



Peer Management Center Help Manual

Copyright (c) 1993-2018 Peer Software, Inc. All Rights Reserved
Updated Friday, August 24, 2018

Table of Contents

Peer Management Center Help	1
Getting Started	1
Terminology	2
Requirements	4
Installation and Initial Configuration	5
Licensing	7
The Peer Management Center User Interface	9
Main View	11
Main View Toolbar	13
Web Interface	14
Menus	20
Jobs View	22
Peer Agent Summary View	25
Alerts View	31
Job Alerts View	32
Preferences	32
Configuring Email Alerts	33
Configuring File Filters	36
Filtering Folders	40
Configuring SNMP Notifications	42
Configuring SMTP Email	44
Configuring Tags	46
Configuring User Management	47
Advanced Configuration	53
Advanced Peer Management Center Configuration	54
Custom TLS Integration	54
Use Existing Certificate	54
Create New Certificate	59
File Filters	65
Filter Expressions	66
Using Tags	68
Advanced Peer Agent Configuration	71
Broker Configuration	72
General	73
Logging	74
Performance	75
VM Options	76
Advanced File Collaboration Configuration	76
Locking	77
Revit Enhancements	79
Scan Manager	81
Advanced Windows Real-time Options	82
Advanced Cloud Sync Configuration	83
Modifying Cloud Sync Settings	84
Creating Cloud Platform Credentials	85
Creating Database Connections	87
Creating Email Alerts	89
Creating File Filters	91

Modifying Cloud Sync Performance Settings	93
Modifying Scan Manager Settings	95
Creating Sync and Retention Policies	96
Real-time Detection	97
MED Configuration	99
NAS Configuration	104
EMC Configuration	104
EMC Prerequisites and Configuration	105
NetApp Configuration	112
NetApp Prerequisites and Configuration	113
Nutanix Configuration	119
Nutanix Prerequisites and Configuration	119
File Collaboration	122
Creating a File Collaboration Job	122
Step 1: Participants	122
Step 2: Security Settings	130
Step 3: Application Support	130
Step 4: Email Alerts	132
Editing a File Collaboration Job	132
Overview	132
Participants and Directories	133
General Settings	135
File Filters	137
File Conflict Resolution	139
Delta Replication	142
File Metadata	144
File Locking	148
Application Support	150
Logging and Alerts	150
Target Protection	154
Email Alerts and SNMP Notifications	155
Tags	157
Save Settings	158
Running and Managing a File Collaboration Job	158
Overview	159
Starting and Stopping	160
File Collaboration Summary View	161
Runtime Reports View	164
Dashboard Summary View	167
Peer Agent Detail Summary View	168
Editing Multiple Jobs	169
Host Connectivity Issues	171
Runtime Job Views	174
Summary View	174
Session View	179
Event Log View	180
File Conflict View	181
Alerts View	185
Participants View	186
Configuration View	188
Cloud Sync	189
Overview	190
Before You Create Your First Cloud Sync Job	190

Creating a Cloud Sync Job	190
Step 1: Job Type and Name.....	191
Step 2: Source Storage Platform	192
Step 3: Management Agent.....	193
Step 4: Storage Information.....	194
NetApp ONTAP Clustered Data ONTAP.....	195
NetApp Data ONTAP 7-Mode	197
Dell EMC Isilon	198
Dell EMC Unity	200
Dell EMC Celerra VNX VNX 2.....	201
Nutanix AFS	203
Step 5: Source Paths.....	204
Step 6: File Filters.....	205
Step 7: Destination.....	206
Step 8: Cloud Platform Credentials.....	207
Step 9: Miscellaneous Options.....	209
Step 10: Sync and Retention Policy.....	211
Step 11: Sync Schedule.....	213
Scheduled Scans.....	213
Continuous Data Protection.....	215
Step 12: Retention	216
Step 13: Email Alerts	217
Step 14: Confirmation.....	219
Running a Cloud Sync Job	221
Starting a Cloud Sync Job.....	221
Stopping a Cloud Sync Job.....	222
Deleting a Cloud Sync Job	223
Monitoring Your Cloud Sync Jobs	224
Views for Monitoring General Progress.....	225
Views for Monitoring Individual Jobs.....	225
Recovering Data from the Cloud	226
Search Options.....	227
Search by Name	228
Search by Snapshot.....	231
Search by Point in Time.....	232
Search by Latest Replication.....	234
Recovery Options.....	234
File Synchronization	237
Creating a File Synchronization Job	238
File Synchronization Configuration.....	238
SMTP Email Configuration.....	239
Email Alerts	241
Integrating Existing PeerSync Instances	244
Requirements	244
How To	244
Deploying New PeerSync Instances.....	245
Requirements	245
How To	245
Step 1: General Information.....	246
Step 2: PeerSync Profile.....	247
Step 3: Jobs Configuration List.....	249
Step 4: Installation Settings.....	250
Logging and Alerts	254
Email Alerts.....	255

Running and Managing a File Synchronization Job	256
Starting and Stopping.....	257
Synchronization Summary	258
Synchronization Dashboard Summary View.....	261
PeerSync Profile Management.....	262
Update the Profile Configuration.....	263
Import Existing Profile.....	264
Edit/Configure Jobs.....	265
Edit Global Settings.....	268
Distribute Profile.....	269
PeerSync Service Management.....	270
Runtime Job Views.....	271
Summary View	272
PeerSync Jobs Stats.....	274
Added Files	274
Updated Files	275
Deleted Files	276
Messages	277
Failed Events View.....	278
Monitoring Log View.....	278
Alerts View	279
Participants View.....	280
Configuration View.....	282
Upgrade/Reprocess Installation.....	282

Index

285

Peer Management Center Help

Using this Help File

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

Trademark Information and Copyright

Copyright (c) 1993-2018 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, and their respective logos are 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 Inc. 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.

Getting Started

The topics in this section provide some basic information about Peer Management Center, including installation and licensing:

- [Terminology](#)
- [Requirements](#)
- [Installation and Initial Configuration](#)
- [Licensing](#)

Terminology

Introduction

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

Terms

File Access Event	An event that is triggered from the opening or closing of a file.
File Change Event	A event that causes a file to be changed in some way, for example, file modify, file delete, file rename, file attribute change.
File Collaboration Job	<p>A specific instance of a Peerlet that can be created, saved, modified, and run. A Peerlet represents a type of job.</p> <p>In the case of File Collaboration, a File Collaboration job represents a single configurable File Collaboration Session. The two terms may be used interchangeably throughout the interface and this document.</p>
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 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.
Heartbeat	A communication mechanism used between the Peer Management Center and all connected Peer Agents to ensure that Peer Agents are alive and responsive. Heartbeats share information about the Peer Agent host server with the Peer Management Center, aid in verifying when a Peer Agents is no longer available, and signal when a disconnected Peer Agent has reconnected. All heartbeat information is sent through the Peer Management Broker.
Job	<ul style="list-style-type: none">• For File Collaboration, a job is a grouping of two or more participating hosts, each with a designated set of folders and files that are to be shared or collaborated on. A File Collaboration job coordinates the primary functions of file locking and synchronization.

	<ul style="list-style-type: none"> For Cloud Sync, a job is a single participating host with a designated set of folders and files that are to be replicated to a cloud storage device.
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.
Participating Host	The Management Agent of a file server that is taking part in a file collaboration session or a Cloud Sync job.
Peer Agent (or Agent)	A lightweight, distributed component that is used to perform operations on the host on which it is running. A Peer Management Center environment will typically contain several Peer Agents, one per participating networked host. Peer Agents invoke the distributed portions of a Peerlet, and will often run near resources of interest, such as collaborated files. The Peer Agent is designed to be purposed across the entire Peer Management Center solution suite, and will normally be directed to perform functions with messages received from Peerlets through the Peer Management Broker.
Peer Management Center	<p>The focal software component where Peerlets are installed, configured and ran. The Peer Management Center can host Peerlets of various types and is where the components of a centralized solution function. The Peer Agent is invoked by Peerlets, distributing components with messages sent through the Peer Management Broker.</p> <p>The Peer Management Center runs as three parts: a Windows Service that is set to run all the time, along with a rich client application and a web server component that both connect to the primary service for configuration and monitoring.</p>
Peer Management Broker	The central messaging system of the Peerlet framework. The Peer Management Broker serves to connect the Peer Management Center and the Peer Agents, forming a Peer Management Center "network" that can be cast over local or wide-area networks via TCP/IP. A Peer Management Center environment will deploy one or more Peer Management Brokers.
Peerlet	A solution built for the Peer Management Center framework. A Peerlet is a distributed application containing various parts, some of which function at a focal point called the Peer Management Center and others invoked at remote points designated as Peer Agents.

Peerlet Editor	<p>A container within the user interface of the Peer Management Center that shows runtime and configuration information for a single file collaboration job. A Peerlet Editor is represented by a single tab, typically in the large center section of the Peer Management Center's interface. The editor itself consists of multiple sub-tabs, with various runtime and configuration information dispersed amongst the sub-tabs. For more information, see Runtime Job Views.</p> <p>Editors for multiple file collaboration jobs can be opened in several different editor tabs, allowing for quick movement between jobs.</p> <p>The Peerlet Editor area of the Peer Management Center is referred to as the File Collaboration Runtime view throughout this document.</p>
Quarantined File	<p>A file that has been removed from a file collaboration session as a result of a file lock conflict that could not be resolved. This file will remain quarantined until the user manually removes it from quarantine.</p>
Source Host	<p>The Management Agent of a file server where a file access or change event originated from.</p>
Target Host	<p>One or more Management Agents of file servers where file access and change events will be propagated to.</p>
Views	<p>Individual sections of the Peer Management Center's user interface, each providing unique information and control.</p> <p>Examples: Main View, Jobs View, Peer Agent Summary View, Alerts View, Job Alerts View.</p>
Watch Set	<p>The configured root folder and all subfolders on a file server that are being scanned and/or monitored by a File Collaboration or Cloud Sync job.</p>

Requirements

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

- [Environmental requirements](#)

- [EMC Isilon prerequisites](#)
- [EMC Unity prerequisites](#)
- [EMC VNX/Celerra prerequisites](#)
- [NetApp 7-Mode prerequisites](#)
- [NetApp cDOT/ONTAP9+ prerequisites](#)
- [Nutanix AFS prerequisites](#)

Installation and Initial Configuration

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

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

Peer Management Center and Peer Management Broker Installation

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

Basic Requirements

See the [Requirements](#) section for more detailed requirements.

Software Installation and Launching

1. Run the PL-Hub_Installer.exe or PL-Hub_Installer64.exe installer and follow all instructions.
2. After the installation finishes, both the Peer Management Center and Peer Management Broker will be installed. The Peer Management Broker will automatically be installed as a running Windows service and set to auto-start. The Peer Management Center is installed in three parts: a Windows service that is set to auto-start, a web service for granting access to the Windows service via web browsers, and a rich client for interacting with the Windows service. The rich client is started as a normal Windows application.

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

Secure Encrypted TLS Connections

By default, the Peer Management Center and Peer Management Broker will be installed on the same host machine that does not require secure TLS communication between each other. To enable a secure TLS connection between the Peer Management Center and Peer Management Broker, first stop the Peer Management Center Service via the Windows Service Panel (services.msc). Once stopped, navigate to the directory **Hub\workspace\prefs**, relative to the installation directory. Within this directory, open the **com.ci.pl.hub.runtime.prefs** file in a text editor. If the file does not contain a line starting with **hub.jms.providerURL**, then add the following line in its entirety:

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

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

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

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

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

Uninstalling

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

Peer Agent Installation

You will need to install a Peer Agent on each server you plan to include in any of your [file collaboration sessions](#).

Basic Requirements

See the [Requirements](#) section for more detailed requirements.

Software Installation and Launching

1. Run the PL-Agent_windows.exe installer on the target server and follow all instructions.
2. During installation you will need to specify the Peer Management Broker Host Name (computer name, fully qualified domain name, or IP Address) of the server where the Peer Management Broker is running, as well as the configured TCP/IP port number (the default port for TLS communication is 61617).
3. After the installation finishes, the Peer Agent will be installed as a Windows service. You will need to verify that the Peer Agent is running, and that it was able to successfully connect to the Peer Management Broker. You can do this by opening Windows Service Panel (services.msc) and making sure that the **Peer Agent Service** is started.
4. Make sure that the Peer Agent was able to successfully connect to the Peer Management Broker by going to the Peer Agent installation folder, opening the output.log text file, and making sure that **Ready** is displayed on the first line.

Secure Encrypted TLS Connections

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

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

Uninstalling

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

Licensing

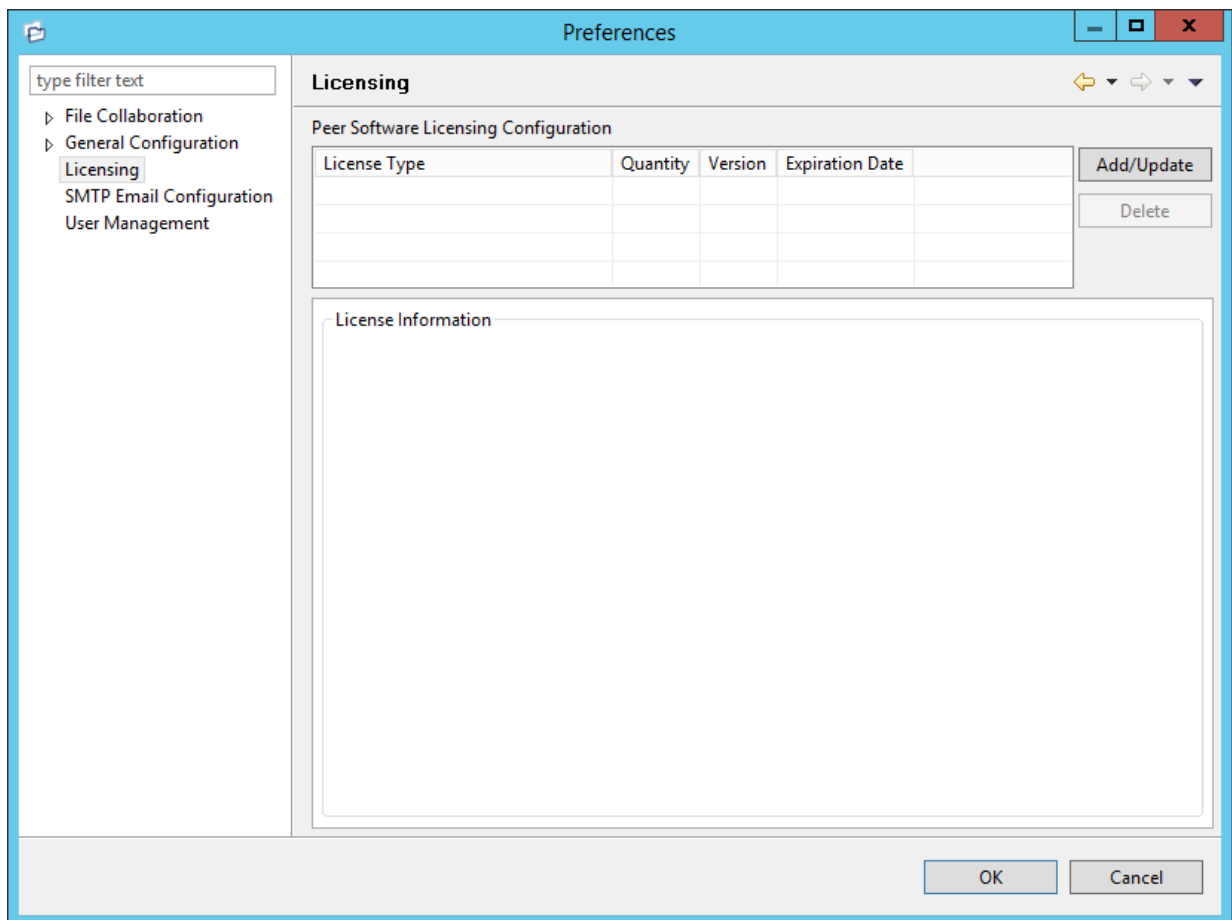
Peer Management Center is licensed by the number of unique [participating hosts](#) and by the number of terabytes in the Watch Set.

Installing or Upgrading a License File

After purchasing or requesting a trial download of Peer Management Center, you will receive a license file representing your purchase or trial.

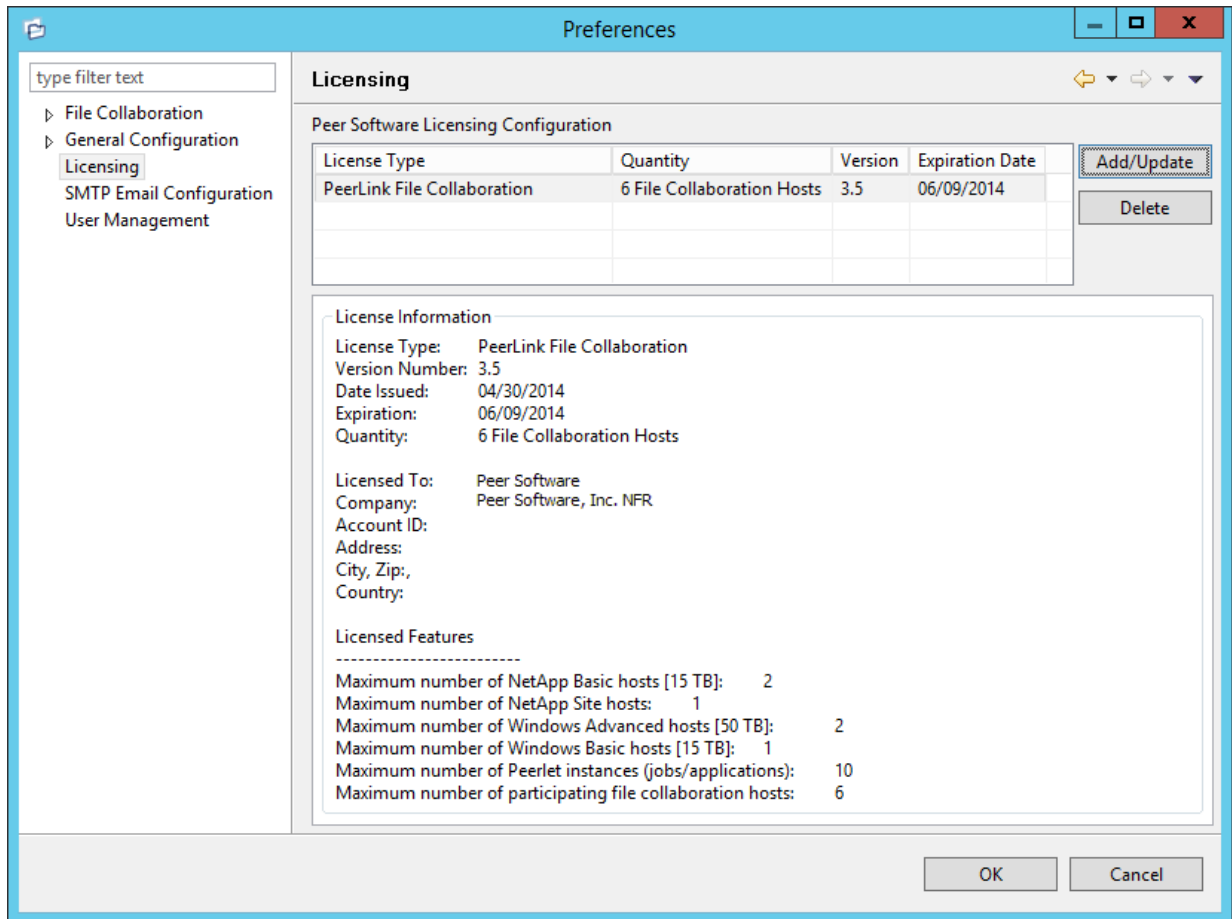
To install a new license file or upgrade an existing license:

1. Navigate to the **Window** menu in the [Peer Management Center](#) and select **Preferences**.
2. Select **Licensing** in the tree on the left of the **Preferences** dialog.



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

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



Deleting a License File

Click the **Delete** button to permanently remove ALL licenses of the selected type (valid and invalid licenses).

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

The Peer Management Center User Interface

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

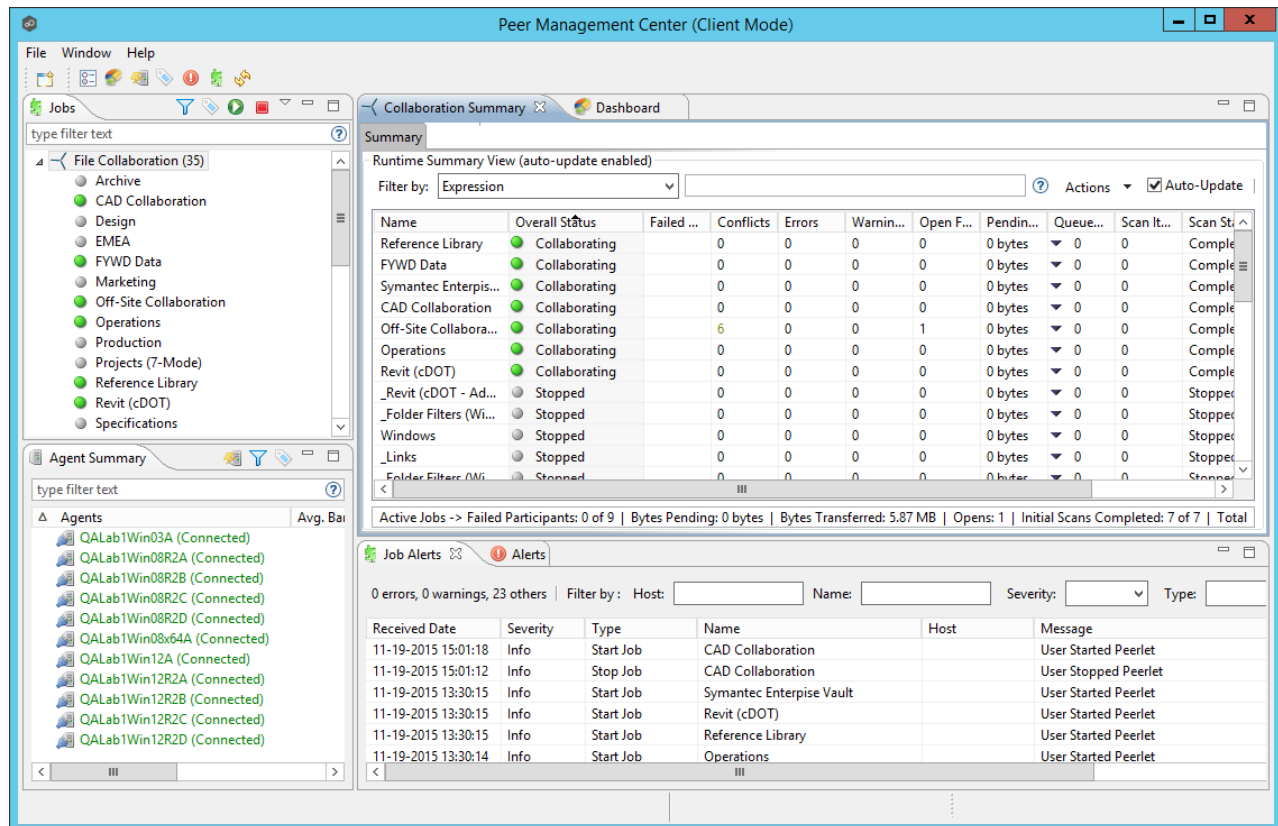
The graphical user interface of the Peer Management Center now has two separate options: a rich client installed and run on the server running the Peer Management Center, and a web service that, when configured, can be accessed from remote systems via a web browser.

For more information, see:

- [Main View](#)
- [Main View Toolbar](#)
- [Web Interface](#)
- [Menus](#)
- [Jobs View](#)
- [Peer Agent Summary View](#)
- [Alerts View](#)
- [Job Alerts View](#)

Main View

After starting up the Peer Management Center Client, the following Main view is displayed:



The [Peer Management Center](#) is made up of the following [views](#):

Jobs View	<p>The Jobs view is a list of all created file collaboration and cloud sync jobs that can be modified, viewed, and started. The list is grouped by Peerlets type, where the primary type is File Collaboration.</p> <p>The following buttons are available within this panel:</p> <ul style="list-style-type: none"> • Start and Stop buttons allow you to start and stop any selected jobs. • View Runtime Summary button displays a table of summary information for all jobs of a selected Peerlet type.
Agent Summary View	<p>The Agent Summary view displays a list of known Peer Agents and connection status for each. Individual Peer Agents can be updated and restarted from this view as well by right-clicking on</p>

	one or more items and selecting the appropriate item from the pop-up menu.
Alerts View	The Alerts view displays a list of Peer Management Center alerts that have occurred with detailed information about each alert. Alerts relating to Peer Agent connection status changes will be reported here.
Job Alerts View	The Jobs Alerts view displays a list of all job-specific alerts that have occurred (including those for file collaboration sessions with detailed information about each alert. Alerts relating to the automatic stopping and restarting of jobs will be reported here.
File Collaboration Runtime View (tabbed View in center of screen)	<p>The Peerlet Editors view is the default location of the Collaboration Summary view, in addition to runtime and configuration subviews for all open jobs.</p> <p>For each open file collaboration job, the following sub-views are available as tabs:</p> <ul style="list-style-type: none"> • Summary tab - Displays current synchronization summary and session statistics. • Session tab - Shows currently opened files, session locks, and files being synchronized. • Event Log tab - Displays a log of recent file activity. • File Conflicts tab - Shows a list of current file conflicts and quarantined files. • Alerts tabs - Displays alerts tied specifically to the selected jobs. • Participants tab - List of currently configured and associated host participants for the selected job, in addition to connection status for each. • Configuration tab - Shows a summary of all configurable items for the selected job.
Dashboard	Shows the Dashboard Summary View panel, which displays metrics and key performance indicators from all running File Collaboration jobs , File Synchronization jobs , and Agents.
Peer Agent Detail	The Peer Agent Detail Summary is a panel that displays a list of all known Peer Agents deployed and their detailed status

Summary	information, which can be used to assess the health of the environment.
----------------	---

Table Detail Viewer

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

The image shows a 'Hub Alert Details' pop-up dialog. It has a title bar with a close button. The dialog contains several fields with labels on the left and values in text boxes on the right:

- Received at: 05-06-2014 15:12:12
- Severity: Info
- Category: Agent
- Host Name: Win12R2a
- Locally Generated at: 05-06-2014 15:12:12
- Name: Connection
- Message: TLSv1.2 Connection to: Win12R2a:61617 (Cipher Suite: TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384)

At the bottom right of the dialog, there is a text link that says 'Click outside of popup to close'.

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

Main View Toolbar

The Peer Management Center toolbar buttons include:

New Job	Initiates the process of creating a new Job based on the Peerlet types that are installed.
User Preferences	Opens the Preferences window allowing the user to configure settings for the Peer Management Center, as well as global settings for file collaboration sessions .
View Dashboard	Shows the Dashboard Summary view, which displays metrics and key performance indicators from all running File Collaboration jobs, Cloud Sync jobs, File Synchronization jobs, and Agents.

View Agent Detail Summary	Shows the Peer Agent Detail Summary view , which displays a list of all known Peer Agents deployed and their detailed status information, and can be used to assess the health of the environment.
Assign Tags	Opens the Assign Tags dialog where resources can be viewed and assigned to tags or categories. Tagging resources helps when managing large number of resources.
View Alerts	Opens the Alert view , which displays Peer Management Center alerts such as Peer Agent connection status changes.
View Job Alerts	Opens the Job Alerts view , which displays alerts such as job restarts.
Refresh the current selected view	Refreshes all current views and tabs.

Web Interface

Peer Management Center now offers a new way to manage and monitor collaboration jobs via a robust web interface. Unlike many other web management consoles, Peer Management Center's web interface is very responsive and is built to mirror the functionality of the rich client (which is still included with the Peer Management Center installer for use by system administrators). When properly configured, the web interface allows system administrators to manage Peer Management Center's collaboration jobs from any location without the need to remotely log in to the Peer Management Center server.

In addition, this web interface includes a role-based log-in system with two out-of-the-box roles: **admin** and **helpdesk**. The former has complete access to all functionality found in the Peer Management Center's rich client, while the latter only has a read-only view of collaboration jobs along with the ability to release conflicts for any running jobs.

How to Set Up

The setup process for the web interface is driven by following screen within the installer for the Peer Management Center:

Setup - Peer Management Center 4.0.1.20151113

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

☒ Local Access - Make Web Service accessible to users on the local server only
☐ Public Access - Make Web Service accessible to users from anywhere on the network
☐ Disable Web Service

Hostname or IP:

☒ Enable HTTP using Port:
☐ Enable HTTPS using Port:

The options on this screen are as follows:

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

Enable HTTP (using Port)	Enables HTTP access to the web interface using the specified port.
Enable HTTPS (using Port)	Enables HTTPS access to the web interface using the specified port and a built-in TLS certificate. More details on changing TLS certificates can be found here .

If you need to make changes to the configuration of the web interface, you will need to stop the **Peer Management Center Web Service** in **services.msc** and use Notepad to modify the **config.ini** file located under **Peer Management Center_INSTALL_FOLDER\Hub\web-configuration** (where **PMC_INSTALL_FOLDER** represents the root installation directory of Peer Management Center). Once modifications are complete, save the file and restart the **Peer Management Center Web Service**. The important items to configure within this file are:

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

Important Notes for the config.ini File

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

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

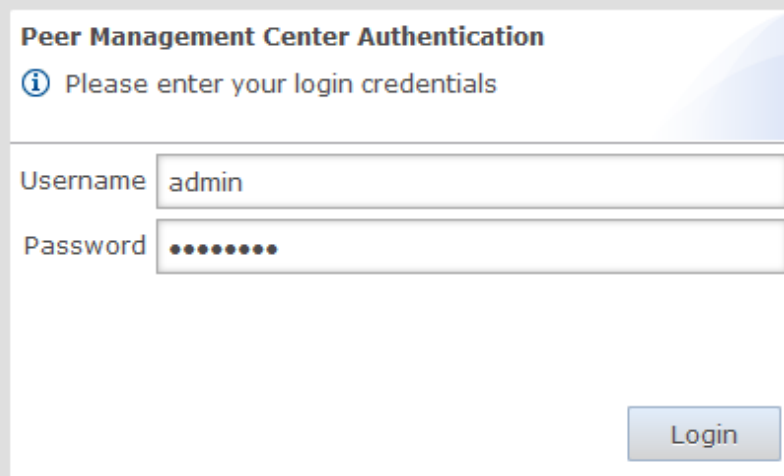
- HTTP and HTTPS are configured independently of one another in the **config.ini** file and as such, can be set to different modes. For example, HTTPS could be configured in a public mode, while HTTP is set to private ("localhost").
- DO NOT modify any other settings in the **config.ini**. Doing so may result in the inability of the web interface to start.
- Duplicate entries in the **config.ini** file may also result in the inability of the web interface to start.

How to Use

Once Peer Management Center has been installed and all services have been started, open up a web browser and enter the following URL:

<http://localhost:8081>

Note that the exact URL will vary depending on the settings you have selected in the [How To Set Up](#) (for example: http vs https, appropriate hostname or IP, and appropriate port). In the page that loads, select the **Peer Management Center Portal** link. The following page will then be displayed:



The image shows a login form titled "Peer Management Center Authentication". Below the title is an information icon and the text "Please enter your login credentials". There are two input fields: "Username" with the text "admin" and "Password" with ten dots. A "Login" button is located at the bottom right of the form.

Peer Management Center Authentication

i Please enter your login credentials

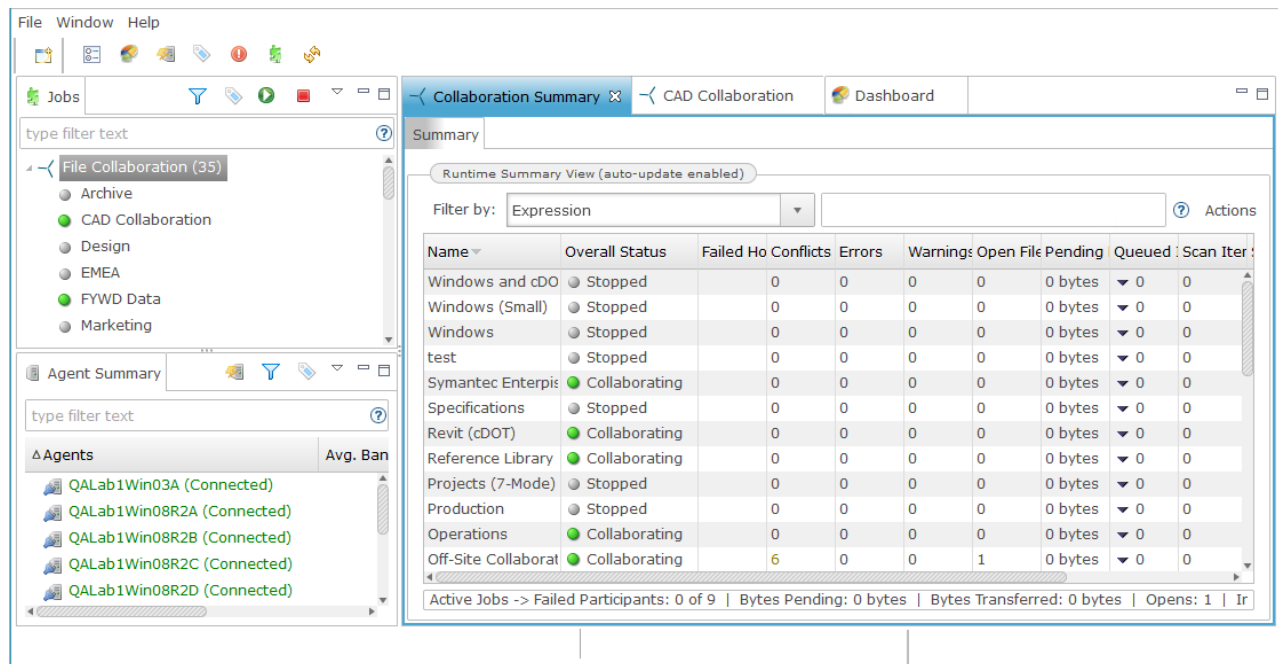
Username

Password

Login

The default user name is **admin** with a default password of **password**. We highly recommend that you change this password. See [User Management](#) for more information on changing account passwords.

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



As mentioned above, those with **admin** accounts will have complete access to the Peer Management Center's user interface. For more details on how to use the full Peer Management Center interface, see [Main View](#).

Those with **helpdesk** accounts are limited to read-only access of the following:

- The [Jobs view](#)
- The [Collaboration Summary view](#)
- The [Summary](#) and [Session](#) tabs of each job.

In addition, these accounts have read-write access to the [File Conflicts](#) tab of each job, with the ability to release conflicts.

How to Secure Access

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

- The default **admin** account password should be changed immediately. For details, see [User Management](#).
- Access to the web interface can be in the form of both HTTP and HTTPS. The latter will ensure that all communication between the client browser and the service hosting the web interface is encrypted. Regardless of which is enabled, the hostname or IP address through which clients can reach the web interface can be configured to limit access. See [How to Set Up](#) for more details.

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

Menus

After starting up the Peer Management Center Client, the following menu and toolbar actions are available:

File Menu

New	Selecting this option will present you with a list of installed Peerlets types from which you can create a new job. The options are based on which Peer Management Center solution is installed. For example, if you installed the File Collaboration solution, then clicking the New menu item will provide you with an option to create a new file collaboration job. The New action is available in the toolbar as well.
Save/Save All	This button will be enabled if any of the open jobs have been modified. Selecting Save will result in the currently open and selected job to be saved to disk. Save All saves all open and modified jobs to disk.
Exit	Selecting this option will exit the Peer Management Center Client application. Note that as long as the Peer Management Center Service remains running, all running jobs will continue to operate.

Window Menu

Open Perspective	Open a predefined layout of views geared towards a specific purpose. For example, one perspective is for job creation and management, while another is for managing Peer Agents.
Reset Perspective...	Selecting this option will reset all current windows, views, and editors to their default size and layout.

View Dashboard	Shows the Dashboard Summary view , which displays metrics and key performance indicators from all running File Collaboration Jobs , File Synchronization Jobs , and Agents.
View Agent Details Summary	The Peer Agent Detail Summary is a panel that displays a list of all known Peer Agents deployed and their detailed status information, which can be used to assess the health of the environment.
View Alerts	Opens the Alerts view , which displays Peer Management Center alerts such as Peer Agent connection status changes.
View Job Alerts	Opens the Job Alerts view , which displays alerts such as job restarts.
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.
Preferences	Opens the Preferences dialog, allowing the user to configure settings for the Peer Management Center, as well as global settings for File Collaboration and Cloud Sync jobs.
View Progress	Opens the Progress view, which displays information pertaining to any running background tasks within the Peer Management Center.
Refresh	Refreshes all current views and tabs.

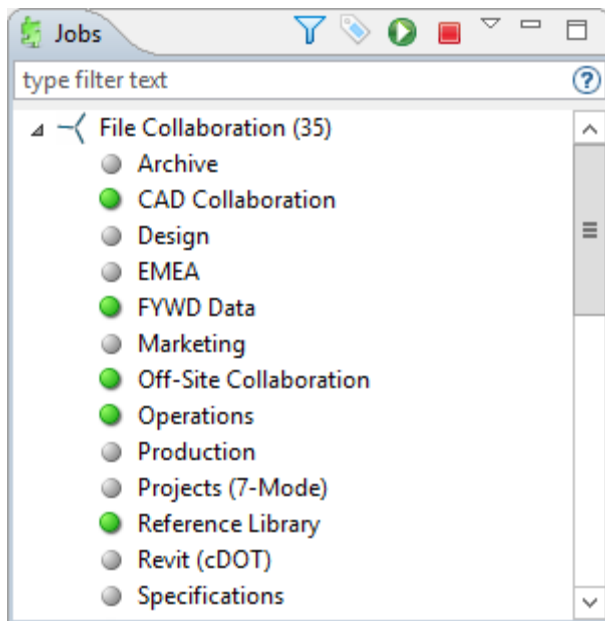
Help Menu

User Guide	Selecting this option will open this help system.
Download Peer Agent Installer	This operation takes you to our website where you can download the Peer Agent installer compatible with this version of the Peer Management Center.
Retrieve Hub/Agent Logs	This operation will collect and retrieve all useful log files for specified Peer Agents, the Peer Management Center, and all configured jobs. All of this information will be assembled into a single encrypted zip file that can optionally be uploaded to our technical support team.

	The collection and retrieval of the log and support files will be performed in the background, which might take awhile depending on content size and network speed. Upon completion, you will be notified and will be able to view the zip file yourself.
Retrieve Broker Statistics	This will display detailed statistical information about all messaging that has transpired for all connections (Peer Agents and the Peer Management Center) to the Peer Management Broker .
Thread Dump	Gives options to generate a thread dump of the running Peer Management Center Client and Service, as well as the running Peer Management Broker service. Both of these can be used by our technical support to debug certain issues.
Generate Memory Dump File	This will generate a memory dump of the running Peer Management Center Client and Service, which can be used by our technical support to debug certain issues.
Compress DB on Restart	Check this option in cases where the database consumes a large amount of disk space. This option will compress the database upon restart of the Peer Management Center Service.
About Peer Management Center	Displays version information about the Peer Management Center along with which components are installed.

Jobs View

The Jobs view is located in the top left section of the Peer Management Center and contains a list of all [Peerlet](#) types and saved instances.



Double-clicking any job will open the selected job in the Main view. Double-clicking the Peerlet type **File Collaboration** will open the [Collaboration Summary view](#) in the open tabs section, while double-clicking the Peerlet type Cloud Sync will open the [Cloud Sync Summary view](#) in the open tabs section.

Context Menu

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

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

Copy	Copy the selected job while assigning it a unique name.
Rename	Rename the selected job.

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

Toolbar

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


Manage, Save and Load Filters	Allows for the selection of built-in or user-defined filters and to save/manage filter expressions . Default Job filters include Failed Jobs, Jobs with Backlog, and Running Scans. For example, filter: "running Scans."
Assign Tags	Opens the Assign Tags dialog where resources can be viewed and assigned to tags or categories. Tagging resources helps when managing large number of resources.
Start Job	Start one or more selected and currently stopped jobs.
Stop Job	Stop one or more selected and currently running jobs.
View Collaboration Summary	View a table of summary information for all jobs of a selected Peerlet type. The view is defined and opened by simply clicking on a selected job (such as Document Collaboration in the image above) or its parent Peerlet type (or File Collaboration in the image above), then pressing the View Runtime Summary button.

Filtering

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

Peer Agent Summary View

The Peer Agent Summary view is located in the bottom left section of the Peer Management Center below the [Jobs view](#). This view contains a list of all known Peer Agents installed in your environment and displays the current connection status for each.



Toolbar

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

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

Filtering

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

Valid connection statuses are:

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

Peer Agent Menu Options

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

Remove	This will remove the selected Peer Agent(s) from the view, but if the Peer Agent is still running or connects again, then it will be added back to the list when the next heartbeat is received.
View Properties	Displays properties for the selected Peer Agent, e.g., heartbeat information, host machine configuration, messaging statistics, performance statistics. See View Peer Agent Properties Dialog for more details.
Edit Configuration	Clicking this menu item will display a dialog where you can edit user configurable properties for the selected Peer Agent.

Restart Agent Service	If the selected Peer Agent is connected, this menu item will restart the Peer Agent Windows service running on the corresponding host. In the event that the Peer Agent is not connected to the Peer Management Center Broker, an attempt will be made to restart the Peer Agent Windows service using the Windows sc command. Please note that this will only work if the user running the Peer Management Center can access the remote Peer Agent system and has the appropriate domain permissions to start and stop services on the remote Peer Agent system.
Remote Desktop	Launch a Windows Remote Desktop connection to the selected Peer Agent.
Edit Agent Configuration	This action displays a dialog through which the selected Peer Agent can be configured. Configurable options include Peer Management Center connectivity, Peer Agent logging, Peer Agent memory usage, among others. For more information, see Advanced Peer Agent Configuration .
Retrieve Log Files	This action retrieves log files for the selected Peer Agent containing information used by our technical support staff to assist in debugging issues. The log files are encrypted and will be located in the support folder of the Peer Management Center installation directory. They can optionally be uploaded to our technical support team.
Test Agent Bandwidth Speed	If the selected Peer Agent is connected, this menu item will start a bandwidth speed test to be performed in the background. You will be notified at completion with the results of the test.
Generate Thread Dump	This will generate a thread dump for the selected Peer Agent, which can be used by our technical support to debug certain issues. The debug file will be located in the Peer Agent installation directory.
Generate Memory Dump	This will generate a memory dump for the selected Peer Agent, which can be used by our technical support to debug certain issues. The debug file will be located in the Peer Agent installation directory.
Memory Garbage Collection	Force a garbage collection operation to attempt to reclaim memory that is no longer used within the Peer Agent's JVM.

Copy File	This action copies a specified file from the Peer Management Center to the designated target folder on each selected Peer Agent. The target folder is relative to the Peer Agent installation directory.
Transfer Rate Report (not available on Web Client)	This action displays a time series performance chart of average transfer rate for the selected Peer Agent over the last 24 hours.

Peer Agent Updates

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

View Peer Agent Properties Dialog

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

View Agent Properties

General | Heartbeat | Machine | Messaging | Performance | JVM Performance

Agent Host Name: QALAB1WIN12R2F

Connection Status: Connected

Custom Description: ☐

Description: Windows

Discovery Time: 03-22-2017 18:27:53

Heartbeat Enabled: ☒

JVM Architecture: amd64

JVM Version: 1.8.0_121-b13

Local Time: 04-25-2017 11:01:36 EDT

Local TimeZone: Eastern Standard Time

SSL Enabled: ☐

Start Time: 04-20-2017 13:02:51

OK Cancel

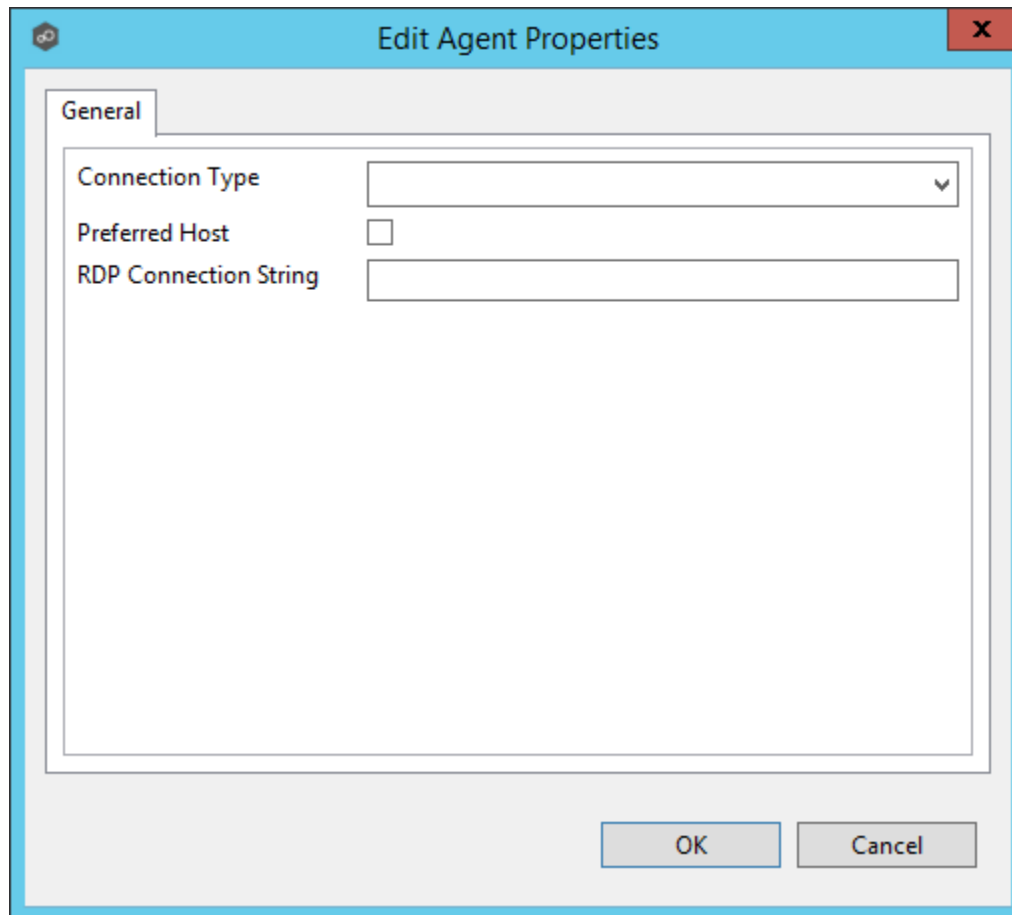
This dialog displays Peer Agent and host machine information across the following categories:

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

Edit Peer Agent Properties Dialog

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



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

Connection Type	Allows for the selection of a connection type between selected Peer Agent and its associated Peer Management Center Broker. When set, optimizations are made to the communication between the two parties based on the selected connection type.
------------------------	--

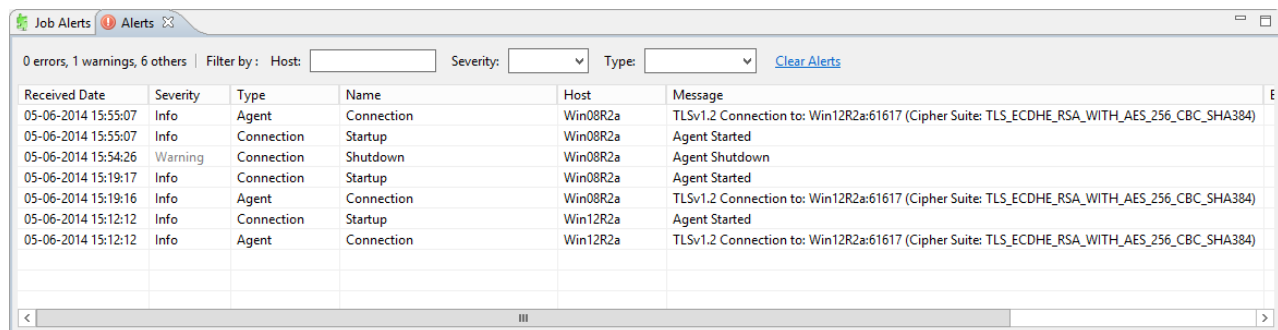
Preferred Host	A best practice optimization for selecting which Peer Agent has the fastest connection to the Peer Management Center Broker (or in appropriate cases, for selecting which Peer Agent are on the same subnet as the Peer Management Center Broker).
RDP Connection String	The connection string to use when activating an RDP session to this Peer Agent.

Alerts View

The Alerts view is automatically displayed when a critical system (Error or Fatal) alert is received. By default, the Alerts View is displayed under the [File Collaboration Runtime view](#). You can close the view at anytime by clicking on the **X** (Close) button on the Alerts tab. You can open the Alerts view at any time by clicking the **View Alerts** button located on the Peer Management Center toolbar or by selecting the **Window** menu, then the **Show View** submenu, followed by the **View Alerts** menu item.

Alert severity is broken down into four main categories: Informational (containing Info, Debug, and Trace), Warning, Error and Fatal. An example of an Informational alert is when an [Peer Agent](#) connects to the [Peer Management Broker](#). If an Peer Agent's network connection is severed, then an Error alert will be logged. All alerts are also logged to the file **hub_alert.log**, available under the 'Hub\logs' sub directory within the Peer Management Center installation directory.

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



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

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

- [Email Alerts](#)
- [File Filters](#)
- [SNMP Notifications](#)
- [SMTP Email Configuration](#)
- [Tags Configuration](#)
- [User Management](#)

Configuring Email Alerts

Overview

The Peer Management Center supports the concept of email alerts, where a single email alert (consisting of a unique name, a selection of event types, and a list of email addresses) can be applied to multiple jobs without requiring repeat entry for each job. When an email alert is applied to a job, an email will be sent to all listed recipients anytime a selected event type is triggered by that job.

To manage these configurations, click the **Window** menu, and then select **Preferences**. In the dialog that appears, select **Email Alerts** from the tree node on the left. A list of defined email alerts is presented, along with buttons to add new ones and to edit, copy, and remove existing ones.

Event Types

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

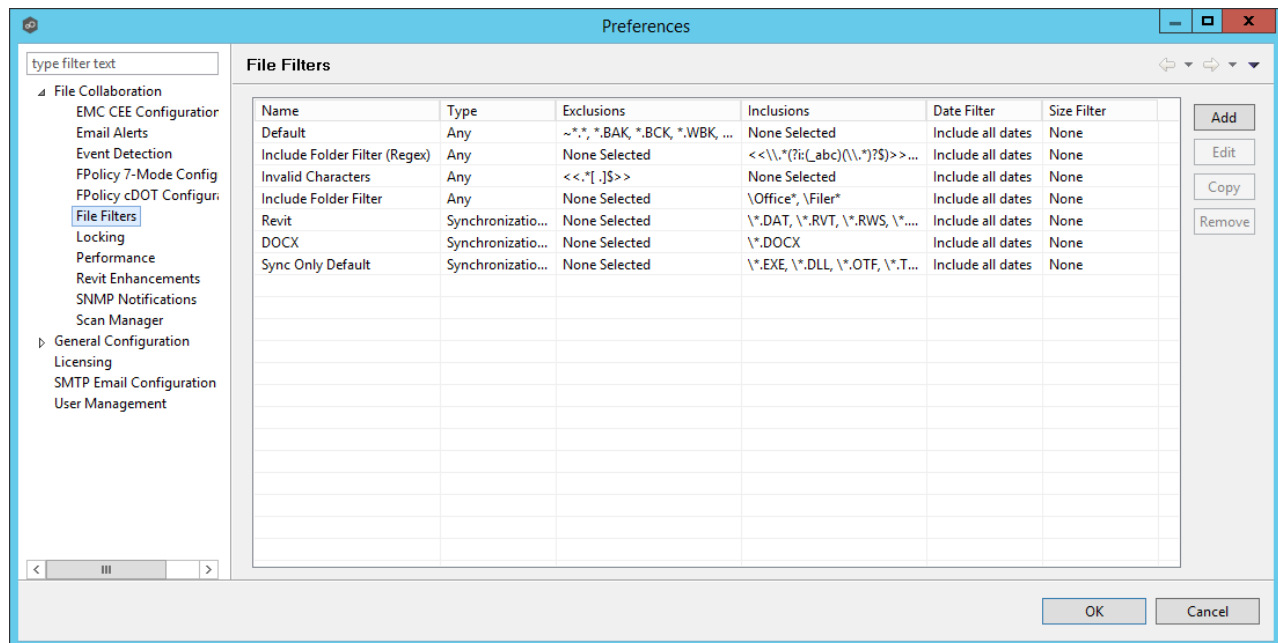
Queue Alerts

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

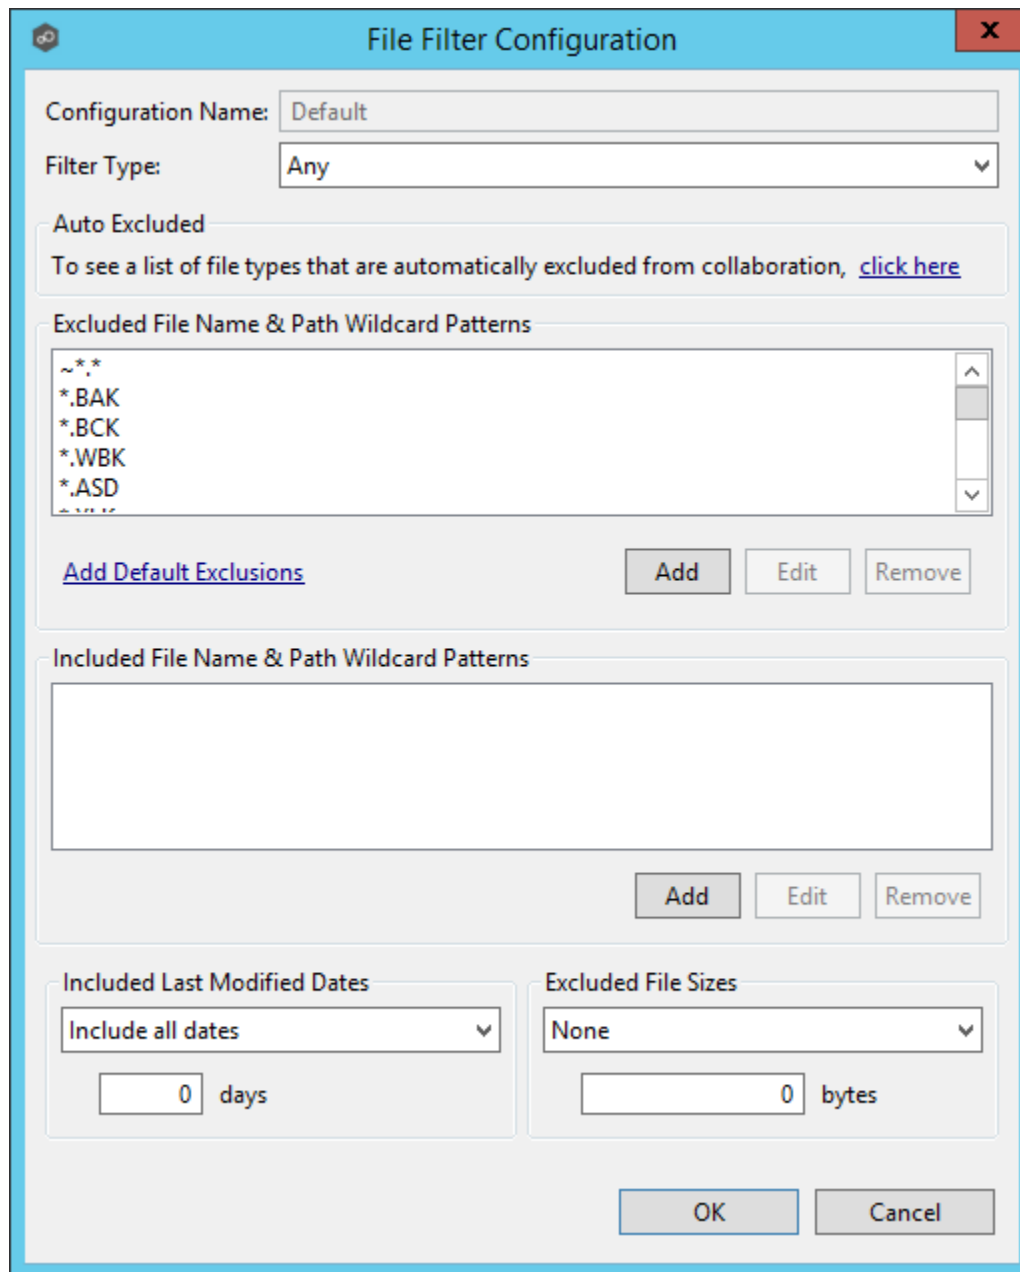
Configuring File Filters

The Peer Management Center supports the concept of file filters, where a single filter (consisting of a unique name, and lists of inclusion and exclusion expressions) can be applied to multiple [file collaboration jobs](#) and Cloud Sync jobs without requiring repeat entry for any job. This capability also allows you to define file filter combinations for use with specific collaboration scenarios.

To manage these configurations, navigate to the **Window** menu of the Peer Management Center, select **Preferences**, then navigate to and select **File Filters** from the tree node on the left. The following screen presents the list of defined file filter configurations, along with buttons to add new ones and to edit, copy, and remove existing ones. To increase flexibility, multiple file filters can be applied to a single job, combining elements of each to form one large filter. For more information on selecting multiple filters, see [File Filters](#) for file collaboration jobs and see [Step 6: File Filters](#) for Cloud Sync jobs.



Upon adding or editing a file filter, the following dialog is displayed:



The image shows a 'File Filter Configuration' dialog box with a blue title bar and a red close button. It contains several sections for configuring file filters:

- Configuration Name:** A text field with 'Default' entered.
- Filter Type:** A dropdown menu set to 'Any'.
- Auto Excluded:** A section with a link: 'To see a list of file types that are automatically excluded from collaboration, [click here](#)'.
- Excluded File Name & Path Wildcard Patterns:** A list box containing:
 - ~**
 - *.BAK
 - *.BCK
 - *.WBK
 - *.ASD
 Below the list are buttons for 'Add', 'Edit', and 'Remove', and a link 'Add Default Exclusions'.
- Included File Name & Path Wildcard Patterns:** An empty text area with 'Add', 'Edit', and 'Remove' buttons below it.
- Included Last Modified Dates:** A dropdown set to 'Include all dates' and a text field with '0' followed by 'days'.
- Excluded File Sizes:** A dropdown set to 'None' and a text field with '0' followed by 'bytes'.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom right.

When creating a file filter, you will generally want to exclude all temporary files created by the applications you use so they are not propagated to the targets hosts. For example, AutoCAD applications should add the following expressions to the Excluded File Name filter table:

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

To do so:

1. Click the **Add** button under the **Excluded File Name Wildcard Pattern** table and enter ***.AC\$**, and then click **OK**.
2. Repeat Step 1 to add ***.SV\$**, ***.DWL*** and ***.BAK**.

Your AutoCAD temporary file exclusion filter configuration is now created and all files ending in *.SV\$ or *.AC\$ or *.DWL or *.BAK will be excluded from collaboration within any running file collaboration job that uses this configuration.

Additionally, complex regular expressions in the format <<regEx>> can be used in both the inclusion and exclusion pattern lists. An example is shown in the dialog screenshot above (<<^.*\\atmp[0-9]{4,}\$>>).

