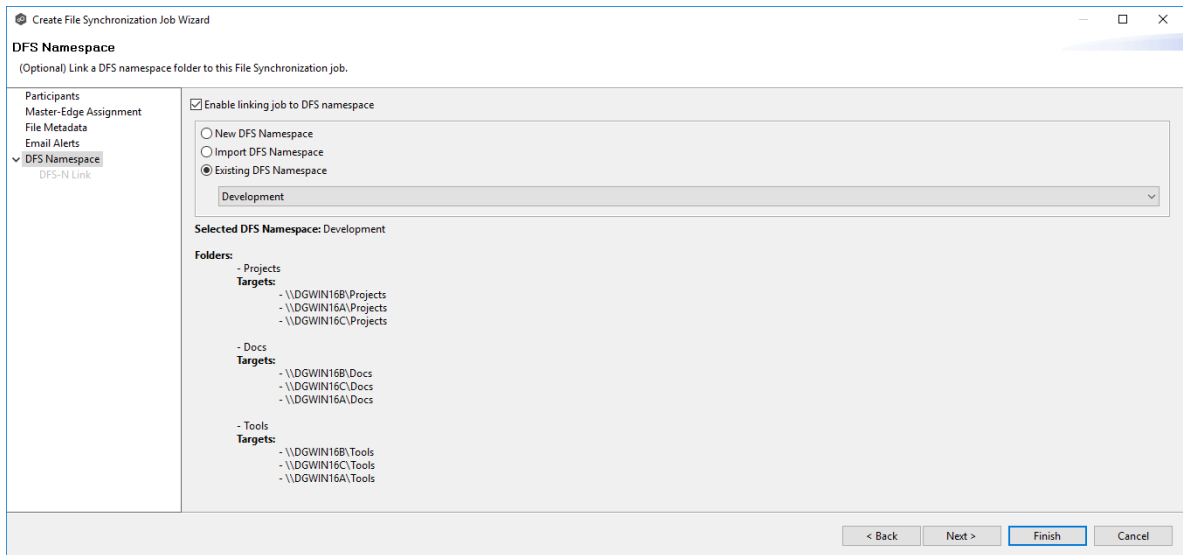
The three options are enabled.



2. Select one of the three options:

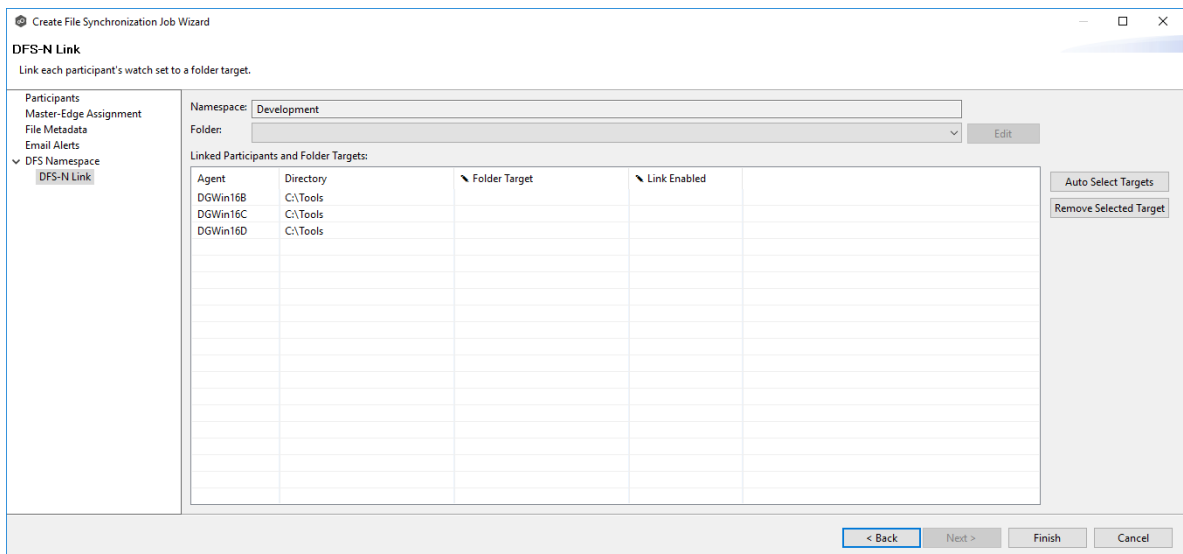
- **New DFS Namespace** - Select this option if you want to [create](#) a new namespace that will automatically be linked to this job. If you select this option, the **Create DFS-N Management Job Wizard** opens. Follow these steps to [create a new namespace](#).
- **Import DFS Namespace** - Select this option if you have a namespace that was created using the Microsoft DFS Management tool and is not currently being managed by a DFS-N Management job. If you select this option, the **Import Existing Namespaces** wizard opens. For detailed instructions, follow these steps to [import an existing namespace](#).

- **Existing DFS Namespace** - Select this option if you want to use an existing namespace that is being managed by a DFS-N Management job. If you select this option, it will display the namespace folder and folders associated with namespace. If you want to make changes to the namespace, you can edit the DFS-N Management job managing that namespace.



Click **Next** if you want to link participants to folder targets on the **DFS Link** page; otherwise continue with Step 3.

For more information about linking participants to folder targets, see [Linking a Namespace Folder to an Existing File Collaboration or Synchronization Job](#).



3. Click **Next** to continue to [Step 7: Save Job](#).

Step 7: Save Job

You are now ready to save the job configuration.

1. If you are satisfied with your job configuration, click **Finish** to save your job. Otherwise, click the **Back** button and make any necessary changes.

Congratulations! You have created a File Synchronization job. A summary of the job configuration is displayed in the runtime view of the job.

See [Running and Managing a File Synchronization Job Running](#) for more information.

The screenshot shows the Peer Management Center Client interface. The left pane displays a tree view of jobs and agents. The right pane shows the configuration summary for a File Synchronization job named FS-4. The configuration is divided into several sections:

- Configuration Summary**
 - Selected Participants and Configurations:** Agent2 (C:\FS-4 Folder A), Agent3 (C:\FS-4 Folder B), Agent4 (C:\FS-4 Folder C).
 - Edge Caching Configuration:** Agent2 (Master), Agent3 (Master), Agent4 (Edge).
 - General Settings:** Job Name: FS-4, Job ID: 145, Job Type: File Synchronization, Transfer Block Size: 2048 KB, Verify Checksum: true, Use Multipart Transfers: false, Global Real-Time Expedited Threads: 20.
 - Locking:** Allow Write Access During Synchronization: true, Exclusive Target Lock: false, Include MS Office User Lock Information: true, Include AutoCAD User Lock Information: false, Source Snapshot Synchronization: false, Source Vss Snapshot Synchronization: false, Snapshot File Extensions: mdb,accdb,zip,psd,ai,indd, Max File Size (MB): 512, Sync. On Save Extensions: none, Sync. On Save Delay: 20.
 - Conflict Resolution:** Latest Modified Time (Truncate Milliseconds=true), Quarantine Multi-edit Conflicts: true, Offline Folder Rename Detection Detection: false, Offline Delete Detection During Scan: true, Delete Detection Master Host: none.
 - Delta Replication Settings:** Enable Block/Byte Synchronization: true, Disable on Session Startup: false, Checksum Transfer Size: 256 KB, Delta Block Transfer Size: 1024 KB, Minimum File Size: 2048 KB, Minimum File Size Percentage Target/Source: 0.3, Excluded File Extensions: none, Excluded File Name Patterns: none.
- Status:** Stopped
- Alerts:** 0 errors, 0 warnings, 0 others. Filter by: Host, Severity, Type. Clear Alerts.

Editing a File Synchronization Job

You can edit a File Synchronization job while it is running; however, any changes will not take effect until the job is restarted.

Overview

When you create a File Synchronization job, the **Create Job** wizard guides you through the process, presenting the most common options for configuration. When editing a job, you have access to all options, allowing you to fine-tune the job configuration. Options not included in the initial job creation include:

- [Application Support](#)
- [Conflict Resolution](#)
- [Delta Replication](#)
- [DFS-N Link](#)
- [File and Folder Filters](#)
- [File Locking](#)
- [General](#)
- [Scheduled Replication Filters](#)
- [SNMP Notifications](#)
- [Target Protection](#)
- [Tags](#)

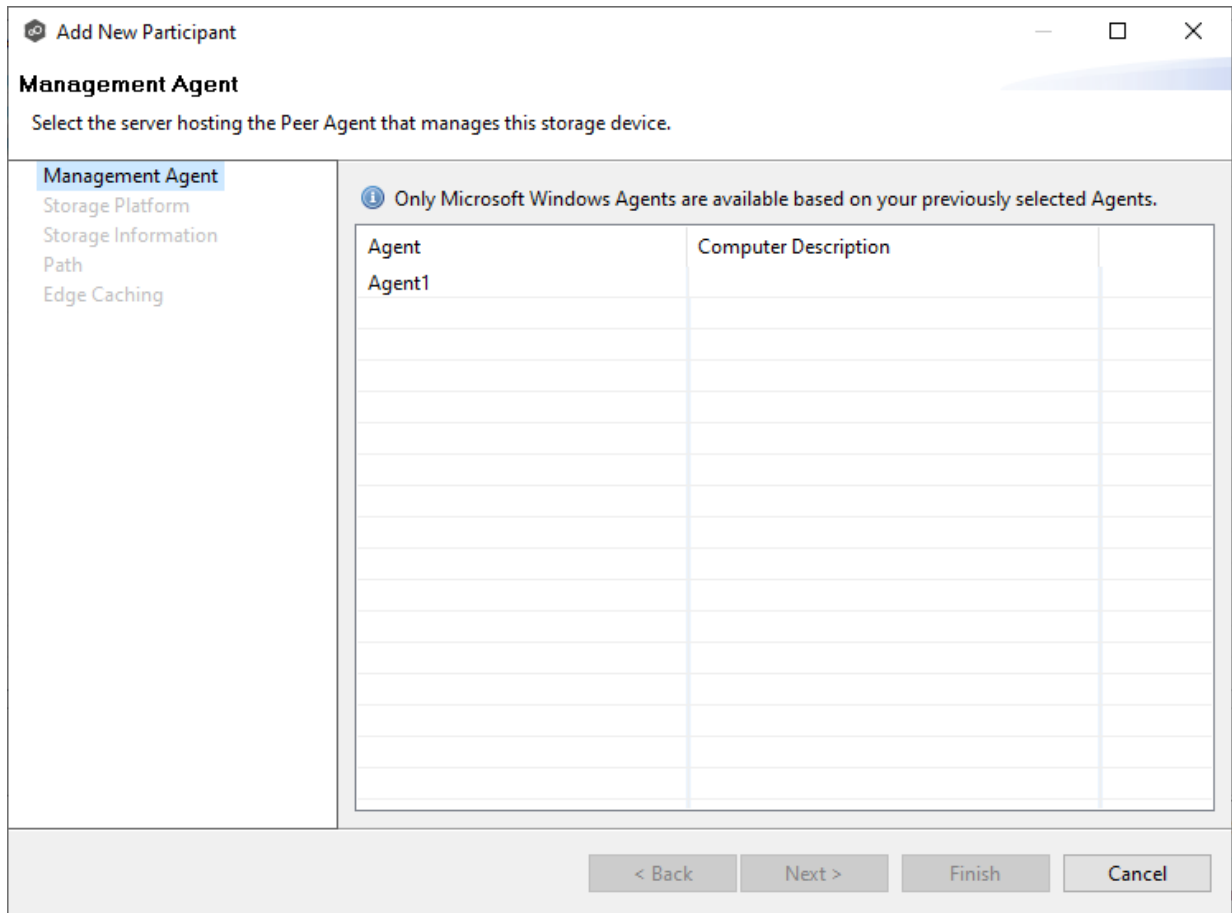
You can edit multiple File Synchronization jobs simultaneously. For information about simultaneously editing multiple jobs, see [Editing Multiple Jobs](#).

Editing a Job

To edit a File Synchronization job:

1. Select the job in the **Jobs** view.
2. Right-click and select **Edit Job**.

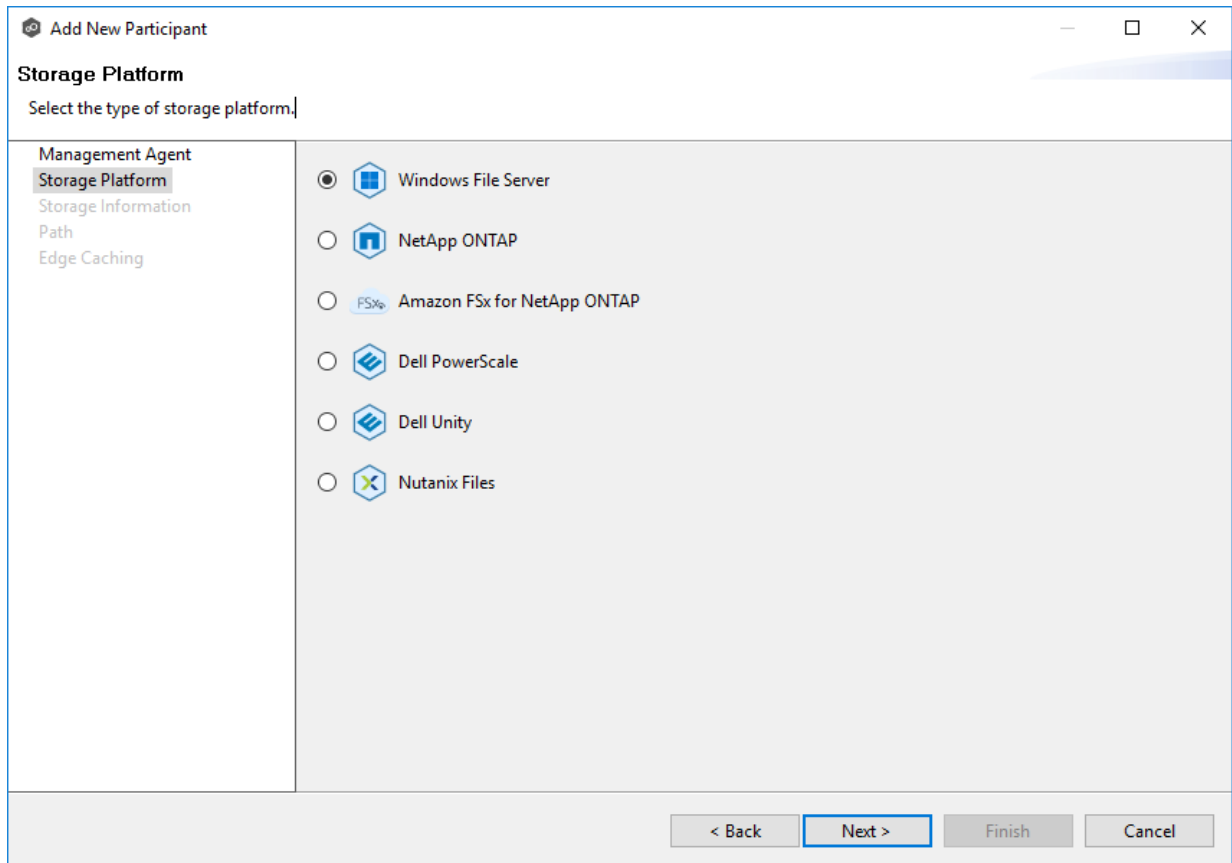
The **Edit File Synchronization Job** dialog appears.



