



Peer Global File Service User Guide

Copyright (c) 1993-2026 Peer Software, Inc. All Rights Reserved.

Updated Thursday, September 17, 2026

Table of Contents

Peer Global File Service Help	1
Terminology	1
Requirements and Prerequisites	7
Installing Peer Management Center	7
Installing Peer Management Center and Agents	8
Windows Installation Instructions	8
Securing Access to the Web Client	14
Installing Peer Agents	15
Updating Peer Management Center and Peer Agents	19
Updating Peer Management Center	19
Updating Peer Agents	21
Uninstalling Peer Management Center and Agents	23
Running Peer Management Center	23
Launching the Rich Client	23
Accessing the Web Client	24
Peer Services Required for Web Client	26
Peer Management Center User Interface	27
Main Window	28
Main Window Menus and Toolbars	29
File Menu	29
Window Menu	30
Tools Menu	32
Help Menu	32
Toolbar	35
Perspectives and Views	38
Views	39
Jobs Perspective	40
Agents View	42
Agents Toolbar	43
Alerts View	43
Brokers View	45
Dashboard	45
Job Alerts View	47
Jobs View	48
Jobs Toolbar	49
Runtime Views	50
Cloud Backup and Replication Job Runtime View	51
DFS-N Management Job Runtime View	52
File Collaboration Job Runtime View	53
Summary Tab	54
Session Tab	56
Event Log Tab	57
Quarantines Tab	58
Retries Tab	59
Alerts Tab	60
Participants Tab	61
Configuration Tab	63

File Replication Job Runtime View	64
File Synchronization Job Runtime View	65
Summary View s	66
Agent Summary View	67
Cloud Summary View	67
Collab, Sync, and Repl Summary View	68
Summary Tab	69
Edge Caching Tab	72
Reports Tab	72
Namespace Summary View	76
Topology Perspective.....	76
Tables	78
Basic Concepts	81
Email Alerts	82
File and Folder Filters	82
Creating and Applying File Filters.....	83
Editing and Deleting File Filters	84
Predefined File Filters.....	84
Updating Predefined File Filters.....	85
Defining Filter Patterns.....	86
Using Wildcards in Filter Patterns	87
Automatically Excluded File Types.....	88
Excluding Temporary Files.....	88
Using Complex Regular Expressions in Filter Patterns.....	90
Filtering on Last Modified Date.....	91
Filtering on File Size	93
Filtering Folders	95
Folder Filter Examples	96
File Filter Usage Notes.....	97
List Filters	98
Creating Basic List Filters.....	98
Creating Complex List Filters.....	99
Predefined List Filters.....	101
Managing List Filters.....	101
Logging and Alerts	104
Retrieving Log Files.....	106
SNMP Notifications	110
Web Client Users	111
Roles for Web Client Users.....	112
Base Roles	112
Base Role Permissions.....	113
Custom Roles	115
Custom Role Permissions	115
Advanced Topics	118
Analytics	118
PeerIQ	119
File System Analytics.....	121
File System Analytics Scan Process.....	122
File Activity Analytics.....	122
Conflicts, Retries, and Quarantines	123
DFS Namespaces	124
Using DFS Namespaces w ith Jobs.....	126
DFS Namespace Failover and Failback.....	127

Edge Caching	127
Handling Symbolic Links	129
File Metadata Synchronization	130
Multi-Protocol (SMB/NFS)	132
SMB Access for Multi-Protocol Jobs	134
Network of Brokers	135
Managing Peer Agents	135
Peer Agent Connection Statuses	138
Re-enabling a Disabled Agent Within a Job	138
Editing an Agent Configuration	139
Broker	142
General	144
Logging	147
Performance	147
Viewing Agent Properties	149
VM Options	151
Editing Agent Properties	152
Updating a Peer Agent	153
PeerGFS API	154
Accessing the PeerGFS API	154
Testing the PeerGFS API	155
Integrating Your Own Tools and Scripts with the PeerGFS API	155
API Quick Reference	156
API Categories	156
Status Codes	156
Proactive Monitoring	157
Scheduled Replication	158
Smart Data Seeding	158
Storage Capacity	160
Tagging Resources	160
Understanding Tags and Categories	161
Creating Tags	161
Assigning Tags	165
Filtering with Tags	173
Associating Tags with Roles	175
Managing Tags	176
Creating Categories	177
Managing Categories	179
Managing TLS Certificates	182
TLS Certificate Terminology	185
Preparing for Certificate Management	185
Managing Broker and Agent Certificates	188
Viewing a Broker Certificate	190
Generating Self-Signed Broker Certificates	192
Importing Broker Certificates	197
Restoring Default Broker Certificates	199
Using Maintenance Mode to Update Broker Certificates	201
Downloading All Broker Certificates	203
Updating Trusted Broker Certificates	205
Managing Web Client Certificates	207
Generating Self-Signed Web Client Certificates	209
Importing Web Client Certificates	213
Restoring Default Web Client Certificates	215
Telemetry	216

Preferences	217
Configuring Preferences	219
Analytics Job Preferences	220
Analytics	221
Email Alerts	222
SNMP Notifications	226
Cloud Backup and Replication Job Preferences	228
Cloud Backup and Replication.....	229
Database Connections.....	231
Destination Credentials.....	234
Email Alerts	236
File Retries and Source Snapshots.....	240
File and Folder Filters	243
Performance	246
Proxy Configuration.....	247
Replication and Retention Policies.....	250
SNMP Notifications	252
Scan Manager	255
Collaboration, Replication, and Synchronization Job Preferences	257
Collab, Sync, and Replication.....	258
Application Enhancements.....	259
DFS-N Management.....	262
Edge Caching	264
Edge Caching	264
Email Alerts	265
Master Data Service	269
Pinning Filters	271
SNMP Notifications	276
Utilization Policies	279
Volume Policies	283
Email Alerts	293
File Retries	298
File and Folder Filters	299
Locking	302
Performance	304
Real-time Event Detection.....	306
SNMP Notifications	308
Scan Manager	312
Scheduled Replication Filters.....	314
DFS-N Management Job Preferences	319
DFS-N Management.....	320
Email Alerts	321
SNMP Notifications	325
Email Configuration	328
Basic Email Configuration.....	331
OAuth2 Email Configuration.....	333
General Configuration	337
General Configuration.....	338
Agent Connectivity.....	340
Broker Configuration.....	343
Email Alerts	343
Proxy Configuration.....	347
Software Updates.....	349
Tags Configuration.....	351

Web and API Configuration	353
Licensing	354
MED Configuration	358
PMC Backup	364
Real-time Event Detection	366
Storage Configuration	368
SMB Access Configuration.....	369
Storage Platforms.....	370
Amazon FSxN Configurations.....	371
Amazon FSxN Advanced Options.....	375
Dell Configurations	376
Dell PowerScale Configurations.....	379
Dell PowerScale Advanced Options.....	382
Dell PowerStore Configurations	386
Dell PowerStore Advanced Options.....	388
Dell Unity Configurations.....	390
Dell Unity Advanced Options.....	393
NetApp ONTAP Configurations.....	396
NetApp ONTAP Advanced Options.....	400
Nutanix Files Configurations	401
Nutanix Files Advanced Options	404
SNMP Configuration	405
Telemetry	407
Proactive Monitoring Setup.....	408
Step 1: General Configuration.....	408
Step 2: Telemetry Options.....	409
Step 3: Agent Locations	412
Step 4: Health Checker Setup.....	413
Set Up the Health Checker.....	415
Select Jobs	420
Schedule Task	421
Update Health Checker.....	422
Step 5: Confirmation	424
Disabling and Re-enabling Proactive Monitoring.....	426
User Management	429
Accessing User Management.....	429
Managing Web Client Users.....	430
Managing Internal Users	431
Creating an Internal User.....	432
Editing an Internal User.....	435
Deleting an Internal User.....	437
Managing Active Directory Users and Groups.....	438
Configuring Active Directory Authentication.....	438
Adding an Active Directory User or Group.....	440
Editing an Active Directory User or Group.....	442
Deleting an Active Directory User and Group.....	443
Managing Roles for Web Client Users.....	443
Creating a Custom Role.....	444
Editing a Custom Role	449
Deleting a Custom Role	450
Assigning Tags to a Role.....	450
Cloud Backup and Replication Jobs	450
Overview	451
Before You Create Your First Cloud Backup and Replication Job	451

Creating a Cloud Backup and Replication Job	451
Step 1: Job Type and Name	452
Step 2: Source Storage Platform	454
Step 3: Management Agent	454
Step 4: Proxy Configuration	456
Step 5: Storage Information	463
Amazon FSx for NetApp ONTAP	464
Dell PowerScale	466
Dell Unity	470
NetApp ONTAP	473
Nutanix Files	475
Step 6: Source Paths	477
Step 7: File and Folder Filters	478
Step 8: Destination	479
Step 9: Destination Credentials	480
Azure Blob Storage Credentials	481
Amazon S3 Credentials	483
NetApp StorageGRID Credentials	484
Nutanix Objects Credentials	486
Wasabi Credentials	487
Step 10: Container or Bucket Details	488
Azure Blob Storage Container Details	489
Amazon S3 Bucket Details	490
NetApp StorageGRID Bucket Details	492
Nutanix Objects Bucket Details	494
Wasabi Bucket Details	496
Step 11: Replication and Retention Policy	497
Step 12: Replication Schedule	498
Scheduled Scans	499
Batched Real-Time	501
Continuous Data Protection	502
Step 13: Retention	503
Step 14: Source Snapshots	504
Step 15: Miscellaneous Options	505
Step 16: Email Alerts	508
Step 17: SNMP Notifications	510
Step 18: Confirmation	511
Deleting a Cloud Backup and Replication Job	513
Running a Cloud Backup and Replication Job	514
Starting a Cloud Backup and Replication Job	514
Stopping a Cloud Backup and Replication Job	516
Monitoring Cloud Backup and Replication Jobs	517
Recovering Data	518
Search Options	521
Search by Name	521
Search by Snapshot	525
Search by Point in Time	527
Search by Latest Replication	529
Recovery Options	530
DFS-N Management Jobs	534
Creating a DFS-N Management Job	535
Step 1: Job Type	536
Step 2: Management Agent	538
Step 3: Agent Verification	539

Step 4: Namespace Name.....	540
Step 5: Namespace Servers.....	541
Step 6: Namespace Settings.....	543
Step 7: Namespace Folders.....	545
Step 8: Email Alerts.....	550
Step 9: SNMP Notifications.....	553
Step 10: Review.....	555
Step 11: Results.....	555
Importing an Existing Namespace	557
Running a DFS-N Management Job	568
Starting a DFS-N Management Job.....	568
Stopping a DFS-N Management Job.....	570
Managing DFS Namespaces	571
Adding a Namespace Server.....	571
Adding a Namespace Folder.....	576
Adding a Folder Target.....	583
Linking a DFS Namespace to File Collaboration and File Synchronization Jobs	588
Creating a File Collaboration or File Synchronization Job from a Namespace Folder	588
Linking an Existing Namespace Folder to an Existing File Collaboration or File Synchronization Job.....	595
File Collaboration Jobs	600
Overview	600
Before You Create Your First File Collaboration Job	601
Creating a File Collaboration Job	601
Step 1: Job Type and Name.....	601
Step 2: Participants.....	603
Management Agent	604
Storage Platform	605
Storage Information	607
Amazon FSx for NetApp ONTAP.....	608
Dell PowerScale	610
Dell PowerStore	614
Dell Unity	617
NetApp ONTAP	621
Nutanix Files	622
Windows File Server.....	624
Windows File Server Advanced Options.....	625
Path	628
Configuring SMB Access	631
Edge Caching	633
Master Data Service	635
Volume Policy	638
Utilization Policy	641
Pinning Filter	643
Step 3: Master-Edge Assignment.....	647
Step 4: File Metadata.....	650
SMB Protocol	651
Multi-Protocol	652
Step 5: Application Support.....	655
Step 6: Email Alerts.....	657
Step 7: DFS Namespace.....	659
Step 8: Save Job.....	663
Editing a File Collaboration Job	664
Participants	666

Adding and Deleting a Participant.....	667
Editing a Participant	673
Master-Edge Assignment.....	676
General	678
File and Folder Filters.....	680
Scheduled Replication Filters.....	681
Conflict Resolution.....	682
Delta Replication	684
File Metadata	686
SMB Protocol	687
Multi-Protocol	689
File Locking	692
Application Support.....	696
Target Protection.....	697
Email Alerts	698
SNMP Notifications.....	701
Tags	702
DFS-N Link	703
Editing Multiple Jobs.....	704
Running and Managing a File Collaboration Job	705
Overview	706
Job Initialization Process.....	706
Initial Synchronization Process.....	706
Starting a File Collaboration Job.....	707
Stopping a File Collaboration Job.....	709
Auto-Restarting a File Collaboration Job.....	709
Host Connectivity Issues.....	711
Removing a File from Quarantine.....	712
Manual Retries	713
File Replication Jobs	714
Overview	714
Before You Create Your First File Replication Job	714
Creating a File Replication Job	714
Step 1: Job Type and Name.....	715
Step 2: Source Agent.....	716
Step 3: Storage Platform.....	718
Step 4: Storage Information.....	720
Amazon FSx for NetApp ONTAP.....	721
Dell Pow erScale	722
Dell Pow erStore	726
Dell Unity	729
Linux File Server	732
Linux File Server Advanced Options.....	733
NetApp ONTAP	734
Nutanix Files	736
Window s File Server	737
Window s File Server Advanced Options.....	738
Step 5: Source Path.....	741
Configuring SMB Access.....	743
Step 6: Destination Agent.....	746
Step 7: Destination Path.....	747
Configuring SMB Access.....	751
SMB Path Details	753
Step 8: File Metadata.....	754

SMB Protocol	754
NFS Protocol	756
Multi-Protocol	757
Step 9: Email Alerts	759
Step 10: Save Job	762
Editing a File Replication Job	763
Participants	765
Adding and Deleting a Participant	765
Editing a Participant	771
General	773
File and Folder Filters	776
Scheduled Replication Filters	778
Conflict Resolution	779
Delta Replication	781
File Metadata	783
SMB Protocol	784
NFS Protocol	786
Multi-Protocol	788
Job Relay	790
File Locking	793
Application Support	797
Target Protection	798
Email Alerts	799
SNMP Notifications	802
Tags	803
Editing Multiple Jobs	803
File Synchronization Jobs	805
Overview	805
Before You Create Your First File Synchronization Job	806
Creating a File Synchronization Job	806
Step 1: Job Type and Name	806
Step 2: Participants	808
Management Agent	809
Storage Platform	810
Storage Information	813
Amazon FSx for NetApp ONTAP	814
Dell Pow erScale	816
Dell Pow erStore	820
Dell Unity	823
Linux File Server	827
Linux File Server Advanced Options	828
NetApp ONTAP	829
Nutanix Files	830
Window s File Server	832
Window s File Server Advanced Options	833
Path	836
Configuring SMB Access	839
Edge Caching	842
Master Data Service	845
Volume Policy	848
Utilization Policy	851
Pinning Filter	853
Step 3: Master-Edge Assignment	858
Step 4: File Metadata	862

SMB Protocol	862
NFS Protocol	864
Multi-Protocol	865
Step 5: Email Alerts	868
Step 6: DFS Namespace	871
Step 7: Save Job	874
Editing a File Synchronization Job	875
Participants	877
Adding and Deleting a Participant	878
Editing a Participant	884
Master-Edge Assignment	888
General	890
File and Folder Filters	893
Scheduled Replication Filters	894
Conflict Resolution	896
Delta Replication	898
File Metadata	900
SMB Protocol	901
NFS Protocol	903
Multi-Protocol	905
File Locking	907
Application Support	909
Target Protection	910
Email Alerts	912
SNMP Notifications	915
Tags	916
DFS-N Link	917
Editing Multiple Jobs	918
Running and Managing a File Synchronization Job	920
Overview	920
Job Initialization Process	921
Initial Synchronization Process	921
Starting a File Synchronization Job	922
Stopping a File Synchronization Job	924
Auto-Restarting a File Synchronization Job	924
Host Connectivity Issues	926
Removing a File from Quarantine	926
Manual Retries	927

Index

928

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) 1993-2026 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, September 17, 2026

PeerGFS Version 6.5

Terminology

Introduction

Before getting started, familiarize yourself with the key concepts and terminology used throughout PeerGFS. This reference can help clarify language used in the interface and documentation. Refer back to it as needed if you encounter unfamiliar terms or concepts.

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.
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.

Term	Definition
file access event	An event that is triggered from the opening or closing of a file.
file change event	An event that causes a file to be changed in some way, for example, file modify, file delete, file rename, file attribute change.
File Collaboration job	A type of job that combines file synchronization with distributed file locking to prevent version conflicts across multiple active file servers.
File Collaboration session	A communication session made up of two or more hosts, each with a designated root of folders and files that are to be shared or collaborated on. A collaboration session coordinates the primary functions of file locking and synchronization.
file filter	A type of filter used to include or exclude specific files from replication and locking.
file lock conflict	A file collaboration condition that exists when two users open a file at the same time, and both hold exclusive locks on the file.
File Replication job	A type of job that involves real-time and/or scheduled copying of files and folders from one file server to another.
File Synchronization job	A type of job that involves multi-directional real-time replication so that two or more file servers are always up to date with each other.
filter	Three types of filters: file, folder, and list.
filter expression	See list filter.
folder filter	A type of filter used to include or exclude specific folders (and the files they contain) from replication and locking.

Term	Definition
heartbeat	A communication mechanism used between Peer Management Center and all connected Peer Agents to ensure that Peer Agents are alive and responsive. Heartbeats share information about the Peer Agent host server with Peer Management Center, aid in verifying when a Peer Agent is no longer available, and signal when a disconnected Peer Agent has reconnected. All heartbeat information is sent through the Peer Management Broker.
host	A server that a Peer Agent is installed upon.
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.
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.
list filter	A type of filter used to show or hide information from various views in Peer Management Center.
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.
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.
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

Term	Definition
	Peer Global File Service solution suite. Typically installed on or alongside the file servers that will be managed by Peer Management Center.
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.
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.
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.
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

Term	Definition
	(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 a 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. 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 folder path or paths specified for monitoring by a PeerGFS job, including all subfolders beneath those paths.

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 PowerStore 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 PowerStore Configuration Guide](#)
- [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 Service 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 6.0.0.20240621

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	Enter the hostname or IP address via which the services can be accessed: • Enter localhost or 127.0.0.1 if you want the services to be accessible only to users of the local server via the loopback interface. • Enter 0.0.0.0 to make the services accessible via all network interfaces. • Enter a specific IP address to restrict access to a specific network interface.
Enable HTTPS Web Access	Select this to enable secure access to the web service, and then enter a port number.
Enable HTTPS 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 as the web service port.

4. If you enabled access to the web client, see [Securing Access to the Web Client](#) for additional information about securing access to the web client.
5. On the final page of the installation wizard, you have several options; we recommend that, at minimum, you select the first option.
6. 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 user name 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.

- To view the installed Peer services, open the **Windows Services** console (services.msc).




8. If you didn't select the option to launch Peer Management Center Client on the final 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.

Peer Management Center should start successfully if both the **Peer Management Center Service** and **Peer Management Broker Service** are running. If not, open the Windows Services console (services.msc) and start both services.

9. 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.




10. Click **Add/Update**.

11. Browse to where the license file is located and select the file.


12. Click **Open**.

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



13. Click **Apply** to save your changes.

The main window of Peer Management Center is displayed.



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

Securing Access to the Web Client

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](#).




The screenshot shows a Windows-style window titled "Setup - Peer Agent 6.0.0.20240621". The main heading is "Peer Management Broker Configuration". Below the heading is a instruction: "Please provide the hostname, protocol and TCP/IP port of the Broker." In the top right corner, there is a green arrow pointing right towards a computer monitor icon. The configuration fields are as follows:

- Hostname:** A text input field with a green question mark icon to its right.
- Protocol:** A dropdown menu currently showing "SSL" with a green question mark icon to its right.
- Port:** A text input field containing "61617" with a green question mark icon to its right.


Below the fields, a note states: "Please note that enabling SSL/TLS may reduce performance." At the bottom right, there are three buttons: "< Back", "Next >" (which is highlighted with a blue border), and "Cancel".

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




The screenshot shows a Windows-style window titled "Setup - Peer Agent 6.0.0.20240621". The window has standard minimize, maximize, and close buttons in the top right corner. Below the title bar, the text "Agent Service Account" is displayed in bold, followed by the instruction "Provide the account under which the Agent Service will run." To the right of this text is a green arrow pointing right, superimposed on a computer monitor icon. Below the instruction, there are two radio button options. The first option, "Enter Domain Credentials (Recommended)", is selected with a blue dot. Below this option are three text input fields: "Domain\UserName" containing the text "PEERTEST\Labadmin", "Password" which is empty and has a blue border, and "Re-enter Password" which is also empty. The second option, "Use Local System account", is unselected. At the bottom right 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 completes, the Peer Agent is installed as a Windows service. Verify that the **Peer Agent Service** is running and has successfully connected to the [Peer Management Broker](#). To do this, open the Windows Services console (services.msc) and confirm that the Peer Agent Service has started.



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:


- 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.



- 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 and managing the application:

- **Rich client:** [Install](#) and run the Peer Management Center Client (rich client application) directly on the server where Peer Management Center is deployed. The rich client is available only on Windows and is not supported on Linux-based systems. For more information, see [Launching the Rich Client](#).
- **Web client:** Access the Peer Management Center web client through a web browser, either locally on the Peer Management Center server or remotely from any system on the network with connectivity to the server. The web client offers the same functionality and responsive interface as the rich client, allowing users to manage and monitor jobs without direct access to the server. For more information, see [Accessing the Web Client](#).

Launching the Rich Client

To use the Peer Management Center Client (rich client application), you must install it on a Windows system that can connect to the Peer Management Center server. Once installed, you can launch the client to access and manage your PeerGFS environment.

To launch the rich client:

1. Double-click the **Peer Management Center Client** shortcut on the desktop. Alternatively, from the Windows **Start** menu, select **Peer Management Center Client**.
2. The application opens and connects automatically to the server configured during installation or setup. If the application fails to open or connect, verify that the required Peer services are running and that the server is reachable.

Once logged in, the Peer Management Center main window opens, allowing you to view system status, manage jobs, and access all administrative features.

Accessing the Web Client

Before accessing the Peer Management Center web client, verify that the required Peer services are running on the Peer Management Center server. For details, see [Peer Services Required for Web Client](#).

Options for Accessing the Web Client

After installing Peer Management Center, configuring the Peer Management Center Web Client Service, and starting the [required Peer services](#), you can access the Peer Management Center web client in the following ways:

- **Using a web browser:** Log in directly via a web browser on the local Peer Management Center server, as described in the next section.
- **From another system on the network:** Access the Peer Management Center web client from any system that has network connectivity to the Peer Management Center server.
- **Through the Peer Management Center Client:** Launch the Peer Management Center Client (rich client application) and select **Open Peer Management Center Web Client** from the **Window** menu. The rich client is available only on Windows.

Accessing Peer Management Center from a Web Browser

To access the Peer Management Center web client from a web browser:


1. Open a web browser.
2. In the address field, enter any one of the following URLs based on how the web service was configured during installation:

If this URL was entered in the Installation Wizard	Access Type	Use this URL
A specific IP address	Local or remote	https://<IP address>:8443/hub Example: https://10.0.0.1:8443/hub Note: You cannot use localhost even when logged directly into the Peer Management Center server.
localhost or 127.0.0.1	Local only	https://localhost:8443/hub https://127.0.0.1:8443/hub
0.0.0.0	Local or remote	https://<PMC server IP address>:8443/hub Example: https://10.0.0.1:8443/hub

Notes:

- The URL for accessing the Peer Management Center web client depends on how the web service was configured during installation. If you're unsure of the correct URL, contact the administrator who performed the installation.
- The default HTTPS port is **8443**. If your environment uses different port, replace **8443** with the appropriate port number.
- Always append **/hub** after the port number to access the PMC web client.
- To modify the web service configuration, go to the [General Configuration](#) page under [Preferences](#). For detailed instructions, see the [Web and API Configuration](#).

3. When the login page appears, enter your credentials.



- **Default local credentials:**


- o **User name:** admin
- o **Password:**
 - For most installations: password
 - For Azure Marketplace deployments: Azure Subscription ID
 - For AWS Marketplace deployments: EC2 instance ID of the PMC virtual appliance
- o For security reasons, immediately change the default admin password after logging in. See [Editing an Internal User](#) for details.

- **Active Directory credentials:** Enter your user name in this format: [username@mydomain.local](#).

4. Click **Login**.

Peer Services Required for Web Client

To use the Peer Management Center web client, certain core Peer services—shown below—must be running on the Peer Management Center server.



If a required service is not running, open the **Windows Services** panel (services.msc) on the PMC server and start the service. You can refresh the **Services** panel to verify that the service has started successfully.

For installation details, see [Installing Peer Management Center and Agents](#) in the [Installing Peer Management Center](#).

Peer Management Center User Interface

Peer Management Center is a comprehensive management console used to configure, deploy, monitor, and manage jobs. This section introduces the user interface and familiarizes you with the primary components and features available throughout the application. Understanding the layout and functionality of the interface will help you navigate Peer Management Center more efficiently and effectively.

User Interface Options

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 information, see [Launching the Rich Client](#).

Note: The rich client is supported only on Windows systems and is not available on Linux-based operating systems, including Ubuntu. Linux users should use the web client.

- **Web client:** The web client can be accessed either locally from the Peer Management Center server or remotely from any network-connected system with access to the server. The web client delivers functionality comparable to the rich client while providing the flexibility of remote access.

Using the web client, users can manage and monitor jobs without needing direct access to the Peer Management Center server. Access to web client functionality is controlled through role-based permissions.

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

User Interface Components


The user interface is organized into several key components:

- [Main window](#) – The primary workspace where you access tools, views, and operational information. Familiarity with the main window layout is essential for efficiently navigating and managing tasks within Peer Management Center.
- [Menus and toolbars](#) – Provide access to commands, navigation options, and frequently used functions throughout the application.
- [Perspectives and views](#) – Perspectives organize related views and tools for specific tasks or workflows. Views display operational data, configuration settings, runtime information, and other relevant details.
- [Tables](#) – Tables present information in a structured format, making it easier to locate, sort, and manage data quickly and efficiently.

To learn how to customize the interface to match your preferences and workflow requirements, see [Preferences](#).

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.



Main Window 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.

Command	Description
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.
Log Out admin	(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.
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.

Command	Description
Show Job Summary	Displays a submenu with the following options: • 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	Displays a submenu with the following options: • Alerts - Displays the Alerts view, which displays Peer Management Center alerts such as Peer Agent connection status changes. • 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 Center API	Opens the Peer Global File Service API in a web browser.

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 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 Guides	Displays a submenu with the following options: • PeerGFS — Opens the PeerGFS User Guide. • PeerIQ — Opens the PeerIQ User Guide.

Command	Description
	By default, the online version of the selected guide opens. If the online version is unavailable, the bundled version opens. In the Windows rich client, press and hold Shift while selecting a guide to open the local version instead.
Knowledge Base	Opens the Peer Software Knowledge Base in a web browser.
Support Portal	Opens the Support Portal in a web browser.
Support Tools	Displays a submenu containing tools for collecting diagnostic information. See Support Tools for details.
Download Peer Agent Installer	Provides the following options for downloading a Peer Agent installer or virtual appliance image compatible with this version of the PMC: • Windows installer • Linux installer • OVA Virtual Appliance for VMware ESX. • QCOW2 Virtual Appliance for Nutanix AHV. • OVA Virtual Appliance for Microsoft Hyper-V
Download Broker Installer	Provides the following options for downloading a Peer Broker installer: • Windows installer • Linux installer • OVA Virtual Appliance for VMware ESX. • QCOW2 Virtual Appliance for Nutanix AHV. • OVA Virtual Appliance for Microsoft Hyper-V See Getting Started with the Broker Virtual Appliance for more information on the Peer Broker Virtual Appliance.

Command	Description
Download PeerIQ	Provides the following options for downloading a PeerIQ virtual appliance image: • OVA Virtual Appliance for VMware ESX. • QCOW2 Virtual Appliance for Nutanix AHV. • OVA Virtual Appliance for Microsoft Hyper-V See Getting Started with the PeerIQ Virtual Appliance for more information on deploying the Peer Agent Virtual Appliance.
Check for Updates	(Rich client only) Checks whether an update to Peer Management Center is available. Minor releases can be downloaded and installed automatically. To obtain a major release and the required new license key, contact Peer Support.
Licenses	Opens the Licensing page in Preferences.
About Peer Management Center	Displays PMC version information and installation information.


Support Tools

The following options are available from the **Support Tools** submenu:

Command	Description
Retrieve PMC/Agent Logs	Collects logs and support files for selected Peer Agents, the PMC, and all jobs, and packages them in a single encrypted ZIP file. Collection occurs in the background and might take some time, depending on the amount of content and network speed. When collection is complete, you are notified and can view the ZIP file or upload it to Peer Support .
Retrieve Broker Statistics	Retrieves detailed messaging statistics for Peer Agent and PMC connections to a Peer Broker. Peer Support can use these statistics to diagnose connection and messaging issues.
Generate Thread Dump	Displays a submenu with the following options: • Generate PMC Service Thread Dump • Generate Broker Service Thread Dump • Generate Topology Service Thread Dump Peer Support can use thread dumps to diagnose problems with the selected service.
Generate Memory Dump	Displays a submenu with the following options: • Generate PMC Service Thread Dump • Generate Broker Service Thread Dump • Generate Topology Service Thread Dump Peer Support can use thread dumps to diagnose problems with the selected service.

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.

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

There are two methods to open a perspective:


- Choose **Open Perspective** from the **Window** menu, and then select a perspective.
- 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 shows the Peer Global File Service interface. The 'Collab, Sync, and Repl Summary' view is active, displaying a table of job status. The table has columns for Name, Overall Status, Job Type, Failed, Ho, Quarant, Retries, Errors, Warnings, Open, File, Pending, I, Queued, I, Backgrou, Scan, Status, Elapsed Time, Session Structure, and Synchronization Pri. The table lists several jobs, including FS-1, FC-2, FR-3, FS-2, FR-2, FR-1, and FC-1, with their respective statuses and details.

Name	Overall Status	Job Type	Failed	Ho	Quarant	Retries	Errors	Warnings	Open	File	Pending	I	Queued	I	Backgrou	Scan	Status	Elapsed Time	Session Structure	Synchronization Pri
FS-1	Halted	File Synchronization	0	0	0	0	0	0	0	0 bytes	0	0	0	0	0	0	Stopped	00:00:00	Size: 934 MB, Files: 2	
FC-2	Halted	File Collaboration	0	0	0	1	0	0	0 bytes	0	0	0	0	0	0	0	Stopped	00:00:00	Size: 2 GB, Files: 4	2
FR-3	Halted	File Replication	0	0	0	1	0	0	0 bytes	0	0	0	0	0	0	0	Stopped	00:00:00	Size: 929.01 MB, F	2
FS-2	Running	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	0	0	0	0	0	Completed - 00:00:02	18:56:09	Size: 1.95 GB, File	2
FR-2	Running	File Replication	0	0	0	0	0	0	0 bytes	0	0	0	0	0	0	0	Completed - 00:00:03	18:56:09	Size: 649 MB, Files: 2	
FR-1	Running	File Replication	0	0	0	0	0	0	0 bytes	0	0	0	0	0	0	0	Completed - 00:00:03	18:56:09	Size: 1.77 GB, File	2
FC-1	Running	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	0	0	0	0	0	Completed - 00:00:04	18:56:20	Size: 2.02 GB, File	2

Active Jobs -> Failed Participants: 0 of 4 | Bytes Pending: 0 bytes | Bytes Transferred: 2.9 GB | Opens: 0 | Initial Scans Completed: 4 of 7 | Total Size: 10.19 GB | Total Files: 218 | Total Directories: 5

Alerts: 0 errors, 14 warnings, 65 others | Filter by: Host: | Severity: | Type: | Clear Alerts

Received Date	Severity	Type	Name	Host	Message	Exception
07-16-2024 01:00	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Maintenance	
07-15-2024 23:59	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Nightly Delete Shadow Copies Job	
07-15-2024 23:30	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Database Maintenance	
07-15-2024 22:37	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Process Retention Policy	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a completed elapsed tin	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Free space on volume C:\ is less than the minimum required size to achieve optim	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\ on host Lab24Win19c completed elapsed tin	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Free space on volume C:\ is less than the minimum required size to achieve optim	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a has started.	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\ on host Lab24Win19c has started.	

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, with each quadrant displaying information in panels called [views](#). These views provide access to job configuration information, runtime activity, environment status, alerts, and operational metrics.

The following illustration identifies the primary quadrants and views available in the Jobs perspective.



Views in the Jobs Perspective

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


Views	Description
Jobs view	Displays a list of all jobs grouped by job type. The toolbar in this view provides options to start and stop jobs. This view also provides access to job-specific configuration, status, and activity information.
Agents and Brokers views	Contains two views: • Agents view – Displays a list of known Peer Agents and the connection status for each. Individual Peer Agents can be updated and restarted by right-clicking one or more Agents and selecting the appropriate action from the shortcut menu. • Brokers view – Displays all known brokers installed in the environment along with the Agents connected to each broker. Actions available in the Agents view can also be performed by right-clicking an Agent within the Brokers view.
Summary and Runtime views	Displays several types of operational views, including a dashboard that provide metrics and key performance indicators, summaries of jobs by job type, an Agent Summary view, and runtime statistics for individual jobs. The Agent Summary view displays all known deployed Peer Agents and detailed status information that can be used to assess the health of the environment. Some views in this area display a combination of Runtime statistics for individual jobs, real-time file I/O activity, history, and configuration information for a specific job, while others display summary information about all jobs of a specific type. Some views in this area display runtime statistics, real-time file I/O activity, history, and configuration information for a specific job, while others display summary information about all jobs of a specific type.
Alerts views	Contains a variety of alert views, including the Alerts view and the Jobs Alerts view: • Alerts view: Displays a list of Peer Management Center alerts with detailed information about each alert. Alerts relating to Peer Agent connection status changes are reported here. • Jobs Alerts view: Displays a list of job-specific alerts that have occurred, include alerts relating to the automatic stopping and restarting of jobs.

Agents View

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.



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.

Alerts View

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

The screenshot shows the Peer Management Center interface. The left sidebar contains navigation options: Cloud Backup and Replication (1), DFS-N Management (1), File Collaboration (2), File Replication (3), and File Synchronization (2). The main area displays a table of alerts. The table has columns: Received Date, Severity, Type, Name, Host, Message, and Exception. The table shows several informational alerts related to scheduled tasks and edge caching updates. A summary bar at the bottom indicates 0 errors, 14 warnings, and 65 other alerts.

Received Date	Severity	Type	Name	Host	Message	Exception
07-16-2024 01:00	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Maintenance	
07-15-2024 23:59	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Nightly Delete Shadow Copies Job	
07-15-2024 23:30	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Database Maintenance	
07-15-2024 22:37	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Process Retention Policy	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a completed elapsed tin	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Free space on volume C:\ is less than the minimum required size to achieve optim	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\ on host Lab24Win19c completed elapsed tin	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Free space on volume C:\ is less than the minimum required size to achieve optim	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a has started.	
07-15-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\ on host Lab24Win19c has started.	

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.

Brokers View

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 interface. The left sidebar shows a hierarchical tree view of the environment, with the 'Brokers' view selected. The main pane shows a table of brokers and their agents. The table has columns: Name, Overall Status, Job Type, Failed, Ho, Quarant, Retries, Errors, Warnings, Open, File, Pending, I, Queued, I, Backgrou, Scan, Status, Elapsed Time, Session Structure, Synchronization Pr. The table lists various brokers and their associated agents. Below the table, there are summary statistics: Active Jobs, Failed Participants, Bytes Pending, Bytes Transferred, Opens, Initial Scans Completed, Total Size, Total Files, and Total Directories. At the bottom, there is an 'Alerts' section showing a list of alerts with columns: Received Date, Severity, Type, Name, Host, Message, and Exception.

Dashboard

The Dashboard is divided into two sections:


- **Collab, Sync, and Repl** - This top section displays a table of metrics and key performance indicators for all running File Collaboration, File Synchronization, and File Replication jobs. It

also contains a link that opens the [Collab, Sync, and Repl Summary view](#). Entries in the table's first column can be double-clicked to display a filtered runtime view of the selected item for additional details.

- **Agents** - The bottom section displays information about Agents. It also contains a link that opens the [Agent Summary view](#).

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

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



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](#).

Job Alerts View

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

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

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

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

The screenshot shows the Peer Management Center interface. The main window displays a table of job alerts. The table has the following columns: Name, Overall Status, Job Type, Failed Hosts, Quarantined, Retries, Errors, Warnings, Open File, Pending I/O, Queued I/O, Background Scan Status, Elapsed Time, Session Structure, and Synchronization Priority. The table lists several jobs, including FS-1, FC-2, FR-3, FS-2, FR-2, FR-1, and FC-1, with their respective statuses and details.

Below the main table, there is a section for 'Alerts' and 'Job Alerts'. The 'Job Alerts' section shows a list of alerts with columns: Received Date, Severity, Type, Name, Host, Message, and Exception. The alerts are filtered by Host, Name, Severity, and Type. A 'Clear Alerts' link is visible in the top right corner of the alerts section.

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

Displaying Job Alerts

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

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

Jobs View

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


The screenshot shows the Peer Management Center interface. The left sidebar contains a tree view of job types: Cloud Backup and Replication (1), DFS-N Management (1), VPeerTest Local Tickets, File Collaboration (2), File Replication (3), and File Synchronization (2). The main pane displays a table of jobs with columns: Name, Overall Status, Job Type, Failed, Ho, Quarant, Retries, Errors, Warnings, Open File, Pending, Queued, Background, Scan Status, Elapsed Time, Session Structure, and Synchronization. Below the table, there are summary statistics: Active Jobs -> Failed Participants: 0 of 4, Bytes Pending: 0 bytes, Bytes Transferred: 2.9 GB, Opens: 0, Initial Scans Completed: 4 of 7, Total Size: 10.19 GB, Total Files: 218, Total Directories: 5. At the bottom, the Job Alerts view is open, showing a table of alerts with columns: Received Date, Severity, Type, Name, Host, Message, and Exception.

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.

Runtime Views

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.


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 **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.

- 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.

Summary

Session

Event Log

Quarantines (0)

Retries (0)

Alerts (0)

Participants (4)

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	File Size	Last Modified	Date Opened	Message	Attribut

Status: Running

Copyright (c) 1993-2026 Peer Software, Inc. All Rights Reserved.

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

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.

Summary Session Event Log Quarantines (0) Retries (0) Alerts (0) Participants (4) Configuration											
Event Log (Auto-Update Disabled)											
0 errors, 0 warnings, 14 others Filter by Severity: Filter by: Actions Auto-Update Refresh 10 seconds Display 1500 Events											
Date	Severity	Type	Host	Is Source File	Comments	Message	Username	File Size	Modified Time	Client IP	
07-15-2024 18:18:30	INFO	Scan Complete		true		Scan completed for path "\", Scan Type					
07-15-2024 18:18:25	INFO	Scan Start		true		Scan Type: Full Directory					
07-15-2024 18:18:08	INFO	Watch Directory	Lab24Win22a	true							
07-15-2024 18:18:08	INFO	Watch Directory	Lab24Win22b	true							
07-15-2024 18:18:08	INFO	Watch Directory	Lab24Win19b	true							
07-15-2024 18:18:08	INFO	Watch Directory	Lab24Win19c	true							
07-15-2024 18:18:06	INFO	Install File Dispo	Lab24Win22a	true							
07-15-2024 18:18:06	INFO	Install File Dispo	Lab24Win19b	true							
07-15-2024 18:18:06	INFO	Install File Dispo	Lab24Win19c	true							
07-15-2024 18:18:06	INFO	Install File Dispo	Lab24Win22b	true							
07-15-2024 18:18:04	INFO	Job Started	Lab24Win22a	true							
07-15-2024 18:18:04	INFO	Job Started	Lab24Win22b	true							
07-15-2024 18:18:04	INFO	Job Started	Lab24Win19c	true							
07-15-2024 18:18:04	INFO	Job Started	Lab24Win19b	true							
Status: Running											

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 screenshot shows the 'Retries' tab in the Peer Global File Service interface. The tab header includes 'Summary', 'Session', 'Event Log', 'Quarantines (0)', 'Retries (0)', 'Alerts (0)', 'Participants (4)', and 'Configuration'. Below the header, there is a section for 'Retries' with a status of '0 Files'. A filter bar allows filtering by file name and selecting a host with the correct file version. A 'Release Conflict' button is also present. The main area is a table with the following columns: Date, File, Cause, Status, and Message. The table is currently empty. At the bottom, a status bar indicates 'Status: Running'.

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.



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](#)

Summary | Session | Event Log | Quarantines (0) | Retries (0) | Alerts (0) | **Participants (4)** | Configuration

Host Participants

Host	Root Path	Status	State	Message	Edge Caching Role	Total Disk Space	Free Disk Space	Trash Bin Size
Lab24Win19b	C:\FC-1 Folder A	Participating	Active		Master	49.40 GB	16.88 GB	0 bytes
Lab24Win19c	C:\FC-1 Folder B	Participating	Active		Edge	49.40 GB	27.41 GB	0 bytes
Lab24Win22a	C:\FC-1 Folder C	Participating	Active		Edge	99.32 GB	67.76 GB	0 bytes
Lab24Win22b	C:\FC-1 Folder D	Participating	Active		Master	99.32 GB	73.56 GB	0 bytes

Host Participant State Change Log

Filter by : Host: Status: State:

Date	Host	Status	State	Message
07-16-2024	Lab24Win19	Participating	Active	
07-16-2024	Lab24Win19	Participating	Active	
07-16-2024	Lab24Win22	Participating	Active	
07-16-2024	Lab24Win22	Participating	Active	
07-16-2024	Lab24Win19	Participating	Active	

Status: ■ Running

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 displays the 'Configuration' tab of the Peer Global File Service interface. The top navigation bar includes links for Summary, Session, Event Log, Quarantines (0), Retries (0), Alerts (0), Participants (4), and Configuration. Below the navigation bar, there is a link to 'Edit this Configuration'.

The main content area is divided into several sections:

- Currently Running Configuration Summary**
 - Selected Participants and Configurations**: Lists participants and their configurations, including Lab24Win19b, Lab24Win19c, Lab24Win22a, and Lab24Win22b.
 - Edge Caching Configuration**: Shows settings for agents (Lab24Win19b, Lab24Win22b, Lab24Win19c) and their roles (Master, Edge). It also includes volume policies, utilization policies, dynamic rules, pinning filters, and masters.
 - General Settings**: Displays job details such as Job Name (FC-1), Job ID (103), Job Type (File Collaboration), Transfer Block Size (2048 KB), Verify Checksum (true), and Verify Full File Checksum (true).
- Selected File Filters**
 - Windows Default (General) - 0**: Shows excluded and included patterns, date filter (Include all dates), and size filter (None).
 - File Collaboration Sync Only (Synchronization Only) - 4**: Shows excluded and included patterns, date filter (Include all dates), and size filter (None).
- Selected Scheduled Replication Filters**: Indicates that no scheduled replication filters are selected.
- Locking**: Lists various locking and synchronization settings, such as Allow Write Access During Synchronization (true), Exclusive Target Lock (false), and Max File Size (MB) (512).
- Conflict Resolution**: Shows settings for latest modified time, quarantine multi-edit conflicts, and offline folder rename detection.


At the bottom left, a status bar indicates the job is 'Running'.

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.



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 Global File Service runtime view for a File Synchronization job. The interface includes a left sidebar with a tree view of jobs, a main pane with tabs for Summary, Session, Event Log, Quarantines, Retries, Alerts, Participants, and Configuration. The Summary tab is active, showing the following details:

- Job Status:** Stopped
- Start Time:** 07:00 PM
- Elapsed Time:** 1
- Watch Set:**
 - Total Files: 18
 - Total Bytes: 934 MB
 - Total Folders: 1
- Activity:**
 - Active Opens: 0
 - Files Pending: 0
 - Bytes Pending: 0 bytes
 - Metadata Pending: 0
 - File Quarantines: 0
 - Renames Pending: 0
 - Deletes Pending: 0
 - File Retries: 0
- Replication Status:**
 - Bytes Transferred: 0 bytes
 - Delta-level Savings: 0

Below the summary, there is a section for Alerts, showing a list of alerts with columns for Received Date, Severity, Type, Name, Host, Message, and Exception. The alerts list shows several entries, including "User Started Peerlet" and "User Stopped Job".

Summary Views

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

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

The screenshot displays the Peer Management Center interface. The left sidebar shows a tree view of jobs and agents. The main pane is titled 'Collab, Sync, and Repl Summary' and contains a table of job details. The bottom pane shows an alerts log.

Name	Overall Status	Job Type	Failed	Ho	Quaranti	Retries	Errors	Warnings	Open	Pending	Queued	Backgrou	Scan	Status	Elapsed Time	Session Structure	Synchronization	Priority
FS-1	⚙️ Halted	File Synchronization	0	0	0	0	0	0	0	0	0	0	0	Stopped	00:00:00	Size: 934 MB, Files: 2		
FC-2	⚙️ Halted	File Collaboration	0	0	0	1	0	0	0	0	0	0	0	Stopped	00:00:00	Size: 2 GB, Files: 4, 2		
FR-3	⚙️ Halted	File Replication	0	0	0	1	0	0	0	0	0	0	0	Stopped	00:00:00	Size: 929.01 MB, F: 2		
FS-2	🟢 Running	File Synchronization	0	0	0	0	0	0	0	0	0	0	0	Completed - 00:00:02	68:01:45	Size: 1.95 GB, File: 2		
FR-2	🟢 Running	File Replication	0	0	0	0	0	0	0	0	0	0	0	Completed - 00:00:03	68:01:45	Size: 649 MB, Files: 2		
FR-1	🟢 Running	File Replication	0	0	0	0	0	0	0	0	0	0	0	Completed - 00:00:03	68:01:45	Size: 1.77 GB, File: 2		
FC-1	🟢 Running	File Collaboration	0	0	0	0	0	0	0	0	0	0	0	Completed - 00:00:04	68:01:56	Size: 2.02 GB, File: 2		

Active Jobs -> Failed Participants: 0 of 4 | Bytes Pending: 0 bytes | Bytes Transferred: 2.9 GB | Opens: 0 | Initial Scans Completed: 4 of 7 | Total Size: 10.19 GB | Total Files: 218 | Total Directories: 5

Alerts

0 errors, 14 warnings, 87 others | Filter by: Host: | Severity: | Type: | Clear Alerts

Received Date	Severity	Type	Name	Host	Message	Exception
07-18-2024 00:59	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Maintenance	
07-17-2024 23:58	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Nightly Delete Shadow Copies Job	
07-17-2024 23:29	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Database Maintenance	
07-17-2024 22:36	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Process Retention Policy	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\ on host Lab24Win19c completed elapsed tim	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Free space on volume C:\ is less than the minimum required size to achieve optin	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\ on host Lab24Win19c has started.	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a completed elapsed tin	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Free space on volume C:\ is less than the minimum required size to achieve optin	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a has started.	


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.



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 displays the 'Summary' tab in the Peer Global File Service interface. The main window is titled 'Summary | Edge Caching | Reports' and contains a table of job status information. The table has columns for Name, Overall Status, Job Type, Failed, Ho, Quarant, Retries, Errors, Warnings, Open File, Pending I, Queued I, Backgrou, Scan Status, Elapsed Time, Session Structure, and Synchronization Priority. The table lists several jobs, including File Synchronization, File Collaboration, and File Replication, with their respective statuses (Halted, Running) and progress details.

Below the table, there is a summary bar showing: Active Jobs -> Failed Participants: 0 of 4 | Bytes Pending: 0 bytes | Bytes Transferred: 2.9 GB | Opens: 0 | Initial Scans Completed: 4 of 7 | Total Size: 10.19 GB | Total Files: 218 | Total Directories: 5.

At the bottom, the 'Alerts' section shows 0 errors, 14 warnings, and 87 others. It includes a table of alerts with columns for Received Date, Severity, Type, Name, Host, Message, and Exception. The alerts list various events such as 'Scheduled Task Started', 'Started Blob Tiering Maintenance', and 'Edge Caching Update'.

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.

The screenshot shows the 'Colab, Sync, and Repl Summary' window. The table displays the following data:

Name	Overall Status	Job Type	Failed	Ho	Quarant	Retries	Errors	Warnings	Open	File	Pending	I	Queued	I	Background	Scan	Status	Elapsed Time	Session Structure	Synchronization Priority
FS-1	Halted	File Synchronization	0	0	0	0	0	0	0	0	0	0	0	0	0	0	Stopped	00:00:00	Size: 934 MB, Files: 2	
FC-2	Halted	File Collaboration	0	0	0	0	1	0	0	0	0	0	0	0	0	0	Stopped	00:00:00	Size: 2 GB, Files: 4	2
FR-3	Halted	File Replication	0	0	0	0	1	0	0	0	0	0	0	0	0	0	Stopped	00:00:00	Size: 929.01 MB, F: 2	
FS-2	Running	File Synchronization	0	0	0	0	0	0	0	0	0	0	0	0	0	0	Completed - 00:00:02	68:04:12	Size: 1.95 GB, File: 2	
FR-2	Running	File Replication	0	0	0	0	0	0	0	0	0	0	0	0	0	0	Completed - 00:00:03	68:04:12	Size: 649 MB, Files: 2	
FR-1	Running	File Replication	0	0	0	0	0	0	0	0	0	0	0	0	0	0	Completed - 00:00:03	68:04:12	Size: 1.77 GB, File: 2	
FC-1	Running	File Collaboration	0	0	0	0	0	0	0	0	0	0	0	0	0	0	Completed - 00:00:04	68:04:23	Size: 2.02 GB, File: 2	

The status bar at the bottom shows: Active Jobs -> Failed Participants: 0 of 4 | Bytes Pending: 0 bytes | Bytes Transferred: 2.9 GB | Opens: 0 | Initial Scans Completed: 4 of 7 | Total Size: 10.19 GB | Total Files: 218 | Total Directories: 5

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 Global File Service interface. The **Edge Caching** tab is selected, displaying a table of jobs. The table has the following columns: Agent/Volumes, Last Update, Temporary Storage Path, Edge Service Status, Volume Utilization (Total/Free), Cache Utilization (Total/Free), Local Files, Pinned Local Files, Stubbied Files, Pinned Stubbied Files, and Scan State. Two jobs are listed: one for Lab24Win19c and one for Lab24Win22a. Below the table is a **Job Details** section and an **Alerts** section showing a list of events with columns for Received Date, Severity, Type, Name, Host, Message, and Exception.

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 shows the 'Collab, Sync, and Repl Summary' window in the Peer Global File Service interface. The window has a sidebar on the left with a tree view of the system hierarchy, including 'Cloud Backup and Replication (1)', 'DFS-N Management (1)', 'File Collaboration (2)', 'File Replication (3)', and 'File Synchronization (2)'. The main area displays a table with the following columns: Name, File Sync Queue, Real-Time Queue, Queued Bytes, Mods, Adds, Metadata, Scan Queue, Deletes, Renames, Event Queue, and Scheduled Replication. The table lists various jobs and their current status. Below the table, there's a section for 'Alerts' showing a list of events with columns: Received Date, Severity, Type, Name, Host, Message, and Exception.

Name	File Sync Queue	Real-Time Queue	Queued Bytes	Mods	Adds	Metadata	Scan Queue	Deletes	Renames	Event Queue	Scheduled Replication
FC-1	0	0	0 bytes	0	0	0	0	0	0	0	0
FS-1	0	0	0 bytes	0	0	0	0	0	0	0	0
FC-2	0	0	0 bytes	0	0	0	0	0	0	0	0
FR-1	0	0	0 bytes	0	0	0	0	0	0	0	0
FR-2	0	0	0 bytes	0	0	0	0	0	0	0	0
FS-2	0	0	0 bytes	0	0	0	0	0	0	0	0
FR-3	0	0	0 bytes	0	0	0	0	0	0	0	0
Totals	0	0	0 bytes	0	0	0	0	0	0	0	0

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 Transfers	The number of events that are queued awaiting an available replication slot.

Actions Menu

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

Option	Description
Filters	Allows you to select predefined or user-defined filters and to save/manage list filters . Default job filters include Failed Jobs, Jobs with Backlog, and Running Scans.
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.
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 displays the Peer Management Center interface. The top menu bar includes 'File', 'Window', 'Tools', and 'Help'. The sidebar on the left contains a 'Jobs' section with a tree view showing 'Cloud Backup and Replication (1)', 'DFS-N Management (1)', 'File Collaboration (2)', 'File Replication (3)', and 'File Synchronization (2)'. Below this is an 'Agents' section with a list of agents: 'Lab24Win19b (Connected)', 'Lab24Win19c (Connected)', 'Lab24Win22a (Connected)', and 'Lab24Win22b (Connected)'. The main pane is titled 'Namespace Summary' and contains a table with the following columns: 'Namespace Path/Folders', 'Management Status', 'State', 'Errors', 'Servers', and 'Total Folders/Targets'. The table lists two namespaces: '\\PeerTest.local\\Tickets' and '\\PeerTest.local\\Bug Reports'. The bottom pane shows a log of alerts with columns: 'Received Date', 'Severity', 'Type', 'Name', 'Host', 'Message', and 'Exception'. The log shows various events such as 'Scheduled Task Started', 'Edge Caching Update', and 'Caching Hydrate scan on volume C:\\'.

Namespace Path/Folders	Management Status	State	Errors	Servers	Total Folders/Targets
\\PeerTest.local\\Tickets	Monitoring	Online	0	lab24win19a, peerTest.local	1 Folder
\\PeerTest.local\\Bug Reports					3 Targets

Received Date	Severity	Type	Name	Host	Message	Exception
07-18-2024 00:59	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Maintenance	
07-17-2024 23:58	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Nightly Delete Shadow Copies Job	
07-17-2024 23:29	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Database Maintenance	
07-17-2024 22:36	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Process Retention Policy	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\\ on host Lab24Win19c completed elapsed time	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Free space on volume C:\\ is less than the minimum required size to achieve optimal	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\\ on host Lab24Win19c has started.	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\\ on host Lab24Win22a completed elapsed time	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Free space on volume C:\\ is less than the minimum required size to achieve optimal	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\\ on host Lab24Win22a has started.	

Topology Perspective


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

In the PMC web client, the Topology perspective is available only to users assigned the **Administrator** or **Power User** role.

The **Topology** perspective contains the following 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 displays a hierarchical tree structure containing the configured sites, brokers, and connected Agents.
Details	Provides information about brokers configured on a selected site from the Sites view. This display shows broker configuration settings, detailed status information, and Agents configured to connect to it.
Map	Provides a visual display of the configured sites 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.

Topology Display Controls

The **NAT**, **TLS**, and **TCP** buttons appear in Details view and control the connection information displayed within the topology diagram.

Button	Description
NAT	Enables or disables the display of Network Address Translation (NAT) connectivity information in the topology display.
TLS	Enables or disables the display of Transport Layer Security (TLS) communication paths between brokers and Agents.
TCP	Enables or disables the display of Transmission Control Protocol (TCP) communication paths and connection status information.

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.

Collab, Sync, and Repl Summary														
Summary Edge Caching Reports														
Runtime Summary (auto-update enabled)														
Filter by: Expression Actions Auto-Update Refresh 10 seconds														
Name	Overall Status	Job Type	Failed Ho	Quaranti	Retries	Errors	Warnings	Open File	Pending I	Queued I	Backgrou	Scan Status	Elapsed Time	Session Structure Synchronization Priority
FC-1	Running	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:04	70:54:53	Size: 2.02 GB, File: 2
FR-1	Running	File Replication	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:03	70:54:42	Size: 1.77 GB, File: 2
FR-2	Running	File Replication	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:03	70:54:42	Size: 649 MB, Files: 2
FS-2	Running	File Synchronizatio	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:02	70:54:41	Size: 1.95 GB, File: 2
FS-1	Halted.	File Synchronizatio	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 934 MB, Files: 2
FC-2	Halted.	File Collaboration	0	0	0	1	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 2 GB, Files: 4 2
FR-3	Halted.	File Replication	0	0	0	1	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 929.01 MB, F 2
4														
Active Jobs -> Failed Participants: 0 of 4 Bytes Pending: 0 bytes Bytes Transferred: 2.9 GB Opens: 0 Initial Scans Completed: 4 of 7 Total Size: 10.19 GB Total Files: 218 Total Directories: 5														
Alerts														
0 errors, 14 warnings, 87 others Filter by: Host: Severity: Type: Clear Alerts														
Received Date	Severity	Type	Name	Host	Message	Exception								
07-18-2024 00:59	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Maintenance									
07-17-2024 23:58	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Nightly Delete Shadow Copies Job									
07-17-2024 23:29	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Blob Tiering Database Maintenance									
07-17-2024 22:36	Info	Agent	Scheduled Task Started	Lab24Win19c	Started Process Retention Policy									
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\ on host Lab24Win19c completed elapsed time									
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Free space on volume C:\ is less than the minimum required size to achieve opti									
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\ on host Lab24Win19c has started.									
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a completed elapsed time									
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Free space on volume C:\ is less than the minimum required size to achieve opti									
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a has started.									

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', 'Agent Summary', 'Cloud Summary', and 'Namespace Summary'. Below these, a 'Runtime Summary (auto-update enabled)' section shows a filter by 'Expression' and an 'Actions' dropdown. A table lists various jobs with columns for Name, Overall Status, Job Type, Failed Hosts, Quarantines, Retries, Errors, Warnings, Open Files, Pending, Queued, Background Scans, Status, Elapsed Time, Session Structure, and Synchronization Priority. A context menu is open over the first row (FR-1), showing options like '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', 'Run Event Detection Analytics...', and a list of columns to display. Below the table, a summary bar shows 'Bytes Transferred: 2.9 GB', 'Opens: 0', 'Initial Scans Completed: 4 of 7', 'Total Size: 10.19 GB', 'Total Files: 218', and 'Total Directories: 5'. At the bottom, an 'Alerts' section shows '0 errors, 14 warnings, 82' and a list of alerts with columns for Received Date, Severity, Agent, Job Type, Host, Message, and Exception.

Name	Overall Status	Job Type	Failed Hosts	Quarantines	Retries	Errors	Warnings	Open Files	Pending	Queued	Background Scans	Status	Elapsed Time	Session Structure	Synchronization Priority
FR-1	Running	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:04	70:55:03	Size: 2.02 GB, File: 2	
FR-2	Running	File Replication	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:03	70:54:52	Size: 1.77 GB, File: 2	
FS-2	Running	File Sync	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:03	70:54:52	Size: 649 MB, Files: 2	
FS-1	Running	File Sync	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:02	70:54:51	Size: 1.95 GB, File: 2	
FC-2	Running	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 934 MB, Files: 2	
FR-3	Running	File Replication	0	0	0	1	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 2 GB, Files: 4	2
			0	0	0	1	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 929.01 MB, F: 2	

Received Date	Severity	Agent	Job Type	Host	Message	Exception
07-18-2024 00:59	Info	Agent	Edge Caching Update	Lab24Win19c	Started Blob Tiering Maintenance	
07-17-2024 23:58	Info	Agent	Edge Caching Update	Lab24Win19c	Started Nightly Delete Shadow Copies Job	
07-17-2024 23:29	Info	Agent	Edge Caching Update	Lab24Win19c	Started Blob Tiering Database Maintenance	
07-17-2024 22:36	Info	Agent	Edge Caching Update	Lab24Win19c	Started Process Retention Policy	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Caching Hydrate scan on volume C:\ on host Lab24Win19c completed elapsed time	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win19c	Free space on volume C:\ is less than the minimum required size to achieve optimal	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a completed elapsed time	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Free space on volume C:\ is less than the minimum required size to achieve optimal	
07-17-2024 22:00	Info	Agent	Edge Caching Update	Lab24Win22a	Caching Hydrate scan on volume C:\ on host Lab24Win22a has started.	

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.

Collab, Sync, and Repl Summary | Agent Summary | Cloud Summary | Namespace Summary

Summary | Edge Caching | Reports

Runtime Summary (auto-update enabled)

Filter by: Expression Actions Auto-Update Refresh 10 seconds

Name	Overall Status	Job Type	Failed Ho	Quarant	Retries	Errors	Warnings	Open File	Pending	Queued I	Backgrou	Scan Status	Elapsed Time	Session Structure	Synchronization Priority
FC-1	Running	File Collaboration	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:04	70:57:29	Size: 2.02 GB, File: 2	
FR-1	Running	File Replication	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:03	70:57:18	Size: 1.77 GB, File: 2	
FR-2	Running	File Replication	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:03	70:57:18	Size: 649 MB, File: 2	
FS-2	Running	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Completed - 00:00:02	70:57:18	Size: 1.95 GB, File: 2	
FS-1	Halted	File Synchronization	0	0	0	0	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 934 MB, File: 2	
FC-2	Halted	File Collaboration	0	0	0	1	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 2 GB, File: 4 2	
FR-3	Halted	File Replication	0	0	0	1	0	0	0 bytes	0	0	Stopped	00:00:00	Size: 929.01 MB, F 2	

PMC Alert Details

Received at: 07-15-2024 18:18:13

Severity: Info

Category: Start Job

Locally Generated at: 07-15-2024 18:18:13

Name: FR-1

Message: User Started Peerlet

Ref ID: 106

Click outside of popup to close

Active Jobs -> Failed Participants: 0 of 4 | Bytes Pending: 0 bytes | Bytes Transferred: 2.9 GB | Opens: 0 | Initial Scans Completed: 4 of 7 | Total Size: 10.19 GB | Total Files: 218 | Total Directories: 5

Alerts Task History

3 errors, 3 warnings, 36 others | Filter by: Host: Name: Severity: Type: Clear Alerts

Received Date	Severity	Type	Name	Host	Message	Exception
07-15-2024 18:18	Info	Start Job	FS-2		User Started Peerlet	
07-15-2024 18:18	Info	Start Job	FR-2		User Started Peerlet	
07-15-2024 18:18	Info	Start Job	FR-1		User Started Peerlet	
07-15-2024 18:18	Info	Start Job	FC-1		User Started Peerlet	
07-15-2024 18:16	Info	Stop Job	FS-2		User Stopped Job	
07-15-2024 18:16	Info	Stop Job	FS-1		User Stopped Job	
07-15-2024 18:16	Info	Stop Job	FR-3		User Stopped Job	
07-15-2024 18:16	Info	Stop Job	FR-2		User Stopped Job	
07-15-2024 18:16	Info	Stop Job	FR-1		User Stopped Job	
07-15-2024 18:16	Info	Stop Job	FC-2		User Stopped Job	

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](#)
- [Web Client Users](#)

Email Alerts

Overview

Email alerts notify recipients when selected events occur (for example: file quarantined, session aborted, host failure, or system alert). Each email alert includes a unique name, one or more event types (which vary by job type), and a list of recipient email addresses.

You can apply multiple email alerts to a job, and an alert can be reused across multiple jobs of the same type. When you create a job, you can select an existing email alert or create a new one and apply it during job setup.