The following regular expression 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
```

Filtering on Last Modified Date

In addition to filtering on file names, file extensions, folder paths, or partial path wildcard pattern matching, you can filter based on a file's last modified date. Peer Management Center only supports filtering on a file's last modified date and does not support filtering on a folders last modified date. In addition, if you have a folder hierarchy that contains files that are all being filtered based on last modified date, then all folders will still be created during the initial scan process on all hosts. If a file is excluded from collaboration based on last modified, then

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

Note that if last modified date filtration is used in a single filter configuration, no other types of filtration can be used in the configuration.

Options for Included Last Modified Date Filter

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

Filtering on File Sizes

Filtration can also be done based on an individual file's size. Peer Management Center does not support filtering on a folder's total size. In addition, if you have a folder hierarchy that contains files that are all being filtered based on size, then all folders will still be created during the initial scan process on all hosts. If a file is excluded from collaboration based on size, then the initial scanning process will not synchronize the file even if the file's last modified time and size do not match, or the file does not exist on all hosts. However, if a user deletes or renames the file on any host, the file will be deleted or renamed from all other hosts where the file exists.

Note that if excluded file size filtration is used in a single filter configuration, no other types of filtration can be used in the configuration.

Options for Excluded File Sizes

None	This is the default option and will include all files regardless file size.
Exclude files greater than or equal to	Exclude all files whose size is greater than or equal to the specified number of bytes. For example, you can configure a job to exclude all files greater than 1GB (1073741824 bytes).

Exclude files less than

Exclude files whose size is less than the specified number of bytes.

Sync-Only and Lock-Only Filters

With sync-only filters, Peer Management Center now supports the ability to exclude file types from being locked when a file open is detected on a participant.

Likewise, with lock-only filters, Peer Management Center now supports the ability to exclude synchronization across an entire job so that only opens and closes are detected and acted on, without any synchronization being performed.

To select one of these two filters, use the **Filter Type** drop-down list.

For more details on these filters and when they should be used, review this Tech Brief: <http://www.peersoftware.com/resources/tech-briefs.html?view=document&id=84>

Filtering Folders

In addition to filtering on files, you can filter on folders using the following syntax: **\Folder** or **\Folder*** or **\Folder***

Presently, Peer Management Center only supports included expressions for a full folder path 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, preferably to no more than ten. 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. Here is a small example that demonstrates this concept:

Example:

Reduce existing jobs down to two:

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

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

After consolidation:

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

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

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

Please note the following regarding regular expressions:

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

Additional Folder Filter Examples

To exclude a specific folder from anywhere within the watch set:

*\FolderName

*\FolderName\FolderName

To exclude a specific folder from the ROOT of the watch set:

\FolderName

\FolderName\FolderName

To exclude folders that END with a specific name from anywhere within the watch set:

*FolderName\

To include a specific folder from the ROOT of the watch set:

\FolderName

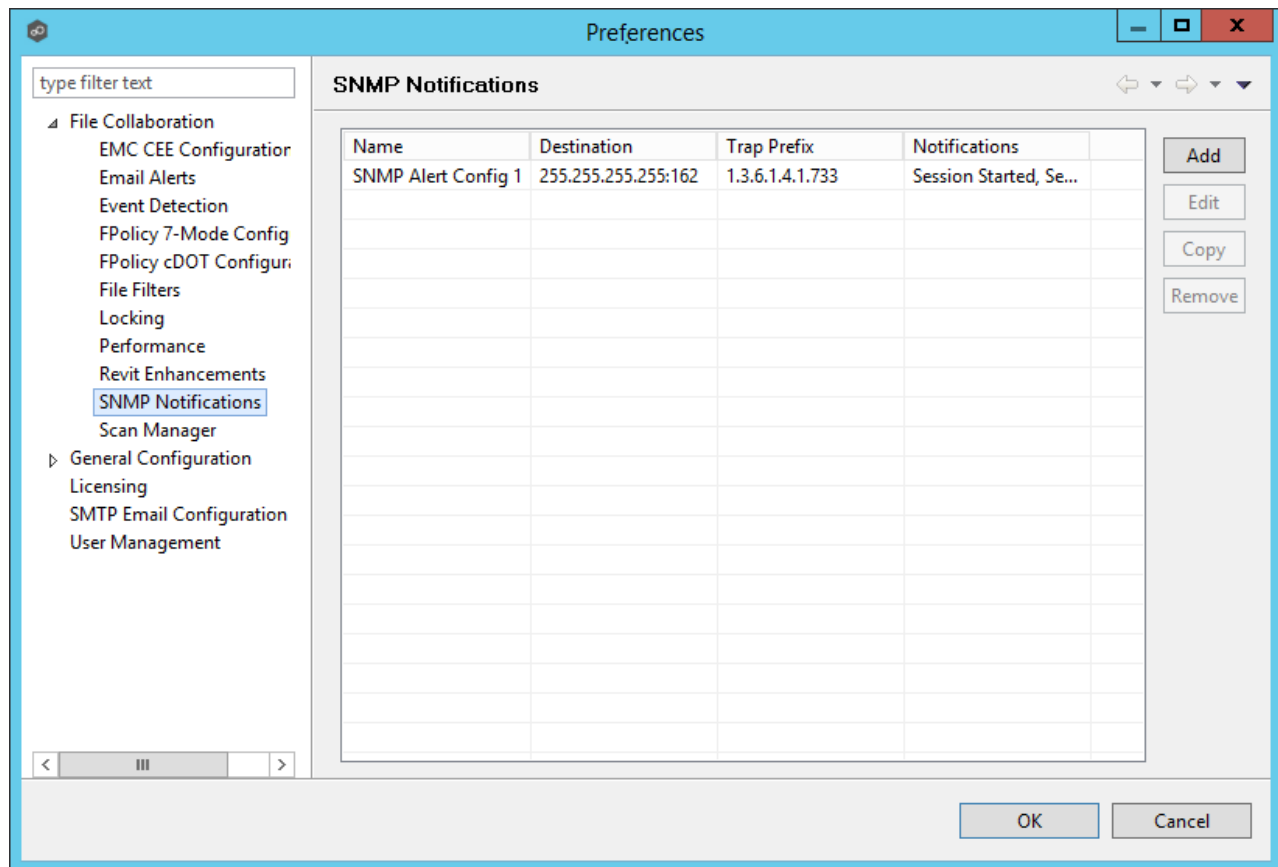
\FolderName\FolderName

Configuring SNMP Notifications

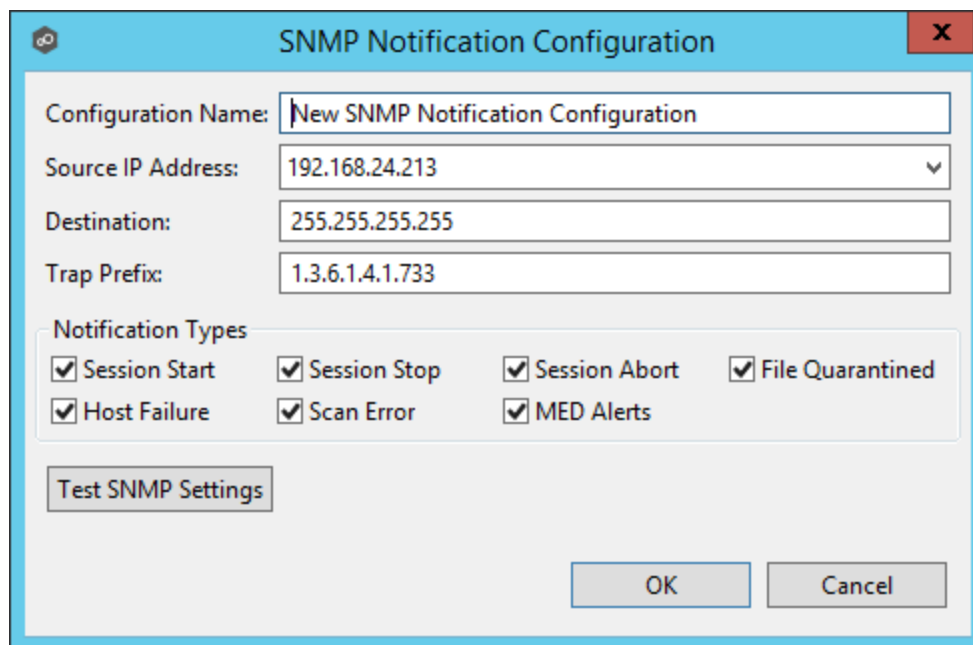
Overview

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

To manage these configurations, click the **Window** menu, select **Preferences**. In the dialog that appears, select **SNMP Notifications** from the tree node on the left. The following screen is displayed. It presents the list of defined SNMP notifications, along with buttons to add new ones and to edit, copy, and remove existing ones.



Upon adding or editing a SNMP notification, the following dialog is displayed:



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

Notification Types

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

Configuring SMTP Email

Before the Peer Management Center can send emails on behalf of any job, a few key SMTP settings must be configured. To set these values, click the **Window** menu, and then select **Preferences**. In the dialog that appears, select **SMTP Email Configuration** from the tree node on the left. The following is displayed:

The screenshot shows the 'Preferences' window with the 'SMTP Email Configuration' tab selected. The sidebar on the left lists various configuration categories, with 'SMTP Email Configuration' highlighted. The main configuration area includes the following settings:

- SMTP Host:** smtp.office365.com
- SMTP Port:** 587
- Encryption:** ☒
- Encryption Type:** TLS
- Username:** user@domain.com
- Password:** [Masked with dots]
- Sender Email:** user@domain.com
- Use Recommended Office365 Settings:** ☒

A 'Test Email Settings' button is located below the configuration fields. The 'OK' and 'Cancel' buttons are at the bottom right of the dialog.

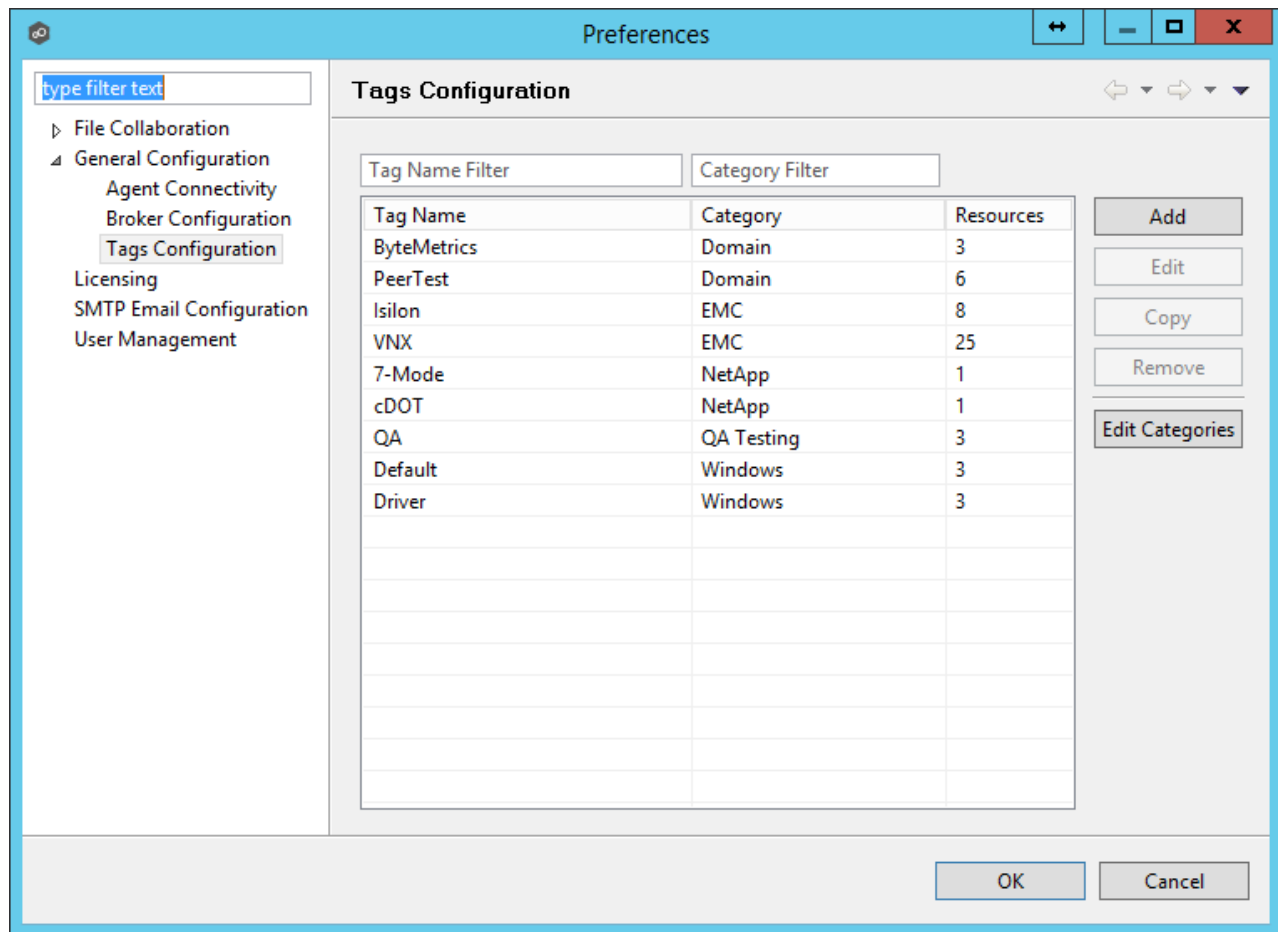
SMTP Host (required)	The host name or IP address of the SMTP mail server through which the Peer Management Center will send emails.
SMTP Port	TCP/IP connection port (default is 25 and 465 for encryption) on which the mail server is hosting the SMTP service. It is recommended that you leave the default setting unless your email provider specifies otherwise.
Encryption	Check this box if the SMTP mail server requires an encrypted connection.
Encryption Type	If encryption is enabled, an encryption method must be selected. TLS and SSL are the available options. If you do not know which one your mail server requires, try one, and then the other.

User	The username to authenticate as on the SMTP mail server (optional).
Password	The password of the username specified above (optional).
Sender Email (required)	The email address that will appear in the From field of any sent emails. This email address sometimes needs to have a valid account on the SMTP mail server.
Use Recommended Office365 Settings	Enable this check box if you are connecting to an Office 365 SMTP server to use recommended settings for the connection. Follow Microsoft's Direct Send recommendations to setup Email configuration with an Office 365 SMTP server.

It is highly recommended that you test your SMTP settings before saving them. To do so, click the **Test Email Settings** button. You will be prompted for an email address to send the test message to. Upon submission, the Peer Management Center will attempt to send a test message using the specified settings.

Configuring Tags

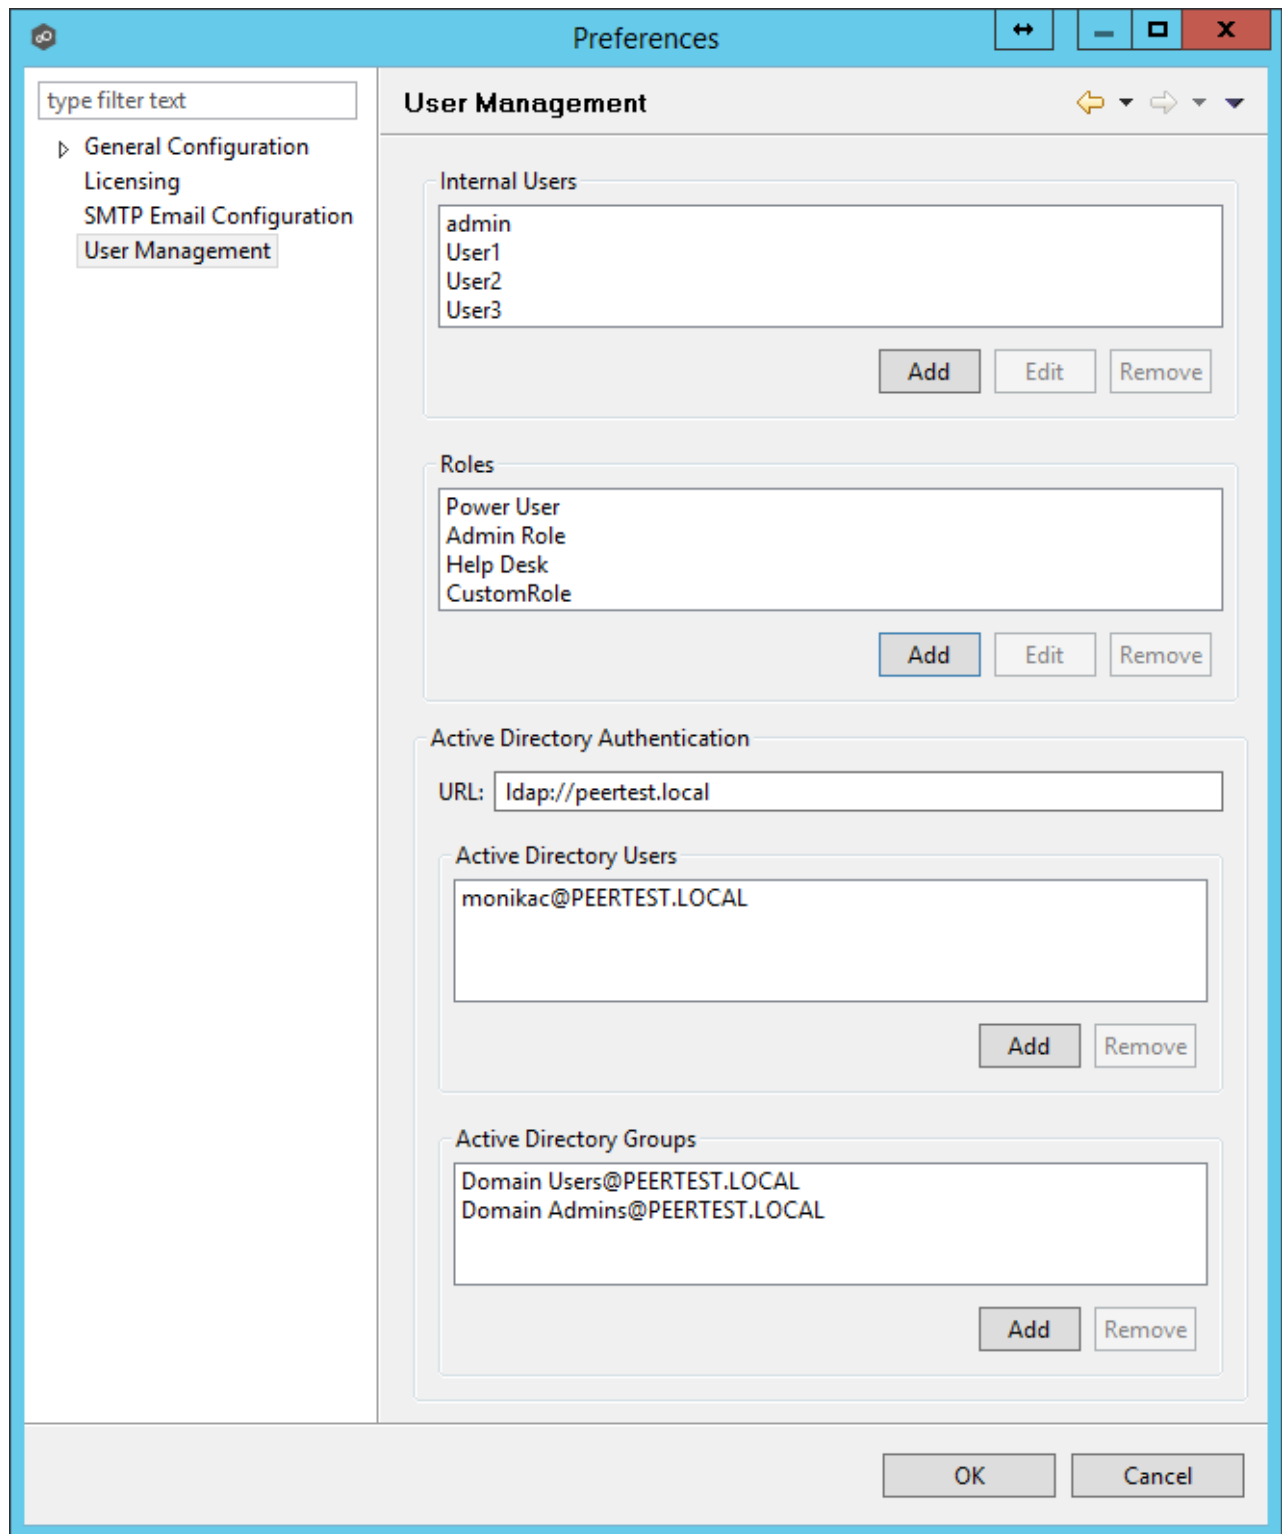
Tags can be used to categorize resources and customize a user's workspace or perspective. Tagging resources helps when managing large number of resources. The **Tags Configuration** window is the starting place for creating tags and categories that can later be assigned to resources. To set these values, click the **Window > Preferences** menu option or click the **User Preferences** button from the [Main View toolbar](#). See [Using Tags](#) for more details.



Configuring User Management

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

To access the **User Management** configuration page, click the **Window** menu, and then select **Preferences**. In the dialog that appears, select **User Management** from the tree node on the left. The following is displayed:



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

Internal Users

Adding an internal account requires a username, a password, an email address, and a selected role. For more details on the available roles, see the [How to Use](#) section. Once an account has been created, its username, password, email address, and role can be changed. The default **admin** user account password is **password**.

Roles

Roles define user permissions to access and edit resources in the Web Interface. There are three predefined roles with specific set of permissions: **Power User**, **Admin Role**, and **Help Desk**. To create a custom role or to edit a role, click the **User Management > Add** or **Edit** button in the **Roles** section. To delete a role, select the role from the **User Management > Roles** section, and then click **Remove**. Adding a custom role requires a name, display name, description, and base role. Roles can also be assigned existing tags to define the resources users with that role can view and edit. The following table outlines the Hub resources that each role can edit and view:

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

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

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

File Sync Advisory Alert View	Edit	Edit	
Expression List Dialog	Edit	Edit	

Active Directory Authentication

In addition to [Internal Users](#), the Peer Management Center also provides Active Directory user and group authentication. To configure Active Directory Authentication, provide the URL of the LDAP server on the network in the format

ldap://MYDOMAIN.LOCAL

or

ldaps://MYDOMAIN.LOCAL

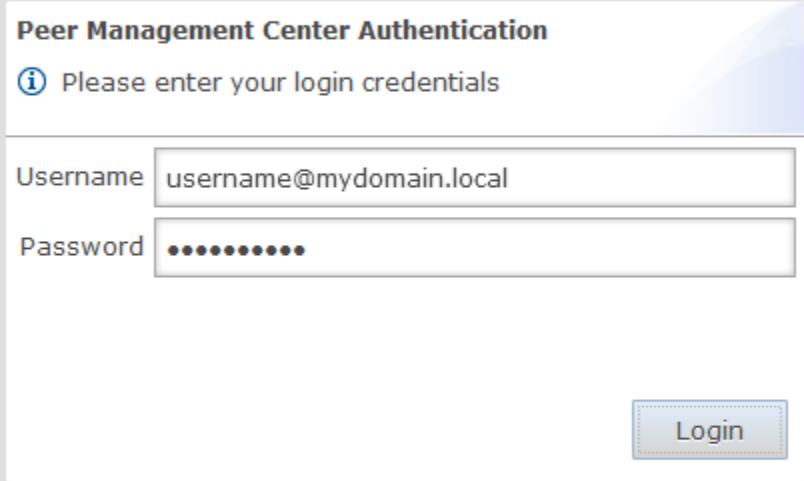
Next, add an Active Directory user or group by clicking the **User Management > Active Directory Users/Active Directory Groups > Add** button. The configuration dialog will require the Domain, Username or Group, and the [Role](#).

To delete an Active Directory User or Group, click the **User Management > Active Directory Users/Active Directory Groups > Remove** button.

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

[Username@MYDOMAIN.LOCAL](#)

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



The image shows a login form titled "Peer Management Center Authentication". Below the title is an information icon and the text "Please enter your login credentials". There are two input fields: "Username" with the text "username@mydomain.local" and "Password" with masked characters (dots). A "Login" button is located at the bottom right of the form.

Advanced Configuration

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

- [Advanced Peer Management Center Configuration](#)
- [Advanced Peer Agent Configuration](#)
- [Advanced File Collaboration Configuration](#)
- [Advanced Cloud Sync Configuration](#)

Advanced Peer Management Center Configuration

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

Topics include:

- [Custom TLS Integration](#)
- [File Filters](#)
- [Tags](#)

Custom TLS Integration

Overview

Peer Management Center supports the ability to use custom or private TLS certificates to connect Peer Agent to the Peer Management Center Broker. The Keytool certificate management utility will be used to store the key and certificate into a keystore file, which protects the private keys with a password.

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

For more information, see:

- [Use Existing Certificate](#)
- [Create New Certificate](#)

Perform the necessary commands using the keytool application bundled with your Peer Management Center or Peer Agent installation (Java 6).

Keytool location on Peer Management Center system:

C:\Program Files\Peer Software\Peer Management Hub\jre\bin

Keytool location on Peer Agent system:

C:\Program Files (x86)\Peer Software\Peer Agent\jre\bin

Peer Management Center Broker and Peer Agent Keystore Generation

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

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

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

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

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

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

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

plBroker4321	The password you assign to the new Broker keystore.
broker.ks	Destination keystore that will be created containing the custom/private certificate.
HubCert.pfx	Custom/private certificate being imported into the new keystore.

PASSWORD	The password of the custom/private certificate, if it exists. If you omit the -srcstorepass command you will be prompted for the certificate password if needed.
ALIAS	The Alias of the custom/private certificate you discovered in Step 1 above.
broker	The Alias of the new keystore containing the custom/private.

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

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

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

plClient4321	The password you assign to the new Broker keystore.
client.ks	Destination keystore that will be created containing the custom/private certificate.
AgentCert.pfx	Custom/private certificate being imported into the new keystore.
PASSWORD	The password of the custom/private certificate, if it exists. If you omit the -srcstorepass command you will be prompted for the certificate password if needed.
ALIAS	The Alias of the custom/private certificate you discovered in Step 1 above.
client	The Alias of the new keystore containing the custom/private.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Copy the Generated Keystore Files into Their Appropriate Location

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

broker.ks

broker.ts

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

client.ks

client.ts

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

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

Perform the necessary commands using the keytool application bundled with your Peer Management Center or Peer Agent installation (Java 6).

Keytool location on Peer Management Center system:

PMC_HUB_INSTALLATION_FOLDER\jre\bin

Keytool location on Peer Agent system:

PEER_AGENT_INSTALLATION_FOLDER\jre\bin

Broker Keystore Generation

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

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

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

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

Example:


```

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

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

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

```

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

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

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

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

Example:

```

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

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

```

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

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

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

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

Example:

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

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

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

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

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

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

Example:

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

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

Optional: List the certificates in the broker keystore.

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

Example:

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

Keystore type: jks
Keystore provider: SUN

Your keystore contains 1 entry

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

*****
*****

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

Verify Client Certificate

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

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

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

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

Example:

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

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

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

Example:

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

Optional: List the certificates in the client keystore.

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

Example:

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

Keystore type: jks
Keystore provider: SUN

Your keystore contains 1 entry

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

*****
*****

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

Copy the Generated Keystore Files into Their Appropriate Location

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

broker.ks

broker.ts

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

client.ks

client.ts

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

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

File Filters

File filters govern the inclusion and exclusion of files under the [Watch Set](#). Included files are subject to scan and event detection, while excluded files are not. Initially, all files are included and no files are excluded, except for the internal expressions listed in Auto Excluded Filter below.

Filtration can be configured with wildcard expressions to more easily cover well-known file extensions or names that follow established patterns. When a single expression is insufficient for configuring filtration, multiple expressions may be supplied. You can also filter file based on a file's last modified time and file size.

Usage Notes

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

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 Watch Set triggers a target addition.

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.

Auto-excluded Filter

The following wildcard expressions are automatically applied as exclusion expressions and cannot be changed:

- 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

Configuration Notes

The excluded and included file name filters take one or more standard wildcard expressions that are combined by performing a logical OR of each wildcard expression.

Standard Wildcard Expressions

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

The following examples show the use of wildcard syntax to enter a file exclusion or inclusion:

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

ext Filter files that contain the string ext

ext* Filter files that start with the string ext

Peer Management Center also supports the use of complex regular expressions, e.g., **<<regEx>>**. These expressions can be used for either included or excluded patterns. For information on where to enter a regular expression, see [Configuring File Filters](#)

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

Filter Expressions

The Peer Management Center provides the ability to filter lists throughout the Peer Management Center interface. Filter expressions can help you quickly find jobs, Agents, and sort through summary reports. The search results of your filter will be displayed in the window below the expression.

Basic

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

Double quotes around the word(s) you are searching for are required if you want an exact word match or the words contain a space. For example, if you want to search for the words "North America" the two words must be contained in double quotes. If you want to search for the word "agent" only without showing "USAgent" or "Agent2015" in results, the word must be contained in double quotes.

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

Refine

Use predefined labels to specify in which field your filter word should appear. Use the following format to take advantage of labels in your filter expression: `<label>:<search string>`. List of possible labels include:

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

Operators

You can create more sophisticated filters by using operators. Operators allow you to combine multiple simple expressions into a single, compound expression. Supported operators are: **OR**, **AND**, and **NOT**. Typing `tag:Americas AND sales` in the Filter Expression will show only Agents with the word Americas in their tag(s) **AND** the word sales in their name, tags, or tag

categories. Parentheses can be used to build more complex expressions by grouping simple expressions.

Remove/Clear Filters

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

Save and Manage Filters

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

Examples:

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

[Sales](#)

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

[cat:Location](#) AND [tag:"North America"](#)

This filter will show all Agents with the word Sales in their name, tag name, and tag category and with a tag that has "North America" in the tag name and "Location" in the tag category.

[Sales](#) AND ([cat:Location](#) AND [tag:"North America"](#))

Using Tags

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

Step 1: Create Tags and Categories

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

Step 2. Assign Tags

During Job Creation

Assign tags during the creation of a brand new job from the [Tags](#) window of the File Collaboration Configuration dialog.

Edit an Existing Job

Tags can be assigned to individual jobs by right-clicking on the job, selecting **Edit Configuration(s)**, and navigating to the **Tags** window of the File Collaboration Configuration dialog.

Assign Tags to One or More Resources

- To assign tags to one or more resources, click the **Assign Tags** button from the [Main view](#), [Jobs view](#), or [Agent Summary view](#) toolbars.
- In the **Assign Tags** dialog, click the **Tags** radio button.
- Select the tag that needs to be assigned to one or more resources.
- Click the **Edit** button to the right.
- From the **Unassigned Resources** table on the left side, select the resources that need to be assigned the selected tag and click the right-arrow button (Add One) to move it to the table on the right side (hold down the Shift key on the keyboard when selecting resources to select more than one).
- Click the **Save** button to commit your changes and close the dialog.
- Repeat the steps above for all the tags that need to be assigned to one or more resources.

Resources to One or More Tags

- To assign resources to one or more tags, click the **Assign Tags** button from the [Main view](#), [Jobs view](#), or [Agent Summary view](#) toolbars.
- In the **Assign Tags** dialog, click the **Resources** radio button.
- On the left hand side, click inside the **Resource Name Filter** or **Type Filter** fields and press the **CTRL + Space** keys on the keyboard to list all possible filters and predefined labels, which can be selected to refine your search quickly.
- Select the resource that needs to be assigned to one or more tags.
- Click the **Edit** button to the right.

- From the **Unassigned Tags** table on the left side, select the tags that need to be assigned the selected Resource and click the right-arrow button (Add one) to move it to the table on the right side (hold down the Shift key on the keyboard when selecting tags to select more than one).
- Click the **Save** button to commit your changes and close the dialog.
- Repeat the steps above for all the Resources that need to be assigned to one or more tags.

Assign Tags to User Roles

User roles can be assigned tags that customize a user's Jobs perspective when they log in via the [Web Interface](#). For example, in a very large deployment scenario, a user that is part of the Help Desk role can be assigned tags that limit their view to only jobs that are part of their region. To achieve this:

- Create tags and categories as outlined in Step 1 above.
- Assign tags to one or more jobs as outlined in Step 2 above.
- Go to [User Management](#) in the [Preferences](#) dialog.
- Select the desired role to which you wish to assign specific job tags.
- Click the **Edit** button.
- In the **Tags** window, from the **Unassigned Tags** table on the left side, select the tags that need to be assigned the selected r and click the right-arrow button (Add one) to move it to the table on the right side (hold down the Shift key on the keyboard when selecting tags to select more than one).
- Click **OK** to commit your changes and close the dialog and close the **Preferences** dialog.
- The user will only see the jobs which were tagged in the user's role.

Step 3. Filter Resources Using Tags

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

Filter Jobs

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

Example:

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

tag:"North America" AND cat:Location

Filter Agents

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

Advanced Peer Agent Configuration

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

To access, right-click any connected Peer Agent, and then select **Edit Agent Configuration**. The **Peer Agent Configuration** dialog will be displayed, with three pages of available configuration items. In order for any configuration change to take effect, the selected Peer Agent must be restarted. If no jobs are running, you will have the option of restarting the Peer Agent at the close of the configuration dialog.

WARNING: Changes to any option on the three pages of this dialog may result in problems when the Peer Agent starts. Ensure all settings are correct before saving the dialog and restarting the selected Peer Agent.

Topics include:

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

Broker Configuration

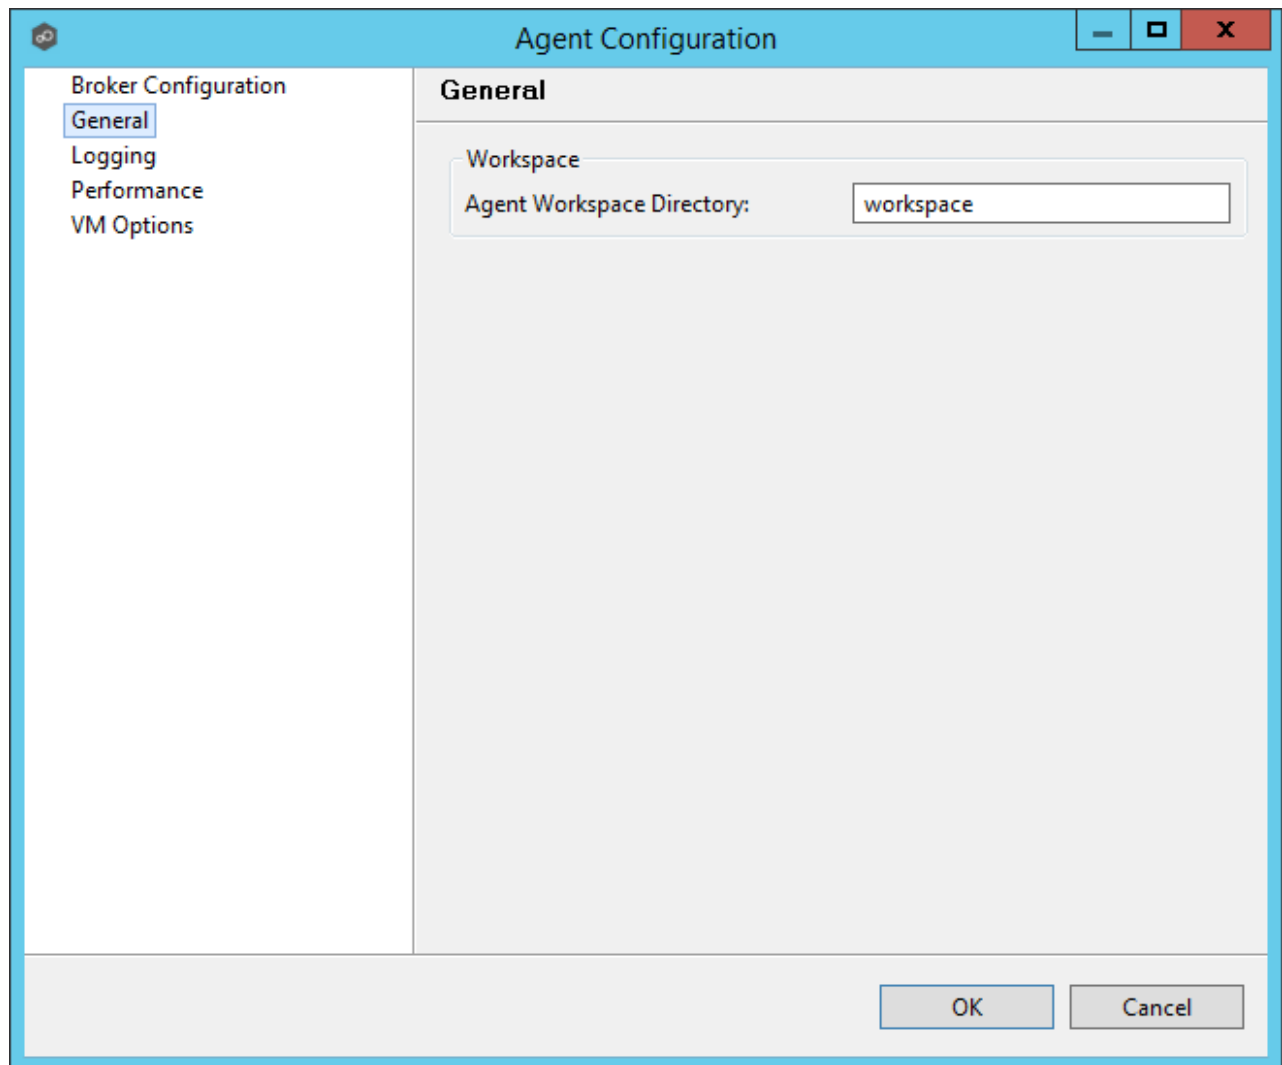
The screenshot shows a window titled "Agent Configuration" with a sidebar on the left containing the following options: "Broker Configuration" (selected), "General", "Logging", "Performance", and "VM Options". The main area is titled "Broker Configuration" and contains a warning message: "WARNING: Changes to this page may make the Agent unable to start. These changes will only take affect after the Agent is restarted." Below the warning are four configuration fields: "Primary Broker Host" with the value "qalab1dc", "Connection Type" with a dropdown menu showing "ssl", "Broker Port" with a value of "61617" and up/down arrows, and "Use Compression" with a checked checkbox. At the bottom right are "OK" and "Cancel" buttons.

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

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

Broker Port	The port on which to communicate with the Peer Management Center Broker.
Use Compression	When enabled, all communication between the selected Peer Agent and the Peer Management Center Broker will be compressed.

General

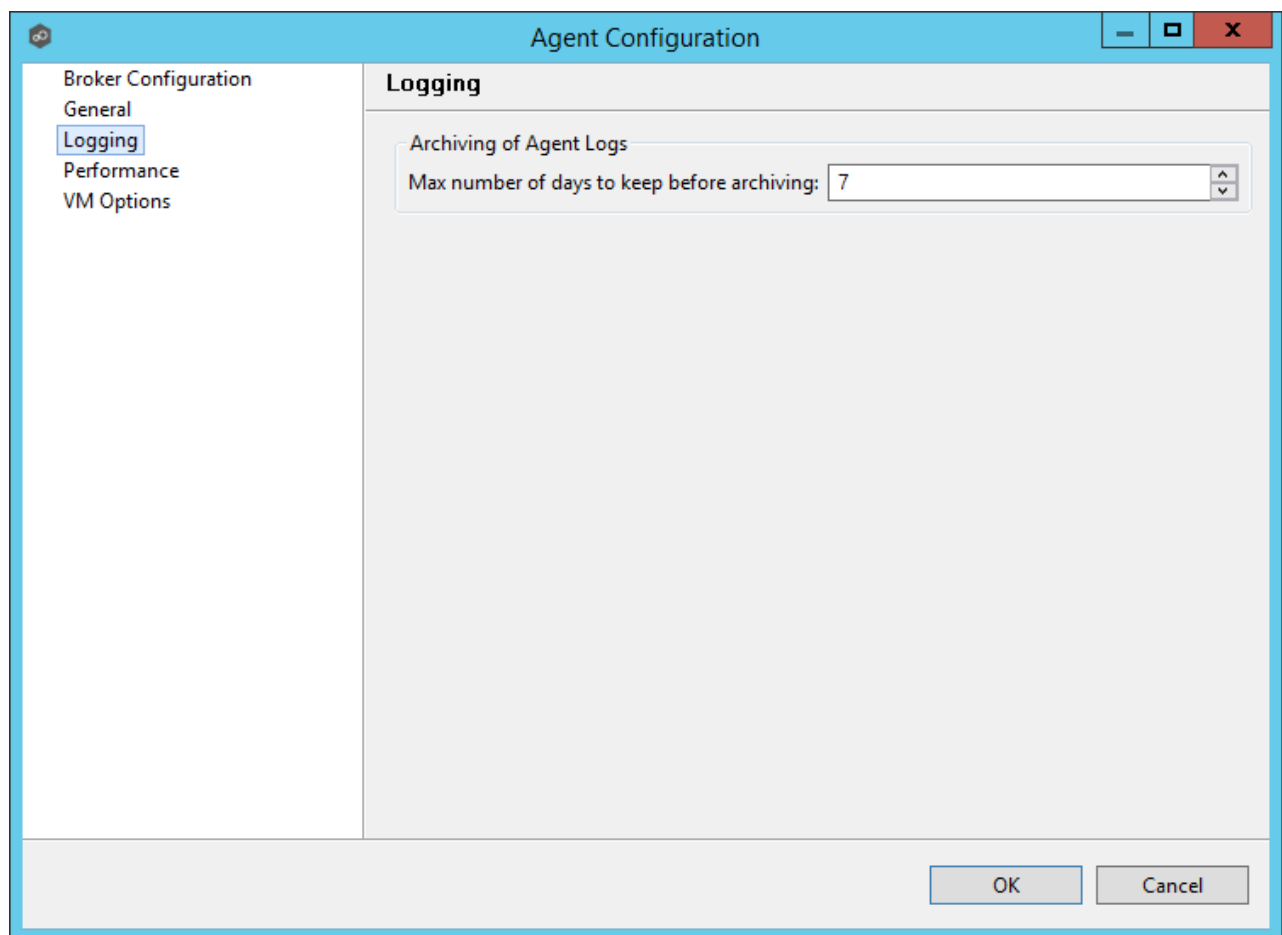


The image shows a screenshot of the 'Agent Configuration' dialog box, specifically the 'General' tab. The dialog has a blue title bar with the text 'Agent Configuration' and standard window controls (minimize, maximize, close). On the left side, there is a vertical list of configuration categories: 'Broker Configuration', 'General' (which is selected and highlighted with a blue border), 'Logging', 'Performance', and 'VM Options'. The main area of the dialog is titled 'General' and contains a single configuration item: 'Agent Workspace Directory:'. This item is displayed as a text label followed by a text input field containing the value 'workspace'. At the bottom right of the dialog, there are two buttons: 'OK' and 'Cancel'.

Workspace

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

Logging

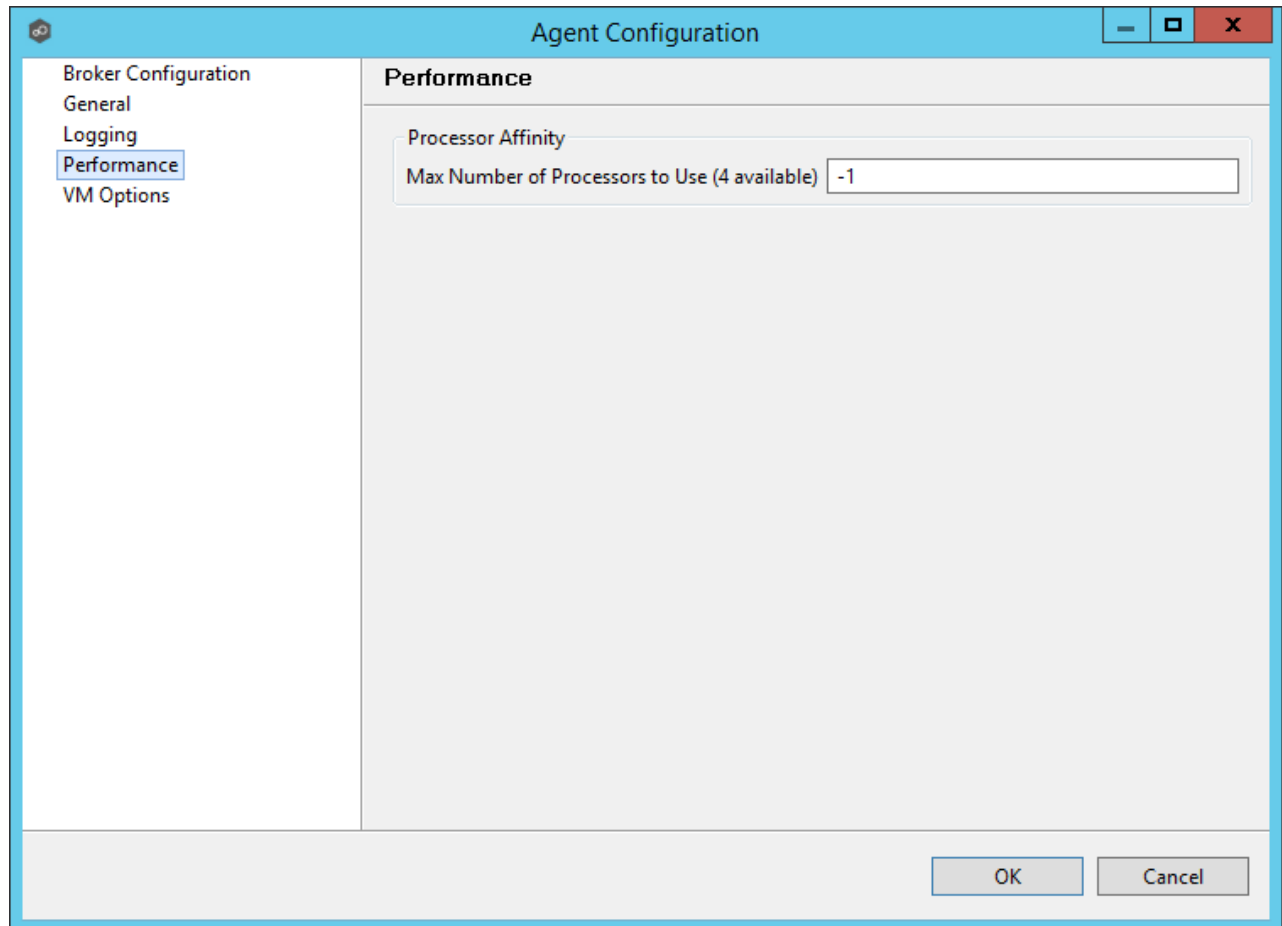


Archiving of Agent Logs

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

keep before archiving	
----------------------------------	--

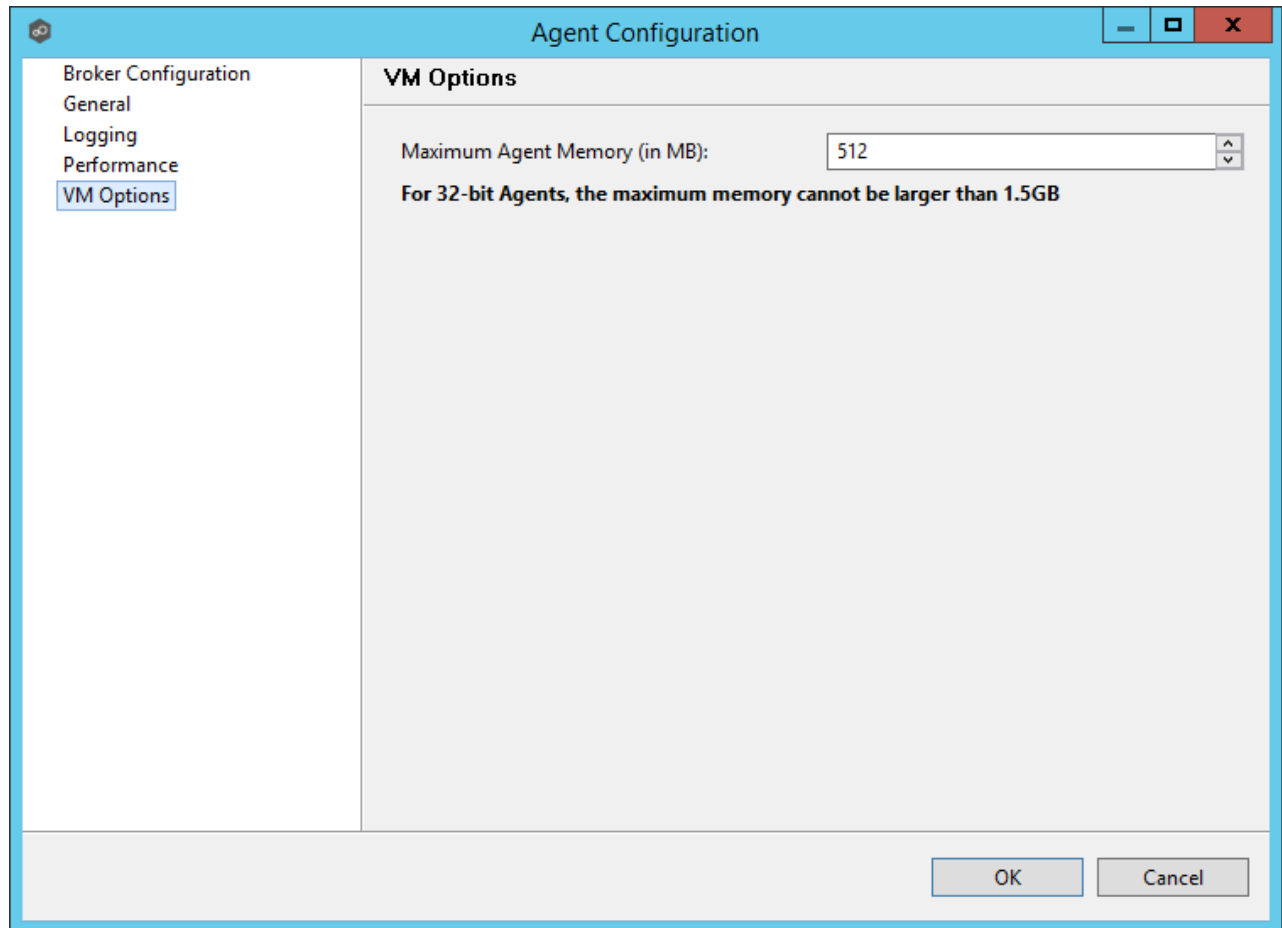
Performance



Processor Affinity

Max Number of Processors to Use (x available)	The maximum number of processors the Peer Agent service will be able to use. If set to -1, all processors will be available. The caption for this setting will display how many total processors are available.
--	---

VM Options



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

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

Advanced File Collaboration Configuration

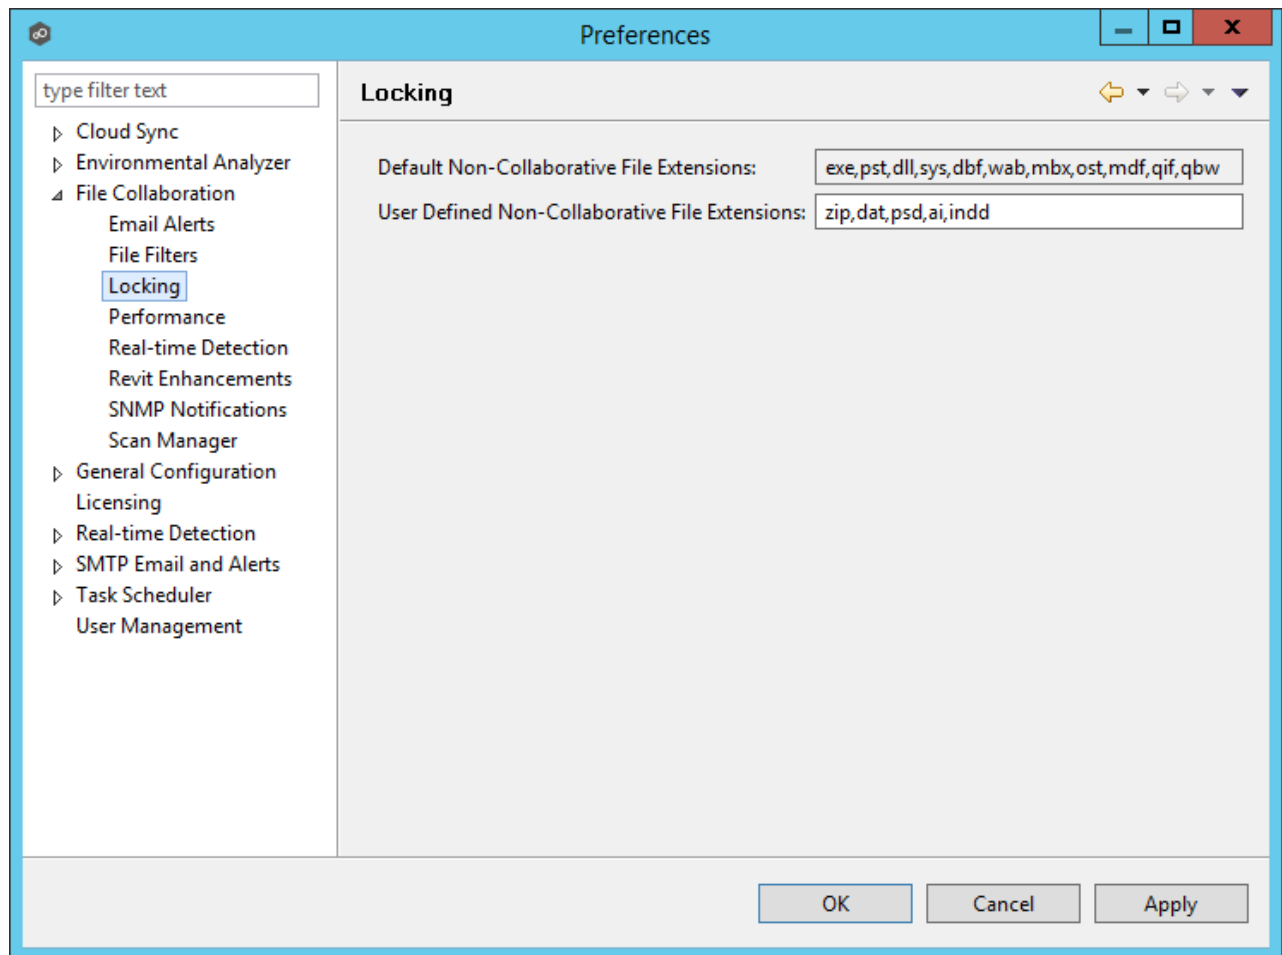
The topics in this section provide information on advanced functionality and configuration options available in File Collaboration.

Topics include:

- [Locking](#)
- [Revit Enhancements](#)
- [Scan Manager](#)
- [Advanced Windows Real-time Options](#)

Locking

An option is available to mark certain file types as non-collaborative, changing the way locks on the specified file types are handled. These settings are configured on a global level for all file collaboration jobs and are critical for certain file types so that the file collaboration solution is able to correctly read any part of these files, ensuring any managed database type files are synchronized in a consistent and usable state. To view and modify these settings, click the **Window** menu from within the Peer Management Center, and select **Preferences**. On the left-hand side of the dialog that pops up, open the tree node titled **File Collaboration** and select **Locking**. The following screen will be displayed.

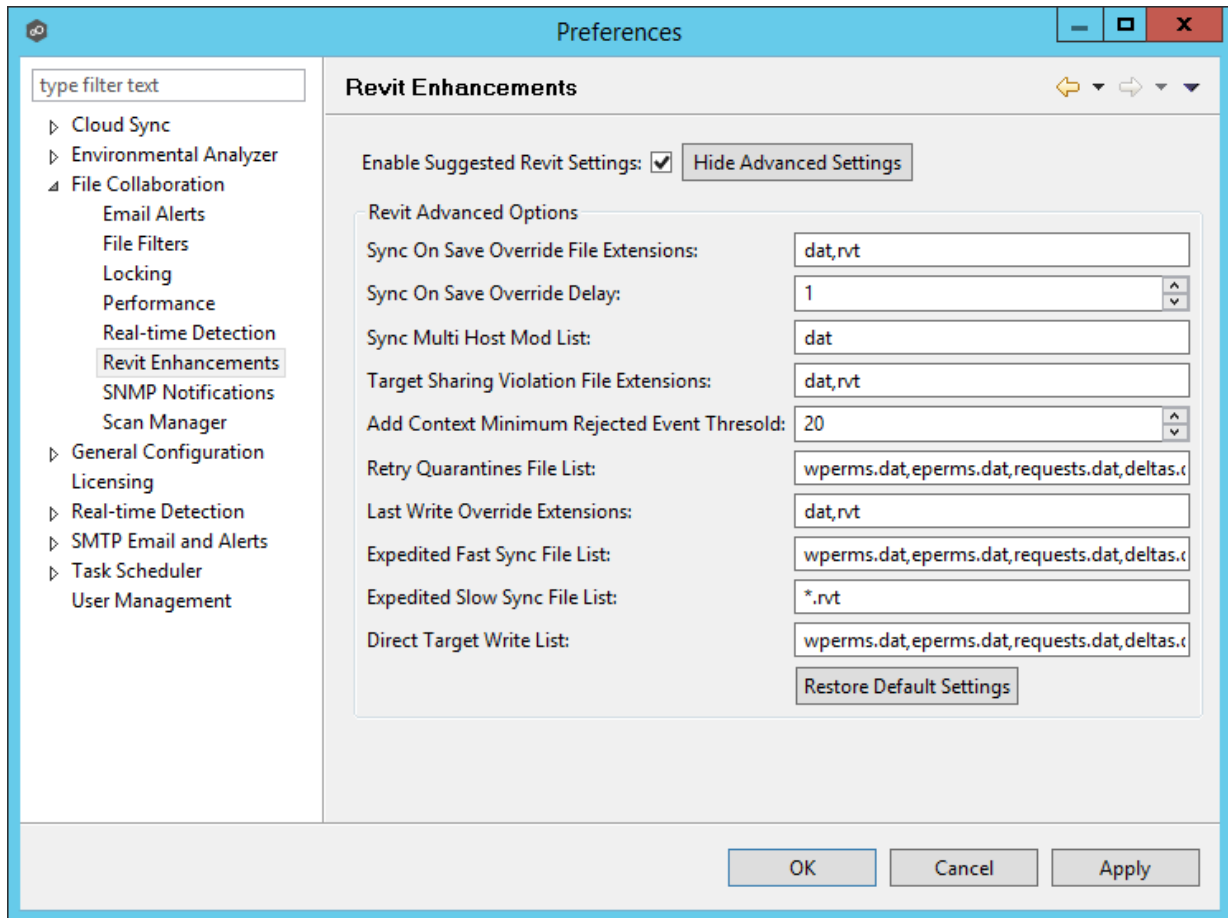


Available options are as follows:

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

Revit Enhancements

Revit Enhancements enable the Expedited Sync Queue for files specified in the Expedited Sync Queue File List.



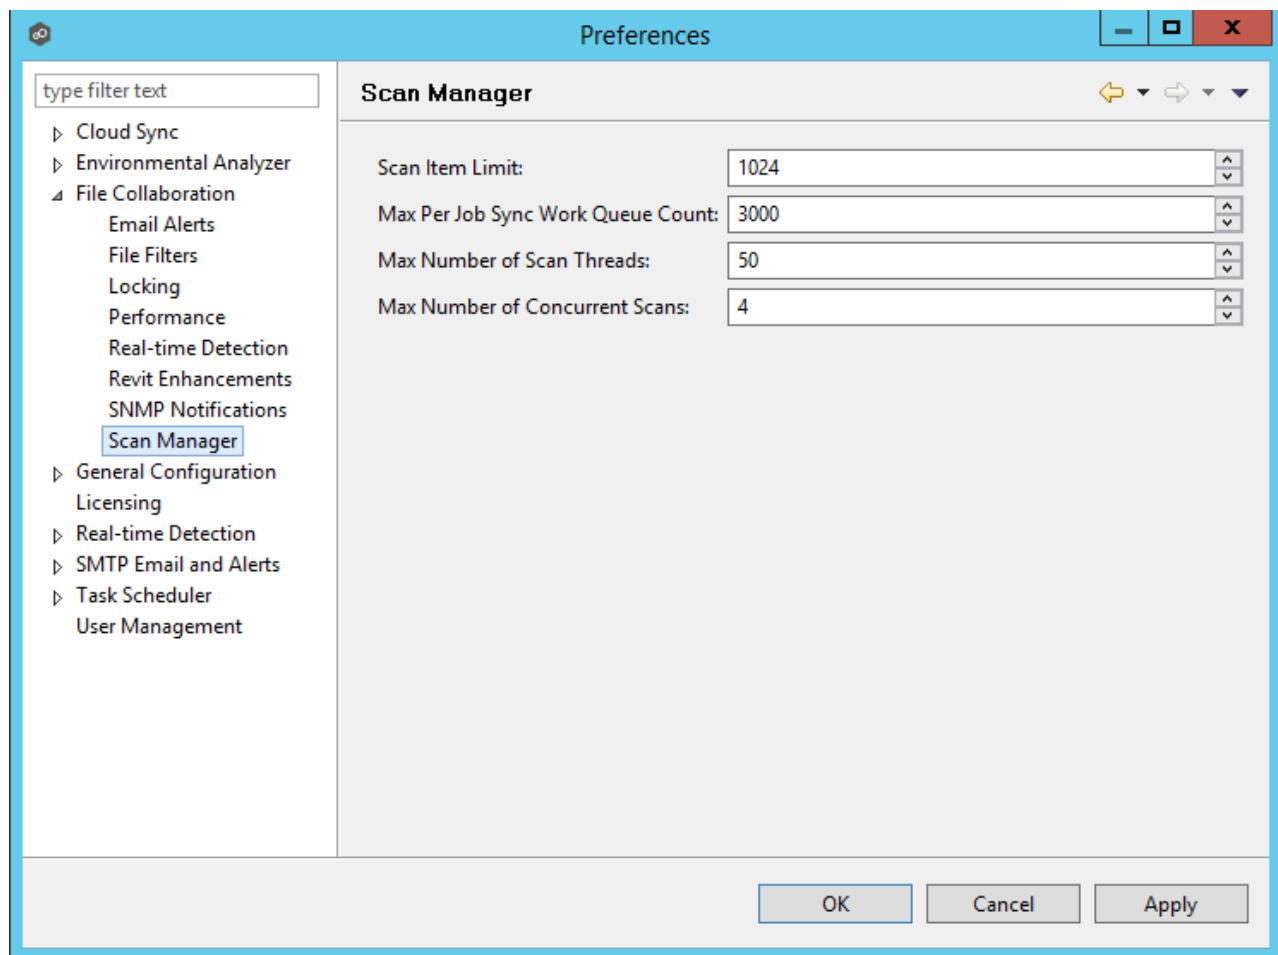
Click the **Show Advanced Settings** button to reveal the recommended settings.

Sync On Save Override File Extension	Extensions configured here will overwrite the "Sync. On Save" values configured in the interface for the job. In addition, these extensions use the delay value in Sync On Save Override Delay setting instead of the delay value configured in the interface. If no delay value is set, it will default to using a one second delay. Extensions configured in this list will still be processed via Sync. On Save even if they also exist in the user defined non-collaborative extension list (under the Window > Preferences menu option). Extensions in the normal Sync. On Save list that also exist in this list will not be processed.
Sync On Save	The "Sync. On Save" delay value in seconds that will only apply to the internal list of extensions listed in the Sync On Save Override File Extension field.

Override Delay	
Sync Multi Host Mod List	Extensions configured here will not be quarantined if they are modified on two hosts simultaneously. The file with the latest modified time stamp will win.
Target Sharing Violation File Extensions	This is an option to retry setting the target lock when receiving error code 32 for the specified list of extensions. This may be useful for file types like .one (OneNote), .rvt (Revit), and .dat (associated Revit files), that don't sustain a handle when the user has the file open.
Add Context Minimum Rejected Event Threshold	The number of bulk add files that PeerLink can process immediately before batching the remainder of the files and process them in a single thread.
Retry Quarantine File List	Quarantined files that are in this list will be automatically removed and flagged as unsynchronized and will be retried every second after a delay period (delay is configured by "fc.retryQuarantinesDelay"). Any change event that is detected for the files will trigger a scan of the files where the newest file will win. This list can contain file names (wperms.dat, eperms.dat, requests.dat, deltas.dat, users.dat) or extensions (*.dat,*.abc).
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.

Scan Manager

A number of options are available to tune the way scans are performed for all [file collaboration jobs](#). These settings are configured on a global level. To view and modify these settings, click on the Window menu within the Peer Management Center, and select **Preferences**. On the left-hand side of the dialog that pops up, open the tree node titled **File Collaboration** and select **Scan Manager**. The following screen will be displayed.



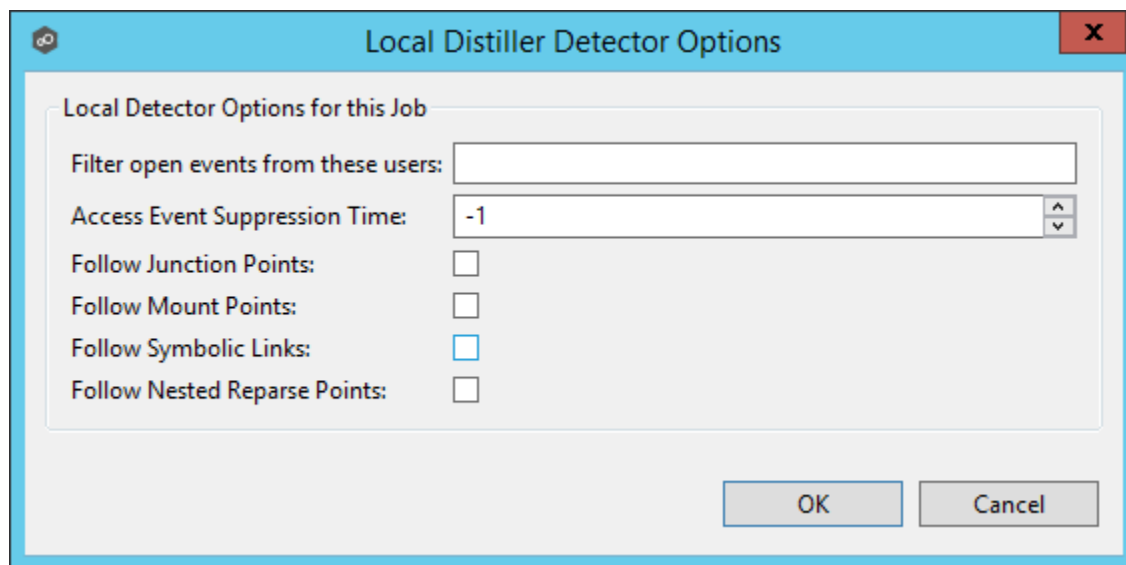
Available options are as follows:

Scan Item Limit	The maximum number of file and folder scan results that are returned in one scan iteration during a job's initial scan. This value is used to constrain the amount of memory used when performing initial scans with a large number of sessions.
------------------------	--

Max Sync Work Queue Count	The maximum number of pending file transfers (as a result of the initial scan) that are queued in memory before pausing the current scan. This value only has an effect on sessions that require a massive amount of initial synchronization.
Max Number of Scan Threads	The maximum number of threads that can be created for use when scanning folders and files. This number should be set to at least the number of jobs that you are running.
Max Number of Concurrent Scans	The maximum number of scan threads that can be actively working at the same time. This differs from the Max Number of Scan Threads in that not all created scan threads can be simultaneously doing work.

Advanced Windows Real-time Options

The real-time detection options available for local Windows file servers can be modified on the **Participants** page of the job configuration dialog by selecting one of the participating hosts configured with the **Default** event detector in the bottom list, then clicking **Edit Detector Settings**. The following dialog will appear.



Available options are as follows:

Filter open events from these users	A comma-separated list of user account names from which all opens and closes will be ignored. Ideal for filtering out events from backup and/or archival services by filtering on the
--	---

	username under which a backup and/or archival service is running.
Access Event Suppression Time	Represents number of seconds an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default vault is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time an open event will be delayed based on the load of the system.
Follow Junction Points	Enables junction point support for the selected Windows file server.
Follow Mount Points	Enables mount point support for the selected Windows file server.
Follow Symbolic Links	Enables symbolic link support for the selected Windows file server.
Follow Nested Reparse Points	Enables support for nested reparse points on the selected Windows file server.

For details on either Junction Point or Symbolic Link support, please contact support@peersoftware.com.

Advanced Cloud Sync Configuration

You can modify the following Cloud Sync settings:

- [Cloud Sync](#)
- [Cloud Platform Credentials](#)
- [Database Connections](#)
- [Email Alerts](#)

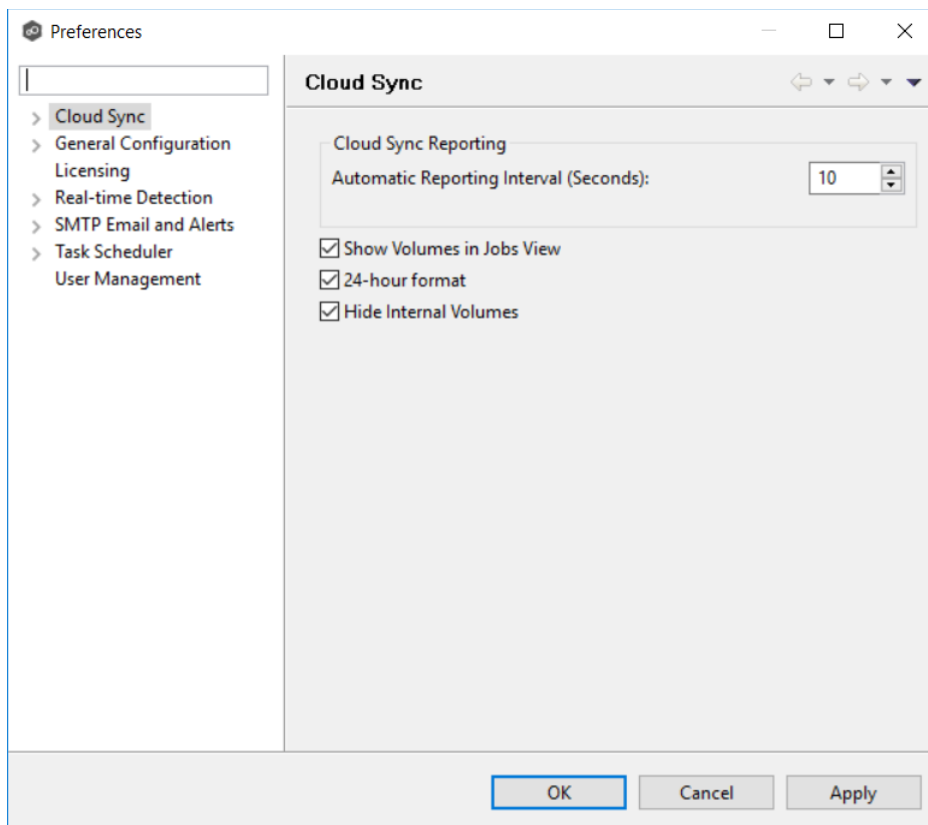
- [File Filters](#)
- [Performance](#)
- [Scan Manager](#)
- [Sync and Retention](#)

Modifying Cloud Sync Settings

Cloud Sync settings control the overall performance of all Cloud Sync jobs. Settings are defined globally in **Preferences** in the **Cloud Sync** page.

To modify the Cloud Sync settings:

1. From the **Window** menu, select **Preferences**.
2. Select **Cloud Sync** in the navigation tree.



3. Modify the settings as needed.

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

Creating Cloud Platform Credentials

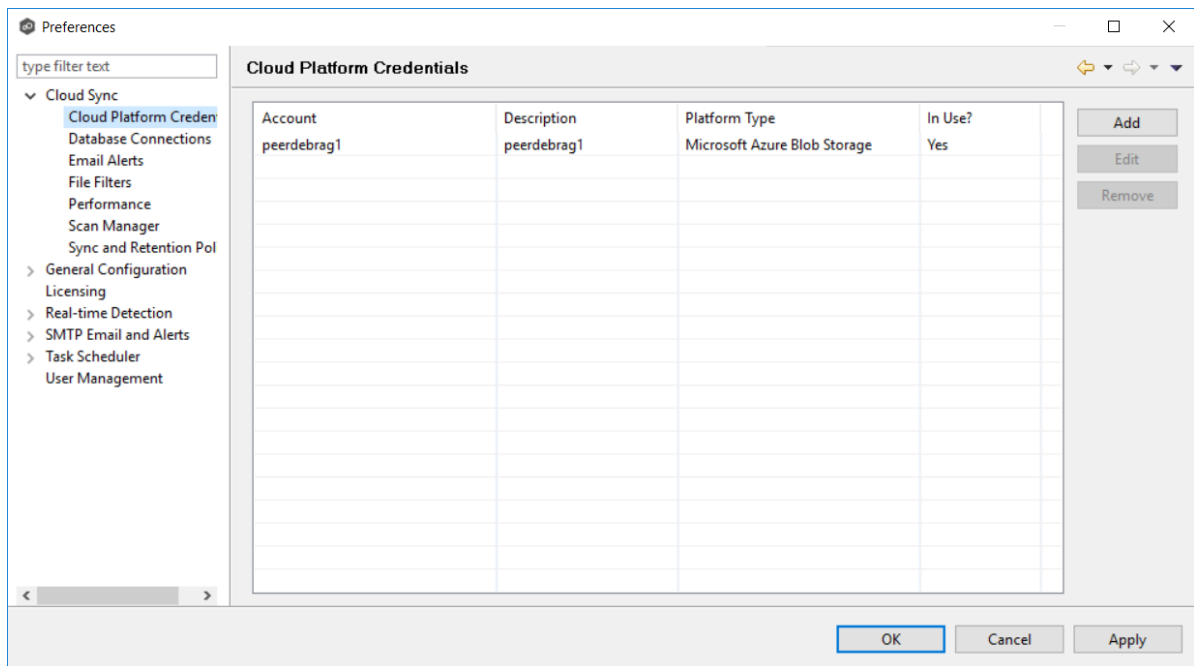
Cloud Platform credentials are defined globally in **Preferences** in the **Cloud Platform Credentials** page. The page lists the existing user names and passwords used to grant access to your Azure Storage account. You can view, add, edit, and remove credentials.

When you create a job, you can select existing credentials to apply to the job or you can create a new filter and apply it to the job.

Note: You cannot modify or delete credentials when they are applied to a job.


To create new cloud platform credentials:

1. From the **Window** menu, select **Preferences**.
2. Expand **Cloud Sync** in the navigation tree, and then select **Cloud Platform Credentials**.



3. Click the **Add** button.

The **Storage Account** dialog appears.



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

Description	Enter a name for the credentials.
Account	Enter the name of the Azure Storage account, which can be found in the Azure Portal.

Shared Key

Enter one of the shared keys of the Azure Storage account, which can be found in the Azure Portal.

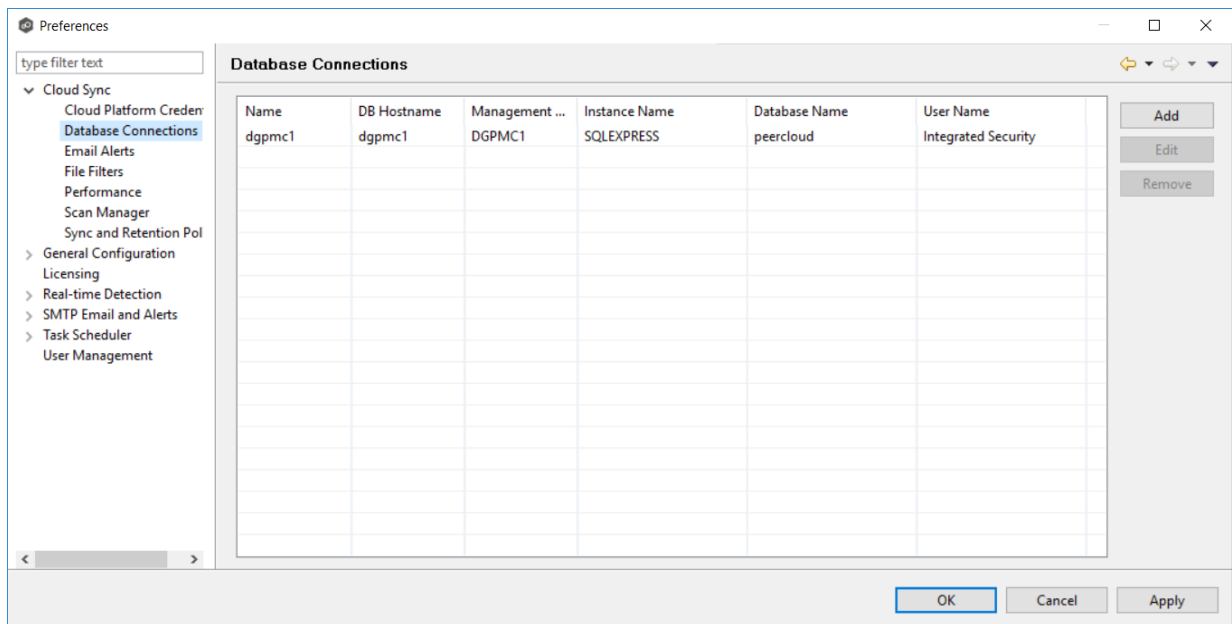
Creating Database Connections

Databases are used to track files and folders that have been replicated, individual file versions, and snapshots. Connections to your Microsoft SQL Server databases are defined globally in **Preferences** in the **Data Sources** page. You can view, add, edit, and remove database connections.

Note: You cannot modify or delete a database connection when it applied to a job.

To create a new database connection:

1. From the **Window** menu, select **Preferences**.
2. Expand **Cloud Sync** in the navigation tree, and then select **Database Connections**.



3. Click the **Add** button.

The **Add Database Connection** dialog appears.

Add Database Connection

Configure connection information for MS SQL Server

Database Connection Name:

Management Agent:

DB Hostname:

Port:

Instance Name:

Database Name:

Authentication: ☒ Integrated ☐ Credentials

User Name:

Password:

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

Database Connection Name	Enter a name for this database connection.
Management Agent	Select the Agent that will use this connection.
DB Hostname	Enter the name of the SQL Server hosting the database. If the database is installed on the Agent server itself, enter the name of the Agent server.
Port	Enter the port to be used to communicate with the specified SQL Server. If not defined, the connection defaults to port 1433.
Instance Name	Enter the database instance name to use on the specified SQL Server. If no named instances are installed on the specified SQL Server, leave this empty.

Database Name	Enter the name of the database that Cloud Sync will create. The default name is "peercloud" but it can be changed to a name that follows your company's naming conventions.
Authentication	Select Integrated if the Agent service account is granted admin rights on the selected SQL instance. Otherwise, select Credentials to enter the user name and password of a database administrator.
User Name	Enter the user name of the database administrator account to be used by Cloud Sync. This can be a locally-defined account such as "sa" or a domain account.
Password	Enter the password of the database administrator account to be used by Cloud Sync.

Creating Email Alerts

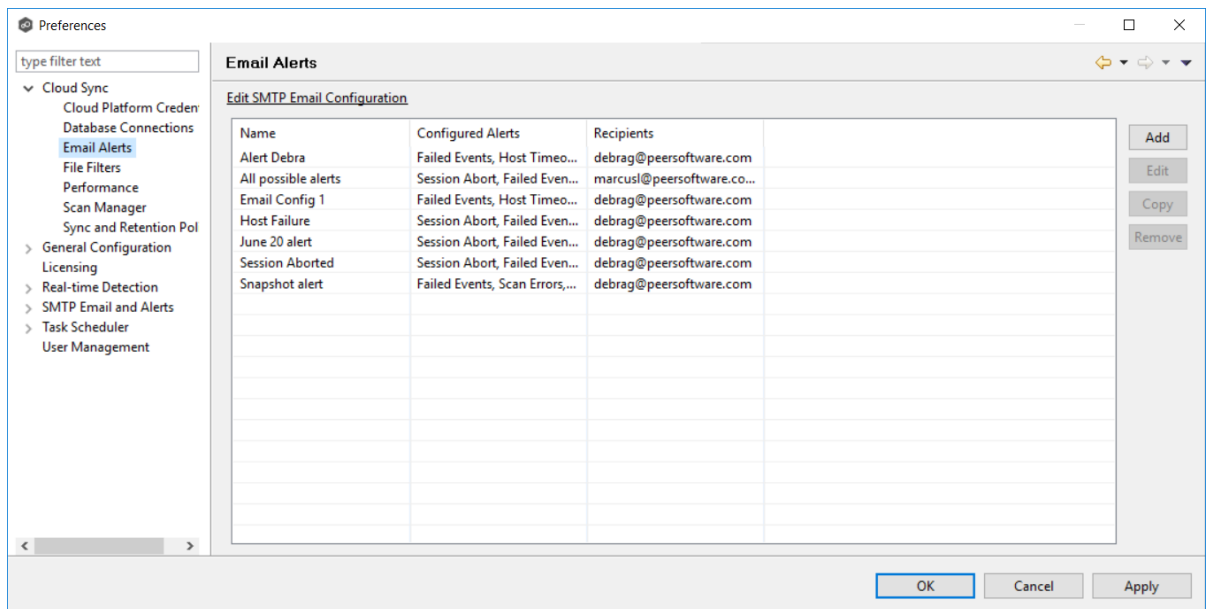
Email alerts can be sent to notify users when a certain type of event occurs, for example, session abort, host failure, system alert. Email alerts are defined globally in **Preferences** in the **Email Alerts** page. You can view, add, edit, and remove alerts. See [Configuring Email Alerts](#) for more information.

When you create a job, you can select an existing email alert to apply to the job or you can create a new alert and apply it to the job.

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

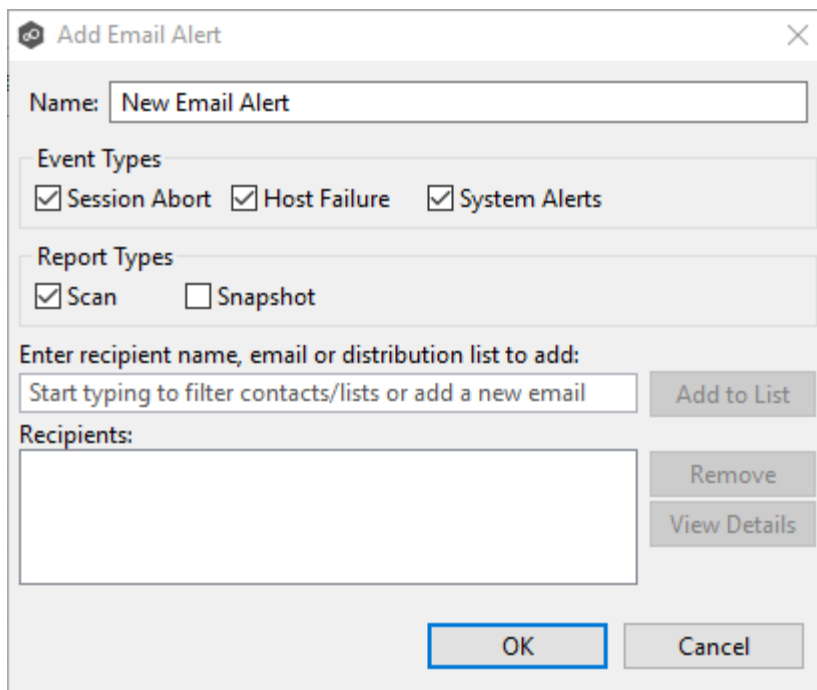
To create an email alert:

1. From the **Window** menu, select **Preferences**.
2. Expand **Cloud Sync** in the navigation tree, and then select **Email Alerts**.



- Click the **Add** button.

The **Add Email Alert** dialog appears.



- Enter the required values, and then click **OK**.

Name	Enter a name for the alert.
-------------	-----------------------------

Event Type	<p>Select an event type:</p> <ul style="list-style-type: none"> • Session Abort: Sends an alert when the Cloud Sync job stops unexpectedly. • Host Failure: Sends an alert when the Management Agent of a Cloud Sync job disconnects or stops responding. • System Alerts: Sends an alert when a system event such as low memory or low hub disk space occurs. <p>The event type determines what will trigger the email alert to be sent.</p>
Report Types	<p>Select a report type:</p> <ul style="list-style-type: none"> • Scan: Sends scan statics after a scan has completed. • Snapshot: Sends the status of the snapshot after the snapshot is triggered.
Recipient name, email or distribution list	Enter an email address for a recipient or search for an existing contact or distribution list.
Recipients	Displays a list of the email addresses, contacts, and distribution lists that will receive the alert and reports.

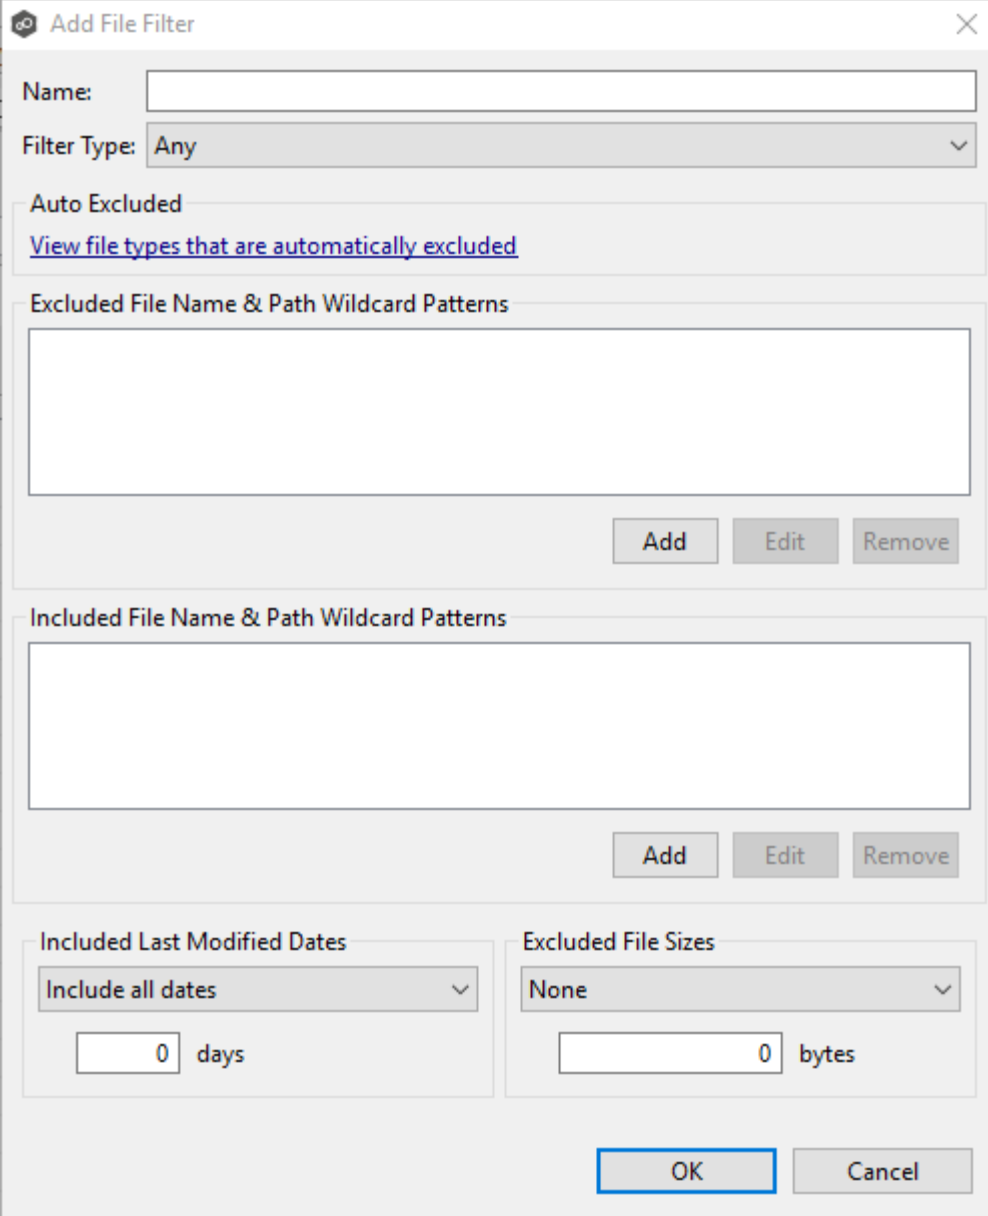
Creating File Filters

File filters allow you to specify files and folders to exclude from or include in a job. File filters are defined globally in **Preferences** in the **File Filters** page. You can view, add, edit, and remove filters. See [File Filters](#) for more information.

When you create a job, you can select an existing file filter to apply to the job or you can create a new filter and apply it to the job.

Note: You cannot modify or delete a file filter when it is being used by a job.

To create a file filter:



The "Add File Filter" dialog box is shown. It has a title bar with a close button. The main area contains several sections: "Name:" with a text input field; "Filter Type:" with a dropdown menu set to "Any"; "Auto Excluded" with a link "View file types that are automatically excluded"; "Excluded File Name & Path Wildcard Patterns" with a large text area and "Add", "Edit", and "Remove" buttons; "Included File Name & Path Wildcard Patterns" with a large text area and "Add", "Edit", and "Remove" buttons; "Included Last Modified Dates" with a dropdown set to "Include all dates" and a "0 days" input; and "Excluded File Sizes" with a dropdown set to "None" and a "0 bytes" input. At the bottom are "OK" and "Cancel" buttons.

Add File Filter

Name:

Filter Type: Any ▾

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

Excluded File Name & Path Wildcard Patterns

Add Edit Remove

Included File Name & Path Wildcard Patterns

Add Edit Remove

Included Last Modified Dates
Include all dates ▾

days

Excluded File Sizes
None ▾

bytes

OK Cancel

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

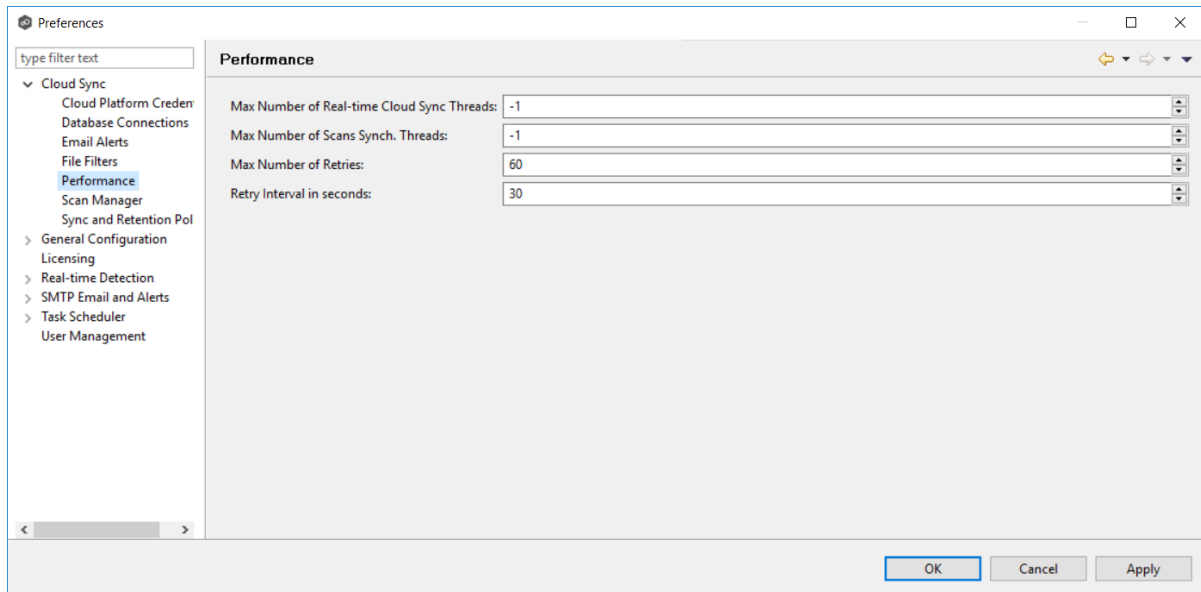
See [Configuring File Filters](#) for information on entering values for the fields.

Modifying Cloud Sync Performance Settings

Performance settings allow you to adjust Cloud Sync performance. Settings are defined globally in **Preferences** in the **Performance** page.

To adjust the performance settings:

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



3. Modify the following settings as needed:

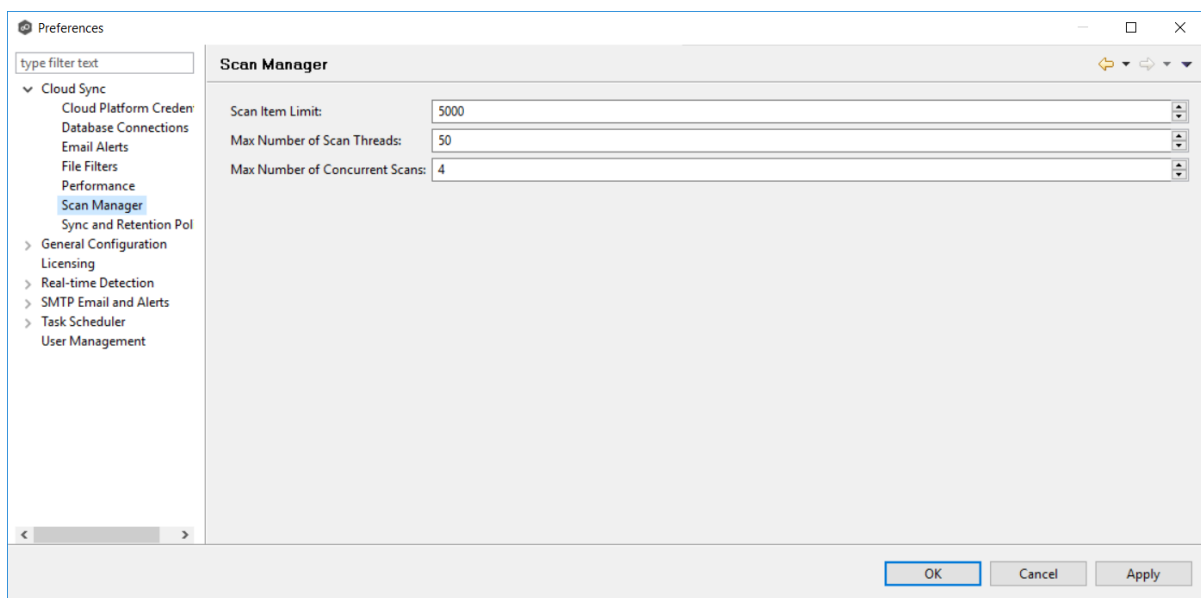
Max Number of Real-time Cloud Sync Threads	Enter the maximum number of threads available for replicating files as they are updated in real-time on the source storage device.
Max Number of Scans Synchronizing Threads	Enter the maximum number of threads available for replicating files during scheduled and on-demand scans of the source storage device.
Max Number of Retries	Enter the maximum number of retries to perform on a file or folder that has failed to be replicated. If the number of retries is exceeded, the file or folder will be added to the Failed Events view and will need to be manually processed.
Retry Interval in seconds	Enter the number of seconds to wait in between retries of the failed replication of a file or folder.

Modifying Scan Manager Settings

The Scan Manager is responsible for handling all scheduled and on-demand scans of the source storage device. Setting are defined globally in **Preferences** in the **Scan Manager** page.

To customize the Scan Manager settings:

1. From the **Window** menu, select **Preferences**.
2. Expand **Cloud Sync** in the navigation tree, and then select **Scan Manager**.



3. Modify the settings as needed.

Scan Item Limit	Enter the maximum number of files and folders to obtain from a folder structure at a time during a scan.
Max Number of Scan Threads	Enter the maximum number of threads available for scanning files and folders. This number should be set to at least the maximum number of jobs running on any single Management Agent.
Max Number of Concurrent Scans	Enter the maximum number of scans that can run in parallel. If the number of active scan threads is greater than this number, scan threads will process on a rotating basis. Increasing this number can increase scan performance but will also increase system memory and CPU utilization.

Creating Sync and Retention Policies

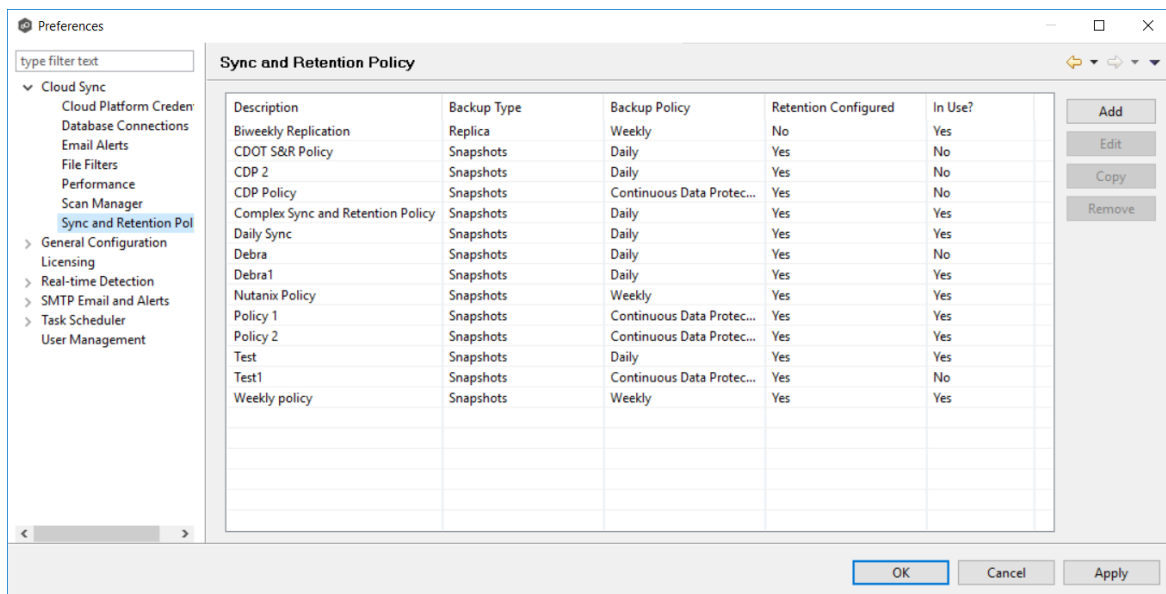
Each Cloud Sync job must have a Sync and Retention Policy applied to it. Policies are defined globally in **Preferences** in the **Sync and Retention Policy** page. You can view, add, edit, and remove policies.

When you create a job, you can select an existing policy to apply to the job or you can create a new policy and apply it to the job.

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

To create a new policy:

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



3. Click the **Add** button.

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

Sync and Retention Policy Wizard

Sync and Retention Policy

Configuration name cannot be empty.

Sync and Retention Policy

Sync Schedule

Retention

*Description:

Job Type:

☒ Snapshots ☐ Replica

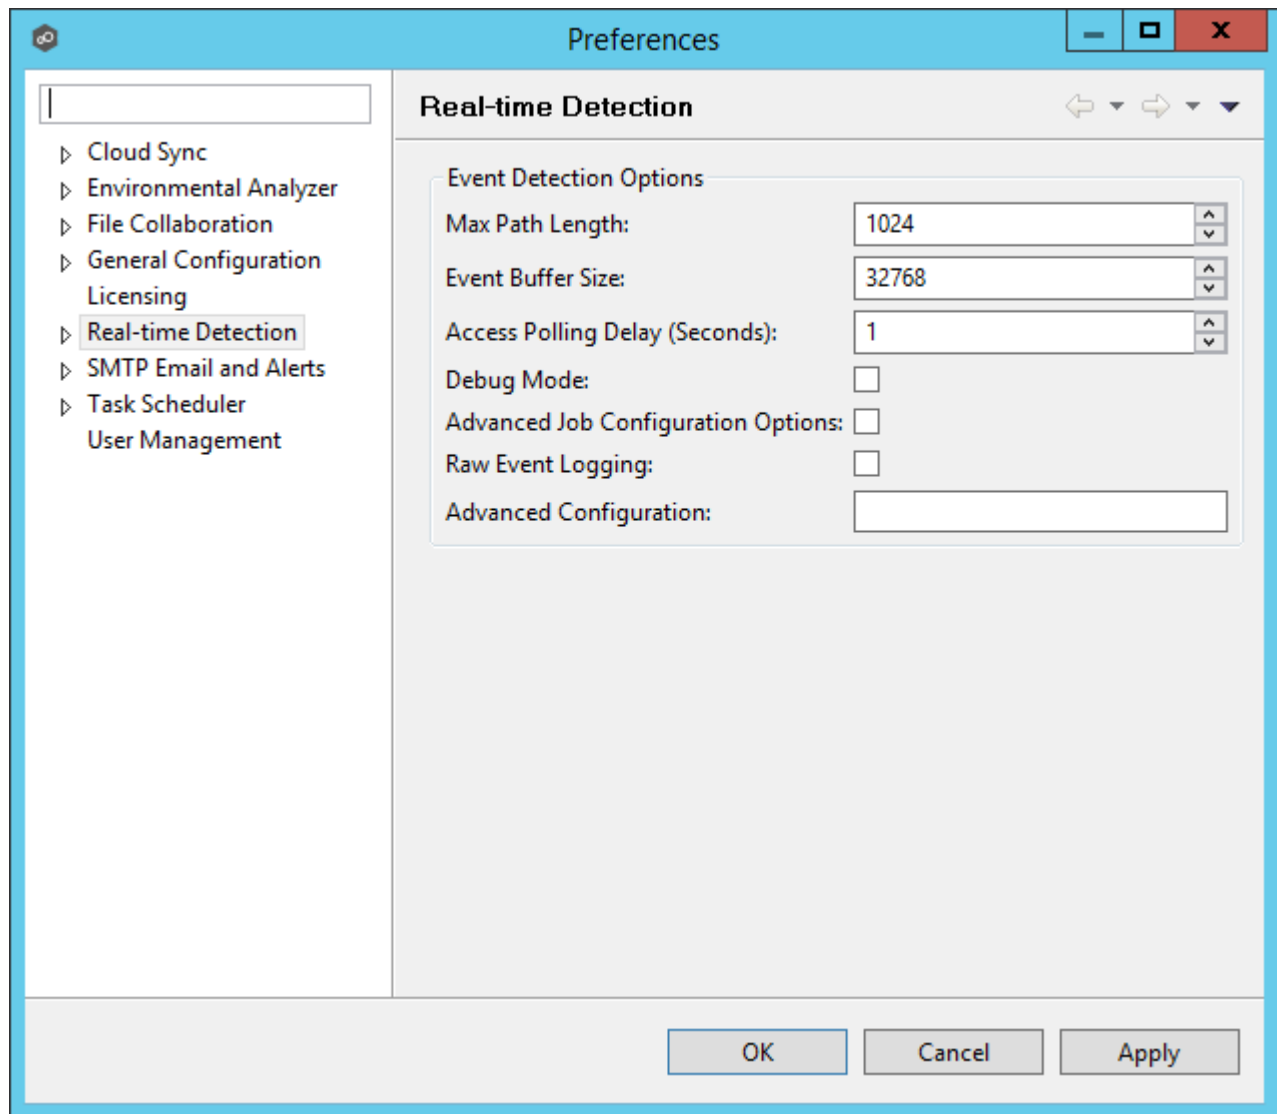
< Back Next > Cancel

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

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

Real-time Detection

A number of options are available to tune the way event detection occurs for all [file collaboration](#). These settings are configured on a global level. To view and modify these settings, click the **Window** menu from within the [Peer Management Center](#), and select **Preferences**. On the left-hand side of the dialog that pops up, select the tree node titled **Real-time Detection**. The following screen will be displayed.



Available options are as follows:

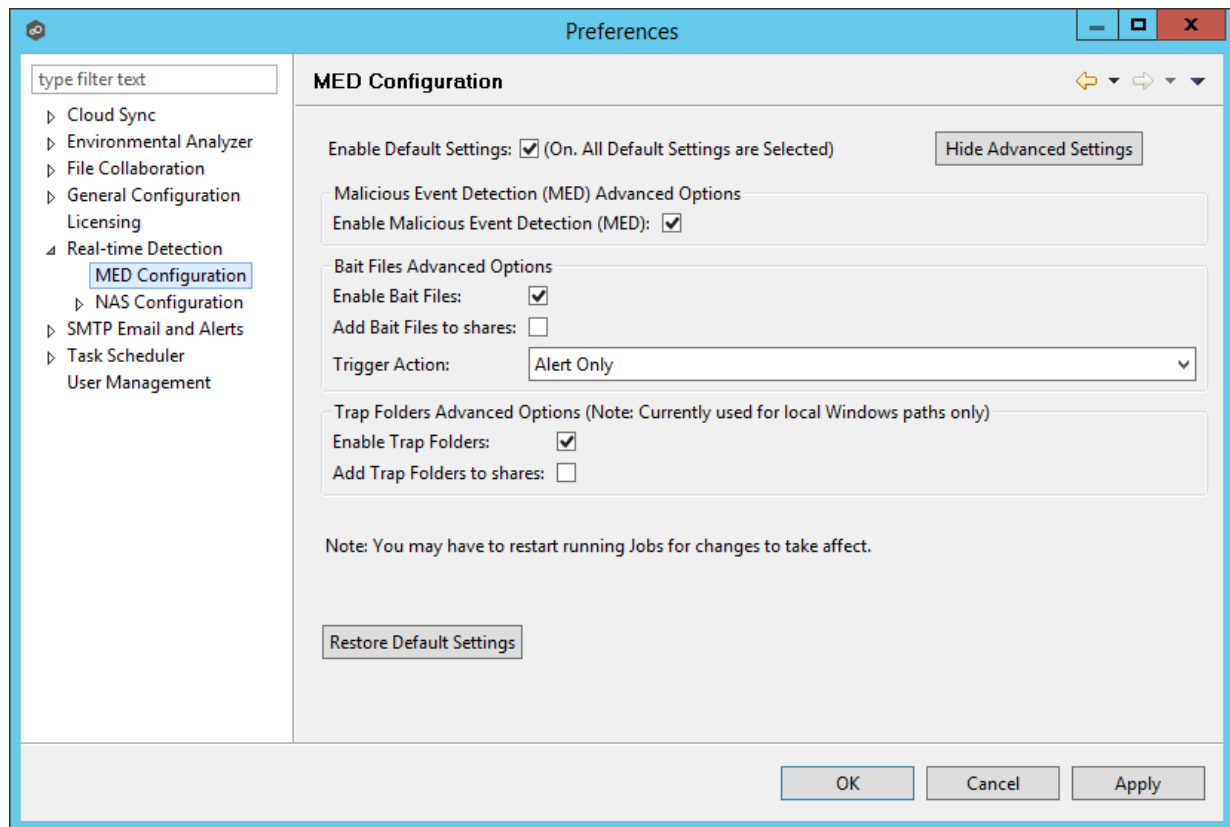
Event Buffer Size	The size in bytes of the buffer used to store real-time events. If you receive Buffer Overflow alerts, then try doubling the size of this buffer to 65536.
Add Context Delay (Seconds)	The number of seconds to wait before scheduling the synchronization of a newly created file.
Change Dispatch Quiet Period (Seconds)	The number of seconds to wait after a file is closed before scheduling the synchronization of the file

Access Polling Delay (Seconds)	The number of seconds between polls of open and closed files.
Debug Mode	Enables advanced debug logging and alerts. Technical support may ask you to enable this feature if you are experiencing certain issues.
Debug Mode w/Full Messages	Enables advanced debug logging and alerts with Full Message information. Technical support may ask you to enable this feature if you are experiencing certain issues.
Advanced Job Configuration Options	Enables configuration for advanced event detection debugging options in job specific host participant detector configuration user interface.
Advanced Configuration	Advanced settings for Event Detection and logging that will override the job settings. Technical support will provide you with a value to put in this field if you are experiencing certain issues.

MED Configuration

With the introduction of Peer Malicious Event Detection (MED), File Collaboration's real-time engine can now be used to spot unwanted activity being executed on storage platforms by ransomware, viruses, malware, hackers, or rogue users. This new technology provides alerting capabilities, as well as the ability to minimize the amount of encrypted or deleted content from being replicated to remote locations. For more information, visit <https://kb.peersoftware.com/tb/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. Once MED is enabled and jobs are restarted, these capabilities will apply to all jobs. To view and modify these settings, click on the **Window** menu from with the [Peer Management Center](#), and select **Preferences**. On the left-hand side of the dialog that pops up, open the tree node titled **Real-time Detection** and select **MED Configuration**. The following screen will be displayed.



The main options are as follows:

Enable Default Settings	Enables/disables Peer MED using default settings. By default, all three MED mechanisms are enabled.
Show/Hide Advanced Settings	Shows/hides options for each of the three MED mechanisms.
Enable Malicious Event 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 Files

Bait files are files of common types, inserted into the file system in a way that hides them from users. Though hidden, these bait files are likely to be accessed by automated processes (like ransomware) or by mass deletions of entire folder structures. As soon as these files are touched, an action is triggered.

The options for Bait Files are:

Enable Bait Files	Enables/Disables bait file creation and monitoring.
Add Bait Files to shares	At the start of each job, create bait files under the root of each participant's configured watch directory. To see the watch directory for a job, review Host Participants and Directories .
Trigger Action	Defines the action to take when MED detects malicious activity on a bait file. See below for more details on available actions.

Trap Folders

On Windows File Servers, Peer MED can be configured to create hidden recursive folders that attempt to trap or slowdown ransomware as it enumerates a folder structure. As with the bait files, these folders cannot be seen by users but will be accessible by automated processes. If bait files (above) are enabled, a bait file will be placed within each trap folder, and an action will be triggered as soon as these files are touched.

Options for trap folders are:

Enable Trap Folders	Enables/disables the creation and monitoring of trap folders.
Add Trap Folders to shares	<p>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.</p> <p>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.</p>

Action Types

For each MED mechanism, one of four actions can be configured on the detection of malicious activity. These actions are:

Alert Only	Triggers an alert in the Peer Management Center. If SMTP email alerts are configured for MED Alerts and enabled for a job, an email will also be sent. For details on SMTP email alerts, see Global Email options . If SNMP traps are configured for MED Alerts and enabled for a job, an SNMP trap will also be sent. For more details on SNMP, see Global SNMP options .
Alert and Disable Host	Triggers an alert while also removing the afflicted Agent from the job in which the malicious activity was detected. Once disabled, Agents will need to be manually re-enabled for collaboration to resume.
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 this job. This option is the most aggressive and will require administrators to re-enable Agents as well as restart jobs.

An example of an alert as displayed in the Peer Management Center is as follows:

Peerlet Advisory Alert Details	
Received Date:	03-12-2018 19:23:26
Severity:	FATAL
Category:	Event Detection
Host Name:	DelIT110a
Locally Created at:	03-12-2018 19:23:26
Message:	<p>Malicious Event Detection (MED) - Bait File Alert (Alert Only: Please check for unwanted activity) Alert Message info=BAIT FILE ALERT appld=113, appSessionId=142 path=See Message Field msg=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\pc-med_bin\Doc_000-med.docx - EventName: RENAME details= Participant Detected=DelIT110a Alert Message=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\pc-med_bin\Doc_000-med.docx - EventName: RENAME Time Detected=Mon Mar 12 19:23:26 EDT 2018 User Detected=MattM IP Detected=Doc_000-med.docx Process Detected=SMBVersion=31 Share Detected=cifs1 Job Session ID=3248744344</p>
Class Name:	WatchDirectoryOperations
App Session Key:	142
Error Code:	2520
Action:	Alert Only

Click outside of popup to close

Re-enabling a Disabled Agent Within a Job

Once disabled within a collaboration 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 bring the afflicted Agent back into collaboration, it will need to be re-enabled on a per job basis. To review the status of an Agent within a job and to re-enable it, navigate to the [Participants View](#) of the job.

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

Summary | Session | Event Log | File Conflicts (0) | Alerts (2) | **Participants (2)** | Configuration

Host Participants

Host	Root Path	Status	State	Message
DellT110a	\\svm9x-1\cifs1\Departments\Sales	Disabled	Disabled	Malicious Event Detection (MED) - Bait File Alert (Alert and Disable Host: Please check for unwanted activity before re-enabling) Alert Message info=BAIT FILE ALERT appId=113, appSessionId=144 path=See Message Field msg=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\pc-med_bin\Doc_001-med.docx - EventName: RENAME details=IParticipant Detected=DellT110aAlert Message=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\pc-med_bin\Doc_001-med.docx - EventName: RENAMETime Detected=Mon Mar 12 19:36:14 EDT 2018User Detected=MattMIP Detected=ActiveCounterValue=IPProcess Detected=SMBVersion=31Share Detected=cifs1Job Session ID=3249149797
DellT3610b	C:\Departments\Sales			

Participant Details

Host Name: DellT110a

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

Status: Disabled

State: Disabled

Monitoring: false

Message: Malicious Event Detection (MED) - Bait File Alert (Alert and Disable Host: Please check for unwanted activity before re-enabling) Alert Message info=BAIT FILE ALERT appId=113, appSessionId=144 path=See Message Field msg=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\pc-med_bin\Doc_001-med.docx - EventName: RENAME details=IParticipant Detected=DellT110aAlert Message=TriggerAlertFileFound: Path=\\svm9x-1\cifs1\Departments\Sales\pc-med_bin\Doc_001-med.docx - EventName: RENAMETime Detected=Mon Mar 12 19:36:14 EDT 2018User Detected=MattMIP Detected=ActiveCounterValue=IPProcess Detected=SMBVersion=31Share Detected=cifs1Job Session ID=3249149797

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

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

Host Participant State Change Log

Filter by: Host:

Date	Host	Status
03-12-2018 ...	DellT110a	Disabled
03-12-2018 ...	DellT3610b	Not Participating
03-12-2018 ...	DellT110a	Disabled
03-12-2018 ...	DellT110a	Not Participating
03-12-2018 ...	DellT3610b	Not Participating
03-12-2018 ...	DellT3610b	Participating

Status: Halted. (Quorum Lost)

Click outside of popup to close

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

NAS Configuration

This section contains detailed information about configuring your NAS:

- [EMC Configuration](#)
- [NetApp Configuration](#)
- [Nutanix Configuration](#)

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more EMC storage devices within a [file collaboration session](#). These EMC devices can be running Isilon, Unity, or VNX.

For more detailed information about prerequisites and configuration, see [EMC Prerequisites and Configuration](#).

EMC Prerequisites and Configuration

Prerequisites

In addition to the standard Peer Management Center [Environmental Requirements](#), the following prerequisites must be met:

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

CEE Server Configuration Guide

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