3. Select a Management Agent, and then click **Next**.

The **Storage Platform** page appears.

4. Select the type of storage platform that hosts the data you want to synchronize, and then click **Next**.



The **Storage Information** page appears; the choices available depend on your selection on the **Storage Platform** page.

5. Enter the requested information for your platform:

[Windows File Server](#)

[NetApp ONTAP](#)

[Amazon FSxN](#)

[Dell PowerScale](#)

[Dell Unity](#)

[Nutanix Files](#)

6. Click **Next**.

The **Path** page appears.

The screenshot shows a dialog box titled "Add New Participant" with a close button (X) in the top right corner. Below the title bar, the section is labeled "Path". A red error icon and the message "You must enter a valid path." are displayed. On the left, a sidebar lists options: "Management Agent", "Storage Platform", "Storage Information", "Path" (which is selected and highlighted), and "Edge Caching". The main area contains an "Enter Path" text input field with a "Browse" button to its right. Below the input field is a checkbox labeled "Seeding Target". At the bottom of the dialog, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

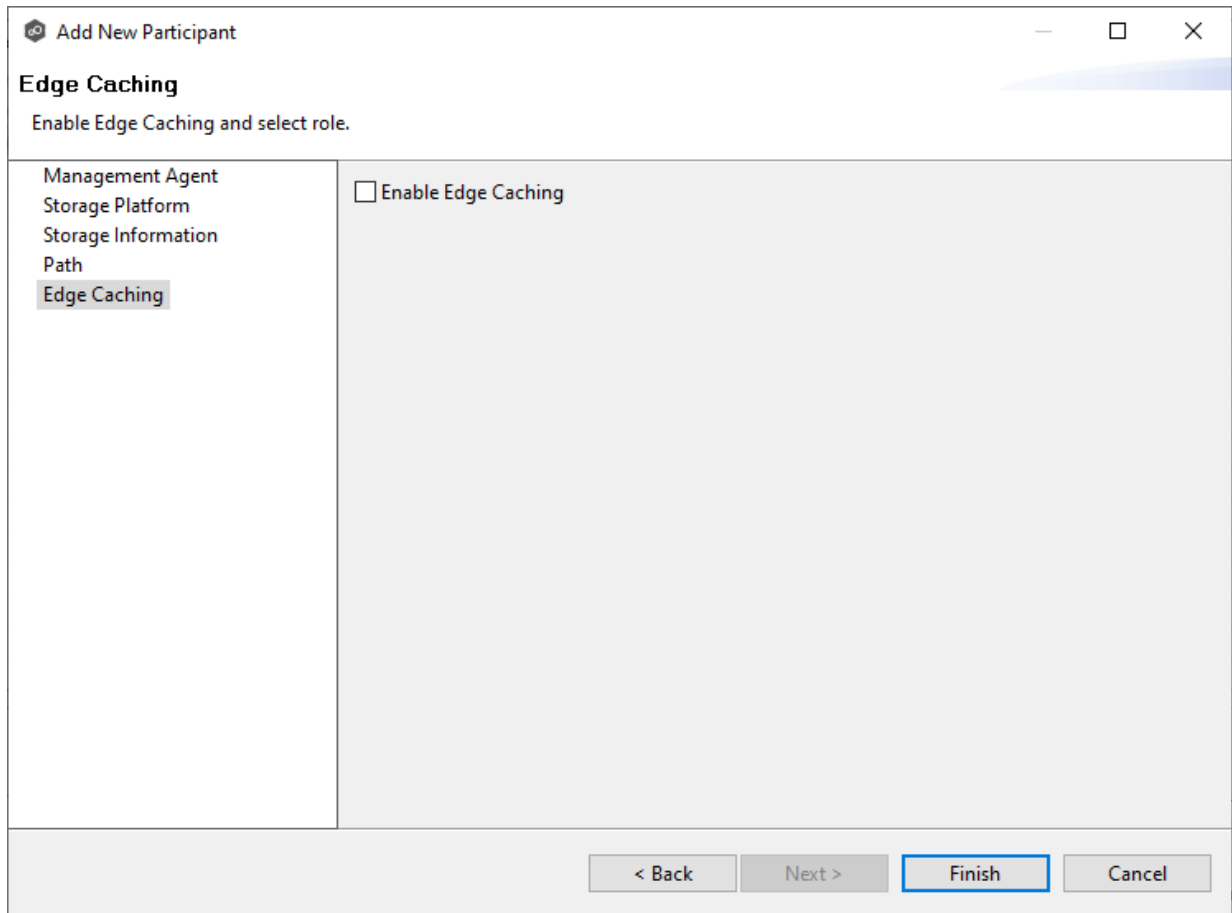
7. Browse to or enter the path to the [watch set](#):

- If the destination storage device is a Windows file server, this path should be a local path such as D:\Data, or it can be the UNC path to any SMB-capable file server.
- If the destination storage device is a Linux-based device, the path should be in NFS format. It should point to an NFS export or a subfolder under the export (e.g. server:/export-name/subfolder).

8. (Optional) Select the **Seeding Target** checkbox, and then click **OK** in the dialog that appears.

9. Click **Next**.

The **Edge Caching** page appears.

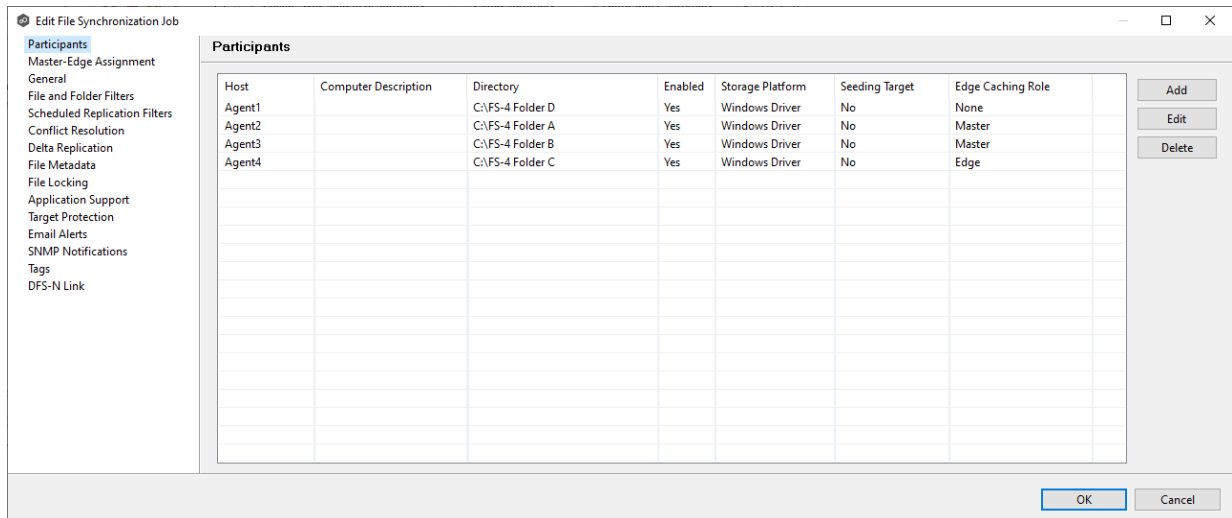


The screenshot shows a dialog box titled "Add New Participant" with a close button in the top right corner. The main heading is "Edge Caching" with the instruction "Enable Edge Caching and select role." Below this, there is a list of roles on the left: "Management Agent", "Storage Platform", "Storage Information", "Path", and "Edge Caching" (which is highlighted). To the right of this list is a checkbox labeled "Enable Edge Caching". At the bottom of the dialog, there are four buttons: "< Back", "Next >", "Finish" (which is highlighted with a blue border), and "Cancel".

10. (Optional) Select the **Enable Edge Caching** checkbox if you want this participant to be able to use Edge Caching; otherwise, click **Finish**.
11. If you enabled Edge Caching, follow the steps outlined in [Step 2: Edge Caching in Creating a File Synchronization Job](#).

For more information about Edge Caching, see [Edge Caching](#) in [Advanced Topics](#).
12. Click **Finish**.

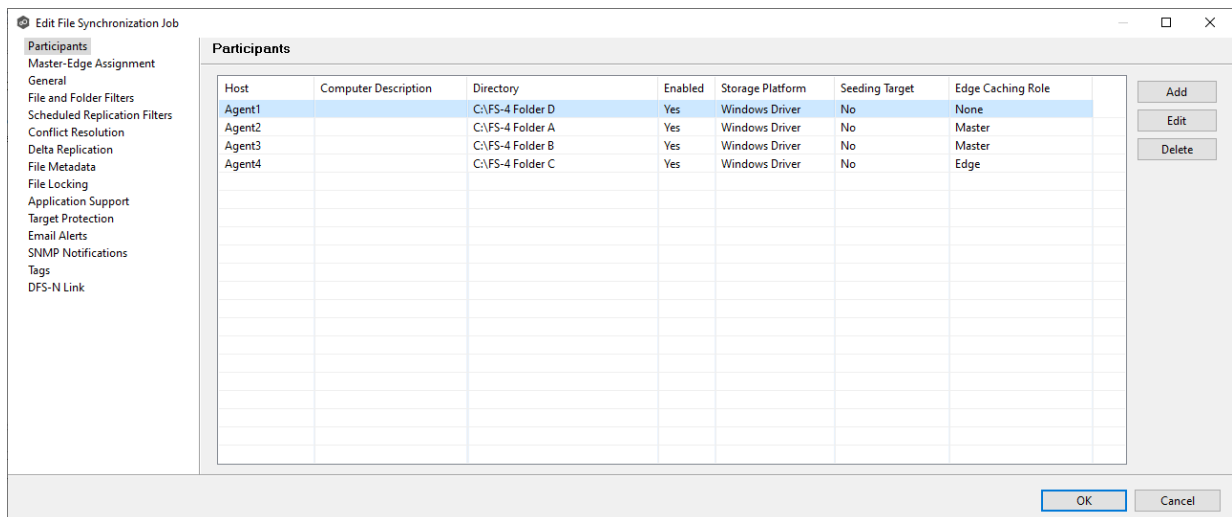
The new participant appears in the **Participants** table.



Deleting a Participant from a File Synchronization Job

To delete a participant from a File Synchronization job:

1. In the **Edit File Synchronization** dialog, select the participant in the **Participants** table you want to remove from the job.



2. Click the **Delete** button.
3. Click **OK** in the **Delete Confirmation** dialog.

The participant is removed from the **Participants** table.


Note: A File Synchronization job requires at least two participants. If you delete a participant and only one remains, you must add another participant to the job.

The screenshot shows a window titled "Edit Participant" with a "General" tab selected. The window contains the following elements:

- General** (Section Header)
- Modify these settings.
- General** (Sub-section Header)
- Edge Caching (Sub-section Header)
- Enabled:
- Host:
- Event Detector: (with a dropdown arrow)
- [Edit Detector Configuration](#) (Link)
- Directory: (with a "Browse" button)
- Seeding Target
- Navigation buttons: < Back, Next >, **Finish**, Cancel

3. To enable or disable the agent, select or deselect the **Enabled** checkbox.
4. To change the directory/folder/share that is replicated, enter the path to the new watch set in the **Directory** field or browse to it.
 - If the destination storage device is a Windows file server, this path should be a local path such as D:\Data, or it can be the UNC path to any SMB-capable file server.
 - If the destination storage device is a Linux-based device, the path should be in NFS format. It should point to an NFS export or a subfolder under the export (e.g. server:/export-name/subfolder).
5. If the storage device that the Agent is managing has changed to a different storage platform, click **Edit Detector Configuration**, and then make the necessary modifications.
6. To change whether the participant is a seeding target, select or deselect the **Seeding Target** checkbox.
7. Click **Next** to edit Edge Caching options; otherwise, click **Finish**, and skip to Step 10.

If you clicked **Next**, the Edge Caching page appears.



8. If you enabled Edge Caching, follow the steps outlined in [Step 2: Edge Caching in Creating a File Synchronization Job](#).
9. Click **OK** to close the Edit wizard or select another configuration item to modify.

Master-Edge Assignment

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

This page appears only when Edge Caching is enabled for the job.

Every edge participant must have at least one master participant assigned to it, so that Edge Caching will know which master participants to use when reading or rehydrating a stub file on

7. Click **OK**.

General

The **General** page in the **Edit File Synchronization Job** dialog presents miscellaneous settings pertaining to a File Synchronization job. You may want to consult with Peer Software's Technical Support team before modifying these values.

To modify these settings:

1. Enter the values recommended by Peer Software Support.

Edit File Synchronization Job

Participants

General

File and Folder Filters

Scheduled Replication Filters

Conflict Resolution

Delta Replication

File Metadata

File Locking

Application Support

Target Protection

Email Alerts

SNMP Notifications

Tags

DFS-N Link

General

Job ID: 179

Job Type: File Synchronization

Job Name: FS-1

Job Description:

Transfer Block Size (KB): 2048

Verify Block Checksums:

Verify Full File Checksums:

Enable Multipart Transfers:

Synchronization Priority: 2

Timeout (Seconds): 180

First Scan Mode: FOLDER_BY_FOLDER

Remove Filtered Files On Folder Delete:

Require All Hosts At Start:

Auto Start:

Auto Start Priority: 1

OK Cancel

Edit File Synchronization Job

Participants

General

File and Folder Filters

Scheduled Replication Filters

Conflict Resolution

Delta Replication

File Metadata

File Locking

Application Support

Target Protection

Email Alerts

SNMP Notifications

Tags

DFS-N Link

General

Job ID: 179

Job Type: File Synchronization

Job Name: FS-1

Job Description:

Transfer Block Size (KB): 2048

Verify Block Checksums:

Verify Full File Checksums:

Enable Multipart Transfers:

Synchronization Priority: 2

Timeout (Seconds): 180

First Scan Mode: FOLDER_BY_FOLDER

Remove Filtered Files On Folder Delete:

Require All Hosts At Start:

Auto Start:

Auto Start Priority: 1

OK Cancel

Option	Description
Job ID	Unique, system-generated job identifier that cannot be edited.

Option	Description
Job Type	Identifies the job type. This cannot be modified.
Job Name	The name of this File Synchronization job. This name must be unique.
Transfer Block Size (KB)	The block size in Kilobytes used to transfer files to hosts. Larger sizes will yield faster transfers on fast networks but will consume more memory in the Peer Management Broker and Peer Agents .
Verify Block Checksums	If selected, each block sent will be checksummed at both the source and target(s) Agents.
Verify Full File Checksums	If selected, the entire file will be checksummed <i>after</i> it has been sent from the source to target Agents. If temp files are enabled, the checksum on the target will be calculated prior to renaming the temp file to the base file's name. If the checksums between source and target(s) do not match, the file will be sent to the retry engine to reattempt the transfer.
Enable Multipart Transfers	If selected, larger files will be broken into smaller chunks and these chunks will be sent over the wire in a parallel fashion. This feature will automatically throttle itself if many other transfers are also waiting to be processed.
Synchronization 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 are NOT seeded, the FOLDER_BY_FOLDER method would be best. For environments where most data are seeded, the BULK_CHECKSUM method will result in a faster first scan.