Requirements and Constraints

- SMTP must be configured before alerts can be sent. See [Email Configuration](#).
- 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.
- Email alerts are defined in the [preferences](#) for each job type.

Managing Email Alerts

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

To manage email alerts:

1. Select **Open Preferences** from the **Tools** 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. For details on creating, editing, or deleting alerts, see the job type-specific preferences topics.

File and Folder Filters

A **file filter** specifies which files and folders are included or excluded from a job's [watch set](#). **Included** files are scanned and monitored for real-time events, while **excluded** files are ignored.

By default, all files are included unless they match [predefined file filters](#) or [automatically excluded file types](#).

Filters can also apply to folders, enabling you to include or exclude entire directories from a job's watch set. For more information, see [Folder Filters](#).

A file filter consists of a unique name and one or more [filter patterns](#). You can also define filters based on [file size](#) or [last modified time](#). For details on creating filter patterns, see [Defining Filter Patterns](#). To define filter patterns that apply to 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** - Available for File Collaboration jobs. Use this type to exclude specific file types from being locked when a file open is detected on a participant in a File Collaboration job.
- **Locking Only** - Available for File Collaboration jobs. Use this type to exclude synchronization across the entire job, allowing only opens and close events to be detected and acted on without any synchronization .

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

Creating and Applying File Filters

You create a file filter in the **File and Folders Filters** page under [Preferences](#) for a specific job type. Once created, the filter can be applied to any job 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 to user-defined filters, some [predefined filters](#) are automatically applied to certain job types. These predefined filters help ensure that common temporary or system files are excluded without requiring manual configuration

For instructions on creating file filters, 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](#)

Editing and Deleting File Filters

You can edit or delete file filters from the **File and Folders Filters** page under [Preferences](#) for the corresponding job type. Changes to a file filter affect only jobs to which that filter is applied.


You cannot edit or delete a file filter while it is applied to a running job. Stop or complete the job before making any changes to the filter.

Deleting a file filter removes it from the list of available filters for that job type. If the filter is currently applied to any jobs, you must first remove it from those jobs before deleting it.

Predefined filters cannot be edited or deleted.

Predefined File Filters

In addition to defining your own file filters, Peer Management Center includes **predefined** file filters that can be applied to jobs. These filters simplify configuration and help ensure that commonly excluded files are automatically ignored. The available predefined filters vary by job type.



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

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


Updating Predefined File Filters

You can update predefined file filters from the **File and Folders Filters** page under [Preferences](#) for the corresponding job type. Updating ensures that the latest filter definitions are installed and available for use with applicable jobs.


To update a predefined file filter:

1. From the **Tools** menu, select **Preferences**.
2. In the navigation tree, expand **Cloud Backup and Replication** or **Collab, Sync, and Repl Summary**, and then select **File and Folder Filters**.


The existing file filters appear in the **File and Folder Filters** table.



3. Select the predefined filter to update, and then click **Upgrade**.
 - If no updates are available, click **OK** to close the message that appears.



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




Predefined filters cannot be renamed or have their type changed.

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 in the watch set. For example, in the following file filter definition, files with names ending in .dotx are excluded and files with names ending in .docx are included.

The following example shows the Excluded Patterns and Included Patterns sections of the Create File Filter dialog:



The "Create File Filter" dialog box is shown with a blue title bar. It contains several sections for configuring file filters. At the top, there is a "Name:" text field and a "Filter Type:" dropdown menu set to "General". Below this is an "Auto Excluded" section with a link to "View file types that are automatically excluded". The "Excluded Patterns" section contains a list box with "*.bak" and "*.tmp", and buttons for "Add", "Edit", and "Delete". The "Included Patterns" section contains a list box with "*.docx" and buttons for "Add", "Edit", and "Delete". At the bottom, there are two sections: "Included Last Modified Dates" with a dropdown set to "Include all dates" and a text field for "0 days", and "Excluded File Sizes" with a dropdown set to "None" and a text field for "0 bytes". "OK" and "Cancel" buttons are at the bottom right.

Create File Filter

Name:

Filter Type:

Auto Excluded

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

Excluded Patterns

- *.bak
- *.tmp

Add Edit Delete

Included Patterns

- *.docx

Add Edit Delete

Included Last Modified Dates

days

Excluded File Sizes

bytes

OK Cancel

Using Wildcards in Filter Patterns

Excluded and included file name filters can contain one or more wildcards. When file filter patterns are specified, a file matches the filter if it matches any one of the expressions.

You can use the following wildcards in a file filter pattern to match file extensions or file names that follow a pattern.

Wildcard	Description
*	Matches zero or more characters.
?	Matches exactly one character.

The following examples show how wildcards can be used:

- *.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**

Automatically Excluded File Types

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

Excluding Temporary Files

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.

A screenshot of the 'Add Exclusion Pattern' dialog box. The dialog has a blue title bar with the text 'Add Exclusion Pattern' and a close button (X) on the right. Inside the dialog, there is a text input field with the placeholder text 'Enter a single file name or a single path pattern:'. Below the input field, there is a message 'String cannot be empty.' in a light gray box. At the bottom right of the dialog, there are two buttons: 'OK' and 'Cancel'.

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\$

Add Edit Delete

Included Patterns

Add Edit Delete

Included Last Modified Dates

 days

Excluded File Sizes

 bytes

OK Cancel

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>.

Filtering on Last Modified Date

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

- Peer Management Center 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

Create File Filter

Name:

Filter Type:

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

Excluded Patterns

Included Patterns

Included Last Modified Dates ☐ days

Excluded File Sizes ☐ bytes

Field	Description
Include all dates	This is the default option and will include all files regardless of last modified date.

Field	Description
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.

Filtering on File Size

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 screenshot shows the 'Create File Filter' dialog box. It has a title bar with 'Create File Filter' and window control buttons. The dialog is divided into several sections:

- Name:** A text input field.
- Filter Type:** A dropdown menu currently set to 'General'.
- Auto Excluded:** A section with a link: [View file types that are automatically excluded](#).
- Excluded Patterns:** A large text area for patterns, with 'Add', 'Edit', and 'Delete' buttons below it.
- Included Patterns:** A large text area for patterns, 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 set to '0 days'.
- Excluded File Sizes:** This section is highlighted with a red rectangle. It contains a dropdown menu set to 'Exclude files less than' and a text input field set to '0 bytes'.

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.

Field	Description
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

Overview

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

Folder Filter Syntax

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

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

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

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 Examples

Folder filters allow you to include or exclude folders from a job based on folder names and paths. The examples in this topic show common folder filter patterns, describe important limitations, and illustrate how folder filtering can be used to reduce the number of jobs required.

Folder Filters Pattern Examples

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

Folder Filter Limitations

Note the following regarding regular expressions:

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

Example of a Simple Folder Filter

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

		Server 1		Server 2	
		Drive D	Drive E	Drive D	Drive F
Old	Job 1	D:\General		D:\General	

Jobs	Job 2		E:\Common		F:\Common
	Job 3	D:\Projects		D:\Projects	
	Job 4		E:\Documents		F:\Documents

After consolidation:

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

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

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

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 **list filters** throughout the interface to help you refine and control the information displayed in lists. List filters make it easier to locate jobs, Agents, and specific data within summary reports, especially in environments with a large amount of information.

A list filter is defined by a **filter expression**, which specifies the criteria used to include or exclude items from a list. The displayed results reflect the criteria defined in the expression. For example, a filter expression might be used to display only jobs with a specific status or only Agents that meet a defined condition.

Filter expressions can range from simple keyword-based searches to more complex expressions that combine multiple conditions. You can save list filters and reuse them later, which is useful when you regularly apply the same criteria or work with complex filter expressions. Saved filters help ensure consistency when reviewing or comparing data. For more information about saving list filters, see [Managing List Filters](#).

Peer Management Center also supports operators and labels that allow you to construct more precise filter expressions.

For more information about list filters, see:

- [Creating Basic List Filters](#)
- [Creating Complex List Filters](#)
- [Predefined List Filters](#)
- [Managing List Filters](#)

Creating Basic List Filters

Basic list filters contain simple words or phrases that narrow the items displayed in a list. When you enter a word in the **filter field**, Peer Management Center displays only the items that contain that word and hides all others.

A basic filter searches item names, tag names, and tag categories. For example, entering the word *sales* matches all items that contain *sales* in their name, tag name, or tag category. Agent attribute fields are not included in these generic searches.

To search for an exact word or a phrase that contains spaces, enclose the search string in double quotation marks. Quoted strings are matched exactly.

For example:

- To search for the phrase *North America*, enter "North America".
- To search for the word *agent* only (and exclude results such as *USAgent* or *Agent2015*), enter "agent".

For information about creating more advanced filter expressions using operators and labels, see [Creating Complex List Filters](#).

Creating Complex List Filters

Complex list filters allow you to refine results by combining multiple conditions in a single filter. You create complex list filters by entering filter expressions that include **operators** and **labels** in the **filter field**.

Using Operators

Operators allow you to combine multiple simple expressions terms into a single compound expression. This lets you refine results by specifying how terms are evaluated together.

The following operators are supported:

Operator	Description
AND	Matches items that meet all specified conditions.
OR	Matches items that meet any of the specified conditions.
NOT	Excludes items that match the specified conditions.

For example, the following filter expression matches only Agents that contain *Americas* in their tag name and also contain *sales* in their name, tag name, or tag category:

```
tag:Americas AND sales
```

You can use parentheses to group terms and control how expressions are evaluated, which is useful when building more complex filters.

Using Labels

Labels specify which field the search string is evaluated against. To use a label, enter it in the following format:

`<label>:<search string>`

The following labels are supported:

L a b e l	Description	Example
name	Matches the specified string in the item name.	name:"Design Data"
tag	Matches the specified string in the tag names.	tag:Americas
cat	Matches the specified string in the category names.	cat:Design
host	Matches jobs that include the specified host in the list of job participants.	host:WIN12R2A
attr	Matches the specified string in the following Agent attribute fields: <i>Connection Status, Operating System, JVM Architecture, and Agent Version</i>	attr:x86
filter	Matches items that have been assigned a default or user-created filter.	

Examples

The following examples illustrate complex filter expressions that combine keywords, labels, and operators.

Purpose	Filter Expressions
Match Agents that contain the word <i>Sales</i> in their name, tag name, or tag category.	Sales
Match Agents with a tag that contains <i>North America</i> in the tag name and <i>Location</i> in the tag category.	cat:Location AND tag:"North America"
Match Agents that contain <i>Sales</i> in their name, tag name, or tag category and also have a tag with <i>North America</i> in the tag name and <i>Location</i> in the tag category.	Sales AND (cat:Location AND tag:"North America")

Predefined List Filters

Peer Management Center includes **predefined list filters** that you can apply from the **filter field** to quickly narrow the items displayed in a list.

- Default job filters include **Failed Jobs**, **Jobs with Backlog**, and **Running Scans**.
- Default Agent filters include **Connected** and **Disconnected**.

To view and select available predefined list filters, press **Ctrl + Space** in the filter field. The available filters are listed and can be selected to refine the results.

Managing List Filters


Managing list filters allows you to save filter expressions for reuse, apply previously saved filters, and clear filters to restore the full list of items. These actions help you work more efficiently when you frequently use the same filter criteria.

Saving List Filters

Saving a list filter allows you to reuse the same filter criteria without re-entering the filter expression. Saved list filters are associated with the view in which they are created and can be reused only within that view.


To save a list filter:

1. Enter a filter expression in the filter field.
2. Click the **Manage, Save, and Load Filters** button.
3. Select **Save filter** from the menu
4. In the dialog that opens, enter a name for the filter and confirm the filter expression and target view.



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

The **Manage, Save, and Load Filters** button is typically located above the filter field or in the **Actions** menu, depending on the view. This button is available in the [Jobs view](#) panel, the [Agent Summary view](#), and the [Summary tab](#) in the [Collab, Sync, and Repl Summary](#) view.



Loading (Applying) a Saved List Filter

Loading a saved list filter applies the saved filter expression to the current list. When a filter is loaded, the filter expression is populated in the filter field and the list updates to reflect the saved criteria. Loading a filter does not modify the saved filter definition.

To apply a saved list filter:

1. Click the **Manage, Save, and Load Filters** button.
2. Select the name of a saved filter from the menu.

Editing a List Filter

Editing a list filter allows you to modify a saved, user-defined filter.

To edit a saved list filter:

1. Click the **Manage, Save, and Load Filters** button.
2. Select **Manage saved filters** from the menu.
3. In the dialog that opens, select a user-defined filter to edit.


Only user-defined list filters can be edited or deleted. Built-in (predefined) list filters are read-only.

After editing a filter, you must save it again to update the saved filter.

Clearing a List Filter

Clearing a list filter removes the active filter expression from the filter field and restores the full, unfiltered list of items. Clearing a filter does **not** delete any saved list filters.

To remove the active list filter, click the pencil icon to the right of the filter field:




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.

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
☐	Lab24Win19b
☐	Lab24Win19c
☐	Lab24Win22a
☐	Lab24Win22b

[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?

Yes

No

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 Lab24Win19a at the following location:

C:\Program Files\Peer Software\Peer Management Center\Hub\workspace\support\PL-Hub-Logs-6.0.0.20240621-Lab24Win19a-07-20-2024.16.22.44.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.

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](#).

Base Roles

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 base 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	

Functionality	Administrator	Power User	Help Desk
Folder Analyzer View	Edit	Edit	View-only
Job Alert View	Edit	Edit	
Job View	Edit	Edit	View-only
Log Dump Action	Edit	Edit	
Memory Dump Action	Edit	Edit	
New Job Action	Edit		
Participant View	Edit	Edit	
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	

Functionality	Administrator	Power User	Help Desk
Tag Resources Dialog	Edit		
Thread Dump Action	Edit	Edit	

Custom Roles

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.

New Role

General
Permissions
Tags

Permissions

Category	Name	Access
Cloud Backup and Replication UI	Job View - Alerts Tab	Full access
Cloud Backup and Replication UI	Job View - Configuration Tab	Full access
Cloud Backup and Replication UI	Job View - Event Log Tab	Full access
Cloud Backup and Replication UI	Job View - Failed Events Tab	Full access
Cloud Backup and Replication UI	Job View - Participants Tab	Full access
Cloud Backup and Replication UI	Job View - Summary Tab	Full access
Cloud Backup and Replication UI	Runtime Summary View	Full access
Collab/Sync/Repl UI	Job View - Alerts Tab	Full access
Collab/Sync/Repl UI	Job View - Configuration Tab	Full access
Collab/Sync/Repl UI	Job View - Event Log Tab	Full access
Collab/Sync/Repl UI	Job View - Participants Tab	Full access
Collab/Sync/Repl UI	Job View - Quarantines Tab	Full access
Collab/Sync/Repl UI	Job View - Retries Tab	Full access
Collab/Sync/Repl UI	Job View - Session Tab	Full access
Collab/Sync/Repl UI	Job View - Summary Tab	Full access
Collab/Sync/Repl UI	Runtime Summary View	Full access
DFS-N UI	Job View - Alerts Tab	Full access
DFS-N UI	Job View - Configuration Tab	Full access
DFS-N UI	Job View - Namespace Servers Tab	Full access
DFS-N UI	Job View - Namespace Tab	Full access
DFS-N UI	Runtime Summary View	Full access
PMC UI	Assign Tags	Full access
PMC UI	Download Broker Installer	Full access
PMC UI	Download Peer Agent Installer	Full access
PMC UI	Download PeerIQ	Full access
PMC UI	Generate PMC Memory Dump File	Full access
PMC UI	Generate PMC and Broker Thread Dump	Full access
PMC UI	Jobs View	Full access
PMC UI	New Job Wizard	Full access
PMC UI	Preferences	Full access
PMC UI	Retrieve Broker Statistics	Full access
PMC UI	Retrieve PMC/Agent Logs	Full access
PMC UI	Show Agent Summary	Full access
PMC UI	Show Progress View	Full access
PMC UI	Task Scheduler and Task History	Full access
PMC UI	Topology	Full access
PMC UI	View Alerts	Full access

OK Cancel

Advanced Topics

This section discusses the following topics:

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

Analytics

PeerGFS offers analytics capabilities for monitoring, analyzing, and optimizing PeerGFS environments and the storage infrastructure they manage. PeerGFS analytics includes **PeerIQ**, **File System Analytics**, and **File Activity Analytics**.

PeerIQ

[PeerIQ](#) is a virtual appliance-based analytics system that provides a self-hosted dashboard and analytics environment for monitoring the health and performance of PeerGFS and the storage environment.

PeerIQ provides real-time monitoring, performance metrics, and operational insights into PeerGFS and the storage environment. It is also the foundation for Peer Software's File System Analytics and File Activity Analytics capabilities.

Access to PeerIQ is determined by the subscription license purchased from Peer Software.

File System Analytics

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

File System Analytics can provide analysis of file systems, including file types, file sizes, data age, data access patterns, and other relevant metrics.

Access to File System Analytics is determined by the subscription license purchased from Peer Software.

File Activity Analytics

[File Activity Analytics](#) is part of PeerIQ and uses real-time data from Peer Agents to provide insights and trends about how unstructured data and storage infrastructure are being used.

File Activity Analytics can provide analysis of user activity over time, including information about the top users by activity.

Access to File Activity Analytics is determined by the subscription license purchased from Peer Software.

For information about enabling and configuring analytics settings, see [Analytics](#) in [Preferences](#).

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.

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.

File System Analytics Scan Process

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.

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 failback** - If a file server goes offline, Peer Management Center can disable the associated folder target in the DFS namespace. This automatically redirects users to another available file server. When the original file server comes back, Peer Management Center will automatically make sure it is brought back in sync, and then enable the associated folder target so users can once again connect to it. See [DFS Namespace Failover and Failback](#) for more information.

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

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

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

Using DFS Namespaces with Jobs

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

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

- You can create a new namespace or import an existing one by first creating a DFS-N Management job and then later [linking the namespace to a File Collaboration or File Synchronization job](#).
- You can create a new namespace or import an existing one when creating a File Collaboration or File Synchronization job, thus automatically linking namespace folder targets to the 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.

Handling Symbolic Links

PeerGFS can synchronize symbolic links in File Collaboration, File Replication, and File Synchronization jobs. The Symbolic Link Strategy setting on the General page determines how symbolic links are handled.

- **Copy:** Copies the symbolic link itself to destination hosts.
- **Follow:** Follows the symbolic link and copies the target file or folder contents.

The default strategy and the options available depend on the job protocol and participant configuration.

Jobs Using SMB

For File Collaboration, File Replication, and File Synchronization jobs that use SMB and contain only Windows File Server participants:

- New jobs default to **Copy**.
- Existing jobs default to **Follow**.
- You can select either **Copy** or **Follow**.

For jobs that use SMB and contain one or more NAS participants:

- New jobs are set to **Follow**, and the setting cannot be changed.
- Existing jobs are assigned **Follow**, and the setting cannot be changed.
- **Copy** is not supported.

If you add a NAS participant to an existing job that uses SMB and whose **Symbolic Link Strategy** is set to **Copy**, PeerGFS displays a warning. Change the strategy to **Follow** before saving the job.

Jobs Using NFS

For File Collaboration, File Replication, and File Synchronization jobs that use NFS:

- New jobs default to **Copy**.
- Existing jobs default to **Follow**.
- You can select either **Copy** or **Follow**.

Jobs Using Multi-Protocol

For File Collaboration, File Replication, and File Synchronization jobs that use Multi-Protocol:

- New jobs default to **Copy**.
- Existing jobs default to **Follow**.
- You can select either **Copy** or **Follow**.

The restrictions that apply to jobs using SMB with NAS participants do not apply to jobs using Multi-Protocol; **Copy** and **Follow** remain available.

Changing the Symbolic Link Strategy

When you change **Symbolic Link Strategy** for an existing job, PeerGFS displays a confirmation message because changing the strategy can affect data that already exists on destination hosts:

- Changing from **Follow** to **Copy** can cause existing regular files or folders to conflict with the symbolic links that PeerGFS attempts to create. Files or folders affected by the change may be quarantined.
- Changing from **Copy** to **Follow** does not remove symbolic links or other data that were previously created on destination hosts.

File Metadata Synchronization

Overview

File metadata consists of additional information associated with a file, including attributes such as ownership, permissions, and auditing settings. The structure and representation of this metadata vary depending on the underlying file system. In this context, file metadata primarily refers to security-related metadata used to control access and auditing.

For Windows-based systems, file metadata is primarily represented using Access Control Lists (ACLs). For NFS and Linux-based systems, metadata includes standard POSIX attributes as well as optional ACLs.

The metadata elements that can be synchronized include:

Windows (NTFS) metadata:

- **Owner:** The user who owns the file or folder and can modify permissions.
- **DACL (Discretionary Access Control List):** Defines which users and groups are granted or denied access.
- **SACL (System Access Control List):** Specifies auditing rules for access attempts.

NFS/Linux metadata:

- **Owner:** The user ID (UID) associated with the file.
- **Group:** The group ID (GID) associated with the file.
- **Permissions:** Standard read, write, and execute permissions for owner, group, and others.
- **ACLs:** Extended permissions using either NFSv4 ACLs or POSIX ACLs, depending on the system configuration.

File Metadata Conflict Resolution

File metadata conflict resolution occurs the first time a file is synchronized and only when one or more security attributes do not match those of the designated master host.

If the file does not exist on the designated master host, conflict resolution is not performed. If a master host is not selected, file metadata synchronization is not performed during the initial synchronization.

ACL Requirements

Enabling ACL synchronization requires that all participating systems can resolve and apply the referenced users and groups defined in the ACLs or file ownership. Failure to do so may result in files being inaccessible on target systems.

All Peer Agents must run under an account with sufficient privileges to read and apply security metadata across all participating systems.

To ensure accurate and consistent metadata propagation, security settings for the watch set must match across all participants. Differences in underlying file systems or permission models (for example, Windows and NFS environments) may affect how permissions are interpreted and applied.

Multi-Protocol (SMB/NFS)

Overview

Peer Software's **multi-protocol** functionality allows both NFS and SMB clients to interact with the same set of data within PeerGFS. This enables seamless data sharing and access between Linux and Windows systems.

Definition of 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 SMB (CIFS) clients simultaneously. This allows the same data to be accessed via different protocols (NFS for Linux clients and SMB for Windows clients).

Multi-Protocol Architecture

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 PeerGFS to monitor the real-time activity of both NFS and SMB clients accessing 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).

Supported Job Types

Multi-protocol functionality is supported for the following PeerGFS job types:

- File Collaboration
- File Replication
- File Synchronization

System Requirements

To use multi-protocol functionality in PeerGFS, ensure your environment meets the following requirements:

- **Supported Clients**

- Multi-protocol support is available for Linux clients only.
- **Supported NAS Platforms**
 - NetApp ONTAP (cDOT)
 - Amazon FSxN
 - Nutanix Files
 - Dell PowerScale (event detection via Syslog only)
 - Dell PowerStore (event detection via CEE only)
 - Dell Unity (event detection via CEE only)
- **File System Configuration**
 - A share and export must point to the same directory.
 - A single security style (Linux or Windows) must be used.

Limitations and Restrictions

While multi-protocol functionality enhances flexibility, the following limitations apply:

- **Transfer Method:** NFS is used for file transfers and maintains name capitalization.
- **File Collaboration Jobs:** Supported only when using Dell PowerScale devices.
- **Job Overlapping:** Job relay (overlapping File Replication jobs) is not supported.
- **File Permissions:** The user defines which file permissions (metadata) are transferred. This ensures predictable scans.
- **Security Style Limitation:** Only one security style (Windows or Linux) can be used at a time—either Linux security or Windows security.
- **Mixed File Permissions:** Mixed Windows and Linux permissions are supported only on Dell PowerScale platforms.

SMB Access for Multi-Protocol Jobs

SMB access configuration enables a Linux Agent to access data on a file server using the SMB file sharing protocol. It is used only in **multi-protocol job scenarios** and is not required for jobs that use SMB-only or NFS-only protocols.

When SMB access is required:

SMB access configuration is required when all of the following conditions apply:

- The job is a multi-protocol job
- A Linux Agent participates in the job
- One or more of the following applies:
 - o Windows-style permissions must be synchronized
 - o Both Windows-style and Linux-style permissions are synchronized (mixed permissions)
 - o The storage platform is Nutanix Files

In these situations, the Linux Agent requires SMB access to correctly process Windows-style permissions while participating in a multi-protocol job.

How SMB access is used:

- SMB access configurations are **created once** and can be **reused across multiple jobs**.
- A single SMB access configuration can be used by multiple Linux Agents, as needed.
- SMB access configuration is **not tied to a specific job**; it provides access credentials that Agents can use when required.
- SMB access configuration does **not control how data is transferred** between systems. It enables Agents to access data using the SMB protocol when required to process Windows-style permissions.

Where SMB access is configured:

SMB access configurations are created and managed in Preferences. Once configured, they can be selected when prompted during job creation or editing.

For instructions on configuration, see [SMB Access Configuration](#).

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	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

Option	Description
	restart the Peer Agent Windows service using the Windows sc command. Note that this works only if the user running the Peer Management Center can access the remote Peer Agent system and has the appropriate domain permissions to start and stop services on the remote Peer Agent system.
Remote Desktop to Agent	(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.
Support Tools	Displays a list of tools that can be used to assist Peer Support.
Run Event Detection	Runs the Event Detection Analytics tool for the selected job, which looks at real-time activity that has been occurring on that specific Agent.

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

Re-enabling a Disabled Agent Within a Job

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

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

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

State Change Alert Details

Timestamp: 07-29-2024 20:21:15

Severity: Unknown

Host Name: Lab24Win22b

Status: Disabled

State: Disabled

Message: Action: Alert and Disable Host - Malicious Event Detection (MED) - Bait File (BF) Alert[Note: Alert and Disable Host: Please check for unwanted activity before re-enabling]Info: BAIT FILE ALERT[App ID: 113|App Session ID: 145|Participant Detected: Lab24Win22b|Time Detected: 30-07-2024 00:21:06-479|User Detected: LabAdmin|IP Detected: N/A|Process Detected: explorer.exe|Share Detected: N/A|Job Session ID: 1005|BF Path: C:\Bug Reports - Centreville\pc-med_bin\Doc_001-med.docx|BF Event Detected: Delete|PM Pattern: N/A|PM Short Version Found: N/A|PM Short Version Deleted: N/A|PM Delete Delay: N/A|PM Write Count: N/A|PM File List: N/A|PM Renamed From: N/A|PM Old Name: N/A|PM New Name: N/A|PM Folder: N/A|Count: N/A]

Ref ID:

Click outside of popup to close

Host	Root Path	Status	State	Message	DFS Namespace	DFS Target Status	Edge Caching Role	Total Disk Space	Free Disk Space	T
Lab24Win19b	C:\Bug Reports - DC	Participating	Active		Tickets	Active	None	49.40 GB	17.81 GB	1
Lab24Win19c	C:\Bug Reports - Tokyo	Participating	Active		Tickets	Active (Auto)	None	49.40 GB	24.09 GB	1
Lab24Win22a	C:\Bug Reports - Mur	Participating	Active		Tickets	Active (Auto)	None	99.32 GB	66.04 GB	1
Lab24Win22b	C:\Bug Reports - Cen	Disabled	Disabled		Tickets [DISABLED]	Disabled (Auto) [DISABLED]	None	99.32 GB	70.47 GB	0

Host Participant State Change Log

Filter by: Host:

Date	Host	Status
07-29-2024	Lab24Win19	Participat
07-29-2024	Lab24Win22	Disabled
07-29-2024	Lab24Win22	Disabled
07-29-2024	Lab24Win19	Participat
07-29-2024	Lab24Win22	Participat

Status: Running

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.



The image shows a screenshot of the 'Agent Configuration' dialog box, specifically the 'Broker' tab. The dialog has a blue title bar with the text 'Agent Configuration' and standard window controls. On the left, there is a sidebar with four tabs: 'Broker' (selected), 'General', 'Logging', and 'VM Options'. The main area is titled 'Broker' and contains a warning message: 'WARNING: Changes to these settings may make the Agent unable to start. These changes will take effect only after the Agent is restarted.' Below the warning, there are three sections: 'Primary Brokers', 'Failover Brokers', and 'Connection Settings'. The 'Primary Brokers' section has fields for 'Broker Hosts' (containing 'Lab24Win19a'), 'Connection Type' (a dropdown menu set to 'ssl'), and 'Broker Port' (a spinner box set to '61617'). The 'Failover Brokers' section has similar fields, but the 'Broker Hosts' field is empty. The 'Connection Settings' section has three checkboxes: 'Use Compression' (checked), 'Use Synchronous Sends' (unchecked), and 'Validate TLS Hostname' (unchecked). At the bottom right of the dialog are 'OK' and 'Cancel' buttons.

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: Lab24Win19a

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

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.

Broker

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.



The image shows a screenshot of the 'Agent Configuration' dialog box, specifically the 'Broker' tab. The dialog has a blue title bar with the text 'Agent Configuration' and standard window controls. On the left, there is a sidebar with four tabs: 'Broker' (selected), 'General', 'Logging', and 'Performance'. Below these is a section for 'VM Options'. The main area of the dialog is titled 'Broker' and contains a warning message: 'WARNING: Changes to these settings may make the Agent unable to start. These changes will take effect only after the Agent is restarted.' Below the warning, there are three sections: 'Primary Brokers', 'Failover Brokers', and 'Connection Settings'. The 'Primary Brokers' section has fields for 'Broker Hosts' (containing 'Lab24Win19a'), 'Connection Type' (a dropdown menu set to 'ssl'), and 'Broker Port' (a spinner box set to '61617'). The 'Failover Brokers' section has similar fields, but the 'Broker Hosts' field is empty. The 'Connection Settings' section has three checkboxes: 'Use Compression' (checked), 'Use Synchronous Sends' (unchecked), and 'Validate TLS Hostname' (unchecked). At the bottom right of the dialog are 'OK' and 'Cancel' buttons.

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: Lab24Win19a

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. 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.
Connection 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.
Connection 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 Compression	Enable to compress all communication between the selected Agent and broker(s).
Use Synchronous Sends	Enable to always send messages from an Agent to broker(s) in synchronous mode. If not enabled, then messages between Agent 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).

General

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

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

The image shows a screenshot of the 'Agent Configuration' dialog box. The 'General' tab is selected in the left-hand sidebar, which also lists 'Broker', 'Logging', 'Performance', and 'VM Options'. The main area of the dialog is titled 'General' and contains three sections: 'Workspace', 'Location Information', and 'Statistics'. The 'Workspace' section has a text field for 'Agent Workspace Directory' with the value 'workspace'. The 'Location Information' section has 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 of the dialog are 'OK' and 'Cancel' buttons.

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.

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.

Logging



The image shows a screenshot of the 'Agent Configuration' window. The window has a blue title bar with the text 'Agent Configuration' and standard window controls. On the left side, there is a vertical navigation pane with the following options: 'Broker', 'General', 'Logging' (which is highlighted with a blue background), 'Performance', and 'VM Options'. The main area of the window is titled 'Logging'. Inside this area, there is a section titled 'Archiving of Agent Logs'. Below this section, there is a label 'Max number of days to keep before archiving:' followed by a numeric spinner control set to the value '7'. At the bottom right of the window, there are two buttons: 'OK' and 'Cancel'.

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.

Performance

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.

The screenshot shows the 'Agent Configuration' dialog box with the 'Performance' tab selected. The left sidebar lists 'Broker', 'General', 'Logging', 'Performance' (highlighted), and 'VM Options'. The main area is titled 'Performance' and contains three sections: 'Network Tuning' with an unchecked 'Auto Calculate Data Streams' checkbox and a 'Bandwidth to Broker (Mbps)' spinner set to -1; 'Advanced Options' with 'Parallel Data Streams' and 'Outstanding Data Packets' spinners set to 4 and 8 respectively; and 'Processor Affinity' with a 'Max Number of Processors to Use (4 available)' dropdown menu set to 'All Available'. 'OK' and 'Cancel' buttons are at the bottom right.

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

Field	Description
	Bandwidth to Broker (Mbps) in the calculation along with latency between the Broker and Agent.
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 Peer Support 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.

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.

The screenshot shows a 'View Agent Properties' dialog box with the following fields and values:

Field	Value
Agent Host Name	Lab24Win19b
Connection Status	Connected
Custom Description	☐
Description	
Discovery Time	07-18-2024 20:12:22
Heartbeat Enabled	☒
JVM Architecture	amd64
JVM Version	17.0.11+9
Local Time	07-21-2024 16:17:28 EDT
Local TimeZone	Eastern Standard Time
SSL Enabled	☒
Start Time	07-15-2024 18:16:54
Username	Labadmin
Version	6.0.0.20240621


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

Tab	Description
General	Displays general Peer Agent run-time information such as discovery time, local time, TLS use, Peer Agent start up time, Peer Agent version, and the user name Peer Agent service is running as.
Heartbeat	Displays heartbeat information and statistics such as heartbeat frequency, average heartbeat time, last heartbeat time, total Peer Agent disconnects, total missing heartbeats.

Tab	Description
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.

VM Options

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.

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.

API Categories

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

Status Codes

The PeerGFS API has three status codes:

- 200 - Success
- 401 - Unauthorized
- 404 - Job(s) not found

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](#) in [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.

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: • Real-time activity on Smart Data Seeding host is disabled. • Real-time events on that host will be quarantined. • Renamed files will be restored.
PASSIVE_WITH_RESTORE	Initial scan will be one-way only with any host in Smart Data Seeding mode: • Real-time activity on Smart Data Seeding host is disabled. • Any activity on that host will be restored to its original state.
ACTIVE_LIMITED	Initial scan will be one-way only with any host in Smart Data Seeding mode: • Real-time activity on Smart Data Seeding host is enabled in a limited mode (real-time file adds are processed). • Unsynchronized file updates will be quarantined. • Unsynchronized file renamed will be restored. • Unsynchronized file deletes will be restored.
ACTIVE_FULL	Initial scan will be one-way only with any host in Smart Data Seeding mode except for updates (updates will be processed as Latest Modified wins): • Real-time activity on Smart Data Seeding host is enabled with latest modified file wins, regardless of whether the latest file is on the Smart Data Seeding host.
REACTIVATION	Initial Scan will be one-way only with any host in Smart Data Seeding mode: • Real-time activity on Smart Data Seeding host is enabled with Quarantine (Added and Updated Files will be quarantined during the scan).

Mode	Description
	• 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%.

Tagging Resources

Tags are labels that can be assigned to **resources** such as jobs and Agents in Peer Management Center. Tags help organize resources and allow users to filter lists to focus on a subset of jobs or Agents. **Categories** group related tags under a broader label, making it easier to organize and browse tags.

Tagging is primarily used to support filtering, making it easier to focus on a subset of jobs or Agents in environments with many managed objects. In this context, **resources** are managed objects that can be tagged, namely jobs and Agents.

The following topics explain how tags and categories work, how to create and manage tags, how to assign tags to resources, and how to use tags to filter lists.

- [Understanding Tags and Categories](#)

- [Creating Tags](#)
- [Assigning Tags](#)
- [Filtering with Tags](#)
- [Associating Tags with Roles](#)
- [Managing Tags](#)
- [Creating Categories](#)
- [Managing Categories](#)

Understanding Tags and Categories

Tags are labels that you assign to **resources** such as jobs and Agents. Once a tag is assigned to a resource, you can use it as a filter to limit the items shown in a list, making it easier to locate the specific resources you need.

When you create a tag you must assign it to a **category**; categories are simply groups that organize related tags so users can find and select the right tag more quickly. For example, with categories like *Location*, *Business Unit*, and *Platform*, filtering on *Location* would display only the resources that have a tag belonging to that category.

Each tag belongs to exactly one category, while a category may contain one or many tags. Tags and categories affect only how items are grouped and displayed—they do not affect job behavior or execution.

Creating Tags

Tags are created in the [Tags Configuration](#) page in **Preferences**. You can also create tags from the **Assign Tags** dialog while managing tag assignments.

Creating Tags from the Tags Configuration Page

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**.


The screenshot shows the 'Preferences' dialog box with the 'Tags Configuration' tab selected. On the left is a tree view of configuration categories, with 'Tags Configuration' highlighted. The main area contains a table with columns 'Tag Name', 'Category', and 'Resources'. It lists three tags: 'EMEA', 'North America', and 'South America', all with the category 'Location' and 0 resources. Above the table are filters for 'Tag Name' and 'Category'. To the right of the table are buttons for 'Create', 'Edit', 'Copy', 'Delete', and 'Edit Categories'. At the bottom are 'Apply and Close', 'Cancel', and 'Apply' buttons.

Tag Name	Category	Resources
EMEA	Location	0
North America	Location	0
South America	Location	0

3. Click **Create**.
4. Enter a tag name.

The screenshot shows the 'New Tag' dialog box. It has two input fields: 'Tag Name' and 'Category Name'. The 'Category Name' field has a dropdown arrow. At the bottom are 'OK' and 'Cancel' buttons.

5. Select a tag category.
6. Click **OK**.



7. Click **Apply and Close**.

Creating Tags from the Assign Tags Dialog

You can also create tags from the **Assign Tags** dialog.

To access the **Assign Tags** dialog:

1. Click the **Assign Tags** icon on the toolbar of the **Main** view, **Jobs** view, or **Agent Summary** view.
2. Select **Tags** for the **Group by** option.

Assign Tags [X]

Tagging resources helps when managing large number of resources. Use this dialog to view/assign tags to resources.

Group by: ☐ Resources ☒ Tags

Tag Name Filter [Pencil] Category Filter [Pencil] Resource Name [Pencil] Type Filter [Pencil]

Tag Name	Category	Resources	Resource Name	Type
APAC	Location	0		
EMEA	Location	0		
North America	Location	0		
Oceania	Location	0		
South America	Location	0		

[Edit] [Create Tag] [Report]

[Close]

- Click **Create Tag**.
- Enter the tag information.

New Tag [X]

Tag Name: [Text Box]

Category Name: [Dropdown]

[OK] [Cancel]

- Click **OK**.


The **Assign Tags** dialog is redisplayed.

6. Click **Close**.


Tags are assigned to resources by associating jobs or Agents with one or more tags. Assigning tags adds metadata to a resource and makes that resource available for tag-based filtering. Tags can be assigned while editing an existing job or by using the Assign Tags dialog. Both approaches create the same tag-resource association.

To assign tags to a job:

1. Right-click the job and select **Edit Job(s)**.
2. Open the **Tags** page.



3. In **Unassigned tags**, select one or more tags to assign to the job, and then click the right-arrow (**Add one**) or double right-arrow (**Add all**) to move them to **Assigned tags**.



4. Click **OK** to apply the changes.

Assigning Tags Using the Assign Tags Dialog

1. Open the **Assign Tags** dialog by doing one of the following:
 - Click **Assign Tags** from the toolbar of the **Main** view, **Jobs** view, or **Agents** view.
 - Right-click a job or Agent and select **Assign Tags**.
2. In the **Group by** options, select **Resources** or **Tags** to determine whether you begin by selecting a resource or a tag.

- When **Resources** is selected, the dialog displays a table of resources on the left and a table of tags on the right.

Assign Tags [X] [Close]

Tagging resources helps when managing large number of resources. Use this dialog to view/assign tags to resources.

Group by: ☒ Resources ☐ Tags

Resource Name	Type
DGUbuntu1	Agent
DGUbuntu2	Agent
DGWin16A	Agent
DGWin16B	Agent
DGWin16C	Agent
DGWin16D	Agent
\\PeerTest.local\\De	Job / DFS-N Management
\\PeerTest.local\\DC	Job / DFS-N Management
\\PeerTest.local\\Do	Job / DFS-N Management
\\PeerTest.local\\Te	Job / DFS-N Management
Development Proje	Job / File Collaboration
FC-1	Job / File Collaboration
FC-2	Job / File Collaboration
FC-3	Job / File Collaboration
FC-Nutanix-Windov	Job / File Collaboration
FC-Nutanix-Windov	Job / File Collaboration

Tag Name	Category
----------	----------

- When **Tags** is selected, the dialog displays a table of tags on the left and a table of resources on the right.

Assign Tags [X] [Close]

Tagging resources helps when managing large number of resources. Use this dialog to view/assign tags to resources.

Group by: ☐ Resources ☒ Tags

Tag Name	Category	Resources
Bonds	Financial	0
IRA	Financial	0
Mutual Funds	Financial	1
Stocks	Financial	0
sws	Financial	0

Resource Name	Type
---------------	------

The table on the right is initially empty. Selecting an item in the table on the left displays the items currently assigned to it.

Assign Tags

Tagging resources helps when managing large number of resources. Use this dialog to view/assign tags to resources.

Group by: ☒ Resources ☐ Tags

Resource Name Filter Type Filter

Resource Name	Type
DGUbuntu1	Agent
DGUbuntu2	Agent
DGWin16A	Agent
DGWin16B	Agent
DGWin16C	Agent
DGWin16D	Agent
\\PeerTest.local\De	Job / DFS-N Management
\\PeerTest.local\DC	Job / DFS-N Management
\\PeerTest.local\Do	Job / DFS-N Management
\\PeerTest.local\Te	Job / DFS-N Management
Development Proje	Job / File Collaboration
FC-1	Job / File Collaboration
FC-2	Job / File Collaboration
FC-3	Job / File Collaboration
FC-Nutanix-Window	Job / File Collaboration
FC-Nutanix-Window	Job / File Collaboration

Tag Name Filter Category Filter

Tag Name	Category
Real Estate	Financial

Edit Report

Close

- (Optional) Use the filter fields to narrow the list of resources or tags.
- Continue with the appropriate procedure below, depending on the **Group by** option selected.

Assigning Tags When Grouping by Resources

- In the **Assign Tags** dialog, select **Resources** for the **Group by** option.
- Select a job or Agent from the resources table.

The tags currently assigned to that resource appear in the table on the right.

Assign Tags

Tagging resources helps when managing large number of resources. Use this dialog to view/assign tags to resources.

Group by: ☒ Resources ☐ Tags

Resource Name	Type	Tag Name Filter	Category Filter
DGUbuntu1	Agent	Bonds	Financial
DGUbuntu2	Agent		
DGWin16A	Agent		
DGWin16B	Agent		
DGWin16C	Agent		
DGWin16D	Agent		
\\PeerTest.local\De	Job / DFS-N Management		
\\PeerTest.local\DC	Job / DFS-N Management		
\\PeerTest.local\Do	Job / DFS-N Management		
\\PeerTest.local\Te	Job / DFS-N Management		
Development Proje	Job / File Collaboration		
FC-1	Job / File Collaboration		

Buttons: Edit, Report, Close

- Click **Edit**.

The **Assign/Unassign resources to (Resource Name)** dialog opens and displays two tables: **Unassigned tags** and **Assigned tags**.

Unassigned tags shows tags that are not currently assigned to the selected resource. **Assigned tags** shows tags that are currently assigned to the selected resource.

Assign/Unassign tags to (DGUbuntu2)

Use this dialog to assign tags to the selected resource.
Move tags from the left table to the right table so the selected tags are assigned to them.

Create Tag

Unassigned tags		Assigned tags	
Tag Name Filter	Category Filter	Tag Name Filter	Category Filter
IRA	Financial	Bonds	Financial
Mutual Funds	Financial		
Real Estate	Financial		
Stocks	Financial		

Buttons: >>, >, <, <<

1 resource selected Save Cancel

- In **Unassigned tags**, select the tags to assign to the resource.

- Click the **right-arrow (Add one)** or **double right-arrow (Add all)** to move the selected tags to **Assigned tags**.

Assign/Unassign tags to (DGUbuntu2)

Use this dialog to assign tags to the selected resource.
Move tags from the left table to the right table so the selected tags are assigned to them.

Create Tag

Unassigned tags

Tag Name Filter	Category Filter
Tag Name	Category
IRA	Financial
Real Estate	Financial
Stocks	Financial

Assigned tags

Tag Name Filter	Category Filter
Tag Name	Category
Bonds	Financial
Mutual Funds	Financial

>>

>

<

<<

2 resources selected

Save Cancel

- Click **Save**, and then **Close**.

Assigning Tags When Grouping by Tags

- In the **Assign Tags** dialog, select **Tags** for the **Group by** option.
- Select a tag from the tags table.

The resources that currently have the selected tag assigned appear in the table on the right.

Assign Tags □ ×

Tagging resources helps when managing large number of resources. Use this dialog to view/assign tags to resources.

Group by: ☐ Resources ☒ Tags

Tag Name Filter
Category Filter

Resource Name
Type Filter

Tag Name	Category	Resources	Resource Name	Type
Bonds	Financial	2	DGWin16B	Agent
IRA	Financial	4	DGWin16C	Agent
Mutual Funds	Financial	2	FC-Nutanix-Windov	Job / File Collaboration
Real Estate	Financial	2	FR-1	Job / File Replication
Stocks	Financial	2		


Edit
Create Tag
Report

Close


- Click **Edit**.

The **Assign/Unassign resources to (Tag Name)** dialog opens and displays two tables: **Unassigned resources** and **Assigned resources**.

Unassigned resources shows resources that do not currently have the selected tag assigned. **Assigned resources** shows resources that currently have the selected tag assigned.



- In **Unassigned resources**, select the resources to which the tag will be assigned.
- Click the **right-arrow (Add one)** or **double right-arrow (Add all)** to move the selected resources to **Assigned resources**.



- Click **Save**, and then **Close**.

Filtering with Tags

Using Tags to Filter Lists

Tags assigned to jobs or Agents can be used as filter criteria by [list filters](#) throughout throughout Peer Management Center, including views such as **Jobs** and **Agent Summary**.

Filtering does not create, modify, or assign tags—it simply uses existing tag assignments to limit which items are displayed.

For information on using tags in list filters, see [Filtering with Tags](#).

You can use tags to filter jobs and Agents. **Filtering** applies to lists displayed in Peer Management Center. It controls which items appear in a list—such as jobs or Agents—based on selected criteria, including tags. Tags do not determine which list an item appears in; they provide filter criteria that limit what is displayed. Filtering affects only the current list and does not modify the underlying items.

Filtering with tags is based entirely on existing tag assignments. Only items that have been tagged appear in tag-based filter results. To filter by a specific tag or category, the tag must first be created and assigned.

Using Tags to Filter Jobs

To filter a large list of jobs using tags:

1. Enter an tag in the **filter field** located below the toolbar buttons in the **Jobs** view.

For example, to show all jobs with a tag that has *North America* in the tag name and *Location* in the tag category, enter the following in the filter field: `tag:"North America" AND cat:Location`

The list of jobs is automatically filtered.

Using Tags to Filter Agents

To filter a large list of Agents:

- Use the filter field located below the toolbar buttons in the **Agents** view .
- Use the **Filter by** field in the **Agent Summary** view:
 - a. Select **Expression** in the **Filter by** dropdown.
 - a. Enter an expression in the filter field.

Filtering Resources

Filtering resources limits which resources are displayed in a view based on selected criteria, such as tags or resource attributes, without modifying the resources themselves. For more details on how to filter resources, see [Creating Complex Filter Expressions](#) in [List Filters](#).

Associating Tags with Roles

Overview

Tags that are assigned to jobs can be **associated** with a [user role](#) to control which jobs are visible to users assigned to that role. When tags are associated with a role, users with that role see only jobs that have those tags assigned.

This allows administrators to limit job visibility so that users see only the jobs relevant to their region, team, or area of responsibility.

Example

For example, an organization might create a **Location** category with tags such as **North America**, **EMEA**, and **APAC**. Jobs are then assigned one or more of these tags.

Job	Tags
Finance	North America
HR	EMEA
Engineering	APAC

A role for the European support team can be associated with the **EMEA** tag. Users assigned to that role will see only the jobs that have the **EMEA** tag assigned.

Before configuring role-based job visibility, ensure that the appropriate tags have been assigned to jobs. For information about assigning tags to jobs or Agents, see [Assigning Tags](#).

To associate tags with a role

1. Open **Preferences**.
2. Select **User Management** from the navigation tree.
3. Select the role and click **Edit**.
4. Click the **Tags** tab.
5. Move the desired tags from **Unassigned Tags** to **Assigned Tags**.
6. Click **OK**.

Managing Tags

After tags are created, they can be modified or removed as your environment changes. You can rename a tag, copy a tag, change the category assigned to a tag, delete tags that are no longer needed, and run a report to review how tags are used across jobs and Agents. You can also run a report summarizing the existing tags.

Editing Tags

You can edit the name of a tag or change the category assigned to the tag.

To edit a tag:

1. From the **Tags Configuration** page, select a tag.

The screenshot shows the 'Preferences' window with the 'Tags Configuration' tab active. The left sidebar contains a list of configuration categories, including Analytics, Cloud Backup and Replication, Collab, Sync, and Replication, DFS-N Management, Email Configuration, General Configuration (with sub-items like Agent Connectivity, Broker Configuration, Email Alerts, Proxy Configuration, Tags Configuration, and Web and API Configuration), Licensing, MED Configuration, PMC Backup, Real-time Event Detection, SNMP Configuration, Storage Configuration, Task Scheduler, and User Management. The 'Tags Configuration' category is selected.

The main area displays a table of tags with the following columns: Tag Name, Category, and Resources. The 'EEA' tag is highlighted. The table contains the following data:

Tag Name	Category	Resources
Accounting	Business Unit	0
Engineering	Business Unit	0
Finance	Business Unit	0
Human Resources	Business Unit	0
Sales	Business Unit	0
APAC	Location	0
EEA	Location	0
EMEA	Location	0
North America	Location	0
Oceania	Location	0
South America	Location	0
Amazon Fsx for NetApp ONTAP	Platform	0
Dell PowerScale	Platform	0
Dell PowerStore	Platform	0
Dell Unity	Platform	0
Linux File Server	Platform	0
NetApp ONTAP	Platform	0
Nutanix Files	Platform	0
Windows File Server	Platform	0

On the right side of the table, there are buttons for 'Create', 'Edit', 'Copy', 'Delete', and 'Edit Categories'. At the bottom of the window, there are buttons for 'Apply and Close', 'Cancel', and 'Apply'.

2. Click **Edit**.
3. Modify the tag name or category.
4. Click **OK**.

Deleting Tags

Tags that are no longer needed can be deleted. If a tag is assigned to jobs or Agents, those assignments are removed when the tag is deleted.

To delete a tag:

1. From the **Tags Configuration** page, select a tag.
2. Click **Delete**.
3. Click **OK**.

Running a Tags Report

You can run a report to review the tags defined in the system, the category assigned to each tag, and the resources that use each tag.

To run the tag reports:


1. Click the **Assign Tags** icon on the toolbar of the **Main** view, **Jobs** view, or **Agent Summary** view.
2. In the **Assign Tags** dialog, select **Tags** for the **Group by** option.
3. Click **Edit**.
4. Click **Report**.
5. In the **Tags** report dialog, click **Export** to export the report to a CSV file.
6. Click **OK** in the confirmation dialog.

Creating Categories

Categories are created from the [Tags Configuration](#) page in **Preferences**.


To create a category:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **General Configuration** in the navigation tree, and then select **Tags Configuration**.




3. Click **Edit Categories**.

The **Edit Tag Categories** dialog is displayed.




- Click **Add**.
- Enter a category name.



A dialog box titled "New Tag Category" with a close button (X) in the top right corner. It contains a text input field labeled "Category name:" and two buttons at the bottom: "OK" and "Cancel".

- Click **OK**.

The **Edit Tag Categories** dialog is redisplayed.



A dialog box titled "Edit Tag Categories" with a close button (X) in the top right corner. It features a "Category Filter" input field with a magnifying glass icon. Below this is a table with three columns: "Category", "Tags", and "Used Tags". The table contains three rows of data. To the right of the table are four buttons: "Add", "Edit", "Delete", and "Cleanup". At the bottom right of the dialog is a "Done" button.

Category	Tags	Used Tags
Business Unit	0 tags	0 tags
Location	0 tags	0 tags
Platform	0 tags	0 tags

- Click **Done**.

The **Tags Configuration** page is redisplayed.

Managing Categories

Overview

Categories let you group related tags so users can find and apply the right tag more quickly. By organizing tags into logical categories you get:

- **Consistent tagging** – Because every tag must be assigned to a category, users always see the same set of category choices when they add a tag, which encourages uniform tag usage.
- **Easier maintenance** – Categories can be renamed or deleted when they're no longer needed, letting you keep the category list current with minimal effort.
- **Cleaner data** – Unused or orphaned categories can be removed, keeping the tag list tidy and improving search performance.

Managing tag categories is therefore a routine part of keeping the system's metadata accurate and useful.

Working with Categories

You can manage tag categories by performing four core actions:

Action	What It Does	Key Points
Add	Creates a new category that will appear when users create or edit tags.	Choose a clear, descriptive name; the new category is immediately available.
Edit	Renames an existing category.	Changing the name does not affect the tags already assigned to it.
Delete	Removes a category that is no longer needed.	All tags assigned to the category must be moved to another category first; the system blocks deletion until this is done.
Cleanup	Purge categories that have no tags (or only orphaned tags) in a single step.	Keeps the category list tidy and improves search performance.


Accessing the Edit Categories Dialog

You can create, rename, delete, and clean up tag categories from the **Edit Categories** dialog:

To access the Edit Categories dialog:

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

- Expand **General Configuration** in the navigation tree, and then select **Tags Configuration**.



- Click **Edit Categories**.

The **Edit Tag Categories** dialog is displayed.

Category	Tags	Used Tags
Business Unit	0 tags	0 tags
Location	0 tags	0 tags
Platform	0 tags	0 tags

4. Make your modifications: Add, edit, delete, or clean up categories

Managing TLS Certificates

Overview

If you are a PeerGFS administrator, you may need to manage Transport Layer Security (TLS) certificates when configuring secure communications, replacing default or expired certificates, or maintaining trusted connections between system components. PeerGFS uses TLS to secure communication between its components, including the Peer Management Center (PMC), brokers, and Agents.

TLS ensures that:

- Connections are encrypted
- Each component can verify the identity of the other

- Data is protected during transmission

Important: SSL vs TCP Configuration

TLS is used only when a component is configured for SSL.

- SSL enabled - Secure, encrypted communication using certificates
- TCP configured - No encryption, no certificates, no TLS handshake

This topic provides a conceptual overview of how TLS certificates are used in PeerGFS and which certificate management tasks are available from the Peer Management Center (PMC) user interface. Procedures for creating, importing, restoring, and updating certificates are covered in the related task topics.

How TLS Certificates Are Used in PeerGFS

TLS certificates are used by the following components:

Component	Purpose
PMC	Secures access to the web user interface and API.
Agents	Secures communication with the PMC and brokers
Brokers	Secures communication and data transfer between other brokers (when using a network of brokers)

TLS certificates provide:

- **Authentication** – Verifies the identity of communicating hosts.
- **Encryption** – Protects data in transit.

PeerGFS provides default certificates for initial deployment. You can replace these certificates with:

- **Self-signed certificates**
- **Certificates signed by a Certificate Authority (CA)**

Certificate Management in the PMC

TLS certificates are managed through the PMC. Certificate management tasks are performed separately for broker certificates and web client certificates. When you initiate a certificate management operation, the PMC distributes the required certificate changes to brokers and Agents and maintains certificate synchronization among the affected components. You do not need to manually update or distribute certificate files on individual hosts.

Broker Certificate Tasks

You can:

- Generate self-signed broker certificates
- Import broker certificates
- Restore default broker certificates
- View broker certificate details
- Download broker certificates
- Update trusted broker certificates
- Replace broker certificates while brokers are in Maintenance Mode

Web Client Certificate Tasks

You can:

- Generate self-signed web client certificates
- Import web client certificates
- Restore default web client certificates

Script-Based Certificate Management

PeerGFS also provides utility scripts for certificate management. For more information about using the scripts, see the knowledge base article [Replace TLS Certificates](#).

TLS Certificate Terminology

The following terms are used throughout the TLS certificate management topics.

Term	Description
TLS	Transport Layer Security. A protocol used to secure communications between PeerGFS components.
Certificate	A digital file used to establish identity and encrypted communication between components.
Broker Certificate	The certificate used by the broker to secure communications with Agents and other brokers.
Agent Certificate	A certificate used by PeerGFS Agents to authenticate and communicate securely with the broker.
Web Client Certificate	A certificate used to secure HTTPS communication for the PMC web client.
Topology Perspective	The PMC perspective used to view and manage topology components, including certificate management operations.
Maintenance Mode	A temporary broker state in which the broker uses the default PeerGFS-installed TLS certificate and keystore instead of the active certificate configuration. Maintenance Mode is intended for recovery scenarios to allow Agents using the default trust relationship to reconnect.

Preparing for Certificate Management

TLS certificate management is performed from the **Topology** perspective in the PMC. Before managing TLS certificates, verify that you have access to the Topology perspective and stop any running jobs that will be affected by the certificate update.

After completing these preparation tasks, continue to:

- [Managing Broker and Agent Certificates](#)
- [Managing Web Client Certificates](#)

Accessing the Topology Perspective

To open the **Topology** perspective:

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

For more information about PMC perspectives, see [Perspectives and Views](#).

Access Restrictions in the PMC Web Client

Although all PMC users can access the Topology perspective in the rich client, access to the Topology perspective in the PMC web client depends on the assigned user role. Because certificate management operations modify topology configuration settings, they require an account with **Admin** privileges. For more information about web client user roles, see [Managing Roles for Web Client Users](#).


Role	Topology Perspective Access
Admin	Full access, including certificate management and topology configuration changes.
Power User	Read-only access to the Topology perspective.
Help Desk	No access to the Topology perspective.

Stop Running Jobs Before Managing Certificates


Before performing certificate management operations, stop any running jobs associated with Agents that will be affected by the operation. Jobs associated with unaffected Agents do not need to be stopped.

To stop running jobs associated with one or more Agent:

1. Open the **Jobs** perspective.
2. In the Agents view, select the Agent or Agents that will be affected.
3. Right-click the selection and select **Stop Running Jobs for Agent**.



4. Click **Stop Jobs** in the dialog that appears.



Managing Broker and Agent Certificates

Overview

Broker and Agent certificates are used to establish trusted TLS communication between brokers and Agents within the PMC environment. Managing these certificates allows you to secure communications, replace existing certificates, distribute trusted certificates to Agents, and maintain certificate trust relationships across sites.

You can manage broker and Agent certificates from the **Topology** perspective in the PMC. Certificate management operations help organizations align PMC security with internal certificate management policies.

Operations that modify broker certificates will temporarily interrupt secure communications while services restart and reload updated certificate information. Agents reconnect automatically after the updated certificates are applied.

Before You Begin

Before managing broker and Agent certificates:

- Verify that you have sufficient permissions to access the [Topology perspective](#).
- Ensure that brokers and Agents are online and communicating normally.
- Notify users of potential temporary communication interruptions during certificate updates.
- If using CA-signed certificates, verify that certificates contain the correct hostnames and required trust chain information.
- Back up existing certificate and keystore files before replacing or importing certificates.

For prerequisite information, see [Preparing for Certificate Management](#).

Common Broker and Agent Certificate Tasks

The following certificate management tasks are available from the **Sites** view in the **Topology** perspective.

Task	Description
View Broker Certificate	Display certificate details for a broker.
Generate Broker Certificates	Create new self-signed broker certificates for secure communications.
Import Broker Certificates	Replace existing broker certificates with user-provided certificates and private keys.
Restore Default Broker Certificates	Restore the default PMC broker certificates.
Download All Broker Certificates	Download broker certificate bundles for use during Agent installation or manual trust configuration.
Update Trusted Broker Certificates	Update the trusted broker certificates used by Agents and brokers.

For more information on performing these tasks, see:

- [Viewing the Broker Certificate](#)
- [Generating Self-Signed Broker Certificates](#)
- [Importing Broker Certificates](#)
- [Restoring Default Broker Certificates](#)
- [Downloading Broker Certificates](#)
- [Updating Trusted Broker Certificates](#)

Viewing a Broker Certificate

Overview

Use this procedure if you need to view a broker certificate to verify details or troubleshoot issues.

When to View the Certificate


You may want to view the broker certificate when:

- Troubleshooting connectivity or trust issues between brokers
- Confirming that a recent update or reset was successful
- Verifying details such as expiration date or issuer
- Providing information to support or security teams
- Performing a security or compliance audit


Viewing the Certificate

To view the current broker key:

1. Open the **Topology** perspective.
2. In the **Sites** view, right-click the broker whose certificate you want to view.
3. Select **Manage Broker Certificates**, and then **View Broker Certificate**.



The **Browser Certificate** dialog displays the details of the selected broker certificate. Information includes the keystore name, certificate alias, certificate owner and issuer, serial number, validity period, encryption algorithm, key size, and certificate fingerprint. These details can be used to verify the identity, validity, and security configuration of the broker certificate.



4. Review the certificate details.
5. Click **Close**.

Generating Self-Signed Broker Certificates

Overview

Use this procedure to generate a new self-signed certificate for a broker when you need to enable or replace TLS/SSL encryption among brokers and Agents.

When to Generate the Certificate

You might generate a new self-signed certificate in the following situations:

- Performing the initial setup for a test or development environment.
- Replacing an expired certificate.
- Regenerating a compromised certificate.
- Updating broker hostnames or domain information.
- Resetting broker trust relationships after environment changes.

Important Considerations

Before generating a new self-signed certificate:

- [Stop all running jobs.](#)

Generating a new self-signed certificate has the following effects:


- Connected Agents and brokers are temporarily disconnected during certificate regeneration.
- The Broker and Agent services restart automatically.
- Certificates generated by this procedure are self-signed and are not issued by a trusted Certificate Authority (CA).

Generate the Certificate

To generate the certificate:

1. Open the **Topology** perspective.
2. In the **Sites** view, right-click the broker whose certificate you want to generate.

3. Select **Manage Broker Certificates**, and then **Generate Self-Signed Key and Certificate**.



4. Enter the certificate details:

Generate Self-Signed Key and Certificate

Common Name (CN)*:

Organization (O):

Organizational Unit (OU):

Locality (L):

State (ST):

Country (C):

Validity Period*:

Years: 1

Months: 0

Advanced Settings

Key Algorithm: RSA

Key Size: 2048


Generate

Cancel

Field	Description
CN (Common Name)	Required. Enter the DNS name or hostname that brokers and Agents use to connect to the broker. This value should match the hostname to avoid certificate validation errors.
O (Organization)	Enter your company or organization name
OU (Organizational Unit)	Enter your department name (for example IT).
L (City/Locality)	Enter the city where your organization is located.
ST (State/Province)	Enter your state or province.

Field	Description
ince)	
C (Country)	Select the two-letter country code (for example, US) where your organization is located.

5. Set the certificate validity period in years and months.
6. (Optional) Under **Advanced Settings**, specify values for the following fields:
 - **Key Algorithm** - Select the algorithm used to generate the key pair, such as RSA or EC.
 - **Key Size** - The default key size is 2048 bits. Optionally, select 4096 bits for stronger encryption.
7. Click **Generate**.
8. Review the confirmation dialog, which lists the certificate details and the brokers and Agents that are affected.



9. Click **Yes** to proceed.

During the operation:

- The Broker and Agent services stop and restart automatically.
- Brokers and Agents temporarily disconnect and reconnect during the operation. Their status changes are displayed in the **Sites** and **Site Details** views.
- A progress indicator appears in the lower-right corner of the **Topology** perspective while the operation is in progress.

After the certificate is regenerated, all connected brokers and Agents automatically receive the new keystore and truststore before reconnecting.