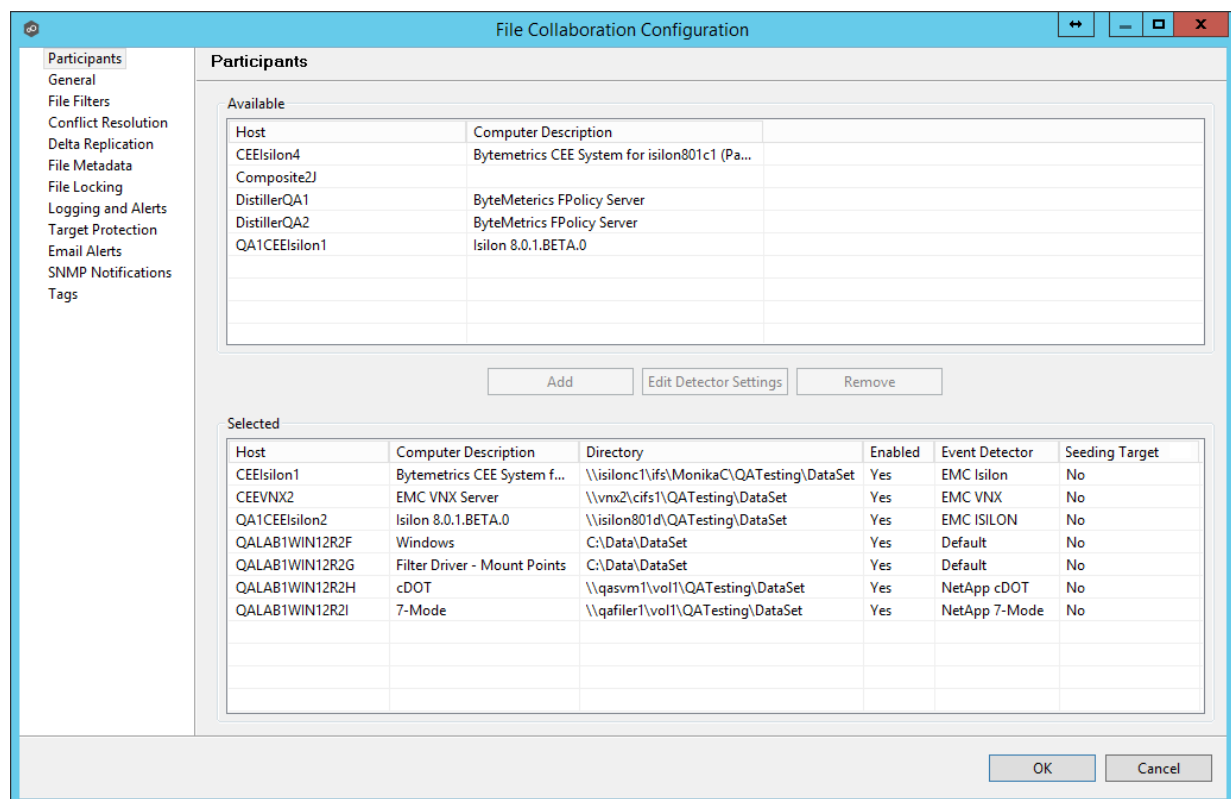
- EMC Isilon CEE Server Configuration Guide: <https://kb.peersoftware.com/tb/emc-isilon-configuration-guide>
- EMC Unity CEE Server Configuration Guide: <https://kb.peersoftware.com/tb/emc-unity-configuration-guide>
- EMC VNX/Celerra CEE Server Configuration Guide: <https://kb.peersoftware.com/tb/emc-vnx-celerra-configuration-guide>

Configuration

1. Review the prerequisites above before beginning the installation and configuration process.
2. Follow the CEE Server Configuration Guide steps.
3. Follow the general [Peer Management Center installation steps](#).
4. Launch the Peer Management Center Client.

Note: Before you can start the Peer Management Center interface/client, the Peer Management Center Service needs to be running. See the [installation](#) section for more information.

5. Install your license within the Peer Management Center. For more information, see the [licensing](#) section. You must contact our sales team to request a license which supports EMC VNX or Isilon. Unless requested, all licenses that are issued do not include EMC support by default.
6. Create a new file collaboration job. For more information, see [Creating a File Collaboration Job](#).
7. During the job configuration process, one or more participating hosts must be configured to interface with EMC. To do so, view the [Participants page](#) of the File Collaboration Configuration dialog, and add the desired available host to the job. After the host is added to the job, enter the UNC path of the appropriate share on the EMC device to the configured directory of the participant that is to act as a CEE Server. Then select **EMC Isilon**, **EMC Unity**, or **EMC VNX** as the participant's configured Event Detector. The example below shows two CEE Servers working with an EMC VNX and an EMC Isilon device.



8. As a result of the selection, a configuration dialog will be displayed requesting additional configuration.
9. With EMC Isilon devices, the following configuration options are available:

EMC Isilon Options

EMC Isilon Options for this Job

Filter open events from these users:

Access Event Suppression Time:

Raw Event Logging: ☐

Advanced Configuration:

Advanced Settings for host: CEEIsilon1 and EMC Isilon:ISILONC1

Filtered IP Addresses:

Nodes:

Cluster IP:

Cluster Port:

Cluster Username:

Cluster Password:

Validate Cluster: ☒

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

OK Cancel

EMC Isilon Options

Filter open events from these users	A comma-separated list of user names to exclude from access event detection . For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default value is -1, which will use a global set value. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.
Raw Event Logging	Enables raw event logging for event detection debugging. Technical support may ask you to enable this feature if you are

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

Advanced Settings for EMC Isilon

Filtered IP Addresses	Events generated from these IP addresses will be filtered. It is recommended that the IP address of the CEE Server is added to this list.
Nodes	Comma-delimited listed of additional node IP address to query for open files. These addresses must be accessible from the CEE Server where the Agent is running.
Cluster IP	The cluster IP address of the Isilon system.
Cluster Port	The cluster port number of the Isilon system. Default value is 8080.
Cluster Username	Username used to sign into the Isilon cluster.
Cluster Password	Password used to sign into the Isilon cluster.
Validate Cluster	If enabled, the Isilon cluster will be validated both on registration and periodically by a maintenance thread.

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

EMC Unity Options

EMC Unity Options for this Job

Filter open events from these users:

Access Event Suppression Time:

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

Filtered IP Addresses:

*Unisphere Management IP:

Unisphere Management Port:

*Unisphere Username:

*Unisphere Password:

Validate Unisphere: ☒

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

OK Cancel

EMC Unity Options

Filter open events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

Advanced Settings for EMC Unity

Filtered IP Addresses	Events generated from these IP addresses will be filtered. It is recommended that the IP address of the CEE Server is added to this list.
------------------------------	---

Unisphere Management IP	The Unisphere Management IP address of the Unity system. This address is used for making API calls to validate configuration.
Unisphere Management Port	The Unisphere Management port number of the Unity system. Default value is 443.
Unisphere Username	The user name used to sign into Unisphere.
Unisphere Password	The password used to connect to Unisphere.
Validate Unisphere	If enabled, Unisphere settings will be validated both on registration and periodically by a maintenance thread.

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

EMC VNX Options

EMC VNX Options for this Job

Filter open events from these users:

Access Event Suppression Time:

Raw Event Logging: ☒

Advanced Configuration:

Advanced Settings for host: CEEVNX2 and EMC VNX:VNX2

Filtered IP Addresses:

Control Station IP:

Control Station Port:

Control Station Username:

Control Station Password:

Validate Control Station: ☒

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

OK Cancel

EMC VNX Options

Filter open events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Access Event Suppression Time	Represents the number of seconds an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default value is -1, which will use a global set value. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.
Raw Event Logging	Enables raw event logging for event detection debugging. Technical support may ask you to enable this feature if you are experiencing certain issues.

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

Advanced Settings for EMC VNX

Filtered IP Addresses	Events generated from these IP addresses will be filtered. It is recommended that the IP address of the CEE Server is added to this list.
Control Station IP	The Control Station IP address of the VNX system.
Control Station Port	The Control Station Port number of the VNX system. Default value is 443.
Control Station Username	Username used to sign into the VNX Control Station.
Control Station Password	Password used to sign into the VNX Control Station.
Validate Control Station	If enabled, the VNX Control Station will be validated both on registration and periodically by a maintenance thread.

10. Once all participating hosts are configured with the appropriate EMC paths and detectors, the file collaboration job may be saved and started.

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more NetApp storage devices within a [file collaboration session](#). These NetApp devices can be running Data ONTAP 7-Mode or clustered Data ONTAP. In order to work with NetApp devices, Peer Management Center leverages the FPolicy API built into the NetApp device. For an in-depth look at how Peer Management Center works with NetApp and the FPolicy system, please email support@peersoftware.com with a request for more information on NetApp support.

For more detailed information about prerequisites and configuration, see [NetAPP Prerequisites and Configuration](#).

NetApp Prerequisites and Configuration

Prerequisites

In addition to the standard Peer Management Center [Environmental Requirements](#), the following prerequisites must be met:

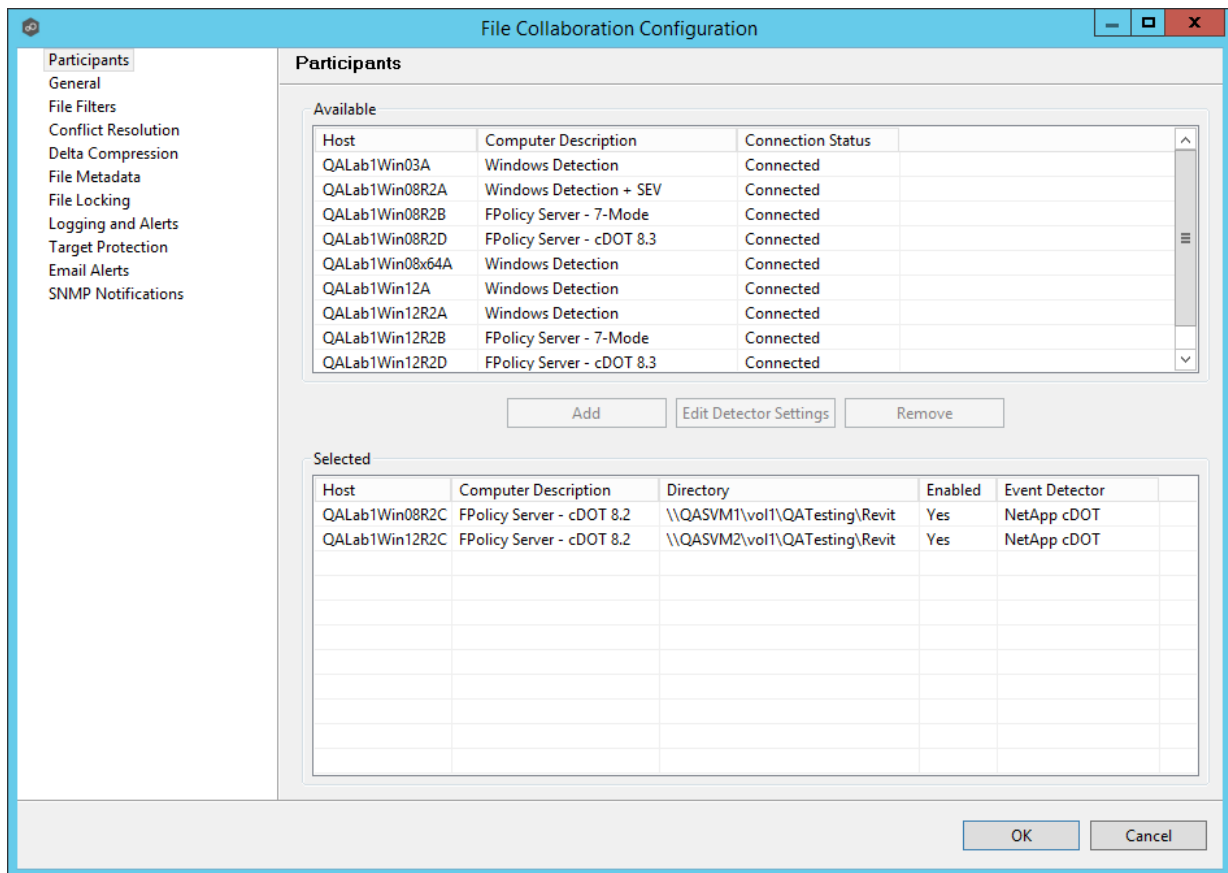
- For NetApp 7-Mode environments: <https://kb.peersoftware.com/tb/netapp-7-mode-prerequisites>
- For NetApp cDOT environments: <https://kb.peersoftware.com/tb/netapp-cdot-prerequisites>

Configuration

1. Review the prerequisites above before beginning the installation and configuration process.
2. Follow the general [Peer Management Center installation steps](#).
3. Launch the Peer Management Center Client.

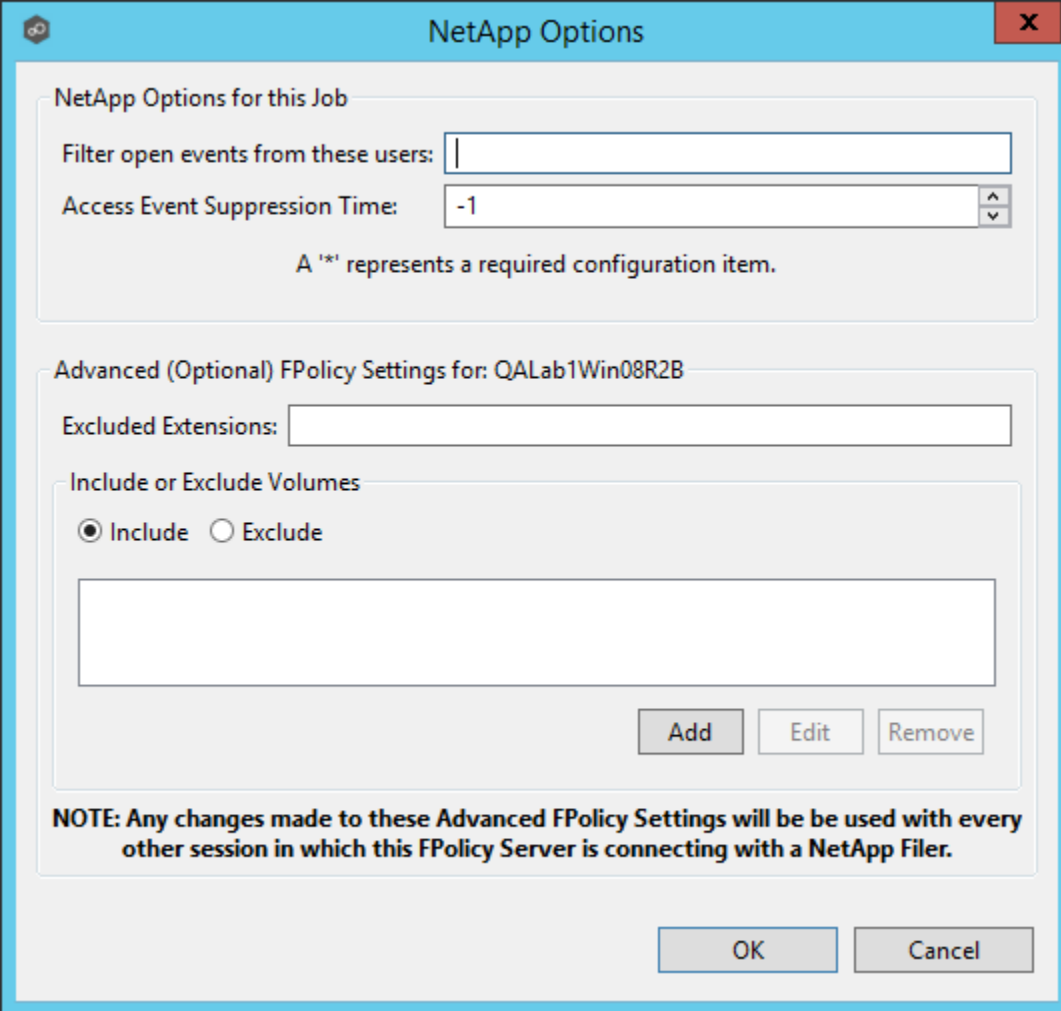
Note: Before you can start the Peer Management Center interface/client, the Peer Management Center Service needs to be running. See the [installation](#) section for more information.

4. Install your license within the Peer Management Center. For more information, see the [licensing](#) section. You must contact our sales team to request a license which supports NetApp. Unless requested, all licenses that are issued do not include NetApp support by default.
5. Create a new file collaboration job. For more information, see [Creating a File Collaboration Job](#).
6. During the job configuration process, one or more [participating hosts](#) must be configured to interface with NetApp. To do so, view the [Participants page](#) of the File Collaboration Configuration dialog, and add the desired available host to the job. After the host is added to the job, enter the UNC path of the appropriate share on the NetApp device to the configured directory of the participant that is to act as a FPolicy Server. Then select **NetApp 7-Mode** or **NetApp cDOT** as the participant's configured Event Detector. This will depend on the edition of Data ONTAP that the NetApp is running. The example below shows an FPolicy Server working with a NetApp 7-Mode device.



7. As a result of the selection, a configuration dialog will be displayed requesting additional configuration for the FPolicy Server.

With NetApp 7-Mode devices, you can just click **OK** and exit the dialog.



NetApp Options

NetApp Options for this Job

Filter open events from these users:

Access Event Suppression Time:

A '*' represents a required configuration item.

Advanced (Optional) FPolicy Settings for: QALab1Win08R2B

Excluded Extensions:

Include or Exclude Volumes

☒ Include ☐ Exclude

Add Edit Remove

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

OK Cancel

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

Filter open events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.
Excluded Extensions	Extensions entered here are excluded from event detection on the NetApp Filer. Values are comma separated and must not

	<p>contain any periods.</p> <p>FPolicy enables you to restrict a policy to a certain list of file extensions by excluding extensions that need to be screened.</p> <p>Note: The maximum length of a file name extension supported for screening is 260 characters. Screening by extensions is based only on the characters after the last period (.) in the file name. For example, for a file named fle1.txt.name.jpg, file access notification takes place only if a file policy is configured for the jpg extension.</p>
Include or Exclude Volumes	<p>List all volumes on the NetApp Filer to exclude or include based on selected choice.</p> <p>FPolicy enables you to restrict a policy to a certain list of volumes by including or excluding volumes that need to be screened.</p> <p>Using the include list, you can request notifications for the specified volume list. Using the exclude list, you can request notifications for all volumes except the specified volume list. However, by default, both the include and exclude list are empty.</p> <p>You can use the question mark (?) or asterisk (*) wildcard characters to specify the volume. The question mark (?) wildcard character stands for a single character. For example, entering vol? in a list of volumes that contain vol1, vol2, vol23, voll4, will result in only vol1 and vol2 being matched.</p> <p>The asterisk (*) wildcard character stands for any number of characters that contain the specified string. Entering *test* in a list of volumes to exclude from file screening excludes all volumes that contain the string such as test_vol and vol_test.</p>

With NetApp cDOT devices, you must fill in the all non-optional settings shown in the table below the screenshot.

NetApp Options
X

NetApp Options for this Job

Filter open events from these users:

Access Event Suppression Time:

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

*SVM Username:

*SVM Password:

SVM Management IP:

*Agent IP for SVM Conn.:

Filtered Extensions:

Admin Share Override:

Disable Share Auto-Detect: ☐

Additional Shares to Include

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

Filter open events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Access Event Suppres	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default vault is -1, which will use a

sion Time	globally set value. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.
SVM Username	The account name of the VSAdmin or similar account on the SVM that has the appropriate access to ONTAPI.
SVM Password	The password of the VSAdmin or similar account on the SVM that has the appropriate access to ONTAPI. This value will be encrypted.
SVM Management IP (optional)	If the primary data LIF for the SVM (whose IP address is registered in DNS) does not support management calls, enter the management IP address of SVM.
Agent IP for SVM Conn.	The IP address over which this Peer Agent will connect to the configured SVM. This MUST be an IP address.
Filtered Extensions	A comma separated list of file extensions to exclude (without a leading asterisk (*)).
Admin Share Override	Enter the administrative-type share that you created on the cDOT SVM. To take advantage of performance improvements when using this option, the share must be created at the root of the SVM's namespace (/). Ideally it should be named to something similar to PMCShare\$ to prevent users from being able to see it.
Disable Share Auto-Detect	Disable the option to auto-detect shares and only use the shares defined in the Participants screen and the Additional Shares to Include option below.
Additional Shares to Include	Specify the shares on the SVM that users can use to access the data that Peer Management Center will be collaborating with. For example, if Peer Management Center is collaborating on data that resides under the Departments share with a local namespace path /departments, but users access data via shares to individual sub folders under the Departments folder (such as Marketing with a local namespace path of /departments/marketing and Sales with a

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

Starting with v3.5.1, Peer Management Center adds a new performance optimization for cDOT environments. To take advantage of this new optimization, **Admin Share Override**, **Disable Share Auto-Detect**, and **Additional Shares to Include** all need to be configured. For more details and before making any changes, please contact Peer Software Support.

Once all participating hosts are configured with the appropriate NetApp paths and detectors, the file collaboration job may be saved and started.

Peer Management Center supports the ability to include content from CIFS/SMB shares on one or more Nutanix Acropolis File Services (AFS) clusters within a [file collaboration session](#).

For more detailed information about prerequisites and configuration, see [Nutanix Prerequisites and Configuration](#).

Nutanix Prerequisites and Configuration

Prerequisites:

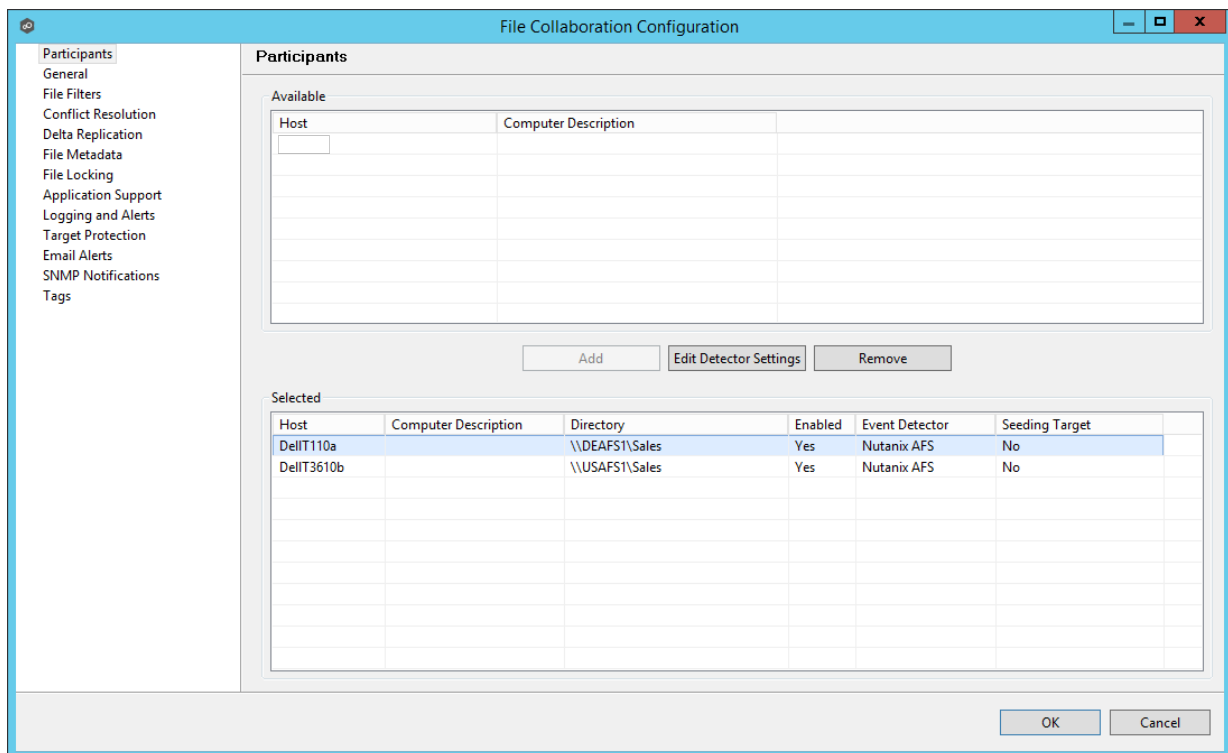
For Nutanix AFS environments, the following up to date prerequisites must be met in addition to the standard Peer Management Center [Environmental Requirements](#):
<https://kb.peersoftware.com/tb/nutanix-afs-prerequisites>

Configuration

1. Review the prerequisites above before beginning the installation and configuration process.
2. Follow the general [Peer Management Center installation steps](#).
3. Launch the Peer Management Center Client.

Note: Before you can start the Peer Management Center interface/client, the Peer Management Service needs to be running. See the [installation](#) section for more information.

4. Install your license within the Peer Management Center. For more information, see the [licensing](#) section. You must contact our sales team to request a license that supports Nutanix AFS.
5. Create a new file collaboration job. For more information, see [Creating a File Collaboration Job](#).
6. During the job configuration process, one or more [participating hosts](#) must be configured to interface with AFS. To do so, view the [Participants page](#) of the File Collaboration Configuration dialog, and add the desired available host to the job. After the host is added to the job, enter the UNC path of the appropriate share on the AFS cluster for the configured directory of the participant that is to act as a AFS Partner Server. Then select Nutanix AFS as the participant's configured Event Detector. The example below shows two Agent Partner Servers working with two different AFS clusters.



7. As a result of the selection, a configuration dialog will be displayed requesting additional configuration for the AFS cluster.
8. With Nutanix AFS clusters, the following configuration options are available:

AFS Options

AFS Options for this Job

Filter open events from these users:

Access Event Suppression Time:

Advanced Settings for host: DellT3610b and AFS Cluster: USAFS1

Partner Server IP:

User Name:

Password:

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

OK Cancel

Nutanix AFS Options

Filter open events from these users	A comma-separated list of user names to exclude from access event detection. For example, if "USER1" is excluded, any access event activity generated by USER1 will be ignored, e.g., file is opened and closed.
Access Event Suppression Time	Represents the number of seconds that an open event will be delayed before being processed. Used to help reduce the amount of chatter generated by Windows 7 clients when mousing over files in Windows Explorer. The default value is -1, which will use a globally set value. A value of 0 will allow for dynamic changes to the amount of time that an open event will be delayed based on the load of the system.

Advanced Settings for Nutanix AFS

Partner Server IP	The IP address over which the configured AFS cluster will send activity to this Peer Agent. This MUST be an IP address.
User Name	User name used to access the APIs on the configured AFS cluster.

Password	Password used to access the APIs on the configured AFS cluster.
-----------------	---

9. Once all participating hosts are configured with the appropriate UNC paths and detectors, the file collaboration job may be saved and started.

File Collaboration

This section provides information about creating, editing, running, and managing a File Collaboration job:

- [Creating a File Collaboration Job](#)
- [Editing a File Collaboration Job](#)
- [Running and Managing a File Collaboration Job](#)
- [Runtime Job Views](#)

Creating a File Collaboration Job

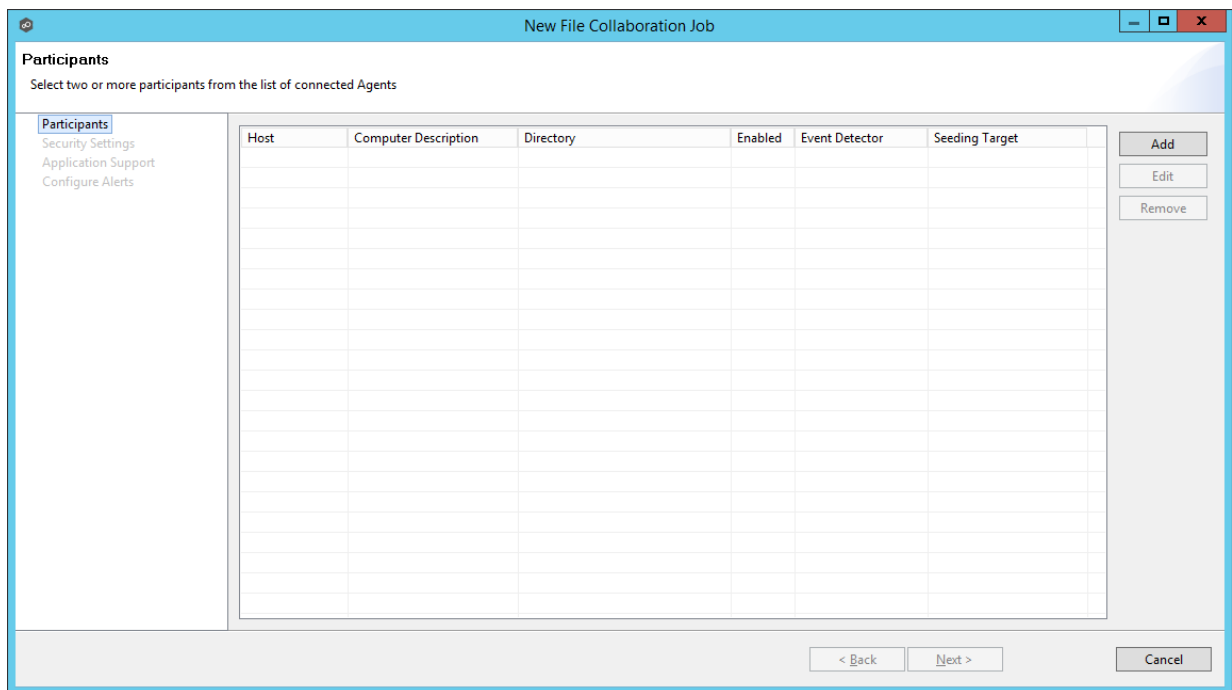
The topics in this section provide some basic information about creating File Collaboration jobs. Starting in Peer Management Center (PMC v4.2), the creation of File Collaboration jobs is accomplished by a new wizard. The following section reviews the steps required in this new wizard.

Step 1: Participants

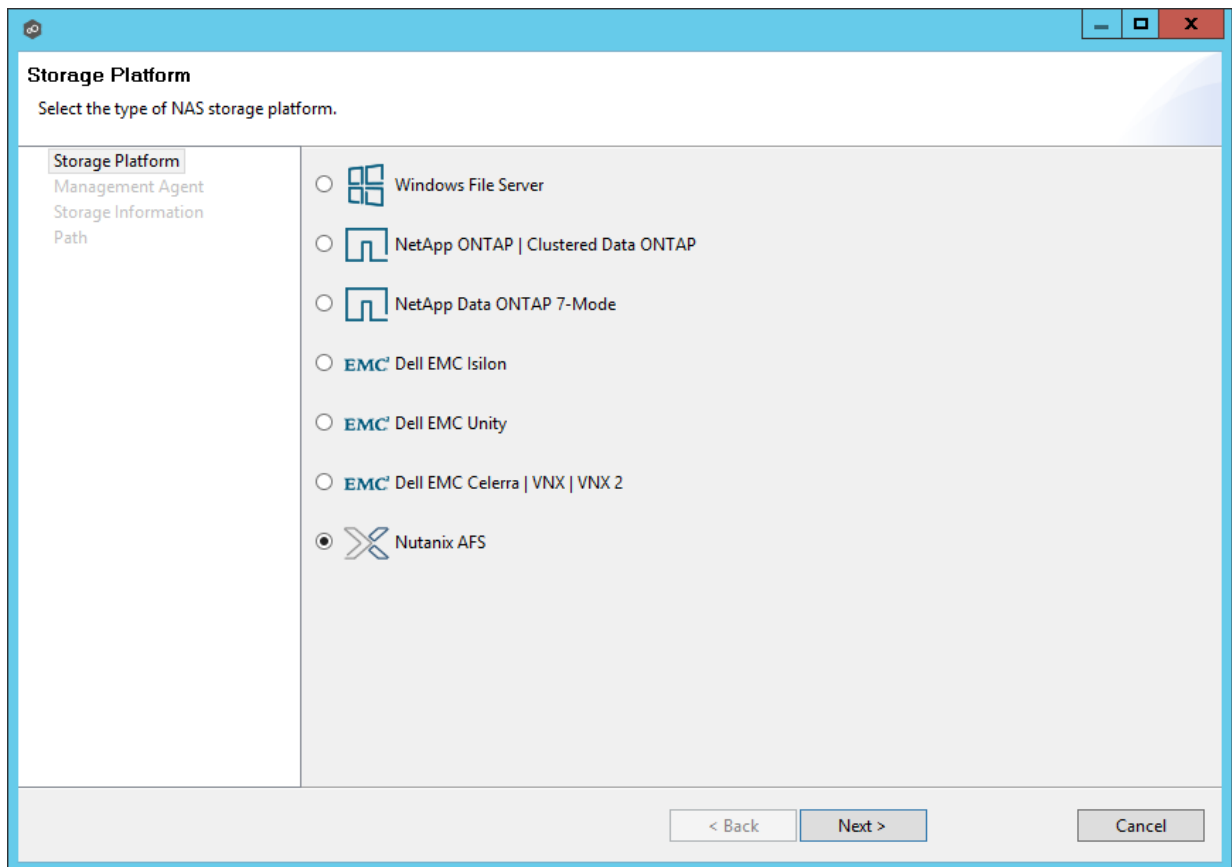
Once [preferences](#) have been configured, create a new file collaboration job by clicking on the Create New button in toolbar of the Peer Management Center, or by selecting New from the File menu. A list of all installed Peerlet types will be displayed. Selecting the File Collaboration option will prompt you for a unique name for the job, then open the New File Collaboration wizard.

The first step in the new File Collaboration job wizard is to select two or more Agents to [participate](#) in this replication relationship.

1. To add a participant, click the **Add** button.

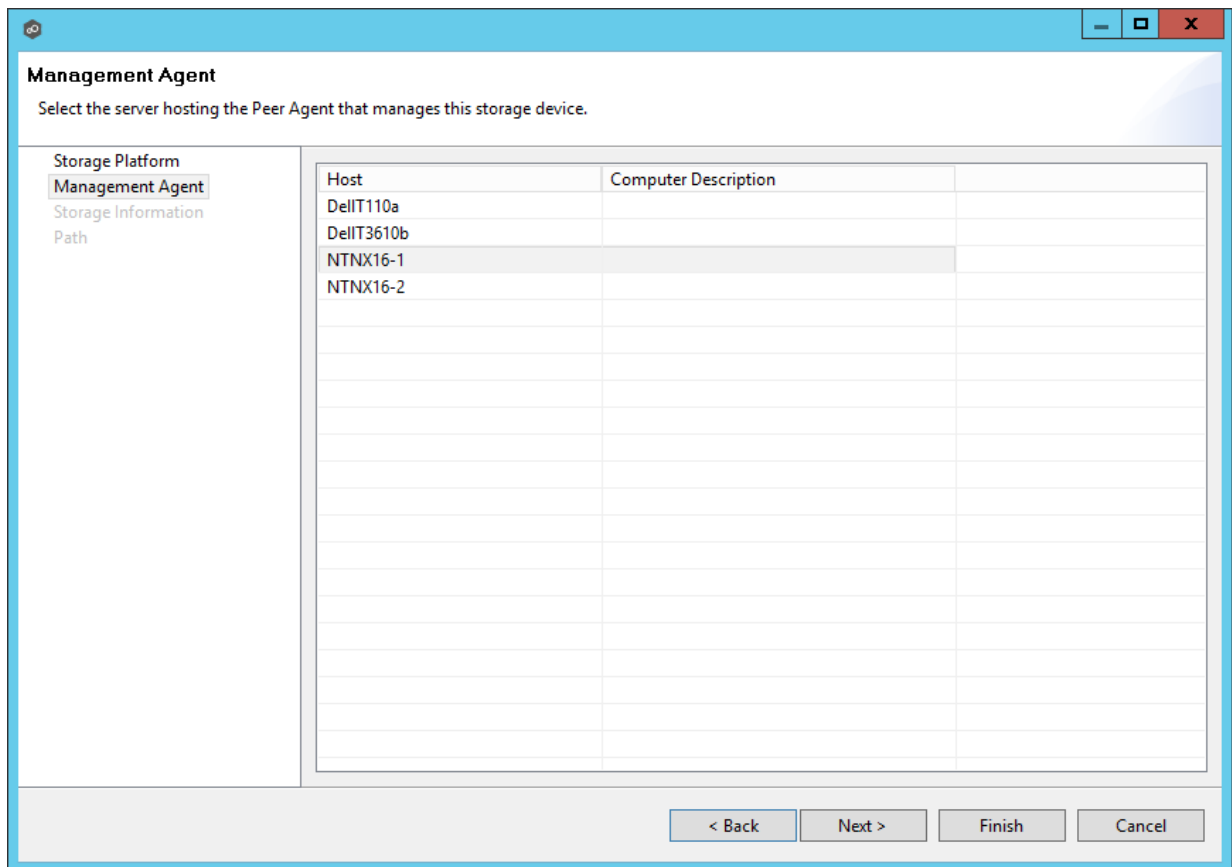


2. Another wizard will appear to guide you through the process of adding an Agent (and an optional NAS device) to this new file collaboration job. The storage platform that you select from this screen will determine what settings are needed later in the wizard.



The screenshot shows a window titled "Storage Platform" with a subtitle "Select the type of NAS storage platform." On the left is a sidebar with a tree view containing "Storage Platform", "Management Agent", "Storage Information", and "Path". The main area contains a list of storage platforms, each with a radio button and an icon: "Windows File Server" (Windows logo), "NetApp ONTAP | Clustered Data ONTAP" (NetApp logo), "NetApp Data ONTAP 7-Mode" (NetApp logo), "EMC Dell EMC Isilon" (EMC logo), "EMC Dell EMC Unity" (EMC logo), "EMC Dell EMC Celerra | VNX | VNX 2" (EMC logo), and "Nutanix AFS" (Nutanix logo). The "Nutanix AFS" option is selected. At the bottom are three buttons: "< Back", "Next >", and "Cancel".

3. Select the appropriate platform, then click **Next**.
The next screen will show a list of connected Agents.



4. Select the one that will directly manage the storage platform selected above, then click **Next**

The required settings on the following screen will vary depending on the storage platform selected above. If Windows was selected, the wizard will skip this page.

5. If the selected Management Agent has never been connected to a NAS device, you will need to enter connection information specific to the selected storage platform. See the [File Collaboration](#) subsection of the **Advanced Configuration** section for more information on the requirements of each storage platform.

If the selected Management Agent has been connected to a NAS device in the past, you can select **Existing Credentials** to automatically use the previously entered settings.

Storage Information
Enter the information required to connect to the storage.

Storage Platform Management Agent
Storage Information
Path

Credentials

☒ New Credentials

*AFS Cluster Name: MMAFS1

*Username: admin

*Password: ••••••••

*Peer Agent IP: 192.168.162.101

Advanced

☐ Existing Credentials

Validate

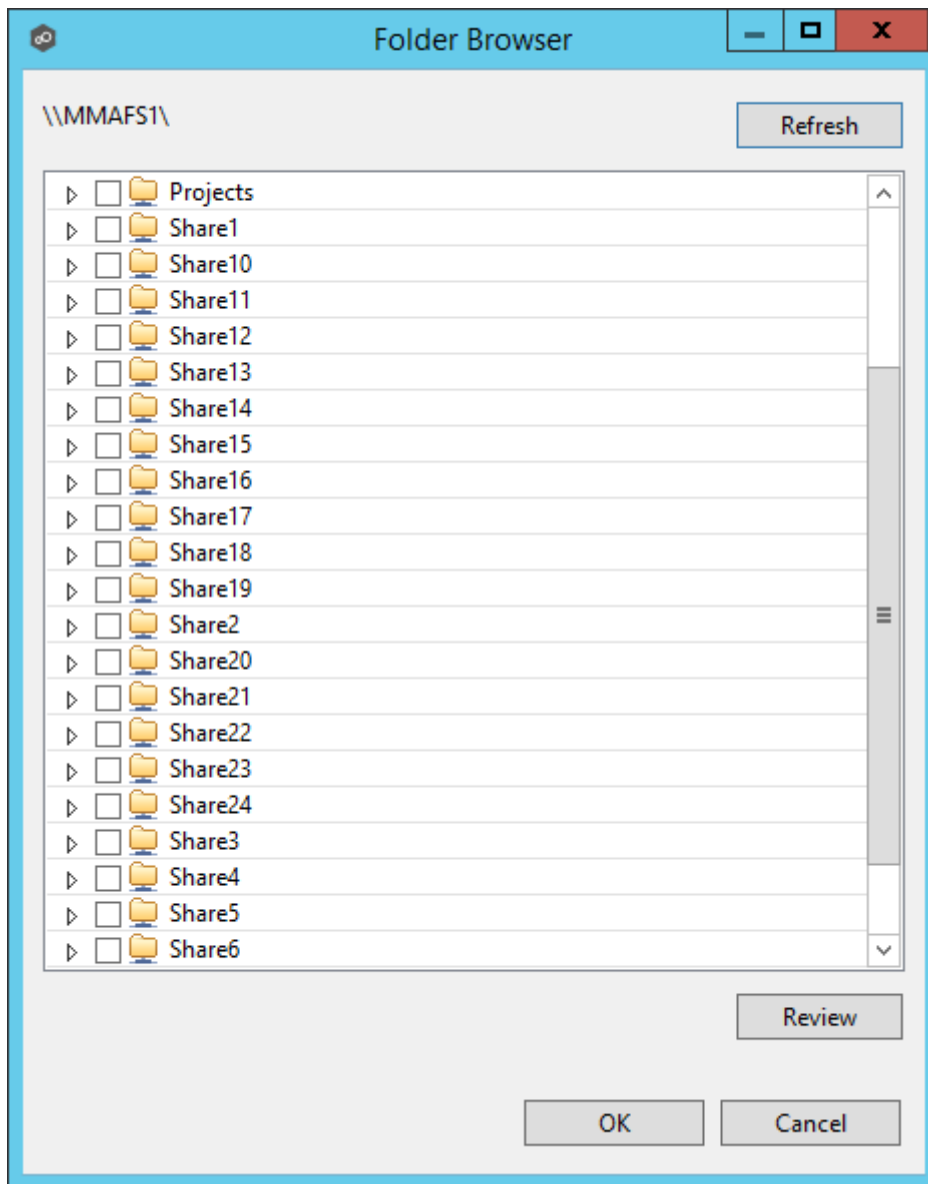
Having trouble connecting? Please verify that all [prerequisites](#) are met for Nutanix AFS environments.

< Back Next > Cancel

6. Click **Next** to continue.
7. Once storage device information has been entered and verified, you will be prompted to browse for or enter a path for the [watch set](#).

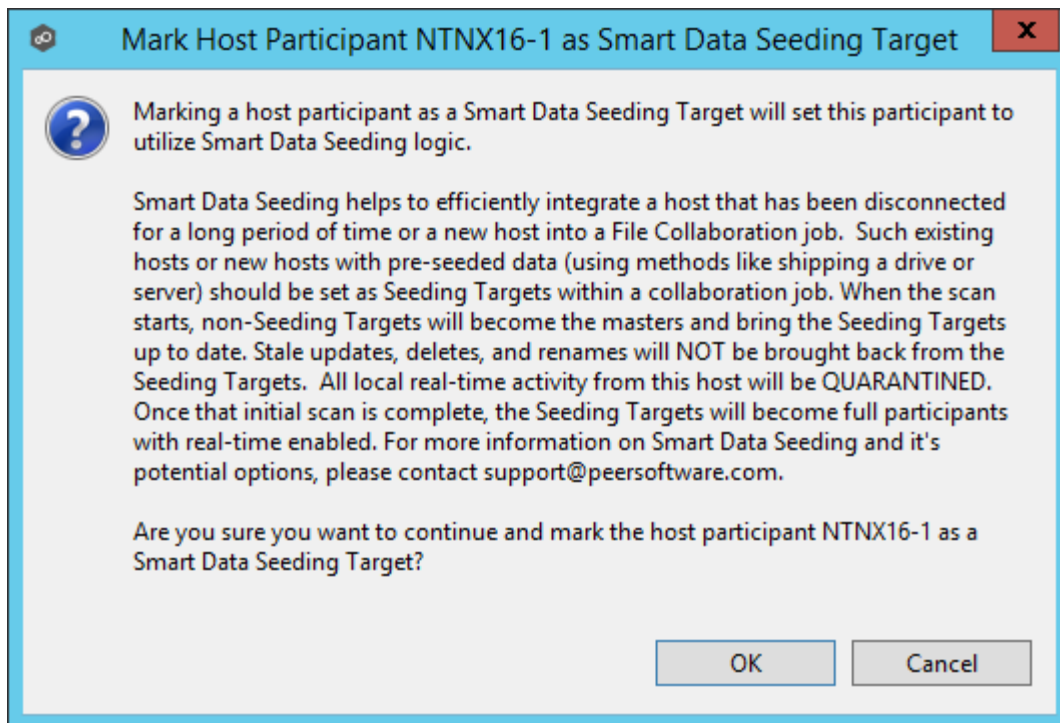
The screenshot shows a window titled "Path" with a subtitle "Browse to or enter a path on the storage device." On the left is a sidebar with a tree view containing "Storage Platform", "Management Agent", "Storage Information", and "Path" (which is selected). The main area contains a text input field with the placeholder text "\\MMAFST\ Enter Path" and a "Browse" button to its right. Below the input field is a "Seeding Target:" label followed by an unchecked checkbox. At the bottom of the window are four buttons: "< Back", "Next >", "Finish", and "Cancel".

If you wish to browse for the desired path, click the **Browse** button. The following dialog will appear.



Check the appropriate volume/share/subfolder, then click **OK**.

Optionally, if the **Seeding Option** is checked, the following warning like the following will appear:



Click **OK** if the described behavior is what you are looking for. For more information, please contact support@peersoftware.com.

8. Click **Finish** to complete the wizard for this participant.
9. Repeat the preceding six steps for each participant.

Step 2: Security Settings

The next step in the new File Collaboration job wizard is to configure NTFS security options.

Security Settings
Configure the replication of NTFS security permissions.

Participants
Security Settings
Application Support
Email Alerts

Synchronize Security Descriptors (ACLs)
☐ Enable synchronizing NTFS security descriptors (ACLs) in real-time
☐ Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan

Synchronize Security Descriptor Options
☒ DACL: Discretionary Access Control List
☐ Owner
☐ SACL: System Access Control List

File Metadata Conflict Resolution
File metadata conflict resolution will only occur the first time a file is synchronized during the initial scan, and only when one or more security descriptors and/or file attributes do not match the designated master host. If the file does not exist on the designated master host, then no conflict resolution will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan. Select a host below to use as the master for resolving file metadata conflicts.
Master Host:

< Back Next > Finish Cancel

For more information on these settings, see the [File Metadata](#) section.

Step 3: Application Support

The new File Collaboration job wizard includes the ability to automatically optimize a job for specific applications.

If an application does not show up on this page, this does not mean that the application is not supported. This list represents the file types that we have known best practices for, including filtering recommendations, the prioritization of certain file types, and the enabling or disabling of file locking.

Check off all applications that will interact with the files and folders under the watch of this job. These optimizations will automatically be applied across all configured watch sets of all participants in this job. You can modify this list later.

New File Collaboration Job

Application Support

Select the applications that will be used with the data that is managed by this job.

Participants
Security Settings
Application Support
Email Alerts

Select below to optimize this job for any of the following file types:

Adobe Products

- ☐ Adobe Illustrator
- ☐ Adobe InDesign
- ☐ Adobe Photoshop

AutoDesk Products

- ☐ AutoDesk Revit
- ☐ AutoDesk AutoCAD
- ☐ AutoDesk Inventor
- ☐ AutoDesk Civil 3D
- ☐ AutoDesk Civil 3D [Sheet Set manager]

Other

- ☐ Dassault Systems CATIA
- ☐ Microsoft Office
- ☐ Microsystems Allegro
- ☐ Newforma Project Center
- ☐ Rhinoceros Rhino3D

< Back Next > Finish Cancel

For exact details as to what each application optimization is, please contact support@peersoftware.com.

prompt you for a unique name for the job, then open the File Collaboration Configuration dialog.

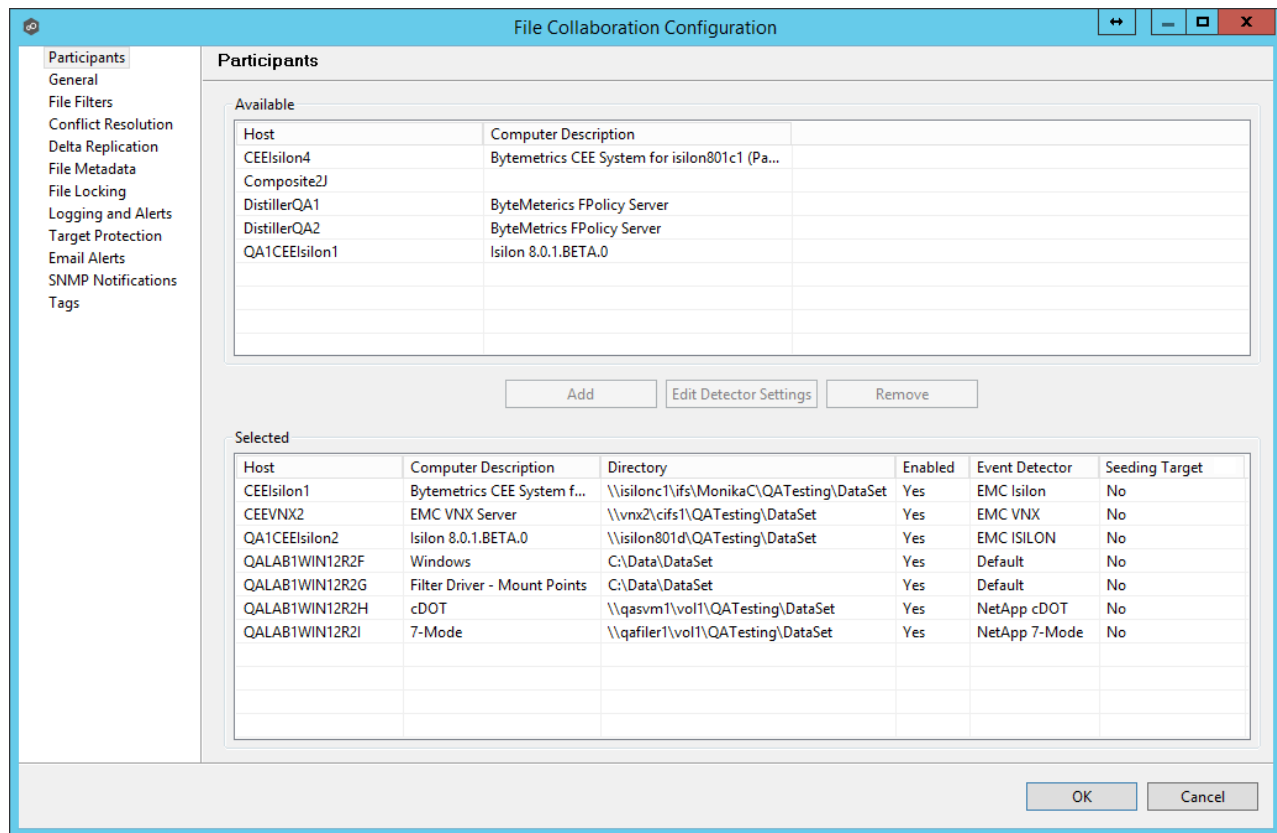
You can edit an existing job by selecting one or more jobs in the Jobs view, right-clicking, and selecting **Edit Configuration(s)**. The Peer Management Center now has support for editing multiple jobs at once. See the [Editing Multiple Jobs](#) section for more details.

Configuring a file collaboration session requires the following steps:

1. [Preferences](#) (important to configure before setting up your first job)
2. [Host Participants & Folders Settings](#) (the beginning process of creating an individual file collaboration job)
3. [General Settings](#)
4. [File Filters Settings](#)
5. [File Conflict Resolver Settings](#)
6. [Delta Replication](#)
7. [File Metadata](#)
8. [File Locking](#)
9. [Logging and Alerts](#)
10. [Target Protection](#)
11. [Email Alerts](#)
12. [Tags](#)
13. [Save Settings](#)

Participants and Directories

The first page of configuration will be for [Host Participants](#) of the file collaboration job. On this page, you will select and configure which hosts will be participating in this job.



Participant configuration steps are as follows:

1. A list of all available hosts will appear in the Available table on the top of the page. Available hosts are any host with a [Peer Agent](#) installed that has successfully connected to the configured [Peer Management Center Broker](#). The name that will be displayed is the computer name of the server that the Peer Agent is running on. If a particular host is not displayed in the list then try restarting the Peer Agent Windows Service on that host, and if it successfully connects to the Peer Management Center Broker, then the list will be updated with the computer name of that host.

Note: Computer Description is defined through Windows on a per-computer basis.

2. Select two or more hosts from the **Available** table and click on the **Add** button to add the hosts to the **Selected** table.
3. For each selected host you will need to type in the path to the root folder, and then press enter. The root folder for all hosts can be identical, or they can have different absolute path names based on your needs.
4. Optionally, if you would like to exclude real time events from certain users, this can be done by selecting the desired host in the **Selected** table and clicking **Edit Detector Settings**. This is helpful if you are trying to prevent events generated from backup and/or archival tools from triggering activity. User names should be separated by commas.

5. If you are properly licensed to and wish to include an enterprise NAS device within a file collaboration job, additional configuration is required for each storage platform. Supported platforms include [EMC Isilon](#), [EMC VNX](#), [NetApp 7-Mode](#), [NetApp cDOT](#), and [Nutanix AFS](#).

Note: From this point on, no other configuration items are mandatory. You can leave the rest of the configuration settings as their default values and move onto to [Save Settings](#). If you wish to continue configuring the jobs, continue to [General Settings](#).

Seeding Target

Smart data seeding helps to efficiently integrate a host that has been disconnected for a long period of time or a new host into a File Collaboration job. Such existing hosts or new hosts with pre-seeded data (using methods like shipping a drive or server) should be set as **Seeding Targets** within a collaboration job. When the scan starts, non-Seeding Targets will become the masters and bring the **Seeding Targets** up to date. Stale updates, deletes, and renames will NOT be brought back from the **Seeding Targets**. All local real-time activity will be quarantined. Once that initial scan is complete, the **Seeding Targets** will become full participants with real-time enabled. For more information on Smart Data Seeding and it's potential options, please contact support@peersoftware.com.

General Settings

The General Settings page contains miscellaneous configuration items pertaining to a file collaboration and is available by selecting **General** from the tree node within the File Collaboration Configuration dialog.

File Collaboration Configuration

General

Job ID: 113

Job Name: Sales Dept Collaboration

Transfer Block Size (KB): 256

File Synchronization Job Priority: 2

Timeout (Seconds): 180

Remove Filtered Files On Folder Delete: ☐

Require All Hosts At Start: ☐

Auto Start: ☐

OK Cancel

Configurable settings for this page are as follows:

Job ID	Unique, system-generated job identifier that cannot be edited.
Job Name	Description of this file collaboration job. This name should be unique.
Transfer Block Size (KB)	The block size in Kilobytes used to transfer files to hosts. Larger sizes will yield faster transfers on fast networks, but will consume more memory in the Peer Management Center Broker and Agents .
File Synchronization Job Priority	Use this to increase or decrease a jobs 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.
Remove Filtered	If enabled, then all child files on target hosts will be deleted when it's parent folder is deleted on another source host. Otherwise,

Files On Folder Delete	filtered files will be left intact on targets when a parent folder is deleted on another source host.
Require All Hosts At Start	This option 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 checked, then this file collaboration session will automatically be started when the Peer Management Center Service is started.

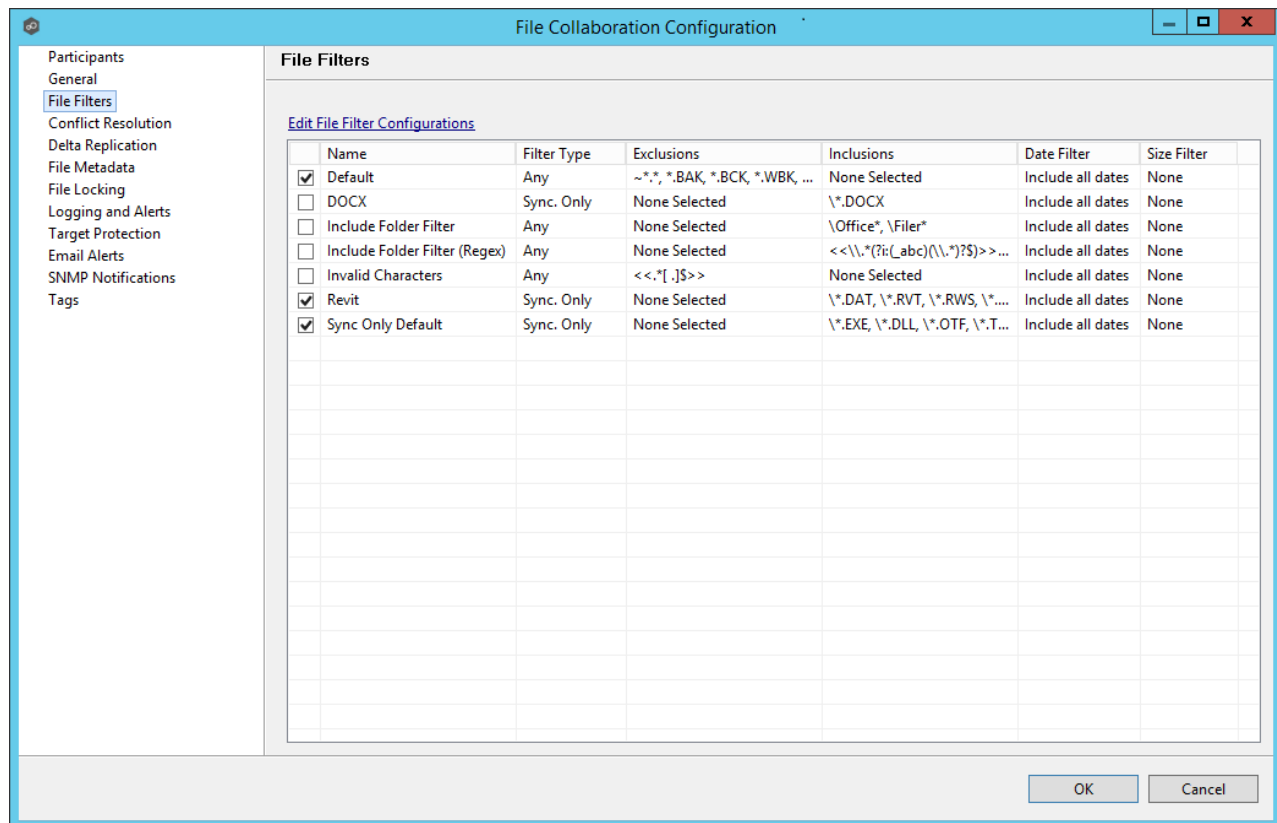
Once all settings are configured to your liking, you can either [save the configuration](#) and exit the dialog, or you can continue with the configuration process by going to [File Filters](#).

File Filters

File filters allows you to specify file and folder path expressions to include and/or exclude from a file collaboration and is available by selecting **File Filters** from the tree node within the File Collaboration dialog.

File filters are configured on a global basis within the [Peer Management Center](#), where individual filters can be applied to multiple jobs without having to manually re-enter each part of the configuration. For more information on what exactly a file filter is, see [File Filters](#). For details on how to configure a file filter within the Peer Management Center, see [Configuring File Filters](#).

The following screenshot shows how individual file filters can be applied to a single job.



Each global file filter will be displayed in the table on this page. If you need to create a new file filter, or edit an existing filter via the [Configuring File Filters](#) screen, click the **Edit File Filter Configurations** link. Once all necessary filters are in place, check all that you would like to apply to the current job. Each checked item will be combined into one large filter when the job is run (by combining all exclusions and inclusions together). In general, you should have at least one default global file filter that is applied to all jobs and possibly other file filters that apply to specific jobs. However, for most environments, only a single default global file filter is necessary.

Once all file filters are set and selected to your liking, you can either [save the configuration](#) and exit the dialog, or you can continue with the configuration process by going to [File Conflict Resolution](#).

Folder Filter Examples

- To exclude a specific folder from anywhere within the Peer Management Center watch set:

*\FolderName

*\FolderName\FolderName

- To exclude a specific folder from the ROOT of the Peer Management Center watch set:

\FolderName

\FolderName\FolderName

- To exclude folders that END with a specific name from anywhere within the Peer Management Center watch set:

*FolderName\

- To include a specific folder from the ROOT of the Peer Management Center watch set:

\FolderName

\FolderName\FolderName

File Conflict Resolution

File conflict resolution allows you to specify the type of file conflict resolution to use during the initial scan when a file conflict exists for a file between two or more hosts. Configuration is available by selecting **Conflict Resolution** from the tree node within the File Collaboration Configuration dialog.

Overview

Conflict resolution is a key feature of file collaboration that is in effect at the start of a session. When a file collaboration job begins, the [host participants'](#) configured folders are synchronized by a scan and merge phase, during which conflicts can be detected. Below, we will define file conflicts, describe our detection scheme, and the configuration options we provide to resolve them.

Defining a Conflict

When a session begins, the participants' folders are first scanned then merged to form a collective view of all participants' content. All files found under the designated folders are subject to collaboration, except for those excluded by filtration (see [Configuring File Filters](#) for more details).

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

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

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

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

In this example, the file '\release-1.0\readme.txt' is found to be in conflict across three hosts. Note that each host can designate varying root folders. Content below the root folder resides under a shared namespace. Conflicts may occur across a partial or total set of participant hosts.

A file conflict can occur for any of the following reasons:

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

Resolving a Conflict

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

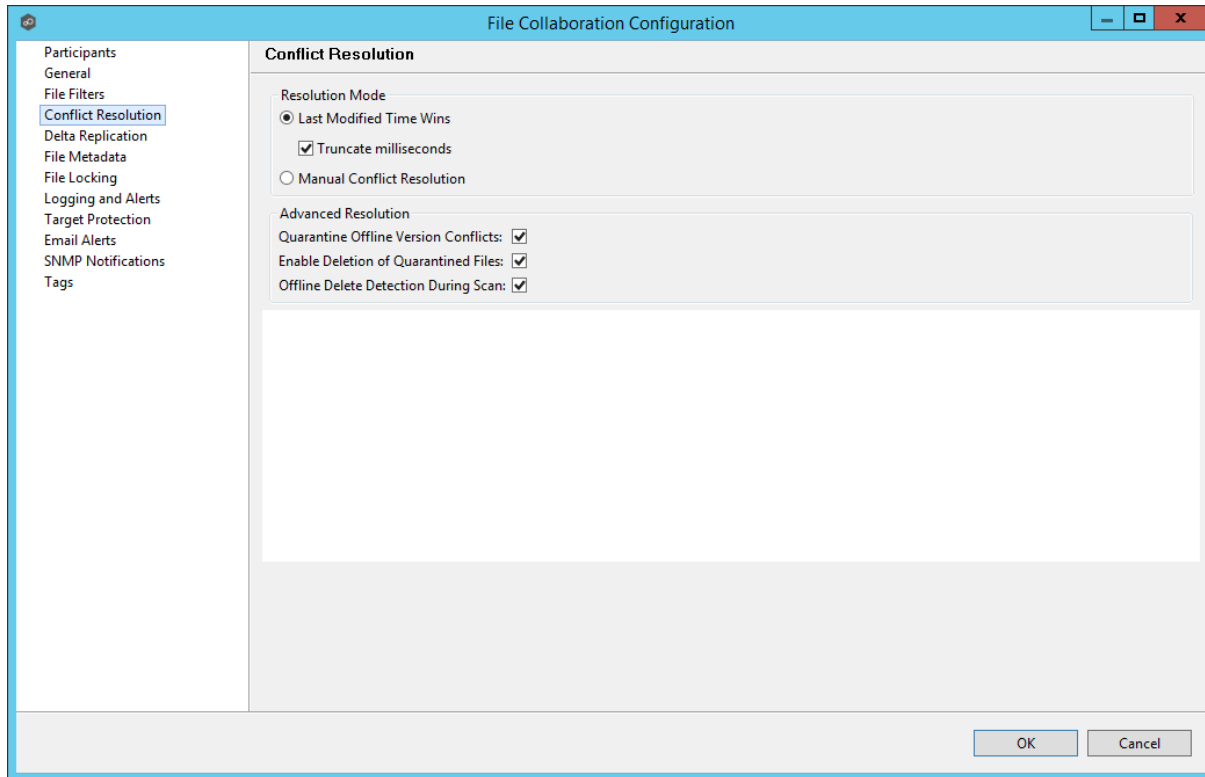
It is important to note that conflict resolution must select a single instance of a file, although it is quite possible that several copies of a file are potential candidates. Drawing from the examples listed in the previous section, if our session was configured to resolve conflicts based on a file's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

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

See the [File Conflict View](#) for a more detailed explanation on how the file conflict process works and how to remove file conflicts and quarantines.

Configuration

The following is a view of the Conflict Resolution configuration page.



The conflict resolution types that are currently available are listed as follows:

Last Modified Time Wins	<p>A file's modification time will be used to designate an instance as a resolution candidate. The later the modification time, the greater the likelihood for a file's selection.</p> <p>Options:</p> <p>Truncate milliseconds: When comparing the time stamps of a file on two or more hosts, truncate the millisecond value from each time stamp.</p>
None (Manual Resolution)	<p>When selected, any file conflicts that are encountered during the initial synchronization process will result in quarantines that are added to the File Conflict List. These file conflicts must be resolved manually by selecting the host with the correct version of the file from the conflict list.</p>

All the types listed above have the potential for producing multiple resolution candidates. A collaboration session can be configured with any one of the available conflict resolvers. If a resolver produces more than one candidate for a conflicted file, a winner will be selected arbitrarily.

Advanced conflict resolution options are listed as follows:

Quarantine Offline Version Conflicts	Enable this option if you want Peer Management Center to quarantine a file that was updated in two or more locations while the collaboration session was not running.
Enable Deletion of Quarantined Files	If a file that is quarantined is deleted, Peer Management Center will process the delete event and remove the quarantine when this option is enabled.
Offline Delete Detection During Scan	If this option is enabled and target protection is enabled, and it can be determined that a file or folder has been deleted since the session was stopped, then the file or folder will be deleted from all hosts. If this option is not enabled then the deleted file or folder will be brought back to any host where it was removed.

Once all File Conflict Resolvers are selected and set to your liking, you can either [save the configuration](#) and exit the dialog, or you can continue with the configuration process by going to [Delta Replication](#).

Delta Replication

Delta-level replication configuration is available by selecting Delta Replication from the tree node within the File Collaboration Configuration dialog.

Overview

Delta-level replication is a byte replication technology that enables block/byte level synchronization for a file collaboration job. Through the use of this feature, Peer Management Center will be 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.

Configuration

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.

The screenshot shows the 'File Collaboration Configuration' dialog box with the 'Delta Compression' tab selected. The left sidebar lists various configuration categories, with 'Delta Replication' highlighted. The main area contains the following settings:

- Enable Block/Byte Synchronization:** A checked checkbox.
- Checksum Transfer Size (KB):** A text box with the value '512'.
- Delta Block Transfer Size (KB):** A text box with the value '512'.
- Minimum File Size (KB):** A text box with the value '100'.
- Minimum File Size Percentage Target/Source:** A text box with the value '0.30'.
- Excluded File Extensions:** A list box containing the following extensions: zip, jpg, jpeg, png, gif, tiff, tif, Z, tgz, gz, gzip, rar, 7z, bz, bz2, and bzip2. Below the list are 'Add' and 'Remove' buttons.
- Excluded File Name Wildcard Patterns:** An empty text box. Below it are 'Add' and 'Remove' buttons.

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

Below is a list of configuration items and their descriptions:

Enable Block/Byte Synchronization	Enables 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.
Disable on Session Startup	Disables delta-level replication during file collaboration session startup where the state of all hosts and files is not known. If enabled, delta encoding would need to be performed between source and each target separately since the state of any files is not known.
Checksum Transfer Size (KB)	The block size 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)	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)	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	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	List of comma-separated wildcard patterns of file extensions to be excluded from delta encoding, e.g., zip, jpg, png. In general, compressed files should be excluded from delta encoding and the most popular compressed file formats are excluded by default.