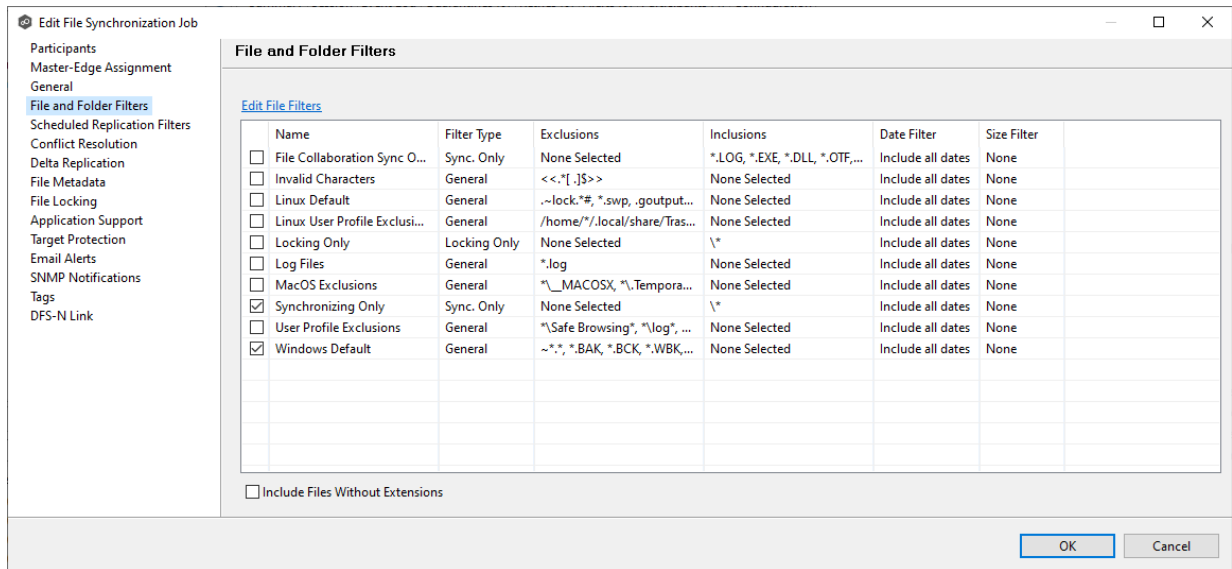
Option	Description
Remove Filtered Files On Folder Delete	If selected, then all child files on target hosts will be deleted when its parent folder is deleted on another source host. Otherwise, filtered files will be left intact on targets when a parent folder is deleted on another source host.
Require All Hosts At Start	If selected, requires all participating hosts to be online and available at the start of the File Synchronization job in order for the job to successfully start.
Auto Start	If selected, then this file synchronization session will automatically be started when the Peer Management Center Service is started.

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

File and Folder Filters

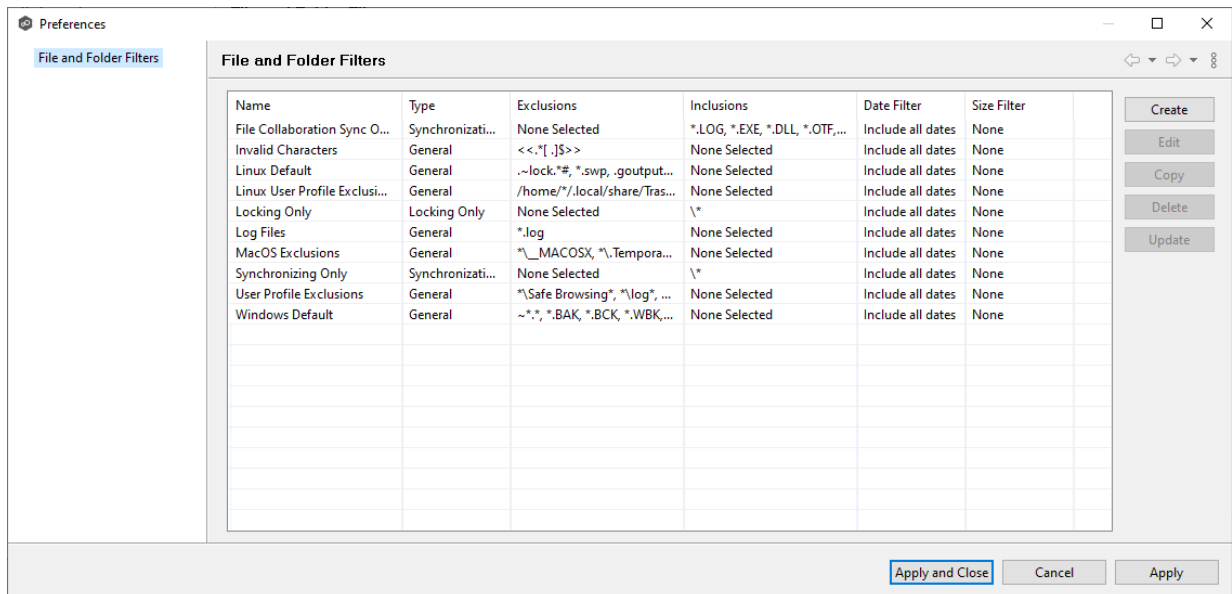
The **File and Folder Filters** page in the **Edit File Synchronization Job** dialog displays a list of [file and folder filters](#). A file or folder filter enables you to exclude and/or include files and folders from the job based on file type, extension, name, or directory path. Any file or folder that matches the filter is excluded or included from replication, depending on the filter's definition. By default, all files and folders selected in the **Source Paths** page will be replicated.

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



- If you want to create a new file or folder filter, modify an existing one, or update a filter, click **Edit File Filters**.

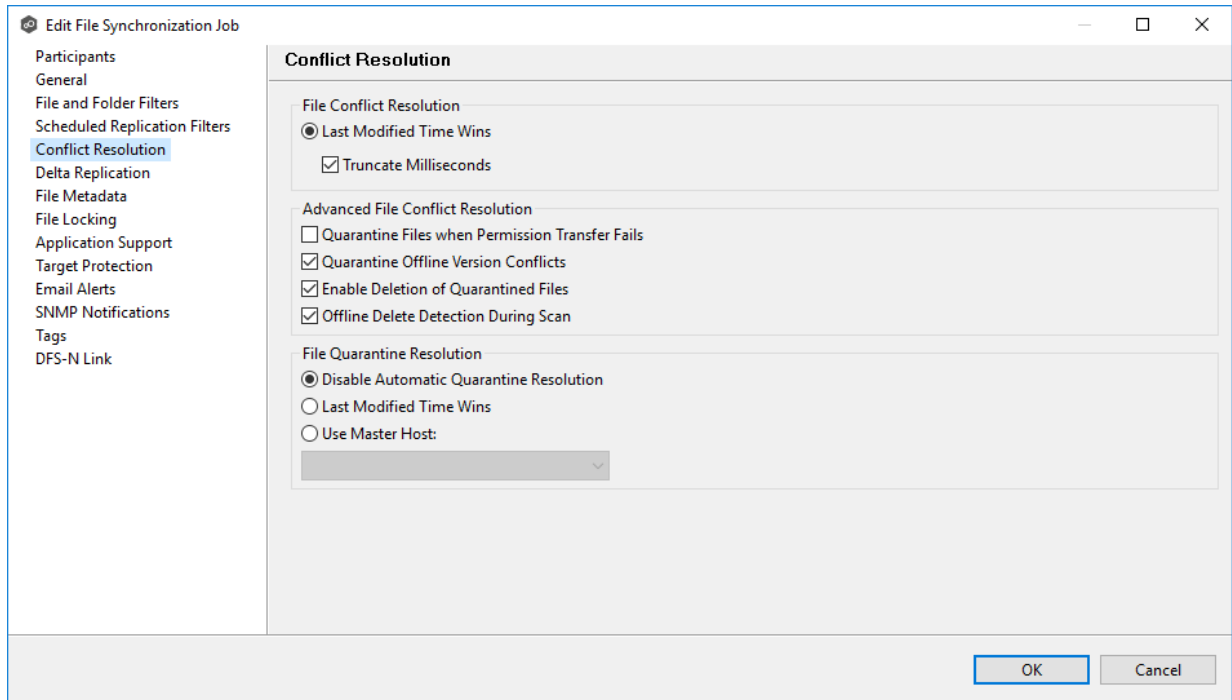
The **File and Folder Filters** dialog appears. You cannot edit predefined filters. See [File Filters](#) in the [Preferences](#) section for information about creating or modifying a file filter.



- Select the **Include Files Without Extensions** checkbox if you want to replicate files that do not have extensions.

Note: Files without extensions are ignored during replication unless you select this checkbox.

As some storage platforms and applications track milliseconds slightly differently, selecting this option will prevent subtle millisecond differences from causing an otherwise in-sync file to be replicated or quarantined.



2. Select the **Advanced File Conflict Resolution** options you want applied:

Option	Description
Quarantine Offline Version Conflicts	Select this option if you want Peer Management Center to quarantine a file that was updated in two or more locations while the collaboration session was not running. If it is not selected, the file with the most recent last modified time will be replicated to all other participants.
Enable Deletion of Quarantined Files	Select this option if you want Peer Management Center to process a delete event for a quarantined file. If it is not selected, the quarantined file is not deleted and remains quarantined.
Offline Delete Detection During Scan	Select this option (and enabled target protection), if you want any files or folders that were deleted while the job was stopped to be deleted from all participants when the job is restarted. If it is not selected, then the deleted file or folder is restored from a participant with the file or folder to any participant where it no longer exists.

3. Select an option for automatically resolving quarantines (this option is intended to be used in environments where a single file server is active for a job):

Option	Description
Disable Automatic Resolution of Quarantines	Select this option if you want to manually resolve quarantines. For more information, see Removing a File from Quarantine .
Last Modified Time	Select this option if you want quarantines automatically resolved by selecting the file with the latest modification time.
Use Master Host	Select this option if you want quarantines automatically resolved by selecting the file on the Master Host.

4. Click **OK** to close the Edit wizard or select another configuration item to modify.

Delta Replication

The **Delta Replication** page in the **Edit File Synchronization Job** dialog allows you to specify the delta-replication options to use for the selected File Synchronization job. Delta-level replication is a byte replication technology that enables block/byte level synchronization for a File Synchronization job. Through this feature, Peer Management Center is able to transmit only the bytes/blocks of a file that have changed instead of transferring the entire file. This results in much lower network bandwidth utilization, which can be an enormous benefit if you are transferring files across a slow WAN or VPN, as well as across a high-volume LAN.

Delta-level replication is enabled on a per File Synchronization job basis and generally affects all files in the [watch set](#). You will only benefit from delta-level replication for files that do not change much between file modifications, which includes most document editing programs.

To modify delta-level replication options:

1. Modify the following the fields as necessary.

Edit File Synchronization Job

- Participants
- Master-Edge Assignment
- General
- File and Folder Filters
- Scheduled Replication Filters
- Conflict Resolution
- Delta Replication**
- File Metadata
- File Locking
- Application Support
- Target Protection
- Email Alerts
- SNMP Notifications
- Tags
- DFS-N Link

Delta Replication

Enable Delta-level Replication:

Checksum Transfer Size (KB): 256

Delta Block Transfer Size (KB): 1024

Minimum File Size (KB): 2048

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

Add Delete

Excluded File Name Patterns

Add Delete

OK Cancel

Field	Description
Enable Delta-Level Replication	Select to enable delta encoded file transfers which only sends the file blocks that are different between source and target(s). If this is disabled, the standard file copy method will be used to synchronize files.
Checksum Transfer Size (KB)	Enter the block in kilobytes used to transfer checksums from target to source at one time. Larger sizes will result in faster checksum transfer, but will consume more memory on the Peer Agents
Delta Block Transfer Size (KB)	Enter the block size in kilobytes used to transfer delta encoded data from source to target at one time. Larger sizes will result in faster overall file transfers but will consume more memory on the Peer Agents.

Field	Description
Minimum File Size (KB)	Enter the minimum size of files in kilobytes to perform delta encoding for. If a file is less than this size, then delta encoding will not be performed.
Minimum File Size Percentage Target/Source	Enter the minimum allowed file size difference between source and target, as a percentage, to perform delta encoding. If the target file size is less than this percentage of the source file size, then delta encoding will not be performed.
Excluded File Extensions	Enter a comma-separated list of file extension patterns to be excluded from delta encoding, e.g., zip, jpg, png. In general, compressed files should be excluded from delta encoding and the most popular compressed file formats are excluded by default.
Excluded File Name Patterns	Enter a list of file name patterns to be excluded from delta encoding. If a file name matches any pattern in this list, then it will be excluded from delta encoding transfers and a regular file transfer will be performed. See File and Folder Filters for more information on specifying wildcard expressions.

2. Click **OK** to close the Edit wizard or select another configuration item to modify.

File Metadata

The **File Metadata** page allows you to specify whether you wish to replicate file security metadata and the types of metadata for synchronization. Additionally, it provides the ability to designate the metadata source (volume, share, export, or folder) to resolve conflicts during the initial synchronization. This designated source, utilized in case of conflicts, is referred to as the [master host](#). This page also provides some additional options not available when creating the job. See [File Metadata Synchronization](#) in [Advanced Topics](#) for more information about file metadata synchronization.