Importing Broker Certificates

Overview

Use this procedure to import a broker certificate and private key for secure TLS communication between brokers and Agents. Importing a broker certificate updates the broker keystore, distributes updated trusted certificate information to connected Agents, and restarts the affected services.

When to Import

You might import a broker certificate in the following situations:

- Importing certificates required for secure TLS communication
- Replacing an expiring broker certificate
- Replacing a self-signed certificate with a certificate signed by a Certificate Authority (CA)
- Updating broker and Agent trust relationships
- Replacing a compromised certificate
- Updating certificate information after broker hostname or domain changes

Before You Begin

Before importing broker certificates:


- Verify that you have sufficient permissions to access the Topology perspective.
- Ensure that brokers and Agents are online and communicating normally.
- Notify users of possible temporary communication interruptions.
- If using CA-signed certificates, verify that certificates contain the correct hostnames and required trust chain information.
- Back up existing certificate and keystore files before replacing certificates.
- Stop all running jobs unless the broker is operating in Maintenance Mode.

For prerequisite information, see [Preparing for Certificate Management](#).


Importing the Broker Certificate

To import a broker certificate:


1. Open the **Topology** perspective.
2. In the **Sites** view, right-click the broker whose certificate you want to replace.
3. Select **Manage Broker Certificates**, and then **Import User Key and Certificate**.



4. In the **Import** dialog, provide the required certificate information:



- **Key File (PKCS12)** – Browse to the PKCS12 certificate file (.p12 or .pfx). After selecting the file, click **OK**. You can also enter the full file path manually.



- **Password** – Enter the keystore password.
- **Alias** – Enter the certificate alias contained in the keystore.

5. Click **Import**.

The PMC validates the password and alias before importing the certificate. If validation fails, a **Keystore Validation Error** message appears. For example:

- Failed to validate keystore info, keystore password was incorrect
- Failed to validate keystore info, Alias not found in keystore: alias

6. In the **Confirmation** dialog, click **Yes** to continue.

During the import operation, a progress indicator appears in the bottom-right corner of the Topology perspective.

The broker keystore is updated, trusted certificate information is distributed to connected Agents, and the broker service restarts automatically. Connected Agents temporarily disconnect and reconnect while the updated certificates are applied.

Restoring Default Broker Certificates

Overview

Use this procedure to permanently replace the broker keystore with the default keystore that was installed with PeerGFS. This operation returns the broker to the default PeerGFS-installed certificate configuration and restarts the broker service.


Important: This operation cannot be undone. Before restoring the default broker key and certificate, make sure this is the action you want to perform.

Note: If you only need to temporarily use the default key and certificate, use Maintenance Mode instead.


Restore the Default Broker Key and Certificate

To restore the default certificate:

1. Open the **Topology** perspective.
2. In the **Sites** view, right-click the broker whose certificate you want to restore.
3. Select **Manage Broker Certificates**, and then **Restore Default Key and Certificate**.



4. In the confirmation message, click **Restore**.



During the restore operation:

- The broker service stops and restarts automatically.
 - Connected Agents are not restarted. Brokers and Agents temporarily disconnect and reconnect during the operation. Their status changes are displayed in the **Sites** and **Site Details** views.
 - A progress indicator appears in the lower-right corner of the Topology perspective while the operation is in progress.
5. After the restore operation completes, verify that the default broker certificate is active by viewing the broker certificate details.

For more information, see [Viewing the Broker Certificate](#).

Using Maintenance Mode to Update Broker Certificates

Overview

Maintenance Mode is intended for recovery scenarios where you need the broker to temporarily use the default TLS certificate and keystore so that Agents can reconnect. For example, an Agent may be unable to reconnect if it missed a broker certificate update.


When Maintenance Mode is enabled:

- The broker temporarily uses the default TLS certificate and keystore.
- The broker's menu option changes to allow exiting Maintenance Mode.

Enter Maintenance Mode


You can enter Maintenance Mode even if a job is running.

1. Open the **Topology** perspective.
2. In the **Sites** view, right-click the broker that you want to place in Maintenance Mode.
3. Select **Manage Broker Certificates**, and then select **Enter Maintenance Mode**.



The broker service restarts. While Maintenance Mode is active:

- The broker label in the **Sites** view includes the suffix - **Maintenance mode**.
- The broker status displays **Maintenance mode** with a yellow indicator.



Exit Maintenance Mode

1. Open the **Topology** perspective.
2. In the **Sites** view, right-click the broker that you want to remove from Maintenance Mode.
3. Select **Manage Broker Certificates**, and then select **Exit Maintenance Mode**.

The broker service restarts.

Downloading All Broker Certificates

Overview

Use this procedure to export all broker certificates to a Java truststore file (brokercerts.ts). The truststore contains the broker certificates required for Agents and remote brokers to establish trusted TLS communication.

When to Download

Downloading broker certificates is useful when:

- Deploying new Agents
- Manually distributing trusted broker certificates to remote hosts.

Before You Begin

Before downloading broker certificates:


- Verify that you have sufficient permissions to access the Topology perspective.
- Ensure that all required brokers are online and reachable.
- If brokers are configured to run in service mode, verify that the appropriate keystore files are available and properly configured.
- Determine where you want to store the downloaded certificate bundle so that it can later be copied to Agent systems if necessary.

For prerequisite information, see [Preparing for Certificate Management](#).

Download Broker Certificates

To download broker certificates:

1. Open the **Topology** perspective.
2. In the **Sites** view, click the **Options** menu.
3. Select **Download Broker Certificates**.



4. In the save dialog, browse to the location where you want to save the certificate bundle.
5. Click **Save**.

The PMC collects the current broker certificates and generates a certificate bundle containing the trusted broker certificates required by Agents.

After the download completes, the generated certificate file is automatically saved to the location that you selected.

Notes

- The generated truststore reflects the currently active broker certificates.
- Browser-based PMC users might receive a standard browser download prompt depending on browser settings.

Updating Trusted Broker Certificates

Overview

Use this procedure to update the trusted broker certificates on selected brokers and Agents. This operation distributes the current trusted broker certificate information to the selected brokers and Agents and updates their trusted certificate stores. Broker identity certificates are not modified.

When to Update Trusted Broker Certificates

Update trusted broker certificates in the following situations:

- An Agent or broker was offline when a new broker key and certificate were installed.
- An Agent or broker was installed after a new broker key and certificate were installed.

Important Considerations

Updating trusted broker certificates restarts the selected broker or Agent services.

Before proceeding:


- All affected components should be online and reachable.
- If any selected Agent has a running job, the operation is canceled and no certificate updates are applied.

Updating trusted broker certificates restarts the selected broker or Agent services. During the update process, communication between affected brokers and Agents is temporarily interrupted. Services reconnect automatically after the updated certificates are applied.


Updating Trusted Broker Certificates

To update trusted broker certificates:

1. Open the **Topology** perspective.
2. In the **Sites** view, select one or more brokers or Agents.
3. Right-click the selection and select **Update Trusted Broker Certificates**.



4. Review the confirmation message that lists the selected components that will be updated.



5. Click **Yes** to proceed.

During the update operation:

- Broker and Agent services stop and restart automatically.
- Selected brokers and Agents temporarily disconnect and reconnect during the operation. Their status changes are displayed in the **Sites** and **Site Details** views
- A progress indicator appears in the bottom-right corner of the Topology perspective while the update operation is in progress.

After the update operation completes, the selected brokers and Agents reconnect automatically and begin using the updated trusted broker certificates.

Managing Web Client Certificates

Overview

Web client certificates are used to establish trusted HTTPS communication between users' web browsers and the PMC web client. Managing these certificates allows you to secure browser-based connections, replace existing certificates, and maintain certificate trust relationships within your PMC environment.

You can manage web client certificates from the **Topology** perspective in the PMC. Certificate management operations help organizations align PMC security with internal certificate management policies.

Web client certificate changes can temporarily interrupt browser-based access while services restart and reload updated certificate information. Users might need to reconnect after the updated certificates are applied.

Before You Begin

Before managing web client certificates:

- Verify that you have sufficient permissions to access the **Topology** perspective.
- Ensure that users are aware of possible temporary reconnect prompts or browser trust warnings.
- If using CA-signed certificates, verify that the certificate matches the PMC hostname used by clients. OR If using CA-signed certificates, verify that certificates contain the correct PMC hostnames and required trust chain information.
- Back up existing certificate and keystore files before replacing or importing certificates.

For prerequisite information, see [Preparing for Certificate Management](#).

Common Web Client Certificate Tasks

The following certificate management tasks are available from the **Sites** view in the **Topology** perspective.

Task	Description
Generate Web Client Certificates	Creates new self-signed web client certificates for secure HTTPS communications.
Import Web Client Certificates	Replaces existing web client certificates with user-provided certificates and private keys.
Restore Web Client Certificates	Restores the default PMC web client certificates

For more information on performing these tasks, see:

- [Generating Self-Signed Web Client Certificates](#)
- [Importing Web Client Certificates](#)
- [Restoring Default Web Client Certificates](#)

Generating Self-Signed Web Client Certificates

Overview

Use this procedure to generate a new self-signed certificate for the web client when you need to enable or replace TLS/SSL encryption for web client connections.

When to Generate the Certificate

You might generate a new self-signed certificate in the following situations:

- Performing the initial setup for a test or development environment.
- Replacing an expired certificate.
- Updating the PMC hostname or domain information.

Important Considerations

Generating a new self-signed certificate affects web client connections that communicate through the broker.

Before generating a new self-signed certificate:

- [Stop all running jobs.](#)

Generating a new self-signed certificate has the following effects:


- Connected web clients might be temporarily disconnected during certificate regeneration.
- The related web services restart automatically.
- Certificates generated by this procedure are self-signed and are not issued by a trusted Certificate Authority (CA).

- The **Manage Web Client Certificates** option is available on a server that hosts both the PMC and a broker.


Generate the Certificate

To generate the certificate:


1. Open the **Topology** perspective.
2. In the **Sites** view, right-click the the PMC broker whose certificate you want to generate.
3. Select **Manage Web Client Certificates**, and then **Generate Self-Signed Key and Certificate**.



4. Click **Yes** in the restart confirmation dialog.



5. Enter the certificate details:




The dialog box titled "Generate Self-Signed Key and Certificate" contains the following fields and controls:

- Common Name (CN)*:
- Organization (O):
- Organizational Unit (OU):
- Locality (L):
- State (ST):
- Country (C):
- Validity Period*: Years: Months:
- Advanced Settings (collapsible section):
 - Key Algorithm:
 - Key Size:
- Buttons:

Field	Description
CN (Common Name)	Required. Enter the DNS name or hostname that users use to access the web client. This value should match the web client hostname to avoid certificate validation errors.
O (Organization)	Enter your company or organization name
OU (Organizational Unit)	Enter your department name (for example IT).
L (City/Locality)	Enter the city where your organization is located.

Field	Description
ST (State/Province)	Enter your state or province.
C (Country)	Select the two-letter country code (for example, US) where your organization is located.

- Set the certificate validity period in years and months
- Optional: Under **Advanced Settings**, specify values for the following fields:
 - Key Algorithm** - Select the algorithm used to generate the key pair, such as RSA or EC.
 - Key Size** - The default key size is 2048 bits. Optionally, select 4096 bits for stronger encryption.
- Click **Generate**.
- Review the confirmation dialog, which lists the certificate details and the affected services.



10. Click **Yes** to proceed.

During the operation:

- The related web services stop and restart automatically.
- Connected web client sessions might be temporarily disconnected.
- A progress indicator appears in the lower-right corner of the **Topology** perspective while the operation is in progress.

After the certificate is regenerated, the updated certificate is available for web client connections.

Importing Web Client Certificates

Overview

Use this procedure to import a web client certificate and private key for secure HTTPS communication between users' web browsers and the PMC web client. Importing a web client certificate updates the PMC web client keystore and restarts the PMC web service.

When to Import

You might import a web client certificate in the following situations:

- Replacing an expiring web client certificate
- Replacing a self-signed certificate with a certificate signed by a Certificate Authority (CA)
- Replacing a compromised certificate
- Updating certificate information after PMC hostname or domain changes

Before You Begin

Before importing web client certificates:


- Verify that you have sufficient permissions to access the Topology perspective.
- Notify users of possible temporary browser disconnects or certificate warnings during the update process.

- If using CA-signed certificates, verify that certificates contain the correct hostnames and required trust chain information.
- Back up existing certificate and keystore files before replacing certificates.

Importing the Web Client Certificate

To import a web client certificate:

1. Open the **Topology** perspective.
2. In the **Sites** view, right-click the **PMC**.
3. Select **Manage Web Client Certificates**, and then **Import User Key and Certificate**.



4. In the **Import** dialog, provide the required certificate information:
 - **Key File (PKCS12)** – Browse to the PKCS12 certificate file (.p12 or .pfx). After selecting the file, click **Open**. You can also enter the full file path manually.
 - **Password** – Enter the keystore password.
 - **Alias** – Enter the certificate alias contained in the keystore.
5. Click **Import**.

The PMC validates the password and alias before importing the certificate. If the password is incorrect or the alias is not found in the PKCS12 file, a Keystore Validation Error message appears. Verify the password and alias, and then try again.

6. In the **Confirmation** dialog, click **Yes** to continue.

During the import operation, a progress indicator appears in the bottom-right corner of the Topology perspective.

The PMC web client keystore is updated and the PMC web service restarts automatically. Users connected to the PMC web client might temporarily disconnect while the updated certificate is applied.

Restoring Default Web Client Certificates

Overview

Use this procedure to permanently replace the current web client keystore with the default keystore that was installed with PeerGFS. This operation returns the web client to the default PeerGFS-installed certificate configuration and restarts the PMC web service.

Important: This operation cannot be undone. Before restoring the default web client key and certificate, make sure this is the action you want to perform.


When to Restore the Default Web Client Key and Certificate

Restore the default web client key and certificate when you want to return to the default PeerGFS-installed web client certificate configuration after replacing the default certificate with a custom certificate.

Restore the Certificate

To restore the default web client key and certificate:

1. Open the **Topology** perspective.
2. In the **Sites** view, right-click the broker whose certificate you want to restore.
3. Select **Manage Web Client Certificates**, and then **Restore Default Key and Certificate**.



4. In the confirmation message, click **Restore**.

The default web client keystore replaces the current web client keystore. The PMC web service restarts automatically to apply the restored certificate configuration.

5. Verify that the default web client certificate is active.

Note: Verify the restored certificate from a web browser by connecting to the web client and viewing the certificate details.

Telemetry

Telemetry allows selected diagnostic, health, and performance information from a PeerGFS environment to be shared with Peer Software. This information helps Peer Software improve PeerGFS, identify trends, and provide better support for customer environments.

PeerGFS telemetry includes **Anonymous Diagnostic Data** and **Proactive Monitoring**.

Anonymous Diagnostic Data

Anonymous diagnostic data helps Peer Software improve PeerGFS by sharing anonymized information about product usage, configuration, and diagnostic conditions. This information helps Peer Software understand how PeerGFS is used and where product improvements may be needed.

Anonymous diagnostic data does not include customer-identifiable information.

Proactive Monitoring

Proactive Monitoring gives the Peer Software Technical Account Manager (TAM) team regular insight into the health and performance of a PeerGFS environment. It is intended to help the TAM team identify potential issues, review environment trends, and provide more informed assistance.

Proactive Monitoring can include information about PeerGFS Agents, jobs, and PMC activity. Collected data is securely uploaded to a Peer Software-owned Microsoft Azure storage account and is accessible only to authorized Peer Software employees.

Note: Proactive Monitoring is available only to customers with dedicated TAMs.

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	• Email Alerts • File and Folder Filters • Proxy Configuration.
File Collaboration, File Replication, and File Synchronization	• Email Alerts • File and Folder Filters
DFS-N Management	• Email Alerts • File and Folder Filters

Configuring Preferences

To modify 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 Job Preferences

The following job-specific preference settings are available for **Analytics** jobs:

- [Analytics](#)


- [Email Alerts](#)
- [SNMP Notifications](#)

Analytics

These settings control analytics reporting to PeerIQ, reporting intervals, and display options for **Analytics** jobs.

To configure these settings:

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



3. Modify the settings as needed.

Option	Description
Enable the sending of analytics data to PeerIQ	Sends analytics data to PeerIQ. Click Learn more for additional information about analytics reporting.
Automatic Reporting Interval (Seconds)	Specifies how often each Peer Agent reports statistics for Analytics jobs to PeerIQ. The default is 10 seconds.
Show Volumes in Jobs View	Displays volumes in the Jobs view.
24-hour format	Displays times in a 24-hour format rather than a 12-hour format.
Hide Internal Volumes	Hides internal volumes when choosing which which volumes to use in Analytics jobs.

4. Click **Apply** to save your changes.

Email Alerts

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


Tip: If you are performing maintenance, you can temporarily disable an email alert by clicking in the **Temporarily Disabled** column in the **Email Alerts** table and selecting **Yes**. No email alerts will be sent for that type of alert until you reenable the alert.

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

To create an email alert:


1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Analytics** in the navigation tree, and then select **Email Alerts**.

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



3. Click **Create**.

The **Create Email Alert** dialog appears.



The image shows a 'Create Email Alert' dialog box. It has a title bar with 'Create Email Alert' and window control buttons. Inside, there is a 'Name' field with 'New Email Alert' entered. Below it is a 'Temporarily Disable' checkbox. The 'Event Types' section contains five checkboxes: 'Job Start', 'Job Stop', 'Job Failure', 'Participant Failure', and 'System Event', all of which are checked. The 'Recipients' section has a text input field with the placeholder 'Enter email, contact, or distribution list:' and a button 'Add to List'. Below this is a list of recipients, currently empty, with 'Delete' and 'View Details' buttons next to it. At the bottom are 'OK' and 'Cancel' buttons.

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 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.

Event Type	Description
System Event	Sends a notification when a system event such as low memory or low hub disk space occurs.

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

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

7. Click **OK**.

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

Preferences

type filter text

- Analytics
 - Email Alerts**
 - SNMP Notifications
 - Cloud Backup and Replication
 - Collab, Sync, and Replication
 - DFS-N Management
 - Email Configuration
 - General Configuration
 - Licensing
 - MED Configuration
 - PMC Backup
 - Real-time Event Detection
 - SNMP Configuration
 - Storage Configuration
 - Task Scheduler
 - Telemetry
 - User Management

Email Alerts

Edit Email Configuration

Name	Temporarily Disabled	Event Types	Recipients
Participant Failure	No	Participant Failure	AriF@peersoftware.com
Job Failure	No	Job Failure	AriF@peersoftware.com
Job Start/Stop	No	Job Start, Job Stop	AriF@peersoftware.com
System Event	No	System Event	debraG@peersoftware.com

Create Edit Copy Delete

Apply and Close Cancel Apply

8. Click **Apply** to save your changes.

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


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 modify 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 **Analytics** in the navigation tree, and then select **SNMP Notifications**.

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



3. Click the **Create** button.

The **Create SNMP Notification** dialog appears.

A screenshot of a 'Create SNMP Notification' dialog box. The dialog has a blue title bar with the text 'Create SNMP Notification' and standard window controls. Inside, there is a text field labeled 'Name:' containing 'New SNMP Notification'. Below it is a checkbox labeled 'Temporarily Disable' which is unchecked. A section titled 'Event Types' contains five checkboxes: 'Job Start', 'Job Stop', 'Job Failure', 'Participant Failure', and 'System Event', all of which are checked. At the bottom are 'OK' and 'Cancel' buttons.

Create SNMP Notification

Name:

☐ Temporarily Disable

Event Types

☒ Job Start ☒ Job Stop ☒ Job Failure


☒ Participant Failure ☒ System Event

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.

5. Click **OK**.

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



6. Click **Apply** to save your changes.

Cloud Backup and Replication Job Preferences

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

The following job-specific preference settings are available for **Cloud Backup and Replication** jobs:

- [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

These settings control reporting intervals and display options for **Cloud Backup and Replication** jobs.

To configure 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.

Option	Description
Automatic Reporting Interval (Seconds)	Each Peer Agent automatically reports its statistics for Cloud Backup and Replication jobs 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 use in Cloud Backup and Replication jobs.

- Click **Apply** to save your changes.

Database Connections


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

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

To create a new database connection:

- Select **Open Preferences** from the **Tools** menu.
- Expand **Cloud Backup and Replication** in the navigation tree, and then select **Database Connections**.

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



3. Click the **Create** button.

The **Create Database Connection** dialog appears.

The 'Create Database Connection' dialog box is shown. It contains the following fields and options:

- Database Connection Name:** Name for this Database Connection configuration
- Management Agent:** (Dropdown menu)
- DB Host Name:** Hostname of the DB server
- Port:** The port to connect to the DB Server
- Instance Name:** SQL Server Instance name for named instance
- Database Name:** peercloud
- Authentication:** ☒ Integrated ☐ Credentials
- Username:** (Text field)
- Password:** (Text field)
- Buttons:** Validate, OK, Cancel


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.
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.

5. 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.



- Click **Apply** to save your changes.


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:


- Select **Open Preferences** from the **Tools** menu.
- 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.



The dialog box titled "Create Destination Credentials" features a blue header bar with standard window controls. Below the header, the "Platform:" section contains six radio button options, each with a corresponding icon: Microsoft Azure Blob Storage (selected), Amazon S3, NetApp StorageGRID, Nutanix Objects, Wasabi, and S3 Compatible. The lower section contains four labeled input fields: "*Description:", "*Account:", "*Shared Key:" (with a toggle icon), and "*Endpoint Type:" (a dropdown menu currently showing "Public"). A "Use SSL" checkbox is checked. At the bottom right are "OK" and "Cancel" buttons.

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** to save your changes.

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

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



3. Click the **Create** button.

The **Create Email Alert** dialog appears.



The image shows a 'Create Email Alert' dialog box with a blue title bar. It contains several sections: 'Name' with a text field containing 'New Email Alert'; a 'Temporarily Disable' checkbox; 'Event Types' with checkboxes for 'Job Start', 'Job Stop', 'Job Failure', 'Participant Failure', 'System Event', and 'Malicious Event'; 'Report Types' with checkboxes for 'Scan' and 'Destination Snapshot'; and 'Recipients' with a text input field, an 'Add to List' button, and a list area with 'Delete' and 'View Details' buttons. 'OK' and 'Cancel' buttons are at the bottom right.

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 .

6. (Optional) Select the report types to be sent.

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

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

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

8. Click **OK**.

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

Preferences

type filter text

- Analytics
- 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
- Collab, Sync, and Replication
- DFS-N Management
- Email Configuration
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- Storage Configuration
- Task Scheduler
- Telemetry
- User Management

Email Alerts

Edit Email Configuration

Name	Temporarily Disabled	Event Types	Recipients
Job Failure	No	Job Failure, Scan Report	debrag@peersoftware.com
Participant Failure	No	Participant Failure	MattM@peersoftware.com

Create Edit Copy Delete

Apply and Close Cancel Apply

9. Click **Apply** to save your changes.

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**.

Preferences

type filter text

- Analytics
- 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
- Collab, Sync, and Replication
- DFS-N Management
- Email Configuration
- General Configuration
 - Licensing
 - MED Configuration
 - PMC Backup
 - Real-time Event Detection
 - SNMP Configuration
- Storage Configuration
- Task Scheduler
- User Management

File Retries and Source Snapshots

Retry Options

Max Number of Retry Threads: 5

Max Number of Retries: 60

Retry Interval in Seconds: 30

Multi-part Upload Retry Count: 3

Multi-part Upload Retry Interval in Milliseconds: 500

Source Snapshot Replication Options

Max Number of Transfer Threads: 3

Max Number of Transfer Retries: 3

Transfer Retry Interval in Minutes: 30

Managed File Extensions: pst,mdf,ldf,ndf

Apply and Close Cancel Apply

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

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

Option	Description
Multi-part Upload Retry Count	Enter the maximum number of retries when performing multi-part upload.
Multi-part Upload Retry Interval in Milliseconds	Enter the number of minutes to wait between retries of multi-part uploads.

4. Modify the **Source Snapshot Replication Options** as needed:

Option	Description
Max Number of Transfer Threads	Enter the maximum number of threads available for replicating files from a source snapshot.
Max Number of Transfer Retries	Enter the maximum number of retries to perform on a file or folder that has failed to be replicated from a source snapshot. If the number of retries is exceeded, the file or folder will be added to the Failed Events view and will need to be manually processed.
Transfer Retry Interval in Minutes	Enter the number of minutes to wait between retries of the failed replication of a file or folder from a source snapshot.
Managed File Extensions	Enter the extensions for managed files that should be read from a source snapshot.

5. Click **Apply** to save your changes.


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.




The "Create File Filter" dialog box is shown with the following fields and controls:

- Name:** A text input field.
- Filter Type:** A dropdown menu currently set to "General".
- Auto Excluded:** A section with a link: [View file types that are automatically excluded](#).
- Excluded Patterns:** A large text area for entering patterns, with "Add", "Edit", and "Delete" buttons below it.
- Included Patterns:** A large text area for entering patterns, 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 with "0" and the unit "days".
- Excluded File Sizes:** A dropdown menu set to "None" and a text input field with "0" and the unit "bytes".
- Buttons:** "OK" and "Cancel" buttons at the bottom.


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.



A dialog box titled "Add Exclusion Pattern" with a blue header bar and a close button (X) in the top right corner. The main area contains the text "Enter a single file name or a single path pattern:" followed by a single-line text input field. Below the input field is a light gray message box that says "String cannot be empty." At the bottom right, there are two buttons: "OK" and "Cancel".

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.



A dialog box titled "Add Inclusion Pattern" with a blue header bar and a close button (X) in the top right corner. The main area contains the text "Enter a single file name or a single path pattern:" followed by a single-line text input field. Below the input field is a light gray message box that says "String cannot be empty." At the bottom right, there are two buttons: "OK" and "Cancel".

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** to save your changes.

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	Enter the maximum number of threads available for replicating files as they are updated in real-time on the source storage device.

Setting	Description
Replication Threads	
Max Number of Scan Replication Threads	Enter the maximum number of threads available for replicating files during scheduled and on-demand scans of the source storage device.
Max Number of Upload Threads	Enter the maximum number of threads available for uploading files to the destination storage device.
Max Number of Restore Threads	Enter the maximum number of threads available for restoring from the destination storage device.

4. Click **Apply** to save your changes.


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 'Create Proxy Configuration' dialog box contains the following fields and controls:

- Agent:** A dropdown menu.
- IP Address:** A text input field.
- Port:** A text input field.
- ☐ **Use Authentication**
- Domain:** A text input field (disabled if 'Use Authentication' is not checked).
- Username:** A text input field (disabled if 'Use Authentication' is not checked).
- Password:** A text input field (disabled if 'Use Authentication' is not checked).
- ☐ **Show Password**
- 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

6. If your proxy server requires authentication, select the **Use Authentication** checkbox, and then supply the necessary values:

Field	Description
Domain	Enter the domain name on the proxy server.
Username	Enter the user name for the proxy server.
Password	Enter the password for the proxy server.

7. Click **OK**.

The new proxy configuration is listed in the **Proxy Configuration** table.



8. Click **Apply** to save your changes.

Replication and Retention Policies

Each Cloud Backup and Replication job must have a Replication and Retention Policy applied to it. When you create a job, you can select an existing policy to apply to the job or you can create a new policy and apply it to the job.

You can modify and delete a policy, however, you cannot:

- Modify a policy while it is applied to a running job.
- Delete a policy while it is applied to any job.

For more information about policies, see [Step 11: Replication and Retention Policy](#) in [Creating a Cloud Backup and Replication Job](#).

To create a new replication and retention policy:

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

Any existing policies are listed in the table.

The screenshot shows the 'Preferences' window with the 'Replication and Retention Policies' section selected in the sidebar. The main area displays a table of existing policies. The table has the following columns: Name, Backup Type, Backup Policy, Retention Configured, and In Use. There is one row listed: Policy-1, Destination Snapshots, Weekly, Yes, and Yes. To the right of the table are buttons for 'Create', 'Edit', 'Copy', and 'Delete'. At the bottom of the window are buttons for 'Apply and Close', 'Cancel', and 'Apply'.

Name	Backup Type	Backup Policy	Retention Configured	In Use
Policy-1	Destination Snapshots	Weekly	Yes	Yes

3. Click the **Create** button.

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

Replication and Retention Policy Wizard

Replication and Retention Policy

✖ You must enter a name for the policy.

Replication and Retention Policy *Name:

Replication Schedule ☒ Enable Backup with Destination Snapshots

Retention

Source 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** to save your changes.


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 **Create SNMP Notification** dialog appears.




The image shows a 'Create SNMP Notification' dialog box. It has a title bar with a close button. Inside, there is a text field for 'Name' containing 'New SNMP Notification'. Below it is a checkbox for 'Temporarily Disable'. A section titled 'Event Types' contains six checkboxes: 'Job Start', 'Job Stop', 'Job Failure', 'Participant Failure', 'System Event', and 'Malicious Event'. All six checkboxes are checked. At the bottom right are 'OK' and 'Cancel' buttons.

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 Event	Sends a notification when a malicious event is detected. For more information, see MED Configuration .

5. Click **OK**.

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




6. Click **Apply** to save your changes.

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.

Setting	Description
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** to save your changes.

Collaboration, Replication, and Synchronization Job Preferences

The following job-specific preference settings are available for **File Collaboration, File Synchronization, and File Replication** jobs:


- [Collab Sync, and Replication](#)
- [Application Enhancements](#)
- [DFS-N Management](#)
- [Edge Caching](#)
- [Email Alerts](#)
- [File Retries](#)
- [File and Folder Filters](#)
- [Locking](#)
- [Performance](#)
- [Real-time Event Detection](#)
- [Revit Enhancements](#)
- [SNMP Notifications](#)
- [Scan Manager](#)
- [Scheduled Replication Filters](#)

Collab, Sync, and Replication

These settings control host reconnect behavior, advanced reporting, and real-time tracking for **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	Enables Event Detection Analytics to track and report common activity processed by Peer Global File Service. If enabled, every 24 hours, an Excel-based report will be written to disk that shows top folders, files, extensions, and users by total processed activity over the previous 24-hour window. These reports are stored under the installation folder of Peer Management Center and can be reviewed by Peer Support when uploading log files.

4. Click **Apply** to save your changes.

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](#) 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**.

Preferences

type filter text

- Analytics
- Cloud Backup and Replication
- Collab, Sync, and Replication
 - Application Enhancements**
 - DFS-N Management
 - Edge Caching
 - Email Alerts
 - File Retries
 - File and Folder Filters
 - Locking
 - Performance
 - Real-time Event Detection
 - SNMP Notifications
 - Scan Manager
 - Scheduled Replication Filters
- DFS-N Management
- Email Configuration
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- Storage Configuration
- Task Scheduler
- User Management

Application Enhancements

Sync On Save Override File Extensions:

Sync On Save Override Delay:

Target Sharing Violation File Extensions:

Bulk Add Minimum Rejected Event Threshold:

Retry Quarantines File List:

Last Write Override Extensions:

Expedited Fast Sync File List:

Expedited Slow Sync File List:

Direct Target Write List:

Apply and Close Cancel Apply

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 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.

- Click **Apply** to save your changes.


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** to save your changes.

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](#)


Edge Caching

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** to save your changes.

Email Alerts

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.


- Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Edge Caching**.
- Select **Email Alerts**.

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



- Click **Create**.

The **Create Email Alert** dialog appears.



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.


Event Type	Description
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.

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

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

8. Click **OK**.

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



9. Click **Apply** to save your changes.


Master Data Service

When you create a job that is Edge Caching-enabled, you specify the settings to be used for the Peer Master Data Service. The Peer Master Data Service handles requests from edge participants for files on a master participant. The Peer Master Data Service is installed as part of the Peer Agent installation process. The **Master Data Service** page displays the existing parameters for the Master Data Service.

To edit the Master Data Service configuration:


1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Edge Caching**.
3. Select **Master Data Service**.

The **Master Data Service** table lists master participants. The **In Use** column identifies whether the participant is currently being used as a master participant in a job.



4. Select a participant, and then click **Edit**.


If you selected a master participant currently being used in a job, the **In Use Configuration** dialog appears. Click **OK** to close the dialog.



Otherwise, the **Edit Master Data Service** dialog appears. The following fields are automatically populated:

- **Protocol:** The protocol used to transfer file content between master and edge participants. HTTPS is currently the only available option because it requires only one open firewall port on each master participant and handles latency effectively over WAN connections.

- **Agent Name:** The Management Agent selected at the beginning of Step 2.



The screenshot shows a dialog box titled "Edit Master Data Service". It contains the following fields and values:

Field	Value
Protocol:	HTTPS
Agent Name:	Lab24Win19b
Agent Alias:	
Port:	8446

At the bottom of the dialog are two buttons: "OK" and "Cancel".

5. (Optional) In the **Agent Alias** field, enter a hostname, FQDN, or IP address if the Agent name cannot be resolved to an IP address through DNS.

The Edge Service on edge participants uses the alias to connect with the Master Data Service. If you do not enter an alias, the Agent name is used.

6. (Optional) In the **Port** field, modify the port on which the Master Data Service listens for this master participant.

The default port, 8446, is set when the Agent installed. If you change the port, the Master Data Service restarts using the new port.

7. Click **OK**.

The revised Master Data Service is listed in the **Master Data Service** table.

8. Click **Apply** to save your changes.

The new settings will be applied to all Edge Caching-enabled jobs.

Pinning Filters

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.



The dialog box titled "Create Pinning Filter" has a blue header bar with a close button (X) in the top right corner. Below the header, there is a text field labeled "*Name:" which is currently empty. Underneath this is a section labeled "Pinning Rules:" containing a table 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". At the bottom right of the dialog are "OK" and "Cancel" buttons.

Path	Pinning State
------	---------------


5. Enter a name for the pinning filter.



The dialog box titled "Create Pinning Filter" is shown again, but now the "*Name:" text field contains the text "AutoCAD Files". The "Pinning Rules:" table remains empty. The "Create", "Edit", "Delete", "OK", and "Cancel" buttons are still present in their respective positions.


Path	Pinning State
------	---------------

6. Click **Create** to add a pinning rule to the filter.



The dialog box titled "Create Pinning Rule" has a blue header bar with standard window controls. Below the header, the text "Enter a file name or path and choose pinning state." is displayed. There are two input fields: "Pattern:" followed by an empty text box, and "Pinning State:" followed by a dropdown menu currently showing "Local at Edge". At the bottom right, there are two buttons: "OK" and "Cancel".


7. Enter a file name or enter a path.



This dialog box is identical to the previous one, but the "Pattern:" text box now contains the text "*.dwg". The "Pinning State:" dropdown remains set to "Local at Edge".

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.




The image shows a 'Create Pinning Filter' dialog box. At the top, there is a text field labeled '*Name:' containing the text 'AutoCAD Files'. Below this, the text 'Pinning Rules:' is followed by a table. The table has three columns: 'Path', 'Pinning State', and an empty column. The first row of the table contains the text '* .dwg' under 'Path' and 'Local at Edge' under 'Pinning State'. To the right of the table are three buttons: 'Create', 'Edit', and 'Delete'. At the bottom right of the dialog box are two buttons: 'OK' and 'Cancel'.

Path	Pinning State	
* .dwg	Local at Edge	

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.



12. Click **Apply** to save your changes.


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 modify 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 **Collab, Sync, and Replication** in the navigation tree, and then select **Edge Caching**.
3. Select **SNMP Notifications**.

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



4. Click the **Create** button.

The **Create Edge Caching SNMP Notification** dialog appears.




5. Select the types of events that will trigger the generation of an SNMP trap:

Event Type	Description
Scan Start/End	Sends a notification when a caching scan is started or stopped.
Scan Fatal Errors	Sends a notification when a fatal error occurs during a caching scan.
Scan Errors	Sends a notification when errors occur during a caching scan.
Cache Size Exceeded	Sends a notification when the amount of volume disk space used by 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.



7. Click **Apply** to save your changes.

Utilization Policies


When you create an Edge Caching-enabled job, you can select an existing utilization policy to apply to the job or create a new utilization policy and apply it to the job. This [Preferences](#) page lists the existing utilization policies. From this page, you can create, edit, and delete utilization policies. However, you cannot edit or delete a utilization policy while it is applied to a job.

A **utilization policy** defines when a file should be stubbed versus fully hydrated across all volumes of an edge participant. The policy parameters are based on the size of the files to be potentially stubbed and when they were last accessed and modified. A utilization policy enables you to balance getting the best performance while keeping the cache as full as possible.

To create a utilization policy:


1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Edge Caching**.
3. Select **Utilization Policies**.

Any existing utilization policies are listed in the **Utilization Policies** table.



4. Click **Create**.

The **Create Utilization Policy** dialog appears.



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

OK Cancel

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 if not accessed within the past X time period	Select this option and specify a time range to use the last accessed date of a file to determine if it should be kept local or stubbed.

9. (Optional) in the **Stubbing Override** section, select the checkbox and a time period if you want to use the last accessed date of all recently hydrated files to determine if they should be kept local or re-stubbed.
10. Click **OK**.

The new policy is listed in the **Utilization Policies** table and can now be applied to jobs.



11. Click **Apply** to save your changes.

Volume Policies

When you create an Edge Caching-enabled job, you specify a volume policy for an edge participant. A **volume policy** specifies how much of the available space on the volume monitored by the edge participant to assign for local (hydrated) files. This space is referred to as a **cache**. The volume refers to the drive letter of the path set on the **Path** page of the wizard (for example, if the participant is configured to monitor D:\Data, the volume policy for this participant would apply to the D volume).

If the Agent you selected is already being used as an edge participant in another job utilizing 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 "touch" the same volume.


From this page, you can edit and delete volume policies. However, note that:

- Any changes made to a volume will not be applied to a running job until the job is restarted.
- You cannot delete a volume policy while it is being used by an edge participant.


To edit a volume policy:


1. Select **Open Preferences** from the **Tools** menu.
2. Expand **Collab, Sync, and Replication** in the navigation tree, and then select **Edge Caching**.
3. Select **Volume Policies**.

Any existing volume policies are listed in the **Volume Policies** table.




4. Select the policy you want to modify, and then click **Edit**.
5. If you have created an Edge Caching email alert or an Edge Caching SNMP notification but not saved it, you will be prompted to save it.





6. If you selected a policy currently being used in a job, the **In Use Configuration** message appears.



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

The **Edit Volume Policy** dialog appears.

Edit Volume Policy (Lab24Win22a - 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 512 MB
Cache usage exceeds 80 % of the cache size

Caching Scan Schedule

☒ Scan every 1 day(s) at 22 : 00 : 00
☐ Define schedule

☒ Run caching scan after job start

*Temporary Storage Path: C:\.PeerTempPath Browse
C:\.PeerTempPath\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 (Lab24Win22a - C:\)

Selecting a different utilization policy will affect all jobs with the same Agent and volume.

Utilization Policy:

Current Configuration

Agent: Lab24Win22a
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

The screenshot shows the 'Edit Volume Policy' window for 'Lab24Win22a - C:\'. The 'Volume Policy' tab is selected, and the 'Cache Size' section is highlighted with a red box. The 'Cache Size' section has two radio buttons: 'Use up to 75 % of this volume' (selected) and 'Use up to 10 GB of this volume'. Below this, the 'Cache Threshold Alerts' section shows 'Send alerts when:' with 'Disk space is less than 512 MB' and 'Cache usage exceeds 80 % of the cache size'. The 'Caching Scan Schedule' section has 'Scan every 1 day(s) at 22 : 00 : 00' selected. At the bottom, the '*Temporary Storage Path' is 'C:\.PeerTempPath' with a 'Browse' button.

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 (Lab24Win22a - 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 GB of this volume

Cache Threshold Alerts
Send alerts when:
Disk space is less than MB
Cache usage exceeds % of the cache size

Caching Scan Schedule
☒ Scan every day(s) at
☐ Define schedule
☒ Run caching scan after job start

*Temporary Storage Path:
C:\.PeerTempPath\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.

Edit Volume Policy (Lab24Win22a - 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 512 MB
 Cache usage exceeds 80 % of the cache size

Caching Scan Schedule
☒ Scan every 1 day(s) at 22 : 00 : 00
☐ Define schedule
☒ Run caching scan after job start

*Temporary Storage Path: C:\.PeerTempPath Browse
 C:\.PeerTempPath\.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 (Lab24Win22a - 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

☒ Run caching scan after job start

***Temporary Storage Path:**

C:\.PeerTempPath\PeerTempPath

13. Click **Next**.

14. (Optional) Select one or more email alerts to be associated with this volume policy, and then click **Next**.

Edit Volume Policy (Lab24Win22a - C:\)

Email Alerts
Select email alerts.

Volume Policy
Email Alerts
SNMP Notifications

Name	Temporarily Disabled	Event Type	Recipients
------	----------------------	------------	------------

Select
Delete
View Details

< Back Next > Finish Cancel

15. (Optional) Select one or more SNMP notifications to be associated with this volume policy.

Edit Volume Policy (Lab24Win22a - C:\)

SNMP Notifications
Select SNMP notification.


Volume Policy
Email Alerts
SNMP Notifications

Name	Temporarily Disabled	Notifications
☐ Low Disk Space	No	Low Disk Space
☒ Fatal Errors in Edge Caching	No	Scan Fatal Errors

< Back Next > Finish Cancel


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** to save your changes.

Email Alerts

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

Tip: If you are performing maintenance, you can temporarily disable an email alert by clicking in the **Temporarily Disabled** column in the **Email Alerts** table and selecting **Yes**. No email alerts will be sent for that type of alert until you reenable the alert.

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

To create an email alert:

1. Select **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.

Preferences

type filter text

- Analytics
- Cloud Backup and Replication
- Collab, Sync, and Replication
 - Application Enhancements
 - DFS-N Management
 - Edge Caching
 - Email Alerts**
 - File Retries
 - File and Folder Filters
 - Locking
 - Performance
 - Real-time Event Detection
 - SNMP Notifications
 - Scan Manager
 - Scheduled Replication Filters
- DFS-N Management
- Email Configuration
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- Storage Configuration
- Task Scheduler
- User Management

Email Alerts

Edit Email Configuration


Name	Temporarily Disabled	Event Types	Recipients
Participant Failure	No	Participant Failure	AriF@peersoftware.com
Participant Reconnect	No	Participant Reconnect	AriF@peersoftware.com

Create Edit Copy Delete

Apply and Close Cancel Apply

3. Click **Create**.

The **Create Email Alert** dialog appears.



The **Create Email Alert** dialog box is shown with the following fields and options:

- Name:** New Email Alert
- ☐ Temporarily Disable
- Event Types**
 - ☐ Job Start
 - ☐ Job Stop
 - ☒ Job Failure
 - ☒ Participant Failure
 - ☒ Participant Reconnect
 - ☒ File Quarantine
 - ☒ Scan Error
 - ☒ Malicious Event
- Queued Items**
 - ☐ Number of Queued Items
 - Exceeds: 5000 Items
 - Recovers Below: 1000 Items
 - ☐ Size of Queued Items:
 - Exceeds: 10240 MB
 - Recovers Below: 1024 MB
 - ☐ Alert on Recovery
- Reports**
 - ☐ Scan
- Batch Email Alerts**
 - ☐ Quarantined Files
- Recipients**
 - Enter email, contact, or distribution list:
 - Start typing to filter contacts/lists or add a new email
 -
 - Recipients:
 -
 -

Buttons:

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 .

- (Optional) Select options in the **Queued Items** section.

Option	Description
Number of Queued Items section	Select this checkbox if you want alerts to be sent regarding 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. If you select this checkbox, you need to enter two values that work in tandem: • Exceeds X Items - Enter the highest number of queued items before an email alert is sent. The default value is 5000. • Recovers Below X Items - Enter a value. The default value is 1000. An alert is sent the first time that the Queued Items counter has items greater than the value set in Exceeds X Items. The counter's value is displayed in the Queued Items column in the Collab, Sync, and Repl Summary view. The counter's value is a combination of the Real-time and File Sync queues. Another alert will not be sent until the Queued Items counter has dropped below the Recovers Below x Items value and then exceeds the Exceeds X Items value again. This prevents multiple or redundant alerts from being sent.
Alert on Recovery	Select this option if you want an alert to be sent when the number of queued items has fallen below the Recovers Below value.
Size of Queued Items section	Select this if you want alerts to be sent based on the total data size of queued items for a job. This is useful to notify you about when there is a queue backlog potentially due to bandwidth issues. If you select this checkbox, you need to enter two values that work in tandem: • Exceeds X MB - Enter the highest number of queued items before an email alert is sent. The default value is 10240 MB. • Recovers Below X MB - Enter a value. The default value is 1024 MB. An alert is sent the first time that the Pending Bytes for a job has items greater than the value set in Exceeds X MB. The counter's value is displayed in the Pending Bytes column in the Collab, Sync, and Repl Summary view. Another alert will not be sent until the Pending Bytes counter has dropped below the Recovers Below x MB value and then exceeds the Exceeds X MB value again. This prevents multiple or redundant alerts from being sent.

7. (Optional) Select the **Scan** checkbox in the **Reports** section if you want scan statistics emailed to you after a scan has completed.
8. (Optional) Select the **Quarantined Files** checkbox in the **Batch Email Alerts** section if you want email alerts about quarantined files sent to you in batches.

The default is a maximum of 1000 alerts with a quiet period of 60 seconds. To change the number of alerts and quiet period time, modify the values for **Batch Email Alerts for Quarantined Files** in [Email Configuration](#).

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

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

10. Click **OK**.

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

The screenshot shows the 'Preferences' dialog box with the 'Email Alerts' section selected. The 'Email Alerts' table is displayed with the following data:

Name	Temporarily Disabled	Event Types	Recipients
File Quarantine	No	File Quarantine, Batch Qu	AriF@peersoftware.com
Participant Failure	No	Participant Failure	AriF@peersoftware.com
Participant Reconnect	No	Participant Reconnect	AriF@peersoftware.com

The 'Edit Email Configuration' button is visible above the table. The 'Create', 'Edit', 'Copy', and 'Delete' buttons are visible to the right of the table. The 'Apply and Close', 'Cancel', and 'Apply' buttons are at the bottom of the dialog.

11. Click **Apply** to save your changes.


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

File Retries

File retries settings enable you to configure the frequency of attempts and the maximum number of attempts. These settings apply to all File Collaboration, File Replication, and File Synchronization jobs. For more information about file retries, see [Conflicts, Retries, and Quarantines](#).

To modify the file retries settings:

1. Select **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.

- Click **Apply** to save your changes.

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:

- Select **Open Preferences** from the **Tools** menu.
- 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.

Preferences

type filter text

- » Analytics
- » Cloud Backup and Replication
- » Collab, Sync, and Replication
 - Application Enhancements
 - DFS-N Management
 - » Edge Caching
 - Email Alerts
 - File Retries
 - File and Folder Filters**
 - Locking
 - Performance
 - Real-time Event Detection
 - SNMP Notifications
 - Scan Manager
 - Scheduled Replication Filters
- » DFS-N Management
- » Email Configuration
- » General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- » Storage Configuration
- » Task Scheduler
- User Management


File and Folder Filters

Name	Type	Exclusions	Inclusions	Date Filter	Size Filter	
Android Open Source Proj	General	*.class, *.lock, *.ninjab	None Selected	Include all date	None	
File Collaboration Sync O	Synchronization	None Selected	*.LOG, *.EXE, *.DLL, *.O	Include all date	None	
Invalid Characters	General	<<.*[.]\$>>	None Selected	Include all date	None	
Linux Default	General	.~lock.*#, *.swp, .goutp	None Selected	Include all date	None	
Linux Kernel Build Excl	General	*./rustup/*, *./nix/*, *./	None Selected	Include all date	None	
Linux User Profile Excl	General	/home/*/.local/share/Tra	None Selected	Include all date	None	
Locking Only	Locking Only	None Selected	*	Include all date	None	
MacOS Exclusions	General	*__MACOSX, *\Tempor	None Selected	Include all date	None	
Relay Filter	General	\Relay Intermediate Child	None Selected	Include all date	None	
Synchronizing Only	Synchronization	None Selected	*	Include all date	None	
User Profile Exclusions	General	*\Safe Browsing*, *\log*	None Selected	Include all date	None	
Windows Default	General	*.BAK, *.BCK, *.WBK, *.	None Selected	Include all date	None	

Create Edit Copy Delete Update

Apply and Close Cancel Apply

3. Click **Create**.



The "Create File Filter" dialog box is shown with a blue title bar. It contains the following fields and controls:

- Name:** A text input field.
- Filter Type:** A dropdown menu currently set to "General".
- Auto Excluded:** A section with a link: [View file types that are automatically excluded](#).
- Excluded Patterns:** A large text area for entering patterns, with "Add", "Edit", and "Delete" buttons below it.
- Included Patterns:** A large text area for entering patterns, with "Add", "Edit", and "Delete" buttons below it.
- Included Last Modified Dates:** A dropdown menu set to "Include all dates" and a numeric input field set to "0" with the label "days".
- Excluded File Sizes:** A dropdown menu set to "None" and a numeric input field set to "0" with the label "bytes".
- Buttons:** "OK" and "Cancel" buttons at the bottom.

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** to save your changes.

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

type filter text

File and Folder Filters

Name	Type	Exclusions	Inclusions	Date Filter	Size Filter
Android Open Source Pro	General	*.class, *.lock, *.ninjab	None Selected	Include all date	None
File Collaboration Sync On	Synchronization	None Selected	*.LOG, *.EXE, *.DLL, *.O	Include all date	None
Invalid Characters	General	<<.*[.]\$>>	None Selected	Include all date	None
Linux Default	General	.~lock.*#, *.swp, .goutp	None Selected	Include all date	None
Linux Kernel Build Exclusion	General	*./rustup/*, */.nix/*, */.	None Selected	Include all date	None
Linux User Profile Exclusion	General	/home/*/.local/share/Tra	None Selected	Include all date	None
Locking Only	Locking Only	None Selected	*	Include all date	None
Log Files	General	*.log	None Selected	Include all date	None
MacOS Exclusions	General	*__MACOSX, *\Tempor	None Selected	Include all date	None
Relay Filter	General	\Relay Intermediate Child	None Selected	Include all date	None
Synchronizing Only	Synchronization	None Selected	*	Include all date	None
User Profile Exclusions	General	*\Safe Browsing*, *\log*	None Selected	Include all date	None
Windows Default	General	*.BAK, *.BCK, *.WBK, *	None Selected	Include all date	None


Apply and Close Cancel Apply

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** to save your changes.

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** to save your changes.

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 Queue 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** to save your changes.


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 modify 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 **Collab, Sync, and Replication** in the navigation tree, and then select **SNMP Notifications**.

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



3. Click the **Create** button.

The **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 the job is aborted because of lack of quorum due to one or more failed participants.
Participant Failure	Sends a notification when a participant timeout occurs, and the participant is taken out of session.
Participant Reconnect	Sends a notification when a participant reconnects to the job and the job resumes with the reconnected participant.
File Quarantine	Sends a notification when a file is marked as quarantined because a file conflict was not able to be resolved.
Scan Error	Sends a notification when an error occurs during the initial synchronization process .
Malicious Event	Sends a notification when Peer MED detects potentially malicious activity. For more information, see MED Configuration .

5. (Optional) Select options in the **Queued Items** section.

Option	Description
Number of Queued Items section	Select this checkbox if you want alerts to be sent regarding 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. If you select this checkbox, you need to enter two values that work in tandem: • Exceeds X Items - Enter the highest number of queued items before an email alert is sent. The default value is 5000. • Recovers Below X Items - Enter a value. The default value is 1000. 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 Collab, Sync, and Repl Summary view. The counter's value is a combination of the Real-time and File Sync queues. Another notification will not be sent until the Queued Items counter has dropped below the Recovers Below x Items value and then exceeds the Exceeds X Items value again. This prevents multiple or redundant alerts from being sent.
Alert on Recovery	Select this option if you want a notification to be sent when the number of queued items has fallen below the Recovers Below value.
Size of Queued Items section	Select this if you want notifications to be sent based on the total data size of queued items for a job. This is useful to notify you about when there is a queue backlog potentially due to bandwidth issues. If you select this checkbox, you need to enter two values that work in tandem. • Exceeds X MB - Enter the highest number of queued items before an email alert is sent. The default value is 10240 MB. • Recovers Below X MB - Enter a value. The default value is 1024 MB. An alert is sent the first time that the Pending Bytes for a job has items greater than the value set in Exceeds X MB. The counter's value is displayed in the Pending Bytes column in the Collab, Sync, and Repl Summary view. Another alert will not be sent until the Pending Bytes counter has dropped below the Recovers Below x MB value and then exceeds the Exceeds X MB value again. This prevents multiple or redundant alerts from being sent.

6. Click **Apply** to save your changes.

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** to save your changes.


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.
2. 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.



3. Click **Create**.




The dialog box is titled "Create Scheduled Replication Filter". It contains the following fields and controls:

- Name:** A text input field.
- Filter Type:** A dropdown menu currently showing "Scheduled Replication".
- Included Patterns:** A large empty text area for listing patterns. Below it are three buttons: "Add", "Edit", and "Delete".
- Scheduling Options:**
 - ☒ **Process every:** A dropdown menu showing "1 minute".
 - ☐ **Process on a schedule**
 - ☐ **Daily**
 - ☒ **Weekly**
 - Day(s):** A list of days with checkboxes: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday.
 - Time:** A dropdown menu showing "None".

At the bottom right are "OK" and "Cancel" buttons.

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](#).

A screenshot of a Windows-style dialog box titled "Add Inclusion Pattern" with a close button (X) in the top right corner. The dialog has a light gray background. Inside, the text "Enter a file name or path." is displayed above a single-line text input field. Below the input field, the text "String cannot be empty." is shown in a smaller font. At the bottom right of the dialog, there are two buttons: "OK" and "Cancel".

Add Inclusion Pattern

Enter a file name or path.

String cannot be empty.

OK Cancel

6. Click **OK**.

The pattern appears in the **Included Patterns** field.

Create Scheduled Replication Filter

Name:

Filter Type:

Included Patterns

*.vsx

Add Edit Delete

Scheduling Options

☒ Process every:

☐ Process on a schedule

☐ Daily ☒ Weekly

Day(s):

☐ Sunday

☐ Monday

☐ Tuesday

☐ Wednesday

☐ Thursday

☐ Friday

☐ Saturday

Time:

OK Cancel


7. Select a scheduling option:

- **Interval** - Process at a specified interval.

- **Schedule** - Process at a scheduled time.

8. After selecting a scheduling option, click **OK**.

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



9. Click **Apply** to save your changes.

DFS-N Management Job Preferences

The following job-specific preference settings are available for **DFS-N Management jobs**:

- [DFS-N Management](#)
- [Email Alerts](#)
- [SNMP Notifications](#)


DFS-N Management

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

These settings control DFS-N root location, namespace checking, timeout behavior, and DFS-N resource display options for **DFS-N Management** jobs.

To modify settings for [DFS-N Management jobs](#):

1. Select **Open Preferences** from the **Tools** menu.
2. Select **DFS-N Management** in the navigation tree.



3. Modify settings as needed.

Setting	Description
DFS Roots Folder	The default folder for namespace roots is C:\DFSRoots\. To use a different folder , 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 detects changes made to a namespace using the Microsoft DFS Management tool. Selecting a low value negatively affects performance but reflects 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 to check 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	Installs Microsoft's DFS-N Management tools when creating or importing a namespace.
Show Resources	Displays individual namespace folders under each namespace in the Jobs view.

- Click **Apply** to save your changes.

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.

Create Email Alert

Name:

☐ Temporarily Disable

Event Types

☒ Job Start ☒ Job Stop ☒ Job Failure ☒ Participant Failure

☒ Participant Reconnect

DFS-N Event Types

☒ Namespace Offline ☒ Namespace Not Found ☒ Folder Target Offline ☒ All Folder Targets Offline

☒ DFS Server Offline

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.
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.



- Click **OK**.

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

- Click **Apply** to save your changes.

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:


- Select **Open Preferences** from the **Tools** menu.
- Expand **DFS-N Management** in the navigation tree, and then select **SNMP Notifications**.

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



3. Click the **Create** button.

The **Create SNMP Notification** dialog appears.



4. Select the types of events that will trigger the generation of an SNMP trap.

Event Type	Description
Job Start	Sends a notification when the DFS-N Management job starts.


Event Type	Description
Job Stop	Sends a notification when the DFS-N Management job stops.
Job Failure	Sends a notification when the DFS-N Management job stops unexpectedly.
Participant Failure	Sends a notification when the Management Agent of the DFS-N Management job disconnects or stops responding.
Participant Reconnect	Sends a notification when the Management Agent of the DFS-N Management job reconnects.

5. Select the DFS-N event types that will trigger the generation of an SNMP trap.

Event Type	Description
Namespace Offline	Sends a notification when a namespace goes offline.
Namespace Not Found	Sends a notification when a namespace is not found.
Folder Target Offline	Sends a notification when a folder target goes offline.
All Folder Target Offline	Sends a notification when all folder targets go offline
DFS Server Offline	Sends a notification when a DFS server goes offline.

- Click **OK**.

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



- Click **Apply** to save your changes.

Email Configuration

Use the **Email Configuration** page to control whether Peer Management Center sends email notifications and to specify how it connects to an email provider. You can also configure how alerts for quarantined files are batched and sent.

The page contains the following groups:

- **Authentication** — Controls which method Peer Management Center uses to authenticate with the email provider and determines which configuration fields are displayed.
- **Email Configuration** — Displays the fields required for the selected authentication method (Basic or OAuth 2.0).
- **Batch Email Alerts for Quarantined Files** — Controls how quarantined-file alerts are batched and delivered.

To configure email settings:

- Select **Open Preferences** from the **Tools** menu.

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

The screenshot shows the 'Email Configuration' window. On the left is a navigation tree with the following items: Analytics, Cloud Backup and Replic, Collab, Sync, and Replic, DFS-N Management, **Email Configuration** (highlighted), General Configuration, Licensing, MED Configuration, PMC Backup, Real-time Event Detectic, SNMP Configuration, Storage Configuration, Task Scheduler, and User Management. The main panel is titled 'Email Configuration' and contains the following settings:

- ☐ Enable Email Notifications
- Authentication: ☒ Basic, ☐ OAuth 2.0
- *SMTP Host: [text field]
- *SMTP Port: [25] [spinners]
- Encryption: ☐
- *Encryption Type: [TLS] [dropdown]
- *Username: [text field]
- *Password: [text field]
- *Sender Email: [text field]
- Use Recommended Office 365 Settings: ☐
- [Send Test Email] button
- Batch Email Alerts for Quarantined Files:
 - Wait Before Sending (seconds): [60] [spinners]
 - Maximum Alerts per Email: [1000] [spinners]

At the bottom of the window are three buttons: 'Apply and Close', 'Cancel', and 'Apply'.

3. At the top of the page, select **Enable Email Notifications**.
 - When this checkbox is **unchecked**, all email-related sections appear dimmed and Peer Management Center does not send any email messages.
 - When it is checked, the authentication and configuration fields become active.
4. Under **Authentication**, select one of the following methods.

Option	Description
Basic	Uses standard user name and password credentials authentication. Selecting this option displays the Email Configuration section for Basic authentication.
OAuth 2.0	Uses modern OAuth2 authentication with provider authorization. Selecting this option displays the Email Configuration section for OAuth 2.0, including provider-specific fields and buttons to authorize and test the connection.

4. Complete the configuration fields for the selected authentication method:

- If you selected **Basic**, see [Basic Email Configuration](#).
- If you selected **OAuth 2.0**, see [OAuth2 Email Configuration](#).

5. (Optional) Complete the **Batch Email Alerts for Quarantined Files** section.

Use this section to define how alerts for quarantined files are batched and sent when email notifications are enabled.

Field	Description
Wait Before Sending (seconds)	Enter the number of seconds Peer Management Center waits after the last alert before sending the batch.
Maximum Alerts per Email	Enter the maximum number of alerts to include in a single email.

6. Click **Apply** to save your changes.

Basic Email Configuration

Select **Basic** under **Authentication** to use standard user name and password credentials for sending email through an SMTP server. Enter the SMTP server settings and credentials required by your email provider.

To configure Basic authentication:

1. Enter values for the following fields:

The screenshot shows the 'Preferences' dialog box with the 'Email Configuration' tab selected. The left sidebar lists various configuration categories, with 'Email Configuration' highlighted. The main area contains the following settings:

- Email Configuration**
 - ☒ Enable Email Notifications
 - Authentication: ☒ Basic, ☐ OAuth 2.0
 - *SMTP Host: [Text Field]
 - *SMTP Port: [Spin Box, value 25]
 - Encryption: ☐
 - *Encryption Type: [Dropdown Menu, value TLS]
 - *Username: [Text Field]
 - *Password: [Text Field]
 - *Sender Email: [Text Field]
 - Use Recommended Office 365 Settings: ☐
 - [Send Test Email Button]
- Batch Email Alerts for Quarantined Files**
 - Wait Before Sending (seconds): [Spin Box, value 60]
 - Maximum Alerts per Email: [Spin Box, value 1000]

At the bottom of the dialog are three buttons: 'Apply and Close', 'Cancel', and 'Apply'.

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 port used by the mail server's SMTP service. The default is 25 , or 465 when encryption is enabled.
Encryption	Select this checkbox if the SMTP server requires an encrypted connection.
Encryption Type	If encryption is enabled, select TLS or SSL . If you are unsure which is required, try one, and then the other.
Username	Enter the user name for the SMTP server account.
Password	Enter the password for the specified user.
Sender Email	Enter the email address to appear in the From field of sent emails. The address must correspond to a valid account on the SMTP server.
Use Recommended Office 365 Settings	Select this checkbox if connecting to an Office 365 SMTP server and you want PeerGFS to apply Microsoft's recommended Direct Send settings automatically. If you use another provider, leave this checkbox unchecked.

- (Recommended) Click **Send Test Email**, enter a destination address, and then click **OK** to send a test message and verify that the configuration works.
- (Optional) Complete the **Batch Email Alerts for Quarantined Files** section.

Use this section to define how alerts for quarantined files are batched and sent when email notifications are enabled. For descriptions of the fields, see [Email Configuration](#).

- Click **Apply** to save your changes.

OAuth2 Email Configuration

Select **OAuth 2.0** under **Authentication** to enable two-factor authentication for sending email notifications through supported providers such as Google Workspace (Gmail) or Microsoft 365. Enter the provider-specific information, update the firewall if required, authorize the account, and verify the configuration before saving.

Important: You must register PeerGFS as an OAuth client (create an OAuth client application in Google or Microsoft 365) **before** completing the **Email Configuration** section. When creating the OAuth client, select an application type that supports redirecting to `http://localhost` (for example, **Desktop/Native** application). This ensures that Peer Management Center can complete the authorization process.