Excluded File Name Wildcard Patterns	A list of file name wildcard patterns to exclude from delta encoding. If a file name matches any wildcard pattern in this list then it will be excluded from delta encoding transfers and a regular file transfer will be performed. See Configuring File Filters for more information on specifying wildcard expressions.

Once all Delta-level Replication settings are set, you can either [save the configuration](#) and exit the dialog, or you can continue with the configuration process by going to [File Metadata](#).

File Metadata

File metadata configuration is available by selecting **File Metadata** from the tree node within the File Collaboration Configuration dialog.

Overview

File metadata is additional information stored as part of the file. The main component of file metadata is Security Descriptor Information, comprised of attributes such as DACLs, SACLs, Owner, Group, ACLs.

By default, enabling real time file metadata synchronization will cause any real-time modifications of metadata to be synchronized with all other target hosts. This alone, however, will not enable synchronizing file metadata during the [initial synchronization process](#). In order to enable file metadata synchronization during the initial synchronization process, you must enable this option and select a MASTER host to use as the conflict winner.

ACL Guidelines and Best Practices

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

More detailed information about [ACL Guidelines](#) can be found at the URL below:

<http://www.peersoftware.com/support/knowledgebase/item/peerlink-acl-guidelines.html>

File Metadata Conflict Resolution

File Metadata Conflict Resolution will only occur the first time a file is synchronized during the initial scan, and only when one or more security descriptors do not match the designated master host. If the file does not exist on the designated master host, then no conflict resolution will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan.

Configuration

The following screen presents available File Metadata configuration options:

The screenshot shows the 'File Collaboration Configuration' dialog box with the 'File Metadata' tab selected. The left sidebar lists various configuration categories, with 'File Metadata' highlighted. The main area contains several sections: 'Synchronize Security Descriptors (ACLs)' with three checkboxes (the third is checked), 'Synchronize Security Descriptor Options' with three checkboxes (the first is checked), 'File Metadata Conflict Resolution' with a text description and a 'Master Host' dropdown, 'File Reparse Point Synchronization' with two input fields, and 'Alternate Data Streams Transfer' with one unchecked checkbox. 'OK' and 'Cancel' buttons are at the bottom right.

File Collaboration Configuration

File Metadata

Synchronize Security Descriptors (ACLs)

- ☐ Enable synchronizing NTFS security descriptors (ACLs) in real-time
- ☐ Enable synchronizing NTFS security descriptors (ACLs) with master host during initial scan
- ☒ Enable prevention of corrupt or blank Owner or DACLS on source or master host from being applied to any target host

Synchronize Security Descriptor Options

- ☒ DACL: Discretionary Access Control List
- ☐ Owner
- ☐ SACL: System Access Control List

File Metadata Conflict Resolution

File metadata conflict resolution will only occur the first time a file is synchronized during the initial scan, and only when one or more security descriptors and/or file attributes do not match the designated master host. If the file does not exist on the designated master host, then no conflict resolution will be performed. If a master host is not selected, then no file metadata synchronization will be performed during the initial scan. Select a host below to use as the master for resolving file metadata conflicts.

Master Host:

File Reparse Point Synchronization

Reparse Tag Name (numerical value only):

Reparse Master Host:

Alternate Data Streams Transfer

- ☐ Enable transfer of file Alternate Data Streams (ADS)

OK Cancel

Below is a list of file metadata options along with their descriptions:

Enable synchronizing NTFS security descriptors (ACLs) in real-time	If enabled, changes to the configured security descriptor component (e.g., DACL, SACL, Owner, etc.) will be transferred to the target host file(s) as they occur.
Enable synchronizing NTFS security descriptors (ACLs) during initial scan	If enabled, changes to the configured security descriptor component (e.g., DACL, SACL, Owner, etc.) will be synchronized during the initial scan (if a Master Host is selected).

Enable prevention of corrupt or blank Owner or DACLS	If enabled, then corrupt or blank Owner or DACLS on source or master host will not be applied on any target host file.
Synchronize Security Descriptor Options	You can select which security descriptor components are synchronized. Choices are DACL, SACL and Owner. In general, you will usually only need to synchronize DACLS. If you need to synchronize SACLs or Owner, then the user that a Peer Agent service is run under on each participating host must have permission to read and write SACLs and Owner.
Master Host	The master host to use for conflict resolution during the initial synchronization process.

File Reparse Point Data Synchronization

This option should only be used if you are utilizing archiving or hierarchical storage solutions that make use of NTFS file reparse points to access data in a remote location, such as, Symantec's Enterprise Vault. Enabling this option will allow synchronizing a file's reparse data, and not the actual offline content, to target hosts, and will prevent the offline file from being recalled from the remote storage device.

Reparse Tag Name	A single numerical value. Must be either empty (reparse synchronization will be disabled), or greater than/equal to 0. The default for Symantec Enterprise Vault is '16'. A value of 0 will enable reparse point synchronization for all reparse file types. If you are unsure as to what value to use, then either contact our technical support, or you can use a value of 0 if you are sure that you are only utilizing one vendor's reparse point functionality.
Reparse Master Host	If a master host is selected then when the last modified times and file sizes match on all hosts, but the file reparse attribute differs (e.g. archived/offline verse unarchived on file server), then the file reparse data will be synchronized to match the file located on the master host. For Enterprise Vault, this should be the server where you run the archiving task on. If the value is left blank, then no reparse data synchronization will be performed, and the files will be left in their current state.

Alternate Data Stream Synchronization

Enable transfer of	If enabled, Alternate Data Streams (ADS) of updated files will be transferred to the corresponding files on target participants as a
---------------------------	--

Alternate Data Streams (ADS)

post process of the normal file synchronization.

Known Limitation: ADS information is only transferred when a modification on the actual file itself is detected. ADS will not be compared between participants. The updated file's ADS will be applied to the corresponding files on target participants.

Once all File Metadata settings are set, you can either [save the configuration](#) and exit the dialog, or you can continue with the configuration process by going to [Logging & Alerts](#).

File Locking

The File Locking settings page contains miscellaneous configuration items pertaining to how source and target files are locked by Peer Management Center, and is available by selecting **File Locking** from the tree node within the File Collaboration Configuration dialog.

The screenshot shows the 'File Collaboration Configuration' dialog box with the 'File Locking' tab selected. The left sidebar contains a tree view with the following items: Participants, General, File Filters, Conflict Resolution, Delta Replication, File Metadata, File Locking (highlighted), Logging and Alerts, Target Protection, Email Alerts, SNMP Notifications, and Tags. The main area of the dialog is titled 'File Locking' and contains the following settings:

- Locking Options**
 - Exclusive Target Lock: ☐
 - Include MS Office User Lock Information: ☒
- Source Snapshot Synchronization**
 - Enable Source Snapshot Sync.: ☐
 - Snapshot File Extensions:
 - Max File Size (MB):
- Sync. On Save**
 - Enable Sync. On Save: ☐
 - Included File Extensions:
 - Synchronization Delay (Seconds):

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

Below are a list of general fields and their descriptions:

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).
Enable Source Snapshot Sync.	If enabled, a snapshot copy of the source file will be created for files that meet the snapshot configuration criteria below, and this copy will be used for synchronization purposes. In addition, no file handle will be held on the source file except while making a copy of the file.
Snapshot File Extensions	A comma separated list of file extensions for which source snapshot synchronization will be utilized.
Max File Size (MB)	The maximum file size 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.
Sync. On Save: Included File Extensions	A comma separated list of file extensions for which to enable the Sync. On Save feature.
Sync. On Save: Synchronization Delay	The number of seconds to wait after a file has been saved before initiating a synchronization of the file.

Application Support

Each job can be automatically optimized to work with specific applications.

If an application does not show up on this page, this does not mean that we do not support the application. This list represents the file types that we have known best practices for, including filtering recommendations, the prioritization of certain file types, and the enabling or disabling of file locking.

Check off all applications that will interact with the files and folders under the watch of this job. These optimizations will automatically be applied across all configured watch sets of all participants in this job.

The screenshot shows a window titled "File Collaboration Configuration" with a sidebar on the left and a main content area. The sidebar lists various configuration options: Participants, General, File Filters, Conflict Resolution, Delta Replication, File Metadata, File Locking, Application Support (which is highlighted with a blue box), Logging and Alerts, Target Protection, Email Alerts, SNMP Notifications, and Tags. The main content area is titled "Application Support" and contains the instruction "Select below to optimize this job for any of the following file types:". Below this instruction are three sections of checkboxes: "Adobe Products" with options for Adobe Illustrator, Adobe Photoshop, and Adobe InDesign; "AutoDesk Products" with options for AutoDesk Revit, AutoDesk AutoCAD, AutoDesk Inventor, AutoDesk Civil 3D, and AutoDesk Civil 3D [Sheet Set manager]; and "Other" with options for Dassault Systems CATIA, Microsoft Office, Microsystems Allegro, Newforma Project Center, and Rhinoceros Rhino3D. At the bottom right of the window are "OK" and "Cancel" buttons.

For details as to what each application optimization is, please contact support@peersoftware.com.

Logging and Alerts

File Event Logging

Various types of file collaboration events can be written to a log file and to the [Event Log](#) tab located within the [File Collaboration Runtime View](#) for the selected file collaboration job. Each

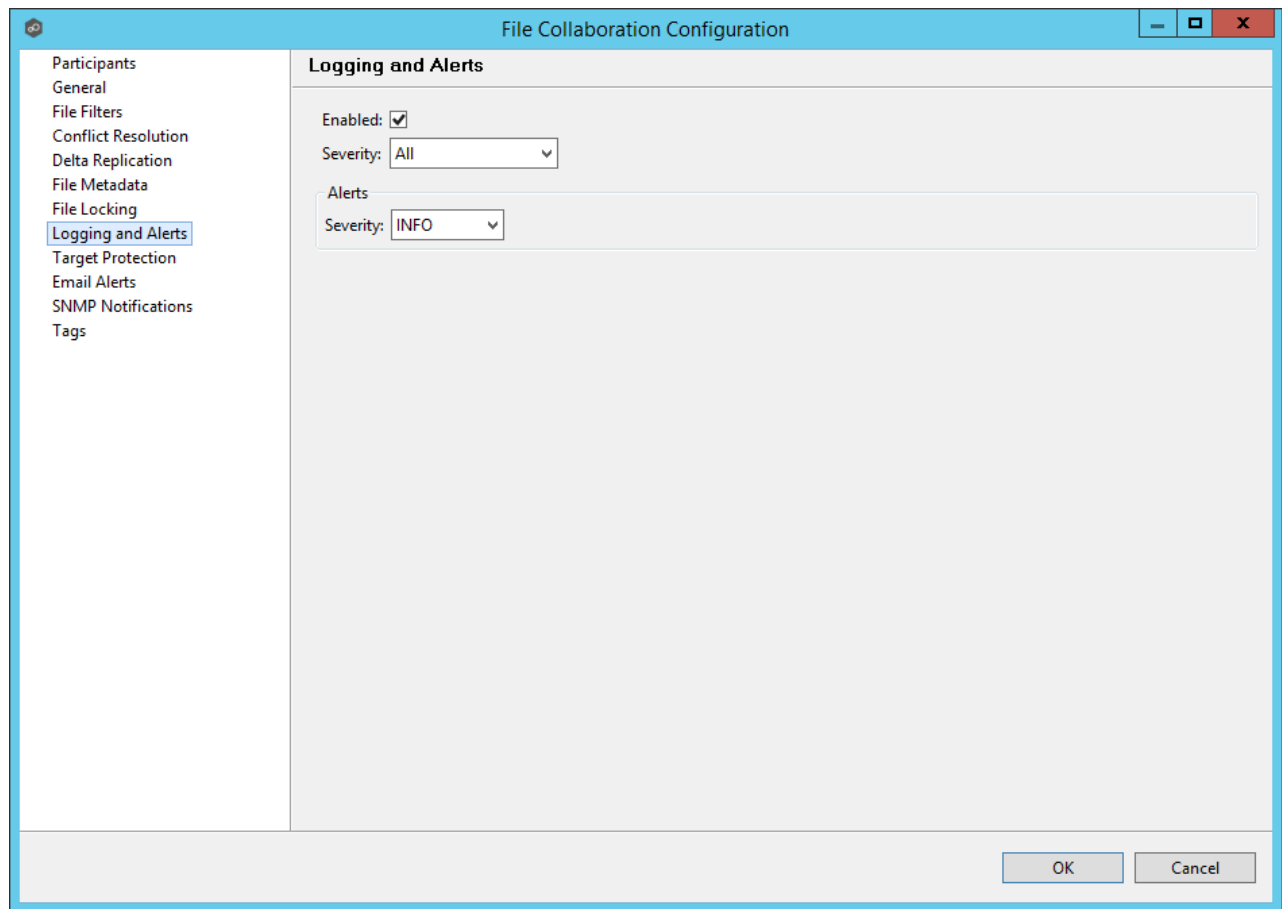
job will log to the **fc_event.log** file located in the 'Hub\logs' subdirectory within the Peer Management Center installation directory. All log files are stored in a tab delimited format that can easily be read by Microsoft Excel or other database applications.

Log Entry Severity Levels

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

Configuration

By default, all file collaboration activity is logged for all severity levels. You can enable or disable file event logging as well as select the level of granularity on what to log through the **Logging and Alerts** page, available by selecting **Logging and Alerts** from the tree node within the File Collaboration Configuration dialog.



Below is a list of logging fields and their descriptions:

Enabled	Checking this option will enable file event logging based on the other settings. Unchecking this option will completely disable all logging.
Severity	Determines what severity levels will be logged. There are two options: <ul style="list-style-type: none">• All (Informational, Warnings, Error, Fatal)• Errors and Warnings (Warnings, Error, Fatal)
Event Types	If checked, the corresponding event type will be logged.
File Open	A file was opened by a remote application on a Source Host .

File Lock	A file lock was acquired on a Target Host by the file collaboration job.
File Close	A file was closed.
File Add	A file was added to the watch set .
File Modify	A file was modified in the watch set.
File Delete	A file was deleted.
File Rename	A file was renamed.
Attribute Change	A file attribute was changed.
Security (ACL) Change	The security descriptor of a file or folder was changed.
Directory Scan	Indicates when a directory was scanned as a result of the initial synchronization process .
File ADS Transfer	The Alternate Data Stream of a modified file was synced to target host(s).

Alerts

Configured in the screen shown above, various types of alerts will be logged to a log file and to the [Alerts](#) table located within the [File Collaboration Runtime View](#) for the selected job. Each file collaboration job will log to the **fc_alert.log** file located in the 'Hub\logs' subdirectory within the Peer Management Center installation directory. All log files are stored in a tab delimited format that can easily be read by Microsoft Excel or other database applications.

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

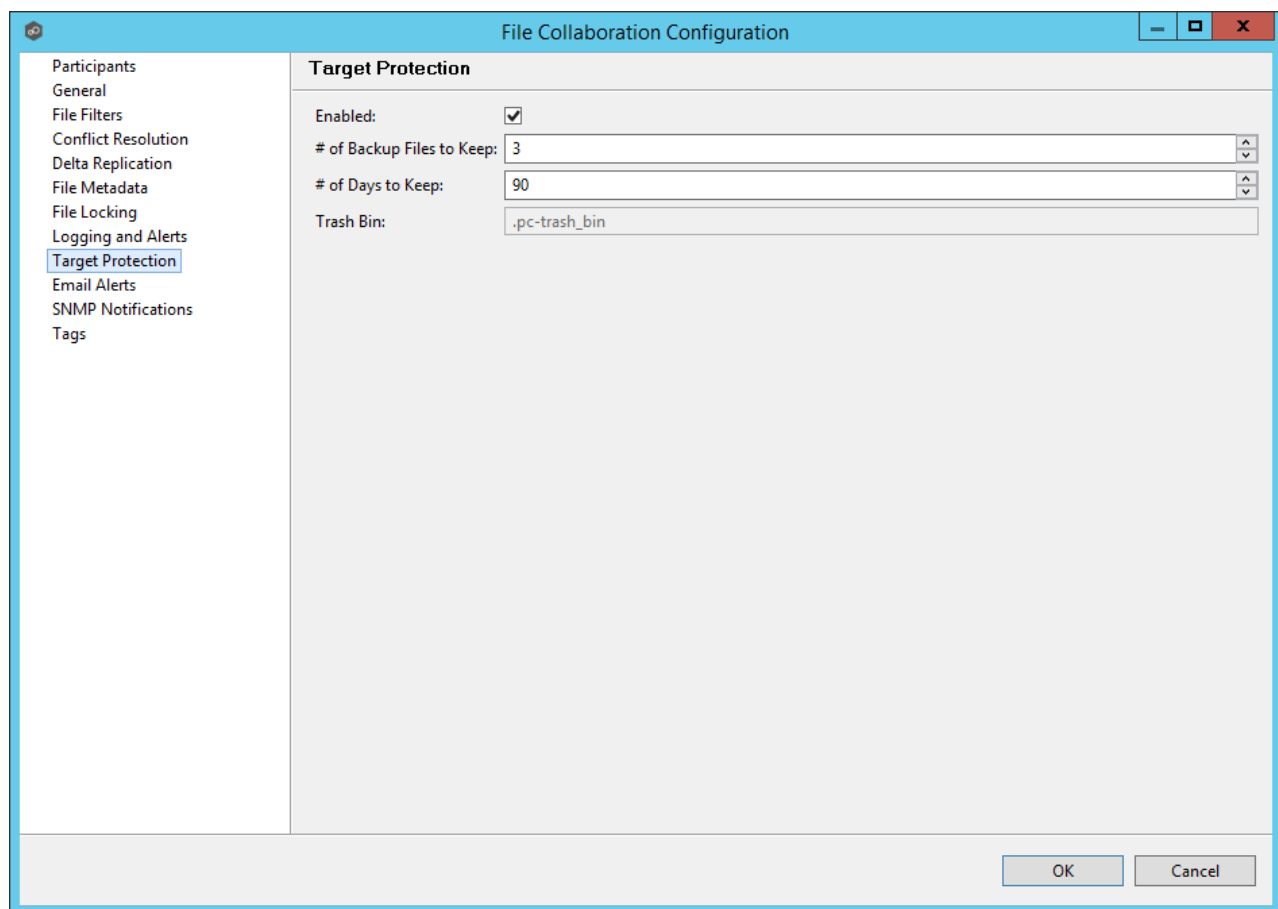
Once all Logging and Alerts settings are set, you can either [save the configuration](#) and exit the dialog, or you can continue with the configuration process by going to [Target Protection](#).

Target Protection

Target Protection is used to protect files on [target hosts](#) by saving a backup copy before a file is either deleted or overwritten on the target host. If enabled, then whenever a file is deleted or modified on the source host, but before the changes are propagated to the targets, a copy of the existing file on the target is moved to the Peer Management Center trash bin.

The trash bin is located in a hidden folder named **.pc-trash_bin** found in the root directory of the [watch set](#) of the target host. A backup file is placed in the same directory hierarchy location as the source folder in the watch set within the recycle bin folder. If you need to restore a previous version of a file then you can copy the file from the recycle bin into the corresponding location in the watch set and the changes will be propagated to all other collaboration hosts.

Target protection configuration is available by selecting **Target Protection** from the tree node within the File Collaboration Configuration dialog.



The screenshot shows the 'File Collaboration Configuration' dialog box with the 'Target Protection' tab selected. The left sidebar lists various configuration categories, with 'Target Protection' highlighted. The main area displays the following settings:

- Enabled:** A checked checkbox.
- # of Backup Files to Keep:** A numeric input field set to 3.
- # of Days to Keep:** A numeric input field set to 90.
- Trash Bin:** A text input field containing '.pc-trash_bin'.

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

Below are a list of general fields and their descriptions:

Enabled	Enables target protection.
# of Backup Files to Keep	The maximum number of backup copies of an individual file to keep in the trash bin before purging the oldest copy.
# of Days to Keep	The number of days to keep a backup archive copy around before deleting from disk. A value of 0 will disable purging any files from archive.
Trash Bin	The trash bin folder name located in the root directory of the Watch Set. This is a hidden folder and the name cannot be changed by the end-user.

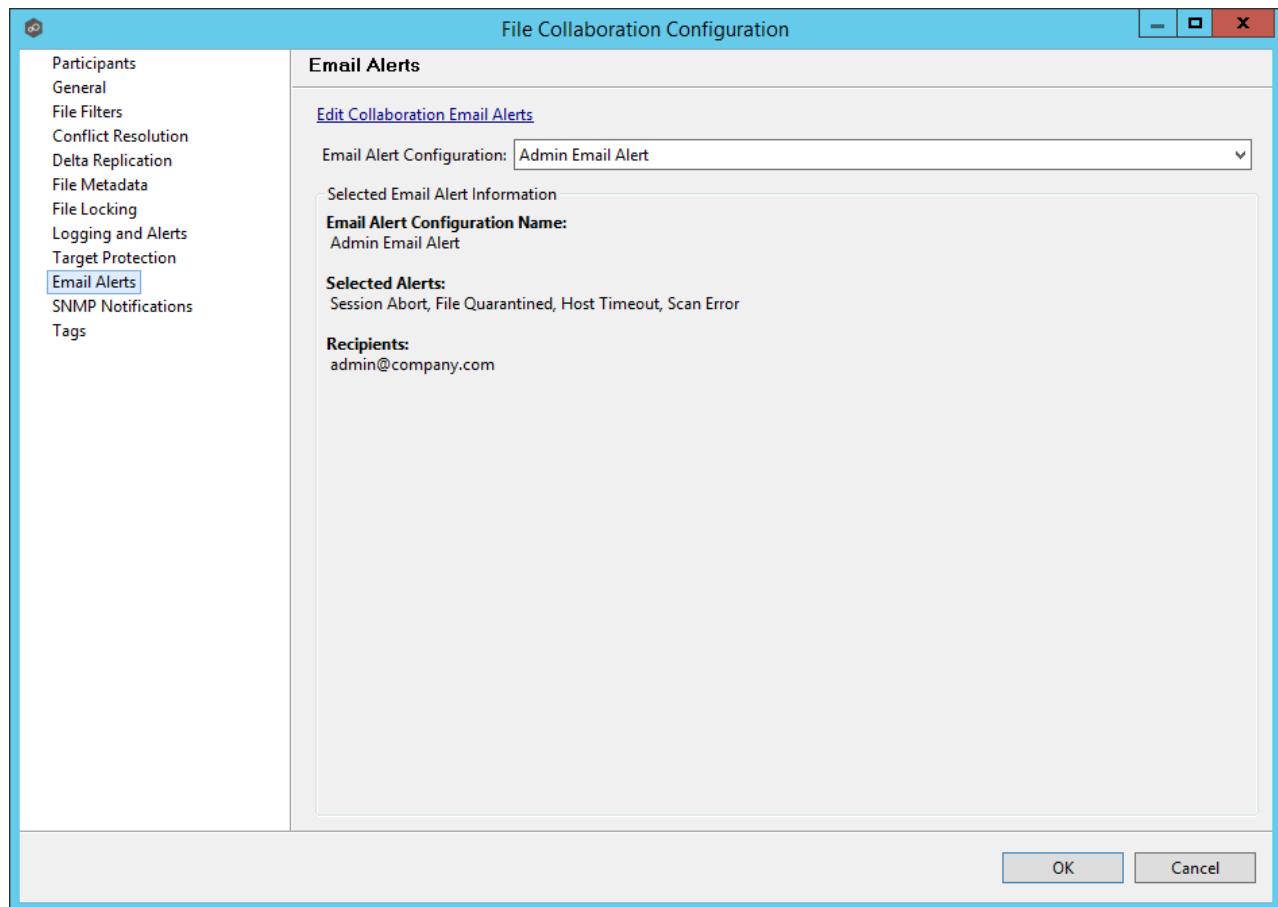
Once all target protection settings are set, you can either [save the configuration](#) and exit the dialog, or you can continue with the configuration process by going to [Email and SNMP Notifications](#).

Email Alerts and SNMP Notifications

Email Alerts

Email alert configuration is available by selecting **Email Alerts** from the tree node within the File Collaboration Configuration dialog.

Email alerts are configured at a global level, then applied to individual file collaboration jobs. The following screen shows how this is accomplished.

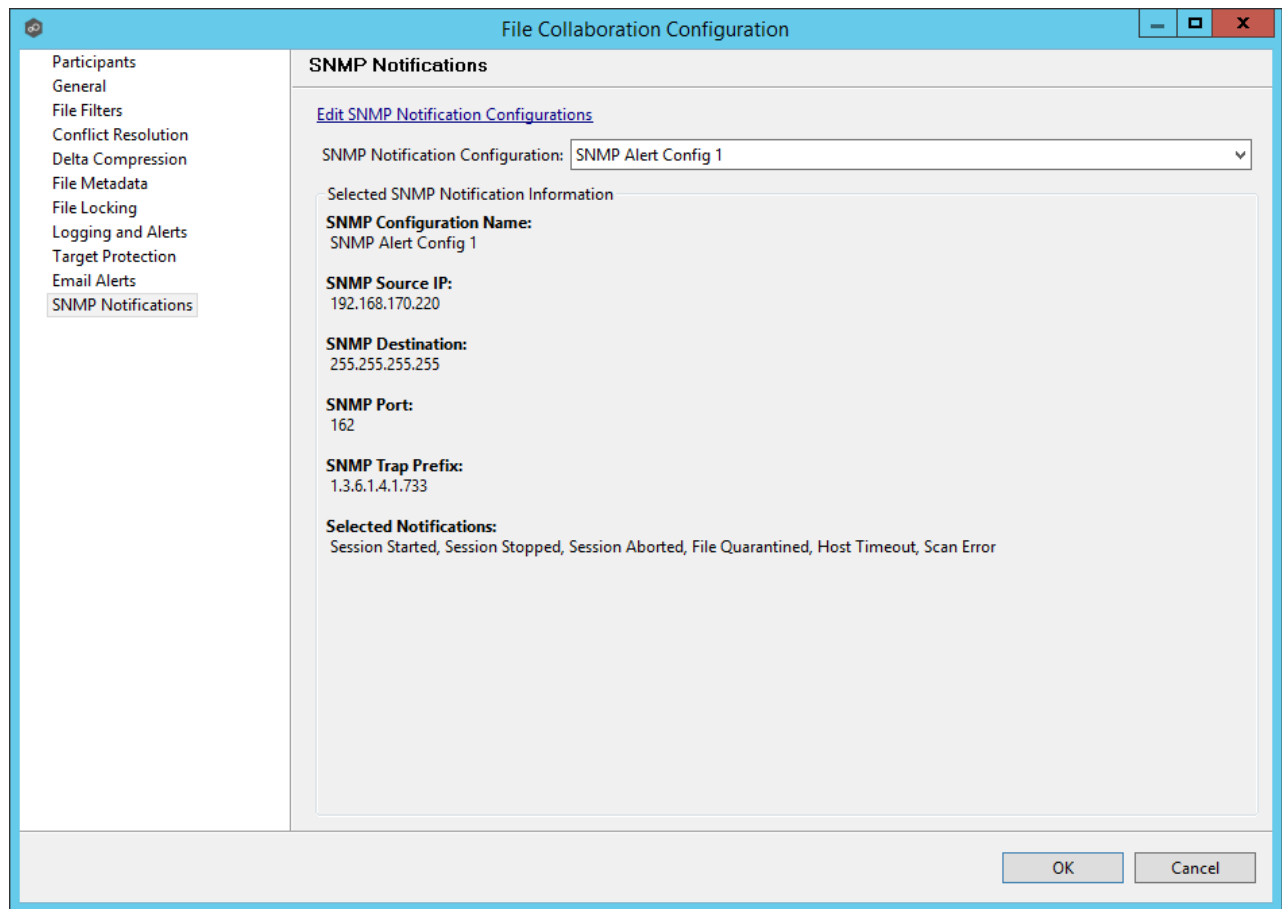


To enable email alerts for this particular job, select an email alert from the drop down list. To disable, select **None - Disabled**. To edit the list of available alerts, select [Configuring Email Alerts](#).

SNMP Notifications

SNMP notification configuration is available by selecting **SNMP Notifications** from the tree node within the File Collaboration Configuration dialog.

SNMP notifications, like email alerts, are configured at a global level, then applied to individual jobs. The following screen shows how this is accomplished:



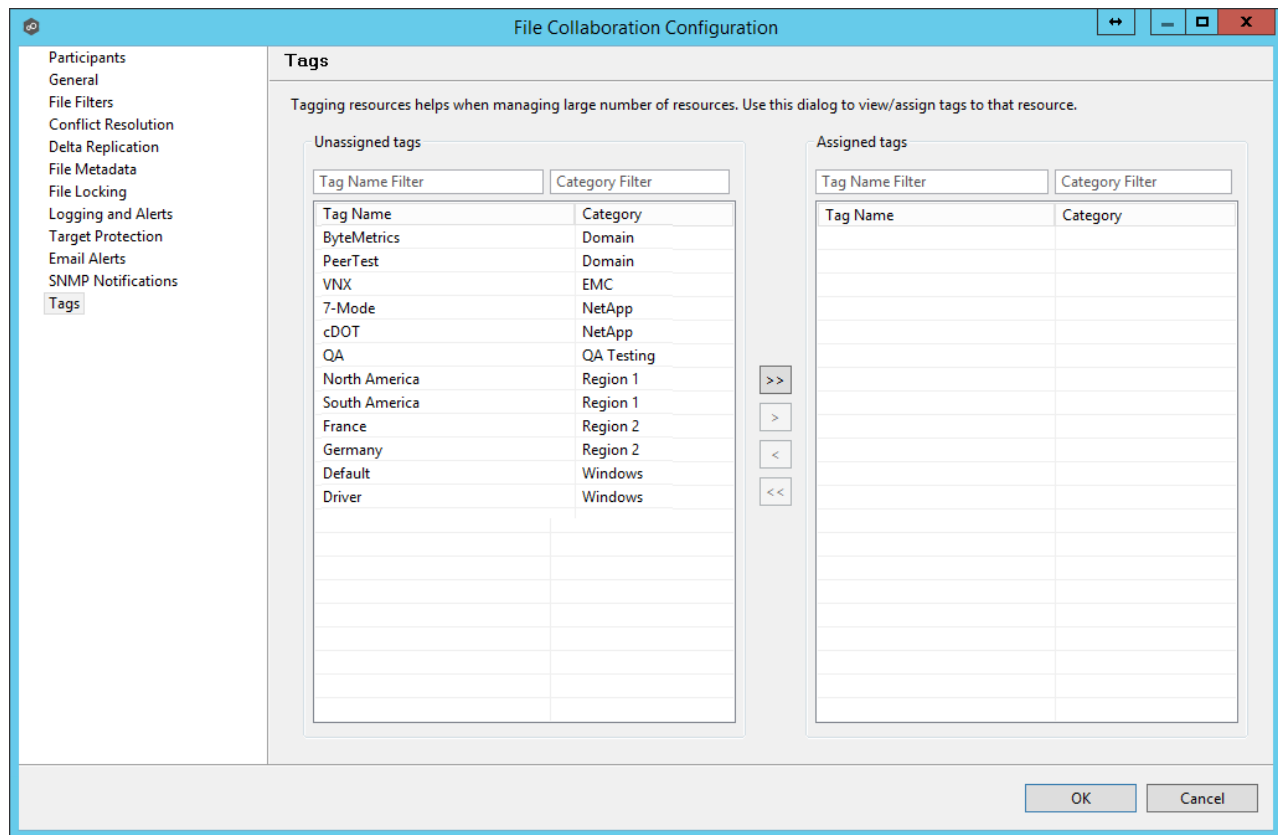
To enable SNMP notifications for this particular job, select an SNMP notification from the drop down list. To disable, select **None - Disabled**. To edit the list of available configurations, select [Configuring SNMP Notifications](#).

Once all email alert and SNMP notification settings are set, you have completed the configuration process and can now [save the configuration](#).

Tags

Use this dialog to assign existing tags and categories to the selected job. This screen is not available during [Multi-Job Editing](#) mode.

To define tags and categories, go to the [Tags Configuration](#) screen available from the **Window > Preferences** menu or by clicking the **User Preferences** button from the [Main View Toolbar](#). See [Using Tags](#) for more details.



Save Settings

Once you have finished configuring the [file collaboration job](#), you will need to save the changes by pressing the OK button at the bottom of the configuration window.

After saving the configuration, the job will be displayed in the [Jobs view](#) in the top left panel of the Peer Management Center. You will also be able to open the job in a tab of the [File Collaboration Runtime View](#).

You are now ready to start the job. See [Running and Managing a File Collaboration job](#) for more information.

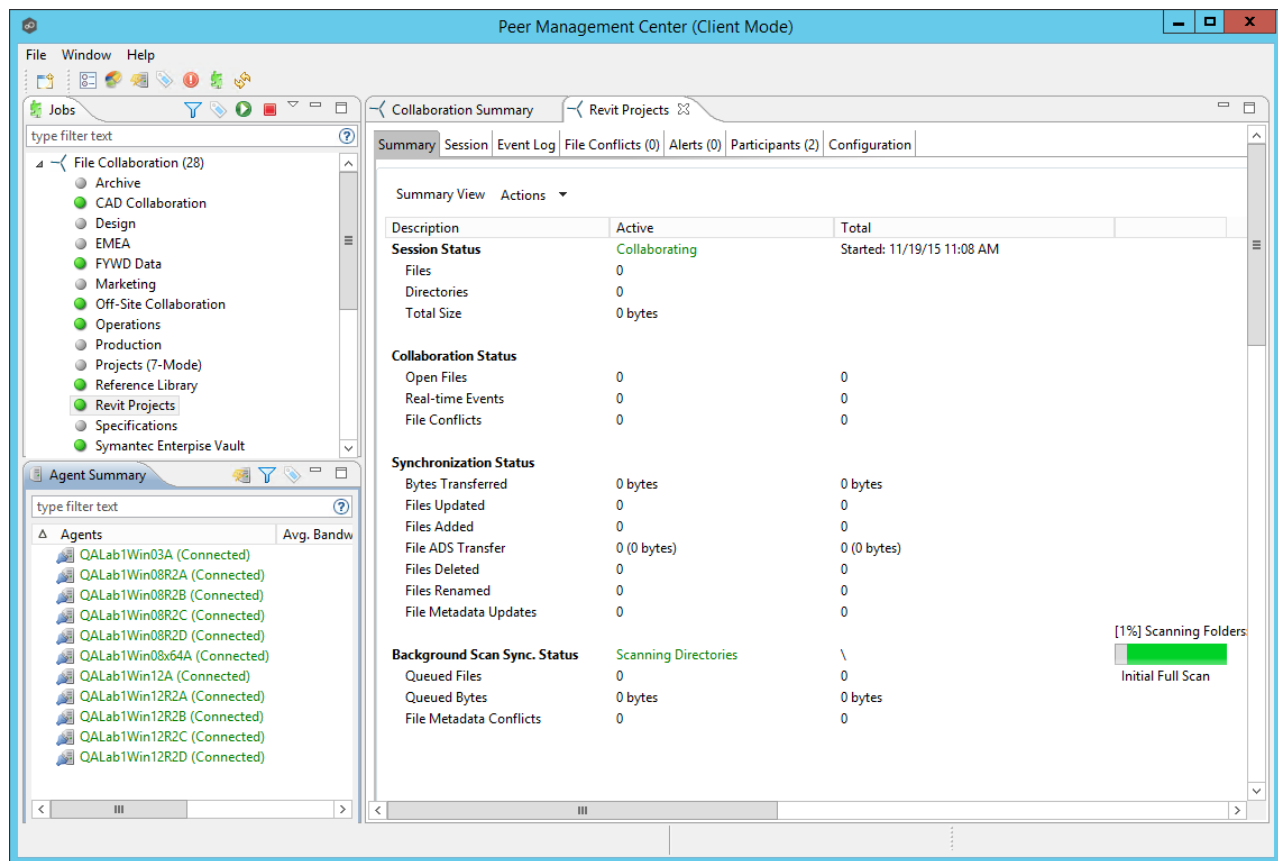
Running and Managing a File Collaboration Job

The topics in this section provide some basic information about starting, stopping and managing File Collaboration jobs.

Overview

[File Collaboration jobs](#) are manually started and stopped in three ways:

- Right-clicking on one or more jobs in the [Jobs view](#).
- Right-clicking on one or more jobs in the [File Collaboration Summary view](#).
- Opening a specific job and clicking the Start/Stop button at the bottom of the job's tab (shown below).



The File Collaboration Runtime View is located in the large center section of the Peer Management Center. It is composed of various tabs (or editors) representing individual file collaboration jobs and/or cross-job [summary information](#). The tabs representing individual jobs consist of the following components:

Runtime View

These tabs allow you to select from the various job-specific views. These views include:

Sub Tabs	<ul style="list-style-type: none"> • Summary (or Status) view - Shows overall statistics for the file collaboration job. The illustration above is displaying the Summary View. • Session view - Shows active open files and files that are currently in transit between participating hosts. • Event Log view - Shows a list of all runtime activity that has occurred within the selected file collaboration job. • File Conflicts view - Shows a list of all files that are quarantined for the session or are in conflict between two or more participating hosts. • Alerts view - Shows a list of all Job Alerts specifically tied to the selected job. • Participants view - Shows a list of all hosts participating in the file collaboration job. • Configuration view - Shows a summary of all configurable options for the selected job.
Job Start/Stop	The button allows you to start and stop the file collaboration job.
Job Status Display	Displays status related messages when the job is running.

Starting and Stopping

Starting a File Collaboration Job

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. The two main options are:

- Have the file collaboration job perform the initial synchronization based on the configured [File Conflict Resolvers](#) strategy.
- 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 will only need to be done the first time that you run the file collaboration job.

If you choose Option 1, click the **Start** button to begin collaboration session initialization. Otherwise, pre-seed each participating host with the necessary data, then click the **Start** button.

Initialization Process

The initialization process consists of the following steps:

1. All participating hosts are contacted to make sure they are online and properly configured.
2. Real-time event detection is initialized on all participating hosts where file locks and changes will be propagated in real-time to all participating hosts. You can view real-time activity and history via the various [Runtime Job Views](#) for the open job.
3. The [initial synchronization process](#) is started, all of the configured root folders on the participating hosts are scanned in the background, and a listing of all folders and files are sent back to the running job.
4. The background directory scan results are analyzed and directory structures compared to see which files are missing from which hosts. In addition, file conflict resolution is performed to decide which copy to use as the master for any detected file conflicts based on the configured [File Conflict Resolvers](#) settings.
5. After the analysis is performed, all files that need to be synchronized are copied to the pertinent host(s).

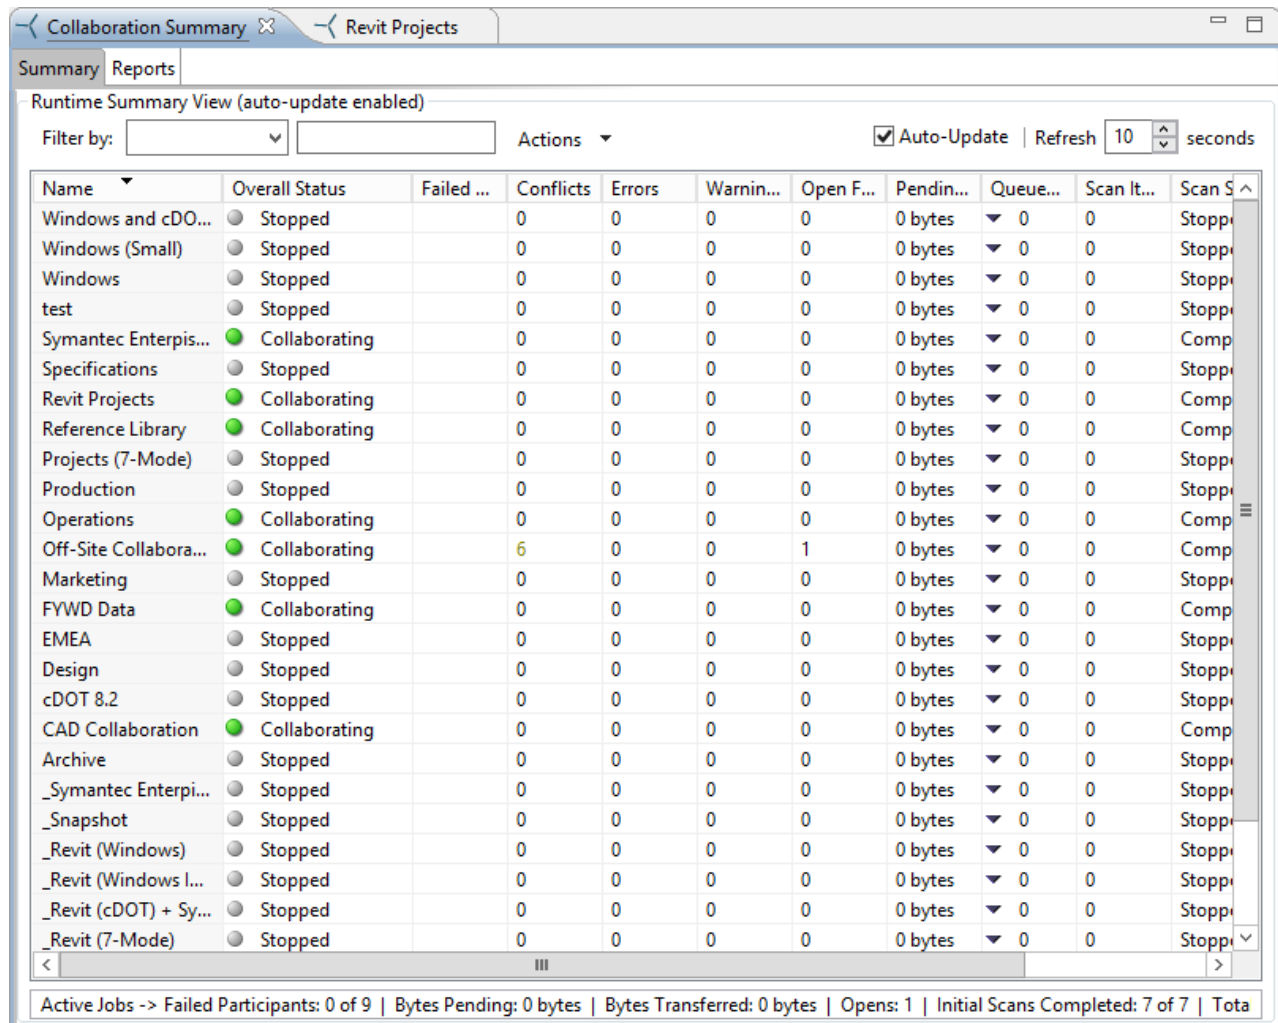
Stopping a File Collaboration Job

You can stop a file collaboration job at any time by pressing the **Stop** button. Doing this will shutdown the real-time file event detection and close all running operations (eg., file transfers).

File Collaboration Summary View

The File Collaboration Summary view aggregates critical status and statistical information from all configured [File Collaboration jobs](#) in a single table view. It is automatically displayed when the [Peer Management Center](#) client is started and can be opened at any other time by double-clicking on the **File Collaboration** parent tree node in the [Jobs view](#) or by clicking on the **View Runtime Summary** icon in the toolbar of the Jobs view. Information in this view can be

sorted and filtered. Operations such as starting, stopping, and editing multiple job at once are available, in addition to the ability to clear [job alerts](#) and purge [file conflicts](#) from stopped jobs.



The screenshot shows the 'Collaboration Summary' window with the 'Summary' tab selected. The window title is 'Revit Projects'. Below the tabs, there's a 'Runtime Summary View (auto-update enabled)' section. It includes a 'Filter by:' dropdown, an 'Actions' dropdown, and a checkbox for 'Auto-Update' which is checked. Next to it is a 'Refresh' button and a spinner set to '10' seconds. The main area is a table with columns: Name, Overall Status, Failed..., Conflicts, Errors, Warnin..., Open F..., Pendin..., Queue..., Scan It..., and Scan S... (partially visible). The table lists various jobs like 'Windows and cDO...', 'Windows (Small)', 'Windows', 'test', 'Symantec Enterpris...', 'Specifications', 'Revit Projects', 'Reference Library', 'Projects (7-Mode)', 'Production', 'Operations', 'Off-Site Collabora...', 'Marketing', 'FYWD Data', 'EMEA', 'Design', 'cDOT 8.2', 'CAD Collaboration', 'Archive', '_Symantec Enterpi...', '_Snapshot', '_Revit (Windows)', '_Revit (Windows I...', '_Revit (cDOT) + Sy...', and '_Revit (7-Mode)'. Each row shows the job name, a status icon (green for Collaborating, grey for Stopped), and numerical values for the other columns. At the bottom, a status bar shows: 'Active Jobs -> Failed Participants: 0 of 9 | Bytes Pending: 0 bytes | Bytes Transferred: 0 bytes | Opens: 1 | Initial Scans Completed: 7 of 7 | Total'.

Name	Overall Status	Failed ...	Conflicts	Errors	Warnin...	Open F...	Pendin...	Queue...	Scan It...	Scan S...
Windows and cDO...	Stopped		0	0	0	0	0 bytes	0	0	Stopp
Windows (Small)	Stopped		0	0	0	0	0 bytes	0	0	Stopp
Windows	Stopped		0	0	0	0	0 bytes	0	0	Stopp
test	Stopped		0	0	0	0	0 bytes	0	0	Stopp
Symantec Enterpris...	Collaborating		0	0	0	0	0 bytes	0	0	Comp
Specifications	Stopped		0	0	0	0	0 bytes	0	0	Stopp
Revit Projects	Collaborating		0	0	0	0	0 bytes	0	0	Comp
Reference Library	Collaborating		0	0	0	0	0 bytes	0	0	Comp
Projects (7-Mode)	Stopped		0	0	0	0	0 bytes	0	0	Stopp
Production	Stopped		0	0	0	0	0 bytes	0	0	Stopp
Operations	Collaborating		0	0	0	0	0 bytes	0	0	Comp
Off-Site Collabora...	Collaborating	6	0	0	0	1	0 bytes	0	0	Comp
Marketing	Stopped		0	0	0	0	0 bytes	0	0	Stopp
FYWD Data	Collaborating		0	0	0	0	0 bytes	0	0	Comp
EMEA	Stopped		0	0	0	0	0 bytes	0	0	Stopp
Design	Stopped		0	0	0	0	0 bytes	0	0	Stopp
cDOT 8.2	Stopped		0	0	0	0	0 bytes	0	0	Stopp
CAD Collaboration	Collaborating		0	0	0	0	0 bytes	0	0	Comp
Archive	Stopped		0	0	0	0	0 bytes	0	0	Stopp
_Symantec Enterpi...	Stopped		0	0	0	0	0 bytes	0	0	Stopp
_Snapshot	Stopped		0	0	0	0	0 bytes	0	0	Stopp
_Revit (Windows)	Stopped		0	0	0	0	0 bytes	0	0	Stopp
_Revit (Windows I...	Stopped		0	0	0	0	0 bytes	0	0	Stopp
_Revit (cDOT) + Sy...	Stopped		0	0	0	0	0 bytes	0	0	Stopp
_Revit (7-Mode)	Stopped		0	0	0	0	0 bytes	0	0	Stopp

Active Jobs -> Failed Participants: 0 of 9 | Bytes Pending: 0 bytes | Bytes Transferred: 0 bytes | Opens: 1 | Initial Scans Completed: 7 of 7 | Total

The Collaboration Summary view is not updated in real-time. This is done for performance reasons. Instead, the table can be set to automatically update itself every few seconds. Checking the **Auto-Update** option will enable this functionality, while the refresh interval (in seconds) can be set right beside the check box. Each refresh cycle will update the totals across all active jobs listed at the bottom of the view. Additional columns can be added to and removed from the table from the right-click context menu.

Selecting one or more items in the table, then right-clicking will bring up a context menu of available actions that can be performed on the selected jobs. The actions that are unique to this table are as follows:

Purge All Conflicts

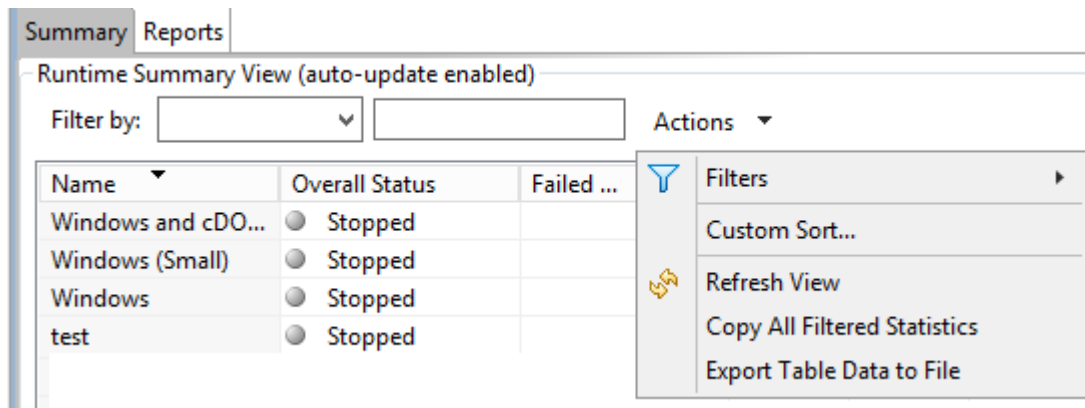
Purges all file conflicts from the selected jobs. This can only be performed on jobs that are not running.

Clear Alerts	Clears all alerts for the selected jobs. This can be performed while a job is running.
Trash-Bin Cleanup	The automatic trash-bin cleanup process runs once daily at 11 PM. Select this option to execute the trash-bin cleanup process on demand.
Show Details	Choose this option to display all the statistics for the selected job in the Runtime Summary Details dialog.
Copy Details	Choose this option to copy detailed information to the system clipboard for the job(s) selected in the table. This information can then be pasted into a document editor.

Double-clicking any item in the table will automatically open the selected File Collaboration job in a tab within the [File Collaboration Runtime View](#), allowing you to drill down and view specific information about that single job.

Items in the summary table can be filtered by a [filter expression](#), built-in states (Running in Good State, Running with Quarantines, Not Running - Stopped, Running with Disconnected Agents, Lost Quorum), job name, [Participant](#), Session Status, or Tags. Select the desired filter or enter your own expression in the text field to the right of the filter drop down list.

Clicking the 9 table menu provides the following options:



Filters	Allows for the selection of built-in or user-defined filters and to save/manage filter expressions . Default job filters include Failed Jobs, Jobs with Backlog, and Running Scans. For example, filter:"Running Scans".
----------------	--

Custom Sort...	Use the Custom Sort option to configure and save how you want the Collaboration Summary view table to be sorted and keep important items visible at the top. For example, you may choose to create a sort level where the Overall Status column is sorted in Ascending order by default.
Refresh View	Refresh all information provided in the table.
Copy All Filtered Statistics	Copy detailed information to the system clipboard for all items current displayed in the table, taking any filters into account. This information can then be pasted into a document editor.
Export Entire Table to File	Dump the entire contents of the table to a text file that can be viewed in any document editor.

Runtime Reports View

The Runtime Reports view aggregates critical statistical information from all configured File Collaboration jobs in a single table view. The **Reports** tab is visible when the global **Enable Advanced Reporting Tab** option is checked. This tab is especially useful to see the number of files that are in the queue waiting to be synchronized (File Sync Queue).

Summary **Reports**

Pending Activity (auto-update enabled)

Filter by: Actions ☐ Auto-Hide ☒ Auto-Update Refresh 10 seconds

Name	File Sync Queue	Real-Time Queue	Queued Bytes	Mods	Adds	Metadata	Backgrc
Specifications	0	0	0 bytes	0	0	0	0
Revit Projects	230	0	257 MB	146	0	0	0
Reference Library	0	0	0 bytes	0	0	0	0
Projects (7-Mode)	0	0	0 bytes	0	0	0	0
Production	0	0	0 bytes	0	0	0	0
Operations	0	0	0 bytes	0	0	0	0
Off-Site Collabora...	0	0	0 bytes	0	0	0	0
Marketing	0	0	0 bytes	0	0	0	0
FYWD Data	0	0	0 bytes	0	0	0	0
EMEA	0	0	0 bytes	0	0	0	0
Design	0	0	0 bytes	0	0	0	0
cDOT 8.2	0	0	0 bytes	0	0	0	0
CAD Collaboration	0	0	0 bytes	0	0	0	0
Archive	0	0	0 bytes	0	0	0	0
_Symantec Enterpi...	0	0	0 bytes	0	0	0	0
_Snapshot	0	0	0 bytes	0	0	0	0
_Revit (Windows)	0	0	0 bytes	0	0	0	0
_Revit (Windows I...	0	0	0 bytes	0	0	0	0
_Revit (cDOT) + Sy...	0	0	0 bytes	0	0	0	0
_Revit (7-Mode)	0	0	0 bytes	0	0	0	0
_Mixed (2)	0	0	0 bytes	0	0	0	0
_Links	0	0	0 bytes	0	0	0	0
_Defrag	0	0	0 bytes	0	0	0	0
Totals	230	0	257 MB	146	0	0	0

Global Event Processor Queue : 0 | Pending Scans: 0 | Running Scans: 0

The Runtime Reports view is not updated in real-time. This is done for performance reasons. Instead, the table can be set to automatically update itself every few seconds. Enabling the **Auto-Update** option will turn on this functionality, while the **Refresh** interval (in seconds) can be set right beside the check box. Each refresh cycle will update the totals across all jobs listed at the bottom of the view by default. Additional columns can be added to and removed from the table from the right-click context menu.

Name	The name of the configured job.
File Sync Queue	The number of files that are in queue waiting to be processed. The number of threads available for this queue is set by the global Performance Real-Time Background Threads option.
Real-Time Queue	The number of open/close events that are in queue waiting to be processed. The number of threads available to process this queue is set by the global Performance Real-Time Expedited Threads option.
Queued Bytes	The number of bytes that are in queue waiting to be processed.

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

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

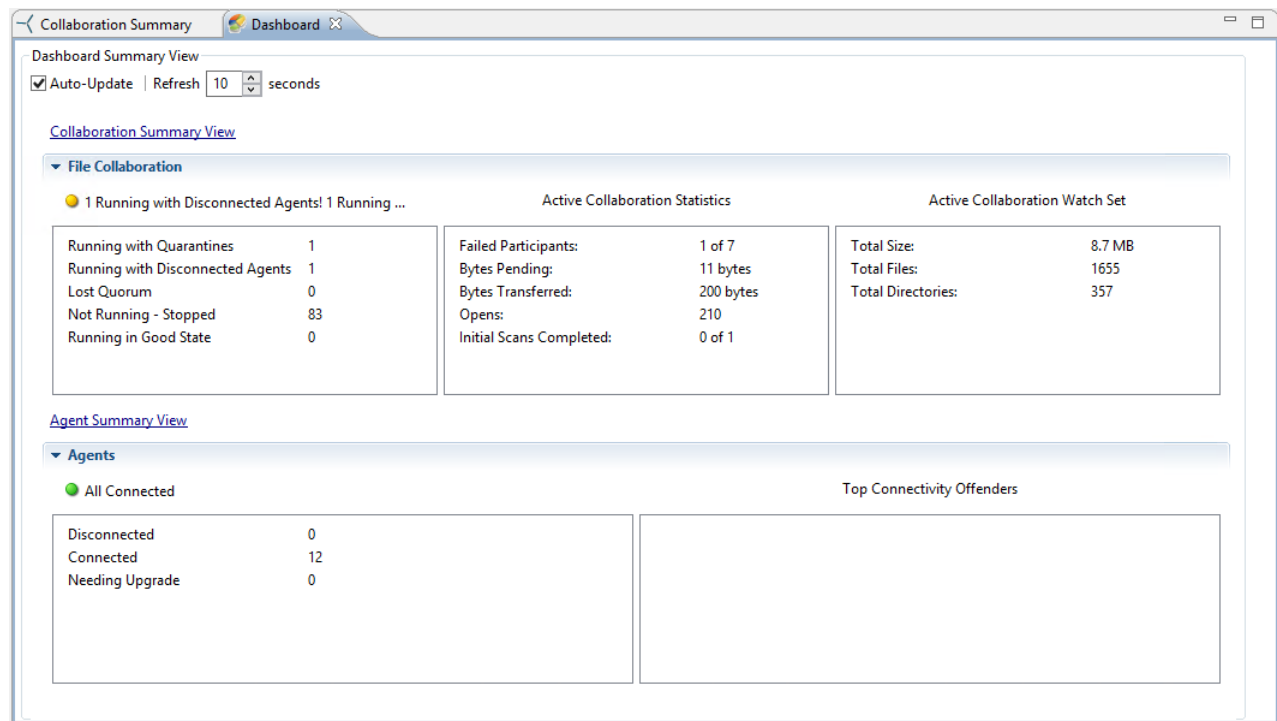
Clicking on the **Actions** table menu provides the following options:

Filters	Allows for the selection of built-in or user-defined filters and to save/manage filter expressions . Default job filters include Failed Jobs, Jobs with Backlog, and Running Scans. For example, filter:"Running Scans".
Custom Sort...	Use the Custom Sort option to configure and save how you want the Runtime Reports table to be sorted and keep important items visible

	at the top. For example, you may choose to create a sort level where the Overall Status column is sorted in Ascending order by default.
Refresh View	Refresh all information provided in the table.
Copy All Filtered Statistics	Copy detailed information to the system clipboard for all items current displayed in the table, taking any filters into account. This information can then be pasted into a document editor.
Export Entire Table to File	Dump the entire contents of the table to a text file that can be viewed in any document editor.
Move Totals Row To Top	Moves the Totals row to the top of the table.
Move Totals Row To Bottom	Moves the Totals row to the bottom of the table.

Dashboard Summary View

The File Collaboration **Dashboard Summary View** is a panel that displays metrics and key performance indicators from all running [File Collaboration jobs](#). It is automatically displayed when the Peer Management Center client is started and can be opened at any other time by selecting **View Dashboard** from the [Windows](#) menu or by clicking on the **View Dashboard** icon in the Peer Management Center [toolbar](#).



The Dashboard is not updated in real-time. This is done for performance reasons. Instead, the table can be set to automatically update itself every few seconds. Enabling the **Auto-Update** option will enable this functionality, while the interval (in seconds) can be set right beside the check box.

Entries in the first column of the **File Collaboration** and **Agents** categories can be double-clicked, which will take the user to a filtered runtime view of the selected item for additional details.

Peer Agent Detail Summary View

The **Peer Agent Detail Summary** view is a panel 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. The Peer Agent Detail Summary View can be opened by selecting **View Agent Detail Summary** from the [Windows](#) menu or by clicking on the **View Agent Detail Summary** icon in the [Peer Management Center](#) or [Peer Agent Summary View](#) toolbars.

Collaboration Summary Dashboard Agents

Agent Summary View

Filter by: Expression

type filter text

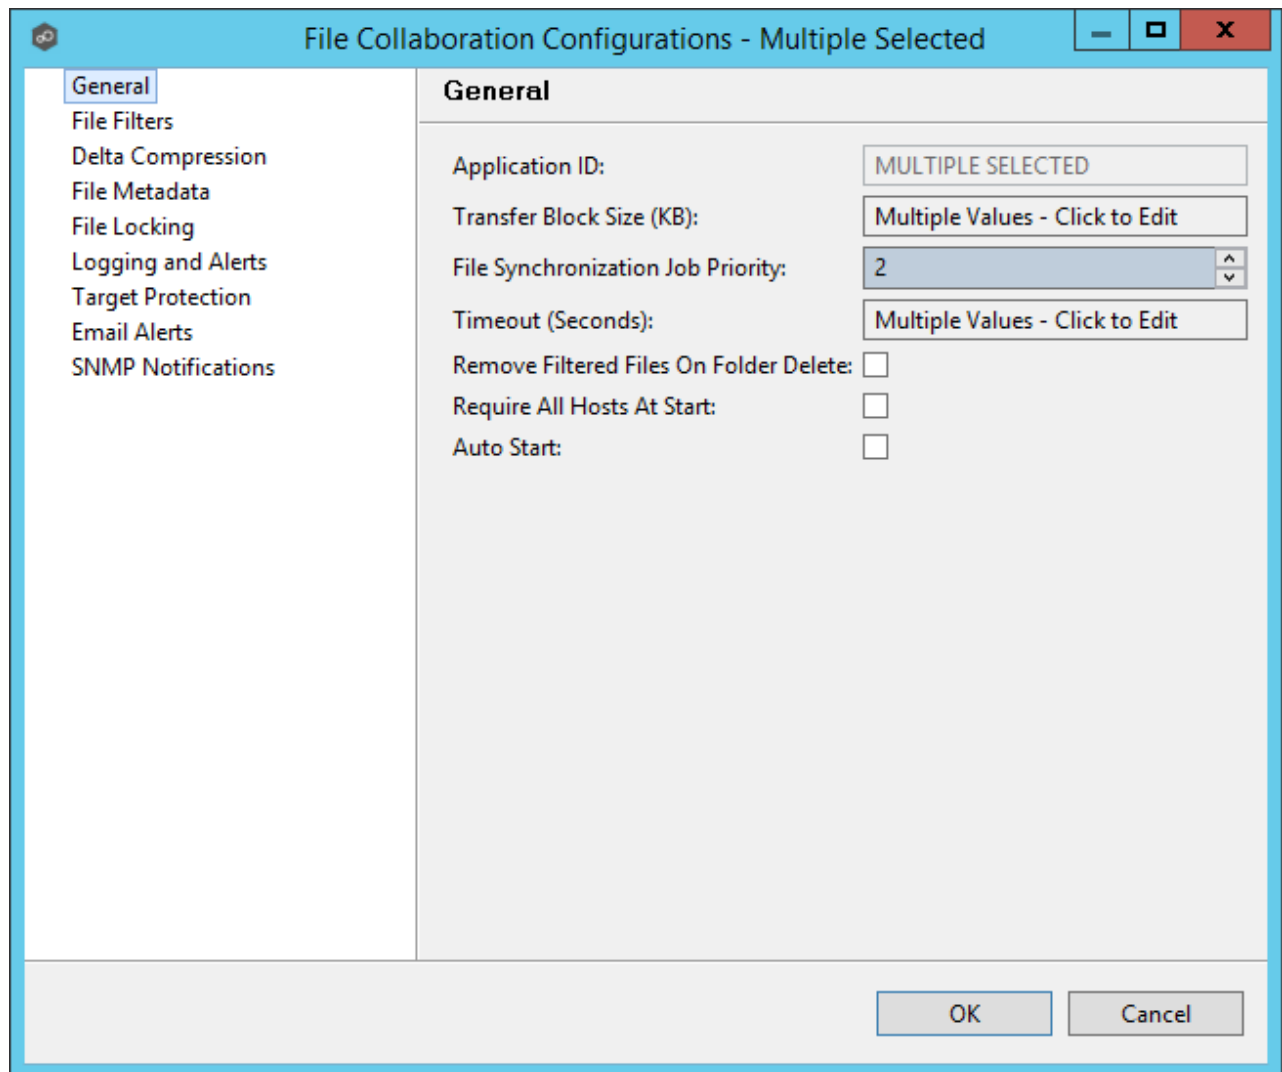
Agents	Version	JVM Architecture	Total Missed Heartbeats	Total Agent Disconnects	Total Pending Disconnects	Mem. Load
CEEIsilon1 (Connected)	4.1.0.20170414	amd64	0	0	0	67
CEEIsilon4 (Connected)	4.1.0.20170414	amd64	0	0	0	33
CEEVNX2 (Connected)	4.1.0.20170414	amd64	0	0	0	72
Composite2J (Connected)	4.1.0.20170414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

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

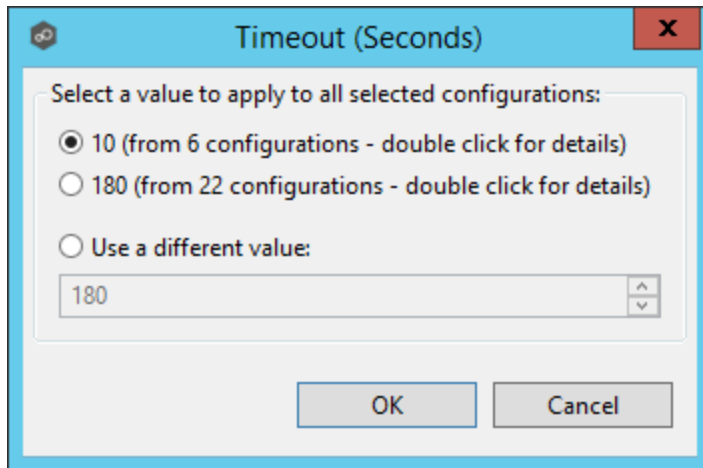
Editing Multiple Jobs

The [Peer Management Center](#) supports Multi-Job Editing, allowing you to quickly and effectively manipulate multiple [File Collaboration jobs](#) at once. For example, you can use this feature to change a single configuration item such as **Auto Start** for any number of already configured jobs in one operation instead of having to change the item individually on each. While this feature does cover most of the options available on a per-job basis, certain options are unavailable in multi-job edit mode, specifically ones tied to [participants](#). See the section on [Creating a File Collaboration job](#) for more details on specific configuration items.

For the most part, the original configuration dialog remains the same with a few minor differences depending on similarities between the selected file collaboration jobs. A sample dialog is as follows:



In this dialog, any discrepancies between multiple selected jobs will generally be illustrated by a read-only text field with the caption **Multiple Values - Click to Edit**. Clicking this field will bring up a dialog similar to the following:



This dialog gives you the option of choosing a value that is already used by one or more selected file collaboration jobs, in addition to the ability to use your own value. Notice that variances in the look and feel of this pop-up dialog above depend on the type of information it is trying to represent (for example, text vs. a check box vs. a list of items).

Upon pressing **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 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 screens operate in a slightly different manner than mentioned above. All of the necessary information is provided at the top of these screens in **bold** writing.

Host Connectivity Issues

Unavailable Hosts

Peer Management Center is designed to be run in an environment where all [participating hosts](#) are highly available and on highly available networks.

If a host becomes unavailable while a File Collaboration job is running, and is unreachable within the configured timeout period (specified within the jobs [General Settings](#)), it may be removed from collaboration. If no response is received while performing a file collaboration operation within the timeout period then the host will be pinged, and if still no response, the host will be taken out of the running session, a FATAL event will be logged, and the [Participants View](#) for the job will be updated to indicate that the host has failed. In addition, if [Email Alerts](#) and/or [SNMP Notifications](#) are configured and enabled for **Host Timeouts**, then the appropriate message(s) will be sent.

If auto-restart support (see below) is not enabled, you will need to stop and start the file collaboration job in order to bring any failed hosts back into the session. As a result, all root

folders on all hosts will need to be scanned again to detect any inconsistencies. Therefore, if you are operating over a WAN with low bandwidth you will want to set the timeout to a higher value on each related jobs.

Quorum

In order for a File Collaboration job to run correctly, a quorum of available hosts must be met. Quorum is currently set to at least 2 hosts, and if quorum is not met then the collaboration session will automatically be terminated. If [Email Alerts](#) and/or [SNMP Notifications](#) are configured and enabled for **Session Aborts**, then the appropriate message(s) will be sent.

Auto-Restart

Peer Management Center includes support for automatically restarting file collaboration jobs that include [participating hosts](#) that have been disconnected, and have reconnected and are once again available. After a host becomes unavailable and quorum is lost on a running file collaboration job, the job will automatically stop running and enter a pending state, waiting for one or more hosts to become available again so that quorum can be met. Once quorum is met, the pending job will automatically be restarted, beginning with a scan of all root folders. In a job where a host becomes unavailable but quorum is not lost, the remaining hosts will continue collaborating. If the unavailable host becomes available once again, it will be brought back into the running job and a background scan will begin on all participating hosts, similar in fashion to the initial background scan at the start of a job.

Configuration

This functionality is enabled on a global level for all file collaboration jobs and is configured by clicking on the **Window** menu within the Peer Management Center, then selecting **Preferences**. Within the opening dialog, select **File Collaboration** in the tree on the left. The following screen will be displayed:

Host Connectivity options are as follows:

Auto Restart Job when Host Available	If checked, auto-restart functionality will be enabled for all running file collaboration jobs.
Minimum Host Reconnect Time (in minutes)	The minimum time in minutes a host must be reconnected before reestablishing the host within any relevant file collaboration jobs.
Enable Advanced Reporting Tab	Check this option to display the Reports tab in the Collaboration Summary screen.

Disabling auto-restart on a per-job and host instance is performed within the Participant View for the desired file collaboration job. For more information on managing and disabling auto-restart at the job level, please see the section on the [Participant View](#).

Runtime Job Views

Each file collaboration job has seven primary Runtime Views used for viewing a combination of real-time file I/O activity, history, and configuration. These views also provide the ability to manage specific collaboration runtime functionality.

The seven views are as follows:

- [Summary view](#)
- [Session view](#)
- [Event Log view](#)
- [File Conflicts view](#)
- [Alerts view](#)
- [Participants view](#)
- [Configuration view](#)