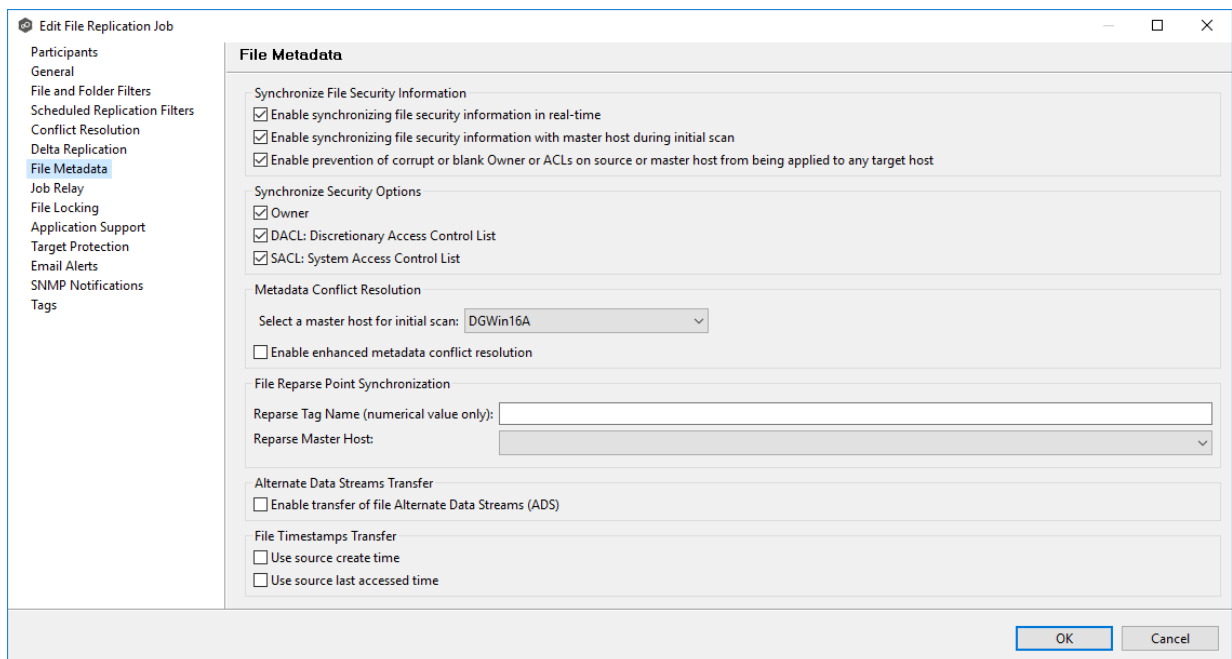
The contents of the File Metadata page vary depending on whether you are using Windows-based or Linux-based Agents for your job:

- If you selected Windows-based, proceed with [SMB File Metadata](#).
- If you selected Linux-based, proceed with [NFS File Metadata](#).

To modify file metadata synchronization settings:

1. **Select when to replicate metadata** (one or both options):

- **Real-time synchronization:** Metadata changes (Owner, DACL, SACL) will be synchronized with the target host as they occur.
- **Initial scan synchronization:** Metadata will be synchronized during the initial scan.



2. **Confirm metadata synchronization.**

If you selected either option, click **OK** in the confirmation message.

3. **Prevent corrupted or blank metadata from propagating (optional):**

- Select **Enable prevention of corrupt or blank Owner or ACLs** to prevent invalid ownership or ACL entries on the source or master host from being applied to any target host. The replication process will skip or correct these entries to maintain data integrity and security.

4. **Select security descriptor components to be synchronized:**

- Choose which components (Owner, DACL, SACL) to synchronize.

Note: To synchronize Owner or SACLs, the user that a Peer Agent service is run under on each participating host must have permission to read and write Owner and SACLs.

5. **Select the master host (if applicable):**

- If synchronizing metadata during the initial scan, choose a master host to resolve conflicts.
- If no master host is selected, metadata synchronization will not occur during the initial scan.
- If security descriptor components differ across participants, conflict resolution will use the master host's permissions.
- If the file does not exist on the master host, a participant will be chosen randomly.

6. Select **Enable enhanced metadata conflict resolution** if you want to prevent the Peer Agent service account from being assigned as the owner of that file or folder when a metadata conflict occurs, and a file or folder is written to a target. If the Peer Agent service account were to be assigned as the owner, the user may not have permission to access the file or folder.

Note: The Peer Agent service account cannot be a local or system administrator. As described in [Peer Global File Service - Environmental Requirements](#), the Peer agent service account should be an actual user.

7. (Optional) Enter values for one or both file reparse point data synchronization options:

- **Reparse Tag Name** - Enter a single numerical value. Must be either blank (if blank, reparse synchronization will be disabled) or greater than or equal to 0. The default for Symantec Enterprise Vault is 16. A value of 0 enables reparse point synchronization for all reparse file types. If you are unsure as to what value to use, then contact Peer Software Support, or you can use a value of 0 if you are sure that you are only utilizing one vendor's reparse point functionality.
- **Reparse Master Host** - Select a master host. If a master host is selected, then when the last modified times and file sizes match on all hosts, but the file reparse attribute differs (e.g., archived/offline versus unarchived on file server), then the file reparse data will be synchronized to match the file located on the master host. For Enterprise Vault, this should be the server where you run the archiving task on. If the value is left blank, then no reparse data synchronization will be performed, and the files will be left in their current state.

Notes:

- Use the **File Reparse Point Data Synchronization** options 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.

- The **File Reparse Point Data Synchronization** options are not available if the job is using Edge Caching.
8. (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.

Note: The **Alternate Data Stream Transfer** option is not available if the job is using Edge Caching.

9. In the **File Timestamps Transfer** section, select one or both options:
 - **Use source create time:** Enable this option if you want to preserve the create time (or creation timestamp) of the source file during the transfer process. Enabling this option ensures that the destination file retains the same create time as the source file after the transfer is complete.
 - **Use source last accessed time:** Enable this option if you want to preserve the last accessed time (or access timestamp) of the source file during the transfer process. Enabling this option ensures that the destination file inherits the same last accessed time as the source file after the transfer is complete.

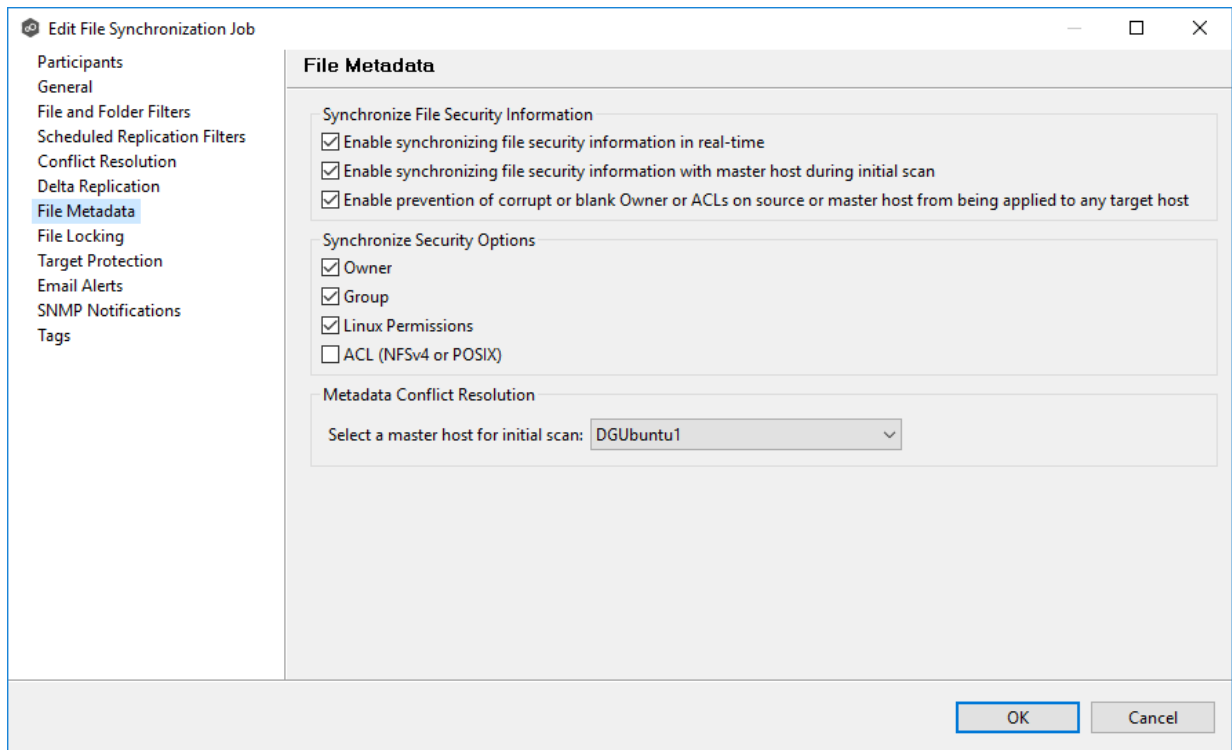
These timestamps are only applied during transfer operations. If the timestamps on the source change but the contents of the file remain the same, they will not be applied to the destination.

10. Click **OK** to close the Edit wizard or select another configuration item to modify.

To modify file metadata synchronization settings:

1. **Select when to replicate metadata** (first or second options):
 - **Real-time synchronization:** Metadata changes (Owner, Group, Linux Permissions, ACL) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.



2. **Confirm metadata synchronization.**

If you selected either of the first option, click **OK** in the confirmation message.

3. **Prevent corrupted or blank metadata from propagating (optional):**

- Select **Enable prevention of corrupt or blank Owner or ACLs** to prevent invalid ownership or ACL entries on the source or master host from being applied to any target host. The replication process will skip or correct these entries to maintain data integrity and security.

4. **Select access controls to be synchronized:**

- Choose which components (Owner, Group, Linux Permissions, ACL) to synchronize.

Note: To synchronize Owner or ACLs, the user that a [Peer Agent](#) service is run under on each participating host must have permission to read and write Owner and ACLs.

5. **Select the master host (if applicable):**

- If synchronizing metadata during the initial scan, choose a [master host](#) to resolve conflicts.

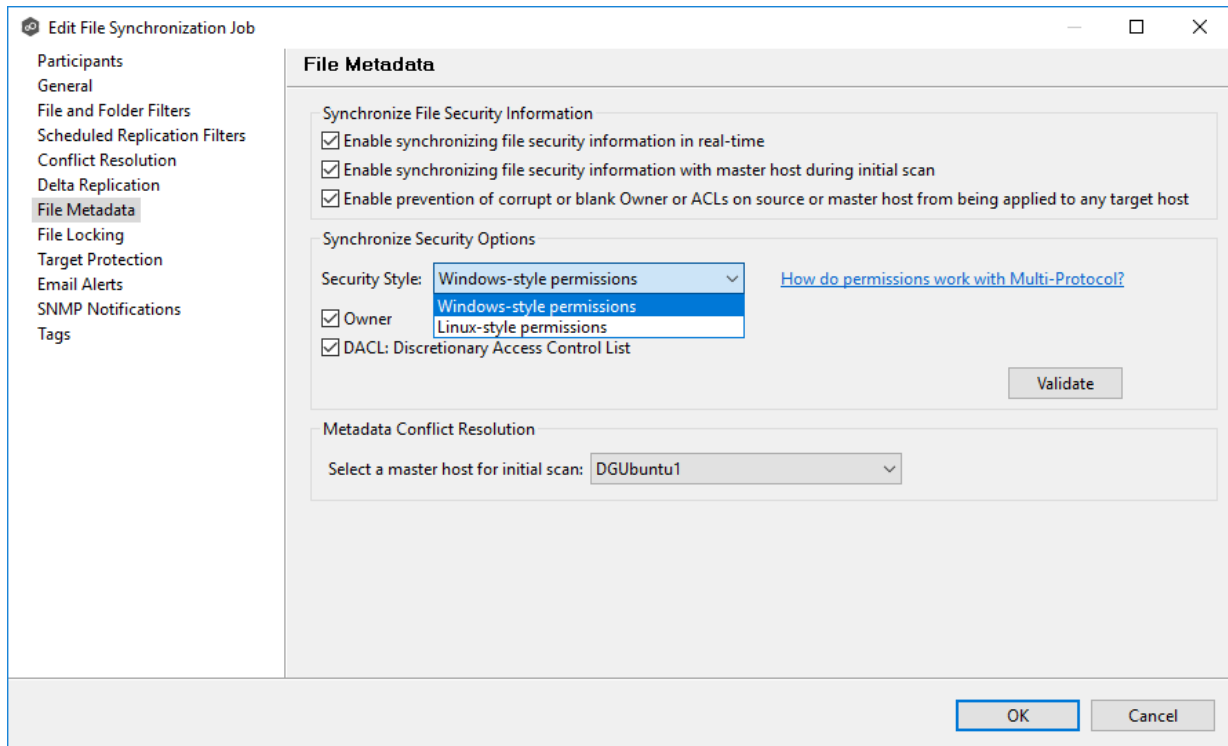
This option is only available when both of the first two options in the **Synchronize File Security Information** section are enabled, and **Owner** is selected under **Synchronize Security and ACL Options**.

- If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If access controls differ across participants, [conflict resolution](#) will use the master host's permissions.
 - If the file does not exist on the master host, a participant will be chosen randomly.
6. Click **OK** to close the Edit wizard or select another configuration item to modify.

To modify file metadata synchronization settings:

1. **Select when to replicate metadata** (one or both options):
 - **Real-time synchronization:** Metadata changes are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.



2. Confirm metadata synchronization.

If you selected either of the first two options, click **OK** in the confirmation message.

3. Prevent corrupted or blank metadata from propagating (optional):

- Select **Enable prevention of corrupt or blank Owner or ACLs** to prevent invalid ownership or ACL entries on the source or master host from being applied to any target host. The replication process will skip or correct these entries to maintain data integrity and security.

4. Select the security style:

- Windows-style permissions
- Linux-style permissions
- Mixed permissions (available only with Dell PowerScale)

5. Select permissions to be synchronized:

For a general discussion of synchronizing permissions in a multi-protocol job, see the Peer Software knowledge base article [How Do Permissions Work with Multi-Protocol?](#)

6. Select the master host (if applicable):

- If synchronizing metadata during the initial scan, choose a [master host](#) to resolve conflicts.
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If access controls differ across participants, [conflict resolution](#) will use the master host's permissions.
 - If the file does not exist on the master host, a participant will be chosen randomly.
7. Click **OK** to close the Edit wizard or select another configuration item to modify.

File Locking

The **File Locking** page in the **Edit File Synchronization Job** dialog provides options for managing how Peer Management Center temporarily locks source and target files.

To modify file locking options:

1. Modify the fields in the **Source Snapshot Synchronization** section as needed:

Edit File Synchronization Job

Participants
Master-Edge Assignment
General
File and Folder Filters
Scheduled Replication Filters
Conflict Resolution
Delta Replication
File Metadata
File Locking
Application Support
Target Protection
Email Alerts
SNMP Notifications
Tags
DFS-N Link

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

Use Storage Snapshots:

Sync. On Save

Enable Sync. On Save:

Included File Extensions: xls,xlsx,doc,docx,dwg,odt,ods,odp,odg

Synchronization Delay (Seconds): 20

OK Cancel

Field	Description
Enable Source Snapshot Sync.	If enabled, a snapshot copy of the source file is created for files that meet the snapshot configuration criteria below, and this copy is used for synchronization purposes. In addition, no file handle is held on the source file except while making a copy of the file.
Snapshot Copy Max File Size (MB)	The maximum file size for which source snapshot synchronization is utilized.
Snapshot Copy File Extensions	A comma-separated list of file extensions for which source snapshot synchronization is utilized.
Use Storage Snapshots	If enabled, a storage volume snapshot is created and used for synchronization purposes. As a result, no file handle is held on the source file. The snapshot is created using either VSS or storage-platform specific snapshot technologies. This option is in

Field	Description
	addition to the Enable Source Snapshot Sync. option above and will only apply to files with pst, mdf, ldf, and ndf extensions.

2. Modify the fields in the **Sync. on Save** section as needed.

Field	Description
Enable Sync. On Save	If enabled, this feature allows supported file types to be synchronized after a user saves a file, rather than waiting for the file to close.
Included File Extensions	A comma-separated list of file extensions for which to enable the Sync. On Save feature.
Synchronization Delay (Seconds)	The number of seconds to wait after a file has been saved before initiating a synchronization of the file.