To configure OAuth2 authentication:

1. Under **Provider**, select the email provider: **Microsoft 365** or **Google**.

If you selected **Microsoft 365**:

Preferences

type filter text

- Analytics
- Cloud Backup and Replic
- Collab, Sync, and Replic
- DFS-N Management
- Email Configuration**
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detectic
- SNMP Configuration
- Storage Configuration
- Task Scheduler
- User Management

Email Configuration

☒ Enable Email Notifications

Authentication: ☐ Basic ☒ OAuth 2.0

Email Configuration

*Provider: Microsoft 365

*SMTP Host: smtp.office365.com

*SMTP Port: 587

Authentication URL: https://login.microsoftonline.com/<Tenant ID>/oauth2/v2.0/authorize

Token URL: https://login.microsoftonline.com/<Tenant ID>/oauth2/v2.0/token

Redirect URI: http://localhost:15155/ Port: 15155

*Client ID:

*Tenant ID:

*Sender Email:

Batch Email Alerts for Quarantined Files

Wait Before Sending (seconds): 60

Maximum Alerts per Email: 1000

If you selected **Google**:

Preferences

type filter text

- Analytics
- Cloud Backup and Replic
- Collab, Sync, and Replic
- DFS-N Management
- Email Configuration**
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detectio
- SNMP Configuration
- Storage Configuration
- Task Scheduler
- User Management

Email Configuration

☒ Enable Email Notifications

Authentication: ☐ Basic ☒ OAuth 2.0

Email Configuration

*Provider:

*SMTP Host:

*SMTP Port:

Authentication URL:

Token URL:

Redirect URI: Port:

*Client ID:

*Client Secret:

*Sender Email:

Batch Email Alerts for Quarantined Files

Wait Before Sending (seconds):

Maximum Alerts per Email:

2. Enter values for the following fields. Some fields are pre-filled and normally do not need to be modified.


Note: When registering PeerGFS as an OAuth client, your provider generates values required in several of the fields below:

- **Microsoft 365: Client ID and Tenant ID**
- **Google Workspace: Client ID and Client Secret**

Field	Description
SMTP Host	Displays the host name of the SMTP server used for outgoing messages. This field is pre-filled based on the selected provider.
SMTP Port	Displays the TCP/IP port used by the SMTP service. This field is pre-filled.
Authentication URL	Displays the URL endpoint used for OAuth2 authentication. This field is pre-filled and should not be changed unless directed by Peer Support.
Token URL	Displays the URL endpoint where the authorization token is exchanged. This field is pre-filled.
Redirect URI	Displays the redirect address used during OAuth2 authentication. This field is pre-filled.
Port	Displays the local port used for redirect communication (default is 15155). This field is pre-filled.
Client ID	Enter the client ID provided when you registered Peer Management Center as an application in your provider's console.
Tenant ID/Client Secret	Enter the appropriate value: • For Microsoft 365, enter the Tenant ID. • For Google, enter the Client Secret generated for the registered application.
Sender Email	Enter the email address that should appear in the From field of sent messages. This must correspond to a valid account within your provider's domain.

3. **For Linux systems:** If Peer Management Center is running on a Linux server that you are configuring through a web browser on another computer, create a temporary SSH tunnel before authorizing the email account. See [OAuth2 Authentication on a Linux PMC](#) for instructions.
4. Click **Authorize Email Account**, and then click **OK** in the message that appears.

- Your default browser opens to the provider's login page so you can sign in and complete two-factor authentication.
- When authentication succeeds, a confirmation message appears in the browser, and a token is returned to the Peer Management Center service running on *localhost*:



5. Close the browser window.

A confirmation message is displayed in Peer Management Center.

6. Click **Send Test Email**, enter a destination address, and then click **OK**. Click **OK** in the confirmation dialog that appears.
7. If a connection error occurs during authorization, click **Update Firewall** to ensure that the local port (default is 15155) is available for authentication, and then retest.
8. Verify that the test email was received.

9. (Optional) Complete the **Batch Email Alerts for Quarantined Files** section.

Use this section to define how alerts for quarantined files are batched and sent when email notifications are enabled. For descriptions of the fields, see [Email Configuration](#).

10. Click **Apply** to save your changes.

OAuth2 Authentication on a Linux PMC

If Peer Management Center is running on a Linux server that you're configuring through a web browser on another computer, you must temporarily create an SSH tunnel before authorizing the email account. The tunnel allows your browser to communicate with the Peer Management Center service on the Linux host when the email provider redirects the authorization token back to *localhost*.

To create the tunnel:

1. On the client computer running the web browser, open a terminal.
2. Run the following command to create an SSH connection with local port forwarding: `ssh <user>@<hostname> -L 15155:localhost:15155`
3. Return to the browser and complete the OAuth2 Email Configuration steps (authorize the account, complete 2FA, and test the email).
4. After successful authentication, close the terminal window to end the SSH session.

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.
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: • 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** to save your changes.

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** to save your changes.

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** to save your changes.

Email Alerts

System email alerts notify recipients when certain types of system events occur, for example, low memory, low disk space, disconnected Agents. This [Preferences](#) page lists the existing system


email alerts. From this page, you can create, edit, and delete system email alerts. You can also disable and enable alerts. See [Email Alerts](#) in the [Basic Concepts](#) section for more information about email alerts.

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

To create a system email alert:


1. Select **Open Preferences** from the **Tools** menu.
2. Expand **General Configuration** in the navigation tree, and then select **Email Alerts**.

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



3. Click **Create**.

The **Create Email Alert** dialog appears.



The image shows a 'Create Email Alert' dialog box with a blue title bar. It contains several sections: a 'Name' text field, a 'Temporarily Disable' checkbox, an 'Event Types' section with four checked checkboxes ('Low Memory', 'Low PMC Disk Space', 'Agent Install Disk Space', 'Agent Disconnects') and one unchecked checkbox ('Licensing Warnings'), and a 'Recipients' section. The 'Recipients' section includes a text prompt 'Enter email, contact, or distribution list:', a search/filter text field with a magnifying glass icon and an 'Add to List' button, and a list box for recipients with 'Delete' and 'View Details' buttons. At the bottom are 'OK' and 'Cancel' buttons.

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** to save your changes.

The new alert is listed in the **Email Alerts** table.



Proxy Configuration

If your PMC lacks direct internet access, you have the option to set up a proxy, enabling log uploads and sending software updates.

To set up a proxy configuration:

1. Select **Open Preferences** from the **Tools** menu.
2. Expand **General** in the navigation tree, and then select **Proxy Configuration**.

The screenshot shows the 'Preferences' window with the 'Proxy Configuration' section selected in the left sidebar. The main area contains the following controls:

- type filter text** (input field)
- Proxy Configuration** (section header)
- ☒ **No proxy**
- ☐ **Manual proxy**
- IP Address:** (text input field)
- Port:** (text input field)
- ☐ **Use Authentication**
- Domain:** (text input field)
- Username:** (text input field)
- Password:** (password input field with an eye icon for toggling visibility)

At the bottom of the window are three buttons: **Apply and Close**, **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** to save your changes.


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 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** to save your changes.

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** to save your changes.

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:

Option	Description
	• 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 to enable HTTP access to the web service using the specified port.
Enable HTTPS REST API	Select to enable HTTPS access to the REST API service using the specified port.

4. Click **Apply** to save your changes.

Licensing

Overview

Peer Global File Service licensing is based on:

- The number of unique [participants](#)
- The number of terabytes configured in the [watch sets](#)

Depending on the license package installed, additional features and functionality may also be enabled.

Use the **Licensing** preferences page to:

- Install or update license files
- View installed license information

- Review license quantities and expiration dates
- Delete license files

Installed licenses are displayed in the **Peer Software Licensing Configuration** table. Select a license to view additional details, including:

- Package name
- License type
- Version number
- Expiration date
- Licensed features
- Capacity and job limits, if applicable


Installing or Upgrading a License File

After purchasing Peer Global File Service or requesting a trial, you receive a license file for your deployment.


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.



- Click the **Add/Update** button to browse for a license file.
- Select the license file, and then click **Open**.



If an existing license is detected, click **Yes** to replace it.



After the license is installed, the license appears in the licensing table with:

- License quantity
- Version
- Expiration date, if applicable

5. Select a license in the table to view additional details:

- The software version changed
- The package level changed
- The new license does not support the number of existing hosts

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

Restart Requirements

Restart existing jobs if any of the following conditions apply:

- The software version changed
- The package level changed
- The new license does not support the number of existing hosts

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.

After a license is removed, job types enabled only by that license are no longer available in 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** to save your changes.

Primary MED Options

The main options are as follows:

Option	Description
Enable Default Settings	Enables/disables Peer MED using default settings. By default, all three MED mechanisms are enabled.
Show/Hide Advanced Settings	Shows/hides options for each of the three MED mechanisms.
Enable Malicious Event Detection (MED)	The master on/off switch for MED. If unchecked, all MED mechanisms will be disabled.
Restore Default Settings	Restores all defaults across the three MED mechanisms.

Bait File Advanced Options

Bait files are files of common types, inserted into the file system in a way that hides them from users. Though hidden, these bait files are likely to be accessed by automated processes (like ransomware) or by mass deletions of entire folder structures. As soon as these files are touched, an action is triggered.

The options for bait files are:

Option	Description
Enable Bait Files	Enables/disables bait file creation and monitoring.
Add Bait Files to shares	At the start of each job, creates bait files under the root of each participant's watch directory. To see the watch directory for a job, review Host Participants and Directories .
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 Job	Triggers an alert, removes the afflicted Agent from the job where the malicious activity was detected, and stops the job. This option is the most aggressive and will require administrators to re-enable Agents as

Action	Description
	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:

PMC Alert Details

Received at: 07-29-2024 20:14:27

Severity: Fatal

Category: Alert Only

Host Name: Lab24Win22b

Locally Generated at: 07-29-2024 20:14:27

Name: FS-2

Message:

Action: Alert Only - Malicious Event Detection (MED) - Bait File (BF) Alert|Note: Alert Only: Please check for unwanted activity|Info: BAIT FILE ALERT|App ID: 113|App Session ID: 140|Participant Detected: Lab24Win22b|Time Detected: 30-07-2024 00:14:27.309|User Detected: LabAdmin|IP Detected: N/A|Process Detected: explorer.exe|Share Detected: N/A|Job Session ID: 1003|BF Path: C:\Bug Reports - Centreville\.pc-med_bin\.Doc_001-med.docx|BF Event Detected: Rename|PM Pattern: N/A|PM Short Version Found: N/A|PM Short Version Deleted: N/A|PM Delete Delay: N/A|PM Write Count: N/A|PM File List: N/A|PM Renamed From: N/A|PM Old Name: N/A|PM New Name: N/A|PM Folder: N/A|Count: N/A

Ref ID: 113

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.

PMC Backup

Preferences

type filter text

Analytics

Cloud Backup and Replication

Collab, Sync, and Replication

DFS-N Management

Email Configuration

General Configuration

Licensing

MED Configuration

PMC Backup

Real-time Event Detection

SNMP Configuration

Storage Configuration

Task Scheduler

User Management

PMC Backup

Backup location:
C:\Program Files\Peer Software\Peer Management Center\Hub\system backup

Scheduled backups limit: 30 0 = Keep all backups

Backup schedule: [Runs every 1 day](#)

Backup Date	Description	PMC Version	Created By
12-22-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
12-23-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
12-24-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
12-25-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
12-26-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
12-27-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
12-28-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
12-29-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
12-30-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
12-31-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-01-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-02-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-03-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-04-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-05-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-06-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-07-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-08-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-09-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-10-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-11-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-12-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler
01-13-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler

Backup Now

Edit

Delete

Export

Import

Restore

No restore scheduled.

Apply and Close

Cancel

Apply

Create PMC Backup □ ×

Timestamp:

PMC Version:

Description:

Create PMC Backup □ ×

Timestamp:

PMC Version:

Description:

PMC Backup ×

 Backup operation started in the background.

Preferences

type filter text

- Analytics
- Cloud Backup and Replication
- Collab, Sync, and Replication
- DFS-N Management
- Email Configuration
- General Configuration
 - Licensing
 - MED Configuration
 - PMC Backup**
 - Real-time Event Detection
 - SNMP Configuration
- Storage Configuration
- Task Scheduler
- User Management

PMC Backup

Backup location:

Scheduled backups limit: 0 = Keep all backups

Backup schedule: [Runs every 1 day](#)

Backup Date	Description	PMC Version	Created By	
12-30-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
12-31-2025 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-01-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-02-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-03-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-04-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-05-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-06-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-07-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-08-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-09-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-10-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-11-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-12-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-13-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-14-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-15-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-16-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-17-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-18-2026 01:00:00 A	Scheduled PMC Backup	6.3.0.20251124	Scheduler	
01-19-2026 01:00:04 A	Scheduled PMC Backup	6.4.0.20260108	Scheduler	
01-20-2026 01:00:00 A	Scheduled PMC Backup	6.4.0.20260108	Scheduler	
01-20-2026 03:45:27 P	Backup before 6.4 GA	6.4.0.20260108	User	

No restore scheduled.

Backup Now Edit Delete Export Import Restore

Apply and Close Cancel Apply

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.

The screenshot shows the 'Preferences' window with the 'Real-time Event Detection' tab selected. The left-hand navigation pane lists various configuration categories, with 'Real-time Event Detection' highlighted in blue. The main pane displays the following settings:

- Event Detection Options**
 - Max Path Length: 1024
 - Event Buffer Size: 1310
 - Access Polling Delay (Seconds): 1
 - Debug Mode: ☐
 - Advanced Job Configuration Options: ☒
 - Enable Real-time Analytics: ☒
 - Raw Event Logging: ☐
 - Advanced Configuration: (empty text box)

At the bottom of the window are three buttons: 'Apply and Close', 'Cancel', and 'Apply'.

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.

Option	Description
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.
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 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 Support.

- Click **Apply** to save your changes.

Storage Configuration

This section includes configuration options for integrating PeerGFS with your environment.

- [Storage Platforms](#) – Provides platform-specific instructions for configuring storage systems such as Amazon FSxN, Dell PowerScale, NetApp ONTAP, and Nutanix Files.
- [SMB Access Configuration](#) – Describes how to configure SMB access for a Linux Agent in multi-protocol job scenarios involving Windows-style permissions, mixed permissions, or Nutanix Files.


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.

A screenshot of a Windows-style dialog box titled "New SMB Access Configuration". It has a blue title bar with a close button. The dialog contains three text input fields: "Description:" (empty), "Domain Username:" (containing "Domain\Username"), and "Domain Password:" (empty with a password icon on the right). At the bottom are "OK" and "Cancel" buttons.

New SMB Access Configuration

Description:

Domain Username:

Domain Password: 

Field	Description
Description	Enter a description of the intended use of the configuration.
Domain Username	Enter the domain username in the format: Domain\Username
Domain Password	Enter the password for the domain account.

5. Click **OK**.
6. Click **Apply** to save your changes.

Storage Platforms

This section contains information about configuring your NAS for use with Peer Global File Service:

- [Amazon FSxN Configurations](#)
- [Dell Configurations](#)
- [NetApp ONTAP Configurations](#)
- [Nutanix Configurations](#)


Amazon FSxN 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 FSx for NetApp ONTAP configuration:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **Storage Configuration** in the navigation tree, and then select **Storage Platforms**.
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.

Management Agent

Select the server hosting the Peer Agent that manages this storage device.


Management Agent
Storage Information

Host	OS	Computer Description
DGUbuntu1	Linux	
DGUbuntu2	Linux	
DGWin16A	Microsoft Windows	
DGWin16B	Microsoft Windows	
DGWin16C	Microsoft Windows	
DGWin16D	Microsoft Windows	

< Back Next > Finish Cancel

5. Select a Management Agent, and then click **Next**.

The **Storage Information** page appears.



The image shows a 'Storage Information' dialog box with a blue header bar. The title bar contains standard window controls (minimize, maximize, close). The main title is 'Storage Information' and the subtitle is 'Enter the required information to connect to the storage'. On the left, there is a sidebar with two tabs: 'Management Agent' and 'Storage Information', with the latter being selected. The main area is divided into two sections. The top section, titled 'Credentials', contains five input fields: '*SVM Name:', '*SVM User Name:', '*SVM Password:', 'SVM Management IP:', and '*Peer Agent IP:'. The '*Peer Agent IP:' field has a dropdown arrow on its right. Below these fields is an 'Advanced' button. The bottom section contains a 'Validate' button followed by the text 'You must enter the name.' Below this is a line of text: 'Having trouble connecting? Please verify that all [prerequisites](#) are met for Amazon FSxN environments.' At the very bottom of the dialog are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Storage Information
Enter the required information to connect to the storage

Management Agent
Storage Information

Credentials

*SVM Name:

*SVM User Name:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

Advanced

Validate You must enter the name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Amazon FSxN environments.

< Back Next > Finish Cancel

6. Enter the required values in the **Credentials** section.

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 is not listed, enter the address manually.

7. (Optional) Click [Advanced](#) and enter the required values.
8. Click **Validate**.
9. Click **Next**.
10. Click **OK**.

Amazon FSxN Advanced Options

The following advanced configuration options are available for Amazon FSxN devices:

Amazon FSx for NetApp ONTAP Options

Advanced FPolicy Settings for host: DGWin16A and SVM: AMAZON SERVER 1

*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

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 Peer Support.

Dell Configurations


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, PowerStore, or Unity. For detailed information about Dell prerequisites, see [Dell Prerequisites](#).

To create a new Dell configuration:

1. Select **Open Preferences** from the **Tools** menu.
2. Select **Storage Configuration** in the navigation tree, and then select **Storage Platforms**.


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

*SmartConnect Zone Name: PowerStore
Unity

*Cluster Management IP:

*Cluster Username:

*Cluster Password:

Access Zones:

Connection Type: ☒ Syslog ☐ RabbitMQ

Syslog

*Agent IP Address:

*Listening Port: 6514

*SSL Certificate Path: Browse

*SSL Private Key Path: Browse

SSL Private Key Password:

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. For platform-specific field descriptions, see:

[Dell PowerScale Configuration](#)

[Dell PowerStore Configuration](#)

[Dell Unity Configuration](#)

8. (Optional) Click the **Advanced** button if you want to specify advanced options, and then enter the required values. For platform-specific field descriptions, see:

[Dell PowerScale Advanced Options](#)

[Dell PowerStore Advanced Options](#)

[Dell Unity Advanced Options](#)

9. Click **Validate**.

10. Click **Next**.

11. Click **OK**.

Dell PowerScale Configurations

To create a new Dell PowerScale configuration:

1. Enter the required values in the **Credentials** section; the required information varies by connection type (**RabbitMQ** or **Syslog**).

RabbitMQ:

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 two sections: "Credentials" and "Advanced". The "Credentials" section contains the following fields: "Device Type" (a dropdown menu set to "PowerScale"), "*SmartConnect Zone Name:" (a text box), "*Cluster Management IP:" (a text box), "*Cluster Username:" (a text box), "*Cluster Password:" (a text box), "Access Zone:" (a text box), and "Connection Type:" (radio buttons for "Syslog" and "RabbitMQ", with "RabbitMQ" selected). An "Advanced" button is located at the bottom right of the "Credentials" section. Below the "Credentials" section, there is a "Validate" button and a message: "You must enter a Cluster Name." At the bottom of the window, there is a message: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments." and four buttons: "< Back", "Next >", "Finish", and "Cancel".

Syslog:

These fields appear in both cases:

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 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.
Access Zone (RabbitMQ) or Access	Optional. Enter the name(s) of the access zone(s) being monitored.

Field	Description
Zones (Syslog)	
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: • Choose RabbitMQ if you're use a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent. • Choose Syslog if the PowerScale device sends notifications directly to the Agent.

2. If you selected **Syslog**, enter values for the following additional fields in the **Syslog** section:


Field	Description
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the PowerScale cluster. The cluster must be able to route traffic to this IP address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the port through 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 the 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 the 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 the Dell PowerScale Syslog Configuration Guide .

3. (Optional) Click [Advanced](#) and enter the required values.

4. Click **Validate**.
5. Click **Finish**.

The following advanced configuration options are available for Dell PowerScale devices; the required information varies by connection type (**RabbitMQ** or **Syslog**).

RabbitMQ:



The image shows a Windows-style dialog box titled "Dell PowerScale Options". It contains a section titled "Advanced Settings for Agent DGWin16A and Dell PowerScale: POWERSCALE1". This section includes several input fields: "*Cluster Management IP:", "Cluster Management Port:", "*Cluster Username:", "*Cluster Password:", "Access Zone:", "Connection Type:" (with radio buttons for "Syslog" and "RabbitMQ", where "RabbitMQ" is selected), "Filtered IP Addresses:" (with a dropdown arrow), "Nodes:", "Audit Cluster Name:", "Validate Cluster:" (with a checked checkbox), and "Update CEE Log Time:" (with a checked checkbox). Below these fields is a note: "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". At the bottom of the dialog are three buttons: "Validate", "OK", and "Cancel".

Dell PowerScale Options

Advanced Settings for Agent DGWin16A and Dell PowerScale: POWERSCALE1

*Cluster Management IP:

Cluster Management Port:

*Cluster Username:

*Cluster Password:

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 OK Cancel

Syslog:

Dell PowerScale Options

Advanced Settings for Agent DGWin16A and Dell PowerScale: POWERSCALE1

*Cluster Management IP:

Cluster Management Port:

*Cluster Username:

*Cluster Password:

Access Zones:

Connection Type: ☒ Syslog ☐ RabbitMQ

Syslog

*Agent IP Address:

*Listening Port: 6514

*SSL Certificate Path: Browse

*SSL Private Key Path: Browse

SSL Private Key Password:

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

ValidateOKCancel

Option	Description
Cluster Management IP	Enter the IP address to use to access the REST-based API integrated into the PowerScale 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.
Access Zone (Rabbit MQ) or Access Zones (Syslog)	Optional. Enter the name(s) of the access zone(s) being monitored.
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: • Choose RabbitMQ if you're use a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent. • Choose Syslog if the PowerScale device sends notifications directly to the Agent.

If you selected **RabbitMQ**, enter values for the following fields:

Option	Description
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 list of additional node IP addresses 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.
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**, enter values for the following required fields in the **Syslog** section:

Field	Description
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the PowerScale cluster. The cluster must be able to route traffic to this IP address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the port through 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 the Dell PowerScale Syslog Configuration Guide .

Field	Description
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 the 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 the Dell PowerScale Syslog Configuration Guide .

Dell PowerStore Configurations

To create a new Dell PowerStore configuration:

1. Enter the required values in the **Credentials** section.

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 two sections: "Credentials" and "CEE Configuration".

Credentials Section:

- Device Type: PowerStore (dropdown menu)
- *NAS Server Name: (text input field)
- *Management IP: (text input field)
- *Management Username: (text input field)
- *Management Password: (password input field)
- Connection Type: ☒ CEE

CEE Configuration Section:

- *Event Pool Name: (text input field)
- *Event Publisher Name: (text input field)
- *Agent IP Address: (dropdown menu)
- *Listening Port: 9843 (spin box)

At the bottom right of the CEE Configuration section is an "Advanced" button.

Below the configuration fields, there is a "Validate" button and a message: "You must enter a NAS Server Name." Below that, a note says: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerStore environments."

At the very bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Field	Description
NAS Server Name	Enter the name of the PowerStore device hosting the data to be replicated.
Management IP	Enter the IP address to use to access the REST-based API integrated into the PowerStore device.
Management Username	Enter the user name for the account managing the PowerStore device.
Management Password	Enter the password for account managing the PowerStore device.
Connection Type	Only option for this storage device. Notifications are sent via CEE to the Agent; no action is required.

2. Enter the required values in the **CEE Configuration** section:

Field	Description
Event Pool Name	Enter the exact name of the event pool on the PowerStore device that publishes CEE notifications. You can find this name in the PowerStore Manager UI under Event Notifications.
Event Publisher Name	Enter the exact name or ID of the event publisher on the PowerStore device that sends CEE notifications to the Agent. You can find this value in the PowerStore Manager UI under Event Notifications.
Agent IP	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the PowerStore device. The device must be

Field	Description
Address	able to route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

3. (Optional) Click [Advanced](#) and enter the required values.
4. Click **Validate**.
5. Click **Finish**.

The following advanced configuration options are available for Dell PowerStore devices:

Dell PowerStore Options

Advanced Settings for Agent DGWin16A and Dell PowerStore: POWERSTORE1

*Management IP:

Management Port:

*Management Username:

*Management Password:

CEE Configuration

*Event Pool Name:

*Event Publisher Name:

*Agent IP Address:

*Listening Port:

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

Option	Description
NAS Server Name	Enter the name of the PowerStore device hosting the data to be replicated.
Management IP	Enter the IP address to use to access the REST-based API integrated into the PowerStore device.
Management Username	Enter the user name for the account managing the PowerStore device.
Management Password	Enter the password for account managing the PowerStore device.

The **CEE Connection** fields are:

Field	Description
Event Pool Name	Enter the exact name of the event pool on the PowerStore device that publishes CEE notifications. You can find this name in the PowerStore Manager UI under Event Notifications.
Event Publisher Name	Enter the exact name or ID of the event publisher on the PowerStore device that sends CEE notifications to the Agent. You can find this value in the PowerStore Manager UI under Event Notifications.
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the PowerStore device. The device must be able to route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

Dell Unity Configurations

To create a new Dell Unity configuration:

1. Enter the required values in the **Credentials** section; the required information varies by connection type (**RabbitMQ** or **CEE**).

RabbitMQ:

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 two sections: "Credentials" and "Advanced". The "Credentials" section contains the following fields: "Device Type" (a dropdown menu showing "Unity"), "*NAS Server Name:" (a text input field), "*Unisphere Management IP:" (a text input field), "*Unisphere Username:" (a text input field), and "*Unisphere Password:" (a text input field). Below these fields is the "Connection Type:" section with two radio buttons: "RabbitMQ" (which is selected) and "CEE". An "Advanced" button is located at the bottom right of the "Credentials" section. Below the "Credentials" section, there is a "Validate" button and a message: "You must enter a NAS Server Name." At the bottom of the window, there is a footer with the text: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments." and four buttons: "< Back", "Next >", "Finish", and "Cancel".

CEE:

These fields appear in both cases:

Field Description	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 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.
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: - Choose RabbitMQ if you're use a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent. - Choose CEE if the Unity device sends notifications directly to the Agent.

- If you selected **CEE**, enter values for the following additional fields in the **CEE Configuration** section.


Field Description	Description
Event Pool Name	Enter the exact name of the event pool on the Unity device that publishes CEE notifications. You can find this name in the Unity Manager UI under Event Notifications.
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the Unity device. The device must be able to route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

- (Optional) Click [Advanced](#) and enter the required values.
- Click **Validate**.

5. Click **Finish**.

The following advanced configuration options are available for Dell Unity devices; the required information varies by connection type (**RabbitMQ** or **CEE**).

RabbitMQ:



The image shows a Windows-style dialog box titled "Dell Unity Options". It contains a section for "Advanced Settings for Agent DGWin16A and Dell Unity: UNITY1". This section includes several input fields: "*Unisphere Management IP:", "Unisphere Management Port:", "*Unisphere Username:", and "*Unisphere Password:". Below these is a "Connection Type:" section with two radio buttons: "RabbitMQ" (which is selected) and "CEE". There is also a "Filtered IP Addresses:" field with a dropdown arrow, and a "Validate Unisphere:" checkbox which is checked. At the bottom of the dialog, there is a "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". At the very bottom, there are three buttons: "Validate", "OK", and "Cancel".

Dell Unity Options

Advanced Settings for Agent DGWin16A and Dell Unity: UNITY1

*Unisphere Management IP:

Unisphere Management Port:

*Unisphere Username:

*Unisphere Password:

Connection Type: ☒ RabbitMQ ☐ CEE

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

CEE:

Dell Unity Options □ ×

Advanced Settings for Agent DGWin16A and Dell Unity: UNITY1

*Unisphere Management IP:

Unisphere Management Port:

*Unisphere Username:

*Unisphere Password:

Connection Type: ☐ RabbitMQ ☒ CEE

CEE Configuration

*Event Pool Name:

*Agent IP Address:

*Listening Port:

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

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 password for the Unisphere account managing the Unity storage device.
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: • Choose RabbitMQ if you're use a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent. • Choose CEE if the Unity device sends notifications directly to the Agent.

If you selected **RabbitMQ**, enter values for the following additional fields:

Option	Description
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.

If you selected **CEE**, enter values for the following additional fields in the **CEE Connection** section:

Field Description	Description
Event Pool Name	Enter the exact name of the event pool on the Unity device that publishes CEE notifications. You can find this name in the Unity Manager UI under Event Notifications.
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the Unity device. The device must be able to route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

NetApp ONTAP Configurations


Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more NetApp storage devices within most available job types. These NetApp devices can be running 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 **Storage Configuration** in the navigation tree, and then select **Storage Platforms**.
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.



Management Agent

Select the server hosting the Peer Agent that manages this storage device.


Management Agent
Storage Information

Host	OS	Computer Description
DGUbuntu1	Linux	
DGUbuntu2	Linux	
DGWin16A	Microsoft Windows	
DGWin16B	Microsoft Windows	
DGWin16C	Microsoft Windows	
DGWin16D	Microsoft Windows	

< Back Next > Finish Cancel

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

*SVM Name:

*SVM User Name:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

Advanced

Validate You must enter the name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for NetApp ONTAP environments

< Back Next > Finish Cancel


6. Enter the required values in the **Credentials** section.

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 is not listed, enter the address manually.

7. (Optional) Click [Advanced](#) and enter the required values.
8. Click **Validate**.
9. Click **Next**.
10. Click **OK**.

NetApp ONTAP Advanced Options

The following advanced configuration options are available for NetApp ONTAP devices:



The image shows a Windows-style dialog box titled "NetApp ONTAP Options". The title bar is blue with a close button (X) on the right. The main content area has a light gray background and contains the following elements:

- A header line: "Advanced FPolicy Settings for host: DGWin16A and SVM: NETAPP1" followed by a horizontal separator line.
- A list of configuration fields, each with a label and a corresponding input field:
 - *SVM Username: [text input]
 - *SVM Password: [password input]
 - SVM Management IP: [text input]
 - *Agent IP for SVM Conn.: [text input with a dropdown arrow on the right]
 - Filtered Extensions: [text input]
 - Admin Share Override: [text input]
 - Additional Properties: [text input]
- A note box with a light gray background and a thin border, containing the text: **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, there are three buttons: "Validate" on the left, and "OK" and "Cancel" on the right.

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 ONTAP 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 Peer Support.


Nutanix Files Configurations

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more Nutanix Files (formerly Acropolis File Services or AFS) clusters within most available job types. For detailed information about Nutanix prerequisites, see [Nutanix Prerequisites](#).

To create a new Nutanix Files configuration:


1. Select **Open Preferences** from the **Tools** menu.
2. Select **Storage Configuration** in the navigation tree, and then select **Storage Platforms**.
3. Select **Nutanix Files Configurations**.

The **Nutanix Files Configurations** page is displayed. It lists any existing configurations.



4. Click the **Create** button.

The **Management Agent** page appears.




Management Agent
Select the server hosting the Peer Agent that manages this storage device.

Management Agent	Host	OS	Computer Description
Storage Information	DGUbuntu1	Linux	
	DGUbuntu2	Linux	
	DGWin16A	Microsoft Windows	
	DGWin16D	Microsoft Windows	
	DGWin16B	Microsoft Windows	
	DGWin16C	Microsoft Windows	

< Back Next > Finish Cancel

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	Credentials
Storage Information	*Nutanix File Server Name:
	*Username:
	*Password:
	*Peer Agent IP: Advanced

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 the **Credentials** section.

Field	Description
Nutani	Enter the name of the Nutanix Files cluster hosting the data to be replicated.

Field	Description
x File Server Name	
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 is not listed, enter the address manually. 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 following advanced configuration options are available for Nutanix Files devices:

Nutanix Files Options

Advanced Settings for Agent DGWin16A and Nutanix File Server Name: NUTANIX1

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

Validate

OK

Cancel

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 is not listed, enter the address manually. 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.

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.

The screenshot shows a 'Preferences' dialog box with a navigation tree on the left and a configuration panel on the right. The navigation tree includes categories like Analytics, Cloud Backup and Replication, Collab, Sync, and Replication, DFS-N Management, Email Configuration, General Configuration, Licensing, MED Configuration, PMC Backup, Real-time Event Detection, **SNMP Configuration** (highlighted), Storage Configuration, Task Scheduler, and User Management. The right panel, titled 'SNMP Configuration', contains three input fields: 'Source IP Address' (a dropdown menu), 'Destination' (a text field containing '255.255.255.255'), and 'Trap Prefix' (a text field containing '1.3.6.1.4.1.58279.2'). Below these fields is a 'Test SNMP Settings' button. At the bottom of the dialog are three buttons: 'Apply and Close', 'Cancel', and 'Apply'.

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** to save your changes.

Telemetry


Telemetry settings determine whether PeerGFS sends anonymous diagnostic data to Peer Software and whether Proactive Monitoring is enabled.

For conceptual information about telemetry, anonymous diagnostic data, and Proactive Monitoring, see [Telemetry](#) in Advanced Topics.

Note: Proactive Monitoring is available only to customers with dedicated Technical Account Managers (TAMs).

To configure Telemetry:

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



3. Modify the settings as needed.

Option	Description
Share anonymous diagnostic data with Peer Software	Select this checkbox to send anonymized diagnostic data to Peer Software. Click Learn more for additional information about anonymous diagnostic data.
Enable Proactive Monitoring	Click this button to enable Proactive Monitoring, and then complete the Proactive Monitoring Setup wizard. Click What is Proactive Monitoring? for additional information about Proactive Monitoring.

4. Click **Apply** to save your changes.

Proactive Monitoring Setup


Proactive Monitoring setup includes several configuration pages.

Complete each page in sequence, starting with with [Step 1: General Configuration](#).

Step 1: General Configuration

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 blue header bar. Below the header, the "General Configuration" tab is selected, showing a list of configuration options on the left: "General Configuration", "Telemetry Options", "Agent Locations", "Health Checker Setup", and "Confirmation". The main area contains three fields: "Subscription ID" (empty), "Environment Name" (filled with "DGWin16A"), and "Upload Interval" (a spinner set to "30" minutes). A link "What is my Subscription ID?" is positioned between the Subscription ID and Environment Name fields. At the bottom, there 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 [Step 2: Telemetry Options](#).

Step 2: Telemetry Options

The **Telemetry Options** page allows you to select detailed telemetry data to upload to be used by the Peer Software Technical Account Management (TAM) 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:

The screenshot shows a window titled "Set Up Proactive Monitoring" with a close button in the top right corner. The window has a left sidebar with a tree view containing "General Configuration", "Telemetry Options" (which is selected and highlighted), "Agent Locations", "Health Checker Setup", and "Confirmation". The main area of the window is titled "Telemetry Options" and contains the instruction "Review and select the types of information to upload for Proactive Monitoring." Below this is a link: "What types of data are uploaded to Proactive Monitoring?". The main area is divided into three sections, each with a heading and a list of checkboxes, all of which are checked:

- 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

At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "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.	• IP Information - The IP address of the server that this PMC is installed on. • Statistical information - In-depth statistics about the queues and performance of PeerGFS's replication engine.
Agent Details	Includes the details about the Agents that are connected to this PMC. Details include service and server memory consumption, replication throughput, uptime, operating system, and disconnect counts.	• IP information - The IP addresses of the servers that the Agents are installed on. • Agent Names - The names assigned to the Agents (typically the name of each Agent's Windows Server). If this option is not checked, random strings will be used in the 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—you must manually enter it. • 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.	• 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. • MED Alerts - If MED alerts are enabled in this PMC, this option will include any alerts for system, the Proactive Monitoring system.
	Per-job details include replication backlog, replication	

2. Click **Next**.

If you selected the **Include Agent Locations** option, continue with [Step 3: Agent Locations](#); otherwise continue with [Step 4: Health Checker Setup](#).

Step 3: Agent Locations

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.

Set Up Proactive Monitoring

Agent Locations
(Optional) Set location details for each Agent for use with dashboard maps.

General Configuration
Telemetry Options
Agent Locations
Health Checker Setup
Confirmation

Only running Agents are listed below.

Agent	Computer Description	Location*	Longitude	Latitude
DGUbuntu1				
DGUbuntu2				
DGWin16A				
DGWin16B				
DGWin16C				
DGWin16D				

*Location identifies the city, state, and country of the Agent (for example: Centreville, VA, US)

< Back Next > Finish Cancel

2. Click **Next** to proceed to [Step 4: Health Checker Setup](#).

Step 4: Health Checker Setup

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

The **Health Checker Setup** page allows to you decide whether you want to set up the Health Checker, a standalone monitoring and reporting tool. The Health Checker is used by the Peer Software Technical Account Management (TAM) team to monitor for replication issues and track the **Real-Time Delta (RTD)**, which indicates how many minutes the replication is out of sync between participants. The RTD is similar to a **Recovery Point Objective (RPO)** but focuses on real-time replication engine rather than the scan engine.

You can install the Health Checker **locally** (on the same server as the Peer Management Center) or **externally** (on a separate server that neither Peer Management Center nor Agents 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.

Set Up Proactive Monitoring

Health Checker Setup

Set up the standalone Health Checker to alert on replication issues and track RTD.

General Configuration
Telemetry Options
Agent Locations
Health Checker Setup
Set up the Health Checker
Select Jobs
Schedule Task
Update Health Checker
Confirmation

[What is the Health Checker?](#)

☐ Configure the Health Checker already installed on a remote server.
[How do I install the Health Checker on a remote server?](#)

☒ Set up Health Checker on this PMC server.

☐ Do not install or configure the Health Checker.

< Back Next > Finish Cancel

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 with 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.



3. 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 [Schedule Task](#) page appears.

Schedule Task

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

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

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

Set Up Proactive Monitoring

Schedule Task
Specify the conditions that will trigger the task.

General Configuration
Telemetry Options
Agent Locations
Health Checker Setup
Set up the Health Checker
Select Jobs
Schedule Task
Update Health Checker
Confirmation

Settings

☒ One Time Start: 06 / 23 / 2026 15 : 51 : 08
☐ Daily
☐ Weekly
☐ Monthly

Advanced Settings

☒ Repeat task every: 4 Hours
☐ Expire: 06 / 23 / 2027 15 : 51 : 08

< Back Next > Finish Cancel

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 [Update 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.

Step 5: Confirmation

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

1. Review the configuration.

The screenshot shows a window titled "Set Up Proactive Monitoring" with a "Confirmation" tab selected. The window contains a summary of the configuration settings. On the left is a sidebar with a tree view of the setup steps: General Configuration, Telemetry Options, Agent Locations, Health Checker Setup (expanded), and Confirmation (selected). The main area displays the following settings:

- General Configuration**
 - Subscription ID: [Redacted]
 - Environment Name: DGWin16A
 - Upload Interval: 30 minutes
- Telemetry Options**
 - PMC Details
 - Include PMC IP: Yes
 - Include Statistical Information: Yes
- Agent Details**
 - Include Agent IP Info: Yes
 - Include Agent Names: Yes
 - Include Agent Locations: Yes
 - Include Agent Storage Information: Yes
- Job Details**
 - Include Job Names: Yes
 - Include MED Alerts: Yes

At the bottom, there is a note: "Note: After clicking Finish, the settings specified in this wizard will be enabled immediately. Information will be uploaded to the Proactive Monitoring dashboard within 30 minutes. You can change these settings or disable the upload of this information by selecting Proactive Monitoring from the Help menu, and then clicking the Disable button in the Proactive Monitoring preference page." Below the note are four buttons: "< Back", "Next >", "Finish", and "Cancel".

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

The settings are displayed in the Telemetry preferences page.

Preferences

type filter text

- Analytics
- Cloud Backup and Replication
- Collab, Sync, and Replication
- DFS-N Management
- Email Configuration
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- Storage Configuration
- Task Scheduler
- Telemetry**
- User Management

Telemetry

Anonymous Diagnostic

☐ Share anonymous diagnostic data with Peer Software. [Learn more](#)

Proactive Monitoring

[What is Proactive Monitoring?](#)

General Configuration

Subscription ID:

[What is my Subscription ID?](#)

Environment Name:

Upload Interval: minutes

Telemetry Options

[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

[Disable Proactive Monitoring](#)

Apply and Close Cancel Apply

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**.

- Click **Apply** to save your changes.
- Notify your Peer Software Technical Account Manager that the setup of Proactive Monitoring is complete.

Disabling and Re-enabling Proactive Monitoring

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**.

The screenshot shows the 'Preferences' window with the 'Telemetry' tab selected. The left sidebar lists various configuration categories, with 'Telemetry' highlighted. The main panel contains the following sections:

- Anonymous Diagnostic**: A checkbox labeled 'Share anonymous diagnostic data with Peer Software.' with a link to 'Learn more'.
- Proactive Monitoring**: A section with a link 'What is Proactive Monitoring?'.
- General Configuration**:
 - Subscription ID**: A text field with a link 'What is my Subscription ID?'.
 - Environment Name**: A text field containing 'DGWin16A'.
 - Upload Interval**: A spinner box set to '30' with the unit 'minutes'.
- Telemetry Options**:
 - 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


At the bottom of the Telemetry section, the button 'Disable Proactive Monitoring' is highlighted with a red rectangle. At the bottom of the window are three buttons: 'Apply and Close', 'Cancel', and 'Apply'.

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** to save your changes.

Re-enabling Proactive Monitoring

To re-enable Proactive Monitoring:

1. Rerun the **Set up Proactive Monitoring** wizard.
2. Click the **Enable Proactive Monitoring** button

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 Replication
- Collab, Sync, and Replication
- DFS-N Management
- Email Configuration
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- Storage 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](#)

Managing Internal Users

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

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' tab selected. The left sidebar contains a list of configuration categories, with 'User Management' highlighted. The main content area is divided into several sections:

- Roles:** A table with three roles: 'Power User', 'Administrator', and 'Help Desk'. Each role has 'Create', 'Edit', and 'Delete' buttons.
- Users:** A section containing 'Internal Users' and 'Active Directory Users and Groups'.
 - Internal Users:** A table with one user: 'admin'. It has 'Create', 'Edit', and 'Delete' buttons.
 - Active Directory Users and Groups:** Two sub-sections: 'Active Directory Users' and 'Active Directory Groups'. Each has an 'Add', 'Edit', and 'Delete' button.
- Active Directory Authentication:** A section with a warning: 'Authentication will not work until the URL and credentials are entered.' It includes fields for 'LDAP Server URLs', 'LDAP Search Domain (Optional)', and 'LDAP Credentials'. There are 'Add/Update LDAP Credentials' and 'Test' buttons.

At the bottom of the window are three buttons: 'Apply and Close', 'Cancel', and 'Apply'.

3. Select the **Create** button for Internal Users.

The **New Internal User** dialog appears.

A screenshot of a 'New Internal User' dialog box. The dialog has a blue title bar with the text 'New Internal User' and a close button (X). Inside, there are several fields: 'Username:' with a text input field and an example 'JohnM' below it; 'Password:' with a text input field; 'Password Confirm:' with a text input field; 'Enabled:' with a checked checkbox; 'Role:' with a dropdown menu showing 'Power User'; and 'Contact:' with a dropdown menu showing 'None'. At the bottom right are 'OK' and 'Cancel' buttons.

New Internal User

Username:
Example: 'JohnM'

Password:

Password Confirm:

Enabled: ☒

Role:

Contact:

OK Cancel

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** to save your changes.

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. Select **Open Preferences** from the **Tools** menu.
2. Select **User Management** from the navigation tree.

The screenshot shows the 'Preferences' dialog box with the 'User Management' tab selected. On the left is a sidebar with a search bar 'type filter text' and a list of configuration categories: Analytics, Cloud Backup and Replication, Collab, Sync, and Replication, DFS-N Management, Email Configuration, General Configuration, Licensing, MED Configuration, PMC Backup, Real-time Event Detection, SNMP Configuration, Storage Configuration, Task Scheduler, and User Management (which is highlighted). The main area is titled 'User Management' and contains several sections: 'Roles' with a list (Power User, Administrator, Help Desk) and buttons (Create, Edit, Delete); 'Users' with a sub-section 'Internal Users' containing a list (admin, Intern) and buttons (Create, Edit, Delete); 'Active Directory Users and Groups' with sub-sections for 'Active Directory Users' and 'Active Directory Groups', each with an 'Add' button and 'Edit'/'Delete' buttons; and 'Active Directory Authentication' with a warning 'Authentication will not work until the URL and credentials are entered.', fields for 'LDAP Server URLs:', 'LDAP Search Domain (Optional):' (with 'Optional' entered), and 'LDAP Credentials:', and buttons for 'Add/Update LDAP Credentials' and 'Test'. At the bottom are 'Apply and Close', 'Cancel', and 'Apply' buttons.

type filter text

- Analytics
- Cloud Backup and Replication
- Collab, Sync, and Replication
- DFS-N Management
- Email Configuration
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- Storage Configuration
- Task Scheduler
- User Management

User Management

Roles

Power User	Create
Administrator	Edit
Help Desk	Delete

Users

Internal Users

admin	Create
Intern	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:

3. Select the user from the list of internal users.
4. Click **Edit**.
5. Make the changes in the **Edit User Information** dialog.

A screenshot of a Windows-style dialog box titled "Edit User Information" with a close button (X) in the top right corner. The dialog contains several fields: "Username:" with a text box containing "Intern" and a hint "Example: 'JohnM'"; "Password:" and "Password Confirm:" with masked password fields; "Enabled:" with a checked checkbox; "Role:" with a dropdown menu showing "Limited User"; and "Contact:" with a dropdown menu showing "None". At the bottom right are "OK" and "Cancel" buttons.

Edit User Information X

Username: Intern
Example: "JohnM"

Password:

Password Confirm:

Enabled: ☒

Role: Limited User ▼

Contact: None ▼

OK Cancel

6. Click **OK**.
7. Click **Apply** to save your changes.

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. Select **Open Preferences** from the **Tools** menu.
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** to save your changes.

Managing Active Directory Users and Groups

An Active Directory (AD) user or group is one whose access to Peer Management Center is authenticated through Active Directory. Adding an Active AD user or group authenticates and authorizes that user or group members to use Peer Management Center. The AD user or group must already exist in Active Directory prior to adding the user or group to Peer Management Center. Active Directory users won't be able to access the web client until [Active Directory authentication is configured](#) in Peer Management Center.

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.

Preferences

type filter text

- Analytics
- Cloud Backup and Replication
- Collab, Sync, and Replication
- DFS-N Management
- Email Configuration
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- Storage 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:

- 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**.

A screenshot of a dialog box titled "LDAP Credentials" with a close button (X) in the top right corner. The dialog contains four input fields: "Domain:", "Username:", "Password:", and "Password Confirm:". Below the "Domain:" field is an example text: "Example: 'MYDOMAIN.local' or 'MYDOMAIN.com'". Below the "Username:" field is an example text: "Example: 'JohnM'". At the bottom of the dialog are two buttons: "OK" and "Cancel".

LDAP Credentials

Domain:
Example: 'MYDOMAIN.local' or 'MYDOMAIN.com'

Username:
Example: 'JohnM'

Password:

Password Confirm:

OK 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.

Preferences

type filter text

- Analytics
- Cloud Backup and Replication
- Collab, Sync, and Replication
- DFS-N Management
- Email Configuration
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- Storage 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:

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.

A screenshot of a Windows-style dialog box titled "Add a new Active Directory User". It has a blue header bar with a close button (X) on the right. The dialog 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 showing "Power User". At the bottom are "OK" and "Cancel" buttons.


Domain:
Example: 'MYDOMAIN.local' or 'MYDOMAIN.com'

Username:
Example: 'JohnM'

Role:

OK Cancel

- For a user group, enter the domain name, group name, and select a role.

A screenshot of a Windows-style dialog box titled "Add a new Active Directory Group". It has a blue header bar with a close button (X) on the right. The dialog 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 showing "Power User". At the bottom are "OK" and "Cancel" buttons.

Domain:
Example: 'MYDOMAIN.local' or 'MYDOMAIN.com'

Group:
Example: 'Domain Admins' or 'Power Users'

Role:

OK Cancel

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](#)

Creating a Custom Role

To create a custom role:

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

The screenshot shows the 'Preferences' dialog box with the 'User Management' tab selected. On the left is a sidebar with a search bar 'type filter text' and a list of configuration categories: Analytics, Cloud Backup and Replication, Collab, Sync, and Replication, DFS-N Management, Email Configuration, General Configuration, Licensing, MED Configuration, PMC Backup, Real-time Event Detection, SNMP Configuration, Storage Configuration, Task Scheduler, and 'User Management' (which is highlighted). The main area is titled 'User Management' and contains several sections: 'Roles' with a table listing 'Power User', 'Administrator', and 'Help Desk', each with 'Create', 'Edit', and 'Delete' buttons; 'Users' with an 'Internal Users' section containing an 'admin' entry with similar buttons; 'Active Directory Users and Groups' with 'Active Directory Users' and 'Active Directory Groups' sections, each with an 'Add', 'Edit', and 'Delete' button; and 'Active Directory Authentication' with a warning 'Authentication will not work until the URL and credentials are entered.', fields for 'LDAP Server URLs', 'LDAP Search Domain (Optional):' (containing 'Optional'), and 'LDAP Credentials', and buttons for 'Add/Update LDAP Credentials' and 'Test'. At the bottom right are 'Apply and Close', 'Cancel', and 'Apply' buttons.

type filter text

- Analytics
- Cloud Backup and Replication
- Collab, Sync, and Replication
- DFS-N Management
- Email Configuration
- General Configuration
- Licensing
- MED Configuration
- PMC Backup
- Real-time Event Detection
- SNMP Configuration
- Storage 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:

- 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 blue title bar and standard window controls. On the left is a sidebar with three tabs: 'General' (selected), 'Permissions', and 'Tags'. The main area is titled 'General' and contains four fields: 'Role Name' with an example of 'projectmanager' or 'helpdesk'; 'Display Name' with an example of 'Administrator' or 'Help Desk'; 'Description' with an example of 'Default administrator role.'; and 'Base Role' which is a dropdown menu currently set to 'Administrator'. At the bottom right are 'OK' and 'Cancel' buttons.

4. Enter the following information for the new role:

- **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:** 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) 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**.

New Role

General
Permissions
Tags

Permissions

Category	Name	Access
Cloud Backup and Replication UI	Job View - Alerts Tab	Full access
Cloud Backup and Replication UI	Job View - Configuration Tab	Full access
Cloud Backup and Replication UI	Job View - Event Log Tab	Full access
Cloud Backup and Replication UI	Job View - Failed Events Tab	Full access
Cloud Backup and Replication UI	Job View - Participants Tab	Full access
Cloud Backup and Replication UI	Job View - Summary Tab	Full access
Cloud Backup and Replication UI	Runtime Summary View	Full access
Collab/Sync/Repl UI	Job View - Alerts Tab	Full access
Collab/Sync/Repl UI	Job View - Configuration Tab	Full access
Collab/Sync/Repl UI	Job View - Event Log Tab	Full access
Collab/Sync/Repl UI	Job View - Participants Tab	Full access
Collab/Sync/Repl UI	Job View - Quarantines Tab	Full access
Collab/Sync/Repl UI	Job View - Retries Tab	Full access
Collab/Sync/Repl UI	Job View - Session Tab	Full access
Collab/Sync/Repl UI	Job View - Summary Tab	Full access
Collab/Sync/Repl UI	Runtime Summary View	Full access
DFS-N UI	Job View - Alerts Tab	Full access
DFS-N UI	Job View - Configuration Tab	Full access
DFS-N UI	Job View - Namespace Servers Tab	Full access
DFS-N UI	Job View - Namespace Tab	Full access
DFS-N UI	Runtime Summary View	Full access
PMC UI	Assign Tags	Full access
PMC UI	Download Broker Installer	Full access
PMC UI	Download Peer Agent Installer	Full access
PMC UI	Download PeerIQ	Full access
PMC UI	Generate PMC Memory Dump File	Full access
PMC UI	Generate PMC and Broker Thread Dump	Full access
PMC UI	Jobs View	Full access
PMC UI	New Job Wizard	Full access
PMC UI	Preferences	Full access
PMC UI	Retrieve Broker Statistics	Full access
PMC UI	Retrieve PMC/Agent Logs	Full access
PMC UI	Show Agent Summary	Full access
PMC UI	Show Progress View	Full access
PMC UI	Task Scheduler and Task History	Full access
PMC UI	Topology	Full access
PMC UI	View Alerts	Full access

OK Cancel

- 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.
- (Optional) Click the **Tags** tab to configure job visibility for the role using tags.

- a. On the **Tags** tab, configure the following options:
 - o **Show all jobs (non-matched jobs will have read-only access):** Select this to allow users assigned to this role to see all jobs. Jobs that do not match the selected tags are displayed with read-only access.
 - o **Filter tags below by job type:** Select this to filter the list of available tags by job type and narrow the tag selection. This option affects only which tags are displayed and does not affect job visibility.
- b. Select the job tags that determine which jobs are visible to users assigned to this role:
 - o In the **Unassigned tags** list, select one or more tags, and then click the right-arrow (**Add One**) or double right-arrow (**Add All**) to move them to the **Assigned** tags list.
 - o Use **Shift** to select multiple tags.

Only jobs that have the selected tags assigned will appear in the user's **Jobs** view. Before configuring role-based job visibility, ensure that the appropriate tags have been assigned to jobs. For information about tagging jobs, see [Assigning Tags](#).

New Role

General
Permissions
Tags

Tags

Assigning tags to a role allows you to restrict user access to only specific jobs matching the selected tags.

☐ Show all jobs (non matched jobs will have read-only access)

☒ Filter tags below by job type: All tags

Unassigned tags

Tag Name Filter Category Filter

Tag Name	Category
----------	----------

Assigned tags

Tag Name Filter Category Filter

Tag Name	Category
----------	----------

>>> > < <<<

OK Cancel

9. Click **OK**.

The new role appears in the **Roles** section.

Editing a Custom Role

You can edit a custom role, changing its Role Name and Display Name, its base role, associated permissions, and tags assigned to the role.

Editing of a base role is much more restricted. It is limited to modifying the tags assigned to the role. You cannot edit its names or associated permissions.

To edit a role, select the role in the **Roles** section in the **User Management** page, and then click **Edit**.

Deleting a Custom Role

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.

Assigning Tags to a Role

Roles do not have tags assigned to them. Instead, roles reference tags that are assigned to jobs to control which jobs are visible to users. For information about tagging jobs,, 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** to proceed to the [Source Storage Platform](#) page.

Step 2: Source Storage Platform

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

1. Select the storage platform type hosting the data you want to replicate.




2. Click **Next** to proceed to the [Management Agent](#) page.

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**. The **Management Agent** page the available [Agents](#) in your environment. You can assign multiple Agents to the same storage device. However, each Agent must be used to access different volumes, shares, exports, or folders on that device.

1. From the **Agents** table, select an Agent that can access the storage device you want to replicate.

Tip: If the Agent you want is not shown in the **Agents** table, 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](#), the **Agents** table will be updated.



Create Cloud Backup and Replication Job Wizard

Management Agent

Select the server hosting the Peer Agent that manages this source storage device.

Source Storage Platform

- Management Agent
- Proxy Configuration
- Storage Information

Source Paths

- File and Folder Filters

Destination


- Azure Storage Credentials
- Container Details
- Replication and Retention Policies
- Miscellaneous Options
- Email Alerts
- SNMP Notifications
- Confirmation

Only Microsoft Windows Agents are available for Cloud Backup and Replication job

Agent	OS	Computer Description	Domain	Database Connection
Lab24Win19b	Windows Server 2019		PEERTEST	N/A
Lab24Win19c	Windows Server 2019		PEERTEST	N/A
Lab24Win22a	Windows Server 2022		PEERTEST	DB-2
Lab24Win22b	Windows Server 2022		PEERTEST	DB-1
Lab24Win22c	Windows Server 2022		PEERTEST	N/A
Lab24Win22d	Windows Server 2022		PEERTEST	N/A

< Back Next > Finish Cancel

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** to proceed to the [Proxy Configuration](#) page.

Step 4: Proxy Configuration

If you do not need a proxy server to connect to outside networks, skip this step and proceed to [Step 5](#).

If you do need a proxy server to connect to outside networks, you have three options:

- **Create a new proxy configuration.**
- **Use the existing proxy configuration.** If there is an existing proxy configuration, details about the configuration will be displayed on the page.


You may have created one in advance through [Cloud Back and Replication Preferences](#) or when you created another Cloud Backup and Replication job. Once a proxy configuration is created for a source storage platform, that proxy configuration is used for all Cloud Backup and Replication jobs using that Agent.

- **Edit an existing proxy configuration.** Click the **Edit Proxy Configuration** link to edit the existing proxy.


If you edit the proxy configuration, it affects other jobs using the same Agent. Editing an existing proxy configuration has the potential to create problems with the other jobs.

If there is not an existing proxy configuration for the selected Management Agent, follow these steps to create a new proxy configuration:

1. Click **Create Proxy Configuration**.




The **Proxy Configuration** page is displayed. Existing proxies are listed in the **Proxy Configuration** table.



2. Click the **Create** button.

The **Create Proxy Configuration** dialog is displayed. The Agent you selected in **Step 3: Management Agent** is preselected.



3. Enter values for the following fields:


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 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.



6. (Optional) Select the proxy you created and click **Validate**.

The **Validate Proxy Configuration** dialog appears.

7. Select the target storage account, and then click **OK**.



Validate Proxy Configuration

Proxy Configuration

Agent: Lab24Win22a

IP Address: 192.168.169.49


Port: 8088

Storage Account:

Close

Once you click OK, PeerGFS tests the connection to the target storage account using the proxy.

8. Click **OK** in the **Validation Result** dialog.
9. Click **Close** in the **Validate Proxy Configuration** dialog.



Validate Proxy Configuration

Proxy Configuration

Agent: Lab24Win22a


IP Address: 192.168.169.49

Port: 8088


Storage Account: NetApp [NetApp StorageGRID]

Close

10. Click **OK** in the **Proxy Configuration** page.



The **Proxy Configuration** page now displays the details about the proxy configuration.



11. Click **Next** to proceed to the [Storage Information](#) page.

Step 5: Storage Information

On the **Storage Information** page, enter the required details for the storage platform selected in [Step 2: Source Storage Platform](#). The fields displayed vary based on the selected platform type.

Follow these steps:

1. If you selected **Windows File Server** as the source storage platform, skip to [Step 6: Source Paths](#). If you selected any other storage platform type, continue with the next step.
2. Select **New Credentials** or **Existing Credentials**.
3. If you selected **Existing Credentials**, choose a credential from the drop-down list, and then click **Next** to proceed to the [Source Paths](#) page.

4. If you selected **New Credentials**, enter the required connection details. The fields vary by platform. For platform-specific field descriptions, see:

[Amazon FSx for NetApp ONTAP](#)

[Dell PowerScale](#)

[Dell Unity](#)

[NetApp ONTAP](#)

[Nutanix Files](#)

5. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
6. Click **Next** to proceed to the [Source Paths](#) page.

Amazon FSx for NetApp ONTAP

To connect to the Amazon FSx for NetApp ONTAP Storage Virtual Machine (SVM) hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

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 Policies
 Miscellaneous Options
 Email Alerts
 SNMP Notifications
 Confirmation

☒ 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

Option	Description
SVM Name	Enter the name, FQDN (Fully Qualified Domain Name), or IP address of the SVM hosting the data to be replicated.
SVM Username	Enter the user name for the account managing the SVM. This must not be a cluster management account.
SVM Password	Enter the password for the SVM management account. 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 SVM Name already allow 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

Option	Description
	desired IP address is not listed, enter the address manually.
Access Path	Use only when experiencing access issues. Contact Peer Support for more information.

- Click **Advanced** if you want to set [advanced options](#).
- Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

- Click **Next** to proceed to the [Source Paths](#) page.

Dell PowerScale

To connect to the Dell PowerScale cluster hosting the data to be replicated:

- If you selected **New Credentials**, enter the required connection details in the **Credentials** section.

The available connection types depend on the operating system of the Agent, as well as the notification delivery method used by the storage platform.

- o **RabbitMQ**: Available only when the Agent is running on Windows. Choose RabbitMQ if you are using a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent.
- o **Syslog**: Available on both Windows and Linux Agents. Choose Syslog if the PowerScale device sends notifications directly to the Agent.

If you selected **Rabbit MQ**, this page appears:

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 Policies
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

Credentials
☒ New Credentials
*SmartConnect Zone Name:
*Cluster Management IP:
*Cluster Username:
*Cluster Password:
Access Zone:
Connection Type: ☐ Syslog ☒ RabbitMQ

☐ 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

If you selected **Syslog**, this page appears:

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 Policies
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

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.

< Back Next > Finish Cancel

These fields appear in both cases:

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.
Cluster Username	Enter the user name for the account managing the PowerScale cluster.
Cluster Password	Enter the password for account managing the PowerScale cluster.

Field	Description
Cluster Access Zone	Optional. Enter 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: - Choose RabbitMQ if you're use a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent. - Choose Syslog if the PowerScale device sends notifications directly to the Agent.

2. If you selected **Syslog**, enter values for the following additional fields in the **Syslog** section:

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 is not listed, enter the address manually.
Listening Port	Enter the port through 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 the 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 the 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 the Dell PowerScale Syslog Configuration Guide .

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** to proceed to the [Source Paths](#) page.

Dell Unity

To connect to the Dell Unity NAS Server hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details in the **Credentials** section.

The available connection types depend on the operating system of the Agent, as well as the notification delivery method used by the storage platform.

- o **RabbitMQ**: Available only when the Agent is running on Windows. Choose RabbitMQ if you are using a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent.
- o **Syslog**: Available on both Windows and Linux Agents. Choose Syslog if the PowerScale device sends notifications directly to the Agent.

If you selected **Rabbit MQ**, this page appears:

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 Policies
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

Credentials

☒ New Credentials

*NAS Server Name:

*Unisphere Management IP:

*Unisphere Username:

*Unisphere Password:

Connection Type: ☒ RabbitMQ ☐ CEE

☐ Existing Credentials

MLUNITY, user:admin

Access Path
Access Path:

You must enter a NAS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments.

If you selected **CEE**, this page appears:

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 Policies
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

Credentials

☒ New Credentials

*NAS Server Name:

*Unisphere Management IP:

*Unisphere Username:

*Unisphere Password:

Connection Type: ☐ RabbitMQ ☒ CEE

CEE Configuration

*Event Pool Name:

*Agent IP Address:

*Listening Port:

☐ Existing Credentials

MLUNITY, user:admin

Access Path
Access Path:

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments.

These fields appear in both cases:

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 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.
Unisphere Password	Enter the password for the Unisphere account managing the Unity storage device.
Access Path	Use only when experiencing access issues. Contact Peer Support for more information.

3. If you selected **CEE**, enter values for the following additional fields in the **CEE Configuration** section.

Field Description	Description
Event Pool Name	Enter the exact name of the event pool on the Unity device that publishes CEE notifications. You can find this name in the Unity Manager UI under Event Notifications.
Agent IP	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the Unity device. The device must be able to

Field Description	Description
Addresses	route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

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.

5. Click **Next** to proceed to the [Source Paths](#) page.

NetApp ONTAP

To connect to the NetApp ONTAP Storage Virtual Machine (SVM) hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details:

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 Policies
 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.

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.
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.

Field	Description
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 is not listed, enter the address manually.
Access Path	Use only when experiencing access issues. Contact Peer Support for more information.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Validate** to test the credentials.

After the credentials are validated, a success message appears.

4. Click **Next** to proceed to the [Source Paths](#) page.

Nutanix Files

To connect to the Nutanix Files cluster hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

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 Policies
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation

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.

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 is not listed, enter the address manually. The IP address should not point to the Files cluster itself.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Validate** to test the credentials.


After the credentials are validated, a success message appears.