The Summary view allows you to view current and cumulative file collaboration and synchronization statistics, as well background synchronization status.

Summary	Session	Event Log	File Conflicts (1)	Alerts (2)	Participants (2)	Configuration
<div>Summary View Actions ▾</div> <div> <input checked="" type="checkbox"/> Auto-Update Refresh 10 seconds </div>						
Description	Active	Total				
Session Status	Collaborating	Started: 11/20/15 10:13 AM				
Files	1569					
Directories	42					
Total Size	3.8 GB					
Collaboration Status						
Open Files	110	614				
Real-time Events	0	0				
File Conflicts	1	1				
Synchronization Status						
Bytes Transferred	21.3 MB	21.3 MB --> Delta Savings 99.49% (4.1 GB)				
Files Updated	110	870				
Files Added	0	0				
File ADS Transfer	0 (0 bytes)	0 (0 bytes)				
Files Deleted	0	0				
Files Renamed	0	0				
File Metadata Updates	0	0				
Background Scan Sync. Status	Completed	6 minutes 31 seconds	100%			
Queued Files	0	259				
Queued Bytes	0 bytes	1.4 GB				
File Metadata Conflicts	0	0				

The Session Summary View is made up of the following sections:

Session Status

This section displays current statistics for all files/folders contained in the running [file collaboration job](#).

Files	Current number of files in the running file collaboration job. Files that are excluded by filtration will not be included in this statistic.
Directories	Current number of subfolders under the watch sets of the running file collaboration job.
Total Size	Cumulative number of bytes of all files in the running file collaboration job.
Start Time	The date and time of the last start of the file collaboration job, manual or automatic.

Collaboration Status

Open Files (Active)	Displays the number of files that are currently being collaborated on, where a user has a file open on the source host and the system is holding locks on all target hosts.
Open Files (Total)	Displays the total number of files that have been opened since the session was started, where locks were propagated to target hosts.
Real-Time Events (Active)	Displays the current number of real-time file events that are pending action.
Real-Time Events (Total)	Displays the total number of real-time event received since the running file collaboration job was started.
File Conflicts (Active)	Displays the current number of files that are in some type of conflicted state.
File Conflicts (Total)	Displays the total number of file conflicts (including pending initial synchronization) that have occurred since the running file collaboration job was started.

Synchronization Status

This section displays current and cumulative statistics for all files that have been added, removed, renamed or modified since the running file collaboration job was started.

Bytes Transferred (Active)	Total number of bytes currently being transferred to target hosts by the running file collaboration job.
Bytes Transferred (Total)	<p>Total number of bytes that have be transferred to target hosts since the file collaboration job was started. If delta-level replication is enabled then the total delta encoding savings will also be displayed as percentage along with the actual cumulative size of the source files.</p> <p>For example a value of 3.9MB --> Delta Savings 47.75% (7.6MB) should be interrupted as a total of 3.9MB were transferred corresponding to the actual total source size of 7.6MB for a savings of 47.75% or 3.7MB.</p>

	Keep in mind that the delta savings also averages in files where delta encoding may not have been used.
Files Updated (Active)	Total number of files currently being updated or that are scheduled to be updated.
Files Updated (Total)	Total number of files that have been modified since the file collaboration job was started.
Files Added (Active)	Total number of files currently being added to the session or that are scheduled to be added.
Files Added (Total)	Total number of files that have been added since the file collaboration job was started.
File ADS Transfer (Active)	Total number and bytes of Alternate Data Streams being synced or scheduled to be synced
File ADS Transfer (Total)	Total number and bytes of Alternate Data Streams synced since the file collaboration job was started.
Files Deleted (Active)	Total number of files currently being deleted or that are scheduled to be deleted.
Files Deleted (Total)	Total number of files that have been deleted since the file collaboration job was started.
Files Renamed (Active)	Total number of files currently being renamed or that are scheduled to be renamed.
Files Renamed (Total)	Total number of files that have been renamed since the file collaboration job was started.
File Metadata Updates (Active)	Total number of files pending file metadata (file attributes and security descriptor) updates.
File Metadata Updates (Total)	Total number of file metadata (file attributes and security descriptor) changes that have occurred since the file collaboration job was started.

Background Synchronization Status

This section displays overall status of the [initial synchronization process](#) performed at the start of the session, as well as current and cumulative statistics for files that needed to be synchronized.

Background Sync. Status	<p>Text label indicating the current status of the initial synchronization process. Valid values are:</p> <ul style="list-style-type: none"> • Stopped: Session is stopped. • Completed: Initial scan and synchronization processes have completed. • Synchronizing Files: Background scan and initial synchronization processes are currently running. • When the status is Synchronizing Files, the Total column will display the directory that is currently being synchronized.
Queued Files (Active)	Total number of files currently being synchronized or that are scheduled to be synchronized.
Queued Files (Total)	Total number of files that have been synchronized by the session as a result of the initial synchronization process.
Queued Bytes (Active)	Total number of bytes currently being synchronized or that are scheduled to be synchronized by the initial synchronization process.
Queued Bytes (Total)	Total number of bytes that have been synchronized by the session as a result of the initial synchronization process.
File Metadata Conflicts (Active)	Total number of file metadata conflicts that are currently being acted on as a result of the initial synchronization process.
File Metadata Conflicts (Total)	Total number of file metadata conflicts that were found as a result of the initial synchronization process.

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

Summary **Session** Event Log File Conflicts (1) Alerts (2) Participants (2) Configuration

Open Files (79)

Session Status: Collaborating | Filter by Host: Filter by: Actions ▾

File Path	Host	Is Source	User Name	Sync. Status	File Size	Last Modified	Date Opened	Message
▶ \FLDR3\FLDR1L1\FILE-new5.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:01	11-20-2015 10:31:50	
▶ \FILE-new5.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:03	11-20-2015 10:28:27	
▶ \FLDR2\FILE-new13.TXT	QALab1W...	true	Session		5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:57	
▶ \FLDR3\FLDR3L1\FILE-new18.TXT	QALab1W...	true	Session		5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:35	
▶ \FLDR2\FLDR3L1\FILE-new18.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:30:59	11-20-2015 10:32:11	
▶ \FILE-new18.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:04	11-20-2015 10:28:27	
▶ \FLDR3\FLDR3L1\FILE-new18.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:46	
▶ \FLDR3\FLDR1L1\FILE-new18.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:01	11-20-2015 10:31:58	
▶ \FLDR3\FLDR1L1\FILE-new18.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:01	11-20-2015 10:31:57	
▶ \FLDR3\FLDR3L1\FILE-new18.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:45	
▶ \FLDR3\FLDR3L1\FILE-new18.TXT	QALab1W...	true	Session	96%	5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:36	
▶ \FLDR2\FLDR3L1\FILE-new18.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:08	
▶ \FILE-new6.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:03	11-20-2015 10:28:27	
▶ \FLDR2\FILE-new9.TXT	QALab1W...	true	Session	8%	5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:19	
▶ \FILE-new17.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:04	11-20-2015 10:28:27	
▶ \FLDR2\FLDR3L1\FILE-new17.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:30:59	11-20-2015 10:32:10	
▶ \FLDR2\FILE-new14.TXT	QALab1W...	true	Session		5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:56	
▶ \FLDR2\FILE-new1.TXT	QALab1W...	true	Session	Scheduled	5.7 MB	11-20-2015 10:30:59	11-20-2015 10:28:57	Pending Synchronizat
▶ \FLDR3\FLDR3L1\FILE-new3.TXT	QALab1W...	true	Session	96%	5.7 MB	11-20-2015 10:31:02	11-20-2015 10:31:37	
▶ \FILE-new3.TXT	QALab1W...	true	MonikaC		5.7 MB	11-20-2015 10:31:03	11-20-2015 10:28:27	
▶ \FLDR3\FLDR1L1\FILE-new8.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:31:01	11-20-2015 10:31:51	
▶ \FLDR2\FILE-new8.TXT	QALab1W...	true	Session	8%	5.7 MB	11-20-2015 10:31:00	11-20-2015 10:32:20	
▶ \FLDR2\FLDR3L1\FILE-new8.TXT	QALab1W...	true	Session	52%	5.7 MB	11-20-2015 10:30:59	11-20-2015 10:32:05	

III

☒ Auto-Update | Refresh 10 seconds

The Session View is made up of the following components:

Session Status	<p>Text label indicating the current status of the session. Valid values are:</p> <ul style="list-style-type: none"> • Stopped: Session is stopped. • Starting: Session is starting up. • Collaborating: Real-time event detection is enabled and session is collaborating. • Stopping: Session is in the process of stopping.
-----------------------	--

Open Files Table	<p>A table showing all currently open files on the source host, any internal file locks being held by the running file collaboration job on the target host(s), and file summary information. This table will also show all file transfers currently in progress along with file summary information, status and overall progress. Clicking on any column headers will sort by that column in ascending or descending order.</p> <p>All items listed in this table are grouped by file path. Each associated lock and/or transfer for each participating host will be available as a hidden child item of a root row. The root row represents the file on the source host. Pressing the + next to the root will show all associated file transfers and/or locks.</p>
Host Filter	A drop down list of participating hosts to filter on. Selecting a specific host will filter the Open Files to just show files on that host.
Filter By Combo	A drop down list of additional filters that can be applied to the Open Files table, including filtering by user name (associated with the opening, adding, deleting, or modification of a file), and by file name.
Actions Menu	<p>Menu items include:</p> <ul style="list-style-type: none"> • Refresh View: Refresh the entire Open Files table to show the latest list of file transfers and locks. • Validate Session Locks: Clicking this link will perform validation of all locks in the session and will report any potential issues. You should perform this action if you believe a file is not open in the session, but the user interface indicates that the file is open, or vice-versa. • File System Report: Generate a text file listing all files and folders being collaborating on within the running file collaboration job.

The Event Log view allows you to view recent file event history for the currently running [file collaboration job](#) based on your [Logging and Alerts](#) settings. You can specify the maximum number of events to store in the table by adjusting the Display Events spinner located in the top right corner of the panel. The maximum number of events that can be viewed is 3,000. If you need to view more events or events from a prior session, then you can use the log files saved in the 'Hub\logs' directory located in the installation directory. The event log files will start with **fc_event.log** and are written in a tab delimited format. Microsoft Excel is a good tool to use to view and analyze a log file. See the [Logging and Alerts](#) settings for more information about log files.

You can click on any column header to sort by the column. For example, clicking on 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 highlighted in light gray, Errors are highlighted in red and Fatal errors are highlighted in orange. Error records will also contain an error message in the Message column.

Summary | Session | **Event Log** | File Conflicts (1) | Alerts (2) | Participants (2) | Configuration

Event Log (Auto-Update Disabled)

0 errors, 2 warnings, 1498 others | Filter by Severity: | Filter by: | Actions ▼

Date	Severity	Type	Host	Is Source	File	Comments	Message	Username	File Size
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\\FLDR1\\FLDR1LT1\\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\\FLDR2\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\\FLDR1\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\\FLDR1\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\\FLDR1\\FLDR1LT1\\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\\FLDR2\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\\FLDR1\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\\FLDR1\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\\FLDR2\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\\FLDR1\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\\FLDR1\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\\FLDR1\\FLDR1LT1\\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\\FLDR2\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\\FLDR1\\FLDR1LT1\\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\\FLDR1\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\\FLDR1\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Modify	QALab1Wi...	false	\\FLDR1\\FLDR1LT1\\...	Updated 0.48...		Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\\FLDR2\\FLDR1LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	true	\\FLDR1\\FLDR2LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Close	QALab1Wi...	false	\\FLDR1\\FLDR2LT1\\...			Session	5.7 MB
11-20-2015 10:3...	INFO	File Lock	QALab1Wi...	false	\\FLDR2\\FLDR1LT1\\...			Session	5.7 MB

☐ Auto-Update | Refresh seconds | Display Events

Clicking on the Actions table menu provides the following options:

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

Introduction

Files conflicts can occur for the following reasons:

- Two or more users open a file at the same time before all files can be locked down by the running file collaboration job.
- A file is already opened by a user when a file collaboration job is started and the file size and timestamp does not match the other target hosts.
- A file is already opened by two or more users when a file collaboration job is started.
- A file was modified on two or more hosts between job restarts or network outages.
- A general I/O failure occurs on the [Source Host](#) after the file has been modified, but before the file is synchronized to all [Target Hosts](#). In this case, the file will automatically be quarantined.

When a file conflict is detected, the file is placed in the File Conflict list (shown below) with a specific status which will determine how the conflict is resolved. The three possible file conflict statuses along with their resolution strategies are as follows:

Conflict Status	Resolution Strategy
Pending Conflict Resolution	<p>This status will be assigned to files that have already been verified or synchronized by the session via the initial synchronization process. When all files in use are closed by users on the source hosts, the files will be analyzed to determine if a file conflict has occurred as follows:</p> <ul style="list-style-type: none"> • If more than one file has been modified then the file will be quarantined by updating the file conflict status to quarantined. • If only one file has been modified then that file will be used as the source, synchronized with all other participating hosts, and removed from the File Conflict list • If no files have been modified then no action will be taken and the file will be removed from the File Conflict list
Pending Initial Synchronization	<p>This status will be assigned to files that have not been verified or synchronized by session via the initial synchronization process. When all files in use are closed by users on the source hosts, then standard file conflict resolution will be performed based on the configured File Conflict Resolvers. However, if the "Quarantine Offline Multi-Edits" option is enabled, then if a file is modified on 2 or more hosts while the collaboration session is not running, and the last modified timestamps are all newer than the last timestamp recorded by the collaboration session, then the file will be quarantined.</p>

Quarantined	A file will be quarantined when a file conflict with "Pending Conflict Resolution" status cannot be resolved or a fatal I/O error occurs. Quarantined files will need to be explicitly removed from the File Conflict list.
--------------------	---

When a file conflict occurs, the status will be set to **Pending Conflict Resolution** if the file has already been verified or synchronized by the initial synchronization process, otherwise the file conflict status will be set to **Pending Initial Synchronization**. If the conflict is a result of a fatal I/O error on the source then the file conflict status will be set to **Quarantined**.

Note: If a file collaboration job is stopped before a file conflict with a status of **Pending Conflict Resolution** is resolved, then that file will automatically be quarantined the next time the file collaboration job is started.

File Conflict and Quarantine Scenarios

- A job is started and Initial Scan Logic is performed on a file

If file has never been synchronized by Peer Management Center and if file sizes and last modified times do not match on all collaboration hosts, or if file does not exist on one or more hosts, then the file will be synchronized based on the configured file conflict resolver, which is typically most recent last modified time. Files that have previously been synchronized by Peer Management Center where just a single file's last modified timestamp is newer than the last recorded timestamp, then that file will be synchronized to all other hosts; however, if two or more files have a more recent last modified timestamp than was last recorded timestamp, then the file will be quarantined (this is the default behavior and can be disabled by deselecting the File Conflict Resolvers "Quarantine Offline Version Conflicts" configuration option).

- A single user has a file opened before starting a collaboration job

A file conflict will be created with a status of "Pending Initial Synchronization". After the user closes the file, if all file sizes and timestamps match then the file conflict is removed and no synchronization is performed. However, if any file last modified times or file sizes do not match, the file will be synchronized or quarantined based on the configured file conflict resolution strategy and according to the initial scan logic detailed above. Once the file is synchronized, the file conflict will be removed.

- Two or more users have a file open before starting a collaboration job

A file conflict will be created with a status of "Pending Conflict Resolution". After the users close all files the conflict will be removed if the last modified timestamp matches on all files, otherwise if the file has never been synchronized by Peer Management Center, then the file conflict will be updated to quarantined. However, if the file has previously been synchronized by Peer Management Center, then the file will be synchronized or quarantined based on the configured file conflict resolution strategy and according to the initial scan logic detailed above.

- Two or more users open a file at the same time

In the rare situation when two users open a file at the same time, or in-and-around the same time and Peer Management Center is unable to obtain corresponding locks on target hosts before this happens (this is dependent on WAN latency and other factors), then a file conflict will be created with a status of "Pending Conflict Resolution." After all users close the files, file lock conflict resolution will be performed as follows:

- If all files last modified timestamps and file sizes match, then the file conflict will be removed.
- If only a single file has been modified, then the file that changed is synchronized or quarantined based on the configured file conflict resolver and according to the initial scan logic detailed above.
- If two or more files have been modified since it was opened, then the file conflict status will be updated to quarantined.

Quarantined Files

Once a file is marked as **Quarantined**, the file will no longer participate in collaboration, and thus changes to any version of the file will not be propagated to other hosts. However, subsequent file activity on a quarantined file will be logged in the event log as a warning so you can determine who modified the file while it was quarantined. Quarantined files are saved to disk and will survive session restarts. The File Conflict list displays the time and date of the quarantine along with an error message indicating the reason for the quarantine (see below). A Quarantined File event is also logged in the Event Log and you can obtain a more detailed reason for the quarantine by analyzing the Event Log file(s). In addition, if [Email Alerts](#) and/or [SNMP Notifications](#) are configured and enabled for **File Quarantines**, then the appropriate message(s) will be sent.

Removing a File from Quarantine

You must explicitly remove a file from quarantine in order to have it participate in the collaboration session once again. To remove a file from quarantine, select the file in the File Conflict list, select the host with the correct version, and press the **Release Conflict** button. After doing this all hosts are checked to make sure the file is not currently locked by anybody. If no locks are found, then locks are obtained on all versions of the file and the targets that are out-of-date are synchronized with the selected source host. You may also chose to perform no action, in which case the file is removed from the File Conflict list but none of the file versions are modified; therefore if the files are not currently in-sync, then the next time the file is modified, changes will be propagated to the other hosts. If an error occurs while removing the file conflict, then the Status field in the File Conflict table is updated to reflect the error.

You may also select multiple files to remove from the conflict list at once.

[illegible]

The right-click context menu for the table contains the following actions that are unique to this particular view:

Refresh View	Refresh all information provided in the table.
Clear Alerts	Clears all alerts for the selected job. This can be performed while a job is running.

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

[illegible]

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.

Copyright (c) 1993-2018 Peer Software, Inc. All Rights Reserved

hosts. If a host has become unavailable, an error message will be displayed next to the failed host in red.

[illegible]

The Participants view also contains a table that displays the most recent host participant state changes, e.g., when a host was removed from collaboration session, or when a host came back online. This functionality is broken down into two parts: right-click context menu items and a subview entitled **Host Participant State Change Log**.

The following unique items are available in a right-click context menu for the top part of the Participants view:

Disable Host Participant	Temporarily disables the selected participant from taking part in the file collaboration job. You might want to do this if the host is experiencing temporary network outages.
Cancel Auto Restart	This menu item is only available if the global auto-restart functionality enabled and the selected host has been removed from the file collaboration job that is currently being viewed. The canceling of the auto-restart functionality for the host will only be in effect until the next time you start the file collaboration job. If quorum has been lost for the job, canceling auto-restart on all unavailable hosts will prevent the job from automatically restarting.

	If quorum has not been lost, canceling auto-restart will simply prevent a host from automatically re-joining collaboration.
--	---

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

The following items are available in the right-click context menu for this table:

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

This view displays a quick summary of all configurable items for the selected [file collaboration job](#). Each page of the [File Collaboration Configuration](#) dialog is represented in its own part of the view and can be collapsed if desired. Clicking **Edit this File Collaboration Configuration** will immediately bring you to the File Collaboration Configuration dialog where you can edit the current configuration.

Summary | Session | Event Log | File Conflicts (1) | Alerts (4) | Participants (2) | Configuration

[Edit this File Collaboration Configuration](#)

Currently Running Configuration Summary

▼ Selected Participants and Configurations

QALab1Win12R2C \\QASVM2\vol1\QATesting\Revit (Detector Type: NetApp cDOT - Local Path: /QASVM2_vol1)

QALab1Win08R2C \\QASVM1\vol1\QATesting\Revit (Detector Type: NetApp cDOT - Local Path: /QASVM1_vol1)

▼ General Settings

Session Name:	Revit Projects
Session ID:	107
Transfer Block Size:	128 KB
Verify Checksum:	true
Global Real-Time Expedited Threads:	30
Global Real-Time Background Threads:	60
File Synchronization Job Priority:	2
Timeout:	180 Seconds
Scan Delay:	10
Remove Filtered Files On Folder Delete:	false
Require All Hosts At Start:	false
Auto Start:	false

▼ Selected File Filters

Default (Any) - 0

Excluded Wildcard Expressions: ~*, *.BAK, *.BCK, *.WBK, *.ASD, *.XLK, *.DWL*, *.ACS, *.SV\$, <<^.*\atmp[0-9]{4,}\$>>, *.SLOG, *.LNK, *.LDB, *.LACCDB

Included Wildcard Expressions:

Date Filter: Include all dates

Cloud Sync

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

- [Overview](#)
- [Before You Create Your First Cloud Sync Job](#)
- [Creating a Cloud Sync Job](#)
- [Running a Cloud Sync Job](#)
- [Monitoring Your Cloud Sync Jobs](#)
- [Recovering Data from the Cloud](#)

Overview

Cloud Sync brings file to object replication into Peer Software's capabilities for enterprise NAS environments. Leveraging the same real-time engine that powers Peer Software's multi-site, multi-vendor replication, Cloud Sync efficiently pushes data into Microsoft Azure Blob storage in an open format that is immediately consumable by other applications and services.

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

Before You Create Your First Cloud Sync Job

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

Creating a Cloud Sync Job

The Create Job Wizard walks you through the process of creating a Cloud Sync job. The process consists of the following steps:

[Step 1: Job Type and Name](#)

[Step 2: Source Storage Platform](#)

[Step 3: Management Agent](#)

[Step 4: Storage Information](#)

[Step 5: Source Paths](#)

[Step 6: Exclusions](#)

[Step 7: Destination](#)

[Step 8: Credentials](#)

[Step 9: Miscellaneous Options](#)

[Step 10: Sync and Retention Policy](#)

[Step 11: Sync Schedule](#)

[Step 12: Retention](#)

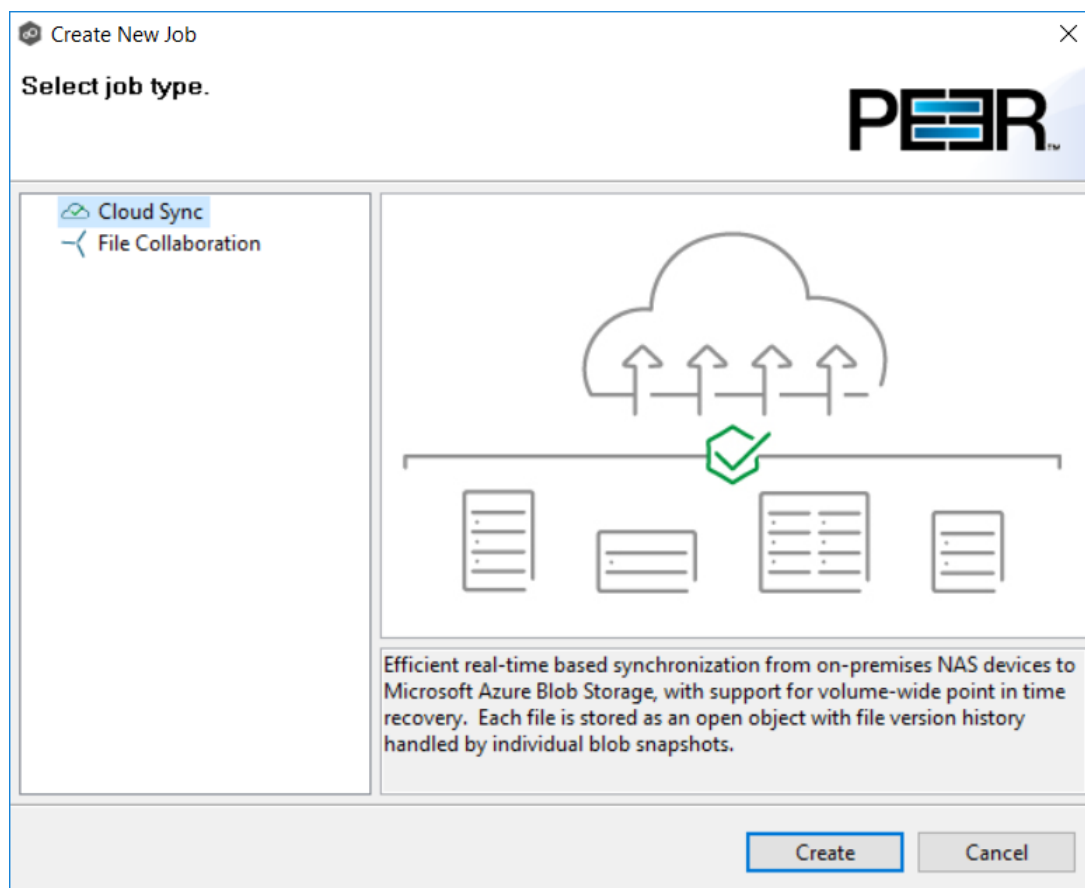
[Step 13: Email Alerts](#)

[Step 14: Confirmation](#)

Step 1: Job Type and Name

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

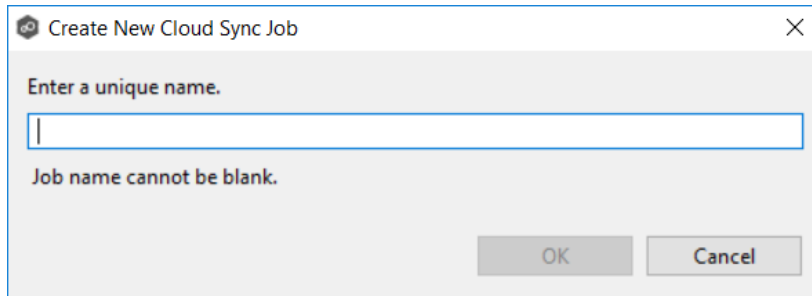
The **New Job** wizard displays a list of job types you can create.



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

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

The job name must be unique.

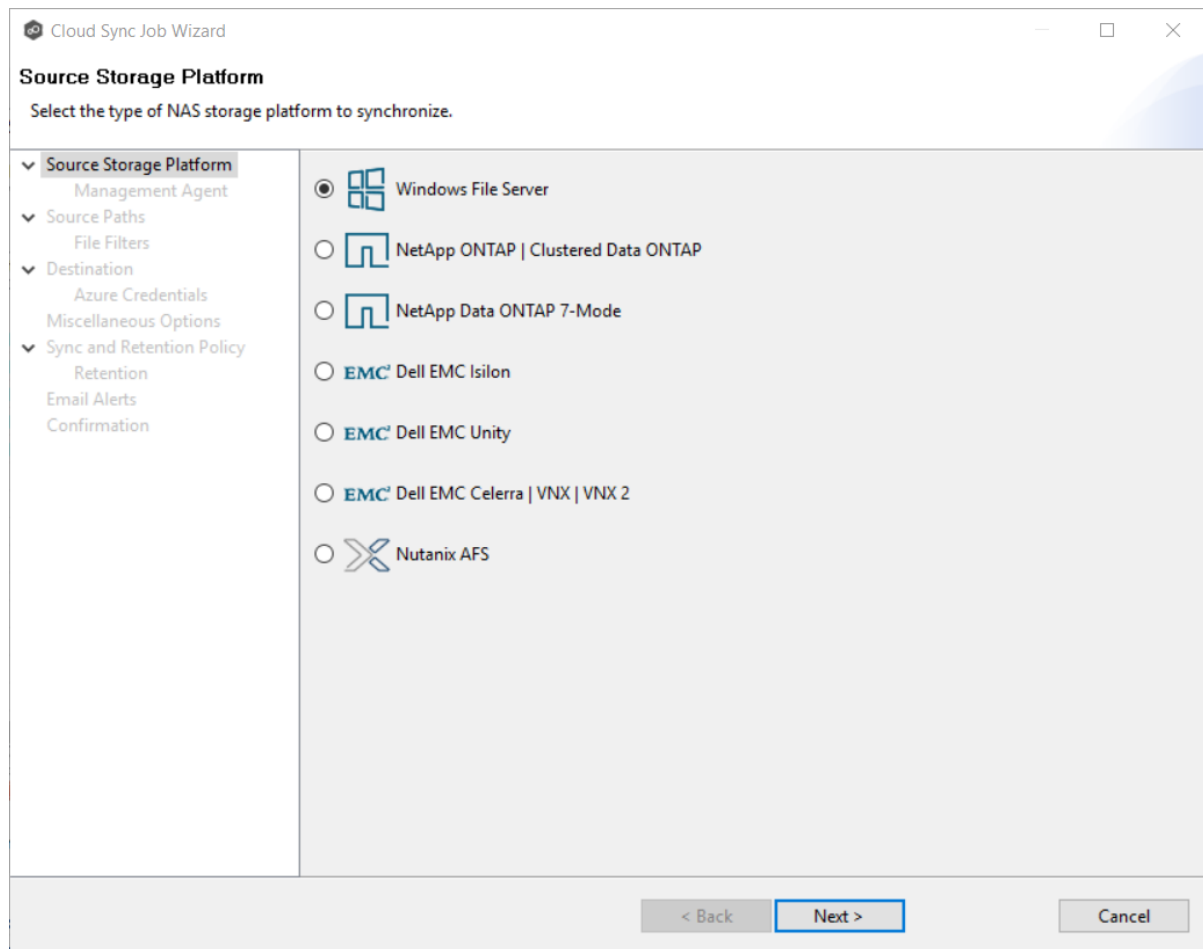
A screenshot of a dialog box titled "Create New Cloud Sync Job". The dialog has a close button (X) in the top right corner. Inside, there is a text input field with the placeholder text "Enter a unique name." and a small cursor. Below the input field, there is a message "Job name cannot be blank." in red. At the bottom right, there are two buttons: "OK" and "Cancel".

5. Click **OK**.

Step 2: Source Storage Platform

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

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



2. Click **Next**.

Step 3: Management Agent

Each storage platform that you want to replicate must have a Peer Agent that manages that device. This associated Peer Agent is also known as the **Management Agent**.

1. Select the server that the Management Agent is installed on.

If you selected **New Credentials**, enter the credentials for connecting to the storage platform. The information you are prompted to enter varies, depending on the type of storage platform:

[NetApp ONTAP | Clustered Data ONTAP](#)

[NetApp Data ONTAP 7-Mode](#)

[EMC Dell EMC Isilon](#)

[EMC Dell EMC Unity](#)

[EMC Dell EMC Celerra | VNX | VNX 2](#)

[Nutanix AFS](#)

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

After the credentials are validated, a success message appears.

4. Click **Next**.

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

Cloud Sync Job Wizard

Storage Information

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

- Source Storage Platform
 - Management Agent
 - Storage Information**
 - Source Paths
 - File Filters
 - Destination
 - Azure Credentials
 - Miscellaneous Options
 - Sync and Retention Policy
 - Retention
 - Email Alerts
 - Confirmation

Credentials

☒ New Credentials

*SVM Name:

*SVM Username:

*SVM Password:

SVM Management IP:

*Peer Agent IP:

☐ Existing Credentials

Access Path

☐ Override Access Path

Access Path:

Having trouble connecting? Please verify that all [prerequisites](#) are met for NetApp ONTAP/cDOT environments.

< Back

Next >

Cancel

SVM Name	Enter the name of the Storage Virtual Machine hosting the data to be replicated.
SVM Username	Enter the user name for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Password	Enter the password for the account managing the Storage Virtual Machine. This must not be a cluster management account.
SVM Management IP	Enter the IP address used to access the management API of the NetApp Storage Virtual Machine. If the data LIFs (Logical Interfaces) corresponding to the SVM Name above already allow management access, this field is not required.

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

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

The screenshot shows the 'Cloud Sync Job Wizard' window, specifically the 'Storage Information' step. The window title is 'Cloud Sync Job Wizard'. Below the title bar, the text 'Storage Information' is displayed, followed by the instruction 'Enter the information required to connect to the source storage platform.'.

On the left side, there is a navigation pane with the following items:

- Source Storage Platform
 - Management Agent
 - Storage Information** (highlighted)
- Source Paths
 - File Filters
- Destination
 - Azure Credentials
 - Miscellaneous Options
- Sync and Retention Policy
 - Retention
 - Email Alerts
 - Confirmation

The main content area is divided into two sections:

- Credentials:** This section contains two radio buttons: 'New Credentials' (selected) and 'Existing Credentials'. Below the 'New Credentials' radio button is a text field labeled '*Filer Name:' and a button labeled 'Advanced'. Below the 'Existing Credentials' radio button is a dropdown menu.
- Access Path:** This section contains a text field labeled 'Access Path:'.

At the bottom of the main content area, there is a 'Validate' button. Below the 'Validate' button, a message reads: 'Having trouble connecting? Please verify that all [prerequisites](#) are met for NetApp 7-Mode environments.'

At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'.

File Name

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

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

The screenshot shows the 'Cloud Sync Job Wizard' window, specifically the 'Storage Information' step. The window title is 'Cloud Sync Job Wizard'. Below the title bar, the text 'Storage Information' is displayed, followed by the instruction 'Enter the information required to connect to the source storage platform.' On the left side, there is a navigation pane with the following items: 'Source Storage Platform' (expanded), 'Management Agent', 'Storage Information' (selected), 'Source Paths', 'File Filters', 'Destination', 'Azure Credentials', 'Miscellaneous Options', 'Sync and Retention Policy', 'Retention', 'Email Alerts', and 'Confirmation'. The main area contains the 'Credentials' section with two radio buttons: 'New Credentials' (selected) and 'Existing Credentials'. Under 'New Credentials', there are four text input fields: '*Cluster Name:', '*Cluster Username:', '*Cluster Password:', and 'Cluster Management IP:'. An 'Advanced' button is located to the right of the 'Cluster Management IP' field. Below the 'Existing Credentials' radio button is a dropdown menu. The 'Access Path' section has a checkbox for 'Override Access Path' and a text input field for 'Access Path:'. At the bottom left of the main area is a 'Validate' button. Below the main area, there is a message: 'Having trouble connecting? Please verify that all [prerequisites](#) are met for EMC Isilon environments.' At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'.

Cluster Name

Enter the name of the EMC Isilon cluster hosting the data to be replicated.

Cluster Username	Enter the user name for the account managing the EMC Isilon cluster.
Cluster Password	Enter the password for account managing the EMC Isilon cluster.
Cluster Management IP	Enter the IP address of the system used to manage the EMC Isilon cluster. Required only if multiple Access Zones are in use on the cluster.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

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

Cloud Sync Job Wizard

Storage Information

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

- Source Storage Platform
 - Management Agent
 - Storage Information**
- Source Paths
 - File Filters
- Destination
 - Azure Credentials
 - Miscellaneous Options
- Sync and Retention Policy
 - Retention
 - Email Alerts
 - Confirmation

Credentials

☒ New Credentials

*NAS Server Name:

*Unisphere Username:

*Unisphere Password:

*Unisphere Management IP:

☐ Existing Credentials

Access Path

☐ Override Access Path

Access Path:

Having trouble connecting? Please verify that all [prerequisites](#) are met for EMC Unity environments.

< Back Next > Cancel

NAS Server Name	Enter the name of the NAS server hosting the data to be replicated.
Unisphere Username	Enter the user name for the Unisphere account managing the Unity storage device.
Unisphere Password	Enter the password for the Unisphere account managing the Unity storage device.
Unisphere Management IP	Enter the IP address of the Unisphere system used to manage the Unity storage device. This should not point to the NAS server.

**Override
Access Path**

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

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

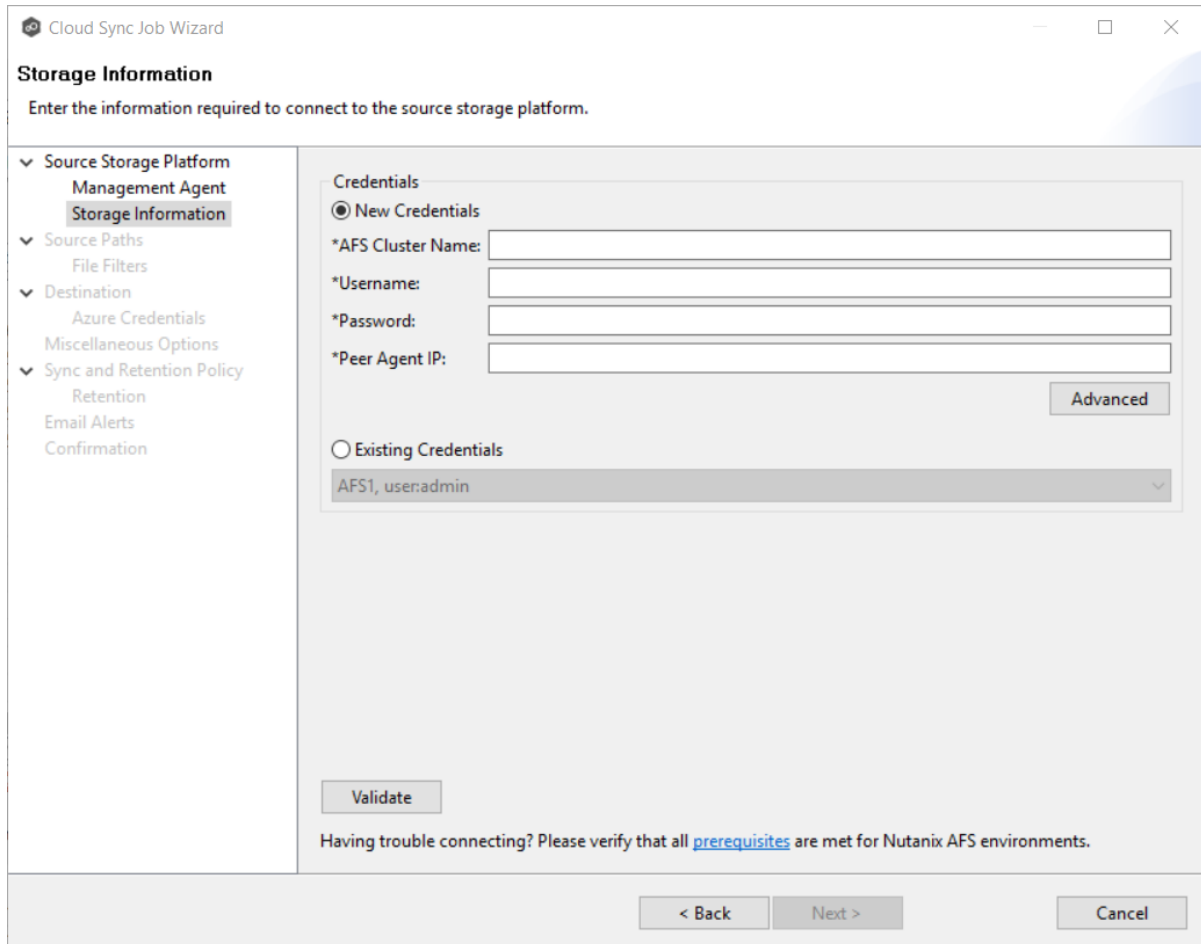
The screenshot shows the 'Cloud Sync Job Wizard' window, specifically the 'Storage Information' step. The window title is 'Cloud Sync Job Wizard'. The main heading is 'Storage Information' with a subtitle 'Enter the information required to connect to the source storage platform.' On the left is a navigation pane with the following items: 'Source Storage Platform' (expanded), 'Management Agent', 'Storage Information' (selected), 'Source Paths', 'File Filters', 'Destination', 'Azure Credentials', 'Miscellaneous Options', 'Sync and Retention Policy', 'Retention', 'Email Alerts', and 'Confirmation'. The main area contains the 'Credentials' section with two options: 'New Credentials' (selected with a radio button) and 'Existing Credentials' (unselected). Under 'New Credentials', there are four text input fields: '*CIFS Server Name:', '*Control Station Username:', '*Control Station Password:', and '*Control Station IP:'. An 'Advanced' button is located to the right of these fields. Below the 'Existing Credentials' option is a dropdown menu. The 'Access Path' section has an 'Override Access Path' checkbox (unchecked) and an 'Access Path:' text input field. At the bottom left is a 'Validate' button. At the bottom right are '< Back', 'Next >', and 'Cancel' buttons. A footer note reads: 'Having trouble connecting? Please verify that all [prerequisites](#) are met for EMC VNX/Celerra environments.'

**CIFS Server
Name**

Enter the name of the CIFS Server hosting the data to be replicated.

Control Station Username	Enter the user name for the Control Station account managing the Celerra/VNX storage device.
Control Station Password	Enter the password for the Control Station account managing the Celerra/VNX storage device.
Control Station IP	Enter the IP address of the Control Station system used to manage the Celerra/VNX storage device. This should not point to the CIFS Server.
Override Access Path	Used only when experiencing access issues. Contact Peer Software support for more information.

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



Cloud Sync Job Wizard

Storage Information

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

- Source Storage Platform
 - Management Agent
 - Storage Information**
 - Source Paths
 - File Filters
 - Destination
 - Azure Credentials
 - Miscellaneous Options
 - Sync and Retention Policy
 - Retention
 - Email Alerts
 - Confirmation

Credentials

☒ New Credentials

*AFS Cluster Name:

*Username:

*Password:

*Peer Agent IP:

Advanced

☐ Existing Credentials

AFS1, user:admin

Validate

Having trouble connecting? Please verify that all [prerequisites](#) are met for Nutanix AFS environments.

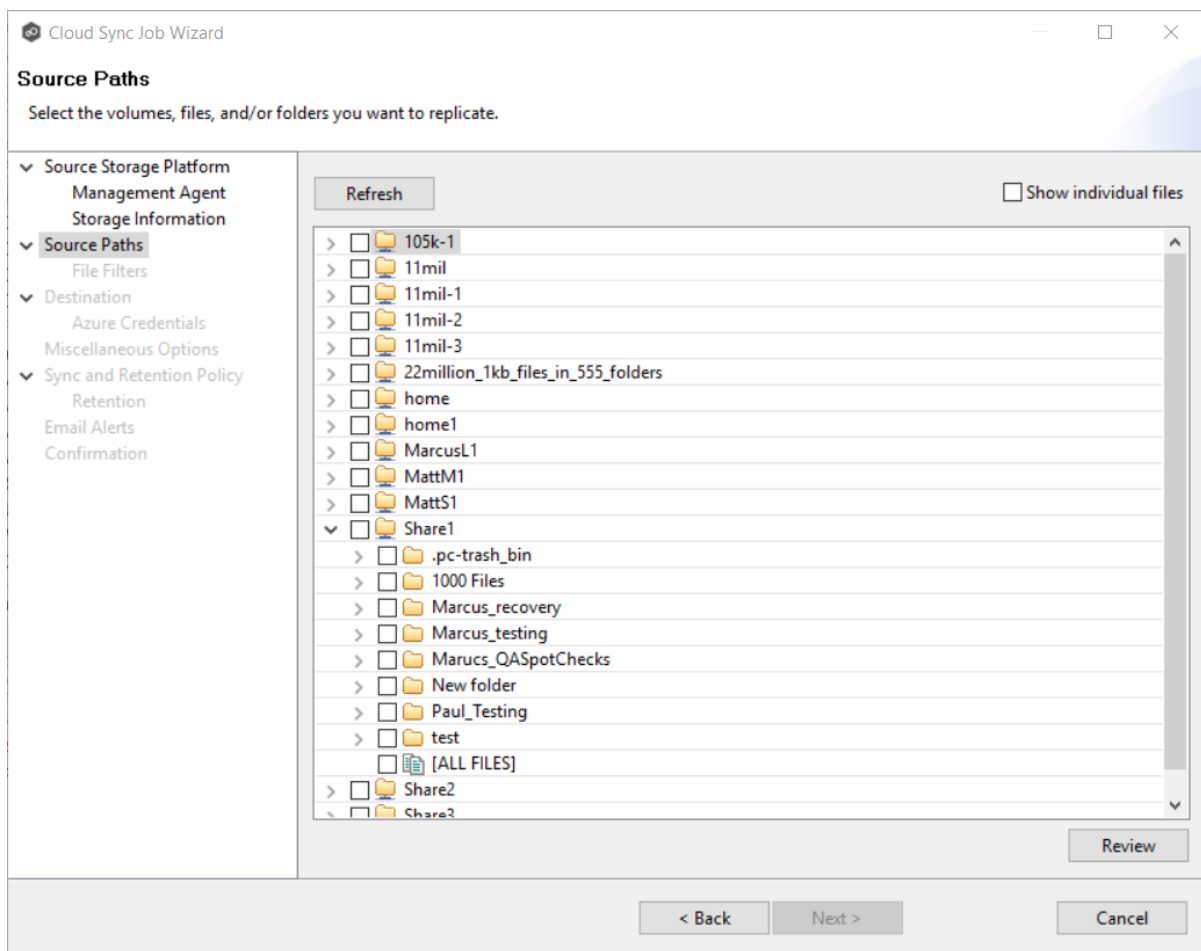
< Back Next > Cancel

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

Step 5: Source Paths

The **Source Paths** page displays a list of available volumes to replicate. You can choose to replicate an entire volume or selectively replicate files and folders. The files and folders selected for replication are referred to as the watch set.

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



To replicate:

The entire volume (all files and folders, including subfolders and their files)

Select the volume check box.

All files at the root level of the volume (but no folders)	Expand the volume, scroll to the bottom of the expanded list, and select All Files .
A specific folder and its content (including subfolders and their files)	Expand the volume, find the desired folder, and select its check box.
All files within a specific folder (but not the folder)	Expand the folder and select All Files .
Specific files and folders	Select the Show individual files check box, expand the folders, and select the files and folders you want to replicate.

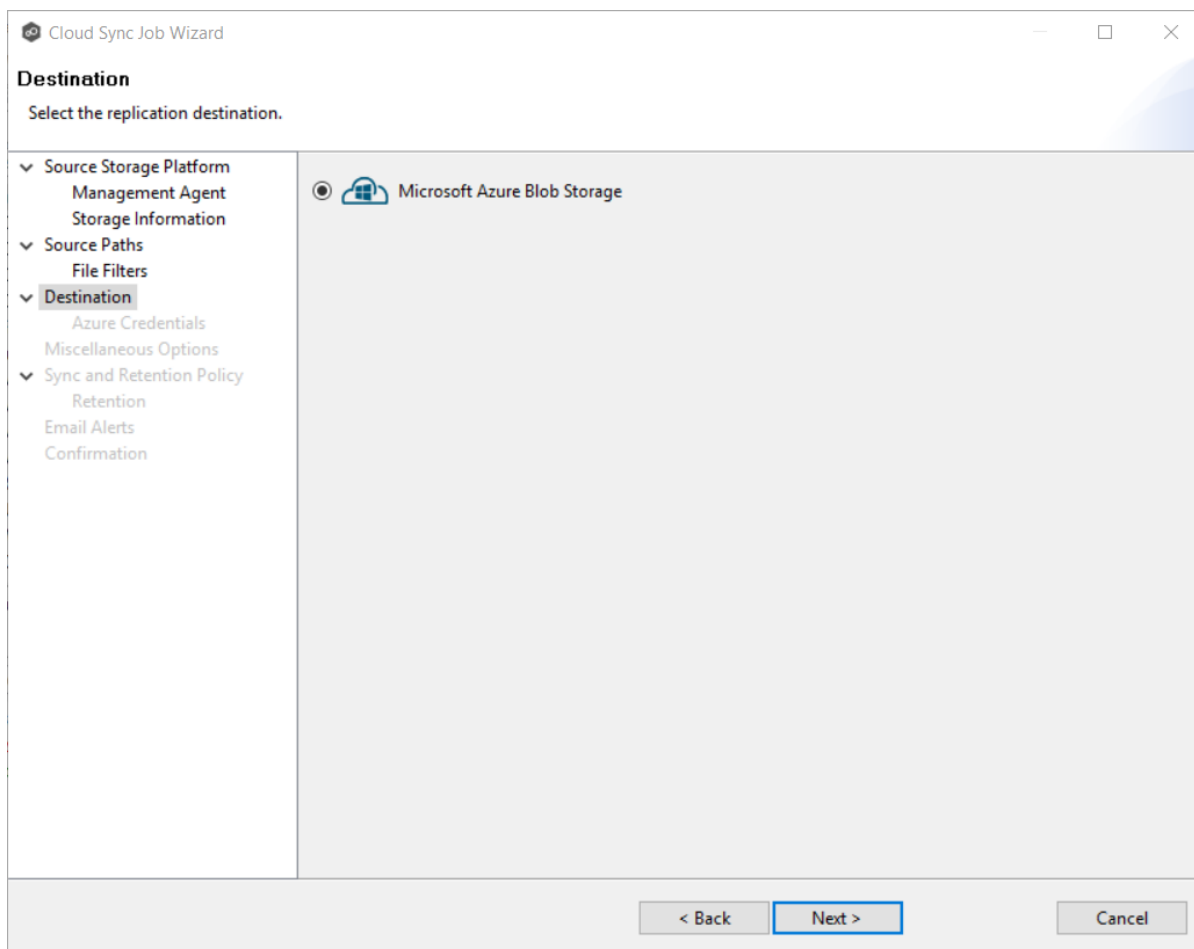
2. (Optional) Click the **Review** button to see your selections.
3. Click **Next**.

Step 6: File Filters

The **File Filters** page displays a list of file filters. A file filter enables you to exclude and/or include files and folders from the job based on file type, extension, name, or directory path. Any file that matches the filter is excluded or included from replication, depending on the filter's definition. By default, all files and folders selected in the **Source Paths** page will be replicated.

See [File Filters](#) for a description of how file filters work.

1. Select the file filters you want to apply to the job. Click **Edit File Filters** if you want to add a new filter or modify an existing one. See [Creating File Filters](#) for details on how to create a file filter.



2. Click **Next**.

Step 8: Cloud Platform Credentials

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

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

Cloud Sync Job Wizard

Azure Credentials

Enter new credentials or select existing credentials.

- Source Storage Platform
 - Management Agent
 - Storage Information
- Source Paths
 - File Filters
- Destination
 - Azure Credentials**
 - Miscellaneous Options
- Sync and Retention Policy
 - Retention
 - Email Alerts
 - Confirmation

Credentials

☐ New Credentials

*Description:

*Account:

*Shared Key: ☐ Show Key

Type: Public

☒ Existing Credentials

peerdebrag1

- If you selected **Existing Credentials**, select a credential from the drop-down list, and then click **Next**. Continue with [Step 8](#).

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

Description	Enter a name for the credentials.
Account	Enter the name of the Azure Storage account, which can be found in the Azure Portal.