3. Click **OK**.

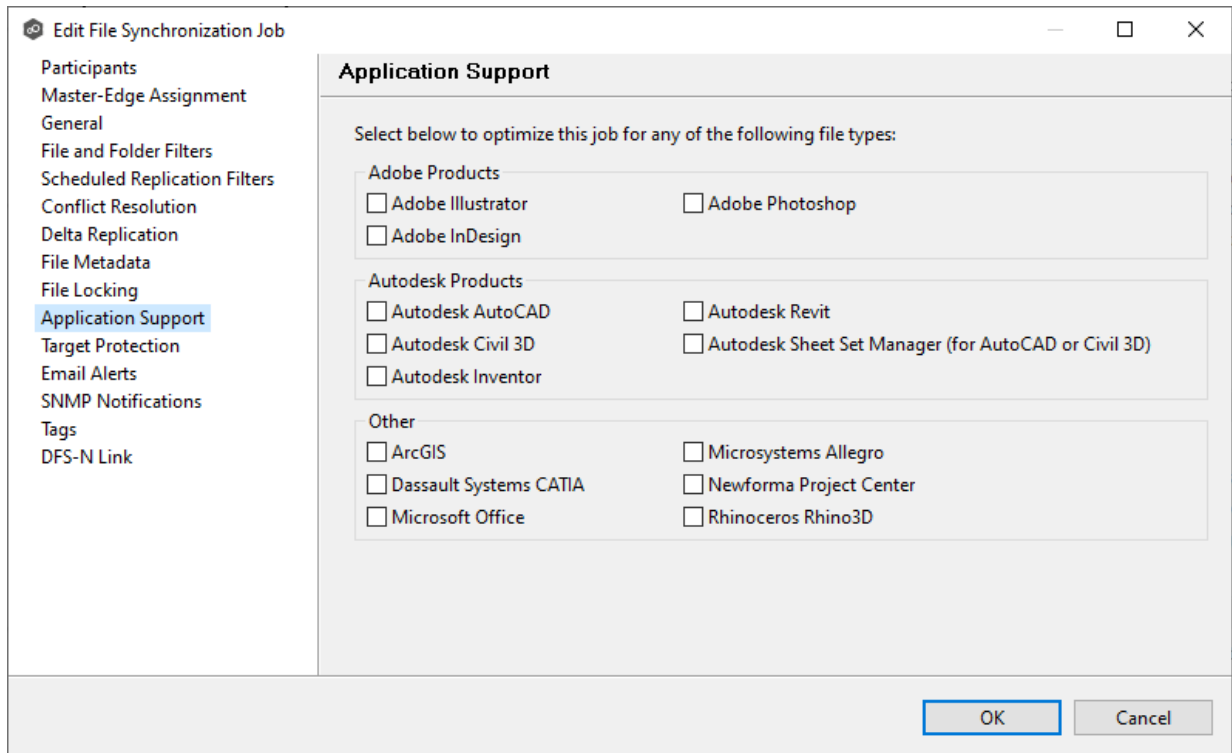
Application Support

Application Support enables automatic optimization of a file synchronization job for files created by certain 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 [Peer Support](#).

To modify which applications are optimized:

1. Select the applications to be optimized.



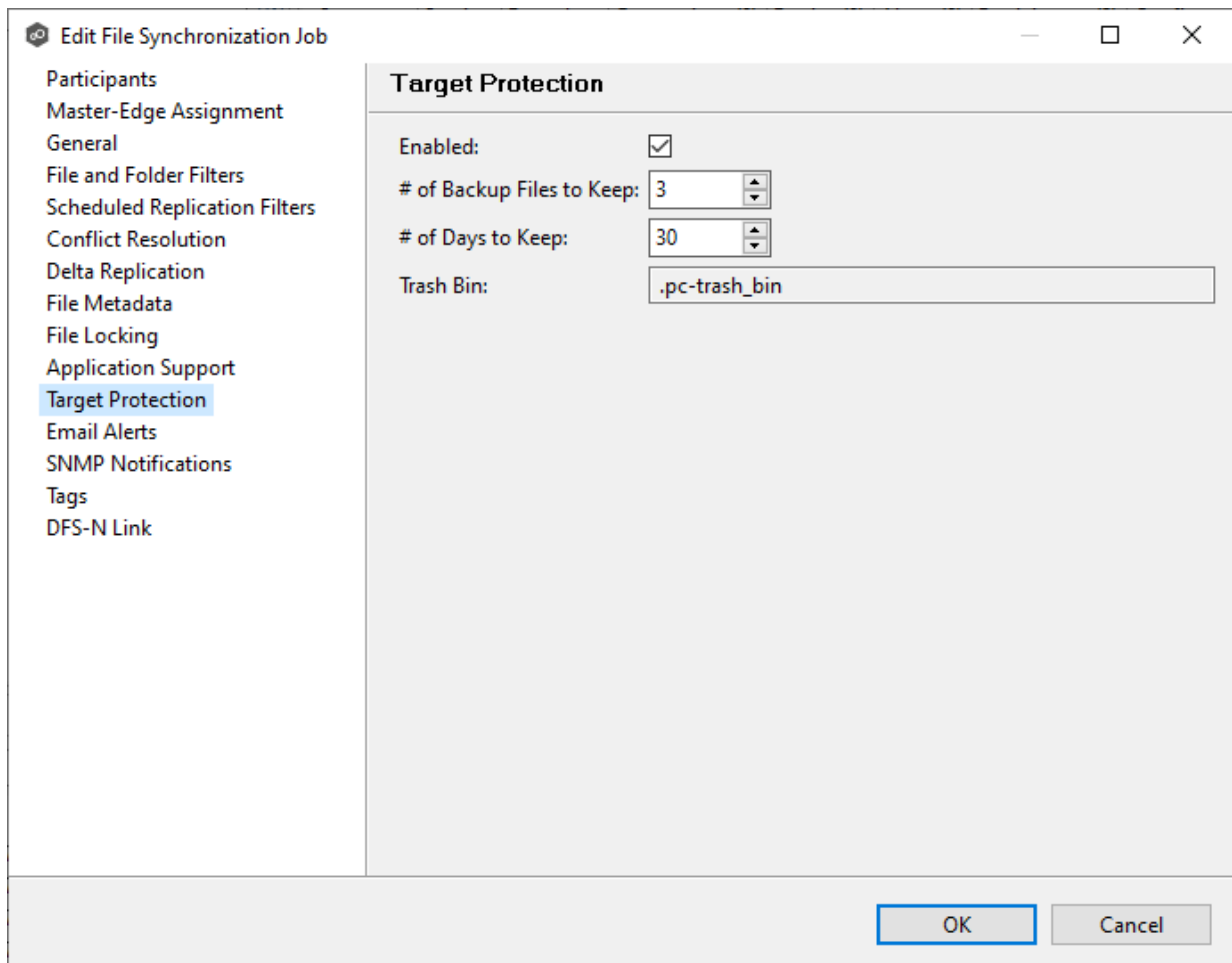
2. Click **OK** to close the Edit wizard or select another configuration item to modify.

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 a job has target protection enabled, then whenever a file is deleted or modified on the source host but before the changes are propagated to the targets, a copy of the existing file on the target is moved to the Peer Management Center trash bin.

The trash bin is located in a hidden folder named **.pc-trash_bin** found in the root directory of the [watch set](#) of the target host. A backup file is placed in the same directory hierarchy location as the source folder in the watch set within the recycle bin folder. If you need to restore a previous version of a file, you can copy the file from the trash bin into the corresponding location in the watch set and the changes will be propagated to all other collaboration hosts.

Target protection configuration is available by selecting **Target Protection** from the tree node within the Edit File Synchronization Configuration dialog.



Modify the fields as needed:

Field	Description
Enabled	Enables target protection.
# of Backup Files to Keep	The maximum number of backup copies of an individual file to keep in the trash bin before purging the oldest copy.
# of Days to Keep	The number of days to keep a backup archive copy around before deleting from disk. A value of 0 will disable purging any files from archive.
Trash Bin	The trash bin folder name located in the root directory of the watch set. This is a hidden folder and the name cannot be changed by the

3. Repeat to add additional alerts to the job.
4. Click **OK** to close the Edit wizard or select another configuration item to modify.

SNMP Notifications

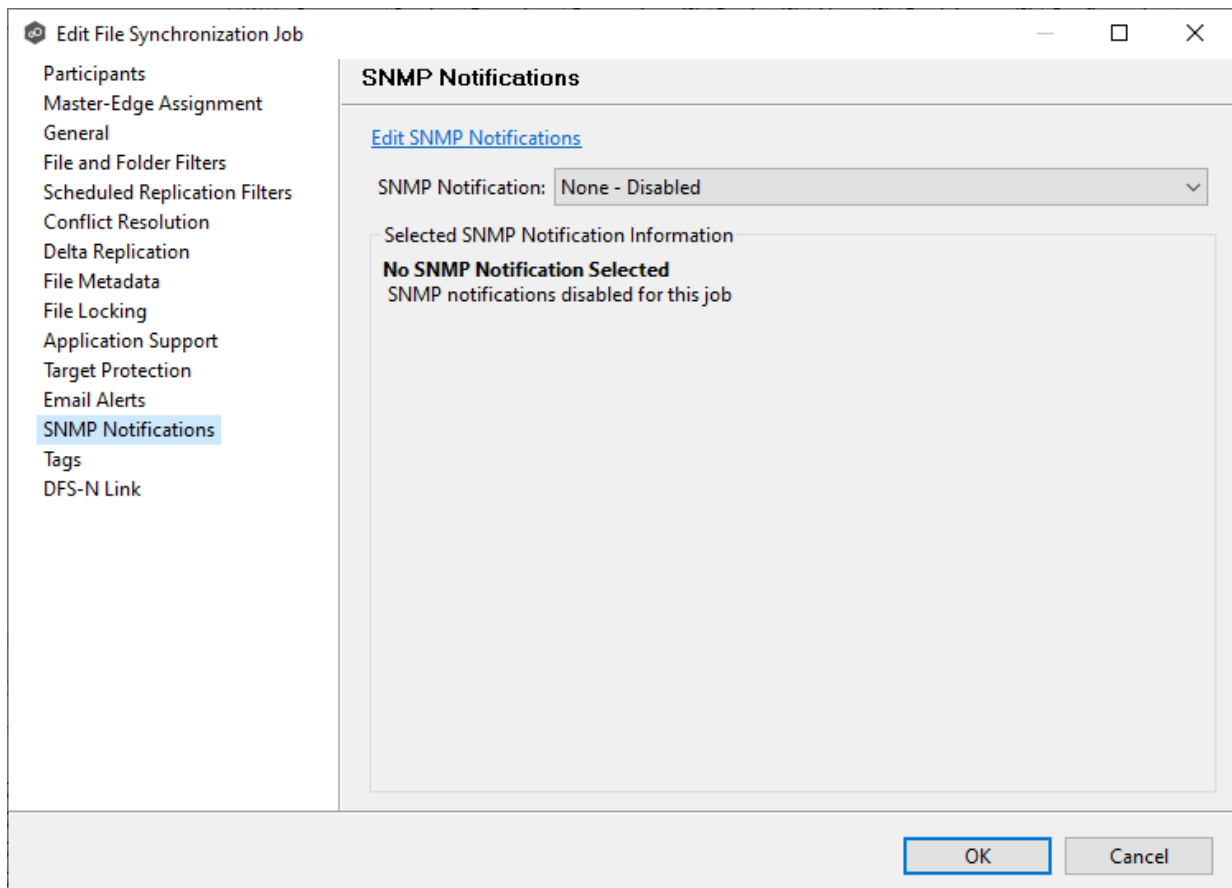
The **SNMP Notifications** page in the **Edit File Synchronization Job** dialog allows you to select which SNMP notification to apply to a File Synchronization job.

SNMP notifications, like email alerts and file filters, are configured at a global level in the [Preferences](#) dialog, then applied to individual jobs. For more information about SNMP Notifications, see [SNMP Notifications](#) in the **Preferences** section.

To enable or disable SNMP notifications for a File Synchronization job:

1. To enable, select an SNMP notification from the drop-down list.

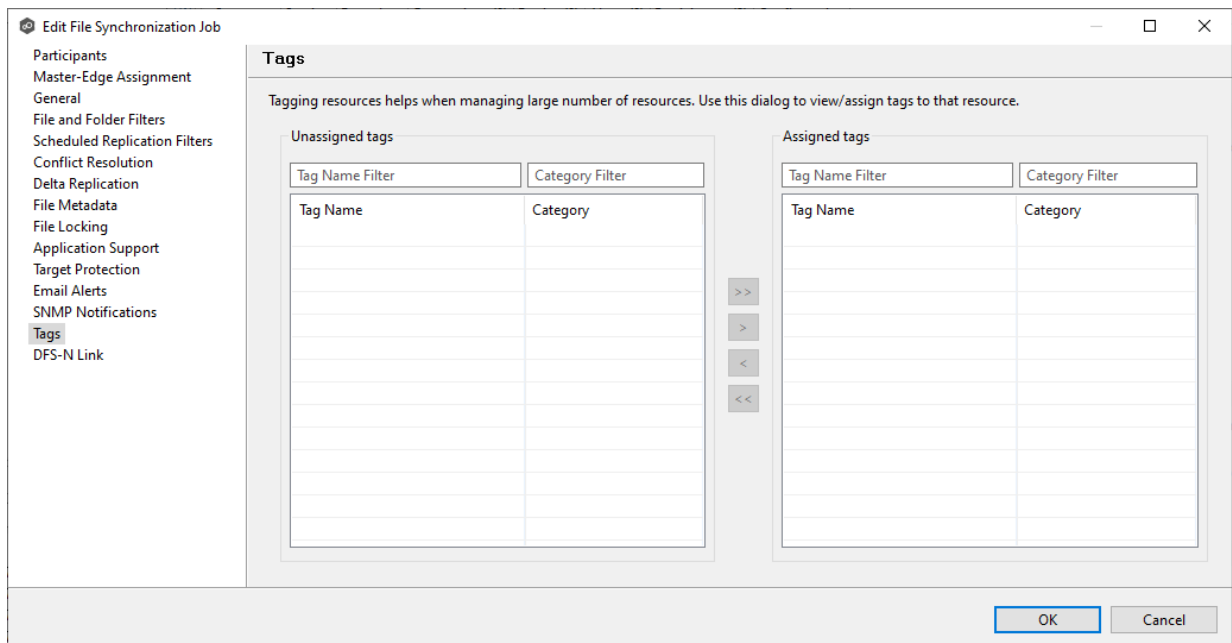
To disable, select **None - Disabled**.