4. Click **Next** to proceed to the [Source Paths](#) page.

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 volumes selected for replication are referred to as the [watch set](#).

1. Select the paths to the volume 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 .
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** to proceed to the [File and Folder Filters](#) page.

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** to proceed to the [Destination](#) page.

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** to proceed to the [Destination Credentials](#) page.

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** to proceed to the **Details** page for the selected destination storage account.

Azure Blob Storage Credentials

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 and Folder 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:

*Endpoint Type:

☒ Use SSL

☐ Existing Credentials

Validate

< Back Next > Finish 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.
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** to proceed to the [Container Details](#) page.

Amazon S3 Credentials

1. Enter the credentials to connect to an Amazon S3 storage account.

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** to proceed to the [Bucket Details](#) page.

NetApp StorageGRID Credentials

1. Enter the credentials to connect to a NetApp StorageGRID storage account.

Create Cloud Backup and Replication Job Wizard

NetApp StorageGRID Credentials

Enter new credentials or select existing credentials.

Source Storage Platform

Management Agent

Proxy Configuration

Storage Information

Source Paths

File and Folder Filters

Destination

NetApp StorageGRID Credentials

Bucket Details

Replication and Retention Policies

Miscellaneous Options

Email Alerts

SNMP Notifications

Confirmation

Credentials

☒ New Credentials

*Description:

*Access Key:

*Secret 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 NetApp StorageGRID account, which can be found in the Tenant Manager.

Field	Description
Secret Key	Enter the secret key of the NetApp StorageGRID account, which can be found in the Tenant Manager.
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** to proceed to the [Container Details](#) page.

Nutanix Objects Credentials

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.

Source Storage Platform
Management Agent
Proxy Configuration
Storage Information
Source Paths
File and Folder Filters
Destination
Nutanix Objects Credentials
Bucket Details
Replication and Retention Policies
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

Credentials

☒ New Credentials

*Description:

*Access Key:

*Secret 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 Nutanix Objects account, which can be found in Prism Central.
Secret Key	Enter the secret key of the Nutanix Objects account, which can be found in Prism Central.
Service Point	Enter the IP or name of the object store.
Use SSL	Select this if you want to use the SSL (Secure Sockets Layer) protocol to communicate with the destination rather than the standard (non-encrypted) HTTP protocol.

2. Click **Validate** to test the connection.
3. Click **Next** to proceed to the [Container Details](#) page.

Wasabi Credentials

1. Enter the credentials to connect to a Wasabi storage account.

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.

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** to proceed to the [Bucket Details](#) page.

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** to proceed to the [Replication and Retention Policy](#) page.

Azure Blob Storage Container Details

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.

Create Cloud Backup and Replication Job Wizard

Container Details
Create a new container or select an existing one.

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

☒ New Container
*Name:
lab24win22a-ycwowpvkp
☒ Automatically name
☐ Existing Container

< Back Next > Finish 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** to proceed to the [Replication and Retention Policy](#) page.

Amazon S3 Bucket Details

1. Select **New Bucket** to create a new bucket or select **Existing Bucket**.

Choose **Existing Bucket** if:

- You (or someone else) already created a bucket you want to use.
- You want to use a bucket that was created outside Peer Management Center.
- You don't have the permissions required to create new buckets and want to use one that someone else will create.

2. If you selected **Existing Bucket**, select a bucket from the drop-down list. If the bucket does not appear in the list because the person who has the permissions to create a bucket has not yet created the bucket, click **Reload** after the bucket is created. The bucket will appear in the updated list.

If you selected **New Bucket**, you have two options. By default, the **Automatically name** checkbox is selected. You can deselect the checkbox and enter a name for the bucket; the bucket name must conform to the following naming rules:

- A bucket name must be unique across all existing bucket names in Amazon S3 (that is, across all AWS customers). For more information, see [Bucket Restrictions and Limitations](#).
- Bucket names must comply with DNS naming conventions. For information about legacy non-DNS-compliant bucket names, see [Bucket Restrictions and Limitations](#).
- A bucket name must start with a lowercase letter or number.
- A bucket name must not contain uppercase characters or underscores.

- A bucket name must be from 3 through 63 characters long.
- A bucket name must be a series of one or more labels. Adjacent labels are separated by a single period (.). Bucket names can contain lowercase letters, numbers, and hyphens. Each label must start and end with a lowercase letter or a number.
- A bucket name must not be formatted as an IP address (for example, 192.168.5.4).
- When you use virtual hosted-style buckets with Secure Sockets Layer (SSL), the SSL wildcard certificate only matches buckets that don't contain periods. To work around this, use HTTP or write your own certificate verification logic. We recommend that you do not use periods (.) in bucket names when using virtual hosted-style buckets.
- Choose a bucket name that reflects the objects in the bucket because the bucket name is visible in the URL that points to the objects that you're going to put in your bucket. After you create the bucket, you cannot change the name, so choose wisely.

For information about naming buckets, see [Rules for Bucket Naming](#) in the Amazon Simple Storage Service Developer Guide.

3. Select the region where you want the bucket to reside.

Important: After you have created a bucket, you cannot change its region.

4. Click **Next** to proceed to the [Replication and Retention Policy](#) page.

NetApp StorageGRID Bucket Details

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
NetApp StorageGRID Credentials
Bucket Details
Replication and Retention Policies
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

☒ New Bucket
☐ Existing Bucket

*Name:
lab24win22b-psnfilpjqvq
☒ 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; 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** to proceed to the [Replication and Retention Policy](#) page.

Nutanix Objects Bucket Details

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 and Folder Filters

Destination
Nutanix Objects Credentials
Bucket Details
Replication and Retention Policy
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

☒ New Bucket

*Name:
lab24win22a-moeidpgeai

☒ Automatically name

☐ Existing Bucket

< Back Next > Finish Cancel

- 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** to proceed to the [Replication and Retention Policy](#) page.

Wasabi Bucket Details

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 and Folder Filters
Destination
Wasabi Credentials
Bucket Details
Replication and Retention Policy
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

☒ New Bucket

*Name:
lab24win22a-lywuqbysdl

☒ Automatically name

☐ Existing Bucket

< 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** to proceed to the [Replication and Retention Policy](#) page.

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**.

Create Cloud Backup and Replication Job Wizard

Replication and Retention Policy

✖ You must enter a name for the policy.

Source Storage Platform
Management Agent
Proxy Configuration
Storage Information
Source Paths
File and Folder Filters
Destination
Wasabi Credentials
Bucket Details
Replication and Retention Policy
Replication Schedule
Retention
Source Snapshots
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

☒ New Policy
*Name:
☒ Enable Backup with Destination Snapshots
☐ Existing Policy

< Back Next > Finish Cancel

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** to proceed to the [Replication Schedule](#) page.

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** to proceed to the [Retention](#) page.

Scheduled Scans

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 and Folder Filters

Destination
Wasabi Credentials
Bucket Details

Replication and Retention Policies
Replication Schedule
Retention
Source Snapshots
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

☒ Scheduled Scans

☒ Daily ☐ Weekly

Times (up to 4)

None	☐ Take Destination Snapshot
None	☐ Take Destination Snapshot
None	☐ Take Destination Snapshot
None	☐ Take Destination Snapshot

☐ Scan at Start

☐ Batched Real-time

☐ Continuous Data Protection

< Back Next > Finish 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.

Create Cloud Backup and Replication Job Wizard

Replication Schedule

✖ You must select at least one day.

Source Storage Platform
Management Agent
Proxy Configuration
Storage Information

Source Paths
File and Folder Filters

Destination
Wasabi Credentials
Bucket Details

Replication and Retention Policies
Replication Schedule
Retention
Source Snapshots
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

☒ Scheduled Scans

☐ Daily ☒ Weekly

Day(s):

☐ Sunday
☐ Monday
☐ Tuesday
☐ Wednesday
☐ Thursday
☐ Friday
☐ Saturday

Time:
None ☒ Take Destination Snapshot

☐ Scan at Start

☐ Batched Real-time
☐ Continuous Data Protection

< Back Next > Finish Cancel

4. Click **Next**.

The [Retention](#) page appears.

Batched Real-Time

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.

Create Cloud Backup and Replication Job Wizard

Replication Schedule

✖ You must select at least one replication time.

☐ Scheduled Scans
☒ Batched Real-time

Times (up to 4)

None	☐ Take Destination Snapshot
None	☐ Take Destination Snapshot
None	☐ Take Destination Snapshot
None	☐ Take Destination Snapshot

☐ Continuous Data Protection

< Back Next > Finish Cancel

4. Click **Next**.

The [Retention](#) page appears.

Continuous Data Protection

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.

Create Cloud Backup and Replication Job Wizard

Replication Schedule

✖ You must select at least one destination snapshot trigger time.

Source Storage Platform
 Management Agent
 Proxy Configuration
 Storage Information
 Source Paths
 File and Folder Filters
 Destination
 Wasabi Credentials
 Bucket Details
 Replication and Retention Policies
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 > Finish 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** to proceed to the [Source Snapshots](#) page.

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** to proceed to the [Miscellaneous Options](#) page.

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.

Create Cloud Backup and Replication Job Wizard

Miscellaneous Options

Select options.

Source Storage Platform
Management Agent
Proxy Configuration
Storage Information
Source Paths
File and Folder Filters
Destination
Azure Storage Credentials
Container Details
Replication and Retention Policies
Miscellaneous Options
Email Alerts
SNMP Notifications
Confirmation

File Metadata

Permissions:

☐ Owner ☐ DACL ☐ SACL

Storage Tiering Options

Storage Tier/Class: Storage Account Default

Rehydrated Data Availability (Days): 7

< Back Next > Finish Cancel

Option	Description
NTFS Permissions	If you want NTFS permissions metadata included in the replication, select the elements to include: • Owner – The NTFS Creator-Owner who owns the object (which is, by default, 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.

Option	Description
	See File Metadata Synchronization for more information about NTFS permissions metadata.
Storage Tier/Class	Only available with Azure Blob Storage with either a General Purpose v2 (GPv2) account or Blob Storage Account. Select a storage tier. If you do not select a tier, it will default to the tier you configured on your Azure Storage account. Azure Storage offers three storage tiers for blob object storage so that you can store your data most cost-effectively depending on how you use it: • Azure Hot Storage Tier is optimized for storing data that is accessed frequently. • Azure Cool Storage Tier is optimized for storing data that is infrequently accessed and stored for at least 30 days. • Azure Archive Storage Tier is optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements (on the order of hours). The archive storage tier is only available at the blob level and not at the storage account level. To read data in archive storage, Cloud Backup and Replication must first change the tier of the blob to hot or cool. This process is known as rehydration and can take up to 15 hours to complete. Rehydrated data remains in hot or cool storage for a specified number of days before Cloud Backup and Replication automatically returns it to archive storage.
Rehydrated Data Availability (Days)	Only available with Azure Blob Storage with either a General Purpose v2 (GPv2) account or Blob Storage Account. Rehydrated data 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.

2. Click **Next** to proceed to the [Email Alerts](#) page.

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.

Create Cloud Backup and Replication Job Wizard

Email Alerts
Select email alerts.

[Edit Email Alerts](#)

Name	Temporarily Disabled	Event Types	Recipients
Participant Failure	No	Participant Failure	AriF@peersoftware.com

Buttons: Select, Delete, View Details

Navigation: < Back, Next >, Finish, Cancel

3. (Optional) Repeat steps 1-3 to apply additional alerts.
4. Click **Next** to proceed to the [SNMP Notifications](#) page.

Step 17: SNMP Notifications

This step is optional.

An [SNMP notification](#) notifies recipients when certain type of event occurs, for example, session abort, host failure, system alert. The **SNMP Notifications** page displays a list of notifications that have been applied to the job. When you first create a job, this list is empty. Like email alerts and file filters, an SNMP notification is defined in [Preferences](#) and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create SNMP notifications in advance. However, from this wizard page, you can select an existing SNMP notification to apply to the job or create new SNMP notifications.

To apply an existing SNMP notification to the job or disable notifications:

1. Select an SNMP notification from the drop-down list.

To disable, select **None - Disabled**.

The screenshot shows the 'Create Cloud Backup and Replication Job Wizard' window. The title bar is blue with the text 'Create Cloud Backup and Replication Job Wizard' and a close button. The window is divided into a sidebar on the left and a main content area on the right. The sidebar contains a tree view of configuration steps: Source Storage Platform, Management Agent, Proxy Configuration, Storage Information, Source Paths, File and Folder Filters, Destination, Wasabi Credentials, Bucket Details, Replication and Retention Policies, Replication Schedule, Retention, Source Snapshots, Miscellaneous Options, Email Alerts, **SNMP Notifications** (highlighted), and Confirmation. The main content area has a header 'SNMP Notifications' and a sub-header 'Select SNMP notification.'. Below this is a link 'Edit SNMP Notifications' and a dropdown menu labeled 'SNMP Notification:' with the value 'None - Disabled'. Underneath, a section titled 'Selected SNMP Notification Information' contains the text 'No SNMP Notification Selected' and 'SNMP notifications disabled for this job'. At the bottom of the window are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

2. Click **Next** to proceed to the [Confirmation](#) page.

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.

The screenshot shows the 'Create Cloud Backup and Replication Job Wizard' window, specifically the 'Confirmation' step. The window has a blue title bar and a sidebar on the left with a tree view of the wizard's steps. The main area displays the job configuration details. At the bottom, there is a checkbox for 'Start job after creation' and navigation buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Create Cloud Backup and Replication Job Wizard

Confirmation
Review your job configuration.

- Source Storage Platform
 - Management Agent
 - Proxy Configuration
 - Storage Information
- Source Paths
 - File and Folder Filters
- Destination
 - Wasabi Credentials
 - Bucket Details
- Replication and Retention Policy
 - Replication Schedule
 - Retention
 - Source Snapshots
 - Miscellaneous Options
 - Email Alerts
 - SNMP Notifications
 - Confirmation**

Source Storage: Nutanix Files
Management Agent: Lab24Win22a
Source Paths: One or more specific paths
Source Items:
Volume: afs
Destination: -1
Item(s) to Include: .\SMB1\AriF\
Destination: Wasabi
Replication and Retention Policy: Policy-2
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 (Policy-2):
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)

☐ 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.



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.

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 Hc	Storage Device	Storage Type	Destination Type	Destination I	Status	Mode	Scan Status	Next Scan R	Scan pending It	Real-time pendi	Re
C:\	CB-1	Lab24Win22a	Lab24Win22a	Windows	Nutanix Objects	lab24win22e	Running	Scheduled Scan	Last Scan: C: 08-10-2024	0 [0 bytes]	0 [0 bytes]	0	0

- 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.

Volume	Job	Management Hc	Storage Device	Storage Type	Destination Type	Destination I	Status	Mode	Scan Status	Next Scan R	Scan pending It	Real-time pendi	Re
C:\	CB-1	Lab24Win22a	Lab24Win22a	Windows	Nutanix Objects	lab24win22e	Running	Scheduled Scan	Last Scan: C: 08-10-2024	0 [0 bytes]	0 [0 bytes]	0	0

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 a blue header bar. Below the header, the title "Volume to Recover" is displayed in bold. Underneath, a subtitle reads "Select the storage device and volume that the data originated from." The main area is divided into two sections. On the left, a vertical list of options is shown: "Recovery Source", "Volume to Recover" (which is highlighted with a dark background), "Search By", "Recover To", "Notifications", and "Confirmation". On the right, there are two dropdown menus. The first is labeled "Storage Device:" and has "Lab24Win22a" selected. The second is labeled "Volume:" and has "C:\\" selected. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

5. Click **Next**.

The **Search By** page is displayed.



The screenshot shows a window titled "Recovery Wizard" with a standard Windows-style title bar (minimize, maximize, close buttons). The window is divided into two main sections. On the left is a vertical sidebar with a list of steps: "Recovery Source", "Volume to Recover", "Search By" (which is highlighted with a dark background), "Recover To", "Notifications", and "Confirmation". The main area on the right is titled "Search By" and contains the instruction "Select a search option." Below this instruction are four radio button options: "Name", "Snapshot", "Point in Time", and "Latest Replication". At the bottom of the window, there is a horizontal bar containing four buttons: "< Back", "Next >", "Finish", and "Cancel".

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](#)

Search by Name

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

Search By
Name
File/Folder Versions

Recover To
Notifications
Confirmation


Name: Any

Name	Sync Date	Last Modified	Size
------	-----------	---------------	------

< Back Next > Finish 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.

Recovery Wizard

File/Folder Versions

Select a version to recover.

Recovery Source

Volume to Recover

Search By

Name

File/Folder Versions

Recover To

Notifications

Confirmation

Selected Sync Item(s):

C:\CB-1\Backup\file_19.txt

<< < Aug, 2024 > >>

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

Select Time

Oldest Available:

08-07-2024 20:04:59

Newest Available:

08-07-2024 20:04:59

< Back

Next >

Finish

Cancel

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.

Recovery Wizard

File/Folder Versions

Select a version to recover.

Recovery Source
Volume to Recover
Search By
Name
File/Folder Versions
Recover To
Notifications
Confirmation

Selected Sync Item(s):
C:\CB-1\Documentation

☒ Snapshots ☐ Date and Time

Aug, 2024

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

Oldest Available:
08-07-2024 20:09:31

Newest Available:
08-07-2024 20:18:23

Select Time

< Back Next > Finish Cancel


- Click **Next** and continue with [Recovery Options](#).

Search by Snapshot

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:

- 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](#).

Search by Point in Time

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.

Recovery Wizard

Point in Time
Select a point in time.

Recovery Source
Volume to Recover

Search By
Point in Time
File/Folder Browser

Recover To
Notifications
Confirmation

Select a date and time:


Aug, 2024						
Su	Mo	Tu	We	Th	Fr	Sa
				1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

Select Time

< Back Next > Finish Cancel

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](#).

Search by Latest Replication

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.

Recovery Wizard

Recover To
Select a recovery location.

Recovery Source
Volume to Recover
Search By
File/Folder Browser
Recover To
Notifications
Confirmation

☒ Another Location ☐ Original Location

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

Recovery Options
If a file to recover already exists:

☒ Recover with unique name
☐ Overwrite if size or timestamps don't match
☐ Overwrite if cloud version is newer
☐ Overwrite always
☐ Skip

Metadata Options

☒ Recover Last Modified Time
☒ Recover Create Time
☐ Recover Permissions
☒ Recover Attributes

< Back Next > Finish Cancel

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.

Recovery Option	Select this option if you want to:
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.

Recovery Wizard

Notifications

Select email notifications.

Recovery Source

Volume to Recover

Search By

File/Folder Browser

Recover To

Notifications

Confirmation

Edit Email Alerts

☐ Send email notification when complete

☐ Only on failure

Recipients

Enter email, contact, or distribution list:

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

Add to List

Recipients:

Delete


View Details

< Back Next > Finish Cancel

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](#)
 - o [Adding a Namespace Server](#)
 - o [Adding a Namespace Folder](#)
 - o [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.



- 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](#).

Agent	Computer Description	Domain	DFS Mgmt. Enabled	Windows Version
Lab24Win19c		PEERTEST	Yes	Windows Server 2019
Lab24Win22a		PEERTEST	Yes	Windows Server 2022
Lab24Win22b		PEERTEST	No	Windows Server 2022

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 Windows-style wizard window titled "Create DFS-N Management Job Wizard". The "Namespace Name" step is selected in the left-hand navigation pane. The main area contains the instruction "Enter a name of the new namespace." and a text box labeled "Namespace Name:" with the example text "Example: Public". A note above the text box states: "This name will appear after the server or domain name in the namespace path, such as \\Server\\Name or \\Domain\\Name." 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**.

The screenshot shows the 'Create DFS-N Management Job Wizard' window, specifically the 'Namespace Servers' step. The window has a blue title bar and a sidebar on the left with the following menu items: Management Agent, Agent Verification, Namespace Name, Namespace Servers (highlighted), Namespace Settings, Namespace Folders, Email Alerts, SNMP Notifications, Review, and Results. The main area contains the text: 'Select one or more servers to host this namespace. The servers you select will be known as namespace servers.' Below this, there is a 'Server Name:' input field with a placeholder 'Enter the fully qualified domain name of a server running the DFS Namespaces service.' and an 'Add' button. Below the input field is a 'Servers:' list box, which is currently empty, and a 'Delete' button. At the bottom of the window are three buttons: '< Back', 'Next >', and 'Cancel'.

The server path is listed in the **Servers** area below.

This screenshot shows the same 'Create DFS-N Management Job Wizard' window, 'Namespace Servers' step, but with a server path added to the 'Servers:' list box. The 'Server Name:' input field is now empty. The 'Servers:' list box now contains the text 'lab24win19a.peertest.local'. The 'Delete' button is still present. The bottom buttons are '< Back', 'Next >', and '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.

Create DFS-N Management Job Wizard

Namespace Settings
Modify the settings of the shared folder.

Management Agent
Agent Verification
Namespace Name
Namespace Servers
Namespace Settings
Namespace Folders
Email Alerts
SNMP Notifications
Review
Results

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

Server Name	DFS Root Local Path	Permissions
lab24win19a.peertest.lc	C:\DFSRoots\Company Files	Everyone Full Access
lab24win19b.peertest.lc	C:\DFSRoots\Company Files	Everyone Full Access

< Back Next > Cancel

2. In the **Permissions** column for the server, select the desired access level from the drop-down menu.

Create DFS-N Management Job Wizard

Namespace Settings
Modify the settings of the shared folder.

Management Agent
Agent Verification
Namespace Name
Namespace Servers
Namespace Settings
Namespace Folders
Email Alerts
SNMP Notifications
Review
Results

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

Server Name	DFS Root Local Path	Permissions
lab24win19a.peertest.lc	C:\DFSRoots\Company Files	Everyone Full Access
lab24win19b.peertest.lc	C:\DFSRoots\Company Files	Everyone Full Access

< Back Next > Cancel


3. (Optional) Modify the path and permissions for other servers.
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 **Namespace Folder Name** dialog appears.

The screenshot shows the 'New Namespace Folder' dialog box with the 'Namespace Folder Name' tab selected. The instruction 'Enter a name for the namespace folder.' is at the top. The 'Folder Name' field contains '\\PeerTest.local\\Company Files\\'. Below it, the 'Preview of folder:' field shows the same path. The left sidebar has three tabs: 'Namespace Folder Name' (selected), 'Folder Targets', and 'Confirmation'. At the bottom are four buttons: '< Back', 'Next >', 'Finish', and '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.

This screenshot shows the same 'New Namespace Folder' dialog box, but the 'Folder Name' field now contains '\\PeerTest.local\\Company Files\\ Project Management'. The 'Preview of folder:' field also updates to show '\\PeerTest.local\\Company Files\\Project Management'. The 'Namespace Folder Name' tab remains selected, and the bottom buttons are the same.

3. Click **Next**.


The **Folder Targets** dialog appears.



The screenshot shows the 'New Namespace Folder' dialog box with the 'Folder Targets' tab selected. The dialog has a title bar with standard window controls. Below the title bar, the 'Folder Targets' section is active, displaying the instruction 'Select one or more folder targets for this folder.' On the left, a sidebar lists 'Namespace Folder Name', 'Folder Targets' (which is highlighted), and 'Confirmation'. The main area contains instructions to 'Enter the UNC path to an existing folder. If the folder is not shared, you must share it first.' Below this, there is a text input field 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 four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.


4. (Optional) To add a folder target, enter the UNC path to a shared folder, and then click **Add**.

If you haven't created your folders target yet, you can skip to Step 6 and [add folder targets to the job](#) later.



This screenshot shows the same 'New Namespace Folder' dialog box, but now a folder target has been added. The 'Path to folder target:' text field contains the UNC path '\\lab24win19c\Company Files\Project_Management - Tokyo', and the 'Add' button is still visible. The 'Folder targets:' list box now contains this path. The 'Delete' button remains to the right of the list box. The rest of the dialog, including the sidebar and bottom navigation buttons, is identical to the previous screenshot.

After the share is validated, it appears in the **Folder targets** area:



New Namespace Folder

Folder Targets

Select one or more folder targets for this folder.

Namespace Folder Name

Folder Targets


Confirmation

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

Path to folder target:

Folder targets:

5. (Optional) Add additional folder targets.



New Namespace Folder

Folder Targets

Select one or more folder targets for this folder.

Namespace Folder Name

Folder Targets

Confirmation


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

Path to folder target:

Folder targets:


6. Click **Next**.

The **Confirmation** dialog appears.




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

The **Folders** page reappears.



9. Expand the tree to view the folders and folder targets you added.



10. Click **Next**.

The [Email Alerts](#) page appears.

Step 8: Email Alerts


This step is optional.

An [email alert](#) notifies recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in [Preferences](#) and can then be applied to multiple jobs of the same type.


Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply. To create a new alert, see [Email Alerts](#) in the [Preferences](#) section.

To apply an existing email alert to the job:

1. Click the **Add** button.




The **Select Email Alert** dialog appears.



2. Select an alert from the **Email Alert** drop-down list, and then click **OK**.

The alert is listed in the **Email Alerts** page.



3. (Optional) Repeat steps 1-3 to apply additional alerts.
4. Click **Next**.

The [SNMP Notifications](#) page appears.

Step 9: SNMP Notifications

This step is optional.


An [SNMP notification](#) notifies recipients when certain type of event occurs, for example, session abort, host failure, system alert. The **SNMP Notifications** page displays a list of notifications that have been applied to the job. When you first create a job, this list is empty. Like email alerts and file filters, an SNMP notification is defined in [Preferences](#) and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create SNMP notifications in advance. However, from this wizard page, you can select an existing SNMP notification to apply to the job or create new SNMP notifications. To create a new alert, see [SNMP Notifications](#) in the [Preferences](#) section.


To apply an existing SNMP notification to the job or disable notifications:

1. Select an SNMP notification from the drop-down list.

If you don't want to send a SNMP notification, select **None - Disabled**.



If you select a notification, details about the notification appear in the Selected SNMP Notification Information section.




2. Click **Next**.

The [Review](#) page appears.

Step 10: Review

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

1. Review the namespace configuration.




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

After you click **Create**, the [Results](#) page appears.

Step 11: Results


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

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




2. Review each task to verify that it was successful; 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.



4. Select the namespace in the **Namespace** tab and then select the namespace folder to view its folder targets.

The folder targets appear in the panel below.




Importing an Existing Namespace

If you have an existing namespace that you want to use in in a File Collaboration or File Synchronization job, you can import the namespace. Importing the namespace automatically creates a new DFS-N Management job with the same name as the imported namespace.

You can then either [link the namespace to an existing File Collaboration or File Synchronization job](#) or [create a new File Collaboration or File Synchronization job](#) that uses the namespace.

To import an existing namespace:

1. Right-click anywhere in the **Runtime Summary** tab of the **Namespace Summary** view, and then select **Import Existing Namespaces** (or right-click the DFS-N Management job type in the **Jobs** view).



The **Import Existing Namespace** wizard appears.

2. Select a Management Agent, and then click **Next**.

Import Existing Namespaces


Management Agent

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

Management Agent	Agent	Computer Description	Domain	DFS Mgmt. Enabled	Windows Version
Agent Verification	Lab24Win19c		PEERTEST	Yes	Windows Server 2019
Namespace	Lab24Win22a		PEERTEST	Yes	Windows Server 2022
Email Alerts	Lab24Win22b		PEERTEST	No	Windows Server 2022
SNMP Notifications					
Confirmation					
Results					


< Back Next > Cancel

3. Click **Start Verification** to verify the Agent environment is set up to manage DFS namespaces.



4. If the DFS PowerShell Management toolkit is not installed, click the **Install** button that will 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


< Back Next > Cancel

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 screenshot shows a web application window titled "Import Existing Namespaces". On the left is a sidebar menu with the following items: "Management Agent", "Agent Verification", "Namespace", "Email Alerts" (which is highlighted with a dark background), "SNMP Notifications", "Confirmation", and "Results". The main content area is titled "Email Alerts" and contains the text "Select email alerts." Below this, there is a link "Edit Email Alerts". A table with four columns is displayed: "Name", "Temporarily Disabled", "Event Types", and "Recipients". The table is currently empty. To the right of the table are three buttons: "Add", "Delete", and "View Details". At the bottom of the window are three navigation buttons: "< Back", "Next >", and "Cancel".

Name	Temporarily Disabled	Event Types	Recipients
------	----------------------	-------------	------------

The **SNMP Notifications** page appears.

10. (Optional) [Select or create an SNMP notification](#) to apply to the job, and then click **Next**.

Import Existing Namespaces

SNMP Notifications

Select a SMNP notification.

Management Agent
Agent Verification
Namespace
Email Alerts
SNMP Notifications
Confirmation
Results

[Edit SNMP Notifications](#)

SNMP Notification: None - Disabled

Selected SNMP Notification Information

No SNMP Notification Selected
SNMP notifications disabled for this job

< Back Next > Cancel

The **Confirmation** page appears.


The screenshot shows a web application window titled "Import Existing Namespaces". The window has a blue header bar with the title and standard window controls. Below the header, the main content area is divided into two sections. On the left is a navigation pane with a list of options: "Management Agent", "Agent Verification", "Namespace", "Email Alerts", "SNMP Notifications", "Confirmation" (which is highlighted with a dark background), and "Results". The main area on the right is titled "Confirmation" and contains the text "Review your configuration prior to importing the namespace." Below this, there is a message: "If the configuration below is correct, click Import. To change the configuration, click Back or select the appropriate page in the navigation panel." The configuration details shown are "Selected Agent: Lab24Win19c" and "Selected Namespaces: - Design". At the bottom right of the window, there are three buttons: "< Back", "Import", and "Cancel".

11. Review the configuration.

12. If you need to modify the job configuration after reviewing it, click **Back** until you reach the appropriate page and make your changes.


13. Once you are satisfied with the job configuration, click **Import**.

The **Results** page appears.



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

A DFS-N Namespace job is created for each namespace you added. The new job(s) are displayed in the **Jobs** view. Runtime views for the jobs are also displayed. The jobs automatically start running. The namespaces can now be [linked to File Collaboration and File Synchronization jobs](#).



Running a DFS-N Management Job


This section describes:

- [Starting a DFS-N Management Job](#)
- [Stopping a DFS-N Management Job](#)


Starting a DFS-N Management Job

To manually start a DFS-N Management job:

1. Choose one of these options:
 - Right-click the job name in the **Jobs** view, and then select **Start**.




- Open the job and then click the **Start/Stop** button in the bottom left corner of the job's **Namespace** tab in the [DFS-N Management runtime view](#). You may need to scroll to the bottom of the tab to see the **Start/Stop** button.



- 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 Global File Service interface. The left pane shows a tree of DFS-N Management jobs, with 'DFS-N Management (3)' expanded. The main pane displays a table of namespace folders and their targets. The bottom pane shows a list of job alerts.

Namespace Path/Folders	Management Status	State	Errors	Servers	Total Folders/Targets
\\PeerTest.local\Company Files	Monitoring	Online	0	lab24win19a.peerTest.local	1 Folder
\\PeerTest.local\Design	Monitoring	Online	0	LAB24WIN19A.PEERTEST.LOCAL	0 Folder
\\PeerTest.local\Tickets	Monitoring	Online	0	lab24win19a.peerTest.local	1 Folder

Received Date	Severity	Type	Name	Host	Message	Exception
08-13-2024 01:00	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Blob Tiering Maintenance	
08-13-2024 01:00	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Blob Tiering Maintenance	
08-12-2024 03:59	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Nightly Delete Shadow Copies Job	
08-12-2024 23:59	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Nightly Delete Shadow Copies Job	
08-12-2024 23:30	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Blob Tiering Database Maintenance	
08-12-2024 23:30	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Blob Tiering Database Maintenance	
08-12-2024 22:37	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Process Retention Policy	
08-12-2024 22:37	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Process Retention Policy	
08-12-2024 20:10	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Nightly Purge Destination Job	
08-12-2024 20:10	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Nightly Purge Destination Job	

The runtime view for the job is displayed.

The screenshot shows the Peer Global File Service interface. The 'Jobs' tab is selected, displaying a tree view on the left and a main pane on the right. The main pane shows a table with columns: Folder Target, Linked to Participant, Status, Reason, AD Site, and Replication Status. Below this, there is a 'Status: Monitoring' section and a 'Job Alerts' section with a table of alerts.


Received Date	Severity	Type	Name	Host	Message	Exception
08-13-2024 01:00	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Blob Tiering Maintenance	
08-13-2024 01:00	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Blob Tiering Maintenance	
08-12-2024 23:59	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Nightly Delete Shadow Copies Job	
08-12-2024 23:59	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Nightly Delete Shadow Copies Job	
08-12-2024 23:30	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Blob Tiering Database Maintenance	
08-12-2024 23:30	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Blob Tiering Database Maintenance	
08-12-2024 22:37	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Process Retention Policy	
08-12-2024 22:37	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Process Retention Policy	
08-12-2024 20:10	Info	Agent	Scheduled Task Started	Lab24Win22a	Started Nightly Purge Destination Job	
08-12-2024 20:10	Info	Agent	Scheduled Task Started	Lab24Win22b	Started Nightly Purge Destination Job	

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

The screenshot shows the Peer Global File Service interface with the 'Namespace Servers' tab selected. The main pane displays a table with columns: Namespace Server, Status, Linked to Participant, Reason, AD Site, and Path. The table shows one entry: lab24win19a.peertest.local, ONLINE, US-HQ, and Path: \\lab24win19a.peertest.local.


Namespace Server	Status	Linked to Participant	Reason	AD Site	Path
lab24win19a.peertest.local	ONLINE			US-HQ	\\lab24win19a.peertest.local

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



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

- 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.

Add DFS Namespace Server

Namespace Servers

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

Namespace Servers

Enter the fully qualified domain name of a server running the DFS Namespaces service.

Server Name:

Servers:

lab24win19c.peertest.local	Delete
----------------------------	---------------------------------------

< Back Next > Cancel

5. Add additional servers if desired.
6. Click **Next**.

The **Namespace Settings** page is displayed.

Add DFS Namespace Server

Namespace Settings

Modify the settings of the shared folder.

Namespace Servers

Namespace Settings

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

Shared Folder: Design


Server Name	DFS Root Local Path	Permissions
lab24win19c.peertest.local	C:\DFSRoots\Design	Everyone Full Access

< Back Next > Cancel

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.


- Click **Next**.

The **Confirmation** page is displayed.




- Review the namespace server configuration.
- Click **Add** if the configuration is correct; otherwise, click **Back** and correct the configuration.

The **Results** page is displayed.



- Click **OK**.

The newly added server is listed in the **Namespace Servers** tab.



Adding a Namespace Folder


You can add a namespace folder to a namespace. At the same time, you can its folder targets or you can [add folder targets](#) later.

Note that PeerGFS does not currently support creation of nested namespace folders. In other words, you cannot add a namespace folder that is a subfolder of a namespace folder, although this can be done when using the Microsoft DFS Management Tool.


A DFS-N Namespace job must be running before you can edit it.

To add a namespace folder to a namespace:


1. Double-click the DFS-N Management job name in the **Jobs** view or in the [Namespace Summary view](#) to open the runtime view for the job.



The runtime view for the job is displayed.




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




The **New Folder** wizard appears.

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




As you enter the folder name, a preview of the folder name and path appears below.



The dialog box is titled "New Namespace Folder". It has a sidebar on the left with four items: "Namespace Folder Name" (selected), "Folder Targets", "Confirmation", and "Results". The main area contains the text "Enter a name for the namespace folder." Below this, there is a "Folder Name:" label followed by a text box containing "\\PeerTest.local\\Company Files\\Security". Below that is a "Preview of folder:" label followed by a text box containing "\\PeerTest.local\\Company Files\\Security". At the bottom right, there are three buttons: "< Back", "Next >", 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 dialog box is titled "New Namespace Folder". It has a sidebar on the left with four items: "Namespace Folder Name", "Folder Targets" (selected), "Confirmation", and "Results". The main area contains the text "Select one or more folder targets for this folder." Below this, there is a "Path to folder target:" label followed by a text box containing "Example: \\Server\\Shared Folder\\Folder Target" and an "Add" button. Below that is a "Folder targets:" label followed by a large empty text box and a "Delete" button. At the bottom right, there are three buttons: "< Back", "Next >", 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' tab selected. The dialog has a title bar with a close button. Below the title bar, the 'Folder Targets' section is active, showing instructions to 'Select one or more folder targets for this folder.' On the left, a sidebar lists 'Namespace Folder Name', 'Folder Targets' (selected), 'Confirmation', and 'Results'. The main area contains a text box for 'Path to folder target:' with the value '\\Lab24Win19c\Company Files\Security - Tokyo' and an 'Add' button. Below this is a list box for 'Folder targets:' which is currently empty, with a 'Delete' button to its right. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.


7. Click **Add**.

The folder target is added to the **Folder targets** section.

This screenshot shows the same 'New Namespace Folder' dialog box, but now the 'Folder targets:' list box contains the path '\\Lab24Win19c\Company Files\Security - Tokyo'. The 'Add' button is still present, and the 'Delete' button is now active next to the listed target. The rest of the interface remains the same as the previous screenshot.

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 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 'New Folder Target' dialog box is shown with the 'Folder Targets' tab selected. The 'Folder Targets' section on the left has three sub-tabs: 'Folder Targets' (selected), 'Confirmation', and 'Results'. The main area contains instructions: 'Enter the UNC path to an existing folder. If the folder is not shared, you must share it first.' Below this, the 'Path to folder target:' field contains '\\Lab24Win22b\\Company Files\\Marketing - Centreville', and the 'Add' button is visible. The 'Folder targets:' list is empty, and the 'Delete' button is disabled. At the bottom, there are '< Back', 'Next >', and 'Cancel' buttons.

4. Click **Add**.


The folder target is added to the **Folder targets** section.



The 'New Folder Target' dialog box is shown with the 'Folder Targets' tab selected. The 'Folder Targets' section on the left has three sub-tabs: 'Folder Targets' (selected), 'Confirmation', and 'Results'. The main area contains instructions: 'Enter the UNC path to an existing folder. If the folder is not shared, you must share it first.' Below this, the 'Path to folder target:' field contains 'Example: \\Server\\Shared Folder\\Folder Target', and the 'Add' button is visible. The 'Folder targets:' list now contains one entry: '\\Lab24Win22b\\Company Files\\Marketing - Centreville', and the 'Delete' button is enabled. At the bottom, there are '< Back', 'Next >', and 'Cancel' buttons.


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:
 - o 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.
 - o 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.


- Open the **Namespace** tab if it is not already displayed.



- 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.

- Click **Next** if you want to create folder targets; otherwise click **Finish** and continue with Step 9.

You can create folder targets later if you wish. See [Adding a Namespace Folder Target](#) for step-by-step instructions.


If you clicked **Next**, the **DFS Link** page appears. The purpose of this page is to link the watch path of each participant of a collaboration/synchronization job (specified in the **Path** page) to the appropriate folder target. Initially, the **Folder Target** column in the **Linked Participants and Folder Targets** table will not contain any folder targets. After linking, a folder target should be listed in this column for each participant.



7. Link the participants to folder targets using one or more of the following methods:


- Click **Auto Select Targets** to have PeerGFS select the targets for you. If you use this method, PeerGFS will try to match the watch path specified for each participant. It will populate the **Folder Target** column with its matches.
- Manually link the participants by typing a path in the **Folder Target** column.
- Use a combination of auto select and manually select. For example, you can use auto select, and if the correct targets are not selected, you can manually enter the folder target path in the Folder Target column.
- If an incorrect target appears in the **Folder Target** column, you can select the incorrect target, and then click **Remove Selected Target**.

After linking targets, the **Linked Participants and Folder Targets** table should contain a folder target for each participant.




8. If the **Link Enabled** column is blank, select **Yes** for each participant.
9. Click **Finish**.

The [runtime view](#) for the newly created File Collaboration or File Synchronization job is displayed.



10. Click the DFS-N Management job tab to display its runtime view.

The **Linked to Participant** column is now populated.




Linking an Existing Namespace Folder to an Existing File Collaboration or File Synchronization Job

You can link an existing DFS namespace folder with an existing File Collaboration or File Synchronization job. These steps require that the DFS namespace has been already created and is being managed by a DFS-N Management job.


To link a namespace folder with an existing File Collaboration or File Synchronization job:

1. Select the File Collaboration or File Synchronization job in the **Jobs** view.




2. Right-click and select **Edit Job**.

The **Edit Job** wizard appears.




3. Select **DFS-N Link** in the navigation tree on the left.

The DFS-N Link page is displayed.




4. In the **Namespace** field, select the namespace you want to link.




Once you've selected a namespace, a list of available namespace folders appears in the **Folder** drop-down list.

5. In the **Folder** field, select the namespace folder.



Once you've selected a namespace and a folder, you need to link each participant to a folder target.



- Click **Auto Select Targets** to have PeerGFS attempt to automatically map the participants with folder target.

After auto selection, the linked participants and folder targets are displayed in the **Linked Participants and Folder Targets** table.

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

Link a folder to an existing file collaboration or file synchronization job.

Namespace: \\PeerTest.local\Design

Folder: Media

Configure Auto Create Targets

Linked Participants and Folder Targets:

Agent	Directory	Folder Target	Link Enabled
Lab24Win22b	C:\Media Assets - Centreville	\\Lab24Win22b\Media - Centreville	Yes
Lab24Win22a	C:\Media Assets - Munich	\\Lab24Win22a\Media - Munich	Yes
Lab24Win19c	C:\Media Assets - Tokyo	\\Lab24Win19c\Media - Tokyo	Yes

Auto Select Targets

Remove Selected Target

Revert

OK Cancel

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.

Auto Create DFS Targets

Auto Create DFS Targets

Would you like to automatically create the following DFS targets for namespace: \\PeerTest.local\Design

Auto Select Targets Press "Auto Select Targets" to start auto matching roots with existing targets or shares on servers

Host	Root Path	DFS-N Folder Target
Lab24Win19c	C:\Media Assets - Tokyo	
Lab24Win22a	C:\Media Assets - Munich	
Lab24Win22b	C:\Media Assets - Centreville	

< Back Apply Cancel

- Review the values in the **Link Enabled** column; if **No** appears, select **Yes** from the drop-down list.
- 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

File Collaboration jobs provide distributed teams with a fast and efficient solution to collaborate with shared project files. These jobs synchronize participant files in real-time, incorporating distributed locking to prevent version conflicts. This functionality ensures data consistency across geographically dispersed locations.

Unlike centralized file collaboration solutions that can result in slow access speeds over wide area networks (WANs), a File Collaboration job replicates shared project files to each office location in a distributed environment. This design guarantees end users experience high-speed local area network (LAN) access to shared regardless of file size, minimizing latency and improving overall responsiveness. Integrated distributed file locking eliminates the risk of version conflicts, maintaining data integrity and simplifying administration.

By keeping frequently accessed ("hot") data local, File Collaboration maximizes user productivity and reduces WAN bandwidth consumption. Proximity to users, their applications, and computer resources delivers optimal performance from a physical perspective, minimizing I/O bottlenecks. Robust file locking guarantees consistent versions, facilitating seamless collaboration and reducing the need for manual conflict resolution.

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: Master-Edge Assignment](#)

[Step 4: File Metadata](#)

[Step 5: Application Support](#)

[Step 6: Email Alerts](#)

[Step 7: DFS Namespace](#)


[Step 8: 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.



- Click **File Collaboration**, 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

On the **Participants** page, you select a file-sharing protocol for the job and start the process of adding participants to your job.

A File Collaboration job requires at least two participants. Each participant consists of a **Management Agent** and a **storage object** (volume, share, export, or folder). After you add a participant, it appears in the **Participants** table. The table lists each participant's Management Agent, storage path, and other key information needed for the File Collaboration job.

Follow these steps:

1. Select the file-sharing protocol that matches the workload for this job: **SMB** or **Multi-Protocol**. The file-sharing protocol you select applies to the entire job and cannot be changed after the job is created. Available protocol options vary based on your environment and license.

Note: Multi-Protocol requires specific license support. Before creating a job that uses both SMB and NFS, verify that your license includes Multi-Protocol support. For more information about the Multi-Protocol feature, see [Multi-Protocol](#) in the [Advanced Topics](#) section.

Agent Type	License Supports Multi-Protocol?	Available Protocol(s)	Default Selection
All Windows-based	N/A	SMB	SMB is preselected and required
All Linux-based	Yes	Multi-Protocol	Multi-Protocol is preselected and required
Mixture of Windows- and Linux-based	Yes	SMB or Multi-Protocol	No default—you must choose

Create File Collaboration Job Wizard

Participants
Add two or more participants to this File Collaboration job.

Participants
File Metadata
Application Support
Email Alerts
DFS Namespace

Protocol: **SMB**

Agent	Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role

Add
Edit
Delete

< Back Next > Finish Cancel

2. Click **Add** to open the **Add New Participant** subwizard, which will walk you through the process of adding a participant.

Management Agent

The **Management Agent** page lists the available [Agents](#) in your environment. After selecting a file-sharing protocol, only compatible Agents will be available:

- **SMB** requires Windows-based Agents.
- **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.

You can assign multiple Agents to the same storage device. However, each Agent must be used to access different volumes, shares, exports, or folders on that device.

To select an Agent:

1. From the **Agents** table, select an Agent that can access the storage device you want to replicate. This Agent will be designated as a Management Agent.

Tip: If the Agent you want is not shown in the **Agents** table, 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](#), the **Agents** table will be updated.

The screenshot shows a window titled "Add New Participant" with a sidebar on the left containing "Management Agent", "Storage Platform", "Storage Information", "Path", and "Edge Caching". The "Management Agent" section is active, displaying the instruction "Select the server hosting the Peer Agent that manages this storage device." and a table of available agents. A message above the table states: "Only Microsoft Windows Agents are available based on job type." The table has three columns: "Agent", "OS", and "Computer Description". It lists four agents: DGWin16A, DGWin16B, DGWin16C, and DGWin16D, all running on "Microsoft Windows". At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Agent	OS	Computer Description
DGWin16A	Microsoft Windows	
DGWin16B	Microsoft Windows	
DGWin16C	Microsoft Windows	
DGWin16D	Microsoft Windows	

2. Click **Next** to continue to the [Storage Platform](#) page.

Storage Platform

The **Storage Platform** page lists the available storage platforms types for your File Collaboration job.

1. Select the storage platform type hosting the data you want to replicate.

The types available depend on the file sharing protocol in the previous step:

- If you selected **SMB**, these are the available storage platform types:

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. Inside the window, the "Storage Platform" tab is selected in a left-hand sidebar. The main area of the window displays the text "Select the type of storage platform." followed by 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 PowerStore
- ☐ Dell Unity
- ☐ Nutanix Files

At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

- If you selected **Multi-Protocol (SMB and NFS)**, these are the available storage platform types:

The screenshot shows a window titled "Add New Participant" with a blue header bar. Below the header, the "Storage Platform" tab is selected in a sidebar on the left. The main area 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 PowerStore, Dell Unity, and Nutanix Files. At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. Click **Next** to proceed to the [Storage Information](#) page.

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**, provide the required configuration details. See [Windows 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. If you selected **Existing Credentials**, choose a credential from the drop-down list, and then click **Next** to proceed to the [Path](#) page.
4. If you selected **New Credentials**, enter the required connection details. The fields vary by platform. For platform-specific field descriptions, see:

[Amazon FSx for NetApp ONTAP](#)

[Dell PowerScale](#)

[Dell PowerStore](#)

[Dell Unity](#)

[NetApp ONTAP](#)

[Nutanix Files](#)

5. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
6. Click **Next** to proceed to the [Path](#) page.

Amazon FSx for NetApp ONTAP

To connect to the Amazon FSx for NetApp ONTAP Storage Virtual Machine (SVM) hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details:

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

*SVM Name:

*SVM User Name:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

Advanced

☐ Existing Credentials

Edit

Validate

You must enter the name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Amazon FSxN environments.

< Back

Next >

Finish

Cancel

Field	Description
SVM Name	Enter the name, FQDN (Fully Qualified Domain Name), or IP address of the 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 SVM management account. This must not be a cluster management account.
SVM Management IP	(Optional) Enter the IP address for accessing the SVM's management API. This is not required if the data LIFs (Logical Interfaces) associated with SVM Name already allow 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 is not listed, enter the address manually.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
4. Click **Next** to proceed to the [Path](#) page.

Dell PowerScale

To connect to the Dell PowerScale cluster hosting the data to be replicated:

1. When you have selected **New Credentials**, enter the required connection details.

The available connection types depend on the operating system of the Agent, as well as the notification delivery method used by the storage platform.

- o **RabbitMQ**: Available only when the Agent is running on Windows. Choose RabbitMQ if you are using a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent.

- o **Syslog**: Available on both Windows and Linux Agents. Choose Syslog if the PowerScale device sends notifications directly to the Agent.

If you selected **Rabbit MQ**, this page appears:

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 with a dark background), "Path", and "Edge Caching". 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, there are two sections for "Credentials". The first section, "New Credentials", is selected with a radio button. It contains five text input fields: "*SmartConnect Zone Name:", "*Cluster Management IP:", "*Cluster Username:", "*Cluster Password:", and "Access Zone:". Below these fields is a "Connection Type:" label with two radio buttons: "Syslog" and "RabbitMQ", where "RabbitMQ" is selected. To the right of the "RabbitMQ" button is an "Advanced" button. The second section, "Existing Credentials", is unselected. It contains a text input field and an "Edit" button. At the bottom of the main area, there is a "Validate" button followed by the text "You must enter a Cluster Name." and a line of text: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments." At the very bottom of the window, there 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

*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.

< Back

Next >

Finish

Cancel

These fields appear in both cases:

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 for accessing the REST-based API integrated into the PowerScale 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.
Access Zone (RabbitMQ) or Access Zones (Syslog)	Optional. Enter the name(s) of the access zone(s) being monitored.
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: - Choose RabbitMQ if you're use a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent. - Choose Syslog if the PowerScale device sends notifications directly to the Agent.

3. If you selected **Syslog**, enter values for the following additional 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 is not listed, enter the address manually.
Listening Port	Enter the port through which the Agent will receive TLS-based syslog events from the PowerScale cluster.
SSL Certificate Path	Enter the path to the certificate used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see the Dell PowerScale Syslog Configuration Guide .
SSL Private Key Path	Enter the path to the private key for the TLS connection. For more information, see the Dell PowerScale Syslog Configuration Guide .
SSL Private Key Password	(Optional) Enter the password for the private key if it's protected. For more information, see the 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 PowerStore

UniTo connect to the Dell PowerStore NAS server hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details details in the **Credentials** section.

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

*NAS Server Name:

*Management IP:

*Management Username:

*Management Password:

Connection Type: ☒ CEE

CEE Configuration

*Event Pool Name:

*Event Publisher Name:

*Agent IP Address:

*Listening Port:

☐ Existing Credentials

You must enter a NAS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerStore environments.

Field	Description
NAS Server Name	Enter the name of the PowerStore server hosting the data to be replicated.
Management IP	Enter the IP address for accessing the REST-based API integrated into the PowerStore server.
Management Username	Enter the user name for the account managing the PowerStore server.
Management Password	Enter the password for account managing the PowerStore server.
Connection Type	Select the method for sending real-time event notifications to the Agent: • CEE: Choose if you are using the CEE framework to dispatch notifications to the Agent.

4. Enter values for the required fields in the **CEE Configuration** section.

Field Description	Description
Event Pool Name	Enter the exact name of the event pool on the PowerStore device that publishes CEE notifications. You can find this name in the PowerStore Manager UI under Event Notifications.
Event Publisher Name	Enter the exact name or ID of the event publisher on the PowerStore device that sends CEE notifications to the Agent. You can find this value in the PowerStore Manager UI under Event Notifications.
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the PowerStore device. The device must

Field Description	Description
	be able to route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

- Click **Advanced** if you want to set [advanced options](#).
- Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
- Click **Next** to proceed to the [Path](#) page.

Dell Unity

To connect to the Dell Unity NAS Server hosting the data to be replicated:

- If you selected **New Credentials**, enter the required connection details details in the **Credentials** section.

The available connection types depend on the operating system of the Agent, as well as the notification delivery method used by the storage platform.

- o **RabbitMQ**: Available only when the Agent is running on Windows. Choose RabbitMQ if you are using a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent.
- o **Syslog**: Available on both Windows and Linux Agents. Choose Syslog if the PowerScale device sends notifications directly to the Agent.

If you selected **Rabbit MQ**, 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

*NAS Server Name:

*Unisphere Management IP:

*Unisphere Username:

*Unisphere Password:

Connection Type: ☒ RabbitMQ ☐ CEE

☐ Existing Credentials

MLUNITY, user:admin

You must enter a NAS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments.

If you selected **CEE**, this page appears:

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), "Path", and "Edge Caching". The main area of the window is titled "Storage Information" and contains the instruction "Enter the information required to connect to the storage device." The main area is divided into two sections. The top section is titled "Credentials" and contains a radio button labeled "New Credentials" which is selected. Below this are four text input fields: "*NAS Server Name:", "*Unisphere Management IP:", "*Unisphere Username:", and "*Unisphere Password:". Below these fields are two radio buttons for "Connection Type": "RabbitMQ" and "CEE", with "CEE" selected. Below the "Connection Type" section is a "CEE Configuration" section with three text input fields: "*Event Pool Name:", "*Agent IP Address:" (with a dropdown arrow), and "*Listening Port:" (with a value of "9843" and a spinner control). To the right of the "CEE Configuration" section is an "Advanced" button. Below the "CEE Configuration" section is a radio button labeled "Existing Credentials". Below this is a text input field containing "MLUNITY, user:admin" and an "Edit" button. At the bottom of the "Credentials" section is a "Validate" button and a message: "You must enter a NAS Server Name." Below the "Validate" button is a message: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments." The bottom of the window has four buttons: "< Back", "Next >", "Finish", and "Cancel".

These fields appear in both cases:

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 managing the Unity storage devices. This should not be the IP address of the NAS server itself.
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. If you selected **CEE**, enter values for the following additional fields in the **CEE Configuration** section.

Field Description	Description
Event Pool Name	Enter the exact name of the event pool on the Unity device that publishes CEE notifications. You can find this name in the Unity Manager UI under Event Notifications.
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the Unity device. The device must be able to route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

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.

NetApp ONTAP

To connect to the NetApp ONTAP Storage Virtual Machine (SVM) hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details:

The screenshot shows a web-based configuration window titled "Add New Participant". The "Storage Information" tab is selected in the left-hand navigation pane, which also includes "Management Agent", "Storage Platform", "Path", and "Edge Caching". The main content area is titled "Storage Information" with 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, indicated by a selected radio button. It 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. Below the "New Credentials" section is the "Existing Credentials" section, which is currently unselected. It features a dropdown menu showing "PTSVM93A, user:vsadmin" and an "Edit" button. At the bottom of the form, 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 are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Field	Description
SVM Name	Enter the name, FQDN (Fully Qualified Domain Name), or IP address of the 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 SVM Name already allow 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 is not listed, enter the address manually.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
4. Click **Next** to proceed to the [Path](#) page.

Nutanix Files

To connect to the Nutanix Files cluster hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details:

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

*Nutanix File Server Name:

*Username:

*Password:

*Peer Agent IP:

Advanced

☐ Existing Credentials

AFS5, user:peeradmin

Edit

Validate

Nutanix File Server Name cannot be empty.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Nutanix Files environments.

< Back

Next >

Finish

Cancel

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 is not listed, enter the address manually. 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

To connect to the Windows File Server hosting the data to be replicated:

1. Select the detector type for the Windows File Server. **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 Support.

The screenshot shows a window titled "Add New Participant" with a blue header bar. Below the header, the "Storage Information" tab is selected, showing a list of options: "Management Agent", "Storage Platform", "Storage Information" (highlighted), "Path", and "Edge Caching". To the right, the "Detector Type" is set to "Windows Driver" in a dropdown menu. An "Advanced" button is located to the right of the dropdown. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".


2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Next** to proceed to the [Path](#) page.

When configuring a job that uses a **Windows File Server**, you can set advanced options to control how the system filters, logs, and follows file activity.

The available options depend on the selected driver. Some options are common to both the **Windows driver** and **Windows Filter driver**, while others are unique to the **Windows Filter driver**.

1. Modify the Windows Detector options as desired.

If you selected **Windows**, this page appears:



The screenshot shows the "Windows Detector Options" dialog box. It has a blue title bar with the text "Windows Detector Options" and standard window controls. The main area contains several settings: a text box for "Filter open/close events from these users:", a spinner box for "Access Event Suppression Time:" set to "-1", and a section titled "Reparse Point Options" containing three checked checkboxes: "Follow Junction Points", "Follow Mount Points", and "Follow Symbolic Links". At the bottom right are "OK" and "Cancel" buttons.

Windows Detector Options

Filter open/close events from these users:


Access Event Suppression Time:

Reparse Point Options

- ☒ Follow Junction Points
- ☒ Follow Mount Points
- ☒ Follow Symbolic Links

OK Cancel

If you selected **Windows Driver**, this page appears:



The screenshot shows the "Windows Driver Detector Options" dialog box. It has a blue title bar with the text "Windows Driver Detector Options" and standard window controls. The main area contains several settings: four text boxes for "Filter open/close events from these users:", "Filter all events from these users:", "Filter events from these IP Addresses:", and "Filter events from these local processes:"; a spinner box for "Access Event Suppression Time:" set to "-1"; three checkboxes: "Enable Local Access Events:" (unchecked), "Enable Remote IP Address Logging:" (checked), and "Enable Close Modify:" (checked); and a text box for "Close Modify Extension Override:". Below these is a section titled "Reparse Point Options" containing three checked checkboxes: "Follow Junction Points", "Follow Mount Points", and "Follow Symbolic Links". At the bottom right are "OK" and "Cancel" buttons.

Windows Driver Detector Options

Filter open/close events from these users:

Filter all events from these users:

Filter events from these IP Addresses:

Filter events from these local processes:

Access Event Suppression Time:

Enable Local Access Events: ☐

Enable Remote IP Address Logging: ☒

Enable Close Modify: ☒

Close Modify Extension Override:

Reparse Point Options

- ☒ Follow Junction Points
- ☒ Follow Mount Points
- ☒ Follow Symbolic Links

OK Cancel

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 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 can be used to filter out events from backup and/or archival services running under specific user accounts.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This can be used to filter out events from backup and/or archival services based on their IP addresses.
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 can be used to filter out events from backup and/or archival services based on the process names.
Access Event Suppression Time	Enter the number of seconds to delay an open event before processing. This reduces chatter from Windows clients when mousing over files in Windows Explorer. The default value is -1, which uses a globally set value. A value of 0 allows dynamic changes to the delay time based on the system load, based on the load of the system.
Enable Local Access Events	Select to enable tracking of open and close events that occur locally on the Agent server.
Enable Remote IP Logging	Select to 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 is triggered when the file is closed.
Close Modify Extension Override	Enter a comma-separated list of exclusions for the Enable Close Modify option. Modify/write events will be detected for these files. This is important for those who rely on sync-on-save functionality.
Follow Junction Points	Select to enable junction point support for the selected Windows file server.

For more information about junction points or symbolic links options, contact [Peer Support](#).

2. Click **OK**.

Path

The **Path** page allows you to specify the path to the volume, share, export, or folder you want to synchronize. This volume, share, export, or folder is referred to as the [watch set](#). A job can monitor only one watch set. To synchronize multiple volumes, shares, exports, or folders, create a separate job for each.

1. Browse to or enter the path to the watch set.


The screenshot shows a window titled "Add New Participant" with a blue header bar. Below the header, the "Path" tab is selected in a sidebar on the left. The main area of the window has a title "Path" and a red error message: "You must enter a valid path." Below this, there 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 window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

If you are entering the path manually rather than browsing, use the appropriate path format based on the file-sharing protocol (**SMB** or **multi-protocol**) and the type of storage device:

Protocol	Storage Type	Path Format	Notes
SMB	Windows File Server	C:\data	Local path
	NAS device	\ \server_name_or_ip\shared_folder\ file_or_directory_path	UNC path
Multi-Protocol	NAS device	host:/volume	Export path

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. 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're not creating a multi-protocol job, **skip to Step 6**.

- If you are creating a multi-protocol job, specify a value for **SMB Access Configuration** by selecting an existing configuration from the **SMB Access Configuration** drop-down or creating a new one.

SMB access enables the Linux Agent to access data on the file server using the SMB protocol. SMB access is required for all **multi-protocol** File Collaboration jobs.

To create a new SMB access configuration, see [Configuring SMB Access](#).

The screenshot shows the 'Add New Participant' dialog box with the 'Path' tab selected. A red error message at the top states: 'SMB access configuration is required for Windows SMB file locking. The root of the file server (/ or /ifs) is not supported. You must specify a export or folder path below /ifs.' The left sidebar contains links for 'Management Agent', 'Storage Platform', 'Storage Information', and 'Path'. The main area has a text field with 'MLPS2ZONE1:/ Enter Path' and a 'Browse' button. Below this is a 'Seeding Target' checkbox. A section titled 'SMB Access' contains a link 'Edit SMB Access Configuration', a dropdown for 'SMB Access Configuration', and radio buttons for 'SMB Access Share'. The 'Use Admin Share (IFS\$)' option is selected. There is also a 'Specify Share Name' option with an adjacent text field. At the bottom are buttons for '< Back', 'Next >', 'Finish', and 'Cancel'.

- Choose an option for **SMB Access Share**.

The selected share will be used by the Agent to access data via the SMB protocol. Any Windows-style permissions associated with the files will be read from and written to the file server via this share.

- Use Admin Share:** Select this option if you have an admin share that you want to use. The same admin share can often be used for multiple jobs.
 - Specify Share Name:** If you want or need a dedicated share for the job, select this option and enter the share name.
- Click **Next** if you want this participant to use [Edge Caching](#); otherwise, click **Finish** to complete the configuration of this participant.

Configuring SMB Access

If your job requires SMB access, you can either use an existing configuration or create a new one.


Recommendation: Create SMB access configurations in advance using the [Storage Platforms](#) section of [Preferences](#).

To create a new SMB access configuration:

1. Click **Edit SMB Access Configuration**.

The screenshot shows the 'Add New Participant' dialog box with the 'Path' tab selected. A red error message at the top states: 'SMB access configuration is required for Windows SMB file locking. The root of the file server (/ or /ifs) is not supported. You must specify a export or folder path below /ifs.' The 'Management Agent' section is expanded, showing 'Storage Platform', 'Storage Information', and 'Path'. The 'Path' section contains a text box with 'MLPS2ZONE1:/ Enter Path' and a 'Browse' button. Below this is a 'Seeding Target' checkbox. The 'SMB Access' section is expanded, showing a link to 'Edit SMB Access Configuration', a dropdown for 'SMB Access Configuration', and two radio buttons for 'SMB Access Share': 'Use Admin Share (IFS\$)' (selected) and 'Specify Share Name:' (with an empty text box). At the bottom are buttons for '< Back', 'Next >', 'Finish', and 'Cancel'.

2. If the dialog that appears, click **Create**.



3. In the **New SMB Access Configuration** dialog, supply the requested information:

The screenshot shows a 'New SMB Access Configuration' dialog box. It has three input fields: 'Description:' (empty), 'Domain Username: Domain\Username' (pre-filled), and 'Domain Password:' (empty with a toggle icon). At the bottom are 'OK' and 'Cancel' buttons.