Shared Key	Enter one of the shared keys of the Azure Storage account, which can be found in the Azure Portal.
Type	

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

- Click **Next**.

Step 9: Miscellaneous Options

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

- Select the options to apply to this job.

Cloud Sync Job Wizard

Miscellaneous Options
Select options.

- Source Storage Platform
 - Management Agent
 - Storage Information
- Source Paths
 - File Filters
- Destination
 - Azure Credentials
 - Miscellaneous Options**
- Sync and Retention Policy
 - Retention
 - Email Alerts
 - Confirmation

Metadata

NTFS Permissions:

☐ Owner ☐ DACL ☐ SACL

Delta-level Replication

☐ Use delta-level replication

Cloud Tiering Options

Storage Tier/Class: Storage Account Default

Rehydrated Data Availability (Days): 7

< Back Next > Cancel

NTFS Permissions

If you want NTFS permissions metadata included in the replication, select the elements to include:

- **Owner** – The NTFS Creator-Owner who owns the object (which is, by default, whomever created it).

	<ul style="list-style-type: none"> • DACL – A Discretionary Access Control List identifies the users and groups that are assigned or denied access permissions on a file or folder. • SACL - A System Access Control List enables administrators to log attempts to access a secured file or folder. It is used for auditing.
Delta-level Replication	<p>If you want delta-level replication, select this check box.</p> <p>Delta-level replication enables Cloud Sync to transmit only the bytes/blocks of a file that have changed instead of transferring the entire file. This results in much lower network bandwidth utilization, which can be an enormous benefit if you are transferring files over a slow link.</p>
Storage Tier/Class	<p>Select a storage tier. If you do not select a tier, it will default to the tier you configured on your Azure Storage account.</p> <p>Azure Storage offers three storage tiers for blob object storage so that you can store your data most cost-effectively depending on how you use it:</p> <ul style="list-style-type: none"> • Azure Hot Storage Tier is optimized for storing data that is accessed frequently. • Azure Cool Storage Tier is optimized for storing data that is infrequently accessed and stored for at least 30 days. • Azure Archive Storage Tier is optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements (on the order of hours). The archive storage tier is only available at the blob level and not at the storage account level. <p>To read data in archive storage, Cloud Sync must first change the tier of the blob to hot or cool. This process is known as rehydration and can take up to 15 hours to complete.</p> <p>Rehydrated data remains in hot or cool storage for a specified number of days before Cloud Sync automatically returns it to archive storage.</p>
Rehydrated Data Availability (Days)	<p>Rehydrated data is automatically returned to archive storage after a specified period. Enter the number of days for rehydrated data to remain in hot or cool storage before returning to archive storage. The default is seven days.</p>

2. Click **Next**.

Step 10: Sync and Retention Policy

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

Each Cloud Sync job must have a Sync and Replication policy. A Sync and Replication policy specifies:

- How often you want to scan the storage device for replication or if you want to replicate in real-time.
- Whether you want to create snapshots of the replicated data.
- How long you want to retain the snapshots.

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

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

The screenshot shows the 'Cloud Sync Job Wizard' window, specifically the 'Sync and Retention Policy' step. A red error message at the top states: 'Configuration name cannot be empty.' The left sidebar contains a tree view with the following items: 'Source Storage Platform' (expanded), 'Management Agent', 'Storage Information', 'Source Paths' (expanded), 'File Filters', 'Destination' (expanded), 'Azure Credentials', 'Miscellaneous Options', 'Sync and Retention Policy' (selected), 'Sync Schedule', 'Retention', 'Email Alerts', and 'Confirmation'. The main area is titled 'Sync and Retention Policy' and contains two radio buttons: 'New Policy' (selected) and 'Existing Policy'. Below 'New Policy' is a text field labeled '*Name:' and a 'Sync Type' section with 'Snapshots' (selected) and 'Replica' radio buttons. Below 'Existing Policy' is a drop-down list. At the bottom of the window are three buttons: '< Back', 'Next >', and 'Cancel'.

2. If you selected **Existing Policy**, select a policy from the drop-down list, and then click **Next**. Continue with [Step 13](#).

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

3. Select the sync type:
 - Select **Snapshots** if you want to replicate what is on premises to the [destination cloud storage platform](#), and, in addition to replication, you want to keep versions of changed files and take snapshots of the watch set at specific points in time.
 - Select **Replica** if you want to replicate what is on premises to the [destination cloud storage platform](#) and do not want to save versions of changed files or take snapshots.

Step 11: Sync Schedule

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

1. Select the frequency of the replication:
 - [Scheduled Scans](#) – Replication is scheduled on a daily or weekly basis.
 - [Continuous Data Protection](#) – Replication occurs in real-time: whenever a file changes, the change is replicated.

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

1. Select **Scan at Start** if you want a baseline replication to be performed.
2. Select **Daily** or **Weekly**:
 - Select **Daily** if you want replications performed every day. You can schedule one to four scans per day.

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

The screenshot shows the 'Cloud Sync Job Wizard' window, specifically the 'Sync Schedule' step. The left sidebar contains a tree view with the following items: 'Source Storage Platform', 'Management Agent', 'Storage Information', 'Source Paths', 'Exclusions', 'Destination', 'Azure Credentials', 'Miscellaneous Options', 'Sync and Retention' (which is expanded), 'Retention', 'Email Alerts', and 'Confirmation'. The 'Sync Schedule' step is highlighted in the sidebar. The main area of the wizard is titled 'Sync Schedule' and includes the instruction 'Select a sync option.' Below this, there are two radio buttons: 'Scheduled' (selected) and 'Continuous Data Protection'. Under the 'Scheduled' option, there are two sub-options: 'Daily' (selected) and 'Weekly'. Below these, there is a section titled 'Times (up to 4)' with four rows, each containing a time dropdown and a 'Trigger snapshot' checkbox. The times are 00:30, 07:00, 11:00, and 22:30. The checkboxes for 00:30 and 22:30 are checked, while the others are unchecked. Below this section is a checkbox for 'Scan at Start'. At the bottom of the wizard, there are three buttons: '< Back', 'Next >' (highlighted with a blue border), and 'Cancel'.

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

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

The screenshot shows the 'Cloud Sync Job Wizard' window, specifically the 'Sync Schedule' step. A red error message at the top states: 'You have to select at least one day to perform the sync.' The left sidebar contains a tree view with the following items: 'Source Storage Platform', 'Management Agent', 'Storage Information', 'Source Paths', 'File Filters', 'Destination', 'Azure Credentials', 'Miscellaneous Options', 'Sync and Retention Policy' (expanded), 'Sync Schedule' (selected), 'Retention', 'Email Alerts', and 'Confirmation'. The main area is titled 'Scheduled Scans' and has two radio buttons: 'Daily' and 'Weekly' (selected). Below these, there is a 'Day(s):' section with checkboxes for 'Sunday', 'Monday', 'Tuesday', 'Wednesday', 'Thursday', 'Friday', and 'Saturday'. A 'Time:' section has a dropdown menu set to 'None' and a checkbox for 'Trigger snapshot'. At the bottom of this section is a checkbox for 'Scan at Start'. Below the 'Scheduled Scans' section is a radio button for 'Continuous Data Protection'. At the bottom of the window are three buttons: '< Back', 'Next >', and 'Cancel'.

3. Click **Next** and continue with [Step 12](#).

If you selected **Continuous Data Protection** for the replication frequency:

1. Enter a value for **Processing Delay** if you want the replication to occur after a slight delay. A delay is useful to ensure that when a file or folder is created and quickly renamed, only the latest copy of the file or folder is replicated.
2. If you chose **Snapshots** as the sync type, choose when snapshots are taken (you must take at least one snapshot). If you chose **Replica** as the sync type, the **Trigger Snapshot** options will not appear.

3. Click **Next** and continue with [Step 12](#).

Step 12: Retention

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

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

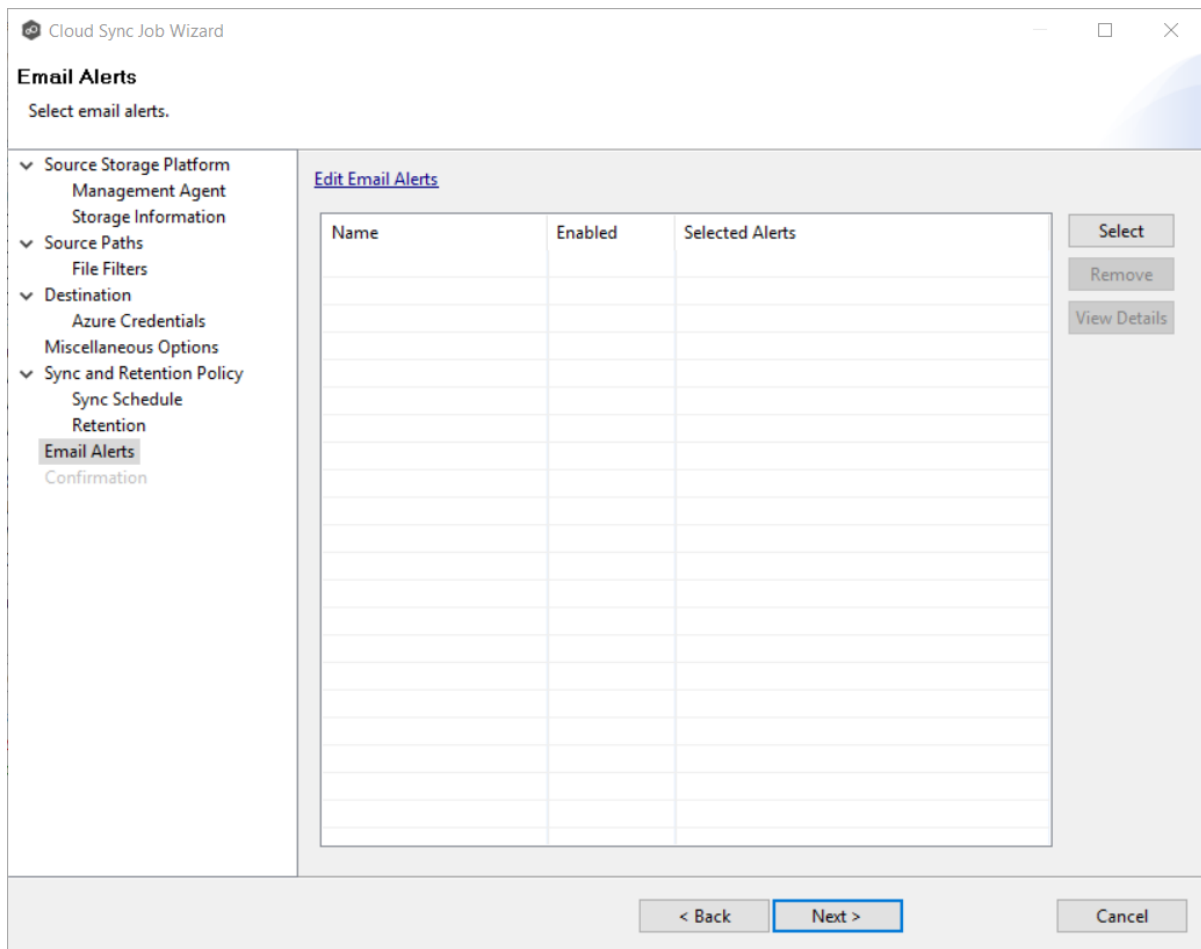
3. Click **Next**.

Step 13: Email Alerts

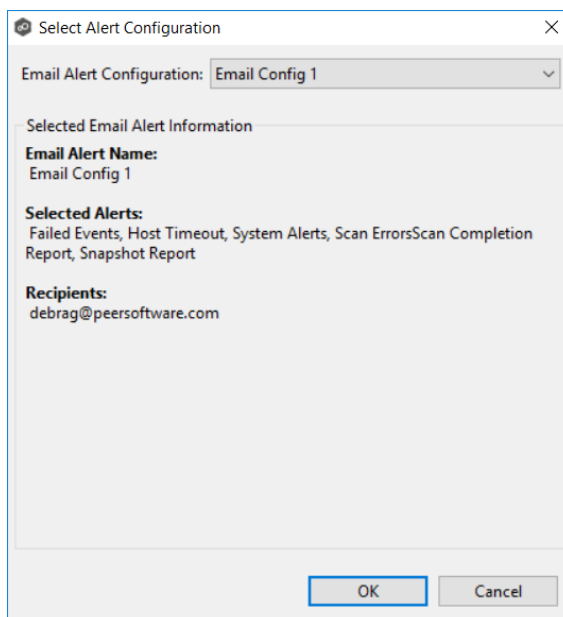
An email alert notifies the recipients when a certain type of event occurs, for example, session abort, host failure, system alert. The **Email Alerts** page displays a list of email alerts that have been applied to the job. When you first create a job, this list will be empty. You can select existing alerts to apply to the job or create new alerts to apply.

See [Configuring Email Alerts](#) for a description of how email alerts work.

1. If you want to apply an alert to the job, click the **Select** button.



The **Select Alert Configuration** dialog appears.



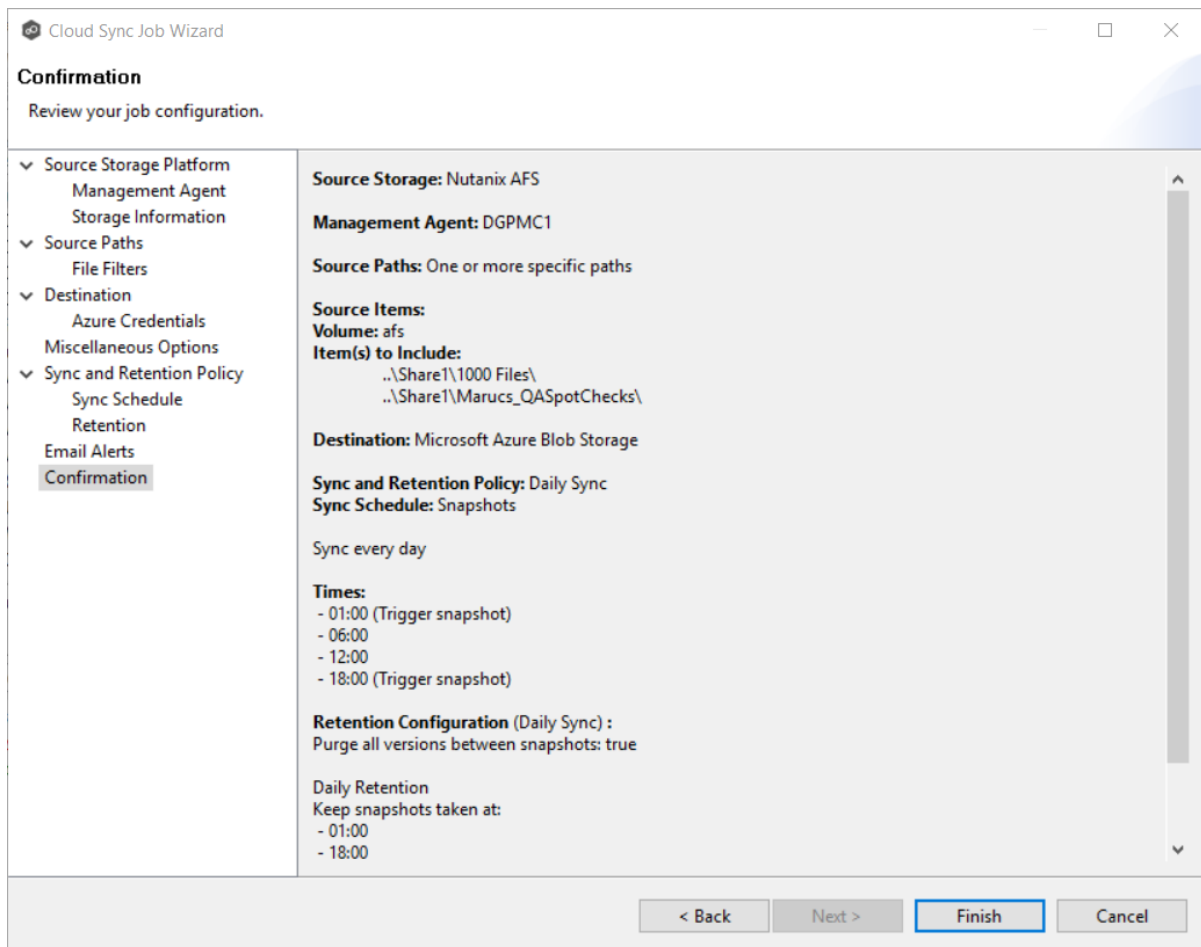
2. From the **Email Alert Configuration** drop-down list, select the email alert to apply to the job, and then click **OK**. If you want to create a new alert or modify an existing one, click the **Edit Email Alerts** link. See [Email Alerts](#) for details on creating an alert.
3. Repeat steps 1-2 if you want to apply additional alerts.
4. Click **Next**.

Step 14: Confirmation

The **Confirmation** page displays the job configuration.

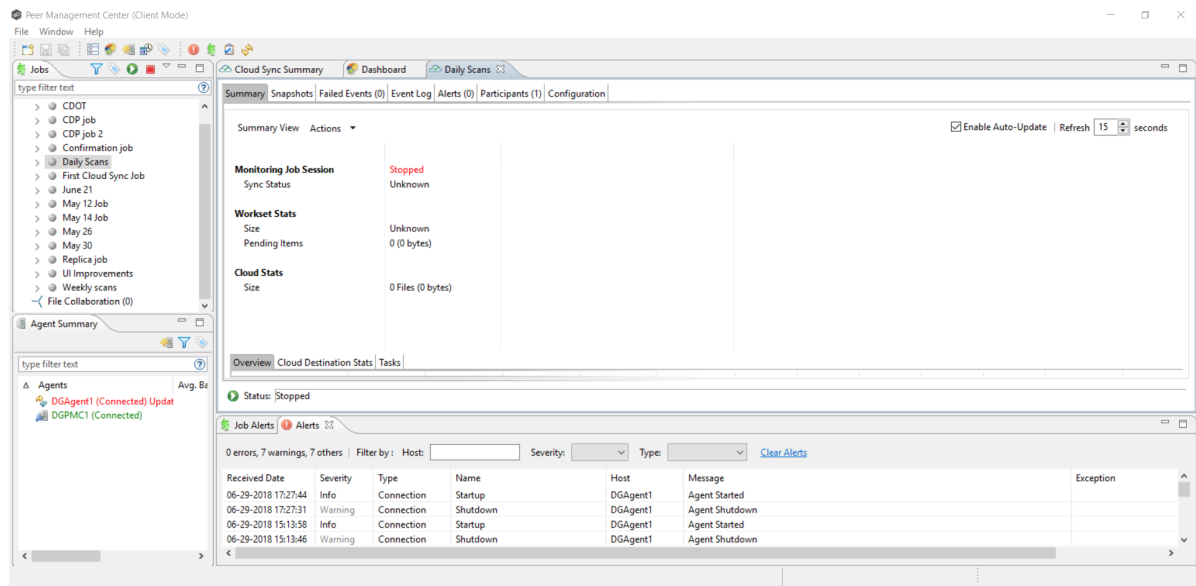
1. Review the job configuration.
2. If you need to modify the job configuration after reviewing it, click **Back** until you reach the appropriate page and make your changes.

Note: You cannot change the job name.



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

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



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

Running a Cloud Sync Job

This topic describes:

- [Starting a Cloud Sync Job](#)
- [Stopping a Cloud Sync Job](#)

Starting a Cloud Sync Job

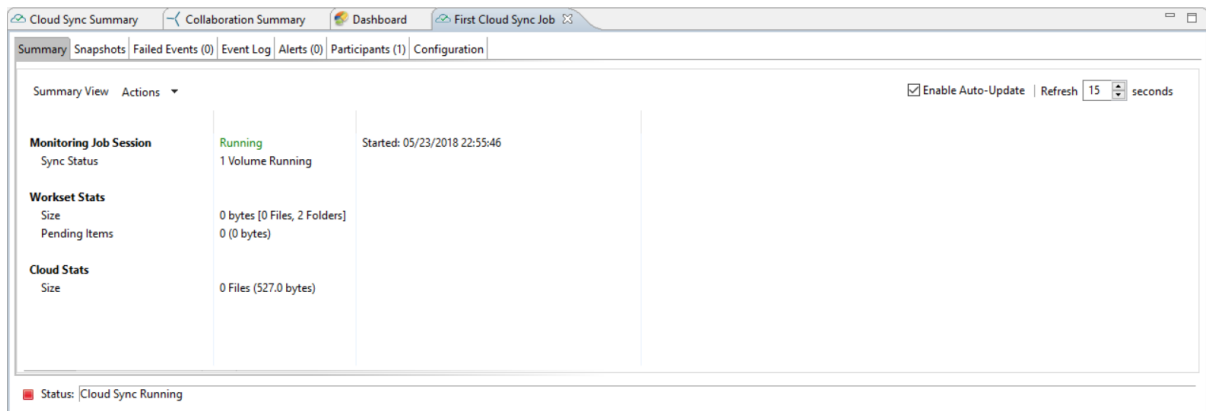
When running a job for the first time, you must manually start it. After the initial run, a job will automatically start, even when the Peer Management Center server is rebooted.

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:

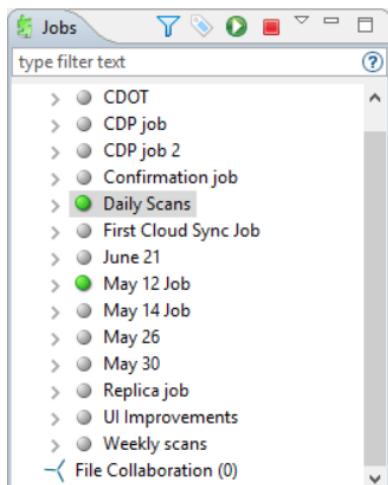
1. Right-click on the job name in the **Jobs** view or in the **Cloud Sync Job Summary** view, and then choose **Start** from the pop-up menu.

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



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

After the job initialization has completed, the job will run. Once the job starts, the icon next to the job name in the **Jobs** view changes from gray to green.



Stopping a Cloud Sync Job

You can stop a Cloud Sync job at any time.

To stop a Cloud Sync job:

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

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


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

The icon next to the job name in the **Jobs** view changes from green to red.

Deleting a Cloud Sync Job


To delete a Cloud Sync job:

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



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.

Monitoring Your Cloud Sync Jobs

Monitoring your Cloud Sync jobs is an important aspect of successfully replicating to the cloud. Monitoring involves checking the execution of a running job, checking the status of a job, reviewing performance statistics, making sure snapshots are created correctly, identifying problems such as a server outage, seeing how much data has been uploaded, and so forth. Cloud Sync provides several views to help you monitor the health and performance of your Cloud Sync jobs.

Many of the views are customizable tables. You can sort the columns in the view, filter by columns, add and subtract columns from the default display, and so forth.

To display a view:

- Double-click **Cloud Sync** in the **Jobs** view to display information about all Cloud Sync jobs. The **Cloud Sync Volume Summary** tab of the **Cloud Sync Summary** view is displayed.

Volume	Job	Management Host	Storage Devi...	Storage Type	Destination Type	Destinatio...	Cloud Syn...	Mode	Scan Status	Next Sc...	Pending Scan I...	Pending Real-t...	Cloud Files	Cloud Bytes
SVM9x1_cifs1	Confirmation job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Stopped	Scheduled ...	Last Scan: I...		0 [0 bytes]	0 [0 bytes]	203	126.88 KB
SVM9x1_cifs1	CDP job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Running	CDP	Last Scan: C...		0 [0 bytes]	0 [0 bytes]	9	60.36 KB
SVM9x1_105k_1	Confirmation job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Stopped	Scheduled ...	Last Scan: I...		0 [0 bytes]	0 [0 bytes]	153	9.59 MB
SVM9x1_105k_1	CDP job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Running	CDP	Last Scan: C...		0 [0 bytes]	0 [0 bytes]	105000	6.43 GB
SVM9x1_105k_1	CDOT	DGPMC1	192.168.171.1...	NETAPP_CDOT	Microsoft Azur...	1921681711...	Running	CDP	Last Scan: C...		0 [0 bytes]	0 [0 bytes]	105000	6.43 GB
afs	Daily Scans	DGPMC1	AFS1	AFS	Microsoft Azur...	afs1-afs-h...	Running	Scheduled ...	Scanning Di...	07/01/2...	4861 [4.75 MB]	0 [0 bytes]	5143	5.96 MB
afs	June 21	DGPMC1	AFS1	AFS	Microsoft Azur...		Stopped							
CA	First Cloud Syn...	DGPMC1	DGPMC1	WINDOWS	Microsoft Azur...	dgpmc1-c-...	Running	CDP	Last Scan: C...		0 [0 bytes]	0 [0 bytes]	0	650.0 bytes
CA	May 12 Job	DGPMC1	DGPMC1	WINDOWS	Microsoft Azur...	dgpmc1-c-...	Running	CDP	Last Scan: C...		0 [0 bytes]	10793 [1.45 TB]	158009	47.54 GB
CA	May 14 Job	DGPMC1	DGPMC1	WINDOWS	Microsoft Azur...		Stopped							
SVM9x1_highbt1_01	Confirmation job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...	svm9x-1-sv...	Stopped	Scheduled ...	Last Scan: I...		0 [0 bytes]	0 [0 bytes]	187	123.06 GB
SVM9x1_cifs1	Replica job	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...		Stopped							
SVM9x1_cifs1	CDP job 2	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...		Stopped							
SVM9x1_cifs1	May 30	DGPMC1	SVM9X-1	NETAPP_CDOT	Microsoft Azur...		Stopped							

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

Summary View		Actions
<div> <input checked="" type="checkbox"/> Enable Auto-Update Refresh 15 seconds </div>		
Monitoring Job Session	Running	Started: 07/01/2018 00:30:47
Sync Status	1 Volume Running	
Workset Stats		
Size	6.41 GB [105000 Files, 10...	
Pending Items	0 (0 bytes)	
Cloud Stats		
Size	105000 Files (6.43 GB)	

Status: Cloud Sync Running

Views for Monitoring General Progress

To monitor the overall health, use the following views:

- **Dashboard** – The **Dashboard** provides a quick status of all Agents.
- **Cloud Sync Summary** – This view is the first place to check to see the status of your Cloud Sync jobs. This view has three tabs:
 - **Cloud Sync Volume Summary** – The view 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.
 - **Cloud Sync Job Summary** – This view displays the status of all Cloud Sync jobs.
 - **Cloud Statistics** – This view displays the total number of files that have been replicated since the first run of the jobs and other statistics.
 - **Task** – This view displays a high-level view of activities such as snapshots, and recovery processes, and background events for all Cloud Sync jobs.

Views for Monitoring Individual Jobs

To monitor a specific job, use the following views:

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

Recovering Data from the Cloud

When you need to recover data from the cloud to on-premises, you can use the **Data Recovery** wizard. To restore data, you must have an existing Cloud Sync job that has been replicating that data.

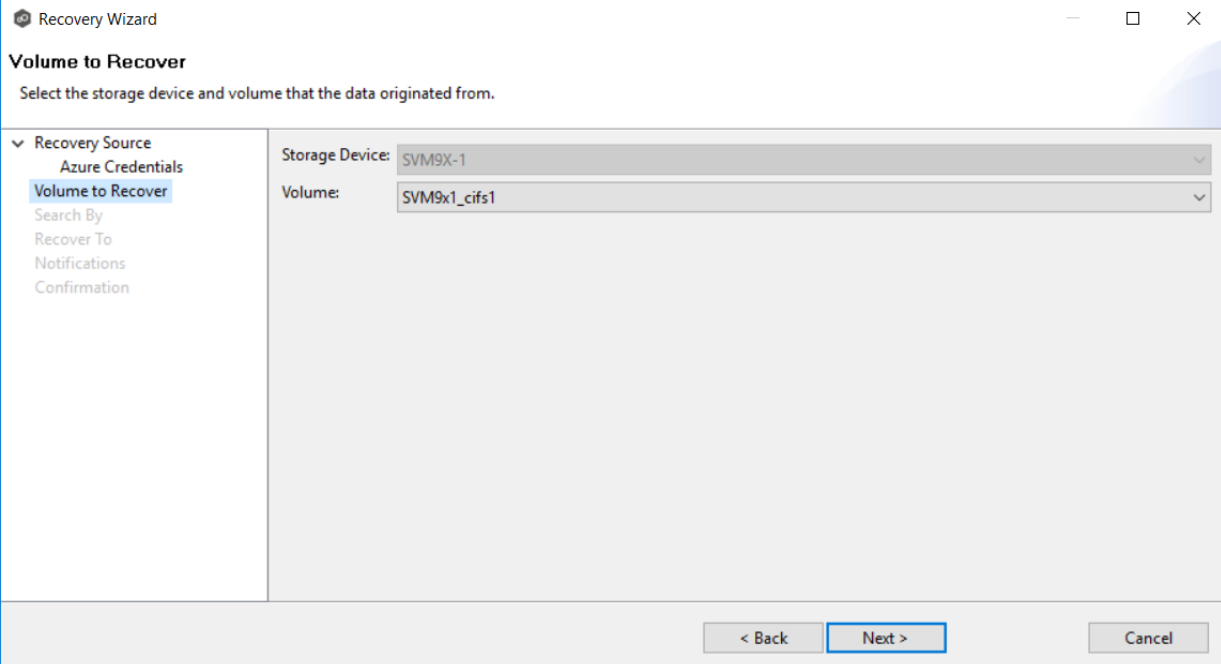
Note: You can recover data from a running job—unless you plan to restore the data to the original location. If so, you should stop the job first.

To recover data:

1. Open the Peer Management Center.
2. In the **Jobs** view, identify the Cloud Sync job that replicated the data you want to restore.
3. Right-click the job name, and then select **Recover Volume/File(s)** from the menu.

The **Recovery Wizard** opens and displays the **Volume to Recover** page. The **Storage Device** field on the page is a read-only field that displays the name of the source storage platform.

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



The screenshot shows a window titled "Recovery Wizard" with a subtitle "Volume to Recover". Below the subtitle is the instruction "Select the storage device and volume that the data originated from." The window is divided into two main sections. On the left is a sidebar with a tree view containing the following items: "Recovery Source", "Azure Credentials", "Volume to Recover" (which is selected and highlighted in blue), "Search By", "Recover To", "Notifications", and "Confirmation". On the right is the main content area. It contains two dropdown menus: "Storage Device:" with the value "SVM9X-1" and "Volume:" with the value "SVM9x1_cifs1". At the bottom of the window is a footer bar with three buttons: "< Back", "Next >" (which is highlighted in blue), and "Cancel".

5. Click **Next**.

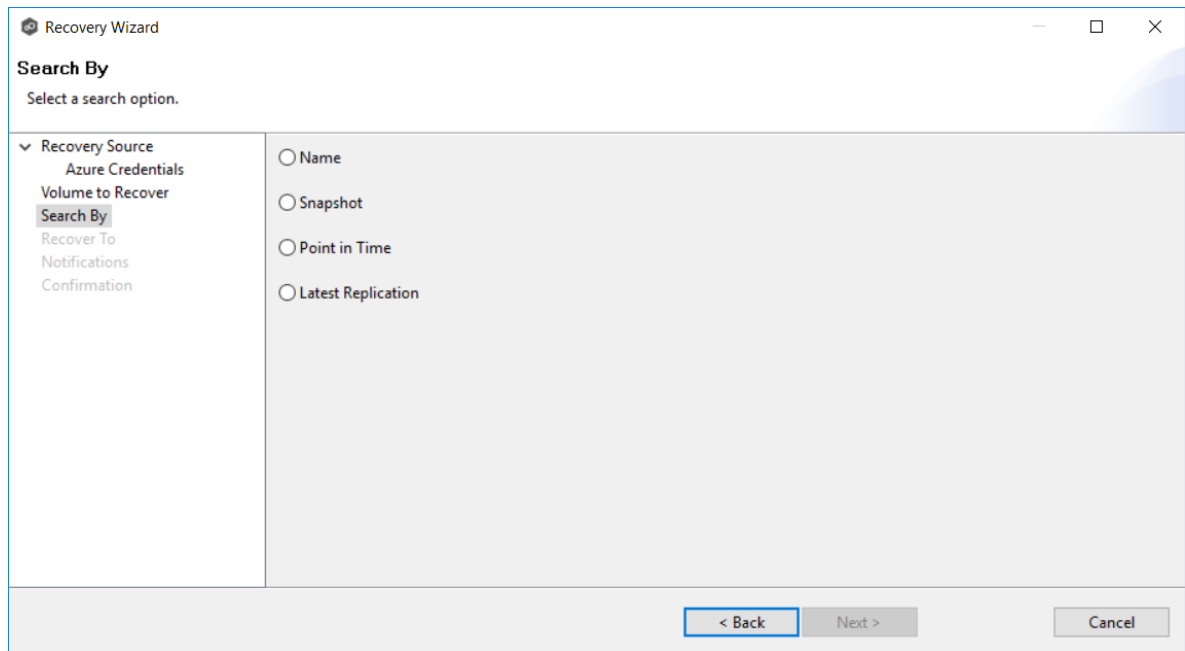
The **Search By** page is displayed. It presents four search options:

[Name](#)

[Snapshot](#)

[Point in Time](#)

[Latest Replication](#)



6. Select one of the search options.

7. Click **Next**.

The search pages vary, depending on the [search option](#) you selected.

Search Options

The four search options are:

- [Name](#)
- [Snapshot](#)
- [Point in Time](#)

- Latest Replication

Use the **Search by Name** option if you know any part of the name of a file or folder but don't know which folder contained it on the original volume on premises.

To search by name:

1. Enter a search string in the **Name** field.

The search string can be a full or partial name and can include wildcards. If you do not enter a search string, all files and folders will be listed in the search results.

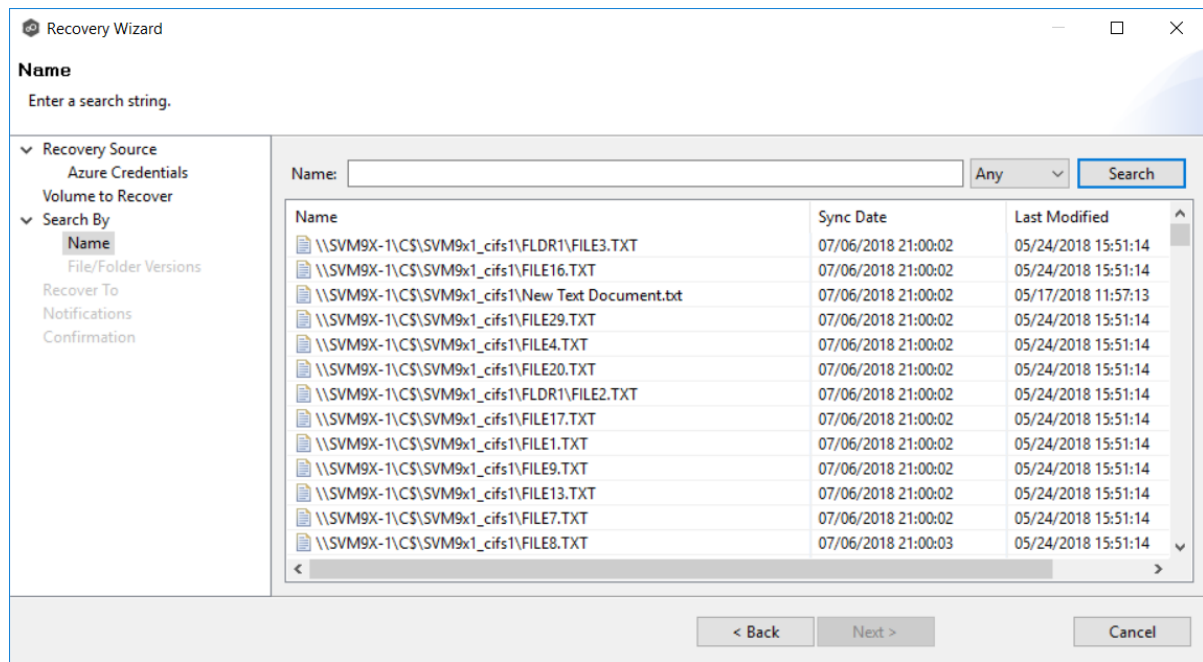
The screenshot shows the "Recovery Wizard" application window. The title bar includes standard Windows window controls (minimize, maximize, close). Below the title bar, there's a section titled "Name" with the instruction "Enter a search string.". To the left of the main content area is a sidebar menu with several expandable sections:

- "Recovery Source" expanded, containing:
 - Azure Credentials
 - Volume to Recover
- "Search By" expanded, containing:
 - Name (highlighted)
 - File/Folder Versions
 - Recover To
 - Notifications
 - Confirmation

The main content area features a search interface at the top with a label "Name:", a text input field, a dropdown menu currently set to "Any", and a "Search" button. Below this is a large table with three columns: "Name", "Sync Date", and "Last Modified". The table has multiple rows, all of which are currently empty. At the bottom right of the table, there are navigation arrows "<" and ">". At the very bottom of the window, there are three buttons: "< Back" (active/highlighted), "Next >", and "Cancel".

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

A list of matching files and/or folders appears. The **Sync Date** column shows the date the file was replicated; the **Last Modified Date** column shows the last known date and time that the file was changed on premises.



4. Select the file or folder to recover.

5. Click **Next**.

The **File/Folder Versions** page appears. Your options will vary, depending on whether you are recovering a file or folder.

6. If you selected a file to recover, all available versions of that file are presented below the calendar. Select the time of the desired version and then click elsewhere in the page.

Recovery Wizard

File/Folder Versions

Select a version to recover.

Recovery Source

- Azure Credentials
- Volume to Recover

Search By

- Name
- File/Folder Versions**
- Recover To
- Notifications
- Confirmation

Selected Sync Item(s):

\\SVM9X-1\CS\SVM9x1_cifs1\FLDR1\FILE2.TXT

Calendar: Jul, 2018

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

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

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

Select Time

< Back Next > 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 in the page.

Recovery Wizard

File/Folder Versions

Select a version to recover.

Recovery Source

- Azure Credentials
- Volume to Recover

Search By

- Name
- File/Folder Versions**
- Recover To
- Notifications
- Confirmation

Selected Sync Item(s):

\\SVM9X-1\CS\SVM9x1_cifs1\22mil\FLDR2L1

☒ Snapshots ☐ Date and Time

Calendar: Jul, 2018

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

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

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

Select Time

< Back Next > Cancel

7. Click **Next** and continue with [Recovery Options](#).

Use the **Search by Snapshot** option if you want to recover data by browsing a previously taken snapshot. All available snapshots will be represented in the calendar widget below.

To search by snapshot:

1. Select the date of the snapshot.

Recovery Wizard

Snapshot
Select a snapshot.

Recovery Source
Azure Credentials
Volume to Recover

Search By
Snapshot
File/Folder Browser

Recover To
Notifications
Confirmation

Available snapshots:

<< < Jul, 2018 > >>

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

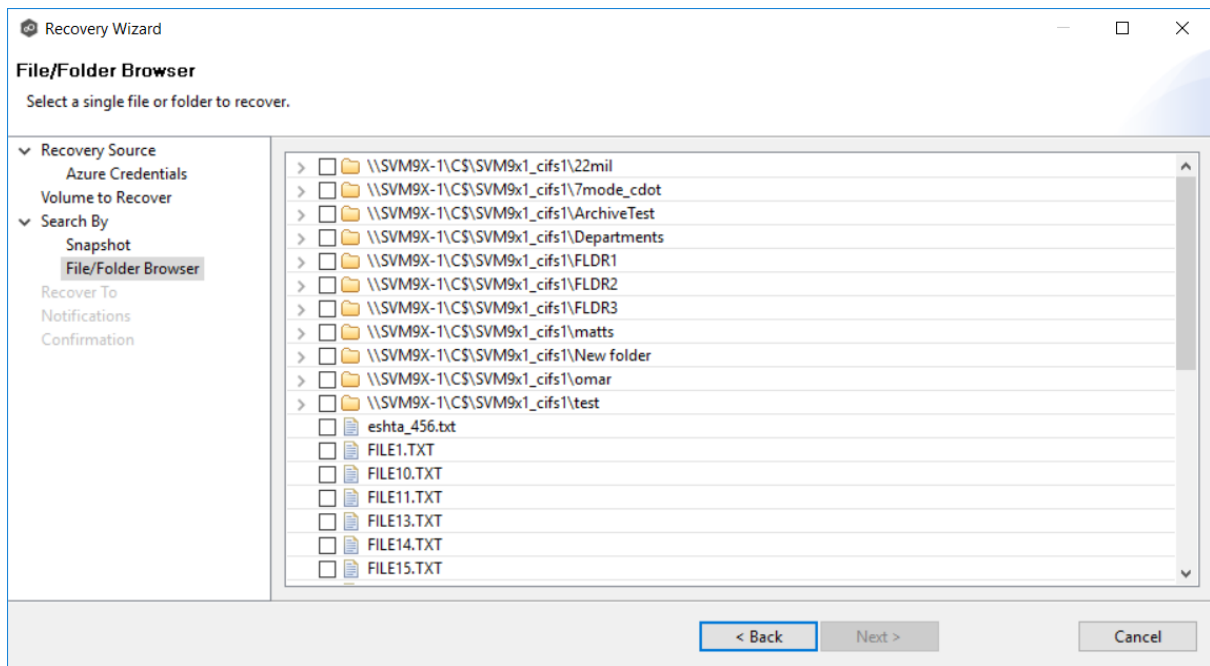
Select Time

Oldest Available: 07/07/2018 00:30:01
Newest Available: 07/10/2018 07:30:01

< Back Next > Cancel

2. Select the time of the snapshot, and then click elsewhere in the page.
3. Click **Next**.

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

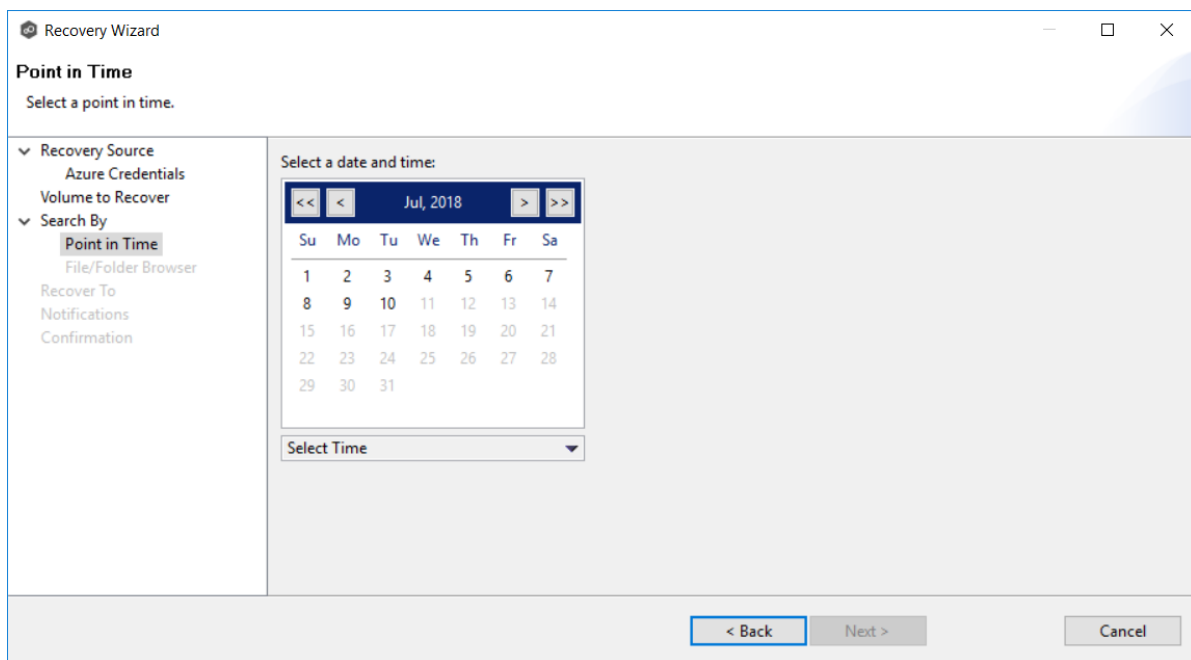


4. Select the file or folder to restore. If no snapshots are available, click **Back** and select a different search option.
5. Click **Next** and continue with [Recovery Options](#).

Use the **Search by Point in Time** option if you want to restore a data from a specific point in time. This option does not require that a snapshot was taken and is very useful if you selected [Continuous Data Protection](#), where replication is performed on an on-going basis

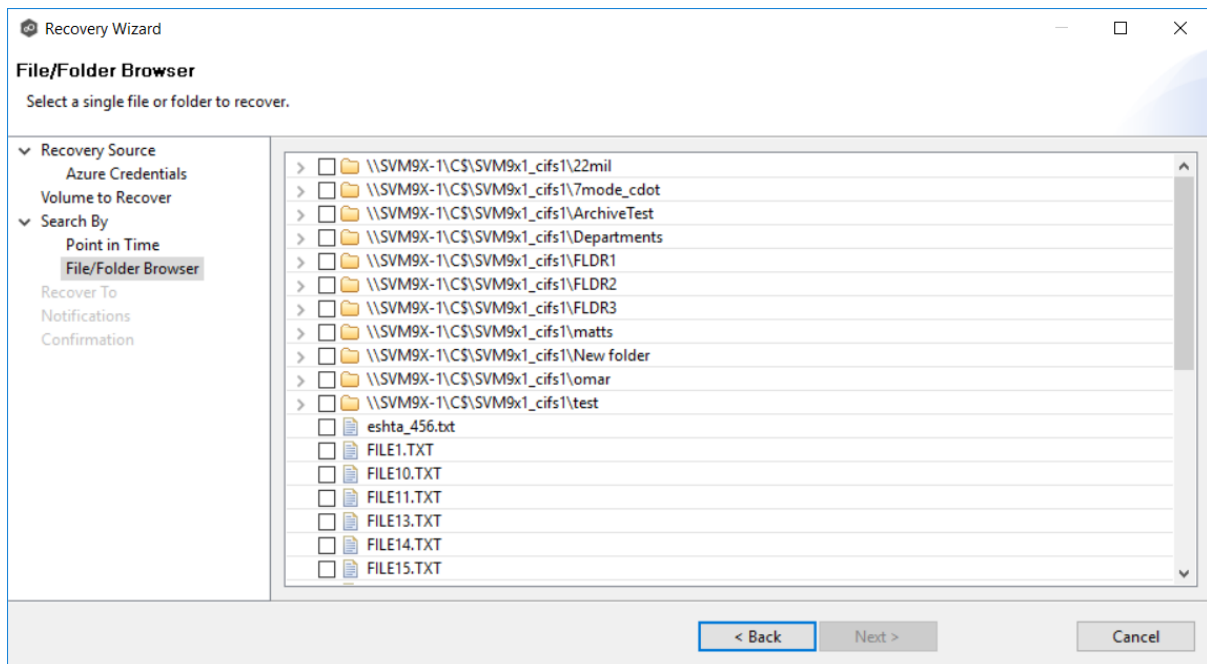
To search by a point in time:

1. Select a date.



2. Select a date and time, and then click elsewhere in the page.
3. Click **Next**.

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




4. Select the file or folder to restore.

5. Click **Next** and continue with [Recovery Options](#).

Use the **Search by Latest Replication** option if you want to restore from the latest replication. For example, you may want to restore data from the last time that replication occurred rather than a snapshot or a point in time. This option is very useful if you selected [Continuous Data Protection](#), where replication is performed on an on-going basis.

To search by latest replication:

1. Select the file or folder to restore.



2. Click **Next** and continue with [Recovery Options](#).

Recovery Options

After you find the data to recover, the **Recover To** page appears.

1. Select the recovery location. You have two options:
 - **Another Location** - Enter the UNC path to a location on another storage device.

- **Original Location** - Browse to a location on the device hosting the management agent. However, we recommend not restoring directly to the original location, especially if the job is currently running.

2. Select the recovery options for if the file to recover already exists in the recovery location:

Recovery Option	Select this option if you want to:
Recover with unique name	Ensure that the existing file is not overwritten with the cloud version.
Overwrite if sizes or timestamps don't match	Overwrite the existing file with the cloud version if the sizes or timestamps the existing file do not match the cloud version.
Overwrite if cloud version is newer	Overwrite the existing file if the cloud version has a more recent modification date.
Overwrite always	Always overwrite the existing file with the cloud version.
Skip	Skip recovering a file if the file already exists.

3. Select the recovery metadata options:

Metadata Option	Select this option if you want to:
Recover Last Modified Time	Set the last modification time of a recovered file to match the last modification time stored at upload rather than the time at which it was recovered.
Recover Create Time	Set the creation time of a recovered file to match the creation time stored at upload rather than the time at which it was recovered.
Recover NTFS Permissions	Set the NTFS permissions of any recovered files and folders to match the original permissions when those files and folders were uploaded.
Recover Attributes	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.

The screenshot shows the 'Recovery Wizard' window with the 'Notifications' step selected in the left sidebar. The main area is titled 'Notifications' with the instruction 'Select email notifications.' Below this, there's a section for 'Edit Email Alerts' which includes a checkbox for 'Send email notification when complete' and a sub-checkbox for 'Only on failure'. A text input field is provided for 'Enter recipient name, email or distribution list to add:' with a placeholder 'Start typing to filter contacts/lists or add a new email'. To the right of this field are buttons for 'Add to List', 'Remove', and 'View Details'. Below the input field is a large empty box labeled 'Recipients:'. At the bottom of the window are navigation buttons: '< Back', 'Next >' (highlighted), and 'Cancel'.

6. (Optional) Select the **Send email notification when complete** check box if you want notifications sent when the recovery process is complete. Select **Only on failure** if you want notifications sent only if the recovery does not successfully complete.
7. Enter recipients and add them to the list.
8. Click **Next**.

The **Confirmation** page is displayed.

9. Review your recovery settings.



10. Click **Finish**.

File Synchronization

This section provides information about creating, editing, running, and managing a File Synchronization job:

- [Creating a File Synchronization Job](#)
- [Running and Managing a File Synchronization Job](#)

Creating a File Synchronization Job

The topics in this section provide some basic information about creating and editing File Synchronization jobs.

Integrating Existing File Synchronization Instances

To integrate existing File Synchronization instances in the Peer Management Center, follow the [step-by-step](#) instructions.

Creating and Deploying New File Synchronization Instances

To create a new job and deploy the PeerSync installation to one or more hosts, click the **Create New** button in toolbar of the Peer Management Center, or you can select the **New** menu item from the **File** menu. A list of all installed Peerlet types will be displayed. Selecting the **File Synchronization** option will open the File Synchronization Configuration dialog. Go to the [Step-by-Step](#) instructions for more information.

When configuring Alerts you will want to configure global settings like SMTP configuration, which is specific to the Peer Management Center. Details on what and how to configure these global options can be found in the [File Synchronization Configuration](#) section.

You can edit an existing job's Alert and logging by selecting one or more jobs in the Jobs view, right-clicking, and selecting **Edit Jobs(s)**. The Peer Management Center now has support for editing multiple jobs at once.

To edit the File Synchronization PeerSync configuration right-click on the job in the Jobs view and select **Edit Configuration**. From the File Synchronization Configuration screen select the Associated Profile Node from the left. For step-by-step, go to [Running and Managing File Synchronization Jobs](#).

- [Integrating Existing PeerSync Instances](#)
- [Deploying New PeerSync Instances](#)
- [File Synchronization Configuration](#)

File Synchronization Configuration

Before configuring the individual aspects of a File Synchronization Session, we first recommend preconfiguring a number of global options that can be applied towards all File Synchronization sessions.

The following configuration items are not always required, but highly recommended:

- [SMTP Email Configuration](#)
- [Email Alerts](#)

Before the Peer Management Center can send emails on behalf of any file synchronization job, a few key SMTP settings must be configured. To set these values, click the **Window** menu from within the Peer Management Center, and select **Preferences**. Within the dialog that pops up, select **SMTP Email Configuration** from the navigation tree. The following screen will be displayed.

The screenshot shows the 'Preferences' dialog box with the 'SMTP Email Configuration' tab selected. The left sidebar contains a search bar labeled 'type filter text' and a list of configuration categories: 'File Synchronization', 'General Configuration', 'Licensing', 'SMTP Email Configuration' (highlighted), and 'User Management'. The main area displays the 'SMTP Email Configuration' settings, which include fields for 'SMTP Host', 'SMTP Port' (set to 25), 'Encryption' (unchecked), 'Encryption Type' (set to TLS), 'Username', 'Password', and 'Sender Email'. A 'Test Email Settings' button is located below these fields. At the bottom right of the dialog are 'OK' and 'Cancel' buttons.

Field	Value
SMTP Host	
SMTP Port	25
Encryption	<input type="checkbox"/>
Encryption Type	TLS
Username	
Password	
Sender Email	

SMTP Host (required)	The host name or IP address of the SMTP mail server through which the Peer Management Center will send emails.
SMTP Port	TCP/IP connection port (default is 25 and 465 for encryption) on which the mail server is hosting the SMTP service. It is recommended that you leave the default setting unless your email provider specifies otherwise.
Encryption	Check this box if the SMTP mail server requires an encrypted connection.

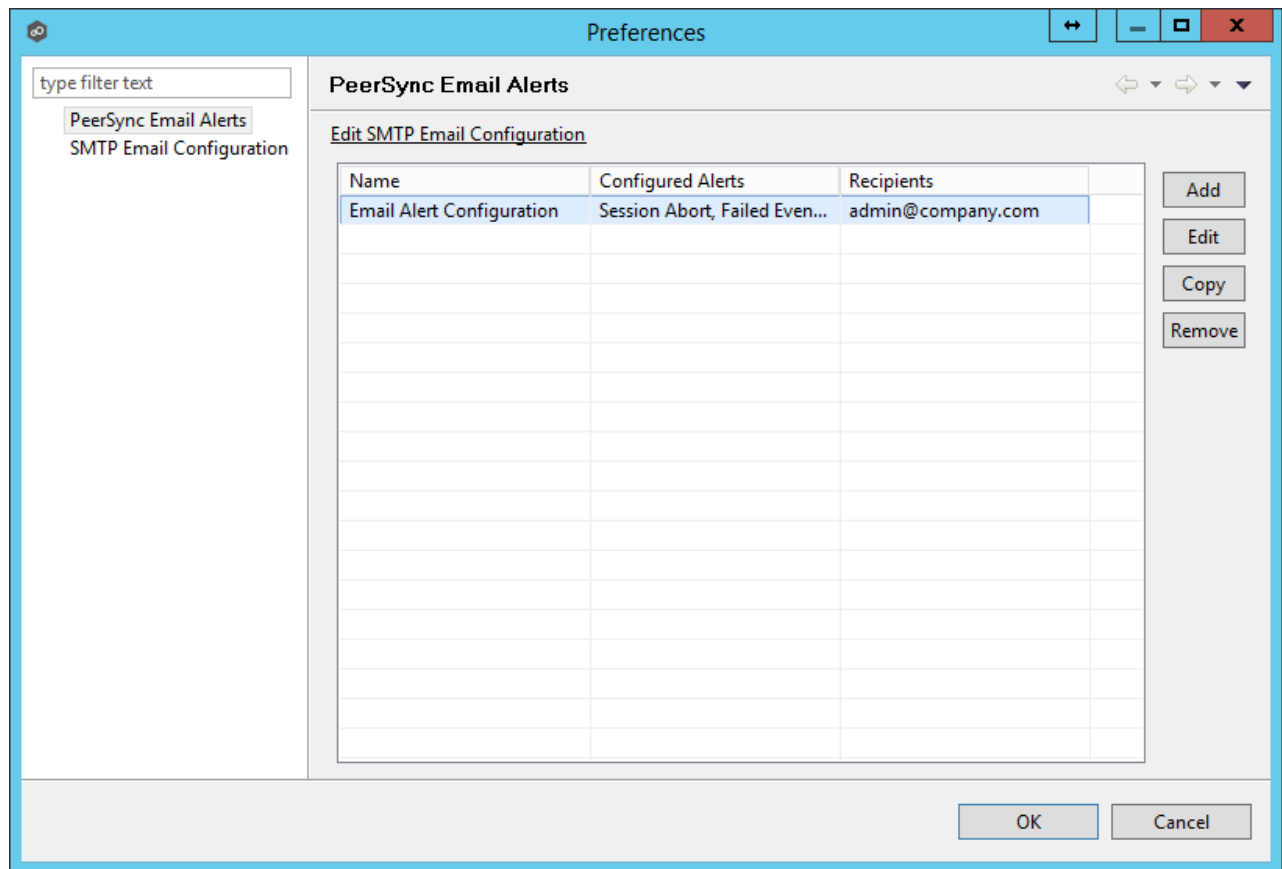
Encryption Type	If encryption is enabled, an encryption method must be selected. TLS and SSL are the available options. If you do not know which one your mail server requires, try one then the other.
User	The username to authenticate as on the SMTP mail server (optional).
Password	The password of the username specified above (optional).
Sender Email (required)	The email address that will appear in the From field of any sent emails. This email address sometimes needs to have a valid account on the SMTP mail server.

It is highly recommended that you test your SMTP settings before saving them. To do so, click the **Test Email Settings** button. You will be prompted for an email address to send the test message to. Upon submission, the Peer Management Center will attempt to send a test message using the specified settings.

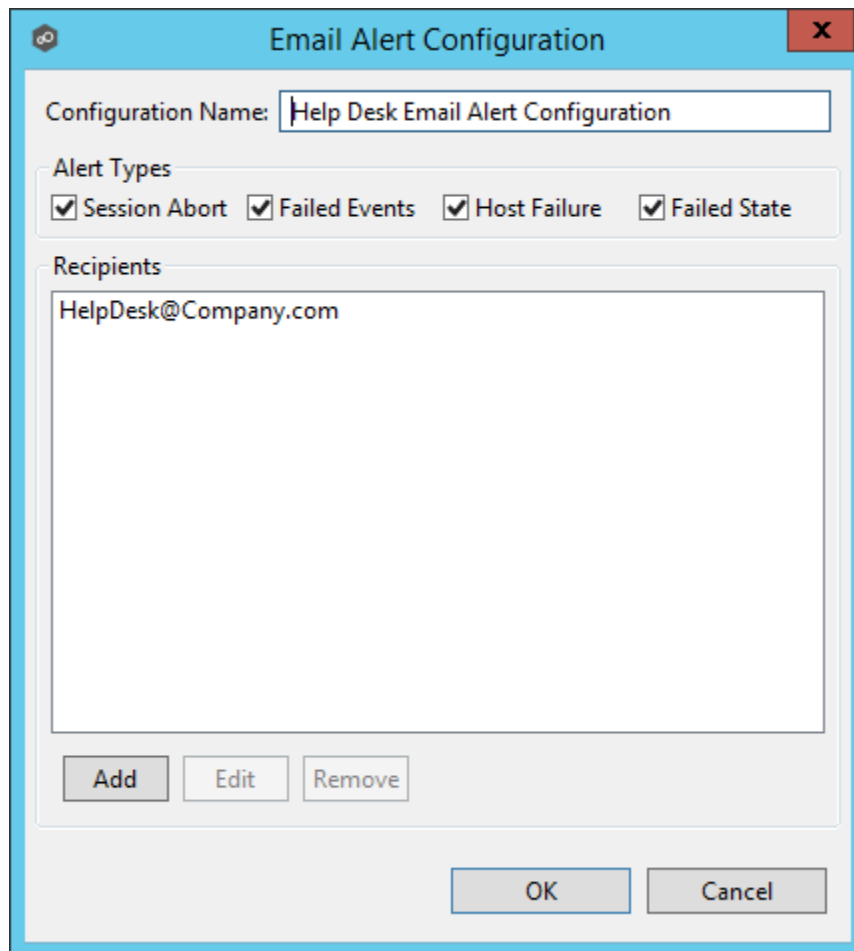
Overview

The Peer Management Center supports the concept of email alerts, where a single alert (consisting of a unique name, a selection of event types along with a list of email addresses) can be applied to multiple file synchronization jobs without requiring repeat entry for each job. When an email alert is applied to a job, an email is sent to all listed recipients anytime a selected event type is triggered by that job.

To manage email alerts, right-click any file synchronization job from the Jobs view and select the **Email Alerts** node from the **Monitoring** node. Click **Edit PeerSync Email Alerts**. The following screen represents the list of defined email alerts, along with buttons to add new ones and edit, copy and remove existing ones.



Upon adding or editing an email alert, the following dialog is displayed:



The image shows a screenshot of the 'Email Alert Configuration' dialog box. The title bar is blue with a close button (X) on the right. Inside the dialog, there is a text field for 'Configuration Name' containing 'Help Desk Email Alert Configuration'. Below this is a section titled 'Alert Types' with four checkboxes: 'Session Abort', 'Failed Events', 'Host Failure', and 'Failed State', all of which are checked. Underneath is a 'Recipients' section with a text area containing 'HelpDesk@Company.com'. At the bottom of the recipients area are three buttons: 'Add', 'Edit', and 'Remove'. At the very bottom of the dialog are 'OK' and 'Cancel' buttons.

Within this dialog, you can select specific event triggers on which an email will be generated and configure the list of email recipients of the alert(s). Event types are defined below.

Event Types

Session Abort	Enables sending an alert when a session is aborted because of lack of quorum due to one or more failed host agents.
Failed Events	Enables sending an alert when a failed event is received from the PeerSync machine.
Host Failure	Enables sending an alert when a host agent timeout occurs or a PeerSync service timeout occurs.
Failed State	Enables sending an alert when the state of the File Synchronization machine changes from Active to "Failed State" indicating that either a

	failed scan or failed event was detected in the latest set of synchronization stats.
--	--

Integrating Existing PeerSync Instances

The topics in this section provide some basic information on how to integrate existing PeerSync instances within the Peer Management Center.

- [Requirements](#)
 - [How to Integrate Existing PeerSync Instances](#)
-
- PeerSync has to be installed as a Service and running version 9.3.0 or newer.
 - Peer Agent has to be installed on the PeerSync machine and connected to the Peer Management Center.

Modify the PeerSync Profile

1. Open the profile on the PeerSync machine with the PeerSync Profiler.
2. Add the argument /LZTAI in Options/Command section.
3. Save the profile.
4. Restart the PeerSync Service.
5. Install the Peer Agent.
6. Start the Peer Agent.

Once the Peer Agent is started and connected to the Peer Management Center, PeerSync will be auto detected and a Peer Management Center file synchronization job will be generated with the name of the machine.

Optionally you can edit the job and add [email alerts](#) and save and restart the File Synchronization job for changes to take effect.

Deploying New PeerSync Instances

The topics in this section provide basic information on how to integrate existing PeerSync instances within the Peer Management Center:

- [Requirements](#)
- [How To](#)

- Peer Agent has to be installed on the machine where PeerSync will be deployed to.
- It is recommended to run the Agent under a domain admin account or account with enough rights to modify registry and service configuration.

The topics in this section provide step-by-step instructions on how to create and deploy new File Synchronization instances of PeerSync software.

- [Step 1: General Information](#)
- [Step 2: PeerSync Profile](#)
- [Step 3: Jobs Configuration List](#)
- [Step 4: Installation Settings](#)

Step 1: General Information

Create a new file synchronization job by clicking the **Create New** button in the toolbar of the Peer Management Center, or by selecting the **New** menu item from the **File** menu. A drop down list of all installed Peerlet types will be displayed. Selecting the **File Synchronization** option will open the **File Synchronization Configuration** dialog.

The first page of configuration will be for general information such as Host Participants and Job name tag.

File Synchronization Configuration

1 of 4 - General Information

Generic information on this PeerSync configuration

Name: [COMPUTERNAME] - EMEA_Region

Available

Host	Computer Description
Win12x64a	

Add Remove

Selected

Host	Computer Description

< Back Next > Finish Cancel

1. The job name will default to the computer name of the host participant. If you wish to group your computers, you can optionally add a name tag in the text box next to the job name (e.g., East Coast, EMEA, Region2). This will help in filtering machines by their given tag.

2. A list of all available hosts that have not yet been configured with a PeerSync installation, will appear in the **Available** table on the top of the page. Available hosts are any host with a Peer Agent installed that has successfully connected to the configured Peer Management Center Broker. The name that will be displayed is the computer name of the server that the Peer Agent is running on. If a particular host is not displayed in the list, then try restarting the Peer Agent Windows Service on that host, and if it successfully connects to the Peer Management Center Broker, then the list will be updated with the computer name of that host.

Note: Computer Description is defined through Windows on a per-computer basis.

3. Select one or more hosts from the **Available** table and click the **Add** button to add the hosts to the **Selected** table. These are the hosts you wish to deploy the PeerSync configuration and installation to.

Step 2: PeerSync Profile

In the second page, choose a preconfigured profile from the available templates, or browse to load a PeerSync profile you may have configured through the PeerSync Profiler and saved as a .snc file on this system.

You may also choose to start from scratch by choosing **Other** from the drop down menu.

Enter or update the **Profile Description** and **Performance Options**.

File Synchronization Configuration

2 of 4 - PeerSync profile

Profile Configuration

Pre-Defined Profile

Backup.snc

Migration.snc

Other

Browse

Profile Description

Performance Options

Maximum number of Job Threads

5

Maximum number of Copy Threads

10

< Back

Next >

Finish

Cancel

Profile Description	A textual description of the current profile.
Maximum number of Job Threads	Maximum number of job scans that can run parallel to one another.
Maximum	Maximum number of events that can be processed parallel to one another.

Profile Description	A textual description of the current profile.
number of Copy Threads	

Step 3: Jobs Configuration List

In this page, modify the loaded PeerSync jobs and/or add new jobs by clicking on the **Add** or **Edit** button to the right of the view.

File Synchronization Configuration

4 of 4 - Installation Settings

Remote Installation Settings

Pre-Defined Settings: EastCoast Name: EMEA

Installation Path: %ProgramFiles%\Peer Software\PeerSync

Existing Exe: peersync93-9_3_1_1007.exe Browse

Imported Exe Version: v9.3.1.1007 [C:\Program Files (x86)\Peer Software\Peer Management Hub\Hub\workspace\peersync\templates\exes\peersync93-9_3_1_1007.exe]

License

User Name: Peer Software Company: Peer Software

Options: Password:

Service

Logon User: bytemetrics\adminandanielad

Password:

< Back
Next >
Finish
Cancel

Pre-Defined Settings	<p>A list of previously used Installation Settings with the Name given at the time of use.</p> <p>Note: If the installation settings have the same path, service user name, license password and installation exe, a new Installation record will not be created, regardless if a new name has been given to the Installation Settings.</p>
Installation Path	<p>Path where PeerSync will be installed. When using the % ProgramFiles% variable, PeerSync will install in the x86 Program Files directory for 64 bit systems, otherwise it will install in the Program Files base directory.</p>

Pre-Defined Settings	<p>A list of previously used Installation Settings with the Name given at the time of use.</p> <p>Note: If the installation settings have the same path, service user name, license password and installation exe, a new Installation record will not be created, regardless if a new name has been given to the Installation Settings.</p>
Existing Exe	A list of PeerSync executables available in the template folder or used in a past installation. This is the PeerSync executable that will be used to install PeerSync.
License User Name	License information provided by Peer Software. Cut and paste the User Name section in this field.
License Company	License information provided by Peer Software. Cut and paste the Company section in this field.
License Options	License information provided by Peer Software. Cut and paste the Options section in this field.
License Password	License information provided by Peer Software. Cut and paste the Password section in this field.