2. Click **OK** to close the Edit wizard or select another configuration item to modify.

Tags

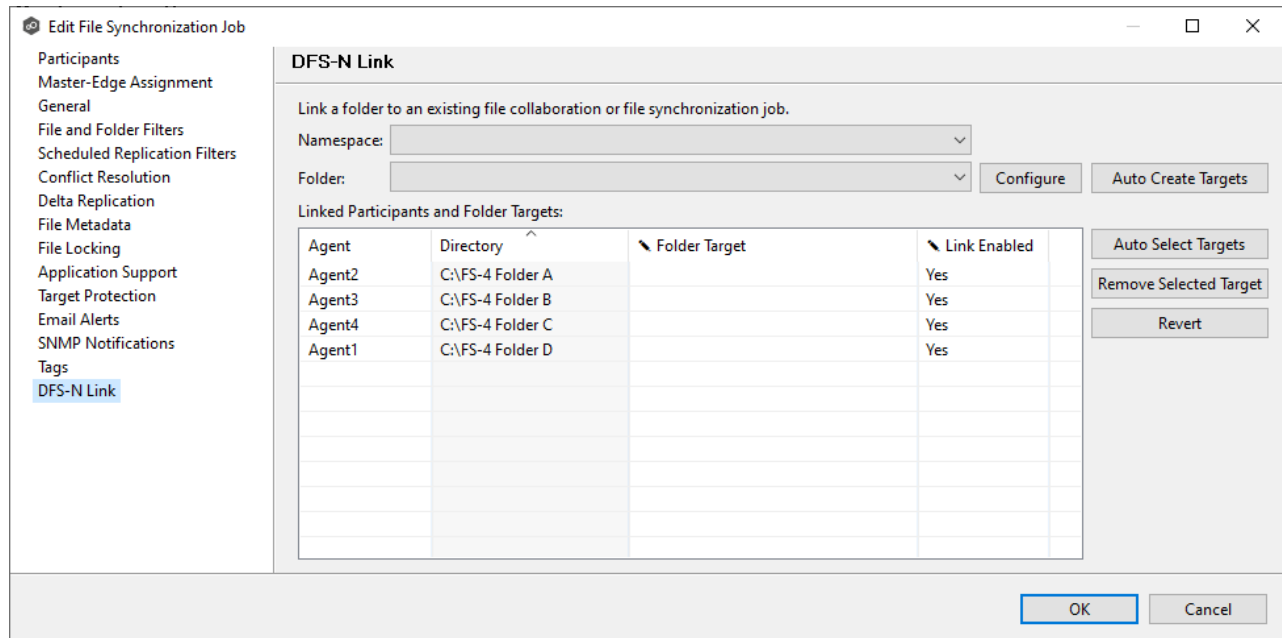
The **Tags** page in the **Edit File Synchronization Jobs** dialog allows you to assign existing tags and categories to the selected job. This page is not available in [Multi-Job Editing](#) mode. For more information about tags, see [Tags](#) in the [Basic Concepts](#) section.



DFS-N Link

Please note that this functionality currently does not support NFS or Multi-Protocol workloads.

The **DFS-N** page in the **Edit File Synchronization Job** dialog presents options for linking a DFS namespace folder to this job. See [Linking a Namespace Folder to an Existing File Collaboration or Synchronization Job](#) for more information.



Editing Multiple Jobs


Peer Management Center supports multi-job editing, allowing you to quickly and effectively manipulate multiple File Synchronization jobs simultaneously. For example, you can use this feature to change a single configuration item such as **Auto Start** for any number of already configured jobs in one operation instead of having to change the item individually for each.

While this feature does cover most of the options available on a per-job basis, certain options are unavailable in multi-job edit mode, specifically ones tied to [participants](#). Configuration of participants must be performed on a per job basis.


To edit multiple jobs simultaneously:

1. Open Peer Management Center.
2. Select the jobs you want to edit in the **Jobs** view.
3. Right-click and select **Edit Jobs**.

For the most part, the original configuration dialog remains the same with a few minor differences depending on similarities between the selected File Synchronization jobs. A sample dialog is as follows:



In this dialog, any discrepancies between multiple selected jobs will generally be illustrated by a read-only text field with the caption **Multiple Values - Click to Edit**. Clicking this field will bring up a dialog similar to the following:



This dialog gives you the option of choosing a value that is already used by one or more selected File Synchronization jobs, in addition to the ability to use your own value. Notice that variances in the look and feel of this pop-up dialog above depend on the type of information it is trying to represent (for example, text vs. a checkbox vs. a list of items).

Upon clicking OK, the read-only text field you originally clicked will be updated to reflect the new value. Any fields that have changed will be marked by a small caution sign. On saving in this multi-job edit dialog, the changed values will be applied to all selected jobs.

Note: Read all information on each configuration page carefully when using the multi-job edit dialog. A few pages operate in a slightly different manner than mentioned above. All of the necessary information is provided at the top of these pages in bold text.

Running and Managing a File Synchronization Job

The topics in this section provide some basic information about starting, stopping, and managing File Synchronization jobs:

- [Overview](#)
- [Starting a File Synchronization Job](#)
- [Starting a File Synchronization Job](#)
- [Auto-Restarting a File Synchronization Job](#)
- [Host Connectivity Issues](#)
- [Removing a File from Quarantine](#)
- [Manual Retries](#)

Overview

This topic describes:

- The [initialization process](#) for a File Synchronization job: What occurs the first time you run a File Synchronization job.
- The [initial synchronization process](#): How files are synchronized the first time you run a File Synchronization job.

The initialization process for a File Synchronization job consists of the following steps:

1. All participating hosts are contacted to make sure they are online and properly configured.
2. [Real-time event detection](#) is initialized on all participating hosts where file locks and changes will be propagated in real-time to all participating hosts. You can view real-time activity and history via the various Runtime views for the open job.
3. The [initial synchronization process](#) is started; all of the configured root folders on the participating hosts are scanned in the background, and a listing of all folders and files are sent back to the running job.
4. The background directory scan results are analyzed, and directory structures compared to see which files are missing from which hosts. In addition, file conflict resolution is performed to decide which copy to use as the master for any detected file conflicts based on the [File Conflict Resolution](#) settings.
5. After the analysis is performed, all files that need to be synchronized are copied to the pertinent host(s).

Before you start a File Synchronization job for the first time, you need to decide how you would like the [initial synchronization](#) to be performed. During the initial synchronization process:

- The watch set is recursively scanned on all participating hosts.
- File conflict resolution is performed.
- Any files that require updating are synchronized with the most current copy of the file.

The two primary options are:

- Have the File Synchronization job perform the initial synchronization based on the [Conflict Resolution](#) settings.
- Pre-seed all [participating hosts](#) with the correct folder and file hierarchy for the configured root folders before starting the session.

If you have a large data set, we strongly recommend that you perform the initial synchronization manually by copying the data from a host with the most current copy to all other participating hosts. This needs to be done only once--before the first time that you run the File Synchronization job.

If you choose the first option, click the **Start** button to begin [synchronization session initialization](#). Otherwise, pre-seed each participating host with the necessary data, then click the **Start** button.

Starting a File Synchronization Job

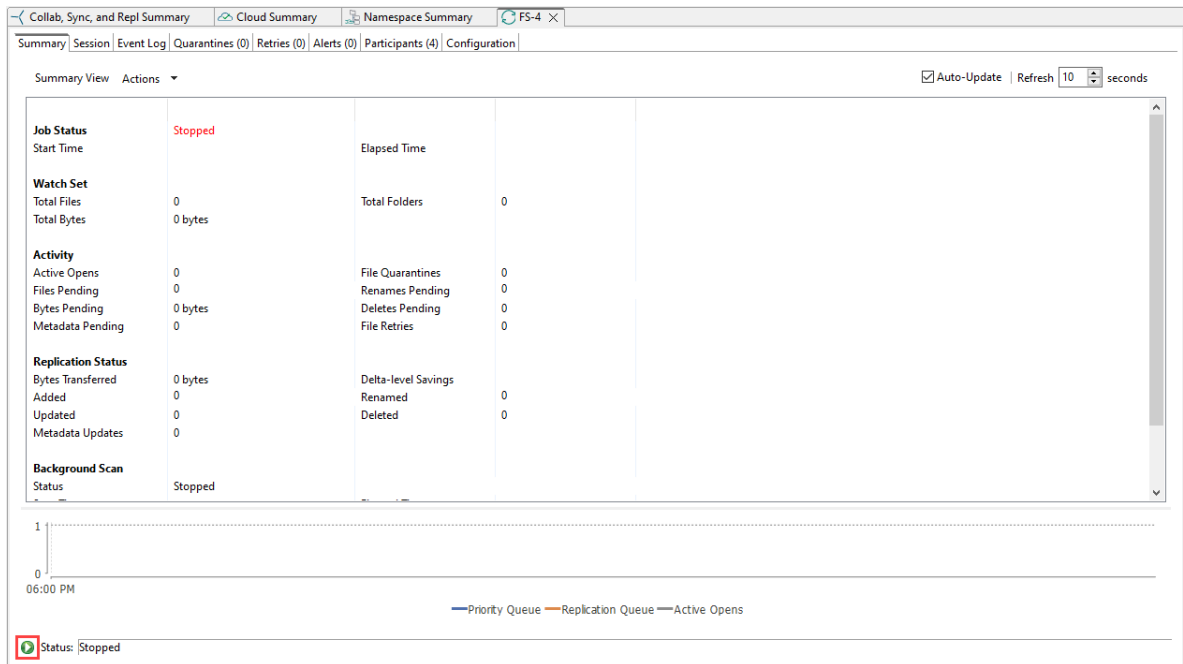
Before starting a File Synchronization job for the first time, make sure that you have decided how you want the [initial synchronization](#) to be performed.

When running a File Synchronization job for the first time, you must manually start it. After the initial run, a job will automatically start, even when the Peer Management Center server is rebooted.

Note: You cannot run two jobs concurrently on the same volume if the [watch sets](#) contain an overlapping set of files and folders.

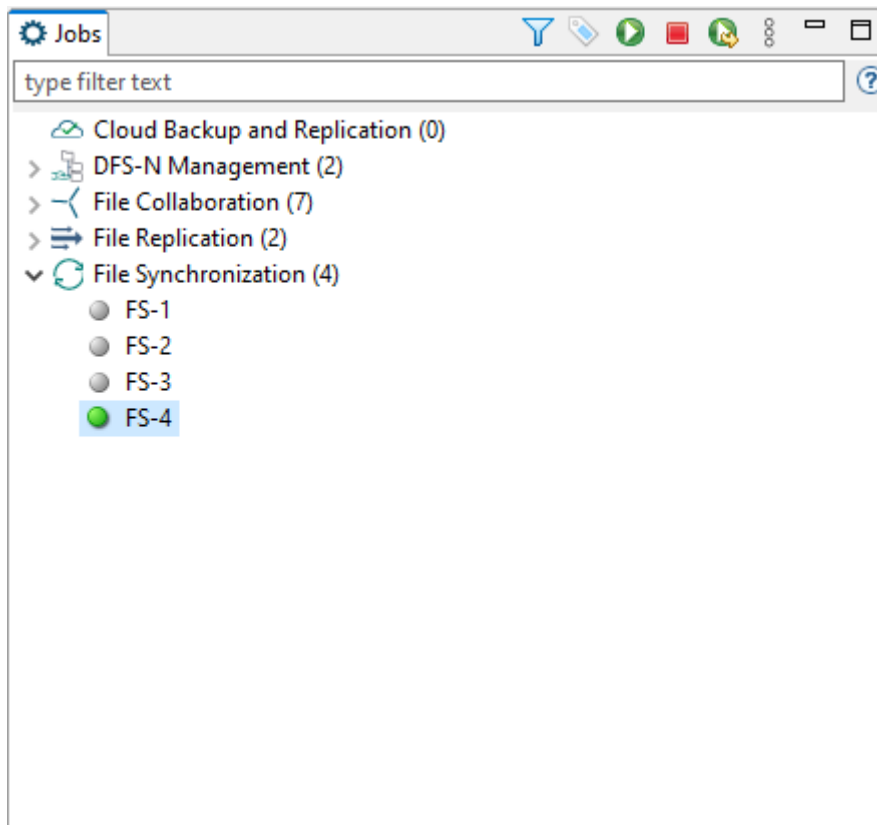
To manually start a job:

1. Choose one of three options:
 - Right-click the job name in the **Jobs** view.
 - Right-click the job name in the **Summary** tab of the **Collab, Sync, and Repl Summary** summary view, and then choose **Start** from the pop-up menu.
 - Open a job and then click the **Start/Stop** button to the left of the **Status** field in the bottom left corner of the job's runtime view (shown below).



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

After the job initialization has completed, the job will run. Once the job starts, the icon next to the job name in the **Jobs** view changes from gray to green.



Stopping a File Synchronization Job

You can stop a File Synchronization job at any time by clicking the **Stop** button on the **Jobs** view toolbar. Doing this shuts down the real-time file event detection and closes all running operations (e.g., file transfers).

Auto-Restarting a File Synchronization Job

Peer Management Center includes support for automatically restarting File Synchronization jobs that include [participating hosts](#) that have been disconnected, have reconnected, and are once again available.

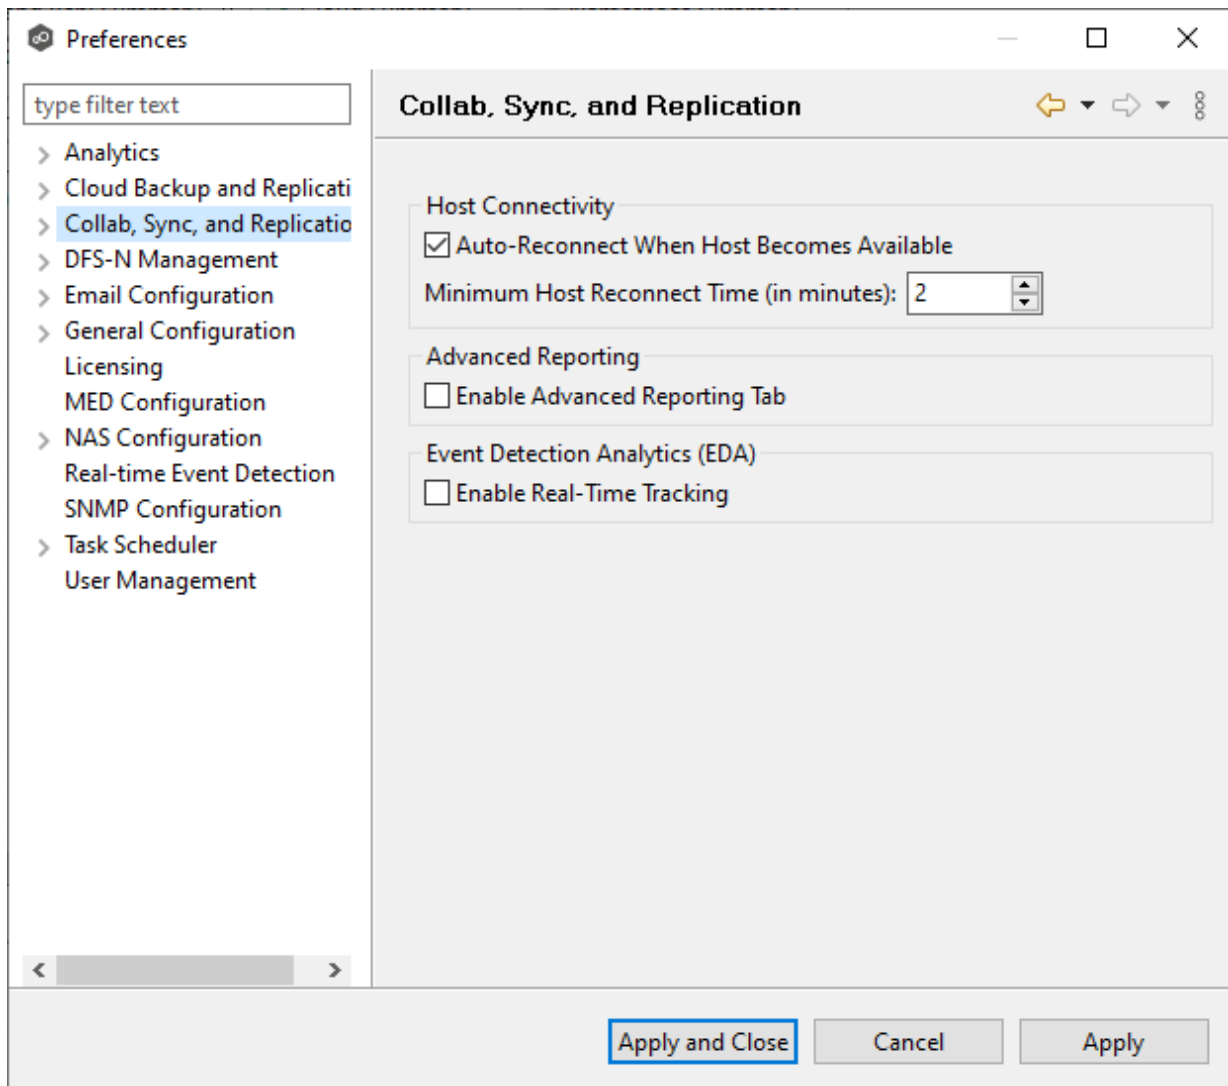
After a host becomes unavailable and the [quorum](#) is lost on a running File Synchronization job, the job automatically stops running and enters a pending state, waiting for one or more hosts to become available again so that the quorum can be met. Once the quorum is met, the pending job will automatically be restarted, beginning with a scan of all root folders.