Field	Description
Description	Enter a name or description to identify the configuration.

Field	Description
Domain Username	Enter the domain username in the format: Domain\Username
Domain Password	Enter the password for the domain account.

- Click **OK** to save the configuration.
- Click **Apply and Close** to apply the configuration to the job and return to the wizard.

Edge Caching

[Edge Caching](#) is a method for conserving space on storage devices by caching files until needed. It saves space by stubbing files and rehydrating them when accessed. If you don't need to conserve space on the storage device managed by the Agent, then you do not need to enable this option.

When Edge Caching is enabled for a participant you must designate the participant as either a **master** or **edge** participant:

- Master participant:** A master participant holds a complete set of files for that job. None of the files are stubbed; they are physically stored on that device.
- Edge participant:** The edge participant stores both stub files, which take up minimal space, and fully rehydrated files. It holds a subset of files from the master participant, with the remaining files being stub files. When a stubbed file is accessed, it is retrieved directly from the master participant and rehydrated, so the full file is stored locally on the edge participant.

A job can have both master and edge participants, as well as participants without either role. If you do not enable Edge Caching for a participant, it will always have a full set of files like a master participant, but it will not serve files to edge participants.


Notes:

- A participant can be a master participant for some jobs and an edge participant for others.
- 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](#).

To enable Edge Caching:

1. Select the **Enable Edge Caching** checkbox to enable Edge Caching for the participant. Otherwise, click **Finish**.



The screenshot shows a window titled "Add New Participant" with a blue header bar. Below the header, the "Edge Caching" tab is selected. The main area contains the text "Enable Edge Caching and select role." and a checkbox labeled "Enable Edge Caching" which is currently unchecked. On the left side, there is a list of roles: "Management Agent", "Storage Platform", "Storage Information", "Path", and "Edge Caching". The "Edge Caching" role is highlighted with a grey background. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

If you enable Edge Caching, the roles are displayed.

2. Choose an Edge Caching role for the participant:
 - **Master** (default): Choose this if the storage device will contain complete copies of all files for the job. Any type of storage platform can be a master participant.
 - **Edge**: Choose this to conserve space on the storage device. Only Windows File Servers can be edge participants.

The screenshot shows a window titled "Add New Participant" with a blue header bar. Inside, the "Edge Caching" tab is selected, showing the instruction "Enable Edge Caching and select role." On the left is a tree view with "Edge Caching" highlighted, and "Master Data Service" listed below it. The main area contains a checked checkbox for "Enable Edge Caching" and a "Select Edge Caching role:" section with two radio buttons: "Master" (selected) and "Edge". At the bottom are four buttons: "< Back", "Next >", "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 following fields on this page are automatically populated:

- **Protocol:** Lists the protocol that will be used to transfer file content between master participants and edge participants. HTTPS is currently the only available option because it requires only one open firewall port on each master participant and handles latency effectively over WAN connections.
- **Agent Name:** Lists the name of the Management Agent that you selected at the beginning of Step 2.

Configure the remaining fields as needed:

1. (Optional) Enter a value for **Agent Alias**. The value can be a hostname, FQDN, 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:
Agent Name:
Agent Alias:
Port:

< 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.

3. Click **Finish**.

[illegible]

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 22 : 00 : 00
☐ Define schedule

☒ Run caching scan after job start

*Temporary Storage Path: \PeerTempPath Browse

< 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**.

- 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 if 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**.

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
☐ Existing Filter

< Back Next > Finish 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 "Add New Participant" dialog box is shown with the "Pinning Filter" tab selected. A red error message at the top states "Name should not be empty." The left sidebar contains a tree view with the following items: Management Agent, Storage Platform, Storage Information, Path, Edge Caching (expanded), Volume Policy, Utilization Policy, and Pinning Filter (selected). The main area is titled "Edit Pinning Filters" and contains two radio buttons: "No Filter" and "New Filter" (selected). Below the "New Filter" radio button is a text field labeled "*Name:". Underneath is a section titled "Pinning Rules:" containing a table with two columns: "Path" and "Pinning State". To the right of the table are three buttons: "Create", "Edit", and "Delete". Below the table 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 "Create Pinning Rule" dialog box is shown. It contains the instruction "Enter a file name or path and choose pinning state." Below this are two fields: "Pattern:" with a text input field, 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.

Agent	OS	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching
Lab24Win19c	Microsoft Windows		C:\Company Files\DevOps - Tok	Yes	Windows Driver	No	Master
Lab24Win22a	Microsoft Windows		C:\Company Files\DevOps - Mui	Yes	Windows Driver	No	Master
Lab24Win22b	Microsoft Windows		C:\Company Files\DevOps - Cer	Yes	Windows Driver	No	Edge

- 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.



2. Click the **Assign** button.

The **Assign Master Participants** dialog appears.



3. Select the master participants you want to assign to the edge participant.

Assign Master Participants

Assign master participants to edge participants and set failover order.

ⓘ A minimum of two master participants per edge participant is recommended.

	Failover Order	Master Participant	Path
☒	1	Lab24Win19c	C:\Company Files\DevOps - Tokyo
☒	2	Lab24Win22a	C:\Company Files\DevOps - Munich

Move Up
Move Down

OK Cancel

4. (Optional) Change the failover order by selecting a master participant and using the **Move Up** and **Move Down** buttons.

Assign Master Participants

Assign master participants to edge participants and set failover order.

ⓘ A minimum of two master participants per edge participant is recommended.


	Failover Order	Master Participant	Path
☒	1	Lab24Win22a	C:\Company Files\DevOps - Munich
☒	2	Lab24Win19c	C:\Company Files\DevOps - Tokyo

Move Up
Move Down

OK Cancel

- Click **OK**.

The **Master Participants** column has been updated for that participant.



- Repeat Steps 1-5 for each edge participant.
- Click **Next** to proceed to the [File Metadata](#) page.

Step 4: File Metadata

This step is optional.

The **File Metadata** page enables you to configure how and when file metadata is synchronized between source and target hosts. Metadata is synchronized in real time and during scans.

Available options include specifying whether to synchronize file permissions (security metadata) and defining the metadata source (volume, share, export, or folder) to use for resolving conflicts during initial synchronization. This designated source, referred to as the [master host](#), determines which version of the metadata is retained when discrepancies occur.

For more information, see [File Metadata Synchronization](#) in [Advanced Topics](#).

The content of the File Metadata page depends on the file-sharing protocol you selected in Step 2:

- See [SMB Protocol](#) for details.
- See [Multi-Protocol](#) for details.

SMB Protocol

To configure file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

Create File Collaboration Job Wizard

File Metadata

Configure the replication of security permissions.

Participants

- File Metadata**
- Application Support
- Email Alerts
- DFS Namespace

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. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. Select the permissions to be synchronized: **Owner**, **DACL**, and **SACL**.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

4. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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 [Application Support](#) page.

Multi-Protocol

To configure file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.
2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. Select the security style to determine how file permissions are interpreted and synchronized across SMB and NFS environments:
 - **Windows-style permissions**
 - **Linux-style permissions**
 - **Mixed permissions** (available only with Dell PowerScale)
4. Select the permissions to be synchronized:

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

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**

The screenshot shows the 'Create File Collaboration Job Wizard' window, specifically the 'File Metadata' step. The window has a blue title bar and a sidebar on the left with the following options: 'Participants', 'File Metadata' (selected), 'Email Alerts', and 'DFS Namespace'. The main area contains a red error message: 'You must select a master host when enabling security or attributes synchronization during initial scan.' Below this, there are three sections: 'Synchronize File Security Information' with two checked checkboxes ('Enable synchronizing file security information in real-time' and 'Enable synchronizing file security information with master host during initial scan'), 'Synchronize Security Options' with a 'Security Style' dropdown set to 'Windows-style permissions' and two checked checkboxes ('Owner' and 'DACL: Discretionary Access Control List'), and 'Metadata Conflict Resolution' with a 'Select master host for initial scan:' dropdown. A 'Validate' button is located to the right of the 'Synchronize Security Options' section. At the bottom of the window are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

- **Linux-style permissions**

Create File Collaboration Job Wizard

File Metadata

You must select a master host when enabling security or attributes synchronization during initial scan.

Participants

- File Metadata
- Email Alerts
- DFS Namespace

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)

Create File Collaboration Job Wizard

File Metadata

✖ You must select a master host when enabling security or attributes synchronization during initial scan.

Participants

- File Metadata**
- Email Alerts
- DFS Namespace

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: Mixed permissions ⚠ [How do permissions work with Multi-Protocol?](#)

Linux

- ☒ Owner
- ☒ Group
- ☒ Linux Permissions
- ☐ ACL (NFSv4 or POSIX)

Windows

- ☒ Owner
- ☒ DACL: Discretionary Access Control List

Validate

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

5. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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 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 applied automatically to all watch sets for all participants in the job.

The **Application Support** page appears only when the **SMB** protocol is selected. This page provides options specific to SMB-based environments, such as support for Microsoft applications that maintain file locks or open handles. If you select **NFS** or **Multi-Protocol**, the page does not appear because these protocols do not use SMB-specific application support features.

The **Application Support** page lists file types for which Peer Software has established best practices. These best practices may include filtering recommendations, prioritization of specific file types, and configuration of file-locking behavior. However, if an application is not listed, this does not necessarily mean it is unsupported. For details about how a specific application is optimized, contact [Peer Support](#).

To configure application support options for a file collaboration job:

1. Select the applications that have files in the job's watch set.

Create File Collaboration Job Wizard

Application Support

Select the applications that will be used with the data that is managed by this job.

Participants
Master-Edge Assignment
File Metadata
Application Support
Email Alerts
DFS Namespace

Select below to optimize this job for any of the following file types:

Adobe Products

- ☐ Adobe Illustrator
- ☐ Adobe Photoshop
- ☐ Adobe InDesign

Autodesk Products

- ☐ Autodesk AutoCAD
- ☐ Autodesk Revit
- ☐ Autodesk Civil 3D
- ☐ Autodesk Sheet Set Manager (for AutoCAD or Civil 3D)
- ☐ Autodesk Inventor

Other

- ☐ ArcGIS
- ☐ Microsystems Allegro
- ☐ Dassault Systems CATIA
- ☐ Newforma Project Center
- ☐ Microsoft Office
- ☐ Rhinoceros Rhino3D

< Back Next > Finish Cancel

2. Click **Next** to proceed to the [Email Alerts](#) page.

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.

Create File Collaboration Job Wizard

Email Alerts
Select email alerts.

[Edit Email Alerts](#)

Name	Temporarily Disabled	Event Types	Recipients
Participant Failure	No	Participant Failure	AriF@peersoftware.com

[Select](#) [Delete](#) [View Details](#)

[< Back](#) [Next >](#) [Finish](#) [Cancel](#)

4. (Optional) Repeat steps 1-3 to apply additional alerts.
5. Click **Next** to continue to the [DFS Namespaces](#) page.

Step 7: DFS Namespace

This step is optional.

The **DFS Namespace** page presents three options for linking a DFS namespace folder to this File Collaboration job.

To link a namespace to this job:

1. Click the **Enable linking job to DFS Namespace** checkbox.

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

☐ Enable linking job to DFS namespace

☒ New DFS Namespace
☒ Import DFS Namespace
☒ Existing DFS Namespace

< Back Next > Finish Cancel

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

(Optional) Link a DFS namespace folder to this File Collaboration job.

DFS Namespace

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

Company Files

Selected DFS Namespace: Company Files

Folders:

- DevOps

Targets:

- \\Lab24Win22a\\Company Files\\DevOps - Munich
- \\Lab24Win19c\\Company Files\\DevOps - Tokyo
- \\Lab24Win22b\\Company Files\\DevOps - Centreville

- Project Management

Targets:

- \\lab24win22b\\Company Files\\Project_Management - Centreville
- \\lab24win19c\\Company Files\\Project_Management - Tokyo
- \\lab24win22a\\Company Files\\Project_Management - Munich

- Security

Targets:

- \\Lab24Win22a\\Company Files\\Security - Munich
- \\Lab24Win19c\\Company Files\\Security - Tokyo

< 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](#).

Create File Collaboration Job Wizard

DFS-N Link
Link each participant's watch set to a folder target.

Participants
Master-Edge Assignment
File Metadata
Application Support
Email Alerts
DFS Namespace
DFS-N Link

Namespace:
Folder:

Linked Participants and Folder Targets:

Agent	Directory	Folder Target	Link Enabled
Lab24Win19c	C:\Company Files\DevOps - Tok		
Lab24Win22b	C:\Company Files\DevOps - Cer		
Lab24Win22a	C:\Company Files\DevOps - Mur		

< Back Next > Finish Cancel

- Continue to [Step 8: Save Job](#).


Step 8: Save Job

You are now ready to save the job configuration.

- If you are satisfied with your job configuration, click **Finish** to save your job. Otherwise, click the **Back** button and make any necessary changes.

Congratulations! You have created a File Collaboration job. A summary of the job configuration is displayed in the runtime view of the job.

See [Running and Managing a File Collaboration job](#) for more information.



Editing a File Collaboration Job

You can edit a File Collaboration job while it is running; however, any changes will not take effect until the job is restarted.

Overview

When you create a File Collaboration job, the **Create Job** wizard guides you through the process, presenting the most common options for configuration. When editing a job, you have access to all options, allowing you to fine-tune the job configuration. Options not included in the initial job creation include:

- [Delta Replication](#)
- [DFS-N Link](#)
- [File and Folder Filters](#)
- [File Metadata](#) - Some options are available only when editing the job.
- [File Locking](#)
- [General](#)
- [Scheduled Replication Filters](#)

- [SNMP Notifications](#)
- [Target Protection](#)
- [Tags](#)

You can edit multiple File Collaboration jobs simultaneously. For information about simultaneously editing multiple jobs, see [Editing Multiple Jobs](#).

Editing a Job

To edit a File Collaboration job:

1. Select the job in the **Jobs** view.
2. Right-click and select **Edit Job**.

The **Edit File Collaboration** dialog appears.

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role
Lab24Win19b		C:\ADS-3	Yes	Windows Driver	No	None
Lab24Win19c		C:\ADS-2	Yes	Windows Driver	No	None

3. Select a configuration item in the navigation tree and make the desired changes:

- [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](#)

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.](#)



Adding and Deleting 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.

The screenshot shows the 'Edit File Collaboration Job' dialog box with the 'Participants' tab selected. The left sidebar lists various settings categories, with 'Participants' highlighted. The main area displays a table of participants. The table has columns for Host, Computer Description, Directory, Enabled, Storage Platform, Seeding Target, and Edge Caching Role. Two participants are listed: Lab24Win19b and Lab24Win19c. The 'Add' button is visible on the right side of the table.

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role
Lab24Win19b		C:\ADS-3	Yes	Windows Driver	No	None
Lab24Win19c		C:\ADS-2	Yes	Windows Driver	No	None

2. Click the **Add** button.

The **Add New Participant** wizard opens; the **Management Agent** page lists the Agents available to be added.

Tip: If the Agent you want is not shown in the **Agents** table, 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](#), the **Agents** table will be updated.

Add New Participant

Management Agent

Select the server hosting the Peer Agent that manages this storage device.

Management Agent

- Storage Platform
- Storage Information
- Path
- Edge Caching

Only Microsoft Windows Agents are available based on your previously selected Agents.

Agent	OS	Computer Description
Lab24Win22a	Microsoft Windows	
Lab24Win22b	Microsoft Windows	

< Back Next > Finish Cancel

3. Select a Management Agent, and then click **Next**.

The **Storage Platform** page appears.

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. Inside the window, the "Storage Platform" tab is selected in a left-hand sidebar. The sidebar also includes "Management Agent", "Storage Information", "Path", and "Edge Caching". The main area of the window displays the instruction "Select the type of storage platform." followed by a list of seven 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 >", "Finish", and "Cancel".

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 PowerStore](#)

[Dell EMC Unity](#)

[Nutanix Files](#)


6. Click **Next**.

The **Path** page appears.

The screenshot shows a window titled "Add New Participant" with a blue header bar. Below the header, the word "Path" is displayed. A red error icon and message state: "You must enter a valid path." On the left, a sidebar lists options: "Management Agent", "Storage Platform", "Storage Information", "Path" (which is highlighted with a dark background), and "Edge Caching". 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, 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.



Add New Participant

Edge Caching
Enable Edge Caching and select role.

Management Agent
Storage Platform
Storage Information
Path
Edge Caching

☐ Enable Edge Caching


< Back Next > Finish 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

A File Collaboration job requires at least two participants. If you delete a participant and only one remains, you must add another participant to the job.

To delete a participant from a File Collaboration job:

1. In the **Edit File Collaboration** dialog, select the participant in the **Participants** table you want to remove from the job.



2. Click the **Delete** button.
3. Click **OK** in the **Delete Confirmation** dialog.

The participant is removed from the **Participants** table.

Editing a Participant

To edit a participant:

1. In the **Edit File Collaboration** dialog, select the participant you want to edit from the **Participants** table.

Edit File Collaboration Job

Participants

Protocol:

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role	
Lab24Win19b		C:\ADS-3	Yes	Windows Driver	No	Master	
Lab24Win19c		C:\ADS-2	Yes	Windows Driver	No	Edge	
Lab24Win22a		C:\ADS-1	Yes	Windows Driver	No	None	

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 'Edit Participant' and standard window controls. Below the title bar, the 'General' tab is highlighted, and the text 'Modify these settings.' is displayed. The main area is divided into two panes. The left pane contains a tree view with 'General' and 'Edge Caching' items. The right pane contains the following settings: 'Enabled' with a checked checkbox, 'Host' with a text field containing 'Lab24Win22a', 'Event Detector' with a dropdown menu showing 'Windows Driver' and a link 'Edit Detector Configuration', 'Path' with a text field containing 'C:\ADS-1' 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'.

Edit Participant

General
Modify these settings.

General
Edge Caching

Enabled: ☒

Host: Lab24Win22a

Event Detector: Windows Driver [Edit Detector Configuration](#)

Path: C:\ADS-1

☐ Seeding Target

< Back Next > Finish Cancel

3. To enable or disable the Agent, select or deselect the **Enabled** checkbox.
4. 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 you are entering the path manually rather than browsing, use the appropriate path format based on the file-sharing protocol (SMB, NFS, or multi-protocol) and the type of storage device:

Protocol	Storage Type	Path Format	Notes
SMB	Windows File Server	C:\data	Local path
	NAS device	\server_name_or_ip\shared_folder\file_or_directory_path	UNC path
NFS	Linux File Server	/mnt/data	Local path (case sensitive)
	NAS device	host:/volume	Export path
Multi-Protocol	NAS device	host:/volume	Export path

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 enable or modify Edge Caching options; otherwise, click **Finish**, and skip to Step 10.

If you clicked **Next**, the Edge Caching page appears.

The screenshot shows a window titled "Edit Participant" with a blue header bar. Below the header, the "Edge Caching" tab is selected, displaying the text "Enable Edge Caching and select role." The main area is divided into two panes: "General" on the left and "Edge Caching" on the right. The "Edge Caching" pane contains a single checkbox labeled "Enable Edge Caching", which is currently unchecked. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

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 **Finish**.
11. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

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.

If you have only one edge participant and one master participant, the master participant is automatically assigned to the edge participant and listed in the Master Participants column. If you have multiple master participants, you need to explicitly assign master participants to each edge participant and then set the failover order.

1. Select an edge participant in the **Assignment** table.



2. Click the **Assign** button.

The **Assign Master Participants** dialog appears.

3. Select the master participants you want to assign to the edge participant.
4. (Optional) Change the failover order by selecting a master participant and using the **Move Up** and **Move Down** buttons.
5. Click **OK**.

The **Master Participants** column has been updated for that participant.

6. Repeat Steps 1-5 for each edge participant.
7. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

General

The **General** page in the **Edit File Collaboration Job** dialog presents miscellaneous settings pertaining to a File Collaboration job. You may want to consult with Peer Software Support before modifying these values.

To modify these settings:

1. Enter the values recommended by Peer Support.

The screenshot shows the 'Edit File Collaboration Job' dialog box with the 'General' tab selected. The left sidebar lists various configuration categories, with 'General' highlighted. The main area contains the following settings:

- Job ID: 507
- Job Type: File Collaboration
- Job Name: FC-4-SMB-Windows Driver & Nutanix-Edge Caching
- Job Description: (Empty text area)
- Transfer Block Size (KB): 2048
- Verify Block Checksums: ☒
- Verify Full File Checksums: ☒
- Enable Multipart Transfers: ☐
- Resume File Transfers: ☐
- Resume Transfer Minimum File Size (MB): 256
- Synchronization Priority: 2
- Timeout (Seconds): 180
- First Scan Mode: FOLDER_BY_FOLDER
- Symbolic Link Strategy: Follow
- Remove Filtered Files On Folder Delete: ☒
- Require All Hosts At Start: ☐
- Auto Start: ☒
- Auto Start Priority: 1

At the bottom right are 'OK' and 'Cancel' buttons.

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.
Job Description	Enter details or notes that help identify the purpose of the job. The description appears on the job's Configuration tab.
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 .

Option	Description
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 after 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.
Resume File Transfers	If selected, interrupted transfers of files that meet the minimum file size attempt to resume after a disconnection or job stoppage. If a transfer cannot be resumed, the entire file is transferred again. Requires Verify Full File Checksums to be selected. Resume file transfer is not used for delta-encoded transfers.
Resume Transfer Minimum Size (MB)	The minimum file size, in megabytes, for which interrupted file transfers will attempt to resume.
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
Symbolic Link Strategy	Specifies how the job handles symbolic links. Select Copy to copy the symbolic link itself to destination hosts. Select Follow to follow the symbolic link and copy the target file or folder contents. The default strategy and the options available depend on the job protocol and participant configuration. For more information, including what happens when you change the strategy, see Handling Symbolic Links .
Remove Filtered Files On Folder Delete	If selected, then all child files on target hosts will be deleted when its parent folder is deleted on another source host. Otherwise, filtered files will be left intact on targets when a parent folder is deleted on another source host.
Require All Hosts At Start	If selected, requires all participating hosts to be online and available at the start of the File Collaboration job in order for the job to successfully start.
Auto Start	If selected, then this file collaboration session will automatically be started when the Peer Management Center Service is started.
Auto Start Priority	Specifies the order in which the job starts relative to other jobs configured to start automatically. Jobs with higher priority values start before jobs with lower priority values. The default is 1 . Available only when Auto Start is selected.


- 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.

- 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.



2. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

Conflict Resolution

By default, any file conflicts that are encountered during the [initial synchronization process](#) are automatically resolved by Peer Management Center. Peer Management Center resolves the conflict by selecting the file with the most recent modification time. Conflicts that cannot automatically be resolved result in the files being quarantined. The **Conflict Resolution** page in the **Edit File Collaboration Job** allows you to select options for resolving file conflicts and quarantines.


However, if you want to resolve the conflicts yourself, you can contact Peer Software to enable manual resolution. With manual resolution, you can select the host with the correct version of the file.

For more information about the cause of file conflicts, see [Conflicts, Retries, and Quarantines](#).

To modify conflict resolution settings for the File Collaboration job:

1. In the **File Conflict Resolution** section, select the **Truncate Milliseconds** option if you want the millisecond value truncated from each time stamp when comparing the time stamps of a file on two or more hosts.

As some storage platforms and applications track milliseconds slightly differently, selecting this option will prevent subtle millisecond differences from causing an otherwise in-sync file to be replicated or quarantined.



2. Select the **Advanced File Conflict Resolution** options you want applied:

Option	Description
Quarantine Offline Version Conflicts	Select this option if you want Peer Management Center to quarantine a file that was updated in two or more locations while the collaboration session was not running. If it is not selected, the file with the most recent last modified time will be replicated to all other participants.
Enable Deletion of Quarantined Files	Select this option if you want Peer Management Center to process a delete event for a quarantined file. If it is not selected, the quarantined file is not deleted and remains quarantined.
Offline Delete Detection	Select this option (and enable target protection), if you want any files or folders that were deleted while the job was stopped to be

Option	Description
During Scan	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
mp3

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.

Copyright (c) 1993-2026 Peer Software, Inc. All Rights Reserved.

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 enables you to configure how and when file metadata is synchronized between source and target hosts. Metadata is synchronized in real time and during scans.

Available options include specifying whether to replicate file security metadata, selecting which metadata types to include, and defining the metadata source (volume, share, export, or folder) to use for resolving conflicts during initial synchronization. This designated source, referred to as the [master host](#), determines which version of the metadata is retained when discrepancies occur.

For more information, see [File Metadata Synchronization](#) in [Advanced Topics](#).

This page also provides additional options not available during the job creation process.

The available options depend on the file-sharing protocol selected for the job:

- See [SMB Protocol](#) for details.
- See [Multi-Protocol](#) for details.

SMB Protocol

To modify file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

The screenshot shows the 'Edit File Collaboration Job' dialog box with the 'File Metadata' tab selected. The sidebar on the left lists various settings categories, with 'File Metadata' highlighted. The main area contains the following sections:

- 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: [Dropdown menu]
 - ☐ Enable enhanced metadata conflict resolution
- File Reparse Point Synchronization:**
 - Reparse Tag Name (numerical value only): [Text input]
 - Reparse Master Host: [Dropdown menu]
- Alternate Data Streams Transfer:**
 - ☐ Enable transfer of file Alternate Data Streams (ADS)

At the bottom right, there are 'OK' and 'Cancel' buttons.

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. (Optional) Enable prevention of corrupt or blank Owner or ACLs.

This ensures that invalid ownership or ACL entries on the source or master host are not applied to target hosts. The replication process skips or corrects these entries to maintain data integrity and security.

4. Select the permissions to be synchronized: **Owner**, **DACL**, and **SAC**.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

5. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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. (Optional) Enable enhanced metadata conflict resolution.

This prevents the Peer Agent service account from being assigned as the owner of a file or folder when a metadata conflict occurs and a file or folder is written to a target. If the Peer Agent service account is 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. The Peer Agent service account should be an actual user, as described in [Peer Global File Service - Environmental Requirements](#).

7. (Optional) In the **File Reparse Point Data Synchronization Options** section, configure reparse point handling.

Enabling these options synchronize only a file's reparse data—not the actual offline content—to target hosts and prevents the offline file from being recalled from remote storage.

- **Reparse Tag Name** - Enter a single numerical value to specify which reparse point type to synchronize.
 - o Leave blank to disable reparse synchronization.
 - o Enter **0** to enable reparse point synchronization for all reparse file types.
 - o For Symantec Enterprise Vault, enter **16**.
 - o If you are unsure which value to use, contact Peer Support. You may also use **0** if you are certain only one vendor's reparse point functionality is in use.
- **Reparse Master Host** - Select a master host.
 - o Leave blank to disable reparse synchronization. Files remain in their current state.

- o If selected, when last modified times and file sizes match across all hosts, but the file reparse attribute differs (for example, archived versus unarchived), the file reparse data is synchronized to match the file on the master host.
- o For Enterprise Vault, select the server where the archiving task runs.

Note: Use these options only if you are utilizing archiving or hierarchical storage solutions that rely on NTFS file reparse points to access data in a remote location (for example, Symantec's Enterprise Vault).

Note: These options are not available if the job is using Edge Caching.

8. (Optional) Enable transfer of file Alternate Data Stream (ADS).

- 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.
- ADS information is transferred only when a modification to the actual file itself is detected. ADS is not compared independently between participants.
- The updated file's ADS is applied to the corresponding files on target participants.

Note: The option is not available if the job is using Edge Caching.

9. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

Multi-Protocol

To modify file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. (Optional) Select whether to prevent corrupted or blank metadata from propagating.

This ensures that invalid ownership or ACL entries on the source or master host are not applied to target hosts. The replication process skips or corrects these entries to maintain data integrity and security.

4. Select the security style to determine how file permissions are interpreted and synchronized across SMB and NFS environments:

- **Windows-style permissions**
- **Linux-style permissions**
- **Mixed permissions** (available only with Dell PowerScale)

5. Select the 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**

The screenshot shows the 'Edit File Collaboration Job' window with the 'File Metadata' tab selected. The left sidebar lists various configuration categories, with 'File Metadata' highlighted. The main panel contains three sections: 'Synchronize File Security Information' with three checked options, 'Synchronize Security Options' with a dropdown set to 'Windows-style permissions' and two checked options ('Owner' and 'DACL: Discretionary Access Control List'), and 'Metadata Conflict Resolution' with a dropdown set to 'lab24ubuntu1'. A 'Validate' button is located below the security options, and 'OK' and 'Cancel' buttons are at the bottom right.

- **Linux-style permissions**

Edit File Collaboration Job

File Metadata

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 Style: Linux-style permissions [How do permissions work with Multi-Protocol?](#)

- ☒ Owner
- ☒ Group
- ☒ Linux Permissions
- ☐ ACL (NFSv4 or POSIX)

Validate

Metadata Conflict Resolution

Select a master host for initial scan: lab24ubuntu1

OK Cancel

- **Mixed permissions** (available only with Dell PowerScale)

Edit File Collaboration Job

File Metadata

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 Style: Mixed permissions ⚠ [How do permissions work with Multi-Protocol?](#)

Linux

- ☒ Owner
- ☒ Group
- ☒ Linux Permissions
- ☐ ACL (NFSv4 or POSIX)

Windows

- ☒ Owner
- ☒ DACL: Discretionary Access Control List

Validate

Metadata Conflict Resolution

Select a master host for initial scan: lab24ubuntu1

OK Cancel

6. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):

- 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 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:

The screenshot shows the 'Edit File Collaboration Job' dialog box with the 'File Locking' tab selected. The left sidebar lists various configuration categories, with 'File Locking' highlighted. The main area contains three sections: 'Locking Options', 'Source Snapshot Synchronization', and 'Sync. On Save'. In 'Locking Options', 'Include MS Office User Lock Information' is checked. In 'Source Snapshot Synchronization', 'Enable Source Snapshot Copy Sync.' is unchecked, 'Snapshot Copy Max File Size (MB)' is set to 512, and 'Snapshot Copy File Extensions' is set to 'mdb,accdb,zip,psd,ai,indd'. In 'Sync. On Save', 'Enable Sync. On Save' is unchecked, 'Included File Extensions' is set to 'xls,xlsx,doc,docx,dwg,odt,ods,odp,odg', and 'Synchronization Delay (Seconds)' is set to 20. 'OK' and 'Cancel' buttons are at the bottom right.

Section	Option	Value
Locking Options	Exclusive Target Lock:	☐
	Include MS Office User Lock Information:	☒
	Include AutoCAD User Lock Information:	☐
Source Snapshot Synchronization	Enable Source Snapshot Copy Sync.:	☐
	Snapshot Copy Max File Size (MB):	512
	Snapshot Copy File Extensions:	mdb,accdb,zip,psd,ai,indd
Sync. On Save	Enable Sync. On Save:	☐
	Included File Extensions:	xls,xlsx,doc,docx,dwg,odt,ods,odp,odg
	Synchronization Delay (Seconds):	20

2. 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 synchronization job for files created by certain applications. Optimization is applied automatically to all watch sets for all participants in the job.

The **Application Support** page appears only when the **SMB** protocol is selected. This page provides options specific to SMB-based environments, such as support for Microsoft applications that maintain file locks or open handles. If you select **NFS** or **Multi-Protocol**, the page does not appear because these protocols do not use SMB-specific application support features.

The **Application Support** page lists file types for which Peer Software has established best practices. These best practices may include filtering recommendations, prioritization of specific file types, and configuration of file-locking behavior. However, if an application is not listed, this does not necessarily mean it is unsupported. For details about how a specific application is optimized, contact [Peer Support](#).

To configure application support options for a file collaboration job:

1. Select the applications to be optimized.

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

Application Support

Select below to optimize this job for any of the following file types:

Adobe Products

☐ Adobe Illustrator ☐ Adobe Photoshop
☐ Adobe InDesign

Autodesk Products

☐ Autodesk AutoCAD ☐ Autodesk Revit
☐ Autodesk Civil 3D ☐ Autodesk Sheet Set Manager (for AutoCAD or Civil 3D)
☐ Autodesk Inventor

Other

☐ ArcGIS ☐ Microsystems Allegro
☐ Dassault Systems CATIA ☐ Newforma Project Center
☐ Microsoft Office ☐ Rhinoceros Rhino3D

OK Cancel

2. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

Target Protection

Target protection helps prevent accidental data loss by saving a backup copy of any file that is deleted or overwritten on a target host. When target protection is enabled for a job, the system moves a copy of the existing file to a hidden `.pc-trash_bin` folder on the target host—**before** applying changes from the source host.


The `.pc-trash_bin` folder is located at the root of the target host's watch set. Within this folder, the file is stored in the same relative directory structure as its original location in the watch set.

For instructions on how to restore a file from the trash bin, see [Restoring a File from the Trash Bin](#).

Enabling Target Protection

To enable and configure target protection for a job:

1. Configure the available fields as needed:



The screenshot shows a window titled "Edit File Collaboration Job" with a sidebar on the left containing a list of configuration categories. The "Target Protection" category is selected and highlighted. The main area of the window displays the configuration for "Target Protection".

Target Protection	
Enabled:	☒
# of Backup Files to Keep:	3
# of Days to Keep:	30
Trash Bin:	.pc-trash_bin

At the bottom right of the window are two buttons: "OK" and "Cancel".

Field	Description
Enabled	Enables target protection for this job.
# of Backup Files to Keep	Specifies the maximum number of backup copies to retain for each file. Older copies will be purged once this limit is reached.
# of Days to Keep	Sets the number of days to retain backup copies in the trash bin. A value of 0 disables automatic purging based on age.
Trash Bin	Displays the name of the trash bin folder used to store backup files. This folder is located in the root of the target host's watch set and is hidden. The folder name is fixed and cannot be changed.

2. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

Restoring a File from the Trash Bin

To restore a deleted or overwritten file from the trash bin:

1. On the **target host**, open the watch set directory and locate the hidden folder named `.pc-trash_bin` at the root level.

Note: The `.pc-trash_bin` folder is hidden by default. Access may require administrative or root privileges, and hidden file visibility must be enabled on the system.

2. Navigate through the folder structure to locate the backup copy. The folder hierarchy mirrors the structure of the original watch set.
3. Copy the desired file from the `.pc-trash_bin` folder.
4. Paste the file back into its original location within the watch set.


The restored file is then synchronized to all other participants in the job.

Email Alerts

The **Email Alerts** page in the **Edit File Collaboration Job** dialog allows you to select which email alerts to apply to a File Collaboration job. Email alerts are defined in the [Preferences](#) dialog and can then be applied to individual jobs. See [Email Alerts](#) in the **Preferences** section for information about creating an email alert for a File Collaboration job.


To apply email alerts to a File Collaboration job while editing the job:

1. Click the **Select** button.




The **Select Email Alert** dialog opens.

2. Select the email alert from the drop-down list, and then click **OK**.



The newly added email alert appears in the **Email Alerts** table.



3. Repeat to add additional alerts to the job.
4. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

SNMP Notifications


The **SNMP Notifications** page in the **Edit File Collaboration Job** dialog allows you to apply SNMP notifications to a File Collaboration job.

SNMP notifications, like email alerts and file filters, are configured at a global level in the [Preferences](#) dialog, then applied to individual jobs. For more information about SNMP Notifications, see [SNMP Notifications](#) in the **Preferences** section.

To enable or disable SNMP notifications for a File Collaboration job:

1. To enable, select an SNMP notification from the drop-down list.


To disable, select **None - Disabled**.



2. Click **OK** to close the Click **OK** to close the **Edit** wizard or select another configuration item to modify. 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 [Tagging Resources](#) in the [Advanced Topics](#) 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.



Editing Multiple Jobs

Management Center supports **multi-job editing**, which allows you to edit multiple File Collaboration jobs at the same time. For example, you can change a single configuration item, such as **Verify Full File Checksum**, for multiple configured jobs in one operation instead of changing the item individually for each job.

While this feature does cover most of the options available on a per-job basis, certain options are unavailable in multi-job edit mode, specifically ones tied to [participants](#). Configuration of participants must be performed on a per job basis.

To edit multiple jobs simultaneously:

1. Open Peer Management Center.
2. Select the jobs you want to edit in the **Jobs** view.
3. Right-click and select **Edit Jobs**.

For the most part, the original configuration dialog remains the same with a few minor differences depending on similarities between the selected File Collaboration jobs. A sample dialog is as follows:

Edit File Collaboration Jobs - Multiple Selected

General

File and Folder Filters
Scheduled Replication Filters
Delta Replication
File Metadata
File Locking
Target Protection
Email Alerts
SNMP Notifications

General

Job ID: MULTIPLE SELECTED

Job Type: MULTIPLE SELECTED

Transfer Block Size (KB): Multiple Values - Click to Edit

Verify Block Checksums: ☒

Verify Full File Checksums: ☒

Enable Multipart Transfers: ☐

Resume File Transfers: ☐

Resume Transfer Minimum File Size (MB): 256

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

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.

Job Initialization Process

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).

Initial Synchronization Process

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. The process consists of the following steps:

[Step 1: Job Type and Name](#)

[Step 2: Source Agent](#)

[Step 3: Source Platform](#)

[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](#)

Additional configuration options, such as applying [file filters](#) and specifying [delta-level replication](#), are available when [editing a File Replication 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.

A screenshot of a Windows-style dialog box titled "Create File Replication Job". The dialog has a blue title bar with a close button (X) in the top right corner. Inside the dialog, the text "Enter a unique name." is displayed above a single-line text input field. Below the input field, a message "Job name cannot be blank." is shown in a light gray background. At the bottom right of the dialog, there are two buttons: "OK" and "Cancel".

5. Click **OK** to proceed to the [Source Agent](#) page.

Step 2: Source Agent

Use the **Source Agent** page to select the file-sharing protocol and Source Agent for the job.

Follow these steps:

1. Select the file-sharing protocol that matches the workload for this job: **SMB**, **NFS**, or **Multi-Protocol**. The file-sharing protocol you select applies to the entire job and cannot be changed after the job is created.

Available protocol options vary based on your environment and license.

Note: Multi-Protocol requires specific license support. Before creating a job that uses both SMB and NFS, verify that your license includes Multi-Protocol support. For more information about the Multi-Protocol feature, see [Multi-Protocol](#) in the [Advanced Topics](#) section.

Agent Type	License Supports Multi-Protocol?	Available Protocol(s)	Default Selection
All Windows-based	N/A	SMB	SMB is preselected and required
All Linux-based	No	NFS	NFS is preselected and required
All Linux-based	Yes	NFS or Multi-Protocol	No default—you must choose
Mixture of Windows- and Linux-based	Yes	SMB, NFS, or Multi-Protocol	No default—you must choose.

Create File Replication Job Wizard

Source Agent
Select the server hosting the Peer Agent that manages the source.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Protocol:

Please select a protocol for the participants.

Agent	Computer Description	Protocol

< Back Next > Finish Cancel

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. Both Source and Destination Agents must use the same protocol

You can assign multiple Agents to the same storage device. However, each Agent must be used to access different volumes, shares, exports, or folders on that device.

Create File Replication Job Wizard

Source Agent
Select the server hosting the Peer Agent that manages the source.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Protocol: **SMB**

Only Microsoft Windows Agents are available based on job type.

Agent	OS	Computer Description
DGWin16A	Microsoft Windows	
DGWin16B	Microsoft Windows	
DGWin16C	Microsoft Windows	
DGWin16D	Microsoft Windows	

< Back Next > Finish Cancel

2. From the **Agents** table, select an Agent that can access the storage device you want to replicate. This Agent will be designated as the Source Agent.

Tip: If the Agent you want is not shown in the **Agents** table, 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](#), the **Agents** table will be updated.

3. Click **Next** to proceed to the [Storage Platform](#) page.


Step 3: Storage Platform

The **Storage Platform** page lists the available storage platforms types for your File Replication job.

1. Select the storage platform type hosting the data you want to replicate.

The types available depend on the file sharing protocol in the previous step:

- If you selected **SMB**, these are the available storage platform types:



Create File Replication Job Wizard


Storage Platform
Select the type of storage platform.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

- ☐ Windows File Server
- ☐ Linux File Server
- ☐ NetApp ONTAP
- ☐ Amazon FSx for NetApp ONTAP
- ☐ Dell PowerScale
- ☐ Dell PowerStore
- ☒ Dell Unity
- ☐ Nutanix Files

< Back Next > Finish Cancel

- If you selected **NFS**, these are the available storage platform types:



Create File Replication Job Wizard


Storage Platform
Select the type of storage platform.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

- ☐ Windows File Server
- ☐ Linux File Server
- ☒ NetApp ONTAP
- ☐ Amazon FSx for NetApp ONTAP
- ☐ Dell PowerScale
- ☐ Dell PowerStore
- ☐ Dell Unity
- ☐ Nutanix Files

< Back Next > Finish Cancel

- If you selected **Multi-Protocol (SMB and NFS)**, these are the available storage platform types:



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. If you selected **Existing Credentials**, choose a credential from the drop-down list, and then click **Next** to proceed to the [Source Path](#) page.
4. If you selected **New Credentials**, enter the required connection details. The fields vary by platform. For platform-specific field descriptions, see:

[Amazon FSx for NetApp ONTAP](#)

[Dell PowerScale](#)

[Dell PowerStore](#)

[Dell Unity](#)

[NetApp ONTAP](#)

[Nutanix Files](#)

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.

Amazon FSx for NetApp ONTAP

To connect to the Amazon FSx for NetApp ONTAP Storage Virtual Machine (SVM) hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

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 You must enter the name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Amazon FSxN environments.

< Back Next > Finish Cancel

Field	Description
SVM Name	Enter the name, FQDN (Fully Qualified Domain Name), or IP address of the 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 SVM management account . This must not be a cluster management account.
SVM Management IP	(Optional) Enter the IP address for accessing the SVM's management API. This is not required if the data LIFs (Logical Interfaces) associated with SVM Name already allow 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 is not listed, enter the address manually.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
4. Click **Next** to proceed to the [Source Path](#) page.

Dell PowerScale

To connect to the Dell PowerScale cluster hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

The available connection types depend on the operating system of the Agent, as well as the notification delivery method used by the storage platform.

- o **RabbitMQ**: Available only when the Agent is running on Windows. Choose RabbitMQ if you are using a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent.
- o **Syslog**: Available on both Windows and Linux Agents. Choose Syslog if the PowerScale device sends notifications directly to the Agent.

If you selected **Rabbit MQ**, this page appears:

The screenshot shows the 'Create File Replication Job Wizard' window, specifically the 'Storage Information' step. The window has a blue title bar and a sidebar on the left with 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' with the instruction 'Enter the information required to connect to the storage device.' Below this, there are two sections for 'Credentials'. The first section, 'New Credentials', is selected with a radio button and contains the following fields: '*SmartConnect Zone Name:', '*Cluster Management IP:', '*Cluster Username:', '*Cluster Password:', and 'Access Zone:'. Below these fields is a 'Connection Type' section with two radio buttons: 'Syslog' and 'RabbitMQ' (which is selected). To the right of the 'RabbitMQ' radio button is an 'Advanced' button. The second section, 'Existing Credentials', is unselected and contains a dropdown menu and an 'Edit' button. At the bottom of the main area, there is a 'Validate' button and a message: 'You must enter a Cluster Name.' Below this message is a link: 'Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments.' At the very bottom of the window are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

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.

< Back Next > Finish Cancel

These fields appear in both cases:

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 for accessing the REST-based API integrated into the PowerScale 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.
Access Zone (RabbitMQ) or Access Zones (Syslog)	Optional. Enter the name(s) of the access zone(s) being monitored.
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: • Choose RabbitMQ if you're use a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent. • Choose Syslog if the PowerScale device sends notifications directly to the Agent.

2. If you selected **Syslog**, enter values for the following additional 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 is not listed, enter the address manually.
Listening Port	Enter the port through which the Agent will receive TLS-based syslog events from the PowerScale cluster.
SSL Certificate Path	Enter the path to the certificate used for the TLS-based syslog connection between the Agent and the PowerScale cluster. For more information, see the Dell PowerScale Syslog Configuration Guide .
SSL Private Key Path	Enter the path to the private key for the TLS connection. For more information, see the Dell PowerScale Syslog Configuration Guide .
SSL Private Key Password	(Optional) Enter the password for the private key if it's protected. For more information, see the Dell PowerScale Syslog Configuration Guide .

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.

Dell PowerStore

To connect to the Dell PowerStore NAS server hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

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

*NAS Server Name:

*Management IP:

*Management Username:

*Management Password:

Connection Type: ☒ CEE

CEE Configuration

*Event Pool Name:

*Event Publisher Name:

*Agent IP Address:

*Listening Port:

☐ Existing Credentials

You must enter a NAS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerStore environments.

Field	Description
NAS Server Name	Enter the name of the PowerStore server hosting the data to be replicated.
Management IP	Enter the IP address for accessing the REST-based API integrated into the PowerStore server.
Management Username	Enter the user name for the account managing the PowerStore server.
Management Password	Enter the password for account managing the PowerStore server.
Connection Type	Select the method for sending real-time event notifications to the Agent: • CEE: Choose if you are using the CEE framework to dispatch notifications to the Agent.

2. Enter values for the required fields in the **CEE Configuration** section.

Field Description	Description
Event Pool Name	Enter the exact name of the event pool on the PowerStore device that publishes CEE notifications. You can find this name in the PowerStore Manager UI under Event Notifications.
Event Publisher Name	Enter the exact name or ID of the event publisher on the PowerStore device that sends CEE notifications to the Agent. You can find this value in the PowerStore Manager UI under Event Notifications.

Field Description	Description
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the PowerStore device. The device must be able to route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

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 Unity

To connect to the Dell Unity NAS Server hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

The required information will vary based on the selected **connection type** (**RabbitMQ** or **CEE**), as these protocols use different communication mechanisms:

- o **RabbitMQ**: Choose **RabbitMQ** if you're use a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent.
- o **CEE**: Choose **CEE** if the Unity device sends notifications directly to the Agent.

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

*NAS Server Name:

*Unisphere Management IP:

*Unisphere Username:

*Unisphere Password:

Connection Type: ☒ RabbitMQ ☐ CEE

☐ Existing Credentials

MLUNITY, user:admin

You must enter a NAS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments.

If you selected **CEE**, 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

*NAS Server Name:

*Unisphere Management IP:

*Unisphere Username:

*Unisphere Password:

Connection Type: ☐ RabbitMQ ☒ CEE

CEE Configuration

*Event Pool Name:

*Agent IP Address:

*Listening Port:

☐ Existing Credentials

MLUNITY, user:admin

You must enter a NAS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments.

These fields appear in both cases:

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 managing the Unity storage devices. This should not be the IP address of the NAS server itself.
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.

2. If you selected **CEE**, enter values for the following additional fields in the **CEE Configuration** section.

Field Description	Description
Event Pool Name	Enter the exact name of the event pool on the Unity device that publishes CEE notifications. You can find this name in the Unity Manager UI under Event Notifications.
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the Unity device. The device must be able to route traffic to this address. If the desired IP address is not listed, enter the address manually.


Field Description	Description
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

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.

Linux File Server

To connect to the Linux File Server hosting the data to be replicated:

1. Select the **Detector Type** for the Linux File Server. The **Detector type** refers to the type of storage platform being monitored for real-time file activity.




The screenshot shows the 'Create File Replication Job Wizard' window, specifically the 'Storage Information' step. The window has a blue title bar and a sidebar on the left with a list of steps: Source Agent, Storage Platform, Storage Information (highlighted), Source Path, Destination Agent, Destination Path, File Metadata, and Email Alerts. The main area contains the instruction 'Enter the information required to connect to the storage device.' and a 'Detector Type' dropdown menu set to 'Linux Connector'. An 'Advanced' button is located to the right of the dropdown. At the bottom, there are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Next** to proceed to the [Source Path](#) page.

Linux File Server Advanced Options

1. Modify the Linux Detector options as desired.



The screenshot shows the 'Linux Detector Options' dialog box. It has a blue title bar and contains four input fields: 'Filter open/close events from these users:', 'Filter all events from these users:', 'Filter events from these IP Addresses:', and 'Access Event Suppression Time:'. The first three fields are empty text boxes, and the last one is a spinner box set to '-1'. At the bottom, there are 'OK' and 'Cancel' buttons.

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 can be used to filter out events from backup and/or archival services running under specific user accounts.
Filter all events from these users	Enter a comma-separated list of user account names from which all file activities will be ignored. This can be used to filter out events from backup and/or archival services running under specific user accounts.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This can be used to filter out events from backup and/or archival services based on their IP addresses.
Access Event Suppression Time	Enter the number of seconds to delay an open event before processing. This reduces chatter from Windows clients when mousing over files in Windows Explorer. The default value is -1, which uses a globally set value. A value of 0 allows dynamic changes to the delay time based on the system load.

2. Click **OK**.

NetApp ONTAP

To connect to the NetApp ONTAP Storage Virtual Machine (SVM) hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

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
 PTSVM93A, user:vsadmin Edit

Validate You must enter the name.
 Having trouble connecting? Please verify that all [prerequisites](#) are met for NetApp ONTAP environments.

< Back Next > Finish Cancel

Field	Description
SVM Name	Enter the name, FQDN (Fully Qualified Domain Name), or IP address of the 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 SVM Name already allow 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 is not listed, enter the address manually.

- Click **Advanced** if you want to set [advanced options](#).

3. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
4. Click **Next** to proceed to the [Source Path](#) page.

Nutanix Files

To connect to the Nutanix Files cluster hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

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:

Advanced

☐ Existing Credentials

AFSS, user:peeradmin

Nutanix File Server Name cannot be empty.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Nutanix Files environments.

< Back Next > Finish Cancel

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 is not listed, enter the address manually. The IP address should not point to the Files cluster itself.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
4. Click **Next** to proceed to the [Source Path](#) page.

Windows File Server

To connect to the Windows File Server hosting the data to be replicated:

1. Select the detector type for the Windows File Server. **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 Support.

The screenshot shows a window titled "Create File Replication Job Wizard" with a blue header bar. Below the header, the "Storage Information" step is selected in a left-hand navigation pane. The main area of the wizard contains the instruction "Enter the information required to connect to the storage device." and a "Detector Type" dropdown menu set to "Windows Driver". An "Advanced" button is located to the right of the dropdown. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Next** to proceed to the [Source Path](#) page.


Windows File Server Advanced Options

When configuring a job that uses a **Windows File Server**, you can adjust advanced options to control how the system filters, logs, and follows file activity.

The available options depend on the selected driver. Some options are common to both the **Windows driver** and **Windows Filter driver**, while others are unique to the **Windows Filter driver**.

1. Modify the Windows Detector options as desired.

If you selected **Windows**, this page appears:



The screenshot shows the "Windows Detector Options" dialog box. It has a blue title bar with the text "Windows Detector Options" and standard window controls. The main area contains several settings: a text box for "Filter open/close events from these users:", a spinner box for "Access Event Suppression Time" set to "-1", and a section titled "Reparse Point Options" containing three checked checkboxes: "Follow Junction Points", "Follow Mount Points", and "Follow Symbolic Links". At the bottom right are "OK" and "Cancel" buttons.

Windows Detector Options

Filter open/close events from these users:


Access Event Suppression Time:

Reparse Point Options

- ☒ Follow Junction Points
- ☒ Follow Mount Points
- ☒ Follow Symbolic Links

OK Cancel

If you selected **Windows Driver**, this page appears:



The screenshot shows the "Windows Driver Detector Options" dialog box. It has a blue title bar with the text "Windows Driver Detector Options" and standard window controls. The main area contains several settings: four text boxes for "Filter open/close events from these users:", "Filter all events from these users:", "Filter events from these IP Addresses:", and "Filter events from these local processes:", a spinner box for "Access Event Suppression Time" set to "-1", a checkbox for "Enable Local Access Events" which is unchecked, a checkbox for "Enable Remote IP Address Logging" which is checked, a checkbox for "Enable Close Modify" which is checked, and a text box for "Close Modify Extension Override:". At the bottom right are "OK" and "Cancel" buttons.

Windows Driver Detector Options

Filter open/close events from these users:

Filter all events from these users:

Filter events from these IP Addresses:

Filter events from these local processes:

Access Event Suppression Time:

Enable Local Access Events: ☐

Enable Remote IP Address Logging: ☒

Enable Close Modify: ☒

Close Modify Extension Override:

Reparse Point Options

- ☒ Follow Junction Points
- ☒ Follow Mount Points
- ☒ Follow Symbolic Links

OK Cancel

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 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 can be used to filter out events from backup and/or archival services running under specific user accounts.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This can be used to filter out events from backup and/or archival services based on their IP addresses.
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 can be used to filter out events from backup and/or archival services based on the process names.
Access Event Suppression Time	Enter the number of seconds to delay an open event before processing. This reduces chatter from Windows clients when mousing over files in Windows Explorer. The default value is -1, which uses a globally set value. A value of 0 allows dynamic changes to the delay time based on the system load, based on the load of the system.
Enable Local Access Events	Select to enable tracking of open and close events that occur locally on the Agent server.
Enable Remote IP Logging	Select to 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 is triggered when the file is closed.
Close Modify Extension Override	Enter a comma-separated list of exclusions for the Enable Close Modify option. Modify/write events will be detected for these files. This is important for those who rely on sync-on-save functionality.
Follow Junction Points	Select to enable junction point support for the selected Windows file server.

For more information about junction points or symbolic links options, contact [Peer Support](#).

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 volume, share, export, or folder is referred to as the [watch set](#). A job can monitor only one watch set. To replicate multiple volumes, shares, exports, or folders, create a separate job for each.

1. Browse to or enter the path 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 Next > Finish Cancel

If you are entering the path manually rather than browsing, use the appropriate path format based on the file-sharing protocol (**SMB**, **NFS**, or **multi-protocol**) and the type of storage device:

Protocol	Storage Type	Path Format	Notes
SMB	Windows File Server	C:\data	Local path
	NAS device	\ \server_name_or_ip\shared_folder\ file_or_directory_path	UNC path
NFS	Linux File Server	/mnt/data	Local path (case sensitive)
	NAS device	host:/volume	Export path
Multi-Protocol	NAS device	host:/volume	Export path

2. If SMB access is not required, **skip to Step 5**.

SMB access enables the Linux Agent to access data on the file server using the SMB protocol. SMB access is **required** when creating a SMB access is File Replication job in the following scenarios:

- Synchronizing Windows-style permissions
 - Using a Nutanix Files storage device
3. If you are creating a multi-protocol job and SMB access is required, specify a value for **SMB Access Configuration** by selecting an existing configuration from the **SMB Access Configuration** drop-down or creating a new one.

To create a new SMB access configuration, see [Configuring SMB Access](#).

3. Choose an option for **SMB Access Share**. The selected share will be used by the Agent to access data via the SMB protocol. Any Windows-style permissions associated with the files will be read from and written to the file server via this share.
 - **Use Admin Share:** Select this option if you have an existing admin share you want to use. The same admin share can often be used for multiple jobs.
 - **Specify Share Name:** Select this option if you prefer or require a dedicated share for the job, and enter the share name.
4. Click **Next** to proceed to the [Destination Agent](#) page.

Configuring SMB Access

If your job requires SMB access, you can either use an existing configuration or create a new one.

Recommendation: Create SMB access configurations in advance using the [Storage Platforms](#) section of [Preferences](#).

To create a new SMB access configuration:

1. Click **Edit SMB Access Configuration**.

Create File Replication Job Wizard

Source Path

✖ The root of the file server (/ or /ifs) is not supported. You must specify a export or folder path below /ifs.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

MLPS2ZONE2:/ Enter Path Browse

— SMB Access —

⚠ **Required only when synchronizing Windows-style permissions.**

[Edit SMB Access Configuration](#)

SMB Access Configuration: DGUbuntu1, Username: peertest\labadmin


SMB Access Share:

☒ Use Admin Share (IFS\$)

☐ Specify Share Name:

< Back Next > Finish Cancel

2. In the dialog that appears, click **Create**.



3. In the **New SMB Access Configuration** dialog, supply the requested information:

The 'New SMB Access Configuration' dialog box contains the following fields and buttons:

- Description:** An empty text input field.
- Domain Username:** A text input field containing the text 'Domain\Username'.
- Domain Password:** A text input field with a toggle icon (an eye) to the right.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom.

Field	Description
Description	Enter a name or description to identify the configuration.

Field	Description
Domain Username	Enter the domain username in the format: Domain\Username
Domain Password	Enter the password for the domain account.

- Click **OK** to save the configuration.
- Click **Apply and Close** to apply the configuration to the job and return to the wizard.

Step 6: Destination Agent

The **Destination Agent** page lists available Agents, excluding the designated Source Agent. This Destination Agent handles the transfer of files and metadata to the destination storage device. No credentials are required since it does not perform real-time monitoring.

- Select the Agent that manages the destination storage device. If the destination is a Windows File Server, the Agent must be installed on it.

Tip: If the Agent you want is not shown in the **Agents** table, 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](#), the **Agents** table will be updated.

Create File Replication Job Wizard

Destination Agent
Select the server hosting the Peer Agent that manages the destination.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Only Microsoft Windows Agents are available based on your previously selected Agents.

Agent	OS	Computer Description
DGWin16A	Microsoft Windows	
DGWin16B	Microsoft Windows	
DGWin16D	Microsoft Windows	

< Back Next > Finish Cancel

2. Click **Next** to proceed to the [Destination Path](#) page.

Step 7: Destination Path

Use the **Destination Path** page to specify the volume, share, export, or folder to which data will be replicated.

1. Browse to or enter the path to the destination path:

Create File Replication Job Wizard

Destination Path

You must enter a valid path.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

Path:

< Back Next > Finish Cancel

If you are entering the path manually rather than browsing, use the appropriate path format based on the file-sharing protocol (**SMB**, **NFS**, or **multi-protocol**) and the type of storage device:

Protocol	Storage Type	Path Format	Notes
SMB	Windows File Server	C:\data	Local path
	NAS device	\server_name_or_ip\shared_folder\file_or_directory_path	UNC path
NFS	Linux File Server	/mnt/data	Local path (case sensitive)
	NAS device	host:/volume	Export path
Multi-Protocol	NAS device	host:/volume	Export path

Note:

- If the path field is empty, clicking **Browse** shows local drives and folders on the Agent server.
- If you enter the start of a UNC path and click **Browse**, the **Folder Browser** dialog attempts to display available shares on the specified file server.
- If you enter the start of an NFS path and click **Browse**, the **Folder Browser** dialog attempts to display available exports on the specified NFS server. You may need to enter the full NFS path manually.

2. If SMB is not required, skip to **Step 5**.
3. Specify the SMB access configuration (required in some cases).

SMB access enables the Linux Agent to access data on the file server using the SMB protocol. It is **required** when creating a multi-protocol File Replication job in the following scenario:

- Synchronizing Windows-style permissions

To create a new one, see [Configuring SMB Access](#).

Create File Replication Job Wizard

Destination Path

Please enter a valid Unix path.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

NFS Path: Enter the path in NFS format host:/path Browse

SMB Access

[Edit SMB Access Configuration](#)

SMB Access Configuration: DGUbuntu1, Username: peertest\labadmin

SMB Access Share:

☐ Use Admin Share

☒ Specify Share Name:

SMB Path: configurable\path

< Back Next > Finish Cancel

4. Enter a share name in **Specify Share Name**.

The selected share will be used by the Agent to access data using the SMB protocol. Any Windows-style permissions associated with the files will be read from and written to the file server via this share.

5. Enter a value for **SMB Path**:

- Provide a UNC path that points to the same data on the destination file server as the NFS Path, but accessed using the SMB protocol.
- The server name is automatically populated based on your destination NFS path.
- The share name is automatically populated based on your **SMB Access Share**.
- Leave this blank if replicating to the root of the share.

Depending on your file server's configuration, you may need to edit the server or share portions of the path. See [SMB Path Details](#) for more details and an example.

6. Click **Next** to proceed to the [File Metadata](#) page

Configuring SMB Access

If your job requires SMB access, you can either use an existing configuration or create a new one.

Recommendation: Create SMB access configurations in advance using the [Storage Platforms](#) section of [Preferences](#).

To create a new SMB access configuration:

1. Click **Edit SMB Access Configuration**.

Create File Replication Job Wizard

Destination Path

Please enter a valid Unix path.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

NFS Path: Enter the path in NFS format host:/path Browse

SMB Access

[Edit SMB Access Configuration](#)

SMB Access Configuration: DGUbuntu1, Username: peertest\labadmin

SMB Access Share:


☐ Use Admin Share

☒ Specify Share Name:


SMB Path: configurable\path

< Back Next > Finish Cancel

2. In the dialog that appears, click **Create**.



3. In the **New SMB Access Configuration** dialog, supply the requested information:



Field	Description
Description	Enter a name or description to identify the configuration.

Field	Description
Domain Username	Enter the domain username in the format: Domain\:Username
Domain Password	Enter the password for the domain account.

- Click **OK** to save the configuration.
- Click **Apply and Close** to apply the configuration to the job and return to the wizard.

SMB Path Details

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 name of the server from the **NFS Path** field.
- SMB Access Share:** Automatically populated with your selected **SMB Access Share**.
- A relative path under the **SMB Access Share**: Depending on the configuration of your file server, this relative path may not be necessary.

Example of Setting the Relative Path

- 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 the 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 enables you to configure how and when file metadata is synchronized between source and target hosts. Metadata is synchronized in real time and during scans.

Available options include specifying whether to synchronize file permissions (security metadata) and defining the metadata source (volume, share, export, or folder) to use for resolving conflicts during initial synchronization. This designated source, referred to as the [master host](#), determines which version of the metadata is retained when discrepancies occur.

For more information, see [File Metadata Synchronization](#) in [Advanced Topics](#).

The content of the File Metadata page depends on the file-sharing protocol you selected in Step 2:

- See [SMB Protocol](#) for details.
- See [NFS Protocol](#) for details.
- See [Multi-Protocol](#) for details.

SMB Protocol

To configure file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata will be synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. Select the permissions to be synchronized: **Owner**, **DACL**, and **SACL**.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

4. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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.

NFS Protocol

To configure metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

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

- ☒ Owner
- ☒ Group
- ☒ Linux Permissions
- ☐ ACL (NFSv4 or POSIX)

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. Select the permissions to be synchronized: **Owner**, **Group**, **Linux Permissions**, and **ACL**.

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

4. Select the [master host](#) for resolving metadata conflicts during the initial scan:

- If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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.

Multi-Protocol

To configure file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) 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.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)
2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. Select the security style to determine how file permissions are interpreted and synchronized across SMB and NFS environments:

- **Windows-style permissions**
- **Linux-style permissions**

4. Select the permissions to be synchronized:

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

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

5. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):

- 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.

Create File Replication Job Wizard

Email Alerts

Select email alerts.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

[Edit Email Alerts](#)


Name	Temporarily Disabled	Event Types	Recipients
------	----------------------	-------------	------------

Select
Delete
View Details

< Back Next > Finish Cancel

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.

Create File Replication Job Wizard

Email Alerts
Select email alerts.

Source Agent
Storage Platform
Storage Information
Source Path
Destination Agent
Destination Path
File Metadata
Email Alerts

[Edit Email Alerts](#)

Name	Temporarily Disabled	Event Types	Recipients
Participant Failure	No	Participant Failure	AriF@peersoftware.com

Select
Delete
View Details

< Back Next > Finish Cancel


4. (Optional) Repeat steps 1-3 to apply additional alerts.
5. Continue to [Step 10: Save Job](#).