Service Logon User*	<p>This is the Service account User Id used to run the PeerSync Service (DOMAIN\USER).</p> <p>Note: This account has to be valid on all included participants for this File Synchronization Configuration.</p>
Service Password	PeerSync Windows Service account Password.

*Note: When using a service account that has not been granted to run as a service on the machine, PeerSync will fail to start return the following Global Alert to the Peer Management Center. This will indicate that PeerSync could not start and you will have to log on to that machine and confirm the credentials to grant access to that account to run as a service.

Hub Alert Details

Received at: 09-30-2015 13:08:52

Severity: Warning

Category: Global Resource

Host Name: Peer Management Center

Locally Generated at: 09-30-2015 13:08:52

Name: Process PeerSyncEvent

Message: pe..... Last Event=PeerSyncEvent [host=Win12x64a, eventType=SERVICE_CMD, description=The service did not start due to a logon failure., exception=null, errorCode=0, coordinationId=null, eventId=66, properties = {}] : The service did not start due to a logon failure.

Click outside of popup to close

Once the Configuration Settings have completed, click Finish and the installation configuration will be sent to the selected Participants.

A File Synchronization job will be auto created for each Participant and set to be in a Pending Installation state. Once the installation completes and PeerSync reports to the Peer Management Center, the state will change to Running/Active.

Synchronization Summary

Runtime Summary View (auto-update enabled)

Filter by: [] Actions [] ☒ Enable Auto

Name	Overall Status	Backup Status	Failed Scans	Failed Events	Messages	Pending Events	Pending Retries	Checked	Updated
Composite4a[East_Coast]	PeerSync Service Not Avail...	N/A							
DDWin12R2a[EastCoast]	Running	Normal	0	0	0	No Pending Items	0	5738	5389
DDWin12R2b[WestCoast]	Running - Failed State	Target Folders ...	1	3	8	No Pending Items	0	258	0
VMSRV2008K32[Asia]	Stopped [Pending Installat...	N/A							
Win12x64a[EMEA_Region]	Running	Normal	0	0	0	No Pending Items	0	26880	0

Install PeerSync

Scheduling PeerSync Operation Task

☐ Always run in background

Run in Background Cancel Details >>

Logging and Alerts

Use the following dialog to enable or disable logging and alerts, including specifying event types to log.

File Synchronization Configuration [DDWin12R2b]

Monitoring

- Logging and Alerts
- Email Alerts

Associated Profile

- Jobs
- Global Settings
- Associated Installation

Logging and Alerts

Enabled: ☒

Severity: All

Event Types

☐ Stats Update ☒ Profile Update ☒ Profile Distribution

☒ PeerSync Service Start ☒ PeerSync Service Stop ☒ Failed Events Reprocess

☒ Restart Detected

Alerts

Severity: WARNING

OK Cancel

Stats Update	Log when PeerSync Stats are received (This could generate large amount of Log Entries).
Profile Update	Log whenever the PeerSync Profile configuration is updated.
Profile Distribution	Log when the PeerSync Profile is distributed to one or more hosts.

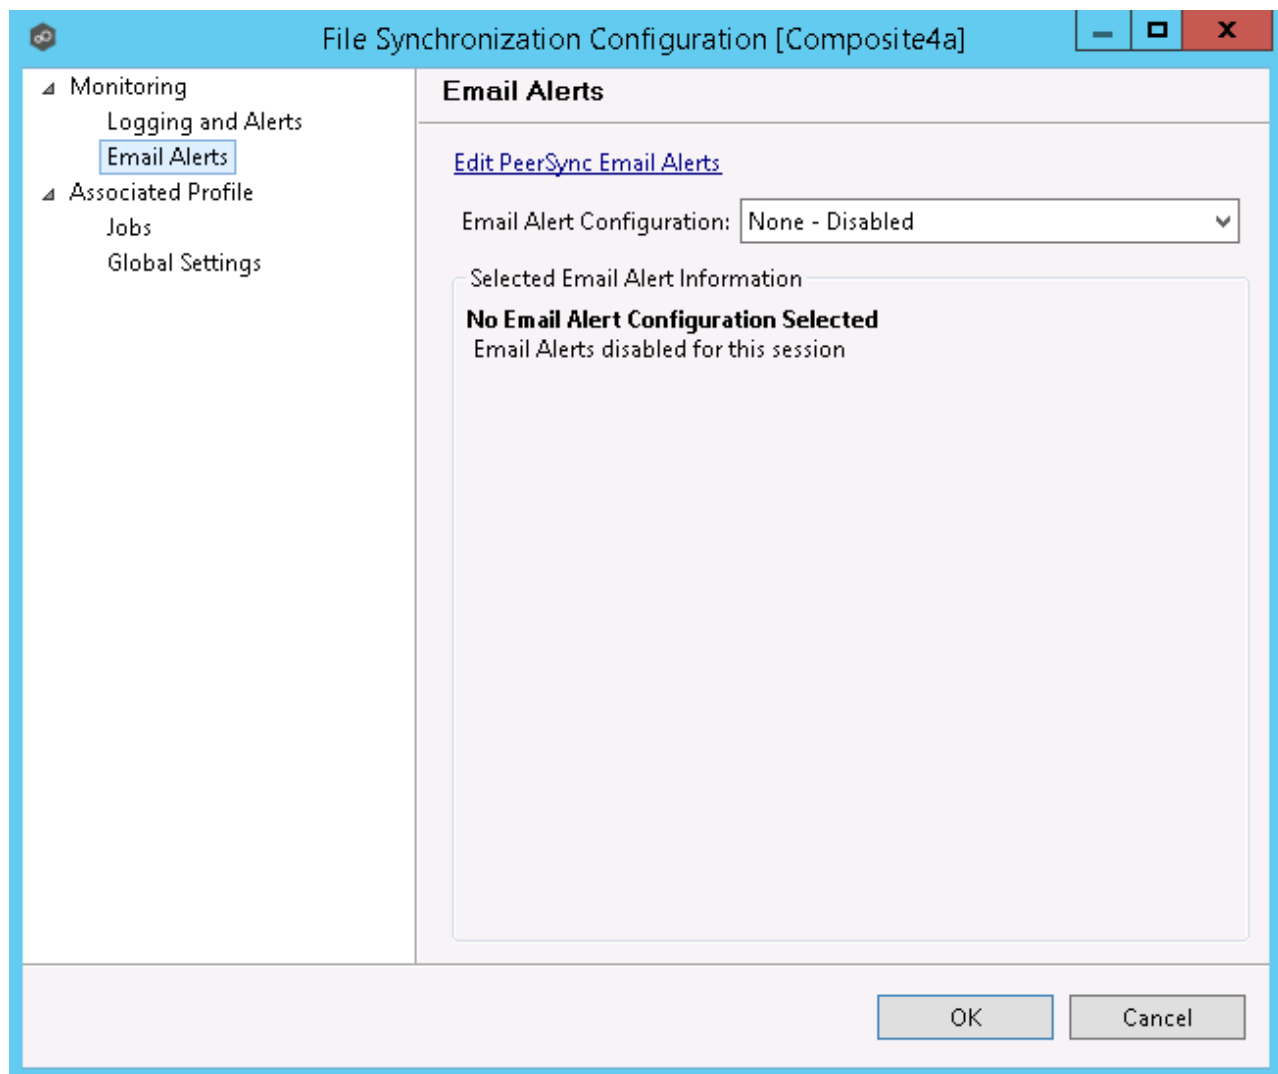
Stats Update	Log when PeerSync Stats are received (This could generate large amount of Log Entries).
PeerSync Service Start	Log when a user initiates a PeerSync Service Start.
PeerSync Service Stop	Log when a user initiates a PeerSync Service Stop.
Failed Events Reprocesses	Log when a user initiates a Failed Event Reprocess.
Restart Detected	Log when Peer Management Center detects that the PeerSync service has been restarted by comparing known Session Id with received one.

Email Alerts

Email Alerts

Email alerts configuration is available by selecting **Email Alerts** from the tree node within the File Synchronization Configuration dialog.

Email alerts are configured at a global level, then applied to individual file synchronization jobs. The following screen shows how this is accomplished.



To enable email alerts for this particular job, select an email alert from the drop down list. To disable, select **None - Disabled**. To edit the list of available configurations, select [Edit PeerSync Email Alerts](#).

Running and Managing a File Synchronization Job

This section includes topics for managing your File Synchronization Jobs.

- [Starting and Stopping](#)
- [Synchronization Summary View](#)
- [Synchronization Dashboard Summary View](#)

- [PeerSync Profile Management](#)
- [PeerSync Service Management](#)
- [Runtime Job Views](#)
- [Upgrade/Reprocess Installation](#)

Starting and Stopping

Starting a File Synchronization Job

Starting the job simply means starting the File Synchronization monitoring/management instance.

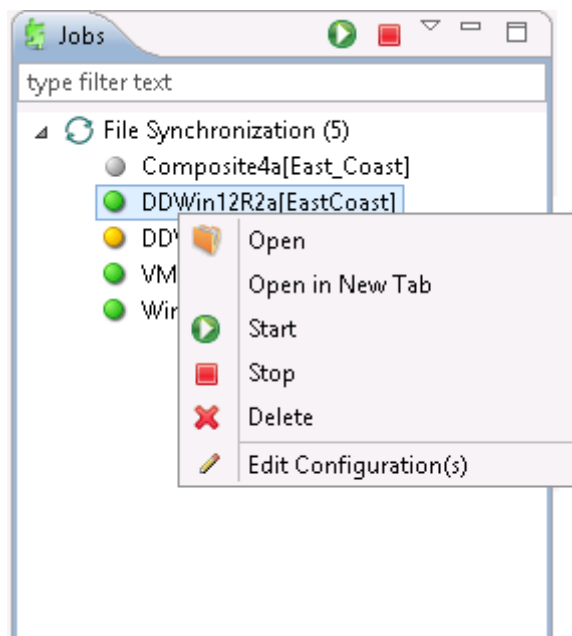
The job is auto started as soon as the Agent connects to the Peer Management Center; normally you will not need to manually start the job.

Click the **Start** button to begin the synchronization session.

Stopping a File Synchronization Job

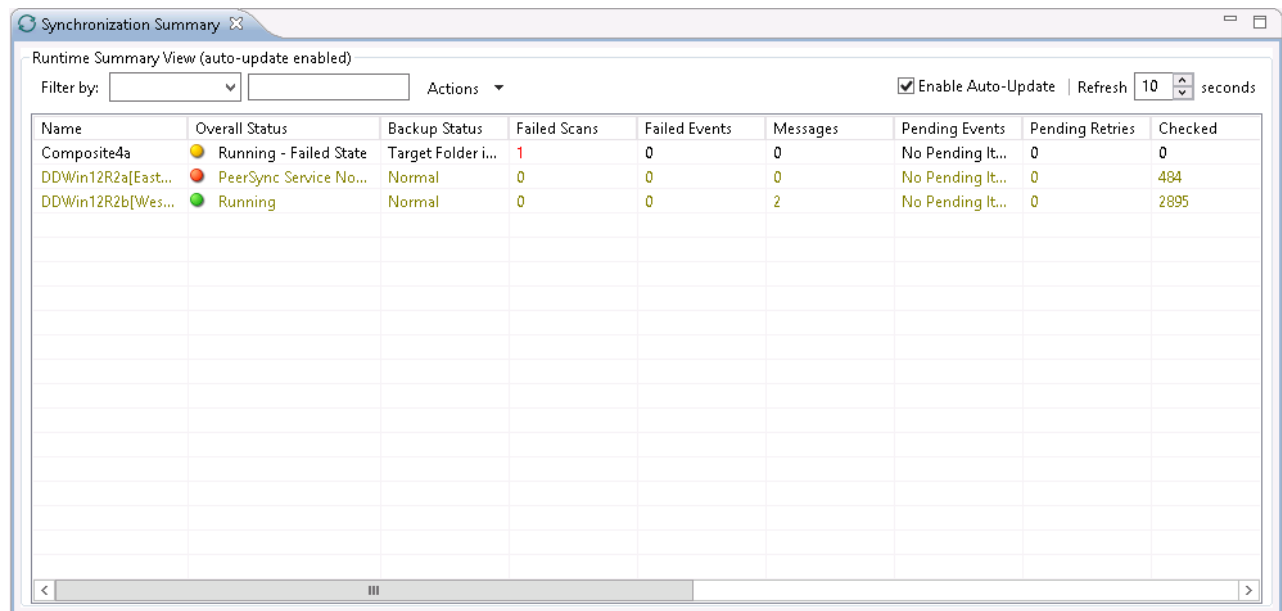
You can stop a file synchronization job at any time by pressing the **Stop** button. Doing this will shutdown the monitoring of the specific PeerSync host(s).

Note: If the job is stopped and the participating host is still running an instance of the PeerSync software, the job will auto start the next time that host agent is restarted or a Reconnect is detected.



Synchronization Summary

The Synchronization Summary view aggregates critical status and statistical information from all configured File Synchronization jobs in a single table view. It is automatically displayed when the Peer Management Center client is started and can be opened at any other time by double-clicking on the **File Synchronization** parent tree node in the Job's View or by clicking the **View Runtime Summary** icon in the toolbar of the Jobs view. Information in this view can be sorted and filtered. Operations such as starting, stopping, and editing multiple jobs at once are available, in addition to the ability to clear Monitoring Alerts, Start PeerSync, Stop PeerSync, Reprocess Failed Events, Request Support Info File and Reprocess/Upgrade Installation.



Unlike other views within the Peer Management Center, the Synchronization Summary View is not updated in real-time. This is done for performance reasons. Instead, the table can be set to automatically update itself every few seconds. Clicking **Enable Auto-Update** will enable this functionality, while the refresh interval (in seconds) can be set right beside the check box. Additional columns can be added to and removed from the table from the right-click context menu.

Double-clicking on any item in the table will automatically open the selected File Synchronization job in a tab within the Runtime Summary view, allowing you to drill down and view specific information about that single job. Items in the summary table can be filtered by job name, overall status, activity state and host participant name.

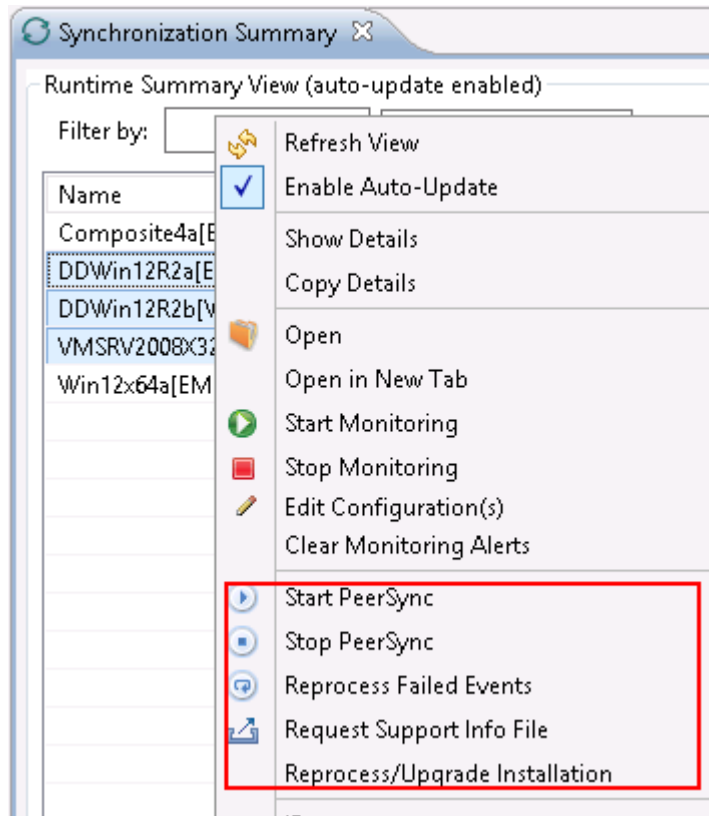
Selecting one or more items in the table, then right-clicking will bring up a context menu of available actions that can be performed on the selected jobs. The actions that are unique to this table are as follows:

Clear Monitoring Alerts	<p>Clears all monitoring alerts for the selected jobs. This can be performed while a job is running.</p>
--------------------------------	--

PeerSync multi-Job global actions:

Clear Monitoring Alerts

Clears all monitoring alerts for the selected jobs. This can be performed while a job is running.



Start PeerSync	Send a Start command to the PeerSync service instance running on the selected jobs' participant. This can be performed while the associated File Synchronization Job is running.
Stop PeerSync	Send a Stop command to the PeerSync service instance running on the selected jobs' participant. This can be performed while the associated File Synchronization Job is running.
Reprocess Failed Events	Send a Reprocess Failed Events command to the PeerSync instance running on the selected jobs' participant. This can be performed while the associated File Synchronization Job is running.
Request Support Info File	Send a request to collect the Support info File from the PeerSync instance running on the selected jobs' participant. This can be performed while the associated File Synchronization Job is running.

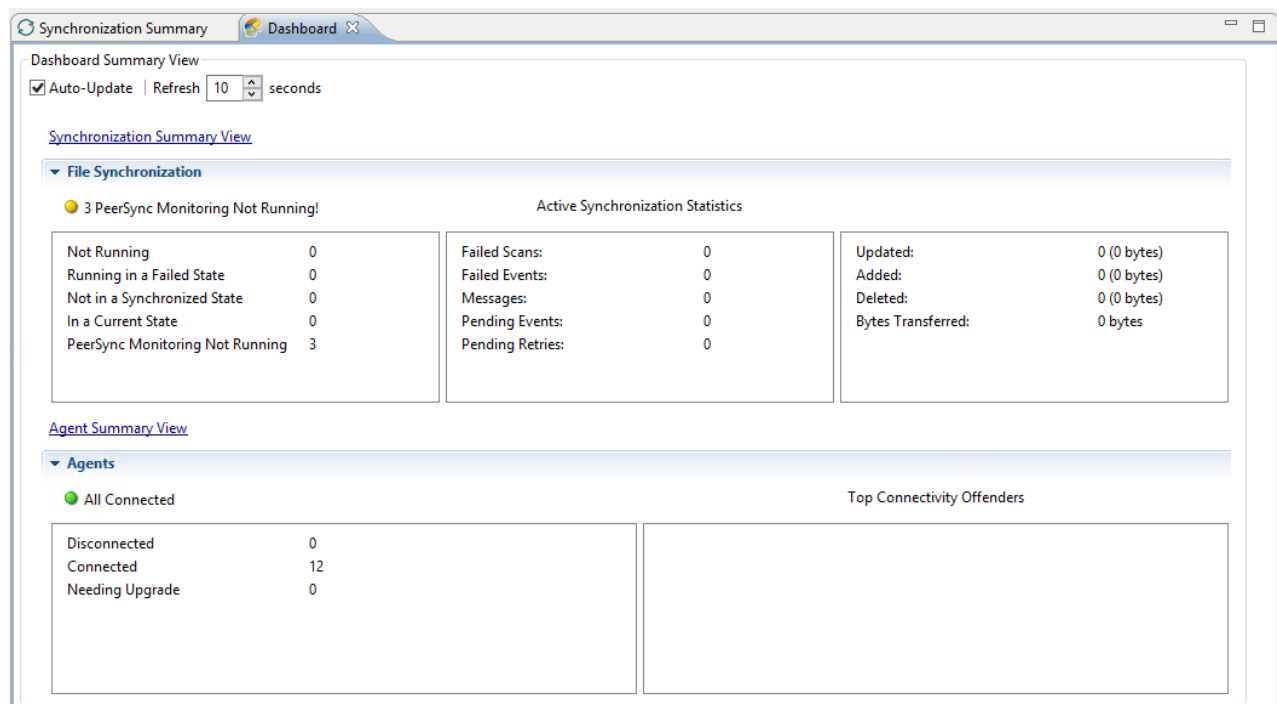
Clear Monitoring Alerts	Clears all monitoring alerts for the selected jobs. This can be performed while a job is running.
Reprocess/Upgrade Installation	Deploy an upgrade or reprocess an existing installation for the selected File Synchronization Job(s). Upgrade/Reprocess Installation

Clicking the **Actions** table menu provides the following options:

Refresh View	Refresh all information provided in the table.
Copy All Filtered Statistics	Copy detailed information to the system clipboard for all items current displayed in the table, taking any filters into account. This information can then be pasted into a document editor.
Export Entire Table to File	Dump the entire contents of the table to a text file that can be viewed in any document editor.

Synchronization Dashboard Summary View

The File Synchronization Dashboard Summary view is a panel that displays metrics and key performance indicators from all running File Synchronization jobs. It is automatically displayed when the [Peer Management Center](#) client is started and can be opened at any other time by selecting View Dashboard from the [Windows](#) menu or by clicking the **View Dashboard** icon in the Peer Management Center [toolbar](#).



The Dashboard is not updated in real-time. This is done for performance reasons. Instead, the table can be set to automatically update itself every few seconds. Enabling the **Auto-Update** option will enable this functionality, while the **Refresh** interval (in seconds) can be set right beside the check box.

Entries in the first column of the **File Synchronization Job** and **Agents** categories can be double-clicked, which will take the user to a filtered Runtime view of the selected item for additional details.

PeerSync Profile Management

The topics in this section provide some basic information on PeerSync Profile Management

- [Update the Profile Configuration](#)
- [Import Existing Profile](#)
- [Distribute Profile](#)

This topic covers information on how to update a PeerSync profile from the Peer Management Center.

We recommend if using the Peer Management Center to manage the PeerSync instances, to make changes through the Peer Management Center. If changes are made directly on the PeerSync machine, they should be [imported](#) in the Peer Management Center job manually to keep the Peer Management Center PeerSync Configuration in sync.

How to Update a PeerSync Profile through the Peer Management Center

- From the Synchronization Summary runtime view (Double-click the File Synchronization jobs node from the left), right-click the machine you wish to modify the profile for and choose **Edit Configuration(s)**. Alternately you can right-click the machine job from the left menu under the **File Synchronization** node and choose **Edit Configuration(s)**.
- You can update the Profile by importing an updated Profile through the **Import** button in **Associated Profile** screen, or manually update the configuration through Jobs and/or Global Settings section.
- If you wish to update the profile outside of the Peer Management Center, export the existing configuration using the **Export** button in **Associated Profile** screen. Make your changes through the PeerSync Profiler and import the updated Profile back in the Hub through the **Import** button.
- After having made your entire configuration changes either through the Hub or by [importing](#) the updated Profile, choose **OK** and close the **Edit Configuration** screen.

Your configuration changes will not reach the PeerSync machine until they are [distributed](#). The updated profile will become active on the machine after the PeerSync service has been restarted.

- [Import Existing Profile](#)
- [Edit/Configure Jobs](#)
- [Edit Global Settings](#)
- [Distribute Profile](#)

Import Existing Profile

In the Associated Profile section of the File Synchronization Configuration screen, you can update the configured profile with one you have saved and configured outside of the Peer Management Center.

Note: If making changes outside of the Peer Management Center, it is recommended to export the profile from the Peer Management Center (by clicking the **Export** button), making necessary changes outside of the Peer Management Center, and finally import the profile back into the Peer Management Center.

Click the **Import** button on the right of the screen to import the profile. To propagate this new updated profile, close the File Synchronization screen, reopen it and distribute to the PeerSync host through the [Distribute](#) button.

File Synchronization Configuration [DDWin12R2a]

Monitoring
Logging and Alerts
Email Alerts
Associated Profile
Jobs
Global Settings
Associated Installation

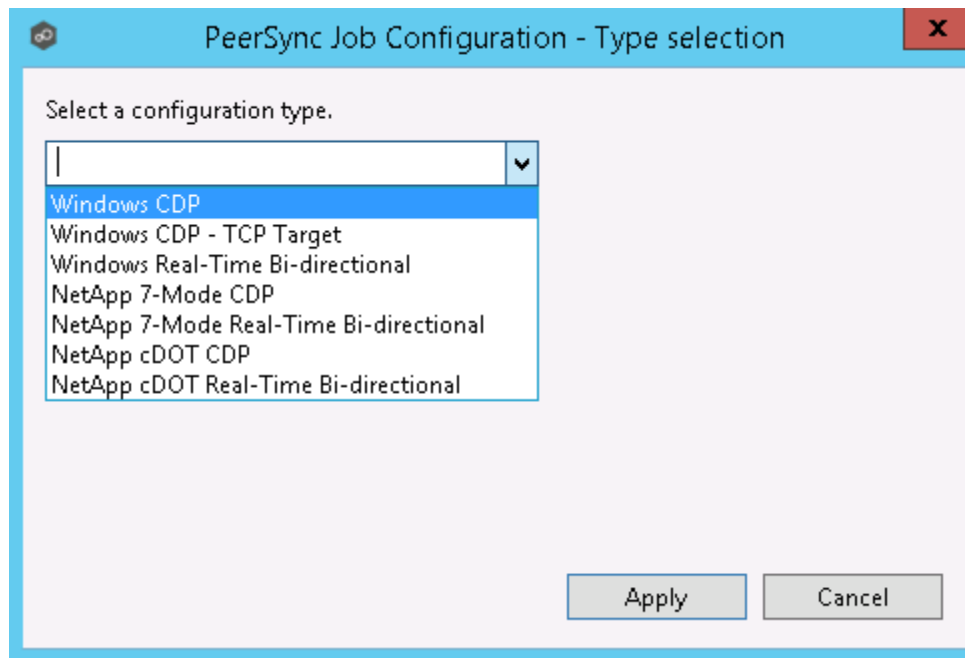
Associated Profile

Description: Backup
File Name: DDWin12R2a.snc
Last Updated On: Aug 24, 2015 10:47:38 PM
Current State: Active

Export
Import
Distribute

Total number of jobs configured in this profile is 1

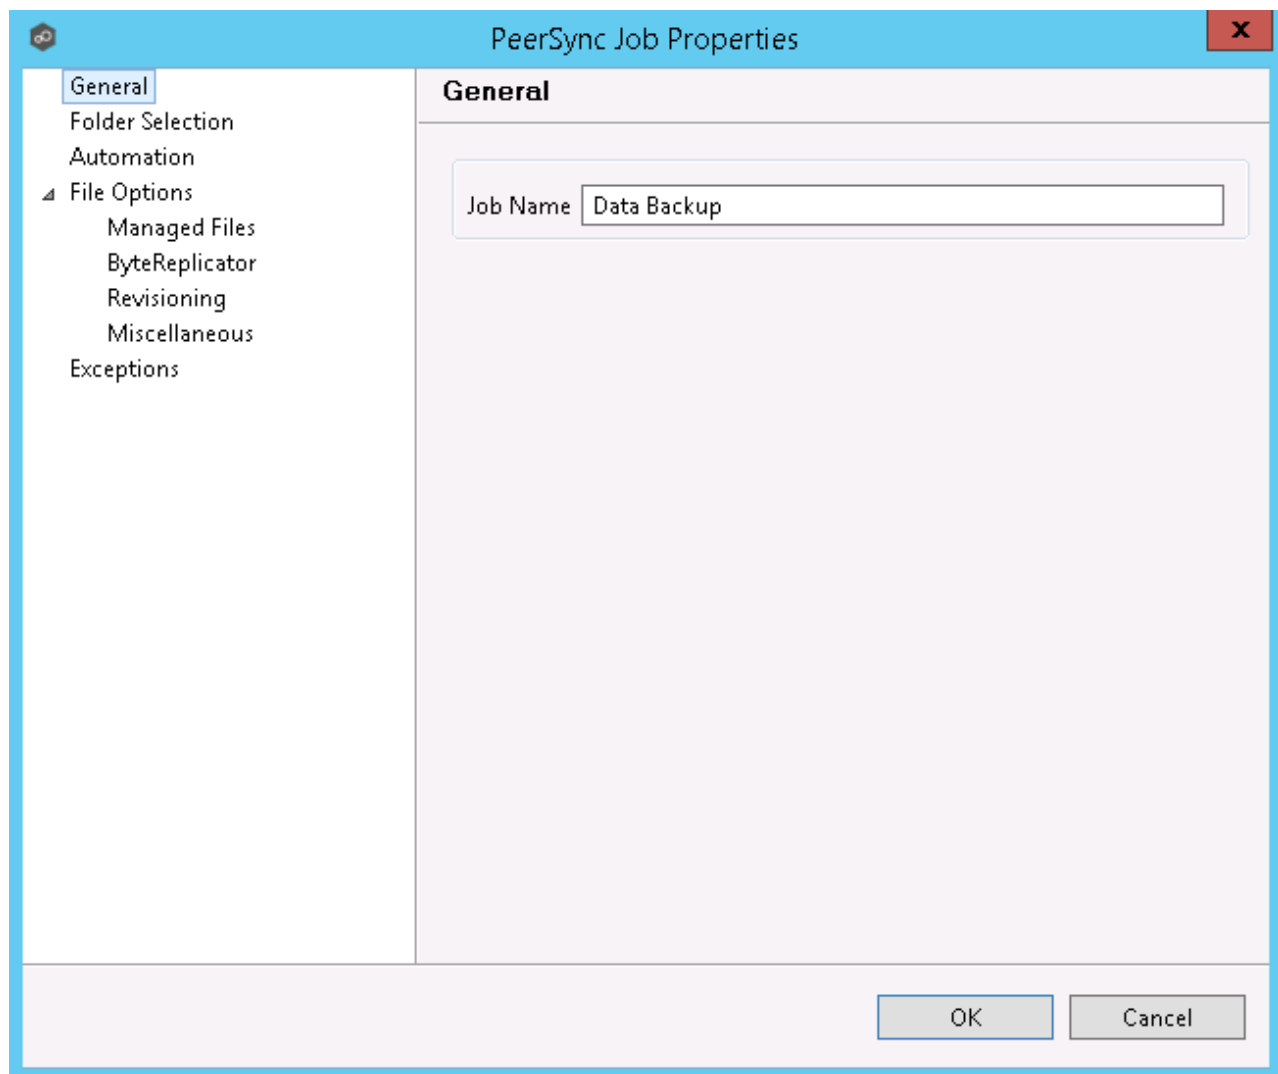
OK Cancel



Once a job type has been selected, click **Apply** and complete the PeerSync Job Configuration wizard screens to complete the job configuration and add the job to the profile.

Edit Existing Job

To edit an existing job, select the job in the **Jobs** view, and then click the **Edit** button on the right. The PeerSync **Job Properties** screen will open with all available settings grouped by category in a left menu tree.



Enable/Disable Job

To enable or disable a job, click the check box to the left of the job name in the **Jobs** view.

To save these changes, click **OK** on the bottom right of the **File Synchronization Configuration** screen.

Copy Job

You can copy an existing job by selecting the job from the **Jobs** view and clicking the **Copy** button on the right. The **PeerSync Job Properties** screen will open, allowing you to make changes to the copied job.

Note: You have to make at least one change to the job settings. If the job settings remain identical, it will not be saved after the **OK** button is clicked.

Remove Job

To remove jobs from the PeerSync Configuration, select one job from the **Jobs** view, click the **Remove** button on the right. Repeat this for any additional jobs you wish to remove.

Edit Global Settings

In the Global Settings of the File Synchronization Configuration screen, you can make changes to settings that apply to all PeerSync Jobs within the profile.

The screenshot shows the 'File Synchronization Configuration [DDWin12R2a]' dialog box with the 'Global Settings' tab selected. The left sidebar contains a tree view with the following items: Monitoring, Logging and Alerts, Email Alerts, Associated Profile, Jobs, Global Settings (highlighted), and Associated Installation. The main area displays the following settings:

- Recovery Options**
 - Retry open/inaccessible files: ☒ 1 times
 - Retry Failed Connection every: ☒ 1 minutes
- Performance Options**
 - Maximum number of Job Threads: 5
 - Maximum number of File Threads: 10
 - Use Enhanced Event Processing: ☒
- Reconnect Options**
 - Reconnect Options: Run a Scan on reconnect (dropdown menu)
 - Application Priority Selection: Normal (dropdown menu)

At the bottom right, there are 'OK' and 'Cancel' buttons.

Recovery Options	These changes will update how we retry failed or inaccessible files as well as the interval in which we retry Failed Connections.
Performance Options	These settings allow you to change the maximum number of job scans that can run parallel to one another and the maximum number of events that can be processed parallel to one another.
Reconnect Options	This setting allows you to choose how PeerSync handles a re-established connection. Options are to Run a Scan on Reconnect or Store missed events and process on reconnect.

Recovery Options	These changes will update how we retry failed or inaccessible files as well as the interval in which we retry Failed Connections.
Application Priority Selection	This setting enables to select the priority level you want PeerSync to have.

Distribute Profile

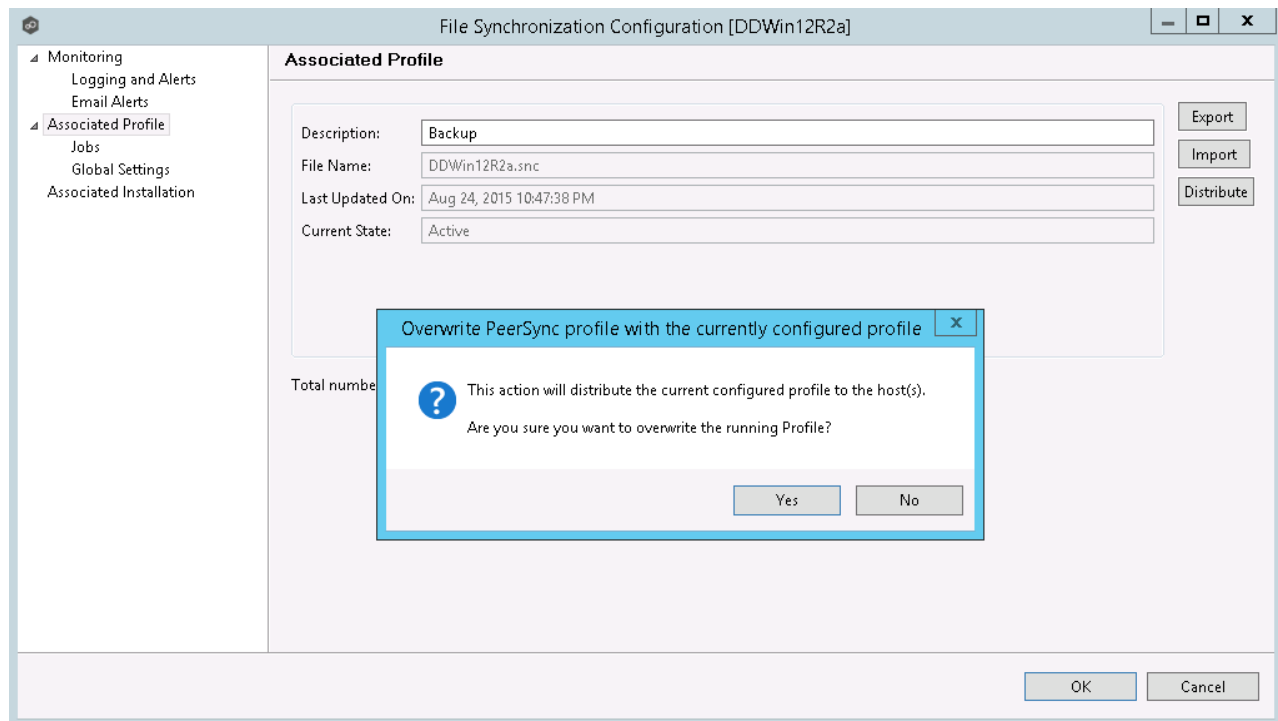
This topic covers information on how to distribute changes to the PeerSync profile from the Peer Management Center

To distribute the PeerSync profile changes, right-click on the File Synchronization job from the **Jobs** view, and then click **Edit Configuration**.

In the File Synchronization Configuration screen, click the **Associated Profile** node and click on the **Distribute** button.

In the event that one or more of your jobs are configured to use a ByteReplicator Relay Server (usually used in NetApp source environments), the Distribute Profile process will also distribute your relay Server configurations by compiling all unique target Hosts and relay servers into a `%profilename%.pls` file. This file will be distributed to the PeerSync machine along side the profile.

Note: This action will distribute the profile to the machine and attempt to stop and start PeerSync Service to commit those changes. If you do not wish to restart PeerSync service, wait to distribute the profile until you are ready to have the service restart.



PeerSync Service Management

The following PeerSync service management actions are available from the [Synchronization Summary view](#) and the [Summary view](#) for a specific file synchronization job.

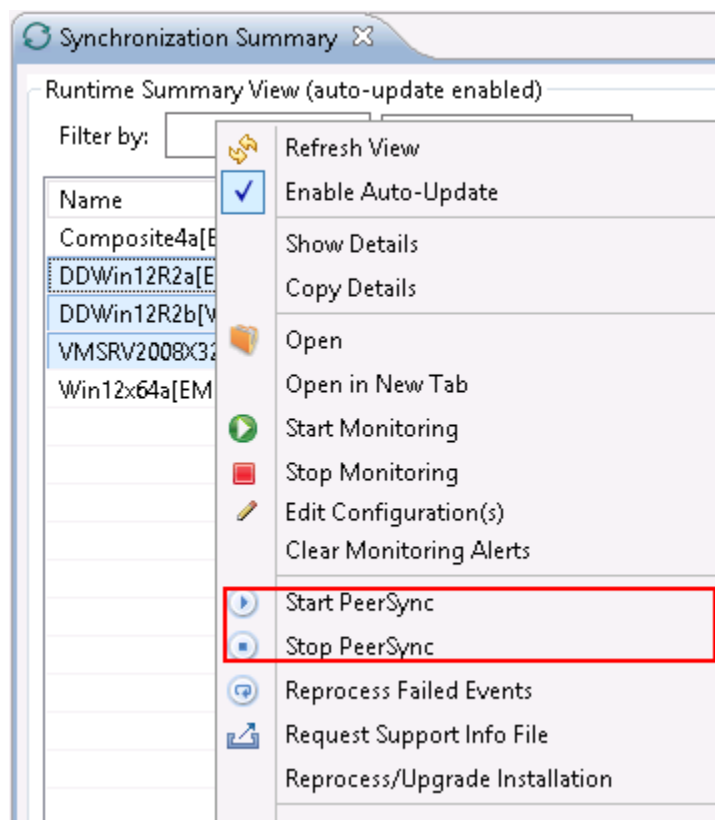
Starting the PeerSync Service

To start the PeerSync service associated with any file synchronization job, right-click the view and choose **Start PeerSync**.

Stopping the PeerSync Service

To stop the PeerSync service associated with any file synchronization job, right-click the view and choose **Stop PeerSync**.

Note: The associated file synchronization job has to be running in order to successfully perform this action.



For information on the additional PeerSync multi-job global actions, see [Synchronization Summary View](#).

Runtime Job Views

Double-clicking on the file synchronization job from the [Synchronization Summary view](#) will open the job-specific runtime views.

- [Summary View](#)
- [Failed Events View](#)
- [Monitoring Log View](#)
- [Alerts View](#)
- [Participants View](#)
- [Configuration View](#)

Summary
Failed Events (2)
Monitoring Log
Alerts (16)
Participants (1)
Configuration

Summary View
Actions

Monitoring Peerlet Session		Failed	Started: 9/30/15 12:23 PM
Stats Timestamp	9/30/2015 1:56:51 PM		
Stats State	Failed		
PeerSync Running Info		Started: 9/30/2015 12:09:46 PM	
Mode	Automatic and Real-Ti...		
Status	Target Folders are Not A...		
PeerSync Real Time Stats			
Real Time Events	128		
Real Time Events Peak	53.318 events/minute		
Real Time Events In Process	0		
Real Time Events Average	1.195 events/minute		

Overview
PeerSync Jobs Stats
Added
Updated
Deleted
Messages

PeerSync Overall Status					
Checked	260	Updated	0	Current Event Status	Real-time Monitoring (Failed Scans: 1) (Failed Events...
Excluded	0	Added	0	Pending Event Status	No Pending Items
Messages	11	Deleted	0	Bytes Transferred	0 bytes
Elapsed Time: 01:47:08 Xfer Rate: N/A					

When double-clicking a file synchronization job, the default selected tab will be the Summary tab. This view will show information received by the PeerSync machine on the status of the File Synchronization session.

Information found in this view is global to the PeerSync profile. To see PeerSync job-specific stats, you can click on the [PeerSync Jobs Stats](#) tab.

Information on this view is received whenever the information changes on the PeerSync machine, normally every 1 minute or so. To auto-refresh this view with the latest data click Enable Auto-Update on the top right of the view, and choose a Refresh cycle. The cycle is the

not the cycle for receiving the information, just to refresh the view with the latest information received by PeerSync.

Summary Failed Events (2) Monitoring Log Alerts (16) Participants (1) Configuration

Summary View Actions

Monitoring Peerlet Session Failed Started: 9/30/15 12:23 PM
Stats Timestamp 9/30/2015 1:56:51 PM
Stats State Failed

PeerSync Running Info Started: 9/30/2015 12:09:46 PM
Mode Automatic and Real-Ti...
Status Target Folders are Not A...

PeerSync Real Time Stats
Real Time Events 128
Real Time Events Peak 53.318 events/minute
Real Time Events In Process 0
Real Time Events Average 1.195 events/minute


Overview PeerSync Jobs Stats Added Updated Deleted Messages

PeerSync Overall Status

Checked	260	Updated	0	Current Event Status	Real-time Monitoring (Failed Scans: 1) (Failed Events...
Excluded	0	Added	0	Pending Event Status	No Pending Items
Messages	11	Deleted	0	Bytes Transferred	0 bytes

Elapsed Time: 01:47:08 Xfer Rate: N/A

On this page, you can right-click to display the PeerSync Actions menu:



On the bottom half of the page, you will find a set of tabs showing more granular information regarding this PeerSync session.

- [PeerSync Jobs Stats](#)
- [Added Files](#)
- [Updated Files](#)
- [Deleted Files](#)
- [Messages](#)

PeerSync Jobs Stats

When clicking on the PeerSync Jobs Stats view, a request goes out to the PeerSync machine to request job-specific stats and return them to the Peer Management Center to be displayed. These stats can only be requested if PeerSync is running on that machine and only if the File Synchronization job is started on the Peer Management Center.

When the stats are received, the view is updated with the job-specific stats and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

Overview

PeerSync Jobs Stats

Added

Updated

Deleted

Messages

Jobs Stats (List Updated on 09-30-2015 13:55:40)

3 jobs

Filter by:

Job Na...	Source	Target	Last Scan (Status: Durati...	Xfer Rate	Event Averages	Upda...	Added	Delet...	Mess...	Fail...	Status
<div></div> Data ...	c:\source\	c:\target\	9/30/2015 12:09:47 PM (...)	N/A	0.00 per min (0...	0 (0 ...	0 (0 ...	0 (0 ...	0	0	Normal
<div></div> Data2...	c:\source2\	c:\target2\	9/30/2015 12:09:47 PM (...)	N/A	0.00 per min (0...	0 (0 ...	0 (0 ...	0 (0 ...	0	0	Normal
<div></div> Data ...	c:\source3\	\\backupserve...	N/A	N/A	0.00 per min (0...	0 (0 ...	0 (0 ...	0 (0 ...	11	3	Target Not Available (Scan on Re...
All Jobs	-----	-----	N/A	N/A	0.00 per min (0...	0 (0 ...	0 (0 ...	0 (0 ...	11	3	Target Folders are Not Available f...

Added Files

When clicking on the Added tab, a request goes out to the PeerSync machine to request a list of latest added files and return it to the Peer Management Center to be displayed. This list can only be requested if PeerSync is running on that machine and only if the File Synchronization job is started on the Peer Management Center.

When the info list is received, the view is updated with the latest added events processed by PeerSync, and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

This information on this table can be filtered by Path or by Job Name.

Overview	PeerSync Jobs Stats	Added	Updated	Deleted	Messages
----------	---------------------	-------	---------	---------	----------

Added Log							
0 errors, 0 warnings, 100 others Filter by: <input type="text"/> <input type="text"/>							
Date	Type	Path	Comments	Message	Job Name	File Size	Modified Time
09-15-2015 17:0...	Event	c:\target\FILE42.TXT	Added [Attrib...		Data Backup	2.3 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE44.TXT	Added [Attrib...		Data Backup	1.9 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE43.TXT	Added [Attrib...		Data Backup	4.0 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE45.TXT	Added [Attrib...		Data Backup	4.0 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE47.TXT	Added [Attrib...		Data Backup	1.7 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE46.TXT	Added [Attrib...		Data Backup	1.6 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE48.TXT	Added [Attrib...		Data Backup	3.2 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE41.TXT	Added [Attrib...		Data Backup	3.3 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE50.TXT	Added [Attrib...		Data Backup	1.4 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE49.TXT	Added [Attrib...		Data Backup	2.7 KB	09-15-2015 17:01:27
09-15-2015 17:0...	Event	c:\target\FILE27.TXT	Added [Attrib...		Data Backup	2.1 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE28.TXT	Added [Attrib...		Data Backup	3.9 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE29.TXT	Added [Attrib...		Data Backup	2.6 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE30.TXT	Added [Attrib...		Data Backup	3.9 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE31.TXT	Added [Attrib...		Data Backup	2.8 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE32.TXT	Added [Attrib...		Data Backup	1.9 KB	09-15-2015 17:01:26
09-15-2015 17:0...	Event	c:\target\FILE33.TXT	Added [Attrib...		Data Backup	1.9 KB	09-15-2015 17:01:26

Updated Files

When clicking on the Updated tab, a request goes out to the PeerSync machine to request a list of latest updated files and return it to the Peer Management Center to be displayed. This list can only be requested if PeerSync is running on that machine and only if the File Synchronization job is started on the Peer Management Center.

When the info list is received, the view is updated with the latest updated events processed by PeerSync, and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

This information on this table can be filtered by Path or by Job Name.

Overview | PeerSync Jobs Stats | Added | **Updated** | Deleted | Messages

Updated Log

0 errors, 0 warnings, 100 others

Filter by:

Date	Type	Path	Comments	Message	Job Name	File Size	Modified Time
09-15-2015 17:0...	Event	c:\target\FLDR3\FILE33....	Updated [Attri...		Data Backup	2.5 KB	09-15-2015 17:03:50
09-15-2015 17:0...	Event	c:\target\FLDR3\FILE35....	Updated [Attri...		Data Backup	1.9 KB	09-15-2015 17:03:50
09-15-2015 17:0...	Event	c:\target\FLDR3\FILE37....	Updated [Attri...		Data Backup	3.8 KB	09-15-2015 17:03:50
09-15-2015 17:0...	Event	c:\target\FLDR3\FILE36....	Updated [Attri...		Data Backup	1.7 KB	09-15-2015 17:03:50
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE13....	Updated [Attri...		Data Backup	3.6 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE16....	Updated [Attri...		Data Backup	1.6 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE12....	Updated [Attri...		Data Backup	1.3 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE14....	Updated [Attri...		Data Backup	2.9 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE15....	Updated [Attri...		Data Backup	2.1 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE10....	Updated [Attri...		Data Backup	3.0 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE11....	Updated [Attri...		Data Backup	1.4 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE17....	Updated [Attri...		Data Backup	2.3 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE18....	Updated [Attri...		Data Backup	1.0 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE19....	Updated [Attri...		Data Backup	2.3 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE23....	Updated [Attri...		Data Backup	1.9 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE21....	Updated [Attri...		Data Backup	1.3 KB	09-15-2015 17:03:48
09-15-2015 17:0...	Event	c:\target\FLDR2\FILE20....	Updated [Attri...		Data Backup	3.0 KB	09-15-2015 17:03:48

Deleted Files

When clicking on the Deleted tab, a request goes out to the PeerSync machine to request a list of latest deleted files and return it to the Peer Management Center to be displayed. This list can only be requested if PeerSync is running on that machine and only if the File Synchronization job is started on the Peer Management Center.

When the info list is received, the view is updated with the latest deleted events processed by PeerSync, and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

This information on this table can be filtered by Path or by Job Name.

Overview | PeerSync Jobs Stats | Added | Updated | **Deleted** | Messages

Deleted Log (List Updated on 10-05-2015 11:57:51)

0 errors, 0 warnings, 10 others

Filter by:

Date	Type	Path	Comments	Message	Job Name	File Size	Modified Time
10-05-2015 11:5...	Data Bac...	c:\target\FLDR3	Scan		Session	0 bytes	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE6.TXT	Scan		Session	1.2 KB	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE4.TXT	Scan		Session	1.3 KB	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE8.TXT	Scan		Session	1.5 KB	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE2.TXT	Scan		Session	3.1 KB	10-05-2015 11:57:49
10-05-2015 11:5...	Data Bac...	c:\target\FILE5.TXT	Scan		Session	0 bytes	10-05-2015 11:57:48
10-05-2015 11:5...	Data Bac...	c:\target\FILE7.TXT	Scan		Session	2.1 KB	10-05-2015 11:57:48
10-05-2015 11:5...	Data Bac...	c:\target\FILE9.TXT	Scan		Session	2.2 KB	10-05-2015 11:57:48
10-05-2015 11:5...	Data Bac...	c:\target\FILE1.TXT	Scan		Session	1.3 KB	10-05-2015 11:57:48
10-05-2015 11:5...	Data Bac...	c:\target\FILE3.TXT	Scan		Session	1.4 KB	10-05-2015 11:57:48

Messages

When clicking on the Messages tab, a request goes out to the PeerSync machine to request a list of messages/errors logged and return it to the Peer Management Center to be displayed. This list can only be requested if PeerSync is running on that machine and only if the File Synchronization job is started on the Peer Management Center.

When the info list is received the view is updated with the messages logged by PeerSync, and the caption on the top of the view will show the date and time the list was last updated.

By right-clicking on the info table, you can choose to hide or show columns.

This information on this table can be filtered by Path, Job Name, or by Message.

Overview	PeerSync Jobs Stats	Added	Updated	Deleted	Messages
Message Log (List Updated on 09-30-2015 13:53:45)					
0 errors, 0 warnings, 13 others Filter by: <input type="text"/> <input type="text"/>					
Date	Path	Message	Job Name		
09-30-2015 13:2...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 13:2...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 13:2...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 12:2...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 12:1...	\\backupserver\da...	Failed Event - Add...	Data 3 Backup		
09-30-2015 12:0...	\\backupserver\da...	Connection Failure	Data 3 Backup		
09-30-2015 12:0...	\\backupserver\da...	Cannot create/got...	Data 3 Backup		

The Failed Events view allows you to see all those events that have failed to be processed by PeerSync. The list is populated when the File Synchronization starts, as well as in real-time as new failures occur. The information can be filtered by File Name.

Summary	Failed Events (2)	Monitoring Log	Alerts (16)	Participants (1)	Configuration
Failed Events					
2 Files Filter by File Name: <input type="text"/>					
Date	File	Cause	Status	Message	
09-30-2015 13:20:48	\\backupserver\data3\FLDR3\FILE...	ADDFILE	Failed Event	Connection Failure (Target Not A...	
09-30-2015 13:20:48	\\backupserver\data3\FILE4 - Co...	ADDFILE	Failed Event	Connection Failure (Target Not A...	

You can right-click on the info table and choose to **Reprocess Failed Events**. This action will send a request to PeerSync to retry all the Failed Events in the list.

The Monitoring Log view allows you to view recent event history for the currently running file synchronization job based on your Logging and Alerts settings. You can specify the maximum number of events to store in the table by adjusting the **Display Events** spinner located in the top right corner of the panel. The maximum number of events that can be viewed is 3,000. If you need to view more events or events from a prior session, then you can use the log files saved in the **Hub\logs** directory located in the installation directory. The event log files will start with **fs_event.log** and are written in a tab delimited format. Microsoft Excel is a good tool to use to view and analyze a log file. See the Logging and Alerts settings for more information about log files.

You can click on any column header to sort by the column. Warnings are highlighted in light gray, Errors are highlighted in red and Fatal errors are highlighted in orange. Error records will also contain an error message in the Message column.

To change what is being logged update the selected Event Types in the [Logging and Alerts](#) settings.

Summary ! Failed Events (2) **Monitoring Log** ! Alerts (16) Participants (1) Configuration

Event Log (Auto-Update Disabled)

0 errors, 0 warnings, 4 others | Filter by Severity: Filter by: Actions ▾

Date	Severity	Type	Host	Is Source	File	Comments	Message
09-30-2015 13:2...	INFO	Failed Events Reprocess	DDWin12R2b	true			Failed Event Repro...
09-30-2015 13:2...	INFO	Failed Events Reprocess	DDWin12R2b	true			Failed Event Repro...
09-30-2015 12:2...	INFO	Watch Directory		true			
09-30-2015 12:2...	INFO	Session Started		true			

Clicking on the **Actions** table menu provides the following options:

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

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

You can click on any column header to sort by that column. For example, clicking on the **Severity** column will sort by alert severity. Warnings are highlighted in light gray, while errors and fatal alerts are highlighted in red. A common error may be the PeerSync service is not running, which will trigger a PeerSync Quorum lost alert.

Summary	Failed Events (2)	Monitoring Log	Alerts (16)	Participants (1)	Configuration
---------	-------------------	----------------	-------------	------------------	---------------

Alert Log				
16 errors, 0 warnings, 0 others Filter by Severity: <input type="text"/>				
Received Date	Severity	Type	Host	Message
09-30-2015 12:09:30	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 12:09:30	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 12:09:30	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 12:09:30	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 12:07:30	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 12:07:30	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 12:07:30	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 12:07:30	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 11:53:59	FATAL	Application		Quorum lost for job DDWin12R2b[WestCoast]. Sessi...
09-30-2015 11:53:59	ERROR	Application	DDWin12R2b	Agent on host DDWin12R2b was restarted while job ...
09-30-2015 11:52:41	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 11:52:41	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 11:51:11	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 11:51:11	ERROR	Application	DDWin12R2b	Service Not Running
09-30-2015 11:50:54	FATAL	Application		PeerSync Quorum lost for job DDWin12R2b[WestCo...
09-30-2015 11:50:54	ERROR	Application	DDWin12R2b	Service Not Running

The following right-click menu items are unique to this particular table:

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

The Participants view shows the currently configured host participant for the selected file synchronization job and contains a column used to display activity status occurring on the hosts. If a host has become unavailable or the PeerSync service stopped, an error message will be displayed next to the failed host in red.

Summary Failed Events (2) Monitoring Log Alerts (16) **Participants (1)** Configuration

Host Participants

Host	Root Path	Status	State	Message
DDWin12R...	C:\Program Files (x86)\P...	Participating	Active	

Host Participant State Change Log

Filter by: Host: Status: State:

Date	Host	Status	State	Message
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Not Participating	Inactive	Job Stopped
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Not Participating	Inactive	Job Stopped
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Not Participating	Inactive	Job Stopped
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Host Resource Av...	Active	PeerSync Service is running
09-30-2015 ...	DDWin12R...	Host Resource Av...	Active	PeerSync Service Available
09-30-2015 ...	DDWin12R...	Host Resource Un...	Inactive	PeerSync Service Not Running
09-30-2015 ...	DDWin12R...	Host Resource Av...	Active	PeerSync Service is running
09-30-2015 ...	DDWin12R...	Host Resource Av...	Active	PeerSync Service Available
09-30-2015 ...	DDWin12R...	Host Resource Un...	Inactive	PeerSync Service Not Running
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Not Participating	Inactive	Job Stopped
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Participating	Active	
09-30-2015 ...	DDWin12R...	Host Unavailable	Inactive	Agent on host DDWin12R2b was restarted while jo

The Participants view also contains a table that displays the most recent host participant state changes, for example, when a host was removed from synchronization session, when a host came back online, or when the PeerSync service was stopped or started. This functionality is broken down into two parts: right-click context menu items and a subview entitled **Host Participant State Change Log**.

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

The following items are available in the right-click context menu for this table:

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

This view displays a quick summary of all configurable items for the selected file synchronization job. Each page of the File Synchronization Configuration dialog is represented in its own part of the view and can be collapsed if desired. Clicking **Edit this Configuration** will immediately bring you to the File Synchronization Configuration dialog where you can edit the current monitoring configuration or the Associated PeerSync profile.

Summary ! Failed Events (2) Monitoring Log ! Alerts (16) Participants (1) Configuration

[Edit this Configuration](#)

Configuration Summary

General Settings

Host Name: DDWin12R2b
 Session Name: DDWin12R2b[WestCoast]
 Session ID: 101
 Alert Severity: WARNING

Configured PeerSync Jobs

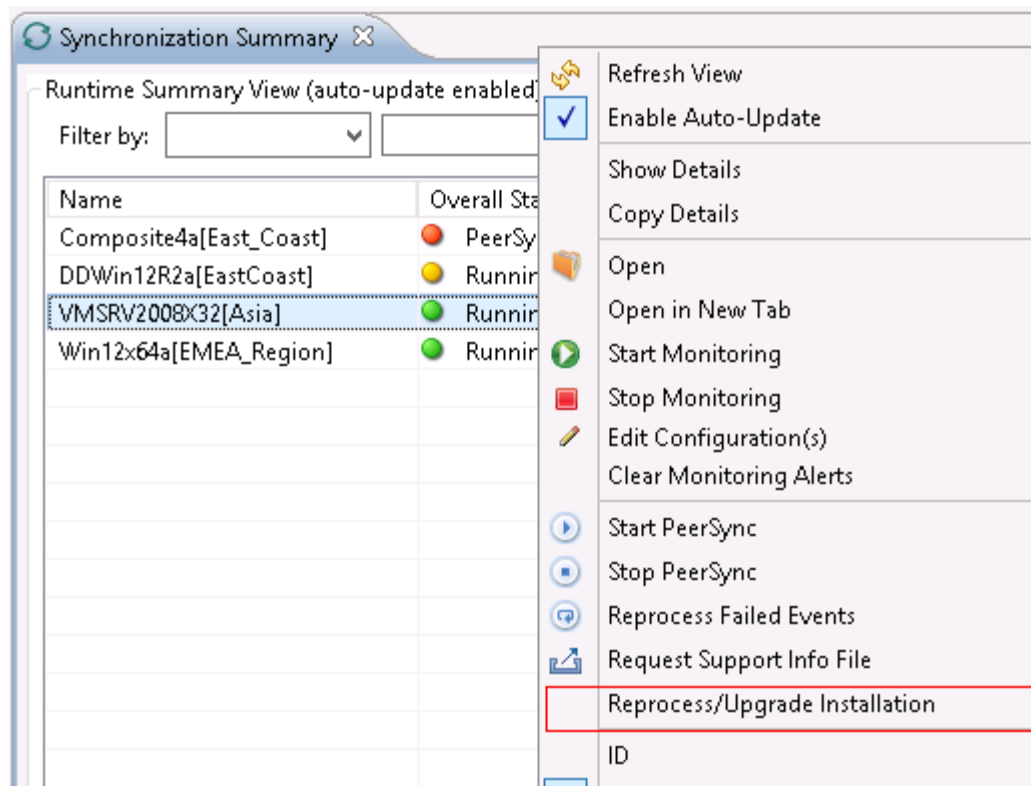
Filter by: ☒ Include Disabled Jobs

Name	#	Source	Target	Type	
<input checked="" type="checkbox"/> Data 3 Backup	3	c:\source3	\\backupserver\data3	Real Time Backup	
<input checked="" type="checkbox"/> Data Backup	1	c:\source	c:\target	Real Time Backup	
<input checked="" type="checkbox"/> Data2 Backup	2	c:\source2	c:\target2	Real Time Backup	

Monitoring Settings

Upgrade/Reprocess Installation

From the [Synchronization Summary view](#), you can click on one or more File Synchronization jobs and choose Reprocess/Upgrade Installation. This option will send a request to the selected PeerSync instances to install/upgrade given the configured settings.



The installation settings should be common for ALL the File Synchronization PeerSync instances in order to successfully install PeerSync.

The image shows a Windows-style dialog box titled "File Synchronization Configuration". It has a blue header bar with standard window controls (minimize, maximize, close) on the right. The main content area is titled "Installation Settings" and contains the following fields and sections:

- Remote Installation Settings:**
 - Pre-Defined Settings:** A dropdown menu set to "EastCoast".
 - Name:** A text field containing "EastCoast".
 - Installation Path:** A text field containing "%ProgramFiles%\Peer Software\PeerSync".
 - Existing Exe:** A dropdown menu (empty) and a "Browse" button.
 - Imported Exe Version:** A text field containing "v9.3.1.1007 [C:\Program Files (x86)\Peer Software\Peer Management Hub\Hub\workspace\peersync\templates\exes\peersync93-9_3_1_1007.exe]".
- License:**
 - User Name:** A text field containing "Peer Software".
 - Company:** A text field containing "Peer Software".
 - Options:** A text field with a yellow background.
 - Password:** A text field with a yellow background.
- Service:**
 - Logon User:** A text field containing "bytemetrics\adminidanielad".
 - Password:** A text field containing ten dots.

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

See [Installation Settings](#) for information on the settings on this page.

Index

- A -

Active Directory 47
Active Directory Authentication 47, 52
Active Directory Groups 47
Active Directory Users 47

- E -

EMC Isilon environment prerequisites 4
EMC VNX/Celerra environment prerequisites 4
Environmental requirements 4

- I -

Internal Users 47

- J -

job 160

- L -

LDAP Settings 47

- N -

NetApp 7-Mode environment prerequisites 4
NetApp cDOT/ONTAP9+ environment prerequisites 4

- O -

Office365 SMTP 44

- P -

Peer Agent 71
Peer Management Center 54
Peer Management Center Broker 133

- R -

Roles 47

- S -

Seeding Target 133

- U -

User Management 47