In a job where a host becomes unavailable, but quorum is not lost, the remaining hosts will continue synchronizing. If the unavailable host becomes available once again, it is brought back into the running job and a background scan begins on all participating hosts, similar in fashion to the initial background scan at the start of a job.

You can enable all File Synchronization jobs to auto-restart. You can also disable auto-restart File Synchronization jobs on a per-job and host instance.

To enable all File Synchronization jobs to auto-restart:

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



3. Select the **Auto Reconnect when Host Becomes Available** checkbox.

4. Enter the minimum number of minutes to wait after an Agent reconnects before re-enabling it in any associated jobs in the **Minimum Host Reconnect Time** field.
5. Click **OK**.

Host Connectivity Issues

Peer Management Center is designed to be run in an environment where all [participating hosts](#) are highly available and on highly available networks. The two primary connectivity issues result from:

- [Unavailable Hosts](#)
- [Quorum Not Met](#)

Unavailable Hosts

If a host becomes unavailable while a File Synchronization job is running and is unreachable within the configured timeout period (specified in the job's [General settings](#)), it may be removed from synchronization. If no response is received while performing a file synchronization operation within the timeout period, Peer Management Center pings the host; if still no response, the host is taken out of the running session, a FATAL event is logged, and the **Participants** tab for the job is updated to indicate that the host has failed. In addition, if [email alerts](#) and/or [SNMP notifications](#) are configured and enabled for **Host Timeouts**, then the appropriate message(s) are sent.

If [auto-restart](#) not enabled, you must stop and start the File Synchronization job to bring any failed hosts back into the session. As a result, all root folders on all hosts will need to be scanned again to detect any inconsistencies. Therefore, if you are operating over a WAN with low bandwidth you will want to set the timeout to a higher value on each related job.

Quorum Not Met

For a File Synchronization job to run correctly, a quorum of available hosts must be met. Quorum is currently set to at least 2 hosts, and if quorum is not met, then the synchronization session is automatically terminated. If [email alerts](#) and/or [SNMP notifications](#) are configured and enabled for **Session Aborts**, then the appropriate message(s) are sent.

Removing a File from Quarantine

Quarantines are a key feature of Peer Global File Service, used for resolving version conflicts. For more detailed information on how quarantines work, see [Conflicts, Retries, and Quarantines](#) in the [Advanced Topics](#) section.

You must explicitly remove a file from quarantine in order to have it participate in the synchronization session once again.

You may also choose to perform no action, in which case, the file is removed from the **Quarantines** table but none of the file versions are modified. Therefore, if the files are not currently in-sync, then the next time the file is modified, changes will be propagated to the other hosts. If an error occurs while removing the quarantine, then the **Status** field in the **Quarantines** table is updated to reflect the error.

To remove a file (or multiple files) from quarantine:

1. Open the job.
2. Open the **Quarantines** tab.
3. Select the file(s) in the Quarantines table.
4. Select the host with the correct version.
5. Click the **Release Conflict** button.

After doing this, all hosts are checked to make sure the file is not currently locked by anyone. If no locks are found, then locks are obtained on all versions of the file and the targets that are out-of-date are synchronized with the selected source host.

Manual Retries

Retries are a key feature of Peer Global File Service, used for automatically handling errors in the collaboration environment that would have otherwise led to a quarantine. For more detailed information on how retries work, see [Conflicts, Retries, and Quarantines](#) in the [Advanced Topics](#) section.

When a file is put in the retry list, it will be automatically retried based on the settings defined in [File Retries](#) in [Preferences](#). If need be, you can also manually force the retry of a file. This can be done from the Retries list of a specific File Synchronization job.

You may also choose to perform no action, in which case, the file is removed from the Retries list but none of the file versions are modified. Therefore, if the files are not currently in-sync, then the next time the file is modified, changes will be propagated to the other hosts. If an

error occurs while forcing the retry, then the **Status** field in the **Retries** table is updated to reflect the error.

To manually force the retry of a file (or multiple files):

1. Select the file(s) in the **Retries** list.
2. Select the host with the correct version.
3. Click the **Release Conflict** button.

After doing this, all hosts are checked to make sure the file is not currently locked by anyone. If no locks are found, then locks are obtained on all versions of the file and the targets that are out-of-date are synchronized with the selected source host.

Index

- A -

Access, PeerGFS API 157
Active Directory authentication 365
Active Directory group 358
Active Directory groups 365, 368, 371
Active Directory user 114, 358
Active Directory users 365, 368, 371
Advanced configuration options, Amazon FSxN 324
Advanced configuration options, Dell PowerScale 332
Advanced configuration options, Dell Unity 339
Advanced configuration options, NetApp ONTAP 347
Advanced configuration options, Nutanix Files 352
Agent availability 295
Agent configuration 141
Agent configuration, Broker 144
Agent configuration, General 146
Agent configuration, Logging 149
Agent configuration, Performance 149
Agent connectivity, preferences 295
Agent properties, editing 154
Agent properties, viewing 152
Agent Summary view 61
Agent, disabled 140
Agent, disconnected 295
Agent, installation 16
Agent, re-enabling 140
Agent, update 21, 23
Agent, updating 156
Agent, VM options 151
Agents view 35, 38
Alert 311
Alerts 103
Alerts view 36
Amazon FSxN prerequisites 8
Amazon FSxN, advanced configuration options 324
Amazon FSxN, configuration 319, 342
Amazon FSxN, credentials 319, 342
Amazon S3 410, 415
Amazon S3, bucket details 419
Analytics, preferences 166
API 156

API access 157
API categories 159
API integration 158
API reference 159
API status codes 159
Application support, File Collaboration job 572, 609
Application support, File Replication job 701
Application support, File Synchronization job 805
Assigning tags to web roles 378
Auto match root 516
Auto-restarting, File Synchronization job 818
Azure Blob Storage 408
Azure Storage, container details 417

- B -

Bait file 311
Base role 115
Base web role 116
Bash 158
Basic concepts 83
Batched emails 254
Batched real-time replication, Cloud Backup and Replication job 431
Broker configuration, preferences 297
Broker, Agent configuration 144
Browsing files and folders 293
Bucket details, Amazon S3 419
Bucket details, Cloud Backup and Replication job 416
Bucket details, NetApp StorageGRID 421
Bucket details, Nutanix Objects 423
Bucket details, Wasabi 425

- C -

CDP (continuous data protection), Cloud Backup and Replication job 432
Cloud Backup and Replication job runtime view 44
Cloud Backup and Replication job, batched real-time replication 431
Cloud Backup and Replication job, bucket details 416
Cloud Backup and Replication job, container details 416
Cloud Backup and Replication job, continuous data protection 432
Cloud Backup and Replication job, creating 379
Cloud Backup and Replication job, deleting 448

- Cloud Backup and Replication job, delta-level replication 435
 - Cloud Backup and Replication job, email alerts 438
 - Cloud Backup and Replication job, file metadata 435
 - Cloud Backup and Replication job, miscellaneous options, 435
 - Cloud Backup and Replication job, monitoring 447
 - Cloud Backup and Replication job, NTFS permission replication 435
 - Cloud Backup and Replication job, proxy configuration 384
 - Cloud Backup and Replication job, recovering data 449
 - Cloud Backup and Replication job, rehydrated data availability 435
 - Cloud Backup and Replication job, replication schedule 428, 429, 431, 432
 - Cloud Backup and Replication job, running 444
 - Cloud Backup and Replication job, snapshot retention 433
 - Cloud Backup and Replication job, SNMP notifications 440
 - Cloud Backup and Replication job, source snapshots 434
 - Cloud Backup and Replication job, starting 444
 - Cloud Backup and Replication job, stopping 446
 - Cloud Backup and Replication job, storage tiering options 435
 - Cloud Backup and Replication job, scheduled scans 429
 - Cloud Backup and Replication jobs 378
 - Cloud Backup and Replication jobs, overview 379
 - Cloud Backup and Replication jobs, preferences 191
 - Cloud Backup and Replication, job, replication and retention policy 427
 - Cloud Summary view 61
 - Collab, Sync, and Repl Summary view, Edge Caching tab 66
 - Collab, Sync, and Repl Summary view, Reports tab 66
 - Collab, Sync, and Repl Summary view, Summary tab 63
 - Collaboration, Synchronization, and Replication Summary view 63
 - Complex regular expressions 93
 - Concepts 83
 - Configuration 9
 - Configuration, Agent 141
 - Configuration, Amazon FSxN 319, 342
 - Configuration, Dell 325
 - Configuration, Dell PowerScale 329
 - Configuration, NetApp ONTAP 342
 - Configuration, Nutanix 348
 - Configuration, PMC 9, 10
 - Configuration, Preferences 165
 - Configuration, Unity 337
 - Conflict resolution, File Collaboration job 598
 - Conflict resolution, File Replication job 684
 - Conflict resolution, File Synchronization job 792
 - Conflict resolution, metadata 602, 687, 688, 796, 797, 799
 - Conflicts 128
 - Connection statuses 140
 - Container details, Azure Storage 417
 - Container details, Cloud Backup and Replication job 416
 - Continuous data protection, Cloud Backup and Replication job 432
 - Create File Collaboration job 516
 - Create file filters 85
 - Create File Synchronization job 516
 - Create folder filters 85
 - Credentials 197
 - Credentials, Amazon FSxN 319, 342
 - Credentials, Dell PowerScale 329
 - Credentials, Dell Unity 337
 - Credentials, NetApp ONTAP 342
 - Credentials, Nutanix 348
 - Custom web role 119
- D -**
- Dashboard view 39
 - Data recovery, Cloud Backup and Replication job 449
 - Database connections 194
 - Delayed replication 159
 - Dell configurations 325
 - Dell PowerScale, configuration 329
 - Dell PowerScale, credentials 329
 - Dell PowerSpace prerequisites 8
 - Dell Unity prerequisites 8
 - Dell Unity, configuration 337
 - Dell Unity, credentials 337
 - Delta-level replication, Cloud Backup and Replication job 435
 - Delta-level replication, File Collaboration job 600
 - Delta-level replication, File Replication job 686
 - Delta-level replication, File Synchronization job 794

Destination credentials 408, 410, 412, 414, 415
 Destination snapshot report 200
 Destination storage credentials 197
 Detector settings 516
 DFS 131
 DFS namespace 516
 DFS namespace folder, adding 504
 DFS namespace root path 280, 469, 499, 516
 DFS namespace server, adding 499
 DFS Namespace, elements 129
 DFS namespace, failover and failback 224
 DFS namespace, File Collaboration job 576
 DFS namespace, File Synchronization job 769
 DFS Namespaces service 469, 499
 DFS namespaces, connecting to jobs 516
 DFS namespaces, managing 499
 DFS-N Link, File Collaboration job 615
 DFS-N Link, File Synchronization job 811
 DFS-N Management job, creating 463
 DFS-N Management job, folder targets 473
 DFS-N Management job, importing existing namespace 486
 DFS-N Management job, Management Agent 465
 DFS-N Management job, namespace folders 473
 DFS-N Management job, namespace name 468
 DFS-N Management job, namespace server 469
 DFS-N Management job, namespace settings 471
 DFS-N Management job, running 496
 DFS-N Management job, starting 496
 DFS-N Management job, stopping 498
 DFS-N Management job, verifying Agent 466
 DFS-N Management jobs 462
 DFS-N Management jobs, preferences 280
 Disabled agent 140
 Disconnected agent 295

- E -

Edge Caching tab, Collab, Sync, and Repl Summary view 66
 Edge caching, File Collaboration job 552
 Edge caching, File Synchronization job 742
 Edge caching, master-edge assignment 567, 592, 757
 Editing multiple File Collaboration jobs 616
 Editing multiple File Replication jobs 708
 Editing multiple File Synchronization jobs 812
 Email alerts, Cloud Backup and Replication job 438

Email alerts, concepts 83
 Email alerts, DFS-N Management job 478
 Email alerts, File Collaboration job 573, 611
 Email alerts, File Replication job 666, 703
 Email alerts, File Synchronization job 767, 808
 Email alerts, preferences 200, 254, 298
 Email configuration preferences 289
 Enhanced metadata conflict resolution 602, 687, 688, 796, 797, 799
 Environmental requirements 8
 Event detection 354
 Expedited Sync Queue 221