Step 10: Save Job

Now that you have completed the first nine steps of the wizard, you are ready to save the job configuration.

1. If you are satisfied with your job configuration, click **Finish** to save your job. Otherwise, click the **Back** button and make any necessary changes.

Congratulations! You have created a File Replication job. A summary of the job configuration is displayed in the runtime view of the job.



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](#)
- [Target Protection](#)

- [Tags](#)

You can edit multiple File Replication jobs simultaneously. For information about simultaneously editing multiple jobs, see [Editing Multiple Jobs](#).

Editing a Job

To edit a File Replication job:

1. Select the job in the **Jobs** view.
2. Right-click and select **Edit Job**.

The **Edit File Replication** dialog appears.

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target
DGWin16C		\\PTPSTORE3\smb1	Yes	Dell PowerStore	No
DGWin16D		C:\User Guides-Cape Town\Test	Yes	None	Yes

3. Select a configuration item in the navigation tree and make the desired changes:

- [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](#)

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.](#)

Adding and Deleting a Participant

This topic describes [adding](#) and [deleting](#) participants in a File Replication job.

Adding a Destination Participant to a File Replication Job

To add a Destination 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.

Edit File Replication Job

Participants

Protocol:

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target
DGWin16C		\\PTPSTORE3\smb1	Yes	Dell PowerStore	No
DGWin16D		C:\User Guides-Cape Town\Test	Yes	None	Yes

2. Click the **Add** button.

The **Add New Participant** wizard opens; the **Destination Agent** page lists the Agents available to be added.

Tip: If the Agent you want is not shown in the **Agents** table, 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](#), the **Agents** table will be updated.

The screenshot shows a window titled "Add New Participant" with a blue header bar. Below the header, the "Destination Agent" tab is selected. The instruction "Select the server hosting the Peer Agent that manages the destination." is displayed. On the left, a sidebar shows "Destination Agent" and "Destination Path". The main area contains a table with columns "Agent", "OS", and "Computer Description". Two rows are visible: "DGWin16A" and "DGWin16B", both with "Microsoft Windows" in the OS column. A message icon and text state: "Only Microsoft Windows Agents are available based on your previously selected Agents." At the bottom, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Agent	OS	Computer Description
DGWin16A	Microsoft Windows	
DGWin16B	Microsoft Windows	

3. Select a Destination Agent, and then click **Next**.

The **Destination Path** page appears.

Add New Participant

Destination Path

✖ You must enter a valid path.

Destination Agent

Destination Path

Path: Enter the path in Windows or UNC format

< Back Next > Finish Cancel

If you are entering the path manually rather than browsing, use the appropriate path format based on the file-sharing protocol (SMB, NFS, or multi-protocol) and the type of storage device:

Protocol	Storage Type	Path Format	Notes
SMB	Windows File Server	C:\data	Local path
	NAS device	\server_name_or_ip\shared_folder\file_or_directory_path	UNC path
NFS	Linux File Server	/mnt/data	Local path (case sensitive)
	NAS device	host:/volume	Export path
Multi-Protocol	NAS device	host:/volume	Export path

- Specify the SMB access configuration (required in some cases). If not required, skip to Step 5.

SMB access enables the Linux Agent to access data on the file server via the SMB file sharing protocol. It is required in the following cases when creating a multi-protocol job:

- Synchronizing **Windows-style** permissions
- Synchronizing **mixed permissions** (both Windows-style permissions and Linux-styled permission)

For instructions, see [Configuring SMB Access](#).

- If SMB access is required, enter a share name in **Specify Share Name**.

The selected share will be used by the Agent to access data via the SMB protocol. Any Windows-style permissions associated with the files will be read from and written to the file server via this share.

- If SMB access is required, enter a value for **SMB Path**: Provide a UNC path that points to the same data on the destination file server as the NFS Path, but accessed via the SMB protocol.
 - The server name is automatically populated based on your destination NFS path.

- The share name is prefilled based on your **SMB Access Share**.

Depending on your file server's configuration, you may need to edit the server or share portions of the path. See [SMB Path Details](#) for more details and an example.

7. Click **Finish**.

The new participant appears in the **Participants** table.


Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target
DGWin16B		C:\User Guides	Yes	None	Yes
DGWin16C		\\PTPSTORE3\smb1	Yes	Dell PowerStore	No
DGWin16D		C:\User Guides-Cape Town\Test	Yes	None	Yes

Deleting a Participant from a File Replication Job

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 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.

Editing a Participant


To edit a participant:

1. In the **Edit File Replication Job** dialog, select the participant you want to edit from the **Participants** table.



2. Click **Edit**.

The **Edit Participant** dialog appears.



- To change the volume, share, export, or folder that is replicated, enter the path to the new watch set in the **Path** field or browse to it.

If you are entering the path manually rather than browsing, use the appropriate path format based on the file-sharing protocol (SMB, NFS, or multi-protocol) and the type of storage device:

Proto col	Storage Type	Path Format	Notes
SMB	Windows File Server	C:\data	Local path
	NAS device	\ \server_name_or_ip\shared_folder\file_or_directory_path	UNC path
NFS	Linux File Server	/mnt/data	Local path (case sensitive)
	NAS device	host:/volume	Export path
Multi - Proto col	NAS device	host:/volume	Export path

- If you are editing a Source Agent, you can edit the Detector Configuration. If it is a Destination Agent, the option is disabled.
- Click **Finish**.
- 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 Support before modifying these values.

To modify these settings:

- Enter the values recommended by Peer Support.

Edit File Replication Job

Participants

- General**
- File and Folder Filters
- Scheduled Replication Filters
- Conflict Resolution
- Delta Replication
- File Metadata
- Job Relay
- File Locking
- Application Support
- Target Protection
- Email Alerts
- SNMP Notifications
- Tags

General

Job ID: 432

Job Type: File Replication

Job Name: FR-7-SMB-PowerStore

Job Description:

Transfer Block Size (KB): 2048

Verify Block Checksums: ☒

Verify Full File Checksums: ☒

Enable Multipart Transfers: ☐

Resume File Transfers: ☐

Resume Transfer Minimum File Size (MB): 256

Synchronization Priority: 2

Timeout (Seconds): 180

First Scan Mode: FOLDER_BY_FOLDER

Symbolic Link Strategy: Follow

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.
Job Description	Enter details or notes that help identify the purpose of the job. The description appears on the job's Configuration tab.
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	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

Option	Description
Checksums	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.
Resume File Transfers	If selected, interrupted transfers of files that meet the minimum file size attempt to resume after a disconnection or job stoppage. If a transfer cannot be resumed, the entire file is transferred again. Requires Verify Full File Checksums to be selected. Resume file transfer is not used for delta-encoded transfers.
Resume Transfer Minimum Size (MB)	The minimum file size, in megabytes, for which interrupted file transfers will attempt to resume.
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.
Symbolic Link Strategy	Specifies how the job handles symbolic links. Select Copy to copy the symbolic link itself to destination hosts. Select Follow to follow the symbolic link and copy the target file or folder contents.


Option	Description
	The default strategy and the options available depend on the job protocol and participant configuration. For more information, including what happens when you change the strategy, see Handling Symbolic Links .
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.
Delete Source File After Successful Transfer	If selected, files are removed from the source host after they have been successfully replicated and a full checksum match is confirmed. This turns the replication job into a “move” operation. The option is only available for two-host replication jobs with full-checksum verification enabled.
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 replication session will automatically be started when the Peer Management Center Service is started.
Auto Start Priority	Specifies the order in which the job starts relative to other jobs configured to start automatically. Jobs with higher priority values start before jobs with lower priority values. The default is 1. Available only when Auto Start is selected.

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 and Folder 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.

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.



2. If you want to create a new filter, modify an existing one, or update a filter, click **Edit File Scheduled Replication Filters**.

The **File and Folder Filters** dialog appears. You cannot edit predefined filters. See [Scheduled Replication Filters](#) in the [Preferences](#) section for information about creating or modifying a scheduled replication filter.



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.

- 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** wizard allows you to specify the delta-level 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:

- Modify the following the fields as necessary.

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

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	Add Delete
jpg	
jpeg	
png	
gif	
tiff	
tif	
Z	
tgz	
gz	
gzip	
rar	
7z	
bz	
bz2	
bzip2	
mn3	

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.
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.

Field	Description
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

File Metadata

The **File Metadata** page enables you to configure how and when file metadata is synchronized between source and target hosts. Metadata is synchronized in real time and during scans.

Available options include specifying whether to synchronize file permissions (security metadata) and defining the metadata source (volume, share, export, or folder) to use for resolving conflicts during initial synchronization. This designated source, referred to as the [master host](#), determines which version of the metadata is retained when discrepancies occur.

For more information, see [File Metadata Synchronization](#) in [Advanced Topics](#).

This page also provides additional options not available during the job creation process.

The content of the File Metadata page depends on the file-sharing protocol you selected for your job:

- See [SMB Protocol](#) for details.
- See [NFS Protocol](#) for details.
- See [Multi-Protocol](#) for details.

SMB Protocol

To modify file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata are synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the 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 panel contains several sections for configuring metadata synchronization:

- Synchronize File Security Information:** Three checkboxes are all checked: 'Enable synchronizing file security information in real-time', 'Enable synchronizing file security information with master host during initial scan', and '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' and 'DACL: Discretionary Access Control List' are checked, while 'SACL: System Access Control List' is unchecked.
- Metadata Conflict Resolution:** A dropdown menu is set to 'DGWin16A', and 'Enable enhanced metadata conflict resolution' is checked.
- File Reparse Point Synchronization:** Fields for 'Reparse Tag Name (numerical value only)' and 'Reparse Master Host' are empty.
- Alternate Data Streams Transfer:** 'Enable transfer of file Alternate Data Streams (ADS)' is unchecked.
- File Timestamps Transfer:** 'Use source create time' and 'Use source last accessed time' are both unchecked.

At the bottom right, there are 'OK' and 'Cancel' buttons.

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. (Optional) Enable prevention of corrupt or blank Owner or ACLs.

This ensures that invalid ownership or ACL entries on the source or master host are not applied to target hosts. The replication process skips or corrects these entries to maintain data integrity and security.

4. Select the permissions to be synchronized: **Owner**, **DACL**, and **SACL**.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

5. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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. (Optional) Enable enhanced metadata conflict resolution.

This prevents the Peer Agent service account from being assigned as the owner of a file or folder when a metadata conflict occurs and a file or folder is written to a target. If the Peer Agent service account is 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. The Peer Agent service account should be an actual user, as described in [Peer Global File Service - Environmental Requirements](#).

7. (Optional) In the **File Reparse Point Data Synchronization Options** section, configure reparse point handling.

Enabling these options synchronizes only a file's reparse data—not the actual offline content—to target hosts and prevents the offline file from being recalled from remote storage.

- **Reparse Tag Name** - Enter a single numerical value to specify which reparse point type to synchronize.
 - o Leave blank to disable reparse synchronization.
 - o Enter **0** to enable reparse point synchronization for all reparse file types.
 - o For Symantec Enterprise Vault, enter **16**.
 - o If you are unsure which value to use, contact Peer Support. You may also use **0** if you are certain only one vendor's reparse point functionality is in use.
- **Reparse Master Host** - Select a master host.
 - o Leave blank to disable reparse synchronization. Files remain in their current state.

- o If selected, when last modified times and file sizes match across all hosts, but the file reparse attribute differs (for example, archived versus unarchived), the file reparse data is synchronized to match the file on the master host.
- o For Enterprise Vault, select the server where the archiving task runs.

Note: Use these options only if you are utilizing archiving or hierarchical storage solutions that rely on NTFS file reparse points to access data in a remote location (for example, Symantec's Enterprise Vault).

Note: These options are not available if the job is using Edge Caching.

8. (Optional) Enable transfer of file Alternate Data Stream (ADS).

- 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.
- ADS information is transferred only when a modification to the actual file itself is detected. ADS is not compared independently between participants.
- The updated file's ADS is applied to the corresponding files on target participants.

Note: The option is not available if the job is using Edge Caching.

9. (Optional) In the **File Timestamps Transfer** section, enable one or both options to apply source file timestamps to the destination file:

- **Use source create time:** Applies the source file's create time to the destination file.
- **Use source last accessed time:** Applies the source file's last accessed time to the destination file.

These timestamps are applied only when the file is transferred. If timestamps on the source change but the file contents remain the same, the updated timestamps are not applied to the destination.

Note: PeerGFS does not synchronize any folder timestamps.

10. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

NFS Protocol

To modify file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:

- **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
- **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

The screenshot shows the 'Edit File Replication Job' window with the 'File Metadata' tab selected. The sidebar on the left lists various configuration options, with 'File Metadata' highlighted. The main panel contains the following sections:

- 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
 - ☒ Group
 - ☒ Linux Permissions
 - ☒ ACL (NFSv4 or POSIX)
- Metadata Conflict Resolution:**
 - Select a master host for initial scan:
- File Timestamps Transfer:**
 - ☒ Use source last accessed time

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

2. If you enable the real-time and initial scan synchronization options, click **OK** when prompted.
3. (Optional) Enable prevention of corrupt or blank Owner or ACLs.

This ensures that invalid ownership or ACL entries on the source or master host are not applied to target hosts. The replication process skips or corrects these entries to maintain data integrity and security

4. Select the permissions to be synchronized: **Owner**, **Group**, **Linux Permissions**, and **ACL**.

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

5. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - 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 Options.

- If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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. (Optional) In the **File Timestamps Transfer** section, enable **Use source last accessed time** to apply the source file's last accessed time to the destination file.

This timestamp is applied only when the file is transferred. If the timestamp on the source changes but the file contents remain the same, the updated timestamp is not applied to the destination.

Note: PeerGFS does not synchronize any folder timestamps.

7. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

Multi-Protocol

To modify file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. (Optional) Enable prevention of corrupt or blank Owner or ACLs.

This ensures that invalid ownership or ACL entries on the source or master host are not applied to target hosts. The replication process skips or corrects these entries to maintain data integrity and security.

4. Select the security style:

- **Windows-style permissions**
- **Linux-style permissions**

5. Select the permissions to be synchronized.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts

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**

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 panel contains the following sections:

- Synchronize File Security Information:** Three checkboxes are checked: 'Enable synchronizing file security information in real-time', 'Enable synchronizing file security information with master host during initial scan', and 'Enable prevention of corrupt or blank Owner or ACLs on source or master host from being applied to any target host'.
- Synchronize Security Options:** The 'Security Style' is set to 'Windows-style permissions'. A link 'How do permissions work with Multi-Protocol?' is present. Two checkboxes are checked: 'Owner' and 'DACL: Discretionary Access Control List'. A 'Validate' button is located at the bottom right of this section.
- Metadata Conflict Resolution:** A dropdown menu is set to 'lab24ubuntu1'.
- File Timestamps Transfer:** The checkbox 'Use source last accessed time' is unchecked.

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

- **Linux-style permissions**

6. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - If synchronizing metadata during the initial scan, choose a [master host](#) to resolve conflicts.
 - If permissions 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. (Optional) In the **File Timestamps Transfer** section, enable **Use source last accessed time** to apply the source file's last accessed time to the destination file.

This timestamp is applied only when the file is transferred. If the timestamp on the source changes but the file contents remain the same, the updated timestamp is not applied to the destination.

Note: PeerGFS does not synchronize any folder timestamps.

8. 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.
3. (Optional) Edit the second job to add additional target participants.

Now you are ready to set up the relay job.

4. 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.)
5. Click **Job Relay** in the navigation tree.

Edit File Replication Job

Participants
General
File and Folder Filters
Scheduled Replication Filters
Conflict Resolution
Delta Replication
File Metadata
Job Relay
File Locking
Application Support
Target Protection
Email Alerts
SNMP Notifications
Tags

Job Relay

☒ Disable Event Relay
☐ Relay Events From All Jobs
☐ Relay Events From The Jobs Selected Below:

Name	Source Agent	Path on Source	Target Agent	Path on Target	
------	--------------	----------------	--------------	----------------	--

OK Cancel

6. 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.

7. Click **OK**.

Relaying Events from Selected Jobs

In the following example, the relay job has two target participants: DGWin16C and DGWin16D.

Edit File Replication Job

Participants

Protocol:

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target
Lab24Win19b		C:\Security Keys Target	Yes	None	Yes
Lab24Win19c		C:\Security Keys Target	Yes	None	Yes
Lab24Win22d		C:\Security Keys Relay	Yes	Windows Driver	No

Buttons: Add, Edit, Delete

Buttons: OK, Cancel

In the Job Relay page, you select which jobs are replicated to the target participants:

Edit File Replication Job

Job Relay

☐ Disable Event Relay
☐ Relay Events From All Jobs
☒ Relay Events From The Jobs Selected Below:

Name	Source Agent	Path on Source	Target Agent	Path on Target
☐ Security Keys	Lab24Win22b	C:\Security Keys	Lab24Win22d	C:\Security Keys Relay

Buttons: OK, Cancel

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 applied automatically to all watch sets for all participants in the job.

The **Application Support** page appears only when the **SMB** protocol is selected. This page provides options specific to SMB-based environments, such as support for Microsoft applications that maintain file locks or open handles. If you select **NFS** or **Multi-Protocol**, the page does not appear because these protocols do not use SMB-specific application support features.

The **Application Support** page lists file types for which Peer Software has established best practices. These best practices may include filtering recommendations, prioritization of specific file types, and configuration of file-locking behavior. However, if an application is not listed, this does not necessarily mean it is unsupported. For details about how a specific application is optimized, contact [Peer Support](#).

To configure application support options for a file replication job:

1. Select the applications to be optimized.

Edit File Replication Job

Participants
General
File and Folder Filters
Scheduled Replication Filters
Conflict Resolution
Delta Replication
File Metadata
Job Relay
File Locking
Application Support
Target Protection
Email Alerts
SNMP Notifications
Tags

Application Support

Select below to optimize this job for any of the following file types:

Adobe Products

☐ Adobe Illustrator ☐ Adobe Photoshop
☐ Adobe InDesign

Autodesk Products

☐ Autodesk AutoCAD ☐ Autodesk Revit
☐ Autodesk Civil 3D ☐ Autodesk Sheet Set Manager (for AutoCAD or Civil 3D)
☐ Autodesk Inventor

Other

☐ ArcGIS ☐ Microsystems Allegro
☐ Dassault Systems CATIA ☐ Newforma Project Center
☐ Microsoft Office ☐ Rhinoceros Rhino3D

OK Cancel

2. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

Target Protection

Target protection helps prevent accidental data loss by saving a backup copy of any file that is deleted or overwritten on a target host. When target protection is enabled for a job, the system moves a copy of the existing file to a hidden `.pc-trash_bin` folder on the target host—**before** applying changes from the source host.

The `.pc-trash_bin` folder is located at the root of the target host's watch set. Within this folder, the file is stored in the same relative directory structure as its original location in the watch set.

For instructions on how to restore a file from the trash bin, see [Restoring a File from the Trash Bin](#).

Enabling Target Protection

To enable and configure target protection for a job:

1. Configure the available fields as needed:

The screenshot shows the 'Edit File Replication Job' window with the 'Target Protection' tab selected. The 'Enabled' checkbox is checked. The '# of Backup Files to Keep' is set to 3, and the '# of Days to Keep' is set to 30. The 'Trash Bin' field contains the text '.pc-trash_bin'. The dialog has 'OK' and 'Cancel' buttons at the bottom right.

Field	Description
Enabled	Enables target protection for this job.
# of Backup Files to Keep	Specifies the maximum number of backup copies to retain for each file. Older copies will be purged once this limit is reached.

Field	Description
# of Days to Keep	Sets the number of days to retain backup copies in the trash bin. A value of 0 disables automatic purging based on age.
Trash Bin	Displays the name of the trash bin folder used to store backup files. This folder is located in the root of the target host's watch set and is hidden. The folder name is fixed and cannot be changed.

2. Click **OK** or **Apply** to save your changes.

Restoring a File from the Trash Bin

To restore a deleted or overwritten file from the trash bin:

1. On the **target host**, open the watch set directory and locate the hidden folder named `.pc-trash_bin` at the root level.

Note: The `.pc-trash_bin` folder is hidden by default. Access may require administrative or root privileges, and hidden file visibility must be enabled on the system.

2. Navigate through the folder structure to locate the backup copy. The folder hierarchy mirrors the structure of the original watch set.
3. Copy the desired file from the `.pc-trash_bin` folder.
4. Paste the file back into its original location within the watch set.


The restored file is then synchronized to all other participants in the job.

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.



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 Replication Job** dialog allows you to select which SNMP notification to apply to a File Replication 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 Replication 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 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 [Tagging Resources](#) in the [Advanced Topics](#) section.



Editing Multiple Jobs

Peer Management Center supports **multi-job editing**, which allows you to edit multiple File Replication jobs at the same time. For example, you can change a single configuration item, such as **Verify Full File Checksum**, for multiple jobs in one operation instead of changing the item individually for each job.

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:

Edit File Replication Jobs - Multiple Selected

General

File and Folder Filters
Scheduled Replication Filters
Delta Replication
File Metadata
File Locking
Target Protection
Email Alerts
SNMP Notifications

General

Job ID: MULTIPLE SELECTED
Job Type: MULTIPLE SELECTED
Transfer Block Size (KB): 2048
Verify Block Checksums: ☒
Verify Full File Checksums: ☒
Enable Multipart Transfers: ☐
Resume File Transfers: ☐
Resume Transfer Minimum File Size (MB): 256
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: Multiple Values - Click to Edit
Auto Start Priority: 1

OK Cancel

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:

Verify Full File Checksum

Select a value to apply to all selected configurations:

☐ false (from: FR-3)
☐ true (from: Security Keys)

OK Cancel

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. The process consists of the following steps:

[Step 1: Job Type and Name](#)

[Step 2: Participants](#)

[Step 3: Master-Edge Assignment](#)

[Step 4: File Metadata](#)

[Step 5: Email Alerts](#)

[Step 6: DFS Namespace](#)


[Step 7: Save Job](#)

Additional configuration options, such as applying [file filters](#) and specifying [delta level replication](#), are available when [editing a File Synchronization job](#).

Step 1: Job Type and Name


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

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



3. Click **File Synchronization**, and then click **Create**.
4. Enter a name for the job in the dialog that appears.

The job name must be unique.



5. Click **OK** to proceed to the [Participants](#) page.

Step 2: Participants

On the **Participants** page, you select a file-sharing protocol for the job and start the process of adding participants to your job.

A File Synchronization job requires at least two participants. Each participant consists of a **Management Agent** and a **storage object** (volume, share, export, or folder). After you add a participant, it appears in the **Participants** table. The table lists each participant's Management Agent, storage path, and other key information needed for the File Synchronization job.

Follow these steps:

1. Select the file-sharing protocol that matches the workload for this job: **SMB**, **NFS**, or **Multi-Protocol**. The file-sharing protocol you select applies to the entire job and cannot be changed after the job is created. Available protocol options vary based on your environment and license.

Note: Multi-Protocol requires specific license support. Before creating a job that uses both SMB and NFS, verify that your license includes Multi-Protocol support. For more information about the Multi-Protocol feature, see [Multi-Protocol](#) in the [Advanced Topics](#) section.

Agent Type	License Supports Multi-Protocol?	Available Protocol(s)	Default Selection
All Windows-based	N/A	SMB	SMB is preselected and required
All Linux-based	No	NFS	NFS is preselected and required
All Linux-based	Yes	NFS or Multi-Protocol	No default—you must choose
Mixture of Windows- and Linux-based	Yes	SMB, NFS, or Multi-Protocol	No default—you must choose.



The **Management Agent** page lists the available [Agents](#) in your environment. After selecting a file-sharing protocol, only compatible Agents will be available:

- You can assign multiple Agents to the same storage device. However, each Agent must be used to access different volumes, shares, exports, or folders on that device.

1. From the **Agents** table, select an Agent that can access the storage device you want to replicate. This Agent will be designated as a Management Agent.

Copyright (c) 1993-2026 Peer Software, Inc. All Rights Reserved.

Add New Participant

Management Agent
Select the server hosting the Peer Agent that manages this storage device.

Management Agent
Storage Platform
Storage Information
Path
Edge Caching

Only Microsoft Windows Agents are available based on job type.

Agent	OS	Computer Description
DGWin16A	Microsoft Windows	
DGWin16B	Microsoft Windows	
DGWin16C	Microsoft Windows	
DGWin16D	Microsoft Windows	

< Back Next > Finish Cancel

2. Click **Next** to proceed to the [Storage Platform](#) page.

Storage Platform

The **Storage Platform** page lists the available storage platforms types for your File Synchronization job.

1. Select the storage platform type hosting the data you want to replicate.

The types available depend on the file sharing protocol in the previous step:

- If you selected **SMB**, these are the available storage platform types:

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. Inside the window, the "Storage Platform" tab is selected in the left-hand navigation pane. The main area displays the instruction "Select the type of storage platform." followed by 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 PowerStore
- ☐ Dell Unity
- ☐ Nutanix Files

At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

- If you selected **NFS**, these are the available storage platform types:

The screenshot shows a software window titled "Add New Participant" with a blue header bar. Inside, the "Storage Platform" tab is selected, displaying the instruction "Select the type of storage platform." On the left, a sidebar lists "Management Agent", "Storage Platform" (highlighted), "Storage Information", and "Path". The main area contains a list of storage platform options, each with a radio button and an icon: "Windows File Server" (blue square icon), "Linux File Server" (blue penguin icon), "NetApp ONTAP" (blue square icon with 'N', selected), "Amazon FSx for NetApp ONTAP" (blue cloud icon with 'FSx'), "Dell PowerScale" (blue square icon with 'P'), "Dell PowerStore" (blue square icon with 'P'), "Dell Unity" (blue square icon with 'U'), and "Nutanix Files" (blue square icon with 'X'). At the bottom right, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

- If you selected **Multi-Protocol (SMB and NFS)**, these are the available storage platform types:

Add New Participant

Storage Platform
Select the type of storage platform.

Management Agent
Storage Platform
Storage Information
Path

- ☐ Windows File Server
- ☐ Linux File Server
- ☒ NetApp ONTAP
- ☐ Amazon FSx for NetApp ONTAP
- ☐ Dell PowerScale
- ☐ Dell PowerStore
- ☐ Dell Unity
- ☐ Nutanix Files

< Back Next > Finish Cancel

2. Click **Next** to proceed to the [Storage Information](#) page.

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. If you selected **Existing Credentials**, choose a credential from the drop-down list, and then click **Next** to proceed to the [Path](#) page.
4. If you selected **New Credentials**, enter the required connection details. The fields vary by platform. For platform-specific field descriptions, see:

[Amazon FSx for NetApp ONTAP](#)

[Dell PowerScale](#)

[Dell PowerStore](#)

[Dell Unity](#)

[NetApp ONTAP](#)

[Nutanix Files](#)

5. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
6. Click **Next** to proceed to the [Path](#) page.

Amazon FSx for NetApp ONTAP

To connect to the Amazon FSx for NetApp ONTAP Storage Virtual Machine (SVM) hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

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

*SVM Name:

*SVM User Name:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

Advanced

☐ Existing Credentials

Edit

Validate

You must enter the name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Amazon FSxN environments.

< Back

Next >

Finish

Cancel

Field	Description
SVM Name	Enter the name, FQDN, or IP address of the SVM hosting the data to be replicated.
SVM User Name	Enter the user name 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 is not required if the data LIFs (Logical Interfaces) associated with SVM Name already allow 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 is not listed, enter the address manually.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
4. Click **Next** to proceed to the [Path](#) page.

Dell PowerScale

To connect to the Dell PowerScale cluster hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

The available connection types depend on the operating system of the Agent, as well as the notification delivery method used by the storage platform.

- o **RabbitMQ**: Available only when the Agent is running on Windows. Choose RabbitMQ if you are using a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent.
- o **Syslog**: Available on both Windows and Linux Agents. Choose Syslog if the PowerScale device sends notifications directly to the Agent.

If you selected **Rabbit MQ**, this page appears:

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), "Path", and "Edge Caching". 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, there are two sections for "Credentials". The first section, "New Credentials", is selected with a radio button. It contains the following fields: "*SmartConnect Zone Name:", "*Cluster Management IP:", "*Cluster Username:", "*Cluster Password:", and "Access Zone:". Below these fields is a "Connection Type" section with two radio buttons: "Syslog" and "RabbitMQ" (which is selected). To the right of the "RabbitMQ" radio button is an "Advanced" button. The second section, "Existing Credentials", is unselected and contains a dropdown menu and an "Edit" button. At the bottom of the window, there is a "Validate" button followed by the text "You must enter a Cluster Name." Below this is a link: "Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerScale environments." At the very 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

*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.

These fields appear in both cases:

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 for accessing the REST-based API integrated into the PowerScale 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.
Access Zone (RabbitMQ) or Access Zones (Syslog)	Optional. Enter the name(s) of the access zone(s) being monitored.
Connection Type	Select the appropriate method for sending real-time event notifications to the Agent: • Syslog: Choose if the storage device directly transmits notifications to the Agent. • RabbitMQ: Choose if you are using the CEE framework to dispatch notifications to the Agent.

2. If you selected **Syslog**, enter values for the following additional fields in the **Syslog** section:

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 is not listed, enter the address manually.
Listening Port	Enter the port through which the Agent will receive TLS-based syslog events from the PowerScale cluster.
SSL Certificate Path	Enter the path to the certificate used for the TLS-based syslog connection between the Agent and the PowerScale cluster. See Dell PowerScale Syslog Configuration Guide for more details.
SSL Private Key Path	Enter the path to the private key for the TLS connection. For more information, see the Dell PowerScale Syslog Configuration Guide .
SSL Private Key Password	(Optional) Enter the password for the private key if it's protected. For more information, see the Dell PowerScale Syslog Configuration Guide .

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 PowerStore

To connect to the Dell PowerStore NAS server hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection detail in the **Credentials** section.

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

*NAS Server Name:

*Management IP:

*Management Username:

*Management Password:

Connection Type:

☒ CEE

CEE Configuration

*Event Pool Name:

*Event Publisher Name:

*Agent IP Address:

*Listening Port:

9843

Advanced

☐ Existing Credentials

PTPSTORE3, user:storageadmin

Edit

Validate

You must enter a NAS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell PowerStore environments.

< Back

Next >

Finish

Cancel

Copyright (c) 1993-2026 Peer Software, Inc. All Rights Reserved.

Field	Description
NAS Server Name	Enter the name of the PowerStore server hosting the data to be replicated.
Management IP	Enter the IP address for accessing the REST-based API integrated into the PowerStore server.
Management Username	Enter the user name for the account managing the PowerStore server.
Management Password	Enter the password for account managing the PowerStore server.
Connection Type	Select the method for sending real-time event notifications to the Agent: • CEE: Choose if you are using the CEE framework to dispatch notifications to the Agent.

2. Enter the required values in the **CEE Configuration** section.

Field Description	Description
Event Pool Name	Enter the exact name of the event pool on the PowerStore device that publishes CEE notifications. You can find this name in the PowerStore Manager UI under Event Notifications.
Event Publisher Name	Enter the exact name or ID of the event publisher on the PowerStore device that sends CEE notifications to the Agent. You can find this value in the PowerStore Manager UI under Event Notifications.
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the PowerStore device. The device must

Field Description	Description
	be able to route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

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 Unity

To connect to the Dell Unity NAS Server hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

The available connection types depend on the operating system of the Agent, as well as the notification delivery method used by the storage platform.

- o **RabbitMQ**: Available only when the Agent is running on Windows. Choose RabbitMQ if you are using a RabbitMQ broker with the CEE framework to dispatch notifications to the Agent.
- o **Syslog**: Available on both Windows and Linux Agents. Choose Syslog if the PowerScale device sends notifications directly to the Agent.

If you selected **Rabbit MQ**, 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

*NAS Server Name:

*Unisphere Management IP:

*Unisphere Username:

*Unisphere Password:

Connection Type: ☒ RabbitMQ ☐ CEE

Advanced

☐ Existing Credentials

MLUNITY, user:admin

Edit

Validate

You must enter a NAS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments.

< Back

Next >

Finish

Cancel

If you selected **CEE**, 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

*NAS Server Name:

*Unisphere Management IP:

*Unisphere Username:

*Unisphere Password:

Connection Type:

☐ RabbitMQ

☒ CEE

CEE Configuration

*Event Pool Name:

*Agent IP Address:

*Listening Port:

9843

Advanced

☐ Existing Credentials

MLUNITY, user:admin

Edit

Validate

You must enter a NAS Server Name.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Dell Unity environments.

< Back

Next >

Finish

Cancel

These fields appear in both cases:

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 managing the Unity storage device. This should not be the IP address of the NAS server itself.
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.

2. If you selected **CEE**, enter values for the following additional fields in the **CEE Configuration** section.

Field Description	Description
Event Pool Name	Enter the exact name of the event pool on the Unity device that publishes CEE notifications. You can find this name in the Unity Manager UI under Event Notifications.
Agent IP Address	Select the IP address of the server hosting the Agent that will receive real-time event notifications for the Unity device. The device must be able to route traffic to this address. If the desired IP address is not listed, enter the address manually.
Listening Port	Enter the TCP port through which the Agent will receive real-time event notifications from CEE.

3. Click **Advanced** if you want to set [advanced options](#).

- Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
- Click **Next** to proceed to the [Path](#) page.

Linux File Server


To connect to the Linux File Server hosting the data to be replicated:

- Select the **Detector Type** for the Linux File Server. **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 blue header bar. Below the header, the "Storage Information" tab is selected, showing a list of tabs: "Management Agent", "Storage Platform", "Storage Information" (highlighted), and "Path". The main area of the dialog is titled "Storage Information" with the instruction "Enter the information required to connect to the storage device." Below this, there is a "Detector Type:" label followed by a dropdown menu currently set to "Linux Connector". To the right of the dropdown is an "Advanced" button. At the bottom of the dialog, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

- Click **Advanced** if you want to set [advanced options](#).
- Click **Next** to proceed to the [Path](#) page.

1. Modify the Linux Detector options as desired.



The image shows a dialog box titled "Linux Detector Options". It contains four input fields: "Filter open/close events from these users:", "Filter all events from these users:", "Filter events from these IP Addresses:", and "Access Event Suppression Time:". The "Access Event Suppression Time" field has a spinner control with the value "-1". At the bottom are "OK" and "Cancel" buttons.

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 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 can be used to filter out events from backup and/or archival services running under specific user accounts.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This can be used to filter out events from backup and/or archival services based on their IP addresses.
Access Event Suppression Time	Enter the number of seconds to delay an open event before processing. This reduces chatter from Windows clients when mousing over files in Windows Explorer. The default value is -1, which uses a globally set value. A value of 0 allows dynamic changes to the delay time based on the system load.

2. Click **OK**.

NetApp ONTAP

To connect to the NetApp ONTAP Storage Virtual Machine (SVM) hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

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 with a dark background), "Path", and "Edge Caching". 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, there are two sections for "Credentials". The first section, "New Credentials", is selected with a radio button. It contains five input fields: "*SVM Name:", "*SVM User Name:", "*SVM Password:", "SVM Management IP:", and "*Peer Agent IP:". The "*Peer Agent IP:" field has a dropdown arrow on its right. To the right of these fields is a button labeled "Advanced". The second section, "Existing Credentials", is unselected. It contains a dropdown menu showing "PTSVM93B, user:vsadmin" and an "Edit" button to its right. 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 line of text: "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".

Field	Description
SVM Name	Enter the name, FQDN, or IP address of the 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 SVM Name already allow 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 is not listed, enter the address manually.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
4. Click **Next** to proceed to the [Path](#) page.

Nutanix Files

To connect to the Nutanix Files cluster hosting the data to be replicated:

1. If you selected **New Credentials**, enter the required connection details.

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

*Nutanix File Server Name:

*Username:

*Password:

*Peer Agent IP:

Advanced

☐ Existing Credentials

Edit

Validate

Nutanix File Server Name cannot be empty.

Having trouble connecting? Please verify that all [prerequisites](#) are met for Nutanix Files environments.

< Back

Next >

Finish

Cancel

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 is not listed, enter the address manually. The IP address should not point to the Files cluster itself.

2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Validate** to test the credentials. If validation is successful, a confirmation message appears.
4. Click **Next** to proceed to the [Path](#) page.

Windows File Server

To connect to the Windows File Server hosting the data to be replicated:

1. Select the detector type for the Windows File Server. **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 Support.

The screenshot shows a window titled "Add New Participant" with a blue header bar. Below the header, the "Storage Information" tab is selected, displaying the instruction "Enter the information required to connect to the storage device." On the left, a sidebar lists "Management Agent", "Storage Platform", "Storage Information" (highlighted), "Path", and "Edge Caching". The main area shows "Detector Type:" with a dropdown menu set to "Windows Driver" and an "Advanced" button. At the bottom, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".


2. Click **Advanced** if you want to set [advanced options](#).
3. Click **Next** to proceed to the [Path](#) page.

When configuring a job that uses a **Windows File Server**, you can set advanced options to control how the system filters, logs, and follows file activity.

The available options depend on the selected driver. Some options are common to both the **Windows driver** and **Windows Filter driver**, while others are unique to the **Windows Filter driver**.

1. Modify the Windows Detector options as desired.

If you selected **Windows**, this page appears:



The screenshot shows the "Windows Detector Options" dialog box. It has a blue title bar with the text "Windows Detector Options" and standard window controls. The main area contains several settings: a text box for "Filter open/close events from these users:", a spinner box for "Access Event Suppression Time:" set to "-1", and a section titled "Reparse Point Options" containing three checked checkboxes: "Follow Junction Points", "Follow Mount Points", and "Follow Symbolic Links". At the bottom right are "OK" and "Cancel" buttons.

Windows Detector Options

Filter open/close events from these users:


Access Event Suppression Time:

Reparse Point Options

- ☒ Follow Junction Points
- ☒ Follow Mount Points
- ☒ Follow Symbolic Links

OK Cancel

If you selected **Windows Driver**, this page appears:



The screenshot shows the "Windows Driver Detector Options" dialog box. It has a blue title bar with the text "Windows Driver Detector Options" and standard window controls. The main area contains several settings: four text boxes for "Filter open/close events from these users:", "Filter all events from these users:", "Filter events from these IP Addresses:", and "Filter events from these local processes:"; a spinner box for "Access Event Suppression Time:" set to "-1"; four checkboxes: "Enable Local Access Events:" (unchecked), "Enable Remote IP Address Logging:" (checked), "Enable Close Modify:" (checked), and "Close Modify Extension Override:" (unchecked); and a section titled "Reparse Point Options" containing three checked checkboxes: "Follow Junction Points", "Follow Mount Points", and "Follow Symbolic Links". At the bottom right are "OK" and "Cancel" buttons.

Windows Driver Detector Options

Filter open/close events from these users:

Filter all events from these users:

Filter events from these IP Addresses:

Filter events from these local processes:

Access Event Suppression Time:

Enable Local Access Events: ☐

Enable Remote IP Address Logging: ☒

Enable Close Modify: ☒

Close Modify Extension Override:

Reparse Point Options

- ☒ Follow Junction Points
- ☒ Follow Mount Points
- ☒ Follow Symbolic Links

OK Cancel

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 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 can be used to filter out events from backup and/or archival services running under specific user accounts.
Filter events from these IP Addresses	Enter a comma-separated list of client IP addresses from which all file activities will be ignored. This can be used to filter out events from backup and/or archival services based on their IP addresses.
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 can be used to filter out events from backup and/or archival services based on the process names.
Access Event Suppression Time	Enter the number of seconds to delay an open event before processing. This reduces chatter from Windows clients when mousing over files in Windows Explorer. The default value is -1, which uses a globally set value. A value of 0 allows dynamic changes to the delay time based on the system load, based on the load of the system.
Enable Local Access Events	Select to enable tracking of open and close events that occur locally on the Agent server.
Enable Remote IP Logging	Select to 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 is triggered when the file is closed.
Close Modify Extension Override	Enter a comma-separated list of exclusions for the Enable Close Modify option. Modify/write events will be detected for these files. This is important for those who rely on sync-on-save functionality.
Follow Junction Points	Select to enable junction point support for the selected Windows file server.

For more information about junction points or symbolic links options, contact [Peer Support](#).

2. Click **OK**.

Path

The **Path** page allows you to specify the path to the volume, share, export, or folder you want to synchronize. This volume, share, export, or folder is referred to as the [watch set](#). A job can monitor only one watch set. To synchronize multiple volumes, shares, exports, or folders, create a separate job for each.

1. Browse to or enter the path to the watch set.

The screenshot shows a window titled "Add New Participant" with a blue header bar. Below the header, the "Path" tab is selected in a sidebar on the left. The main area of the window has a title "Path" and a red error message: "You must enter a valid path." Below this, there is a text input field labeled "Enter Path" and a "Browse" button to its right. A checkbox labeled "Seeding Target" is located below the input field. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".


If you are entering the path manually rather than browsing, use the appropriate path format based on the file-sharing protocol (**SMB**, **NFS**, or **multi-protocol**) and the type of storage device:

Protocol	Storage Type	Path Format	Notes
SMB	Windows File Server	C:\data	Local path
	NAS device	\server_name_or_ip\shared_folder\file_or_directory_path	UNC path
NFS	Linux File Server	/mnt/data	Local path (case sensitive)
	NAS device	host:/volume	Export path
Multi-Protocol	NAS device	host:/volume	Export path

Note:

- If the path field is empty, clicking **Browse** shows local drives and folders on the Agent server.
 - If you enter the start of a UNC path and click **Browse**, the **Folder Browser** dialog attempts to display available shares on the specified file server.
 - If you enter the start of an NFS path and click **Browse**, the **Folder Browser** dialog attempts to display available exports on the specified NFS server. You may need to enter the full NFS path manually.
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 SMB access is not required, **skip to Step 5**.

SMB access enables the Linux Agent to access data on the file server using the SMB protocol. SMB access is **required** when creating a **multi-protocol** File Synchronization job in any of the following scenarios

- Synchronizing Windows-style permissions
 - Synchronizing mixed permissions (Windows-style and Linux-style; supported only with PowerScale)
 - Using a Nutanix Files storage device
4. If you are creating a multi-protocol job and SMB access is required, specify a value for **SMB Access Configuration** by selecting an existing configuration from the **SMB Access Configuration** drop-down or creating a new one.

To create a new SMB access configuration, see [Configuring SMB Access](#).

4. Choose an option for **SMB Access Share**. The selected share will be used by the Agent to access data via the SMB protocol. Any Windows-style permissions associated with the files will be read from and written to the file server via this share:
 - **Use Admin Share:** Select this option if you have an existing admin share that you want to use. The same admin share can often be used for multiple jobs.
 - **Specify Share Name:** Select this option if you prefer or require a dedicated share for the job, and enter the share name.
5. Click **Next** if you want this participant to use [Edge Caching](#); otherwise, click **Finish** to complete the configuration of this participant.

Configuring SMB Access

If your job requires SMB access, you can either use an existing configuration or create a new one.

Recommendation: Create SMB access configurations in advance using the [Storage Platforms](#) section of [Preferences](#).

To create a new SMB access configuration:

1. Click **Edit SMB Access Configuration**.

The screenshot shows the 'Add New Participant' dialog box with the 'Path' tab selected. A red error message at the top states: 'The root of the file server (/ or /ifs) is not supported. You must specify a export or folder path below /ifs.' The left sidebar contains 'Management Agent', 'Storage Platform', 'Storage Information', and 'Path'. The main area has a text field 'MLPS2ZONE1:/ Enter Path' with a 'Browse' button. Below it is a 'Seeding Target' checkbox. The 'SMB Access' section is expanded, showing a warning: 'Required only when synchronizing Windows-style permissions.' with a link to 'Edit SMB Access Configuration'. Below the link is an 'SMB Access Configuration' dropdown menu. The 'SMB Access Share' section has two options: 'Use Admin Share (IFS\$)' (selected) and 'Specify Share Name:' with an empty text field. At the bottom are buttons for '< Back', 'Next >', 'Finish', and 'Cancel'.

Add New Participant

Path

The root of the file server (/ or /ifs) is not supported. You must specify a export or folder path below /ifs.

Management Agent
Storage Platform
Storage Information
Path

MLPS2ZONE1:/ Enter Path

☐ Seeding Target

— SMB Access —

Required only when synchronizing Windows-style permissions.


[Edit SMB Access Configuration](#)

SMB Access Configuration:

SMB Access Share: ☒ Use Admin Share (IFS\$)

☐ Specify Share Name:

2. If dialog that appears. click **Create**:



3. In the **New SMB Access Configuration** dialog, supply the requested information:

The 'New SMB Access Configuration' dialog box contains the following fields and buttons:

- Description:
- Domain Username:
- Domain Password:
- Buttons: OK, Cancel

Field	Description
Description	Enter a name or description to identify the configuration.

Field	Description
Domain Username	Enter the domain username in the format: Domain\:Username
Domain Password	Enter the password for the domain account.

- Click **OK** to save the configuration.
- Click **Apply and Close** to apply the configuration to the job and return to the wizard.

Edge Caching

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 blue header bar. Below the header, the "Edge Caching" tab is selected. The main area contains the text "Enable Edge Caching and select role." and a list of roles on the left: "Management Agent", "Storage Platform", "Storage Information", "Path", and "Edge Caching". The "Edge Caching" role is highlighted. To the right of the list is a checkbox labeled "Enable Edge Caching". At the bottom of the window 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 blue header bar. Inside, the "Edge Caching" tab is selected, showing a list of components on the left: Management Agent, Storage Platform, Storage Information, Path, Edge Caching (highlighted), and Master Data Service. The main area contains the text "Enable Edge Caching and select role." followed by a checked checkbox for "Enable Edge Caching" and a section titled "Select Edge Caching role:" with two radio buttons: "Master" (selected) and "Edge". At the bottom are four buttons: "< Back", "Next >", "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.

Add New Participant

Edge Caching
Enable Edge Caching and select role.

Management Agent
Storage Platform
Storage Information
Path
Edge Caching
Volume Policy
Utilization Policy
Pinning Filter

☒ Enable Edge Caching

Select Edge Caching role:

☐ Master
☒ Edge

< Back Next > Finish 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 following fields on this page are automatically populated:

- **Protocol:** Lists the protocol that will be used to transfer file content between master participants and edge participants. HTTPS is currently the only available option because it requires only one open firewall port on each master participant and handles latency effectively over WAN connections.
- **Agent Name:** Lists the name of the Management Agent that you selected at the beginning of Step 2.

Configure the remaining fields as needed:

1. (Optional) Enter a value for **Agent Alias**. The value can be a hostname, FQDN, 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: Lab24Win22a
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.

The screenshot shows a window titled "Add New Participant" with a close button in the top right corner. The window has a left-hand navigation pane and a main content area. The navigation pane lists several categories: "Management Agent", "Storage Platform", "Storage Information", "Path", "Edge Caching", and "Master Data Service". The "Master Data Service" item is currently selected and highlighted. The main content area is titled "Master Data Service" and contains the instruction "Configure access to the Master Data Service." Below this, there is a link that says "Edit Master Data Service". Underneath the link, a box displays the "Current Master Data Service (Lab24Win22a)" with the following details: "Protocol: HTTPS", "Agent Name: Lab24Win22a", "Agent Alias:", and "Port: 8446". At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

3. Click **Finish**.

The **Participants** page reappears. The participant is listed in the **Participants** table with the **Master** role.

Agent	OS	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching
Lab24Win22a	Microsoft Windows		C:\Company Files\Training - Mui	Yes	Windows Driver	No	Master

- 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 22 : 00 : 00

☐ Define schedule

☒ Run caching scan after job start

*Temporary Storage Path: Browse

\\.PeerTempPath

< Back Next > Finish Cancel

2. In the **Cache Threshold Alerts** section, set threshold values for automatic alerts about free disk space and cache usage.

Peer Management Center will automatically display alerts in the **Alerts** tab when:

- The amount of free disk space on the volume falls below the specified value. For example, if a 1 TB volume has 500 MB of free space and the threshold is set to 512 MB, an alert will be sent.
- Cache usage on the volume exceeds the specified percentage of the cache size. For example, if the cache size is set to 80%, equating to 750 GB, 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.

3. In the **Caching Scan Schedule**, set the frequency that the volume should be scanned to optimize the balance of fully hydrated vs stubbed files.

This scan can be run daily at a specified time, or you can define a more customized schedule.

4. In the **Temporary Storage Path** field, enter a path or browse to the location you want to be used as temporary storage space.


The temporary storage space will be used to store the content of stub files as they are 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 if 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**.

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
☐ Existing Filter

< Back Next > Finish 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.

Add New Participant

Pinning Filter

✖ Name should not be empty.

Management Agent
Storage Platform
Storage Information
Path
Edge Caching
Volume Policy
Utilization Policy
Pinning Filter

Edit Pinning Filters

☐ No Filter

☒ New Filter

*Name:

Pinning Rules:

Path	Pinning State
------	---------------

☐ Existing Filter


Create
Edit
Delete

< Back Next > Finish Cancel

3. Enter a name for the filter.

4. Click **Create**.

The **Create Pinning Rule** dialog appears.



The image shows a 'Create Pinning Rule' dialog box with a blue title bar and a close button (X) in the top right corner. The dialog contains the following fields and controls:

- Directory Patterns:** A text input field.
- Pinning State:** A dropdown menu currently showing 'Local at Edge'.
- ☐ Unless a file's size is greater than: A numeric input field with '1' and a unit dropdown menu showing 'GB'.
- ☐ Unless a file has not been accessed in last: A numeric input field with '3' and a unit dropdown menu showing 'Months'.
- A red error message: 'Directory patterns must not be empty.'
- At the bottom right, there are 'OK' and 'Cancel' buttons.

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

✖ Name should be unique.

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

7. (Optional) Create additional pinning rules.

8. Click **Finish**.

The **Participants** page reappears. The participant is listed in the **Participants** table with the **Edge** role.

Create File Synchronization Job Wizard

Participants
Add two or more participants to this File Synchronization job.

Agent	OS	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching
Lab24Win22a	Microsoft Windows		C:\Company Files\Training - Mui	Yes	Windows Driver	No	Master
Lab24Win22b	Microsoft Windows		C:\Company Files\Training - Cer	Yes	Windows Driver	No	Edge
Lab24Win22c	Microsoft Windows		C:\Company Files\Training - Tok	Yes	Windows Driver	No	Master

Navigation: Back Next > Finish Cancel

9. Continue adding more participants if applicable or continue with [Step 3: Master-Edge Assignment](#)

Step 3: Master-Edge Assignment

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

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.

1. Select an edge participant in the **Assignment** table.

Create File Synchronization Job Wizard

Master-Edge Assignment
Assign one or more master participants to each edge participant.

Participants

- Master-Edge Assignment
- File Metadata
- Email Alerts
- DFS Namespace

Edge Participant	Path	Master Participants
Lab24Win22b	C:\Company Files\Training - Centreville	0
Lab24Win22d	C:\Company Files\Training - DC	0


Assign

< Back Next > Finish Cancel

2. Click the **Assign** button.


The **Assign Master Participants** dialog appears.

4. (Optional) Change the failover order by selecting a master participant and using the **Move Up** and **Move Down** buttons.



5. Click **OK**.

The **Master Participants** column has been updated for that participant.



6. Repeat Steps 1-5 for each edge participant.

7. Click **Next**.

The [File Metadata](#) page appears.

Step 4: File Metadata

This step is optional.

The **File Metadata** page enables you to configure how and when file metadata is synchronized between source and target hosts. Metadata is synchronized in real time and during scans.

Available options include specifying whether to synchronize file permissions (security metadata) and defining the metadata source (volume, share, export, or folder) to use for resolving conflicts during initial synchronization. This designated source, referred to as the [master host](#), determines which version of the metadata is retained when discrepancies occur.

For more information, see [File Metadata Synchronization](#) in [Advanced Topics](#).

The content of the File Metadata page depends on the file-sharing protocol you selected in Step 2:

- See [SMB Protocol](#) for details.
- See [NFS Protocol](#) for details.
- See [Multi-Protocol](#) for details.

SMB Protocol

To configure file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata will be synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. Select the permissions to be synchronized: **Owner**, **DACL**, and **SACL**.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

4. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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.

NFS Protocol

To configure file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

The screenshot shows a window titled "Create File Synchronization Job Wizard" with a close button in the top right corner. The "File Metadata" step is active, indicated by a blue header and a red error message: "You must select a master host when enabling security or attributes synchronization during initial scan." On the left, a sidebar lists "Participants" with "File Metadata" selected and "Email Alerts" below it. The main area contains three sections: "Synchronize File Security Information" with two checked checkboxes ("Enable synchronizing file security information in real-time" and "Enable synchronizing file security information with master host during initial scan"); "Synchronize Security Options" with three checked checkboxes ("Owner", "Group", "Linux Permissions") and one unchecked checkbox ("ACL (NFSv4 or POSIX)"); and "Metadata Conflict Resolution" with a dropdown menu labeled "Select master host for initial scan:". At the bottom, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. Select the permissions to be synchronized: **Owner**, **Group**, **Linux Permissions**, and **ACL**.

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

4. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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.

Multi-Protocol

To configure file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **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.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. Select the security style to determine how file permissions are interpreted and synchronized across SMB and NFS environments:
 - **Windows-style permissions**
 - **Linux-style permissions**
 - **Mixed permissions** (available only with Dell PowerScale)

4. Select the 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 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
- DFS Namespace

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

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

- **Linux-style permissions**

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
- DFS Namespace

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)

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
- DFS Namespace

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: Mixed permissions ⚠ [How do permissions work with Multi-Protocol?](#)

Linux

- ☒ Owner
- ☒ Group
- ☒ Linux Permissions
- ☐ ACL (NFSv4 or POSIX)

Windows

- ☒ Owner
- ☒ DACL: Discretionary Access Control List

Validate

Metadata Conflict Resolution

Select master host for initial scan:

< Back Next > Finish Cancel

5. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - 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 5: Email Alerts

This step is optional.


An email alert notifies recipients when a certain type of event occurs, for example, session abort, host failure, or system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list is empty. Email alerts are defined in [Preferences](#) and can then be applied to multiple jobs of the same type.

Peer Software recommends that you create email alerts in advance. However, from this wizard page, you can select existing alerts to apply to the job or create new alerts to apply.

To create a new alert, see [Email Alerts](#) in the [Preferences](#) section.


To apply an existing email alert to the job.

1. Click the **Select** button.



The **Select Email Alert** dialog appears.

2. Select an alert from the **Email Alert** drop-down list.



3. Click **OK**.

The alert is listed on the **Email Alerts** page.

Create File Synchronization Job Wizard

Email Alerts
Select email alerts.

Participants
Master-Edge Assignment
File Metadata
Email Alerts
DFS Namespace

[Edit Email Alerts](#)

Name	Temporarily Disabled	Event Types	Recipients
Participant Failure	No	Participant Failure	AriF@peersoftware.com

Select
Delete
View Details

< Back Next > Finish Cancel

- (Optional) Repeat steps 1-3 to apply additional alerts.
- Click **Next** to proceed to the [DFS Namespace](#) page.

Step 6: DFS Namespace

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

This step is optional.

The **DFS Namespace** page presents three options for linking a DFS namespace folder to this File Synchronization job.

To link a namespace to this job:

- Click the **Enable linking job to DFS Namespace** checkbox.

The screenshot shows the 'Create File Synchronization Job Wizard' window, specifically the 'DFS Namespace' step. The window has a blue title bar and a sidebar on the left with the following menu items: 'Participants', 'Master-Edge Assignment', 'File Metadata', 'Email Alerts', and 'DFS Namespace' (which is highlighted). The main area contains the text '(Optional) Link a DFS namespace folder to this File Synchronization job.' Below this text is a checkbox labeled 'Enable linking job to DFS namespace', which is currently unchecked. To the right of the checkbox is a group box containing three radio button options: 'New DFS Namespace', 'Import DFS Namespace', and 'Existing DFS Namespace'. At the bottom of the window are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

The three options are enabled.

This screenshot is identical to the one above, but the checkbox 'Enable linking job to DFS namespace' is now checked. The rest of the interface, including the sidebar, options, and buttons, remains the same.

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.

Create File Synchronization Job Wizard

DFS Namespace
(Optional) Link a DFS namespace folder to this File Synchronization job.

Participants
Master-Edge Assignment
File Metadata
Email Alerts
DFS Namespace
DFS-N Link

☒ Enable linking job to DFS namespace

☐ New DFS Namespace
☐ Import DFS Namespace
☒ Existing DFS Namespace

Company Files

Selected DFS Namespace: Company Files

Folders:
- Training
Targets:
- \\Lab24Win22c\Company Files\Training - Tokyo
- \\AB24WIN22d\Company Files\Training - DC
- \\Lab24Win22b\Company Files\Training - Centreville
- \\Lab24Win22a\Company Files\Training - Munich

- DevOps
Targets:
- \\Lab24Win22a\Company Files\DevOps - Munich
- \\Lab24Win22d\Company Files\DevOps - DC
- \\Lab24Win22b\Company Files\DevOps - Centreville
- \\Lab24Win22c\Company Files\DevOps - Tokyo

- Project Management
Targets:

< Back Next > Finish Cancel

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](#).

Create File Synchronization Job Wizard

DFS-N Link
Link each participant's watch set to a folder target.

Participants
Master-Edge Assignment
File Metadata
Email Alerts
DFS Namespace
 DFS-N Link

Namespace:
Folder:

Linked Participants and Folder Targets:

Agent	Directory	Folder Target	Link Enabled
Lab24Win22a	C:\Company Files\Training - Munich		
Lab24Win22c	C:\Company Files\Training - Tokyo		
Lab24Win22b	C:\Company Files\Training - Centreville		
Lab24Win22d	C:\Company Files\Training - DC		

< Back Next > Finish Cancel

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.



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.

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role
Lab24Win22a		C:\Company Files\Training - Munich	Yes	Windows Driver	No	Master
Lab24Win22b		C:\Company Files\Training - Centreville	Yes	Windows Driver	No	Edge
Lab24Win22c	-	C:\Company Files\Training - Tokyo	Yes	Windows Driver	No	Master
Lab24Win22d		C:\Company Files\Training - DC	Yes	Windows Driver	No	Edge

3. Select a configuration item in the navigation tree and make the desired changes:

- [Participants](#)
- [General](#)
- [File and Folder Filters](#)
- [Scheduled Replication Filters](#)

- [File 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 Synchronization Job** dialog allows you to:

- [Add and remove participants from a job.](#)
- [Edit a participant.](#)

Edit File Synchronization Job

Participants

Protocol:

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role
Lab24Win22a		C:\Company Files\Training - Munich	Yes	Windows Driver	No	Master
Lab24Win22b		C:\Company Files\Training - Centreville	Yes	Windows Driver	No	Edge
Lab24Win22c	-	C:\Company Files\Training - Tokyo	Yes	Windows Driver	No	Master
Lab24Win22d		C:\Company Files\Training - DC	Yes	Windows Driver	No	Edge

Buttons: Add, Edit, Delete

Buttons: OK, Cancel

Adding and Deleting a Participant

This topic describes [adding](#) and [deleting](#) participants in a File Synchronization job.

Adding a Participant to a File Synchronization Job

To add a participant to a File Synchronization job:

1. Select the job in the **Jobs** view; right click and select **Edit Job**.

The **Edit File Synchronization** dialog opens; the **Participants** page displays the current job participants.

Edit File Synchronization Job

Participants

Protocol:

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role
Lab24Win22a		C:\Company Files\Training - Munich	Yes	Windows Driver	No	Master
Lab24Win22b		C:\Company Files\Training - Centreville	Yes	Windows Driver	No	Edge
Lab24Win22c	-	C:\Company Files\Training - Tokyo	Yes	Windows Driver	No	Master
Lab24Win22d		C:\Company Files\Training - DC	Yes	Windows Driver	No	Edge

Buttons: Add, Edit, Delete

Buttons: OK, Cancel

2. Click the **Add** button.

The **Add New Participant** wizard opens; the **Management Agent** page lists the Agents available to be added.

Tip: If the Agent you want is not shown in the **Agents** table, 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](#), the **Agents** table will be updated.

The screenshot shows a window titled "Add New Participant" with a blue header bar. Below the header, the "Management Agent" tab is selected in a sidebar on the left. The main area displays a message: "Only Microsoft Windows Agents are available based on your previously selected Agents." Below this message is a table with three columns: "Agent", "OS", and "Computer Description". The table contains two rows of data. At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Agent	OS	Computer Description
Lab24Win19b	Microsoft Windows	
Lab24Win19c	Microsoft Windows	

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 screenshot shows a window titled "Add New Participant" with a sidebar on the left and a main content area. The sidebar contains the following items: "Management Agent", "Storage Platform" (which is highlighted with a dark background), "Storage Information", "Path", and "Edge Caching". The main content area has the heading "Storage Platform" and the instruction "Select the type of storage platform." Below this, there is 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
- ☐  Dell PowerStore
- ☐  Nutanix Files

At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

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 window titled "Add New Participant" with a blue header bar. Below the header, the title "Path" is displayed. A red error icon and the message "You must enter a valid path." are shown. On the left, a sidebar lists the steps: Management Agent, Storage Platform, Storage Information, Path (highlighted), and Edge Caching. 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, 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 window titled "Add New Participant" with a close button in the top right corner. Inside the window, the "Edge Caching" tab is selected. Below the tab name is the instruction "Enable Edge Caching and select role." On the left side, there is a list of roles: "Management Agent", "Storage Platform", "Storage Information", "Path", and "Edge Caching". The "Edge Caching" role is highlighted with a grey background. To the right of this list is 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".

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.

Edit File Synchronization Job

Participants

Protocol:

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role
Lab24Win19b		C:\Company Files\Training - New York	Yes	Windows Driver	No	None
Lab24Win22a		C:\Company Files\Training - Munich	Yes	Windows Driver	No	Master
Lab24Win22b		C:\Company Files\Training - Centreville	Yes	Windows Driver	No	Edge
Lab24Win22c	-	C:\Company Files\Training - Tokyo	Yes	Windows Driver	No	Master
Lab24Win22d		C:\Company Files\Training - DC	Yes	Windows Driver	No	Edge

Buttons: Add, Edit, Delete, OK, Cancel

Deleting a Participant from a File Synchronization Job

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.

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.

Edit File Synchronization Job

Participants

Protocol:

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role
Lab24Win19b		C:\Company Files\Training - New York	Yes	Windows Driver	No	None
Lab24Win22a		C:\Company Files\Training - Munich	Yes	Windows Driver	No	Master
Lab24Win22b		C:\Company Files\Training - Centreville	Yes	Windows Driver	No	Edge
Lab24Win22c	-	C:\Company Files\Training - Tokyo	Yes	Windows Driver	No	Master
Lab24Win22d		C:\Company Files\Training - DC	Yes	Windows Driver	No	Edge

Buttons: Add, Edit, Delete, OK, Cancel

2. Click the **Delete** button.
3. Click **OK** in the **Delete Confirmation** dialog.

The participant is removed from the **Participants** table.

Editing a Participant

To edit a participant:

1. In the **Edit File Synchronization Job** dialog, select the participant you want to edit from the **Participants** table.

The screenshot shows the 'Edit File Synchronization Job' dialog box with the 'Participants' tab selected. On the left is a sidebar with various configuration options. The main area contains a table of participants. The first row is highlighted. To the right of the table are 'Add', 'Edit', and 'Delete' buttons. At the bottom right are 'OK' and 'Cancel' buttons.