- F -

File and folder filters, Cloud Backup and Replication job 405
 File and folder filters, File Collaboration job 596
 File and folder filters, File Replication job 681
 File and folder filters, File Synchronization job 789
 File and folder filters, preferences 205, 261
 File Collaboration job, application support 572, 609
 File Collaboration job, auto-restarting 622
 File Collaboration job, conflict resolution 598
 File Collaboration job, create job from namespace folder 516
 File Collaboration job, creating 529
 File Collaboration job, delta-level replication 600
 File Collaboration job, DFS namespace 576
 File Collaboration job, DFS-N Link 615
 File Collaboration job, edge caching 552
 File Collaboration job, editing 581
 File Collaboration job, editing participants 590
 File Collaboration job, email alerts 573, 611
 File Collaboration job, file and folder filters 596
 File Collaboration job, file locking 604
 File Collaboration job, file metadata 570, 602
 File Collaboration job, general 594
 File Collaboration job, host connectivity issues 624
 File Collaboration job, Management Agent 533
 File Collaboration job, manual retries 626
 File Collaboration job, master data service 555
 File Collaboration job, participants 531, 583
 File Collaboration job, pinning filter 563
 File Collaboration job, quarantined file 625
 File Collaboration job, running 618
 File Collaboration job, runtime view 46
 File Collaboration job, runtime view, Alerts tab 54

- File Collaboration job, runtime view, Configuration tab 57
- File Collaboration job, runtime view, Event Log tab 51
- File Collaboration job, runtime view, Participants tab 55
- File Collaboration job, runtime view, Quarantines tab 52
- File Collaboration job, runtime view, Retries tab 53
- File Collaboration job, runtime view, Session tab 49
- File Collaboration job, runtime view, Summary tab 47
- File Collaboration job, SNMP notifications 613
- File Collaboration job, starting 620
- File Collaboration job, stopping 622
- File Collaboration job, tags 614
- File Collaboration job, target protection 609
- File Collaboration job, utilization policy 561, 751
- File Collaboration job, volume policy 558
- File Collaboration jobs 527
- File Collaboration jobs, editing multiple jobs 616
- File Collaboration jobs, overview 528
- File Collaboration jobs, preferences 219
- File conflicts 128
- File extensions 311
- File filters, concepts 84
- File filters, create 85
- File filters, predefined 85
- File filters, usage notes 100
- File locking 265
- File locking, File Collaboration job 604
- File locking, File Replication job 696
- File locking, File Synchronization job 803
- File menu 73
- File metadata 133
- File metadata, Cloud Backup and Replication job 435
- File metadata, conflict resolution 602, 687, 688
- File metadata, File Collaboration job 570, 602
- File metadata, File Replication job 687, 688
- File metadata, File Synchronization job 761, 796, 797, 799
- File quarantine 128
- File Replication job, application support 701
- File Replication job, conflict resolution 684
- File Replication job, delta-level replication 686
- File Replication job, editing participants 677
- File Replication job, email alerts 666, 703
- File Replication job, file and folder filters 681
- File Replication job, file locking 696
- File Replication job, file metadata 687, 688
- File Replication job, General 679
- File Replication job, job relay 694
- File Replication job, participants 672
- File Replication job, runtime view 46, 58
- File Replication job, scheduled replication filters 682
- File Replication job, SNMP notifications 706
- File Replication job, tags 707
- File Replication job, target protection 702
- File Replication jobs 626
- File Replication jobs, editing multiple 708
- File Replication jobs, preferences 219
- File retries 128, 260
- File Retries and Source Snapshots, preferences 203
- File Synchronization job, application support 805
- File Synchronization job, auto-restarting 818
- File Synchronization job, conflict resolution 792
- File Synchronization job, create job from namespace folder 516
- File Synchronization job, creating 711
- File Synchronization job, delta-level replication 794
- File Synchronization job, DFS namespace 769
- File Synchronization job, DFS-N Link 811
- File Synchronization job, edge caching 742
- File Synchronization job, editing 670, 772
- File Synchronization job, editing participants 782
- File Synchronization job, email alerts 767, 808
- File Synchronization job, file and folder filters 789
- File Synchronization job, file locking 803
- File Synchronization job, file metadata 761, 796, 797, 799
- File Synchronization job, General 786
- File Synchronization job, host connectivity issues 820
- File Synchronization job, Management Agent 714
- File Synchronization job, Master Data Service 745
- File Synchronization job, participants 713, 775
- File Synchronization job, pinning filter 753
- File Synchronization job, running 814
- File Synchronization job, runtime view 46, 59
- File Synchronization job, SNMP notifications 810
- File Synchronization job, starting 816
- File Synchronization job, stopping 818
- File Synchronization job, tags 811
- File Synchronization job, target protection 806
- File Synchronization job, volume policy 748
- File Synchronization jobs 710

- File Synchronization jobs, editing multiple 812
- File Synchronization jobs, preferences 219
- Files and folders, browsing 293
- Filter expressions 102
- Filter patterns 88
- Filter patterns, complex regular expressions 93
- Filter patterns, excluding temporary files 90
- Filter patterns, using wildcards 89
- Filter, scheduled replication 159
- Filtering by date 93
- Filtering by file size 96
- Filtering folders 98
- Filters, file 100
- Filters, folders 98
- Filters, list 101
- Folder filters, concepts 84
- Folder filters, create 85
- Folder target, adding 511
- Folders, filtering 98

- G -

- General Configuration, preferences 292, 293
- General, Agent configuration 146
- General, File Collaboration job 594
- General, File Replication job 679
- General, File Synchronization job 786

- H -

- Heartbeats 295
- Help menu 76
- Hidden items 293

- I -

- Importing existing namespace 486
- Initial synchronization process 619, 815
- Initialization process 619, 815
- Installation 9
- Installation, Peer Agent 16
- Installation, PMC 9, 10
- Internal user 114, 358
- Internal users 358, 360, 361, 364

- J -

- Job alert 311
- Job Alerts view 40
- Job initialization process 619, 815
- Job relay, File Replication job 694
- Job, manual retries 821
- Jobs views 41
- Jobs, Cloud Backup and Replication 378
- Jobs, DFS-N Management 462
- Jobs, File Collaboration 527
- Jobs, File Replication 626
- Jobs, File Synchronization 710

- K -

- Keytool certificate management utility 163

- L -

- Last modified date 93
- LDAP administrator 365
- LDAP settings 358
- License 162
- License file 309
- Licensed storage capacity 162
- Licensing 309
- Licensing preferences 309
- Link namespace folder to job 523
- Link namespace to job 516
- List filter expressions 102
- List filters, concepts 101
- List filters, examples 102
- List filters, managing 103
- List filters, removing 103
- Locking 265
- Log files 105
- Logging 103
- Logging, Agent configuration 149

- M -

- Malicious Event Detection 311
- Malware 311
- Management Agent 383
- Management Agent, DFS-N Management job 465

Management Agent, File Collaboration job 533
 Management Agent, File Synchronization job 714
 Managing web roles 372
 Master Data Service, File Collaboration job 555
 Master Data Service, File Synchronization job 745
 Master-edge assignment, edge caching 567, 592, 757
 MED 311
 MED configuration, preferences 311
 Menus 29
 Menus, File 73
 Menus, Help 76
 Menus, Tools 75
 Menus, Window 73
 Metadata 133
 Metadata conflict resolution 796, 797, 799
 Miscellaneous options, Cloud Backup and Replication job 435

- N -

Namespace elements 129
 Namespace failback 131
 Namespace failover 131
 Namespace folder 516
 Namespace folder target, adding 511
 Namespace folder target, disable 224
 Namespace folder target, renewable 224
 Namespace folder, adding 504
 Namespace folder, linking to job 523
 Namespace server 469
 Namespace server, adding 499
 Namespace Summary view 70
 Namespace, linked to job 516
 NAS configuration, Amazon FSxN 319, 342
 NAS configuration, Dell 325
 NAS configuration, Dell PowerScale 329
 NAS configuration, Dell Unity 337
 NAS configuration, NetApp ONTAP 342
 NAS configuration, Nutanix 348
 NAS configuration, preferences 319
 NetApp ONTAP prerequisites 8
 NetApp ONTAP, advanced configuration options 347
 NetApp ONTAP, configuration 342
 NetApp ONTAP, credentials 342
 NetApp StorageGRID 412
 NetApp StorageGRID, bucket details 421

NTFS permission replication, Cloud Backup and Replication job 435
 NTFS permissions 570, 602, 761, 796, 797
 NTFS permissions, File Replication job 687, 688
 Nutanix Files prerequisites 8
 Nutanix Files, advanced configuration options 352
 Nutanix Objects 414
 Nutanix Objects, bucket details 423
 Nutanix, configuration 348
 Nutanix, credentials 348

- P -

Participants, adding to File Collaboration job 584
 Participants, adding to File Replication job 672
 Participants, adding to File Synchronization job 775
 Participants, editing 590, 677, 782
 Participants, File Collaboration job 531, 583
 Participants, File Replication job 672
 Participants, File Synchronization job 713, 775
 Participants, removing from File Collaboration job 584
 Participants, removing from File Replication job 672
 Participants, removing from File Synchronization job 775
 Peer Agent, installation 16
 Peer Agent, update 21, 23
 Peer Agents, managing 137
 Peer Global File Service license 162, 309
 Peer Management Broker 583, 672, 775
 Peer Management Center, update 21
 PeerGFS API 157
 PeerGFS license 162, 309
 PeerGFS preferences 163
 Performance, Agent configuration 149
 Performance, preferences 208, 266
 Permissions 115
 Permissions, base web role 116
 Permissions, custom web role 119
 Pinning filter, File Collaboration job 563
 Pinning filter, File Synchronization job 753
 PMC configuration 9, 10
 PMC installation 9, 10
 PMC user interface 27, 32
 PMC, update 21
 PowerScale, advanced configuration options 332
 PowerScale, configuration 329
 PowerScale, credentials 329

Powershell 158
Predefined file filters 85
Preferences 163
Preferences, Agent connectivity 295
Preferences, Analytics 166
Preferences, Broker configuration 297
Preferences, Cloud Backup and Replication 191
Preferences, configuring 165
Preferences, database connections 194
Preferences, DFS-N Management jobs 280
Preferences, email alerts 200, 254
Preferences, Email configuration 289
Preferences, file and folder filters 205, 261
Preferences, File Collaboration jobs 219
Preferences, File Replication jobs 219
Preferences, File Retries and Source Snapshots 203
Preferences, File Synchronization jobs 219
Preferences, General configuration 292, 293
Preferences, licensing 309
Preferences, MED configuration 311
Preferences, NAS configuration 319
Preferences, performance 208, 266
Preferences, real-time event detection 221, 267, 354
Preferences, replication and retention policies 213
Preferences, Scan Manager 217, 273
Preferences, SNMP configuration 356
Preferences, SNMP Notifications 215, 237, 269
Preferences, Software updates 304
Preferences, system alerts 298
Preferences, tags 306
Preferences, User management 358
Prerequisites, Dell PowerSpace 8
Prerequisites, Dell Unity 8
Prerequisites, NetAPP ONTAP 8
Prerequisites, Nutanix Files 8
Pre-Seeding 161
Properties, Agent 152, 154
Proxy configuration 384

- Q -

Quarantined file 254, 821
Quarantined file, File Collaboration job 625
Quarantines 128

- R -

Real-time event detection 267
Real-time event detection, preferences 221, 267, 354
Real-time replication 159
Reconnect attempts 295
Recovering data, Cloud Backup and Replication job 449
Re-enable agent 140
Rehydrated data availability, Cloud Backup and Replication job 435
Removing file from quarantine 625, 821
Replication and Retention Policies, preferences 213
Replication schedule and retention policy, Cloud Backup and Replication job 427
Replication schedule, Cloud Backup and Replication job 428, 429, 431, 432
Reports tab, Collab, Sync, and Repl Summary view 66
Reports, destination snapshot 200
Reports, scan 200
Requirements, environmental 8
Retries 128
Retrying a File Collaboration job 626
Retrying a job 821
REVIT 221
Role, Administrator 115
Role, base 115
Role, Help Desk 115
Role, Power User 115
Roles 358
Runtime view, Cloud Backup and Replication job 44
Runtime view, File Collaboration job 46
Runtime view, File Replication job 46
Runtime view, File Synchronization job 46
Runtime views 43

- S -

Scan Manager, preferences 217, 273
Scan report 200
Scheduled replication 159
Scheduled replication filter 159
Scheduled replication filters, File Replication job 682

Scheduled scans, Cloud Backup and Replication job 429
 Security 163
 Seeding Target 583, 672, 775
 Smart Data Seeding 160, 161
 SNMP configuration, preferences 356
 Snapshot retention, Cloud Backup and Replication job 433
 SNMP notifications, Cloud Backup and Replication job 440
 SNMP notifications, DFS-N Management job 481
 SNMP notifications, File Collaboration job 613
 SNMP notifications, File Replication job 706
 SNMP notifications, File Synchronization job 810
 SNMP notifications, overview 110
 SNMP Notifications, preferences 215, 237, 269
 Software updates, preferences 304
 Source paths 403
 Source snapshots, Cloud Backup and Replication job 434
 Source storage platform 382
 Starting, File Synchronization job 816
 Status codes, API 159
 Stopping, File Synchronization job 818
 Storage capacity 162
 Storage information 390
 Storage platform credentials 197
 Storage tiering options, Cloud Backup and Replication job 435
 Summary tab, Collab, Sync, and Repl Summary view 63
 Summary views 60
 System alerts 293, 298
 System folders 293

- T -

Tables 81
 Tag categories 111
 Tags overview 111
 Tags, assigning 111
 Tags, File Collaboration job 614
 Tags, File Replication job 707
 Tags, File Synchronization job 811
 Tags, filtering resources 113
 Tags, preferences 306
 Tags, web roles 378
 Target protection, File Collaboration job 609
 Target protection, File Replication job 702

Target protection, File Synchronization job 806
 Temporary files 90
 Terminology 1
 Testing API access 157
 TLS certificates 163
 Toolbar 29, 79
 Tools menu 75
 Trap folder 311
 Trigger action 311

- U -

Uninstall 24
 Unity, advanced configuration options 339
 Unity, configuration 337
 Unity, credentials 337
 Update Peer Agent 21, 23
 Update Peer Management Center 21
 Update PMC 21
 Updating software 304
 Updating, Agent 156
 Upgrade 309
 User interface 27, 32
 User management 358
 User management, preferences 358
 User, Active Directory 114
 User, internal 114
 Users, internal 360, 361, 364
 Utilization policy, File Collaboration job 561, 751

- V -

View, File Collaboration job runtime 46
 View, File Replication job runtime 46
 View, File Synchronization job runtime 46
 Views, Agent Summary 61
 Views, Agents 35, 38
 Views, Alerts 36
 Views, Cloud Backup and Replication jobs runtime 44
 Views, Cloud Summary 61
 Views, Collaboration, Synchronization, and Replication 63
 Views, dashboard 39
 Views, Job Alerts 40
 Views, Jobs 41
 Views, Namespace Summary 70
 Views, runtime 43

Views, summary 60
VM options 151
Volume policy, File Collaboration job 558
Volume policy, File Synchronization job 748

- W -

Wasabi, bucket details 425
Web client interface 114
Web client services 20
Web client user 114, 115, 116, 119
Web client, accessing 20, 25
Web client, secure access 16
Web role 114
Web role, custom 115, 119
Web role, permissions 119
Web role, standard 115
Web roles 115
Web roles permissions 116
Web roles, create custom 373
Web roles, delete 377, 378
Web roles, edit 377, 378
Web roles, managing 372
Wildcards, filter patterns 89
Window menu 73

- Y -

YAML file 158