Host	Computer Description	Directory	Enabled	Storage Platform	Seeding Target	Edge Caching Role
Lab24Win19b		C:\Company Files\Training - New York	Yes	Windows Driver	No	None
Lab24Win22a		C:\Company Files\Training - Munich	Yes	Windows Driver	No	Master
Lab24Win22b		C:\Company Files\Training - Centreville	Yes	Windows Driver	No	Edge
Lab24Win22c	-	C:\Company Files\Training - Tokyo	Yes	Windows Driver	No	Master
Lab24Win22d		C:\Company Files\Training - DC	Yes	Windows Driver	No	Edge

2. Click **Edit**.

The **Edit Participant** dialog appears.

The screenshot shows a Windows-style dialog box titled "Edit Participant". It has a blue header bar with standard window controls (minimize, maximize, close). Below the header, the "General" tab is selected, and the text "Modify these settings." is displayed. On the left side, there is a vertical pane with two tabs: "General" (which is active) and "Edge Caching". The main area of the dialog contains the following settings:

- Enabled:** A checkbox that is checked.
- Host:** A text input field containing "Lab24Win19b".
- Event Detector:** A dropdown menu currently showing "Windows Driver". To the right of this dropdown is a link that says "Edit Detector Configuration".
- Path:** A text input field containing "C:\Company Files\Training - New York". To the right of this field is a "Browse" button.
- Seeding Target:** An unchecked checkbox.

At the bottom of the dialog, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

If the job is a multi-protocol job, this dialog will be displayed

The screenshot shows the 'Edit Participant' dialog box with the 'General' tab selected. The dialog has a title bar with 'Edit Participant' and standard window controls. Below the title bar, the 'General' tab is active, showing the text 'Modify these settings.' and a list of configuration options. The 'Enabled' checkbox is checked. The 'Host' field contains 'lab24ubuntu1'. The 'Event Detector' field contains 'Dell PowerScale', with a link to 'Edit Detector Configuration' to its right. The 'Path' field contains 'MLPS2ZONE1:/ ifs/Zone1/MP1/AriF/FS-1 A', with a 'Browse' button to its right. There is an unchecked 'Seeding Target' checkbox. Below this is an 'SMB Access' section with a link to 'Edit SMB Access Configuration'. The 'SMB Access Configuration' dropdown is set to 'SMB Config A, Username: peertest\Labadmin'. The 'SMB Access Share' section has two radio buttons: 'Use Admin Share (IFS\$)' (selected) and 'Specify Share Name:' (unselected). At the bottom of the dialog are four buttons: '< Back', 'Next >', 'Finish', and 'Cancel'.

Edit Participant

General

Modify these settings.

General

Enabled: ☒

Host: lab24ubuntu1

Event Detector: Dell PowerScale [Edit Detector Configuration](#)

Path: MLPS2ZONE1:/ ifs/Zone1/MP1/AriF/FS-1 A [Browse](#)

☐ Seeding Target

SMB Access

[Edit SMB Access Configuration](#)

SMB Access Configuration: SMB Config A, Username: peertest\Labadmin

SMB Access Share:

☒ Use Admin Share (IFS\$)

☐ Specify Share Name:

< Back Next > Finish Cancel

3. To enable or disable the Agent, select or deselect the **Enabled** checkbox.
4. 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 you are entering the path manually rather than browsing, use the appropriate path format based on the file-sharing protocol (SMB, NFS, or multi-protocol) and the type of storage device:

Protocol	Storage Type	Path Format	Notes
SMB	Windows File Server	C:\data	Local path
	NAS device	\server_name_or_ip\shared_folder\file_or_directory_path	UNC path
NFS	Linux File Server	/mnt/data	Local path (case sensitive)
	NAS device	host:/volume	Export path
Multi-Protocol	NAS device	host:/volume	Export path

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.

The screenshot shows a window titled "Edit Participant" with a blue header bar. Below the header, the "Edge Caching" tab is selected, showing a sub-header "Edge Caching" and the instruction "Enable Edge Caching and select role." A checkbox labeled "Enable Edge Caching" is present and unchecked. The left sidebar shows "General" and "Edge Caching" tabs, with "Edge Caching" being the active one. At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

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 Synchronization Job](#).
10. Click **Finish**.
11. 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 an edge participant. For each edge participant, you want to assign the master participant that is the fastest and closest to it. This will increase the speed of rehydrating a stub file.

It is highly recommended that you assign, if possible, more than one master participant to an edge participant, so that if one master participant is not available, another master participant is able to provide the file content to the edge participant. You can set the order in which the master participants are consulted.

If you have only one edge participant and one master participant, the master participant is automatically assigned to the edge participant and listed in the Master Participants column. If you have multiple master participants, you need to explicitly assign master participants to each edge participant and then set the failover order.

1. Select an edge participant in the **Assignment** table.



2. Click the **Assign** button.

The **Assign Master Participants** dialog appears.

3. Select the master participants you want to assign to the edge participant.
4. (Optional) Change the failover order by selecting a master participant and using the **Move Up** and **Move Down** buttons.

- Click **OK**.

The **Master Participants** column has been updated for that participant.

- Repeat Steps 1-5 for each edge participant.
- 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 Support before modifying these values.

To modify these settings:

- Enter the values recommended by Peer Support.

The screenshot shows the 'Edit File Synchronization Job' dialog box with the 'General' tab selected. The left sidebar lists various settings categories, and the main area displays the following configuration options:

Setting	Value
Job ID	363
Job Type	File Synchronization
Job Name	FS-3-NFS-NetApp cDOT
Job Description	
Transfer Block Size (KB)	2048
Verify Block Checksums	☒
Verify Full File Checksums	☒
Enable Multipart Transfers	☐
Resume File Transfers	☐
Resume Transfer Minimum File Size (MB)	256
Synchronization Priority	2
Timeout (Seconds)	180
First Scan Mode	FOLDER_BY_FOLDER
Symbolic Link Strategy	Follow
Remove Filtered Files On Folder Delete	☒
Require All Hosts At Start	☐
Auto Start	☒
Auto Start Priority	1

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

Option	Description
Job ID	Unique, system-generated job identifier that cannot be edited.
Job Type	Identifies the job type. This cannot be modified.

Option	Description
Job Name	The name of this File Synchronization job. This name must be unique.
Job Description	Enter details or notes that help identify the purpose of the job. The description appears on the job's Configuration tab.
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 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.
Resume File Transfers	If selected, interrupted transfers of files that meet the minimum file size attempt to resume after a disconnection or job stoppage. If a transfer cannot be resumed, the entire file is transferred again. Requires Verify Full File Checksums to be selected. Resume file transfer is not used for delta-encoded transfers
Resume Transfer Minimum Size (MB)	The minimum file size, in megabytes, for which interrupted file transfers will attempt to resume.


Option	Description
Timeout (Seconds)	Number of seconds to wait for a response from any host before performing retry logic.
First Scan Mode	Determines which scan type will be used when the job is first started. For environments where most data 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.
Symbolic Link Strategy	Specifies how the job handles symbolic links. Select Copy to copy the symbolic link itself to destination hosts. Select Follow to follow the symbolic link and copy the target file or folder contents. The default strategy and the options available depend on the job protocol and participant configuration. For more information, including what happens when you change the strategy, see Handling Symbolic Links.
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.
Auto Start Priority	Specifies the order in which the job starts relative to other jobs configured to start automatically. Jobs with higher priority values start before jobs with lower priority values. The default is 1. Available only when Auto Start is selected.

- 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.



- If you want to create a new filter, modify an existing one, or update a filter, click **Edit File Scheduled Replication Filters**.

The **File and Folder Filters** dialog appears. You cannot edit predefined filters. See [Scheduled Replication Filters](#) in the [Preferences](#) section for information about creating or modifying a scheduled replication filter.



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 Synchronization Job** allows you to select options for resolving file conflicts and quarantines.

However, if you want to resolve the conflicts yourself, you can contact Peer Software to enable manual resolution. With manual resolution, you can select the host with the correct version of the file.

For more information about the cause of file conflicts, see [Conflicts, Retries, and Quarantines](#).

To modify conflict resolution settings for the File Synchronization job:

1. In the **File Conflict Resolution** section, select the **Truncate Milliseconds** option if you want the millisecond value truncated from each time stamp when comparing the time stamps of a file on two or more hosts.

As some storage platforms and applications track milliseconds slightly differently, selecting this option will prevent subtle millisecond differences from causing an otherwise in-sync file to be replicated or quarantined.

2. Select the **Advanced File Conflict Resolution** options you want applied:

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

Option	Description
	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):

Delta Block Transfer Size (KB):

Minimum File Size (KB):

Minimum File Size Percentage Target/Source:

Excluded File Extensions

zip
jpg
jpeg
png
gif
tiff
tif
Z
tgz
gz
gzip
rar
7z
bz
bz2
bzip2
mp3

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 enables you to configure how and when file metadata is synchronized between source and target hosts. Metadata is synchronized in real time and during scans.

Available options include specifying whether to synchronize file permissions (security metadata) and defining the metadata source (volume, share, export, or folder) to use for resolving conflicts during initial synchronization. This designated source, referred to as the [master host](#), determines which version of the metadata is retained when discrepancies occur.

For more information, see [File Metadata Synchronization](#) in [Advanced Topics](#).

This page also provides additional options not available during the job creation process.

The content of the File Metadata page depends on the file-sharing protocol you selected for your job:

- See [SMB Protocol](#) for details.
- See [NFS Protocol](#) for details.

- See [Multi-Protocol](#) for details.

SMB Protocol

To modify file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.

Note: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

Edit File Synchronization Job

File Metadata

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)

OK Cancel

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. (Optional) Enable prevention of corrupt or blank Owner or ACLs.

This ensures that invalid ownership or ACL entries on the source or master host are not applied to target hosts. The replication process skips or corrects these entries to maintain data integrity and security.

4. Select the permissions to be synchronized: **Owner**, **DACL**, and **SACL**.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

5. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):

- If no master host is selected, metadata synchronization will not occur during the initial scan.
- If permissions 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. (Optional) Enable enhanced metadata conflict resolution.

This prevents the Peer Agent service account from being assigned as the owner of a file or folder when a metadata conflict occurs and a file or folder is written to a target. If the Peer Agent service account is 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. The Peer Agent service account should be an actual user, as described in [Peer Global File Service - Environmental Requirements](#).

7. (Optional) In the **File Reparse Point Data Synchronization Options** section, configure reparse point handling.

Enabling these options synchronizes only a file's reparse data—not the actual offline content—to target hosts and prevents the offline file from being recalled from remote storage.

- **Reparse Tag Name** - Enter a single numerical value to specify which reparse point type to synchronize.
 - o Leave blank to disable reparse synchronization.
 - o Enter **0** to enable reparse point synchronization for all reparse file types.
 - o For Symantec Enterprise Vault, enter **16**.
 - o If you are unsure which value to use, contact Peer Support. You may also use **0** if you are certain only one vendor's reparse point functionality is in use.
- **Reparse Master Host** - Select a master host.

- o Leave blank to disable reparse synchronization. Files remain in their current state.
- o If selected, when last modified times and file sizes match across all hosts, but the file reparse attribute differs (for example, archived versus unarchived), the file reparse data is synchronized to match the file on the master host.
- o For Enterprise Vault, select the server where the archiving task runs.

Note: Use these options only if you are utilizing archiving or hierarchical storage solutions that rely on NTFS file reparse points to access data in a remote location (for example, Symantec's Enterprise Vault).

Note: These options are not available if the job is using Edge Caching.

8. (Optional) Enable transfer of file Alternate Data Stream (ADS).

- 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.
- ADS information is transferred only when a modification to the actual file itself is detected. ADS is not compared independently between participants.
- The updated file's ADS is applied to the corresponding files on target participants.


Note: The option is not available if the job is using Edge Caching.

9. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

NFS Protocol

To modify Linux-style file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **Real-time synchronization:** Metadata changes (such as ownership, permissions, and ACLs) are synchronized with the target host as they occur.
 - **Initial scan synchronization:** Metadata is synchronized during the initial scan.



2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. (Optional) Enable prevention of corrupt or blank Owner or ACLs.

This ensures that invalid ownership or ACL entries on the source or master host are not applied to target hosts. The replication process skips or corrects these entries to maintain data integrity and security

4. Select the permissions to be synchronized: **Owner**, **Group**, **Linux Permissions**, and **ACL**.

Note: Nutanix does not support Access Control Lists (ACL) in NFSv4 or POSIX formats.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts.

5. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - 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 Options**.
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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.

Multi-Protocol

To modify file metadata synchronization settings:

1. Select whether to synchronize metadata in real time and during the initial scan:
 - **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: Enable both synchronization options together or disable both. Using only one option is not supported. For more information, see [When should I synchronize permissions in real-time but not during the scan?](#)

2. If you enable real-time and initial scan synchronization, click **OK** when prompted.
3. (Optional) Enable prevention of corrupt or blank Owner or ACLs.

This ensures that invalid ownership or ACL entries on the source or master host are not applied to target hosts. The replication process skips or corrects these entries to maintain data integrity and security.

4. Select the security style to determine how file permissions are interpreted and synchronized across SMB and NFS environments:
 - **Windows-style permissions**
 - **Linux-style permissions**
 - **Mixed permissions** (available only with Dell PowerScale)

5. Select the permissions to be synchronized.

Note: To synchronize ownership or auditing information, the account under which the Peer Agent service runs must have sufficient privileges to read and modify ownership, permissions, and ACLs on all participating hosts

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**

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 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 Style: Windows-style permissions [How do permissions work with Multi-Protocol?](#)

- ☒ Owner
- ☒ DACL: Discretionary Access Control List

Validate

Metadata Conflict Resolution

Select a master host for initial scan: lab24ubuntu1

OK Cancel

- **Linux-style permissions**

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 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 Style: Linux-style permissions [How do permissions work with Multi-Protocol?](#)

- ☒ Owner
- ☒ Group
- ☒ Linux Permissions
- ☐ ACL (NFSv4 or POSIX)

Validate

Metadata Conflict Resolution

Select a master host for initial scan: lab24ubuntu1

OK Cancel

- **Mixed permissions** (available only with Dell PowerScale)

6. Select the [master host](#) for resolving metadata conflicts during the initial scan (if applicable):
 - If no master host is selected, metadata synchronization will not occur during the initial scan.
 - If permissions 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 addition to

Field	Description
	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 applied automatically to all watch sets for all participants in the job.

The **Application Support** page appears only when the **SMB** protocol is selected. This page provides options specific to SMB-based environments, such as support for Microsoft applications that maintain file locks or open handles. If you select **NFS** or **Multi-Protocol**, the page does not appear because these protocols do not use SMB-specific application support features.

The **Application Support** page lists file types for which Peer Software has established best practices. These best practices may include filtering recommendations, prioritization of specific file types, and configuration of file-locking behavior. However, if an application is not listed, this does not necessarily mean it is unsupported. For details about how a specific application is optimized, contact [Peer Support](#).

To configure application support options for a file synchronization job:

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 helps prevent accidental data loss by saving a backup copy of any file that is deleted or overwritten on a target host. When target protection is enabled for a job, the system moves a copy of the existing file to a hidden `.pc-trash_bin` folder on the target host—**before** applying changes from the source host.

The `.pc-trash_bin` folder is located at the root of the target host's watch set. Within this folder, the file is stored in the same relative directory structure as its original location in the watch set.

For instructions on how to restore a file from the trash bin, see [Restoring a File from the Trash Bin](#).

Enabling Target Protection

To enable and configure target protection for a job:

1. Configure the available fields 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

Target Protection

Enabled: ☒

of Backup Files to Keep: 3

of Days to Keep: 30

Trash Bin: .pc-trash_bin

OK Cancel

Field	Description
Enabled	Enables target protection for this job.
# of Backup Files to Keep	Specifies the maximum number of backup copies to retain for each file. Older copies will be purged once this limit is reached.
# of Days to Keep	Sets the number of days to retain backup copies in the trash bin. A value of 0 disables automatic purging based on age.
Trash Bin	Displays the name of the trash bin folder used to store backup files. This folder is located in the root of the target host's watch set and is hidden. The folder name is fixed and cannot be changed.

- Click **OK** or **Apply** to save your changes.

Restoring a File from the Trash Bin

To restore a deleted or overwritten file from the trash bin:

1. On the **target host**, open the watch set directory and locate the hidden folder named `.pc-trash_bin` at the root level.

Note: The `.pc-trash_bin` folder is hidden by default. Access may require administrative or root privileges, and hidden file visibility must be enabled on the system.

2. Navigate through the folder structure to locate the backup copy. The folder hierarchy mirrors the structure of the original watch set.
3. Copy the desired file from the `.pc-trash_bin` folder.
4. Paste the file back into its original location within the watch set.


The restored file is then synchronized to all other participants in the job.

Email Alerts

The **Email Alerts** page in the **Edit File Synchronization Job** dialog allows you to select which email alerts to apply to a File Synchronization job. Email alerts are defined in the [Preferences](#) dialog and can then be applied to individual jobs. See [Email Alerts](#) in the **Preferences** section for information about creating an email alert for a File Synchronization job.


To apply email alerts to a File Synchronization job while editing the job:

1. Click the **Select** button.




The **Select Email Alert** dialog opens.

2. Select the email alert from the drop-down list, and then click **OK**.



The newly added email alert appears in the **Email Alerts** table.



3. Repeat to add additional alerts to the job.
4. Click **OK** to close the **Edit** wizard or select another configuration item to modify.

SNMP Notifications


The **SNMP Notifications** page in the **Edit File Synchronization Job** dialog allows you to select which SNMP notification to apply to a File Synchronization job.

SNMP notifications, like email alerts and file filters, are configured at a global level in the [Preferences](#) dialog, then applied to individual jobs. For more information about 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 [Tagging Resources](#) in the [Advanced Topics](#) 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**, which allows you to edit multiple File Synchronization jobs at the same time. For example, you can change a single configuration item, such as **Verify Full File Checksum**, for multiple jobs in one operation instead of changing the item individually for each job.

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:



Edit File Synchronization Jobs - Multiple Selected

General

- File and Folder Filters
- Scheduled Replication Filters
- Delta Replication
- File Metadata
- File Locking
- Target Protection
- Email Alerts
- SNMP Notifications

General

Job ID: MULTIPLE SELECTED

Job Type: MULTIPLE SELECTED

Transfer Block Size (KB): 2048

Verify Block Checksums: ☒

Verify Full File Checksums: ☒

Enable Multipart Transfers: ☐

Resume File Transfers: ☐

Resume Transfer Minimum File Size (MB): 256

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

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:



Verify Checksum

Select a value to apply to all selected configurations:

☒ false (from: Training)

☐ true (from: Media Assets)

OK Cancel

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.

Job Initialization Process

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).

Initial Synchronization Process

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 154
- Active Directory authentication 438
- Active Directory group 429
- Active Directory groups 438, 440, 442, 443
- Active Directory user 111, 429
- Active Directory users 438, 440, 442, 443
- Active-passive configuration
 - automatic failover 135
 - Network of Brokers 135
- Advanced configuration options, Amazon FSxN 375
- Advanced configuration options, Dell PowerScale 382
- Advanced configuration options, Dell PowerStore 388
- Advanced configuration options, Dell Unity 393
- Advanced configuration options, NetApp ONTAP 400
- Advanced configuration options, Nutanix Files 404
- Agent availability 340
- Agent configuration, Broker 142
- Agent configuration, editing 139
- Agent configuration, General 144
- Agent configuration, Logging 147
- Agent configuration, Performance 147
- Agent connectivity, preferences 340
- Agent location, Proactive Monitoring 412
- Agent properties, editing 152
- Agent properties, viewing 149
- Agent Summary view 67
- Agent, destination 746
- Agent, disabled 138
- Agent, disconnected 340
- Agent, installation 15
- Agent, re-enabling 138
- Agent, source 716
- Agent, update 19, 21
- Agent, updating 153
- Agent, VM options 151
- Agents
 - database connection 454
 - Management Agent 454
 - proxy configuration 456
- Agents view 42, 45
- Alert 358
- Alerts 104
 - File System Analytics transfer status 122
- Alerts view 43
- All Files
 - source paths 477
- Amazon FSx for NetApp ONTAP
 - Cloud Backup and Replication credentials 463
 - credentials 607, 720, 813
- Amazon FSxN
 - multi-protocol 132
- Amazon FSxN prerequisites 7
- Amazon FSxN, advanced configuration options 375
- Amazon FSxN, configuration 371, 396
- Amazon FSxN, credentials 371, 396
- Amazon S3 483, 487
 - bucket details 488
 - Cloud Backup and Replication 451
 - Cloud Backup and Replication destination 479
 - credentials 480
- Amazon S3, bucket details 490
- Analytics 118
 - File Activity Analytics 118
 - File System Analytics 118
 - licensing 118
 - PeerIQ 118
 - preferences 118
- Analytics preferences
 - configuring 118
 - enabling 118
 - PeerIQ 119
- Analytics, preferences 220, 221
- Anonymous Diagnostic Data 216, 407
- Anonymous Diagnostic Data, customer-identifiable information 216
- Anonymous Diagnostic Data, product usage and configuration 216
- API 154
- API access 154
- API categories 156
- API integration 155
- API reference 156
- API status codes 156
- Application support, File Collaboration job 655, 696
- Application support, File Replication job 797
- Application support, File Synchronization job 909
- Assigning tags to web roles 450

Authentication 182
 Auto match root 588
 Auto-restarting, File Synchronization job 924
 Azure Blob Storage 481
 container details 488
 credentials 480
 Azure Storage, container details 489

- B -

B 190
 Backup PMC
 automatic failover 135
 Bait file 358
 Base role 112
 Base web role 113
 Baseline replication
 Scan at Start 499, 501
 Bash 155
 Basic concepts 81
 Batched emails 222, 293
 Batched real-time
 Cloud Backup and Replication 498
 Cloud Backup and Replication jobs 501
 destination snapshots 501
 real-time monitoring 498
 replication frequency 501
 Scan at Start 501
 scheduled replication 498
 Broker configuration, preferences 343
 Broker, Agent configuration 142
 Browsing files and folders 338
 Bucket Details
 Cloud Backup and Replication jobs 488
 existing bucket 488
 new bucket 488
 Bucket details, Amazon S3 490
 Bucket details, NetApp StorageGRID 492
 Bucket details, Nutanix Objects 494
 Bucket details, Wasabi 496
 Buckets
 Cloud Backup and Replication 488
 creating 488
 selecting existing 488

- C -

Capacity consumption
 unique TBs 160
 Cloud Backup and Replication
 file-to-object replication 451
 object storage 451
 overview 451
 point-in-time recovery 451
 real-time replication 451
 Cloud Backup and Replication job
 delta-level replication 505
 file metadata 505
 miscellaneous options, 505
 NTFS permission replication 505
 recovering data 518
 rehydrated data availability 505
 runtime view 51
 storage tiering options 505
 Cloud Backup and Replication job, creating 451
 Cloud Backup and Replication job, deleting 513
 Cloud Backup and Replication job, email alerts 508
 Cloud Backup and Replication job, monitoring 517
 Cloud Backup and Replication job, running 514
 Cloud Backup and Replication job, SNMP notifications 510
 Cloud Backup and Replication job, starting 514
 Cloud Backup and Replication job, stopping 516
 Cloud Backup and Replication jobs 450
 batched real-time 501
 bucket details 488
 container details 488
 Continuous Data Protection 502
 daily replication 499
 destination 479
 destination credentials 480
 destination snapshot retention 503
 destination statistics 67
 destination storage platform 479
 job status 67
 Management Agent 454
 object storage 451
 overview 451
 proxy configuration 456
 real-time replication 498, 502
 Replication and Retention policy 497
 replication schedule 498

- Cloud Backup and Replication jobs 450
 - retention 503
 - scheduled replication 498, 501
 - scheduled scans 499
 - source credentials 463
 - source paths 477
 - source snapshots 504
 - source storage platform 454, 463
 - storage information 463
 - storage platform types 454
 - summary view 67
 - tasks 67
 - volume status 67
 - watch set 477
 - weekly replication 499
- Cloud Backup and Replication jobs, preferences 228
- Cloud Backup and Replication preferences
 - proxy configuration 456
- Cloud services
 - Cloud Backup and Replication 451
- Cloud Summary
 - Destination Statistics tab 67
 - Job Summary tab 67
 - opening 67
 - startup display 67
 - Tasks tab 67
 - Volume Summary tab 67
- Collab, Sync, and Repl Summary
 - Edge Caching tab 68
 - opening 68
 - reports tab 68
 - startup display 68
 - summary tab 68
- Collab, Sync, and Repl Summary view, Edge Caching tab 72
- Collab, Sync, and Repl Summary view, Reports tab 72
- Collab, Sync, and Repl Summary view, Summary tab 69
- Complex regular expressions 90
- Concepts 81
- Configuration 7
 - Configuration, Agent 139
 - Configuration, Amazon FSxN 371, 396
 - Configuration, Dell 376
 - Configuration, Dell PowerScale 379
 - Configuration, Dell PowerStore 386
 - Configuration, NetApp ONTAP 396
- Configuration, Nutanix 401
- Configuration, PMC 8
- Configuration, Preferences 219
- Configuration, Unity 390
- Conflict resolution, File Collaboration job 682
- Conflict resolution, file metadata 650, 686, 687, 754, 783, 862, 900, 901
- Conflict resolution, File Replication job 779
- Conflict resolution, File Synchronization job 896
- Conflict resolution, metadata 651, 689, 754, 784, 862, 903, 905
- Connection statuses, Peer Agents 138
- Consistent backups
 - source snapshots 504
- Container Details
 - Cloud Backup and Replication jobs 488
 - existing container 488
 - new container 488
- Container details, Azure Storage 489
- Continuous Data Protection
 - Cloud Backup and Replication 498
 - Cloud Backup and Replication jobs 502
 - destination snapshots 502
 - Processing Delay 502
 - real-time monitoring 498
 - real-time replication 498, 502
- Create File Collaboration job 588
- Create file filters 83
- Create File Synchronization job 588
- Create folder filters 83
- Credentials 234
 - Cloud Backup and Replication 463
 - Cloud Backup and Replication jobs 480
 - destination storage 480
 - existing 463, 607, 720, 813
 - existing credentials 480
 - new 463, 607, 720, 813
 - new credentials 480
 - SMB access 134
 - validating 463, 607, 720, 813
- Credentials, Amazon FSxN 371, 396
- Credentials, Dell PowerScale 379
- Credentials, Dell PowerStore 386
- Credentials, Dell Unity 390
- Credentials, NetApp ONTAP 396
- Credentials, Nutanix 401
- Custom web role 115
- Customer-identifiable information 216

- D -

- Daily scans
 - destination snapshots 499
 - scheduled replication 499
- Dashboard view 45
- Data aging
 - File System Analytics 121
- Data consistency
 - File Collaboration 600
- Data integrity
 - File Synchronization 805
- Data recovery
 - Cloud Backup and Replication job 518
 - destination snapshots 497
- Data sharing 407
- Data uploads 407
- Database connections 231
 - Management Agent 454
- Delayed replication 158
- Dell configurations 376
- Dell PowerScale
 - Cloud Backup and Replication credentials 463
 - credentials 607, 720, 813
 - multi-protocol 132
- Dell PowerScale, configuration 379
- Dell PowerScale, credentials 379
- Dell PowerSpace prerequisites 7
- Dell PowerStore
 - credentials 607, 720, 813
 - multi-protocol 132
- Dell PowerStoreconfiguration 386
- Dell PowerStorecredentials 386
- Dell Unity
 - Cloud Backup and Replication credentials 463
 - credentials 607, 720, 813
 - multi-protocol 132
- Dell Unity prerequisites 7
- Dell Unity, configuration 390
- Dell Unity, credentials 390
- Delta-level replication
 - Cloud Backup and Replication job 505
- Delta-level replication, File Collaboration job 684
- Delta-level replication, File Replication job 781
- Delta-level replication, File Synchronization job 898
- Destination
 - Cloud Backup and Replication jobs 479
 - storage platform 479
- Destination Agent 746
- Destination Agents
 - File Replication 714
- Destination credentials 481, 483, 484, 486, 487
 - Cloud Backup and Replication jobs 480
- Destination path 747, 836
- Destination snapshot report 236
- Destination snapshots
 - after replication 501
 - batched real-time 501
 - Cloud Backup and Replication 497
 - Continuous Data Protection 502
 - daily retention 503
 - daily schedule 499
 - data recovery 497
 - monthly retention 503
 - retention 497, 503
 - scheduled scans 499
 - source snapshots 504
 - weekly retention 503
 - weekly schedule 499
 - yearly retention 503
- Destination statistics
 - Cloud Backup and Replication 67
- Destination storage
 - credentials 480
- Destination storage credentials 234
- Destination storage platform
 - Cloud Backup and Replication jobs 479
- Detector settings 588
- DFS namespace 588
- DFS namespace failback
 - automatic 124, 127
 - re-enabling folder targets 127
 - rescanning 127
- DFS namespace failover
 - automatic 124, 127
 - disabling folder targets 127
- DFS namespace folder, adding 576
- DFS namespace folders
 - folder targets 124
 - overview 124
- DFS namespace root path 319, 320, 541, 571, 588
- DFS namespace roots
 - overview 124
- DFS namespace server, adding 571
- DFS namespace servers

- DFS namespace servers
 - overview 124
- DFS namespace, failover and failback 262
- DFS namespace, File Collaboration job 659
- DFS namespace, File Synchronization job 871
- DFS namespaces
 - domain-based namespaces 124
 - failback 124, 127
 - failover 124, 127
 - folder targets 124
 - importing 126
 - linking to jobs 126
 - namespace folders 124
 - namespace roots 124
 - namespace servers 124
 - overview 124
 - PeerGFS management 124
 - using with jobs 126
- DFS Namespaces service 541, 571
- DFS namespaces, connecting to jobs 588
- DFS namespaces, managing 571
- DFS preferences
 - namespaces 126
- DFS-N Link, File Collaboration job 703
- DFS-N Link, File Synchronization job 917
- DFS-N Management job
 - running 568
 - runtime view 52
- DFS-N Management job, creating 535
- DFS-N Management job, folder targets 545
- DFS-N Management job, importing existing namespace 557
- DFS-N Management job, Management Agent 538
- DFS-N Management job, namespace folders 545
- DFS-N Management job, namespace name 540
- DFS-N Management job, namespace server 541
- DFS-N Management job, namespace settings 543
- DFS-N Management job, starting 568
- DFS-N Management job, stopping 570
- DFS-N Management job, verifying Agent 539
- DFS-N Management jobs 534
 - creating 126
 - creating namespaces 124
 - DFS namespaces 124
 - importing namespaces 126
 - linking namespaces 126
 - managing namespaces 124
 - status 76

- summary view 76
- DFS-N Management jobs, preferences 319, 320
- DFS-N Management preferences
 - automatic failback 127
 - automatic failover 127
- Diagnostic data 407
- Diagnostic data, anonymous 216
- Disabled agent 138
- Disconnected agent 340
- Distributed file locking
 - File Collaboration 600
 - version conflicts 600
- Domain-based namespaces
 - PeerGFS 124

- E -

- Edge Caching tab, Collab, Sync, and Repl Summary view 72
- Edge caching, File Collaboration job 633
- Edge caching, File Synchronization job 842
- Edge caching, master-edge assignment 647, 676, 858
- Editing multiple File Collaboration jobs 704
- Editing multiple File Replication jobs 803
- Editing multiple File Synchronization jobs 918
- Email alerts, Cloud Backup and Replication job 508
- Email alerts, concepts 82
- Email alerts, DFS-N Management job 550
- Email alerts, File Collaboration job 657, 698
- Email alerts, File Replication job 759, 799
- Email alerts, File Synchronization job 868, 912
- Email alerts, preferences 222, 236, 293, 343
- Email configuration preferences 328
- Enhanced metadata conflict resolution 650, 686, 687, 689, 754, 783, 784, 862, 900, 901, 903, 905
- Environment name 338
- Environment name, Telemetry 408
- Environmental requirements 7
- Event detection 366
- Expedited Sync Queue 259
- Exports
 - multi-protocol 132

- F -

- FAA 118
- Failover

- Failover
 - Network of Brokers 135
 - Peer Management Center 135
- File Activity Analytics 118
 - historical analysis 122
 - licensing 118
 - NFS 122
 - PeerIQ 118, 122
 - real-time data 118
 - real-time statistics 122
 - reporting 122
 - SMB 122
 - user activity 118, 122
- File Activity Analyzer
 - File Activity Analytics 122
- File age
 - PeerIQ 119
- File and folder filters, Cloud Backup and Replication job 478
- File and folder filters, File Collaboration job 680
- File and folder filters, File Replication job 776
- File and folder filters, File Synchronization job 893
- File and folder filters, preferences 243, 299
- File Collaboration
 - data consistency 600
 - distributed environment 600
 - distributed file locking 600
 - local access 600
 - overview 600
 - real-time synchronization 600
 - WAN bandwidth 600
- File Collaboration job
 - editing 664
 - running 705
 - runtime view 53
- File Collaboration job, application support 655, 696
- File Collaboration job, auto-restarting 709
- File Collaboration job, conflict resolution 682
- File Collaboration job, create job from namespace folder 588
- File Collaboration job, creating 601
- File Collaboration job, delta-level replication 684
- File Collaboration job, DFS namespace 659
- File Collaboration job, DFS-N Link 703
- File Collaboration job, edge caching 633
- File Collaboration job, editing participants 673
- File Collaboration job, email alerts 657, 698
- File Collaboration job, file and folder filters 680
- File Collaboration job, file locking 692
- File Collaboration job, file metadata 650, 651, 652, 686, 687, 689
- File Collaboration job, general 678
- File Collaboration job, host connectivity issues 711
- File Collaboration job, Management Agent 604
- File Collaboration job, manual retries 713
- File Collaboration job, master data service 635
- File Collaboration job, participants 603, 666
- File Collaboration job, pinning filter 643
- File Collaboration job, quarantined file 712
- File Collaboration job, running 706
- File Collaboration job, runtime view, Alerts tab 60
- File Collaboration job, runtime view, Configuration tab 63
- File Collaboration job, runtime view, Event Log tab 57
- File Collaboration job, runtime view, Participants tab 61
- File Collaboration job, runtime view, Quarantines tab 58
- File Collaboration job, runtime view, Retries tab 59
- File Collaboration job, runtime view, Session tab 56
- File Collaboration job, runtime view, Summary tab 54
- File Collaboration job, SNMP notifications 701
- File Collaboration job, starting 707
- File Collaboration job, stopping 709
- File Collaboration job, tags 702
- File Collaboration job, target protection 697
- File Collaboration job, utilization policy 641, 851
- File Collaboration job, volume policy 638
- File Collaboration jobs 600
 - DFS namespaces 124, 126
 - distributed locking 600
 - linking namespaces 126
 - multi-protocol 132
 - overview 600
 - real-time synchronization 600
 - storage information 607, 813
 - storage platform 605
 - storage platform credentials 607, 813
 - storage platform types 605
 - summary view 68
 - version conflicts 600
- File Collaboration jobs, editing multiple jobs 704
- File Collaboration jobs, preferences 257
- File extensions 358
 - File System Analytics 121

- File filters, concepts 82
- File filters, create 83
- File filters, predefined 84
- File filters, usage notes 97
- File locking 302
 - File Synchronization 805
 - temporary locking 805
- File locking, File Collaboration job 692
- File locking, File Replication job 793
- File locking, File Synchronization job 907
- File menu 29
- File metadata
 - Cloud Backup and Replication job 505
- File metadata, conflict resolution 651, 687, 754, 784, 862, 901
- File metadata, File Collaboration job 650, 651, 652, 686, 687, 689
- File metadata, File Replication job 754, 756, 757, 783, 784, 786, 788
- File metadata, File Synchronization job 862, 864, 865, 900, 901, 903, 905
- File permissions
 - multi-protocol 132
- File Replication
 - one-way replication 714
 - overview 714
 - source storage device 714
 - target storage devices 714
- File Replication job
 - editing 763
 - runtime view 53, 64
- File Replication job, application support 797
- File Replication job, conflict resolution 779
- File Replication job, creating 714
- File Replication job, delta-level replication 781
- File Replication job, editing participants 771
- File Replication job, email alerts 759, 799
- File Replication job, file and folder filters 776
- File Replication job, file locking 793
- File Replication job, file metadata 754, 756, 757, 783, 784, 786, 788
- File Replication job, General 773
- File Replication job, job relay 790
- File Replication job, participants 765
- File Replication job, scheduled replication filters 778
- File Replication job, SNMP notifications 802
- File Replication job, tags 803
- File Replication job, target protection 798
- File Replication jobs 714
- multi-protocol 132
- overview 714
- source storage device 714
- storage information 720
- storage platform 718
- storage platform credentials 720
- storage platform types 718
- summary view 68
- target storage devices 714
- File Replication jobs, editing multiple 803
- File Replication jobs, preferences 257
- File retries 298
- File Retries and Source Snapshots, preferences 240
- File Synchronization
 - multi-directional synchronization 805
 - non-collaborative workloads 805
 - overview 805
 - real-time synchronization 805
 - temporary locking 805
 - user home directories 805
 - user profiles 805
- File Synchronization job
 - editing 875
 - runtime view 53, 65
- File Synchronization job, application support 909
- File Synchronization job, auto-restarting 924
- File Synchronization job, conflict resolution 896
- File Synchronization job, create job from namespace folder 588
- File Synchronization job, creating 806
- File Synchronization job, delta-level replication 898
- File Synchronization job, DFS namespace 871
- File Synchronization job, DFS-N Link 917
- File Synchronization job, edge caching 842
- File Synchronization job, editing participants 884
- File Synchronization job, email alerts 868, 912
- File Synchronization job, file and folder filters 893
- File Synchronization job, file locking 907
- File Synchronization job, file metadata 862, 864, 865, 900, 901, 903, 905
- File Synchronization job, General 890
- File Synchronization job, host connectivity issues 926
- File Synchronization job, Management Agent 809
- File Synchronization job, Master Data Service 845
- File Synchronization job, participants 808, 877
- File Synchronization job, pinning filter 853
- File Synchronization job, running 920

- File Synchronization job, SNMP notifications 915
 - File Synchronization job, starting 922
 - File Synchronization job, stopping 924
 - File Synchronization job, tags 916
 - File Synchronization job, target protection 910
 - File Synchronization job, volume policy 848
 - File Synchronization jobs 805
 - DFS namespaces 124, 126
 - linking namespaces 126
 - multi-protocol 132
 - overview 805
 - participants 805
 - real-time synchronization 805
 - storage platform 810
 - storage platform types 810
 - summary view 68
 - temporary locking 805
 - File Synchronization jobs, editing multiple 918
 - File Synchronization jobs, preferences 257
 - File System Analytics 118
 - data access patterns 118
 - data age 118
 - data aging 121
 - file extensions 121
 - file sizes 118
 - file system data 121
 - file types 118, 121
 - historical analysis 121
 - licensing 118
 - NFS 121
 - PeerIQ 118, 121
 - reporting 121
 - scan process 122
 - scan results 122
 - scan schedule 121
 - scan timing 122
 - scheduled scans 121
 - SMB 121
 - File System Analytics scans
 - duration 122
 - file system traversal 122
 - metadata collection 122
 - results 122
 - schedule 122
 - File System Analyzer
 - File System Analytics 121
 - File system metadata
 - File System Analytics 122
 - File types
 - File System Analytics 121
 - PeerIQ 119
 - Files
 - source paths 477
 - Files and folders, browsing 338
 - Filter expressions 99
 - Filter patterns 86
 - Filter patterns, complex regular expressions 90
 - Filter patterns, excluding temporary files 88
 - Filter patterns, using wildcards 87
 - Filter, scheduled replication 158
 - Filtering by date 91
 - Filtering by file size 93
 - Filtering folders 95
 - Filters, file 97
 - Filters, folders 95
 - Filters, list 98
 - Folder filters, concepts 82
 - Folder filters, create 83
 - Folder target, adding 583
 - Folder targets
 - automatic disabling 127
 - automatic enabling 127
 - DFS namespaces 124
 - failback 124, 127
 - failover 124, 127
 - linking to watch sets 126
 - referrals 124
 - Folders
 - source paths 477
 - Folders, filtering 95
 - FSA 118
- ## - G -
- General Configuration, preferences 337, 338
 - General, Agent configuration 144
 - General, File Collaboration job 678
 - General, File Replication job 773
 - General, File Synchronization job 890
- ## - H -
- Health and performance monitoring 407
 - Health Checker, setup 413
 - Heartbeats 340

Help menu 32
 Hidden items 338
 High availability
 Peer Management Center 135

Historical data
 File Activity Analytics 122
 File System Analytics 121
 PeerIQ 119

Hot data
 File Collaboration 600

Hyper-V
 PeerIQ 119

- I -

Importing existing namespace 557
 Initial scan
 batched real-time replication 498
 Initial synchronization process 706, 921
 Initialization process 706, 921
 Installation 7
 Installation, Peer Agent 15
 Installation, PMC 8
 Internal user 111, 429
 Internal users 429, 431, 432, 435, 437

- J -

Job
 File Collaboration 600
 Job alert 358
 Job Alerts view 47
 Job health
 Cloud Backup and Replication 67
 Collab, Sync, and Repl Summary 68
 Job initialization process 706, 921
 Job relay
 multi-protocol not supported 132
 Job relay, File Replication job 790
 Job, manual retries 927
 Jobs
 DFS-N Management 534
 File Replication 714
 File Synchronization 805
 Jobs view
 Cloud Summary 67
 Collab, Sync, and Repl Summary 68

Namespace Summary 76
 Jobs views 48
 Jobs, Cloud Backup and Replication 450

- L -

LAN
 File Collaboration 600
 local access 600
 Last modified date 91
 LDAP administrator 438
 LDAP settings 429
 License file 354
 License utilization
 PeerIQ 119
 Licensed storage capacity
 alerts 160
 PeerGFS 160
 usage 160
 Licensing 354
 analytics 118
 storage capacity 160
 Licensing preferences 354
 Link namespace folder to job 595
 Link namespace to job 588
 Linux
 multi-protocol 132
 Peer Management Center rich client not supported 23
 Linux Agents
 SMB access 134
 Linux File Server
 storage information 720, 813
 List filter expressions 99
 List filters, concepts 98
 List filters, examples 99
 List filters, managing 101
 List filters, removing 101
 List filters, saving 101
 Local access
 File Collaboration 600
 Locked files
 source snapshots 504
 Locking 302
 Log files 106
 Logging 104
 Logging, Agent configuration 147

- M -

- Malicious Event Detection 358
- Malware 358
- Management Agent
 - Cloud Backup and Replication jobs 454
 - database connection 454
 - multiple Agents 454
 - proxy configuration 456
 - selecting 454
 - storage devices 454
- Management Agent, DFS-N Management job 538
- Management Agent, File Collaboration job 604
- Management Agent, File Synchronization job 809
- Management Status
 - DFS-N Management jobs 76
- Managing web roles 443
- Master Data Service, File Collaboration job 635
- Master Data Service, File Synchronization job 845
- Master host 650, 651, 652, 686, 687, 754, 756, 757, 783, 784, 786, 788, 862, 864, 865, 900, 901
- Master-edge assignment, edge caching 647, 676, 858
- MED 358
- MED configuration, preferences 358
- Menus 28
- Menus, File 29
- Menus, Help 32
- Menus, Tools 32
- Menus, Window 30
- Metadata conflict resolution 650, 651, 686, 687, 689, 754, 783, 862, 900, 901, 903, 905
- Metadata synchronization 650, 651, 652, 686, 687, 754, 756, 757, 783, 784, 786, 788, 862, 864, 865, 900, 901
- Microsoft Azure
 - Cloud Backup and Replication 451
 - Cloud Backup and Replication destination 479
- Microsoft Azure, telemetry storage 216
- Miscellaneous options
 - Cloud Backup and Replication job 505
- Mixed permissions
 - SMB access 134
- Multi-job editing, File Collaboration jobs 704
- Multi-job editing, File Replication jobs 803
- Multi-job editing, File Synchronization jobs 918
- Multi-Protocol

- architecture 132
- File Collaboration storage platforms 605
- File Replication storage platforms 718
- File Synchronization storage platforms 810
- limitations 132
- SMB/NFS 132
- supported job types 132
- system requirements 132
- Multi-protocol jobs
 - SMB access 134

- N -

- Namespace folder 588
- Namespace folder target, adding 583
- Namespace folder target, disable 262
- Namespace folder target, renewable 262
- Namespace folder, adding 576
- Namespace folder, linking to job 595
- Namespace folders
 - linking to watch sets 126
- Namespace referrals
 - folder targets 124
- Namespace server 541
- Namespace server, adding 571
- Namespace Summary
 - importing namespaces 126
 - management status 76
 - namespace state 76
 - opening 76
 - startup display 76
- Namespace, linked to job 588
- Namespaces
 - Not Found 76
 - Offline 76
 - Online 76
 - state 76
 - status 76
 - Unknown 76
- NAS configuration, Amazon FSxN 371, 396
- NAS configuration, Dell 376
- NAS configuration, Dell PowerScale 379
- NAS configuration, Dell PowerStore 386
- NAS configuration, Dell Unity 390
- NAS configuration, NetApp ONTAP 396
- NAS configuration, Nutanix 401
- NetApp ONTAP
 - Cloud Backup and Replication credentials 463

NetApp ONTAP
 credentials 607, 720, 813
 multi-protocol 132

NetApp ONTAP prerequisites 7

NetApp ONTAP, advanced configuration options 400

NetApp ONTAP, configuration 396

NetApp ONTAP, credentials 396

NetApp StorageGRID 484
 bucket details 488
 Cloud Backup and Replication destination 479
 credentials 480

NetApp StorageGRID, bucket details 492

Network of Brokers
 active-passive configuration 135
 automatic failover 135
 multiple brokers 135
 overview 135
 replication traffic 135

NFS
 exports 132
 File Activity Analytics 122
 File Replication storage platforms 718
 File Synchronization storage platforms 810
 File System Analytics 121
 multi-protocol 132
 NFSv3 132
 NFSv4 132

NTFS permission replication
 Cloud Backup and Replication job 505

NTFS permissions 650, 651, 652, 686, 687, 689, 754, 757, 783, 784, 788, 862, 865, 900, 901, 905

Nutanix AHV
 PeerIQ 119

Nutanix Files
 Cloud Backup and Replication credentials 463
 credentials 607, 720, 813
 multi-protocol 132
 SMB access 134

Nutanix Files prerequisites 7

Nutanix Files, advanced configuration options 404

Nutanix Objects 486
 bucket details 488
 Cloud Backup and Replication destination 479
 credentials 480

Nutanix Objects, bucket details 494

Nutanix, configuration 401

Nutanix, credentials 401

- O -

Object storage
 Cloud Backup and Replication 451

On-premises data
 Cloud Backup and Replication 451

Open files
 source snapshots 504

OVA
 PeerIQ 119

- P -

Participants
 Agents 805
 File Synchronization 805
 storage capacity 160
 storage paths 805

Participants, adding to File Collaboration job 667

Participants, adding to File Replication job 765

Participants, adding to File Synchronization job 878

Participants, editing 673, 771, 884

Participants, File Collaboration job 603, 666

Participants, File Replication job 765

Participants, File Synchronization job 808, 877

Participants, removing from File Collaboration job 667

Participants, removing from File Replication job 765

Participants, removing from File Synchronization job 878

Path, destination 747, 836

Path, source 741

Peer Agent location, Proactive Monitoring 412

Peer Agent Service
 Management Agent 454
 troubleshooting 454

Peer Agent, configuration 139

Peer Agent, connection to Peer Management Broker 138

Peer Agent, installation 15

Peer Agent, properties 149

Peer Agent, update 19, 21

Peer Agents
 analytics data 118
 destination Agents 714
 DFS namespace failback 127
 DFS namespace failover 127

- Peer Agents
 - File Replication 714
 - File Synchronization 805
 - File System Analytics scans 122
 - Management Agent 454
 - PeerIQ monitoring 119
 - scan results 122
 - source Agent 714
 - storage device access 454
 - zip file creation 122
- Peer Agents connection statuses 138
- Peer Agents, managing 135
- Peer Global File Service license 354
- Peer Management Broker 666, 765, 877
 - Agent connection 454
 - Network of Brokers 135
- Peer Management Center
 - DFS namespace failback 127
 - DFS namespace failover 127
 - DFS namespace management 124
 - failover 135
 - File System Analytics scan results 122
 - Network of Brokers 135
 - rich client 23
 - running 23
 - web client 23
 - zip file transfer 122
- Peer Management Center, update 19
- Peer services 26
- PeerGFS API 154
- PeerGFS environment, health and performance 216
- PeerGFS license 354
- PeerGFS preferences 217
- PeerIQ 118
 - Agent information 119
 - analytics 118
 - capacity planning 119
 - dashboards 119
 - deployment 119
 - File Activity Analytics 119, 122
 - File Activity Analytics reports 122
 - File System Analytics 119, 121
 - File System Analytics reports 121
 - File System Analytics scan results 122
 - historical data 119
 - job information 119
 - license utilization 119
 - licensing 118
 - monitoring 118, 119
 - operational insights 118
 - performance metrics 118
 - PMC information 119
 - setup 119
 - volume information 119
 - web interface 119
- Performance Metrics 118
- Performance, Agent configuration 147
- Performance, preferences 246, 304
- Permissions 112
- Permissions, base web role 113
- Permissions, custom web role 115
- Pinning filter, File Collaboration job 643
- Pinning filter, File Synchronization job 853
- PMC configuration 8
- PMC installation 8
- PMC user interface 27, 40
- PMC, update 19
- Point-in-time recovery
 - Cloud Backup and Replication 451
- POSIX permissions 650, 652, 686, 689, 754, 756, 757, 783, 786, 788, 862, 864, 865, 900, 903, 905
- PowerScale, advanced configuration options 382
- PowerScale, configuration 379
- PowerScale, credentials 379
- Powershell 155
- PowerStoreadvanced configuration options 388
- PowerStoreconfiguration 386
- PowerStorecredentials 386
- Predefined file filters 84
- Preferences 217
 - SMB access 134
- Preferences, Agent connectivity 340
- Preferences, Analytics 220, 221
- Preferences, Broker configuration 343
- Preferences, Cloud Backup and Replication 228
- Preferences, configuring 219
- Preferences, database connections 231
- Preferences, DFS-N Management jobs 319, 320
- Preferences, email alerts 222, 236, 293
- Preferences, Email configuration 328
- Preferences, file and folder filters 243, 299
- Preferences, File Collaboration jobs 257
- Preferences, File Replication jobs 257
- Preferences, File Retries and Source Snapshots 240
- Preferences, File Synchronization jobs 257

- Preferences, General configuration 337, 338
- Preferences, licensing 354
- Preferences, MED configuration 358
- Preferences, performance 246, 304
- Preferences, real-time event detection 259, 306, 366
- Preferences, replication and retention policies 250
- Preferences, Scan Manager 255, 312
- Preferences, SNMP configuration 405
- Preferences, SNMP Notifications 226, 252, 276, 308
- Preferences, Software updates 349
- Preferences, storage platform configuration 370
- Preferences, system alerts 343
- Preferences, tags 351
- Preferences, Telemetry 407
- Preferences, User management 429
- Prerequisites, Dell PowerSpace 7
- Prerequisites, Dell Unity 7
- Prerequisites, NetAPP ONTAP 7
- Prerequisites, Nutanix Files 7
- Pre-Seeding 159
- Primary PMC
 - automatic failover 135
- Proactive Monitoring 216, 407, 426
- Proactive Monitoring, collected data 216
- Proactive Monitoring, eligibility 216
- Proactive Monitoring, environment health and performance 216
- Proactive Monitoring, Health Checker 413
- Proactive Monitoring, Microsoft Azure storage 216
- Proactive Monitoring, setup 408
- Proactive Monitoring, Technical Account Managers 216
- Proactive Monitoring, telemetry options 409
- Processing Delay
 - Continuous Data Protection 502
 - WAN usage 502
- Properties, Agent 149, 152
- Proxy authentication
 - domain 456
 - password 456
 - username 456
- Proxy Configuration
 - Cloud Backup and Replication jobs 456
 - creating 456
 - editing 456
 - existing configuration 456
 - user authentication 456

- validating 456
- Proxy servers
 - authentication 456
 - Cloud Backup and Replication 456
 - IP address 456
 - port 456
- PST files
 - source snapshots 504
- Purge all versions between snapshots
 - retention 503

- Q -

- QCOW2
 - PeerIQ 119
- Quarantined file 222, 293, 926
- Quarantined file, File Collaboration job 712

- R -

- Real-time activity
 - source Agent 714
- Real-time event detection 306
- Real-time event detection, preferences 259, 306, 366
- Real-time replication 158
 - Cloud Backup and Replication 497
 - Continuous Data Protection 502
- Real-time statistics
 - File Activity Analytics 122
- Reconnect attempts 340
- Recovering data
 - Cloud Backup and Replication job 518
- Re-enable agent 138
- Rehydrated data availability
 - Cloud Backup and Replication job 505
- Removing file from quarantine 712, 926
- Replication and Retention Policies, preferences 250
- Replication and Retention policy
 - Cloud Backup and Replication jobs 497
 - destination snapshots 497
 - existing policy 497
 - new policy 497
 - real-time replication 497
 - replication schedule 497
 - snapshot retention 497
 - source snapshots 497

- Replication frequency
 - batched real-time 501
- Replication schedule
 - batched real-time 498
 - Cloud Backup and Replication 497
 - Cloud Backup and Replication jobs 498
 - Continuous Data Protection 498
 - retention options 503
 - scheduled scans 498
- Replication traffic
 - Network of Brokers 135
- Reports tab, Collab, Sync, and Repl Summary view 72
- Reports, destination snapshot 236
- Reports, scan 236
- Requirements, environmental 7
- Rescans
 - DFS namespace failback 127
- Retention
 - Cloud Backup and Replication jobs 503
 - daily retention 503
 - destination snapshots 503
 - monthly retention 503
 - purging versions 503
 - weekly retention 503
 - yearly retention 503
- Retrying a File Collaboration job 713
- Retrying a job 927
- Review
 - source path selections 477
- REVIT 259
- Rich client
 - launching 23
 - Windows 23
- Role, Administrator 112
- Role, base 112
- Role, Help Desk 112
- Role, Power User 112
- Roles 429
- Runtime view
 - Cloud Backup and Replication job 51
 - DFS-N Management job 52
 - File Collaboration job 53
 - File Replication job 64
 - File Replication job 53
 - File Synchronization job 53, 65
- Runtime views 50

- S -

- s 207
- S3-compatible storage
 - Cloud Backup and Replication destination 479
- Scan at Start
 - baseline replication 499, 501
 - batched real-time 501
 - scheduled scans 499
- Scan Manager, preferences 255, 312
- Scan report 236
- Scheduled replication 158
- Scheduled replication filter 158
- Scheduled replication filters, File Replication job 778
- Scheduled scans
 - Cloud Backup and Replication 498
 - Cloud Backup and Replication jobs 499
 - daily 498
 - daily scans 499
 - destination snapshots 499
 - File System Analytics 121
 - Scan at Start 499
 - weekly 498
 - weekly scans 499
- Security 182
- Security style
 - file permissions 132
 - multi-protocol 132
- Seeding Target 666, 765, 877
- Shared project files
 - File Collaboration 600
- Shares
 - multi-protocol 132
- Show individual files
 - source paths 477
- Smart Data Seeding 159
- SMB
 - File Activity Analytics 122
 - File Collaboration storage platforms 605
 - File Replication storage platforms 718
 - File Synchronization storage platforms 810
 - File System Analytics 121
 - multi-protocol 132
 - shares 132
- SMB access 631, 741, 743, 747, 751, 765, 836, 839, 884

- SMB access 631, 741, 743, 747, 751, 765, 836, 839, 884
 - configuration 134
 - credentials 134
 - Linux Agents 134
 - mixed permissions 134
 - multi-protocol jobs 134
 - Nutanix Files 134
 - Windows-style permissions 134
- SMB access configuration 631, 741, 743, 747, 751, 765, 836, 839, 884
 - Preferences 134
- SMC access configuration 628
- SNMP configuration, preferences 405
- Snapshot retention
 - Cloud Backup and Replication 503
 - destination snapshots 497
- Snapshots
 - replication schedule 498
- SNMP notifications, Cloud Backup and Replication job 510
- SNMP notifications, DFS-N Management job 553
- SNMP notifications, File Collaboration job 701
- SNMP notifications, File Replication job 802
- SNMP notifications, File Synchronization job 915
- SNMP notifications, overview 110
- SNMP Notifications, preferences 226, 252, 276, 309
- Software updates, preferences 349
- Source Agent 716
- Source path 741
- Source Paths
 - Cloud Backup and Replication jobs 477
 - files and folders 477
 - reviewing selections 477
 - volumes 477
 - watch set 477
- Source snapshots
 - always use when replicating 504
 - Cloud Backup and Replication 497
 - Cloud Backup and Replication jobs 504
 - consistent backups 504
 - destination snapshots 504
 - disabled 504
 - locked files 497, 504
 - open files 497, 504
- Source storage device
 - File Replication 714
 - source Agent 714
- Source storage platform
 - Cloud Backup and Replication jobs 454
 - credentials 463
 - storage information 463
- Starting, File Synchronization job 922
- State
 - namespaces 76
- Status codes, API 156
- Stopping, File Synchronization job 924
- Storage accounts
 - Cloud Backup and Replication 479
 - destination credentials 480
 - proxy validation 456
- Storage Analytics 118
- Storage capacity
 - alerts 160
 - capacity consumption 160
 - licensed capacity 160
 - unique TB model 160
- Storage capacity alerts
 - licensed capacity 160
- Storage containers
 - Cloud Backup and Replication 488
 - creating 488
 - selecting existing 488
- Storage devices
 - Management Agent 454
- Storage Information
 - Cloud Backup and Replication jobs 463
 - credentials 463, 607, 720, 813
 - existing credentials 463, 607, 720, 813
 - new credentials 463, 607, 720, 813
 - source storage platform 463
 - validating credentials 463, 607, 720, 813
- Storage Platform
 - Cloud Backup and Replication jobs 454
 - File Collaboration jobs 605
 - File Replication jobs 718
 - File Synchronization jobs 810
 - multi-protocol 605, 718, 810
 - NFS 718, 810
 - SMB 605, 718, 810
 - source storage platform 454
- Storage platform configuration, preferences 370
- Storage platform credentials 234
- Storage platform types
 - Cloud Backup and Replication jobs 454
 - File Collaboration jobs 605

- Storage platform types
 - File Replication jobs 718
 - File Synchronization jobs 810
- Storage platforms
 - connection details 607, 720, 813
 - credentials 607, 720, 813
 - File Synchronization 805
- Storage tiering options
 - Cloud Backup and Replication job 505
- Summary tab, Collab, Sync, and Repl Summary view 69
- Summary views 66
- Symbolic Link Strategy
 - changing strategy 129
 - Copy 129
 - Follow 129
 - Multi-Protocol jobs 129
 - NAS participants 129
 - NFS jobs 129
- Symbolic links
 - changing strategy 129
 - Copy strategy 129
 - Follow strategy 129
 - Multi-Protocol 129
 - NAS participants 129
 - NFS 129
 - quarantined files and folders 129
 - SMB 129
- System alerts 338, 343
- System folders 338

- T -

- Tables 78
- Tags, File Collaboration job 702
- Tags, File Replication job 803
- Tags, File Synchronization job 916
- Tags, organizing resources 161
- Tags, preferences 351
- Tags, web roles 450
- TAM 407
- TAMs, Proactive Monitoring 216
- Tape replacement
 - Cloud Backup and Replication 451
- target protection 697, 798, 910
- Target protection, File Collaboration job 697
- Target protection, File Replication job 798
- Target protection, File Synchronization job 910

- Target storage devices
 - destination Agents 714
 - File Replication 714
- Tasks
 - Cloud Backup and Replication 67
- Technical Account Managers (TAMs) 216
- Telemetry 216, 407
- Telemetry options, Proactive Monitoring 409
- Telemetry, Anonymous Diagnostic Data 216
- Telemetry, environment name 408
- Telemetry, preferences 407
- Telemetry, Proactive Monitoring 216
- Temporary files 88
- Terminology 1
- Testing API access 155
- TLS certificates 182
- Toolbar 28, 35
- Tools menu 32
- Trap folder 358
- Trigger action 358

- U -

- Uninstall 23
- Unique TB model
 - storage capacity 160
- Unity, advanced configuration options 393
- Unity, configuration 390
- Unity, credentials 390
- Unstructured data 118
 - analytics 118
- Update Peer Agent 19, 21
- Update Peer Management Center 19
- Update PMC 19
- Updating software 349
- Updating, Agent 153
- Upgrade 354
- User activity 118
 - File Activity Analytics 122
 - PeerIQ 119
- User home directories
 - File Synchronization 805
- User interface 27, 40
- User management 429
- User management, preferences 429
- User profiles
 - File Synchronization 805
- User, Active Directory 111

User, internal 111
 Users, internal 431, 432, 435, 437
 Utilization policy, File Collaboration job 641, 851

- V -

Validate Proxy Configuration
 connection test 456
 storage account 456
 Version conflicts
 File Collaboration 600
 VHD
 PeerIQ 119
 View
 File Collaboration job runtime 53
 File Replication job runtime 53
 File Synchronization job runtime 53
 Views
 Cloud Backup and Replication jobs runtime 51
 DFS-N Management jobs runtime 52
 File Replication jobs runtime 64
 File Synchronization jobs runtime 65
 Views, Agent Summary 67
 Views, Agents 42, 45
 Views, Alerts 43
 Views, dashboard 45
 Views, Job Alerts 47
 Views, Jobs 48
 Views, runtime 50
 Views, summary 66
 VM options 151
 VMware
 PeerIQ 119
 Volume policy, File Collaboration job 638
 Volume policy, File Synchronization job 848
 Volumes
 Cloud Backup and Replication 477
 Cloud Backup and Replication status 67
 entire volume 477

- W -

WAN
 bandwidth 600
 File Collaboration 600
 WAN usage
 Processing Delay 502

Wasabi
 bucket details 488
 Cloud Backup and Replication destination 479
 credentials 480
 Wasabi, bucket details 496
 Watch set
 Cloud Backup and Replication 477
 Watch sets
 namespace folder targets 126
 Web client
 accessing 23
 remote access 23
 web browser 23
 Web client interface 111
 Web client required Peer services 26
 Web client user 111, 112, 113, 115
 Web client, accessing 24, 26
 Web client, secure access 14
 Web role 111
 Web role, custom 112, 115
 Web role, permissions 115
 Web role, standard 112
 Web roles 112
 Web roles permissions 113
 Web roles, create custom 444
 Web roles, delete 449, 450
 Web roles, edit 449, 450
 Web roles, managing 443
 Weekly scans
 destination snapshots 499
 scheduled replication 499
 Wildcards, filter patterns 87
 Window menu 30
 Windows
 multi-protocol 132
 Peer Management Center rich client 23
 Windows File Server
 Cloud Backup and Replication storage information 463
 storage information 607, 720, 813
 Windows-style permissions
 SMB access 134

- Y -

YAML file 155

- Z -

Zip files

File System Analytics scan results 122
transfer to PeerIQ 122